


TM

Supervising the Network

N O V E L L[®]

NetWare[®]

NETWORK COMPUTING PRODUCTS



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 to 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 any obligation to notify any person or entity of such changes.

trademarks

Novell, NetWare, and the N-Design are registered trademarks and the NetWare Logotype is a trademark of Novell, Inc.

A complete list of trademarks and their respective owners appears in "Trademarks" on page 861.

Copyright © 1993–1994 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.

**Novell, Inc.
122 East 1700 South
Provo, UT 84606
U.S.A.**

**Supervising the Network Volume I
December 1994
Part Number 100-002074-001**

HP7...: 95

85499



Contents

How to Use This Manual

Checklist for Installing and Setting Up Your Network	xiii
Managing Your Network	xiv

1 Managing NetWare Directory Services Objects

Default Objects and Rights for NetWare 4.1	2
Setting Up Administration Utilities	5
Creating Client Install Diskettes from CD-ROM	6
Installing an MS Windows Workstation and Starting NetWare Administrator	9
Installing an OS/2 Workstation and Starting NetWare Administrator	12
Installing a DOS Workstation and Starting NETADMIN	15
Rights Needed to Create and Manage Objects	17
Object Rights	18
Property Rights	20
Managing Trustee Assignments to Objects	22
Using NetWare Administrator to Manage Trustee Assignments	23
Using NETADMIN to Manage Trustee Assignments	24
Creating Container Objects	25
Types of Container Objects	25
Naming Container Objects	28
Creating Searchable Container Objects	29
Creating Container Objects Using INSTALL	30
Creating Container Objects Using NetWare Administrator	30
Creating Container Objects Using NETADMIN	32
Creating Leaf Objects	35
How to Use Leaf Objects	36
Naming Leaf Objects	41
Creating Searchable Leaf Objects	43
Creating Leaf Objects Using NetWare Administrator	44
Creating Leaf Objects Using NETADMIN	45
Managing Groups of User Objects	48
Managing Group Objects	49
Managing Profile Objects	59

Managing Organizational Role Objects	63
Managing User Templates	67
Searching for Objects	71
Searching for Objects Using NetWare Administrator.	72
Searching for Objects Using NETADMIN	74
Searching for Objects Using NLIST	75
Moving Objects in the Directory Tree	76
Moving Leaf Objects Using NetWare Administrator	77
Moving Leaf Objects Using NetWare Administrator (Drag-and-Drop)	78
Moving Leaf Objects Using NETADMIN	79
Moving Container Objects Using NetWare Administrator	81
Moving Container Objects Using NETADMIN	83
Deleting Objects from the Directory Tree	86
Cautions When Deleting NetWare Server Objects.	86
Cautions When Deleting User Objects	87
Cautions When Deleting Alias Objects	88
Deleting Objects Using NetWare Administrator	88
Deleting Objects Using NETADMIN	89
Renaming Leaf and Container Objects	90
Renaming Objects Using NetWare Administrator	91
Renaming Objects Using NETADMIN	92
Changing Object Property Values	94
Changing Object Property Values Using NetWare Administrator	94
Changing Object Property Values Using NETADMIN	99

2 Managing Directories, Files, and Applications

Planning Directory Structures	105
System-Created Directories	105
High Capacity Storage System (HCSS) Directories	106
Workstation Operating System Directories.	107
Application Directories	107
Data Directories	107
Sample Directory Structure.	108
Creating Directories and Copying Files	109
Creating Directories Using NetWare Administrator	109
Creating Directories Using FILER	111
Copying or Moving Files Using NetWare Administrator	112
Copying or Moving Files Using FILER	114
Loading Operating Systems and Applications onto the Network	117
Loading DOS onto the Network	117
Loading MS Windows onto the Network	119
Loading Other Applications onto the Network	123
Assigning Trustee Rights to Operating System and Application Directories	124

- Creating and Using Directory Map Objects 125
 - How to Use Directory Map Objects 125
 - Creating a Directory Map Object 127
- Making the File System Secure and Accessible 131
 - Understanding File and Directory Rights 131
 - Adding a Trustee to a Directory or File 134
 - Deleting a Trustee from a Directory or File 137
 - Modifying a Trustee's Rights to a Directory or File 140
 - Viewing/Modifying the Inherited Rights Filter for Directories and Files 144
 - Changing Attributes of a Directory or File 147
 - Changing the Owner of a Directory or File 150
- Viewing Effective Rights and Other Information 153
 - Viewing a Trustee's Effective Rights to a Directory or File 153
 - Viewing Other Information about a Directory or File 156
- Salvaging and Purging Deleted Files and Directories. 161
 - Salvaging Deleted Files 161
 - Purging Files 165

3 Creating Login Scripts

- About Login Scripts 169
 - Four Types of Login Scripts 170
 - Deciding Which Login Scripts to Create. 171
 - Using Login Scripts from Other NetWare Versions 174
- Creating, Modifying, Copying, and Printing Login Scripts. 175
 - Login Script Hints 176
 - Creating or Modifying a Login Script Using NetWare Administrator. 177
 - Creating or Modifying a Login Script Using NETADMIN 180
 - Copying a Login Script Using NetWare Administrator. 183
 - Copying a Login Script Using NETADMIN 185
 - Printing Login Scripts 187
- Login Script Commands and Variables 187
 - # (Execute External Program). 189
 - ATTACH. 191
 - BREAK 192
 - CLS 192
 - COMSPEC 193
 - CONTEXT 194
 - DISPLAY 195
 - DOS BREAK 196
 - DOS SET 197
 - DOS VERIFY 197
 - DRIVE 198
 - EXIT 199

FDISPLAY	201
FIRE PHASERS	202
GOTO	203
IF...THEN	204
INCLUDE.	208
LASTLOGINTIME	210
MACHINE	211
MAP	212
NO_DEFAULT	220
NOSWAP.	221
PAUSE	222
PCCOMPATIBLE.	222
PROFILE	223
REMARK	224
SCRIPT_SERVER	225
SET.	226
SET_TIME	228
SHIFT	229
SWAP	231
TEMP SET	232
WRITE	233
Identifier Variables	236
Using Identifier Variables.	240
Using LOGIN Parameters with %n Variables	241
Examples of Login Scripts	243
Default Login Script	243
Container Login Script	245
Profile Login Script	248
User Login Script	250

4 Creating Menus

Creating, Converting, and Modifying Menu Files	251
Getting Acquainted with NMENU	252
NMENU Syntax.	252
What Menus Look Like	252
What Makes Menus Work	254
Planning Your Menus	259
Using the NMENU Commands.	260
NMENU Organizational Commands and Options	260
NMENU Control Commands and Options	265
Creating a Menu File	272
Example Menu Programs	272
Making Menus Work	275

Setting Up the User Environment	275
Converting Old Menu Files	277

5 Managing the NetWare Directory Tree

About NetWare Directory Services	281
Creating and Managing Directory Services Partitions	282
About Partitions and Replicas	282
Planning and Using Partitions and Replicas	283
Restricting Access to Partition Manager	286
Creating a New Partition	286
Merging Partitions	290
Moving a Partition	293
Aborting a Partition Operation	299
Creating a Replica	302
Deleting a Replica	306
Sending Updates to Other Replicas	309
Receiving Updates from Other Replicas	312
Viewing a List of Partitions Stored on a NetWare Server	315
Viewing a List of Replicas of a Partition	318
Viewing a List of Partitions in a Directory Tree	320
Changing a Replica's Type	321
Removing NetWare Directory Services from a Server	325
Considerations Before Removing NDS	325
Removing NDS from a Server	326
Reinstalling NDS on a Server	328
Deleting a NetWare Server Object from the NDS Database	329
Considerations Before Deleting NetWare Server Objects	329
Deleting a NetWare Server Object Using NetWare Administrator	330
Deleting a NetWare Server Object Using PARTMGR	331
Importing User Information into the Directory Services Database	333
Understanding the Import Process	333
Preparing to Import User Information into the Directory Services Database	334
UIMPORT Field Name Changes	354
Examples of Files Used by UIMPORT	355
Creating User Objects with the UIMPORT Utility	359
Repairing the NetWare Directory Database	361
DSREPAIR Overview	361
DSREPAIR Options	363
Running an Unattended Full Repair	364
Checking Time Synchronization (DSREPAIR)	365
Checking Replica Synchronization	368
Viewing and Editing the Repair Log File	371
Using the Advanced Options	371

Merging NDS Trees	373
DSMERGE Overview.	373
Preparing the Source and Target Trees	376
Time Synchronization before the Merge	377
DSMERGE Options	378
Checking Servers in the Tree	379
Checking Time Synchronization (DSMERGE)	381
Merging Two Trees.	382
Renaming the Tree.	385
Completing the Tree Merge	387
Viewing and Managing NDS Synchronization Status	389

6 Migrating Data Using the High Capacity Storage System

What is the High Capacity Storage System?	391
HCSS Concepts	392
How the HCSS File System Is Structured	392
What Happens When Data Migrates	394
Media Access Control	396
Media Export and Import Operations.	396
Restrictions.	397
Overview of HCSS Setup and Management.	398
How the HCSS File System Is Managed.	399
Installing and Setting Up HCSS	400
Installing Hardware.	400
Loading HCSS Drivers	401
Creating an HCSS Volume.	402
Loading HCSS on the Server	403
Setting Up a Workstation to Manage HCSS	404
Setting Up the HCSS File System	405
Managing the HCSS File System	415
Exporting Media	415
Reloading Media	418
Deleting an HCSS Domain Directory.	418
Reformatting Media	420
Checking Media Status	420
Managing HCSS Parameters	421
What the Parameters Do	421
436Setting Parameters Using NetWare Administrator	427
Setting Parameters at the Server Console	434
Using HCSS with SFT III.	438
Loading HCSS Drivers with SFT III	438
Loading HCSS on the MEngine	440

7 Maintaining the NetWare Server

Common Management Tasks	444
Sending Console Messages to Workstations	444
Bringing Down a Server	446
Booting the Server with the EXIT Command	447
Loading a NetWare Loadable Module	448
Loading and Binding LAN Drivers	449
Viewing and Adding Server Search Paths	454
Installing, Uninstalling, and Configuring a Server Product	455
Copying NetWare Files	456
Extracting NetWare Files from the Installation CD-ROM	457
Creating or Editing a Server Batch (.NCF) File	459
Editing Text Files from the Server Console	460
Setting a Server's Language	462
Viewing and Setting Server Time and Time Zone	469
Changing a Server's Name or IPX Internal Network Number	471
Managing Server Connections	473
Disabling and Enabling Logins	473
Clearing a Workstation Connection	473
Monitoring Workstation Connections	475
Increasing the Maximum Number of Packet Receive Buffers	477
Increasing the Minimum Number of Packet Receive Buffers	478
Avoiding Bottlenecks During Peak Usage	480
Monitoring and Optimizing the Server	481
Assessing Server RAM	481
Controlling Memory Allocation	482
Speeding Directory Searches	484
Prioritizing Server Processes	486
Checking a Server's Processor Utilization	488
Viewing a Server's Packet Traffic	489
Viewing LAN Driver Statistics	489
Improving Server Performance	490
Improving Disk Writes	492
Improving Disk Reads	496
Controlling Resource Allocation with Locks	499
Displaying Alert Messages	502
Viewing the Server Error Log	506
Managing Error Log Files	508
Checking for Disk Errors	508
Testing for Bad Disk Blocks	511
Maintaining Volumes	515
Creating Volumes	515
Mounting and Dismounting Volumes	518

Deleting Volumes	519
Renaming Volumes	521
Setting Up a Volume to Store Non-DOS Files	522
Repairing Volumes	525
Adding a Segment to an Existing Volume	528
Modifying the Size of a Volume	529
Using a CD-ROM as a NetWare Volume.	530
Managing Server Hard Disks	535
Checking Available Disk Space with NDIR.	535
Using File Compression	536
Enabling File Compression on a Volume.	537
Setting Compression for a File or Directory	539
Suspending File Compression	543
Using SET to Control File Compression	544
Purging Files from a Disk.	549
Adding Optical Storage for File Migration	551
Adding a Hard Disk to the NetWare Server	551
Replacing a Hard Disk	556
Loading Disk Drivers	559
Creating NetWare Disk Partitions	561
Deleting NetWare Disk Partitions	563
Mirroring and Duplexing a Hard Disk.	565
Unmirroring Hard Disks	567
Formatting a Hard Disk.	569
Disaster Prevention and Recovery.	570
Securing the Server Console.	570
Preventing Virus Infection	571
Preventing Packet Forgery	572
Activating UPS Monitoring	579
Protecting the Operating System's Memory	581
Protecting Database Integrity with TTS	589
Mirroring and Duplexing to Protect Data	591
Recovering Data from an Unmirrored (Out of Sync) Hard Disk	591
Managing Network Time Synchronization	593
SET Parameters for TIMESYNC	594
Editing the TIMESYNC.CFG File.	595
Creating a Custom Time Source Configuration	597
Adjusting a Fast or Slow Reference Time Server	600
Configuring Time Synchronization for Large Networks	601
Using Remote Console to Manage a Server.	606
Understanding Connection Types	607
Running a Remote Console Session Through a Direct Connection	607
Running a Remote Console Session Over a Modem	613
Loading Remote Modules on Bootup	619

Upgrading Files on a Remote Server	621
Installing or Reconfiguring Applications on a Remote Server	622
Rebooting a Remote Server	624
Administering Accounting	626
Setting Up Accounting	626
Calculating Charge Rates for Accounting	628
Viewing Accounting Totals	629

8 Auditing Network Events

What Is NetWare Auditing?	631
Enabling Auditing	634
Creating a Directory and an Auditor User Object	634
Enabling Auditing for a Volume	636
Enabling Auditing for a NetWare Directory Services Container	637
Auditor Procedures	639
Setting Passwords	640
Logging In	643
Setting Up Your Work Environment	646
Selecting Items to Audit	650
Creating Audit Reports	655
Creating Report Filters	655
Sending Reports to a File	658
Viewing a Report on the Screen	660
Maintaining Audit Files	662
Auditing Examples	666
Examples of Selecting Items to Audit	667
Examples of Creating a Report Filter	672
Example of Using a Filter without Saving the Changes	676

9 Backing Up and Restoring Data

What SBACKUP Does	679
Backup and Restore Concepts	680
How SBACKUP Presents Choices	680
Rules for Using SBACKUP	682
Using SBACKUP	682
Using Target Service Agents	684
Installing and Setting Up SBACKUP	685
Supported Storage Devices and Drivers	685
Memory	687
Loading Files Needed to Run SBACKUP	687
Loading Procedures for Drivers	688
Loading Procedures for Servers	689

Loading Procedures for Workstations	690
Loading Procedures for NetWare Directory Services	694
Loading Procedures for SFT III Users	695
LOAD Commands Summary	696
Backing Up Data	698
Introduction to Backing Up Data	698
Prerequisites to Backing Up	710
Procedures for Full, Differential, or Incremental Backups	711
Procedures for Custom Backups	721
Restoring Data	736
Introduction to Restoring Data	736
Prerequisites to Restoring	746
Procedures for a Restore.	747
Log and Error Files	762
Introduction to Log and Error Files	762
Viewing a Log File	764
Viewing an Error File	766
Setting Location of Log and Error Files	768
Deleting the Log Files for Overwritten Sessions	770
Storage Device Administration	772
Viewing Media Lists	772
Checking Status of a Device	773
Checking Status of Media in a Device	774
Media Utilities	776
Renaming a Storage Device	777
Managing Storage Media.	778
Changing Targets	779
Unloading SBACKUP Files	781
Enhancing SBACKUP Performance	782
Compressed Files	782
Enhancing Host Server Performance	783
Troubleshooting	786

A Troubleshooting the Network

Troubleshooting Hardware and Network Problems	789
Making Computer Memory Available for Network Drivers	789
Resolving Server Hardware Problems	790
Resolving Network Problems.	791
Resolving Cabling Problems	792
Power Supply Errors	792
Troubleshooting the NetWare Server	794
Resolving Abends	794
Resolving Server I/O Errors	796

Resolving Server Disk Space Problems	797
Resolving Server Memory Problems	798
Resolving Locked Device Errors	799
Resolving File I/O Errors	799
Resolving Volume I/O Errors	800
Resolving Event Control Block Allocation Errors	801
Troubleshooting Workstations	802
Resolving Workstation Memory Problems	802
Resolving Workstation Disk Space Problems	803

B Managing NetWare Server for OS/2

NetWare Server for OS/2 and NetWare 4.	805
Loading NetWare Server for OS/2	806
Using NWOS2.EXE to Load the Server from the Command Line.	806
Using the OS/2 Graphical Interface to Load the Server.	806
Navigating in NetWare Server for OS/2	807
Using Loadable Modules in NetWare Server for OS/2	807
Monitoring NetWare Server for OS/2	808
Loading PMMON	808
Using PMMON	810
Setting Performance Tuning.	811
Setting Memory Sharing.	812
Releasing Server Memory.	815
Supporting Workstations	816
Configuring NetWare Server for OS/2	816
NET.CFG File Format.	817
Conventions for Entering NET.CFG Options	817
Editing the NET.CFG File Using the Installation Program.	819
Editing the NET.CFG File Using an OS/2 Text Editor.	821
NET.CFG Options and Settings.	822
The NetWare for OS/2 Option.	822
The NetWare Requester Option.	829

C SFT III Management Tips

Server Synchronization	832
Detecting Server Synchronization Errors	832
Handling Secondary Restarts Immediately After Synchronization	832
Notifying Users of Server Synchronization	833
Reducing the Time for Resynchronizing and Remirroring.	833
Responding to "Invalid Mirrored Server Initialization" Messages	833
Responding to "Mirrored Server Engine Already Loaded" Messages.	834
Mirrored Server Link (MSL)	834

Determining Appropriate MSL Cable Lengths	834
Handling MSL Cable or MSL Board Failure	834
Responding to MSL Communications Error Messages	835
Responding to "MSL Deadlock Delivering Data" Messages	835
Server Configuration (.NCF) Files	836
Server Memory	837
Adding Memory Without Bringing the MSEngine Down	837
Responding to "Secondary Server Is Missing RAM" Messages	838
Server Consoles	840
Logging and Viewing All Console Messages	840
Handling Lost Interrupt Alerts on the Console	841
Responding to "Loader Cannot Find..." Messages	841
Shortening the Console Prompt	841
Server Hard Disks	842
Controlling the Size of Log Files on Volume SYS:	842
Handling Primary Hard Disk Failure	842
Mounting a CD-ROM as a NetWare 4.1 SFT III Volume	843
Notifying Users of Disk or Server Failure.	845
Recovering an Orphaned Partition	845
Network Clients	847
Finding a Server from an Ethernet Workstation	847
Finding the SFT III Server First.	847
Handling DOS IPX Session Timeouts During Server Switchover	847
Logging in ARCnet or Token Ring Clients Immediately After a Halt	848
Responding to "LAN Driver Loopback Error Detected" Messages	848
Turning Off the IPX Checksum Option for Token Ring.	848
Network Performance	848
Handling Frequent Ups and Downs of the IPX Internet	848
Responding to "IPX Network No Longer Returning Status" Messages	849
Handling Primary Network Board Failure	850
Troubleshooting	851

Trademarks

Novell Trademarks	861
Third-party trademarks	863

Index



How to Use This Manual

Supervising the Network has two purposes: to help you set up your network after completing the installation procedure, and to help you manage your network after you have installed or upgraded servers to NetWare® 4.1.

This manual is to be used by network supervisors responsible for maintaining all or part of a NetWare 4.1 network.

The checklist in the following section explains the process of planning and setting up a network, and refers you to the manual containing the information and instructions you need to complete that step in the process.

Checklist for Installing and Setting Up Your Network

To set up your NetWare 4.1 network, use the manuals in your NetWare 4.1 set in the order shown in the following checklist. (Chapter-only references refer to chapters within *this* manual.)

Checklist



- Plan the network. *New Features* gives you an overview of NetWare 4.1 features and the planning information you need.
- Install or upgrade servers to NetWare 4.1. See *Installation or Upgrade* for instructions.
- Set up NetWare Directory Services™ (NDS) objects, such as organizations, groups, and user accounts. Chapter 1, “Managing NetWare Directory Services Objects,” gives you information and procedures.
- Set up your file system and install network applications. Chapter 2, “Managing Directories, Files, and Applications,” contains planning guidelines and instructions for creating directories and installing applications.

- ❑ Plan and create login scripts. Chapter 3, “Creating Login Scripts,” explains the types of login scripts you can use and gives you instructions for setting them up.
- ❑ Set up printing. See *Print Services* for information and instructions.
- ❑ Set up workstations. See *NetWare Client for DOS and MS Windows User Guide* or *NetWare Client for OS/2 User Guide* for information and instructions.



For a list of the command format conventions used in Novell® documentation, see Chapter 1, “Command Syntax,” of *Utilities Reference*.

Managing Your Network

After setting up your network, use this manual to help you keep it updated and running efficiently.

Use this manual to help you perform such routine tasks as managing your NetWare Directory database and its associated objects, maintaining your network file system, and setting up servers and keeping your network running efficiently.

You can also use it as a reference for backing up and restoring network data, using optical storage devices to increase and manage your available disk space, and using the new NetWare DynaText* online documentation.

To get an overview of how this manual can help you manage your network, review the tasks listed in the Table of Contents (see page i).



The printed version of this manual is divided into two volumes. Volume I contains the Table of Contents and Chapters 1 through 6. Volume II consists of Chapters 7 through 9, Appendixes A through C, and the Index.



chapter

1

Managing NetWare Directory Services Objects

This chapter helps you start managing your network after installing the NetWare® 4.1 operating system.

The first section in this chapter explains what objects exist in the Directory tree immediately after you have installed NetWare 4.1 or upgraded to NetWare 4.1.

The rest of the sections explain basic “getting started” tasks, such as

- ◆ Installing a workstation and logging in to the network for the first time
- ◆ Creating and managing objects, such as Users, Groups, Organizational Roles, and Profiles
- ◆ Searching for objects
- ◆ Moving, deleting, and renaming objects
- ◆ Changing properties of objects

You can use either the NetWare Administrator graphical utility or the NETADMIN text utility to do those tasks. Both methods are explained. In some cases, a third method is available, which is also explained.

Default Objects and Rights for NetWare 4.1

Table 1-1 explains what objects exist in a Directory tree immediately after you install NetWare 4.1 or upgrade to NetWare 4.1, and what default rights those objects have.

Table 1-1
Default Objects and Rights in a Directory Tree

Default Objects after NetWare 4.1 Installation	Default Objects after NetWare 4.1 Upgrade	Default Rights after NetWare 4.1 Installation or Upgrade
NetWare Server object for the server on which NetWare 4.1 was installed	NetWare Server object for the server you upgraded to NetWare 4.1	User object ADMIN has the Supervisor object right to the NetWare Server object, which means that User object ADMIN also has the Supervisor right to the root directory of the file system of all volumes attached to the server.
Volume object SYS	Volume object SYS for volume SYS: on the server you upgraded	The container object for Volume object SYS is granted Read and File Scan rights to the volume's SYS:PUBLIC directory. This means that when users are created, they can access utilities located in the SYS:PUBLIC directory.
Volume objects for any other volumes on the server's disk besides SYS: that you created during installation	Volume objects for every other volume besides SYS: on the server you upgraded	[Root] is granted the Read property right to the Host Server Name and Host Resource properties on all Volume objects. This means that all objects in the Directory tree have access to the physical volume name and the physical server name. User object ADMIN has the Supervisor right to the root directory of the file systems on these volumes.

Table 1-1 *continued*

Default Objects and Rights in a Directory Tree

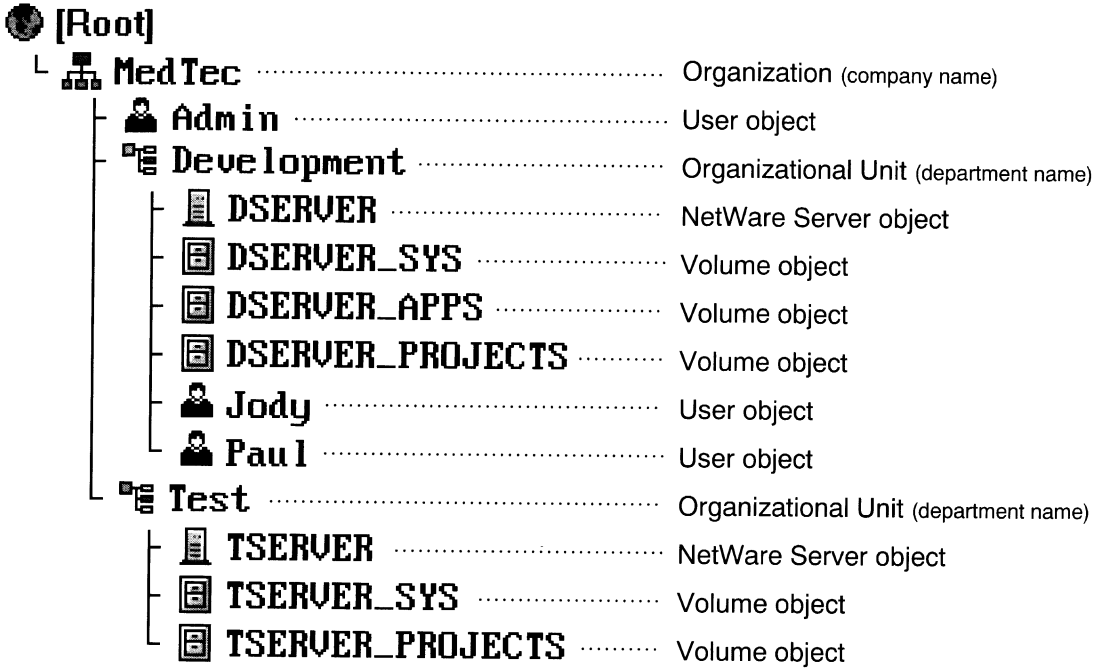
Default Objects after NetWare 4.1 Installation	Default Objects after NetWare 4.1 Upgrade	Default Rights after NetWare 4.1 Installation or Upgrade
<p>User object ADMIN</p> <p>When User object ADMIN is first created, by default it is placed in the Organization container object. This may not be the same context in which you installed the server.</p>	<p>User object ADMIN</p> <p>When User object ADMIN is first created, by default it is placed in the Organization container object. This may not be the same context in which you installed the server.</p> <p>All other objects that were on the server's bindery are placed in the same container as the server that was upgraded.</p>	<p>User object ADMIN receives a trustee assignment of Supervisor object right to the Root object of the Directory tree. This means that User object ADMIN has all rights to all the objects in the Directory tree.</p> <p>User object ADMIN also has the Supervisor object right to the NetWare Server object, which means that User object ADMIN has the Supervisor right to the root directory of the file system of all volumes attached to the server.</p> <p>[Public] has Browse object right at the root. This means that users who are attached to a NetWare Directory Services™ server, but who are not authenticated, can browse the tree.</p> <p>The network supervisor can remove [Public] browse rights so that users cannot do this.</p> <p>Users are granted all rights to their home directories, which are migrated when you migrate User objects.</p>



The upgrade utility converts most existing bindery objects to NetWare Directory Services (NDS) objects. For more information, see *Upgrade*.

Figure 1-1 shows objects in a Directory tree after you have upgraded one server (DSERVER) and installed a NetWare 4.1 server (TSERVER).

Figure 1-1
**Example of a New
 Directory Tree**



DSERVER was upgraded from NetWare 3.11 and placed in the context DEVELOPMENT.MEDTEC (that is, in Organizational Unit DEVELOPMENT, which is in Organization MEDTEC).

The volumes and users on the NetWare 3.11 server (DSERVER_SYS, DSERVER_APPS, DSERVER_PROJECTS, JODY, and PAUL) were converted into Directory objects when the server was upgraded to NetWare 4.1.

TSERVER is a newly installed NetWare 4.1 server that was placed in the context TEST.MEDTEC (that is, in Organizational Unit TEST, which is in Organization MEDTEC). Volume object TSERVER_SYS was created by default, and Volume object TSERVER_PROJECTS was defined by the network supervisor during installation.

Setting Up Administration Utilities

The first time the network supervisor logs in, it must be as User object ADMIN. This is the only object that has rights to create and manage objects immediately after installation or upgrade.

However, before you can log in, you must first use the instructions in this section to install a single client workstation, from which you can run either the NetWare Administrator or NETADMIN utility to start creating objects on your network.

You can install and customize other workstations later, using the appropriate client manual (*NetWare Client for DOS and MS Windows User Guide* or *NetWare Client for OS/2 User Guide*).


If you have installed the NetWare server from CD-ROM, you also need to create installation diskettes for NetWare clients (DOS, MS Windows, and OS/2) before you can install a workstation. Then proceed with the workstation operation.

If you want to	Go to
Create workstation installation diskettes from CD-ROM	"Creating Client Install Diskettes from CD-ROM" on page 6
Install an MS Windows workstation	"Installing an MS Windows Workstation and Starting NetWare Administrator" on page 9
Install an OS/2 workstation	"Installing an OS/2 Workstation and Starting NetWare Administrator" on page 12
Install a DOS workstation	"Installing a DOS Workstation and Starting NETADMIN" on page 15

Creating Client Install Diskettes from CD-ROM

You can install the CD-ROM either as a DOS volume or as a network device.

Procedure



1. For DOS or MS Windows, format five high-density diskettes using the DOS FORMAT command. For OS/2, format seven diskettes.
2. (DOS volume only) Install the CD-ROM drive as a DOS device according to manufacturer's instructions.
 - 2a. Go to the drive corresponding to the CD-ROM.
 - 2b. Change to the CLIENT directory.
 - 2c. Skip to Step 4.
3. (Network device only) Map a drive to a server directory containing the NetWare 4.1 files.
 - 3a. Change to the SYS:PUBLIC\CLIENT directory.
 - 3b. Continue with Step 4.
4. Change to one of the following subdirectories.

If creating diskettes for	Go to this subdirectory
DOS	DOSWIN
MS Windows	DOSWIN
OS/2	OS2

5. If you are installing a language version other than English (the default language), set the “nwlanguage” environment variable by typing the following at the workstation command line:

```
SET NWLANGUAGE=language <Enter>
```

Replace *language* with the appropriate language found in the NLS subdirectory under either CLIENT\OS2 or CLIENT\DOSWIN.

6. Type

```
MAKEDISK drive_letter: <Enter>
```

Replace *drive_letter* with the letter representing the drive you will insert the blank formatted diskettes into. For example, type either

```
MAKEDISK A: <Enter>
```

or

```
MAKEDISK B: <Enter>
```

The MAKEDISK utility copies the client installation files from the CD-ROM directory to the diskettes. It prompts you to insert new blank formatted diskettes.

7. Attach a label to each diskette, and then label the diskettes as described in the following table.

**Table 1-2
Client Diskette Labels**

Diskette #	Label for DOS/MS Windows	Label for OS/2
1	<i>NetWare Client for DOS and MS Windows Disk 1</i>	<i>WSOS2_1</i>
2	<i>NetWare Client for DOS and MS Windows Disk 2</i>	<i>WSOS2_2</i>
3	<i>NetWare Client for DOS and MS Windows Disk 3</i>	<i>WSOS2_3</i>
4	<i>NetWare Client for DOS and MS Windows Disk 4</i>	<i>OSUTIL1</i>
5	<i>NetWare Client for DOS and MS Windows ODI LAN drivers</i>	<i>OS2DOC_X</i> 1=English 2=French 3=German 4=Italian 5=Spanish
6		<i>WSDRV_1</i>
7		<i>VLMBOOT</i> <i>(optional)</i>

8. Save these diskettes until you are ready to install the client software.

Installing an MS Windows Workstation and Starting NetWare Administrator

If you want to use the NetWare Administrator graphical utility on an MS Windows workstation to set up your network, follow the instructions in this section.


Prerequisites



Checklist

- A workstation cabled to the network and running Microsoft* Windows (MS Windows) 3.1
- At least one NetWare 4.1 server installed
- The *WSDOS_1*, *WSWIN_1*, *WSDRV_1*, and *WSDRV_2* diskettes
- 1.2 MB of free disk space on the workstation

Procedure



Procedure

1. Insert the *WSDOS_1* diskette into your workstation's disk drive.

2. Load the client installation program by typing

```
drive:install <Enter>
```

Replace *drive* with the letter of your workstation disk drive.

3. Follow the installation instructions on the workstation screen.



Note

Use the arrow keys to move between the fields. Press <Enter> to modify a particular field.

Because this workstation is being installed to get you started with the network setup, you can do a standard installation with the default settings now and configure other options later.

4. Exit the installation program.

5. Reboot the workstation.

In order for your modifications or new installation to function, you need to reboot the workstation.

- 6. From network drive F:, use the LOGIN command to log in to the network as ADMIN.**

The LOGIN command must include your context. Depending on the Directory tree structure you set up for the server during installation, use the sequence of Common Name, Organizational Unit, Organization, and Country.

For example, if this server were installed in an Organization called ABC_INC, with no Organizational Unit or Country defined in the Directory tree, you would type

```
LOGIN .ADMIN.ABC_INC <Enter>
```

For more information on the LOGIN syntax and options, see "LOGIN" in *Utilities Reference*.

- 7. Enter a password if prompted; then press <Enter>.**
- 8. Map the next network drive to the PUBLIC directory of volume SYS: by typing**

```
MAP N SYS:PUBLIC <Enter>
```

- 9. Change to the network drive that is mapped to the PUBLIC subdirectory.**

For example, if you mapped network drive Z:, you'd type

```
z: <Enter>
```

- 10. Start MS Windows.**
- 11. Choose the NetWare program group or, if you prefer, choose another program group you want to start NetWare Administrator from.**
- 12. From the Program Manager "File" menu, choose "New."**
- 13. Select "Program Item."**
- 14. Choose "OK."**
- 15. Enter "NWADMIN" in the "Description" field and press <Tab>.**

16. Choose “BROWSE.”

17. From the “Drives” drop-down list, select the drive that points to SYS:PUBLIC.

18. From the list under “File Name,” select NWADMIN.EXE and choose “OK.”

The path to the executable file is placed in the “Command Line” text box.

19. Again, choose “OK,” and then choose “Yes.”

NetWare Administrator is created as a program item icon in the group you selected.

You can now use the NetWare Administrator graphical utility to complete your network setup.

20. To start NetWare Administrator, choose the “NWADMIN” icon.

21. Before exiting Windows, from the “Options” menu, select “Save Settings on Exit.”

Additional Information

For more information about	Refer to
Installing MS Windows workstation software	<i>NetWare Client for DOS and MS Windows User Guide</i>
Logging in	“LOGIN” in <i>Utilities Reference</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Installing an OS/2 Workstation and Starting NetWare Administrator


If you want to use the NetWare Administrator graphical utility on an OS/2 workstation to set up your network, follow the instructions in this section.

Prerequisites



- A workstation cabled to the network and running OS/2* v2.x and Microsoft Windows or WIN-OS/2.
- At least one NetWare 4.1 server installed
- The *WSOS2_1*, *WSOS2_2*, *WSOS2_3*, *OSUTIL_1*, *OS2DOC_X*, *VLMBOOT*, and *WSDRV_1* diskettes

Procedure



1. **Start the OS/2 workstation and open an OS/2 full screen or window.**
2. **Insert the *WSOS2_1* diskette into your workstation's disk drive.**
3. **Change to the drive where you inserted the diskette and type**
INSTALL <Enter>
4. **Select "Requester on Client" from the "Installation" menu.**
5. **Follow the online help and the screen prompts to finish installing the Requester.**

See *NetWare Client for OS/2 User Guide* if you need more information.

Because this workstation is being installed to get you started with the network setup, you can do a standard installation with the default settings now and configure other options later.

On the drive you booted OS/2 from, a directory called NETWORKE is created for the Requester files during installation. A program group called Novell® is also created on the OS/2 desktop.

6. Install VLM Boot.

For instructions, see “Installing and Using NetWare VLM Boot” in Chapter 3 of *NetWare Client for OS/2 User Guide*.

7. Exit VLM Boot installation.

8. Remove the diskette from the drive; then select the OS/2 Shutdown feature and reboot the workstation.

9. After your workstation reboots, from the Novell folder launch NetWare VLM Boot by choosing the “VLMBOOT” icon.

10. From network drive F:, use the LOGIN command to log in to the network as ADMIN.

The LOGIN command must include your context. Depending on the Directory tree structure you set up for the server during installation, use the sequence of Common Name, Organizational Unit, Organization, and Country.

For example, if this server were installed in an Organization called ABC_INC, with no Organizational Unit or Country defined in the Directory tree, you would type

```
LOGIN .ADMIN.ABC_INC <Enter>
```

For more information on the LOGIN syntax and options, see “LOGIN” in *Utilities Reference*.

11. Enter a password if prompted; then press <Enter>.

12. Map the next network drive to the PUBLIC directory of volume SYS: by typing

```
MAP N SYS:PUBLIC <Enter>
```

13. Change to the network drive that is mapped to the PUBLIC subdirectory.

For example, if you mapped network drive Z:, you’d type

```
z: <Enter>
```

14. Start MS Windows or WIN-OS/2.

15. Choose the NetWare program group or, if you prefer, choose another program group you want to start NetWare Administrator from.
16. From the Program Manager “File” menu, choose “New.”
17. Select “Program Item” and choose “OK.”
18. Enter “NWADMIN” in the “Description” field and press <Tab>.
19. Choose “BROWSE.”
20. From the “Drives” drop-down list, select the drive that points to SYS:PUBLIC.
21. From the list under “File Name,” select NWADMIN.EXE and choose “OK.”

The path to the executable file is placed in the “Command Line” text box.

22. Again, choose “OK,” and then choose “Yes.”

NetWare Administrator is created as a program item icon in the group you selected.

You can now use the NetWare Administrator graphical utility to complete your network setup.

23. To start NetWare Administrator, choose the “NWADMIN” icon.
24. Before exiting Windows, from the “Options” menu, select “Save Settings on Exit.”

Additional Information

For more information about	Refer to
Installing OS/2 workstation software	<i>NetWare Client for OS/2 User Guide</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>
Logging in	“LOGIN” in <i>Utilities Reference</i>

Installing a DOS Workstation and Starting NETADMIN


If you want to use text utilities, such as NETADMIN, to set up your network, follow these instructions to install DOS workstation software.

Prerequisites



- A workstation cabled to the network and running DOS 3.30 or later
- At least one NetWare 4.1 server installed
- The *WSDOS_1*, *WSWIN_1*, *WSDRV_1*, and *WSDRV_2* diskettes
- 1.2 MB of free disk space on the workstation

Procedure



1. Insert the *WSDOS_1* diskette into the workstation's disk drive.

2. Load the *INSTALL.EXE* program by typing

```
drive: INSTALL <Enter>
```

Replace *drive* with the drive letter of your disk drive.

3. Follow the instructions on your screen.



Use the arrow keys to move between the fields. Press <Enter> to modify a particular field.

Because this workstation is being installed to get you started with the network setup, you can do a standard installation with the default settings now and configure other options later.

4. Exit the installation program.

5. Reboot your workstation.

In order for your modifications or new installation to function, you need to reboot your workstation.

6. From network drive F:, use the LOGIN command to log in to the network as ADMIN.

The LOGIN command must include your context. Depending on the Directory tree structure you set up for the server during installation, use the sequence of Common Name, Organizational Unit, Organization, and Country.

For example, if this server were installed in an Organization called ABC_INC, with no Organizational Unit or Country defined in the Directory tree, you would type

```
LOGIN .ADMIN.ABC_INC <Enter>
```

For more information on the LOGIN syntax and options, see "LOGIN" in *Utilities Reference*.

7. Enter a password if prompted; then press <Enter>.

8. Map the next network drive to the PUBLIC directory of volume SYS: by typing

```
MAP N SYS:PUBLIC <Enter>
```

9. Change to the network drive that is mapped to the PUBLIC subdirectory.

For example, if you mapped network drive Z:, you would type

```
z: <Enter>
```

You can now use text utilities such as NETADMIN to complete your network setup.

10. To start NETADMIN, type

```
NETADMIN <Enter>
```

Additional Information

For more information about	Refer to
Installing DOS workstation software	<i>NetWare Client for DOS and MS Windows User Guide</i>
Logging in	“LOGIN” in <i>Utilities Reference</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Rights Needed to Create and Manage Objects

As User object ADMIN, you have all rights to all objects on the Directory tree. However, to allow various parts of the Directory tree to be managed by other users, you need to give those users the rights necessary for them to manage their sections of the tree.

When you give an object, such as a User, rights to another object, such as a container, you have made a *trustee assignment*. That User is now a *trustee* of that container. An object with a trustee assignment to another object is a *trustee* of that object.

Each object contains a list of trustee assignments called a *trustee list*. This list tells who can access that object. An object’s trustee list is stored in its Access Control List (ACL) property.

Four kinds of rights exist in NetWare 4.1:

- ◆ **Object rights:** Control what a trustee can do with an object. These rights control the object as a single piece in the Directory tree, but don’t allow access to information stored within that object (unless the Supervisor object right is granted).
- ◆ **Property rights:** Control a trustee’s access to information stored within the object—that is, the information stored in the object’s properties. Each object has several properties.

- ◆ **Directory rights:** Control what a trustee can do with a directory. Directory rights also apply to files in the directory only if file rights aren't granted and the file's Inherited Rights Filter doesn't block the directory rights.
- ◆ **File rights:** Control what a trustee can do with a file.

In previous versions of NetWare, you could assign directory and file rights. In NetWare 4.1, you can also assign rights to an object and to properties belonging to an object.

This section covers only object rights and property rights.

Directory rights and file rights apply only to the file system. For a discussion of these rights, see "Loading Operating Systems and Applications onto the Network" on page 117.

Object and property rights are assigned separately so that you can control access to the pieces of information (or properties) contained in the object.

Any object to which you grant sufficient rights can make trustee assignments, using the NetWare Administrator or NETADMIN utility.

Object Rights

Object rights control what trustees can do with the object that they are trustees of. Object rights control the object as a single piece in the Directory tree, but do not allow the trustee to access information stored in that object's properties (unless the Supervisor object right is granted).

Table 1-3 lists and describes the object rights that you can assign to a trustee.

Table 1-3
Object Rights

Right	Description
Supervisor	<p>Gives you all rights to the object and to all its properties. However, the Supervisor object right <i>can</i> be blocked by the Inherited Rights Filter (IRF) below the object where the Supervisor right is granted.</p> <p>(An IRF is a list of rights that can be created for any object. It controls the rights that a trustee can inherit from container objects.)</p>
Browse	<p>Allows you to see the object in the Directory tree. Also, when you perform a search for a value that matches the object, the Browse right to an object allows that object to be listed.</p>
Create	<p>Allows you to create a new object within a container object in the Directory tree. This right applies only to container objects because leaf objects cannot contain other objects.</p>
Delete	<p>Allows you to delete an object from the Directory tree. However, you cannot delete a container object unless all the objects in the container are deleted first. See “Moving Container Objects Using NETADMIN” on page 83, or “Moving Container Objects Using NetWare Administrator” on page 81.</p>
Rename	<p>Allows you to change the name of the object, in effect changing the naming property. This changes what the object is called when it’s a part of complete names. See “Renaming Leaf and Container Objects” on page 90.</p>

Property Rights

While object rights allow you to see the object, delete the object, create a new object, etc., they do not allow you to see the information stored in the object's properties.

To read the information in an object's properties, you must have the correct property rights. Property rights control access to each property in an object.

For example, if you include a private telephone number as a property for a User object, you can use property rights to prevent anyone else from seeing that telephone number. At the same time, you can allow other properties, such as Address or Fax Number, to be viewed.

Table 1-4 lists and describes property rights that you can assign to a trustee.

Table 1-4
Property Rights

Right	Description
Supervisor	Gives you all rights to the property. You can block the Supervisor property right with an Inherited Rights Filter. See “Security” in <i>Concepts</i> .
Compare	Allows you to compare any value to an existing value of the property. The comparison can return True or False, but cannot give the value of the property.
Read	Allows you to read the values of the property. This right includes the Compare right; that is, if the Read right is given, Compare operations are also allowed.
Write	Allows you to add, change, or remove any values of the property. The Write right includes the Add or Delete Self right The Write right to the ACL property is the same as giving the Supervisor right to the object.
Add or Delete Self	Allows you to add or remove yourself as a value of the property, but you cannot change any other values of the property. This right is only used for properties where your User object can be listed as a value, such as group membership lists or mailing lists. The Write right includes the Add or Delete Self right.

See the following references for more detailed information.

Additional Information

For more information about	Refer to
Access Control List	“Access Control List” in <i>Concepts</i>
Container and leaf objects	“Object” in <i>Concepts</i>
Trustees and rights	“Security” in <i>Concepts</i>

Managing Trustee Assignments to Objects

An object granted rights to work with another object is called a *trustee* of the object.

Through *trustee assignments*, you determine what level of rights you want trustees to have on the objects they are assigned to.

For example, to make user KSMITH a trustee of Organizational Unit MARKETING, go to the Organizational Unit MARKETING and add KSMITH as a trustee, giving KSMITH whatever object and property rights you think are sufficient.

KSMITH's trustee assignment to MARKETING overrides any other rights that KSMITH may have inherited or received through security equivalence.

You can manage trustee assignments to objects using either the NetWare Administrator graphical utility or the NETADMIN text utility. Both procedures are documented in this section.

For more information about	Refer to
Inherited Rights Filter	"Inherited Rights Filter" in <i>Concepts</i>
Rights	"Rights" in <i>Concepts</i>
Security	"Security" in <i>Concepts</i>
Trustees	"Trustee" in <i>Concepts</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>

Using NetWare Administrator to Manage Trustee Assignments

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Supervisor object right to the object to which trustees are to be assigned

Procedure



1. **From the MS Windows Program Manager or the OS/2 desktop, choose the “NetWare Administrator” icon.**
2. **Select the object you want to assign trustees to.**
3. **Click once on the right button of your mouse and choose “Trustees of this Object,” or, from the “Object” menu, choose “Trustees of this Object.”**

All the current trustees of this object are listed in the “Trustees” box.

From the “Trustees” dialog, you can

- ◆ View or change a trustee’s effective rights to this object
 - ◆ Add a trustee to this object
 - ◆ Delete a trustee from this object
 - ◆ Change the object rights of a trustee
 - ◆ Change the property rights of a trustee
 - ◆ View or change the Inherited Rights Filter (IRF) of this object
4. **For instructions on completing any of the above procedures, press the “Help” button and then select one of the procedures from the list.**


Using NETADMIN to Manage Trustee Assignments

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Supervisor object right to the object to which trustees are to be assigned

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Browse the Directory tree until you see the object you want to assign trustees to.

Use the instructions at the bottom of the screen to browse the directory. Press <F1> if you need help.

4. When the desired object appears in the “Object, Class ” list, select it and press <F10>.

5. Select “View or Edit the Trustees of This Object.”

From the “Trustees of This Object” menu you can

- ◆ Set the “Inherited Rights Filter” (IRF), which affects all trustees that are not explicitly assigned to this object.
- ◆ Add or change trustee assignments (which override the IRF)
- ◆ View the “Effective Rights” for a trustee

6. From the “Trustees of This Object” menu, select one of the options.

7. For instructions on completing a procedure, press <F1>.

Creating Container Objects

You can create container objects by using either the NetWare Administrator graphical utility or the NETADMIN text utility. Each of these methods is explained in this section.

Considerations for naming container objects and suggestions for creating searchable objects are also covered.

Types of Container Objects

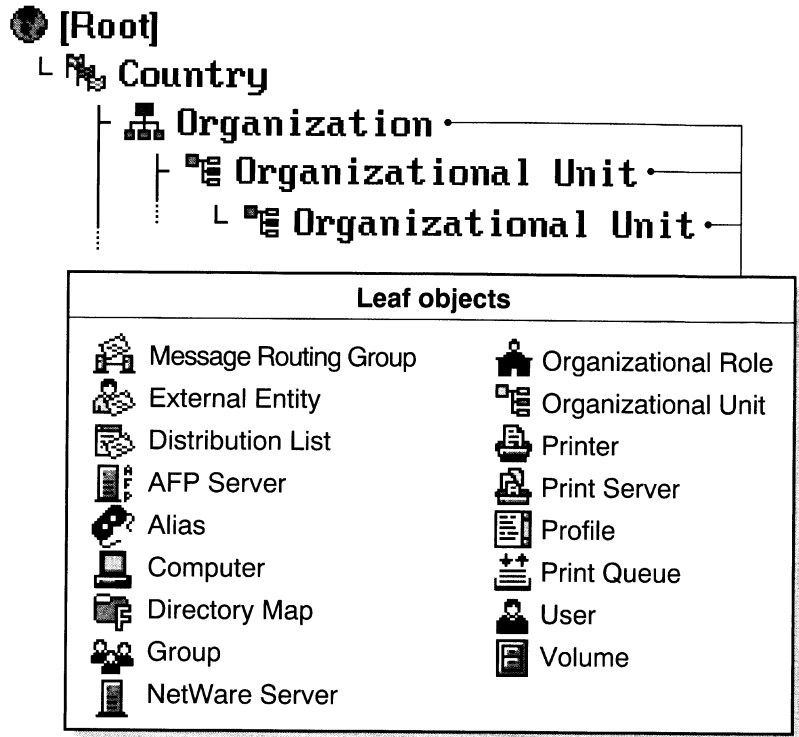
The kinds of container objects you can create are Country, Organization, and Organizational Unit. The top object, called Root, is created by default and is placed at the top of the Directory tree when NetWare 4.1 is installed.

Container objects form the top levels of the Directory tree. Use them to manage and organize the Directory by relating groups of objects, both container objects and leaf objects.

For more information about planning the top levels of your Directory tree, see Chapter 5, "Planning NetWare Directory Services Implementation," of *Introduction to NetWare Directory Services*.

Figure 1-2 illustrates the hierarchy of container objects and leaf objects in NetWare Directory Services. (The icons represent the leaf objects as they appear in the NetWare Administrator graphical utility.)

Figure 1-2
Hierarchy of
Objects



Note The first three leaf objects in the figure—Message Routing Group, External Entity, and Distribution List—are NetWare Message Handling Service™ (NetWare MHS) objects. They appear in NetWare Administrator only if you have installed NetWare MHS™ Services on your NetWare server.

You can create leaf objects only under the Organization and Organizational Unit container objects.

Table 1-5 describes each type of container object you can create and when to use it.

**Table 1-5
Container Objects You Can Create**

Container Object	Description	When to Use It
Country	<p>Designates the countries where your network resides and organizes other objects within the country.</p> <p>You must always include the name type of the object in complete names when you include the Country container object in your Directory tree. Even when you refer to objects located in the same container object, you must designate the name type (CN, OU, or O) of the object.</p>	<p>This object is optional. You do not have to create a Country object.</p> <p>If you choose to create a Country object, you can use it to represent the country where your organization headquarters reside or, if you have a multinational network, to represent each country that is a part of your network.</p> <p>You can create a Country container object only under the Root object.</p>
Organization	<p>Allows you to organize other objects in the Directory, to set defaults in a login script, and to create a user template for User objects you create in this container.</p> <p>For example, you can use an Organization object to designate a corporation.</p>	<p>The Organization container object is mandatory. The Directory tree must contain at least one.</p> <p>You can create an Organization object only under the Root or Country object.</p>
Organizational Unit	<p>Allows you to organize leaf objects in the Directory tree, to set defaults in a login script, and to create a user template for User objects you create in this container.</p>	<p>You can use an Organizational Unit object to designate a division, a business unit, or a project team.</p> <p>You can create multiple levels of Organizational Units.</p> <p>You can create Organizational Units in Organization objects and in other Organizational Unit objects.</p>

Naming Container Objects

Try to keep container object names short and simple. This makes it easier for users to change context and to remember their own context.

The following rules apply to most objects. For specific rules about naming leaf objects, see “Naming Leaf Objects” on page 41.

Object Naming Rules

- ◆ The name must be unique in the branch (container) of the Directory tree where the object is located.
- ◆ The object name can be up to 64 characters in length.
- ◆ You can use any special characters. But if the object needs to be accessed from a client running a version of NetWare earlier than NetWare 4™, you should avoid using special characters. (For a list of these characters, see “Object Name Restrictions for Bindery Services” on page 29.)
- ◆ You can enter object names in either upper or lowercase. Object names are displayed with uppercase and lowercase letters as they were first entered, but they are not case sensitive. Therefore, “ManagerProfile” and “MANAGERPROFILE” are considered to be identical names.
- ◆ You can use both spaces and underscores, but they are both displayed as spaces. Therefore, “Manager_Profile” and “Manager Profile” are considered to be identical names.
- ◆ Country objects can have only a two-character name.



If you anticipate managing objects created from different code pages, you must limit object names and properties to those characters common to all the applicable code tables.

Nondisplayable Unicode* characters for your code page are represented by an ASCII 3 character (a “heart” symbol). For more information, see “Unicode” in *Concepts*.

Object Name Restrictions for Bindery Services

When you create objects to be accessed from a client running a version of NetWare earlier than NetWare 4, the names of the objects must follow bindery naming rules or the non-NetWare 4.1 client will not recognize them. Object names in bindery services are interpreted as the following:

- ◆ Spaces in object names are replaced by underscores.
- ◆ Object names longer than 47 characters are cut off after the 47th character.

You cannot use the following characters in an object name that must be accessed from a client running a version of NetWare earlier than NetWare 4:

- / Slash
- \ Backslash
- : Colon
- , Comma
- * Asterisk
- ? Question mark

Creating Searchable Container Objects

When you create a container object, you can enter various types of information about that object into its properties, such as location and telephone number. If you enter data into the containers' properties in a consistent format, it is easier to search the Directory database for a particular type of information.

Many container object properties are optional; you are not required to enter information about the property in order to create the object. However, information in objects' properties can help you track and manage container objects.

After you have created container objects, you can use the NetWare Administrator, NETADMIN, or NLIST utility to search for and list these objects. You can also search for their various properties.

Creating Container Objects Using INSTALL

When you install NetWare 4.1 on a server, you are required to type a context (the pathname from the container object to the Root object) in which the NetWare Server object is placed.

If you create a new context, several events happen by default:

- ◆ An Organization or Organizational Unit container object is created, depending on the context you create.
- ◆ Bindery services are set for that container object, so that the server you installed in the container object is running bindery services.
- ◆ A read/write replica of the Directory partition that the container object is in is stored on the server that you just installed.

For more information on how to create container objects using the installation process, see Chapter 2, "Simple Installation," or Chapter 3, "Custom Installation," of *Installation*.


Creating Container Objects Using NetWare Administrator

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Create object right to the container that will contain the new container object

Procedure



- 1. From the MS Windows Program Manager or the OS/2 desktop, choose the "NetWare Administrator" icon.**
- 2. Select the object that will contain the new container object.**

For information on moving around in the browser and choosing objects, press <F1>.

- 3. Choose "Create" from the "Object" menu.**

4. Select the new container object class that you want from the “New Object” dialog.

If the container object class you want to create does not appear under “New Object,” you cannot create that object in the selected container. Choose “Cancel” to return to the browser; then select a different container type.

5. Choose “OK.”

The “Create Object” dialog appears.

6. Type a name for the object in the box provided.

7. (Optional) Select “Define Additional Properties.”

Select this option if you want to enter more information for the new container object.

Examples of additional properties are Login Script, Intruder Detection, Postal Address, and See Also. The Login Script and Intruder Limit apply to User objects in the new container object.

8. (Optional) Select “Define User Defaults.”

Select this option if you want to use the same default information in the new container as was present in the parent container. This default information is used whenever you create a new user.

The user default information for each container is actually stored in a User object named USER_TEMPLATE.

9. Choose “Create.”

10. (Optional) Choose “Yes” if you want to inherit user template properties from the parent container, or “No” if you want to define a new user template.

11. (Optional) Add information to the object dialog pages.

If you chose “Define Additional Properties,” add the information now. Press <F1> to get help on each field.

12. Choose “OK” to save the properties you have just entered in the dialog pages.

Additional Information

For more information about	Refer to
Objects	"Object" in <i>Concepts</i>
Rights	"Rights" in <i>Concepts</i>
User templates	"Managing User Templates" on page 67
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>
Using object dialog pages in NetWare Administrator	"NetWare Administrator" in <i>Utilities Reference</i>


Creating Container Objects Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later
- The Create object right to the object that will contain the new container object

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose "Manage Objects" from the "NETADMIN Options" menu.

3. Select the container object that will contain the new container object.

The objects in the selected container are listed.

To see if you're in the right context, look at the title bar on the screen. Press <F1> if you need help.

4. Press <Insert>.

5. Select the container object class that you want to create from the "Select an Object Class" screen.

If the container object class you want to create does not appear, you cannot create that object in the selected container. Press <Esc> to return to the browser; then select a different container type.

6. Type the new container object name.

7. Enter a Mailbox Location and press <Enter>.



If you are creating a Country object, you are not prompted to define a Mailbox Location or create a user template.

8. If you want to create a user template to be applied to new User objects created in this container, type "Y" and press <Enter>.

A user template contains default information that you can apply to users that you create to give them default property values.

The user template is actually a User object named USER_TEMPLATE.

If you do not want to create a user template, type "N" and press <Enter>.

9. Press <F10> to save the information.

- 10. If you want to create another container object, choose “Yes.” If you do not, choose “No” and then press <Enter>.**

If you chose “Yes,” you are prompted to type the new container object name. Repeat Steps 6 through 9, and then continue with Step 12.

If you chose No, then the container object is displayed in the Directory tree. Continue with Step 12.

- 11. To edit this object, press <F10>.**

A menu appears from which you can choose to view or edit properties of this object and make trustee assignments to this object and to files and directories.

- 12. Choose an option from the “Actions” menu and add any necessary information.**

- 13. To exit, press <Esc> until you return to the “NETADMIN Options” menu.**

Additional Information


















For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Object properties	“Property” in <i>Concepts</i>
User templates	“Managing User Templates” on page 67
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Creating Leaf Objects

Leaf objects represent network resources, such as users, computers, printers, and lists. They do not contain any other objects.

You create leaf objects within a container object. Figure 1-3 lists the leaf objects you can create. (The icons represent the leaf objects as they appear in the NetWare Administrator graphical utility.)

Figure 1-3
Leaf Objects You
Can Create

Leaf objects	
 Message Routing Group	 Organizational Role
 External Entity	 Organizational Unit
 Distribution List	 Printer
 AFP Server	 Print Server
 Alias	 Profile
 Computer	 Print Queue
 Directory Map	 User
 Group	 Volume
 NetWare Server	



Note The first three leaf objects in the figure—Message Routing Group, External Entity, and Distribution List—are NetWare MHS Services objects. They appear in NetWare Administrator only if you have installed NetWare MHS Services on your NetWare server.

How to Use Leaf Objects

Table 1-6 describes each leaf object you can create and when to use it.

Table 1-6
Leaf Objects You Can Create

Leaf Object	Description	When to Use
AFP Server	Represents an AppleTalk* Filing Protocol-based server that is operating as a node on your NetWare network (and possibly also acting as a NetWare router to, and the AppleTalk* server for, several Apple* Macintosh* workstations).	<p>Create this object when you have an AFP server that you need to represent on the network. Use it to store information about this server, such as its description, location, and network address.</p> <p>This object has no effect on the operation of the network; it only stores information about the AFP server.</p>
Alias	<p>Points to another object in the Directory tree and makes it appear as if the object that it points to actually exists in the Directory tree where the Alias object is.</p> <p>Although an object appears both where it was actually created and where an Alias referring to it was created, only one copy of the object really exists.</p> <p>If you delete or rename an Alias, the Alias itself (not the object it is pointing to) is deleted or renamed.</p>	<p>Use this object to allow access to an object that is in another context.</p> <p>For example, you can use an Alias to represent a resource, such as a special printer, that most users in the tree need to access.</p> <p>Also, when you move or rename a container object in a Directory tree, you have the option of creating an alias in place of the moved or renamed object. If you select this option, NetWare Administrator automatically creates the alias for you and assigns it the same name as the original object.</p> <p>Creating an alias in place of a moved or renamed container object allows users to continue logging into the network and to see the container object (and the objects it contains) in its original Directory location</p>

Table 1-6 *continued*

Leaf Objects You Can Create

Leaf Object	Description	When to Use
Computer	Represents a nonserver computer on the network, such as a workstation or a router.	<p>Use this object to store information about a nonserver computer, such as its network address, or serial number, or the person the computer is assigned to.</p> <p>This object has no effect on the operation of the network; it only stores information about the computer.</p>
Directory Map	<p>Represents a particular directory in the file system. Directory Map objects can be especially useful in login scripts by pointing to directories that contain applications or other frequently used files.</p> <p>For example, if you have a directory that contains DOS 5.0, you will probably map a search drive to that directory in any login scripts you create. Later, if you upgrade to DOS 6.0 and rename the directory, you would have to change the mapping in every login script where that search mapping appears.</p> <p>With a Directory Map object, you only change the information in that one object.</p>	<p>Use this object to avoid making changes to many login scripts when the location of applications changes; instead, you change only the Directory Map object.</p> <p>For more information on Directory Map objects, see “Loading Operating Systems and Applications onto the Network” on page 117.</p>
Distribution List	Represents a list of mail recipients.	<p>Use this object to simplify sending mail to recipients.</p> <p>For example, you could create a Distribution List object called Recreation Committee. Anyone wanting to send a message to all the members in the Recreation Committee can simply address the message to Recreation Committee, rather to each member.</p>

Table 1-6 *continued*

Leaf Objects You Can Create

Leaf Object	Description	When to Use
External Entity	<p>Represents a nonnative NDS™ (NetWare Directory Services) object that is imported into NDS or registered in NDS.</p> <p>NetWare MHS Services use External Entity objects to represent users from non-NDS directories to provide an integrated address book for sending mail.</p>	<p>If your messaging environment contains non-MHS servers, such as SMTP hosts, SNADS nodes, or X.400 MTAs, you might choose to add users and lists at these servers to your NetWare database as External Entities.</p> <p>Adding these objects to the database as External Entities adds them to the address books of your messaging applications. When addressing messages, local users can choose non-MHS users and lists from a directory list.</p>
Group	<p>Assigns a name to a list of User objects that can be located anywhere in the Directory tree.</p>	<p>Create a Group when you have many User objects that need the same trustee assignments. Rather than making many trustee assignments, you make just one trustee assignment to all the users who belong to the group, by making the trustee assignment to the Group object itself.</p>
Message Routing Group	<p>Represents a group of messaging servers that can transfer messages directly with each other.</p>	<p>Create a Message Routing Group when you have two or more messaging servers that need to communicate with each other.</p>
Messaging Server	<p>Represents a messaging server that resides on a NetWare server. A Messaging Server object is automatically created in the NDS tree when you install NetWare MHS Services on a NetWare server.</p>	<p>Create a Messaging Server (by installing NetWare MHS Services on a NetWare server) when you need a server to handle messaging between users and groups on the network.</p>

Table 1-6 *continued*

Leaf Objects You Can Create

Leaf Object	Description	When to Use
NetWare Server	<p>Represents a server running NetWare on your network.</p> <p>Store information about the server in the NetWare Server object's properties, such as the server's address, the physical location of the server, and what services it provides.</p> <p>In addition to storing information about the NetWare server, the NetWare Server object affects the network in that it is referred to by several other objects.</p>	<p>Use the NetWare Server object to tie the physical server on the network to the Directory tree. Without this object, you cannot access file systems that are on that server's volumes.</p> <p>If you have a non-NetWare 4.1 server, you must create this object to be able to access those non-NetWare 4.1 volumes. When you create a NetWare Server object for a non-NetWare 4.1 server, the non-NetWare 4.1 server must be running.</p>
Organizational Role	<p>Defines a position or role within an organization.</p>	<p>Create an Organizational Role object so that you can assign rights to a particular position rather than to the person who may occupy that position. The occupant may change frequently, but the responsibilities of that position do not.</p> <p>You can assign any user to be an occupant of the Organizational Role object because every occupant receives the same rights that you granted to the Organizational Role object.</p>
Print Queue	<p>Represents a print queue on the network.</p>	<p>You must create a Print Queue object for every print queue on the network.</p> <p>This object cannot be created with NETADMIN.</p> <p>See <i>Print Services</i> for more information.</p>
Print Server	<p>Represents a network print server.</p>	<p>You must create a Print Server object for every print server on the network.</p> <p>This object cannot be created with NETADMIN.</p> <p>See <i>Print Services</i> for more information.</p>

Table 1-6 *continued*

Leaf Objects You Can Create

Leaf Object	Description	When to Use
Printer	Represents a physical printing device on the network.	<p>You must create a Printer object for every printer on the network.</p> <p>This object cannot be created with NETADMIN.</p> <p>See <i>Print Services</i> for more information.</p>
Profile	Contains a profile script (login script). When the Profile object is listed as a User object's property, the Profile object's login script is executed when that User object logs in. The Profile login script executes after the system login script and before the user login script.	Create a Profile object for a set of users who need to share common login script commands but who are not located in the same container in the Directory tree, or who are a subset of users in the same container.
User	<p>Represents a person who uses your network.</p> <p>In the User object properties, you can set login restrictions, intruder detection limits, password and password restrictions, security equivalences, etc.</p>	<p>You must create a User object for every user who needs to log in to the network. When you create a User object, you can create a home directory for that user who will have default rights to that home directory. When you create User objects, you can also choose to apply a user template to the user that provides default property values.</p> <p>For users who have NetWare 4.1 workstations, you can create the User objects anywhere in the Directory tree, but the users must know their context in order to log in. You should create User objects in the container where the users will typically log in.</p> <p>For users who have non-NetWare 4.1 workstations, you must create the User objects in the container at which the bindery services context is set for the server that they need to log in to. (Bindery services is set by default for every NetWare 4.1 server that is installed.) Non-NetWare 4.1 users do not need to know their context because they log in to the server rather than to the Directory tree.</p>

Table 1-6 *continued*

Leaf Objects You Can Create

Leaf Object	Description	When to Use
Volume	<p>Represents a physical volume on the network.</p> <p>In the Volume object's properties, you can enter identification information, such as the Host server, volume location, etc. You can also set restrictions for use of the volume, such as space limits for users.</p>	<p>You can create a Volume object for every physical volume on the network. (During installation of NetWare 4.1 on a server, Volume objects are created for every physical volume on that server.)</p> <p>When you create a Volume object, you are prompted for the server name and the volume name on the server. That information is placed in the Volume object's properties.</p> <p>You can use the Volume object to display information about the directories and files on that volume.</p>

Naming Leaf Objects


The object naming rules described in the next section apply to most objects. Special rules applying to NetWare Server objects and objects viewed through bindery services are described in separate sections. For rules about naming container objects, see "Naming Container Objects" on page 28.

Object Naming Rules

- ◆ The name must be unique in the branch (container) of the Directory tree where the object is located.
- ◆ The object name can be up to 64 characters in length.
- ◆ You can use any special characters. But if the object needs to be accessed from a client running a version of NetWare earlier than NetWare 4, you should avoid using special characters. (For a list of these characters, see "Object Name Restrictions for Bindery Services" on page 29.)

- ◆ You can enter object names in either upper or lowercase. Object names are displayed with uppercase and lowercase letters as they were first entered, but they are not case sensitive. Therefore, “ManagerProfile” and “MANAGERPROFILE” are considered to be identical names.
- ◆ You can use both spaces and underscores, but they are both displayed as spaces. Therefore, “Manager_Profile” and “Manager Profile” are considered to be identical names.

(If you use a space in a name, you must place quote marks around that text string whenever you use a command line utility that includes that text string.)



If you anticipate managing objects created from different code pages, you must limit object names and properties to those characters common to all the applicable code tables.

Nondisplayable Unicode characters for your code page are represented by an ASCII 3 character (a “heart” symbol). For more information, see “Unicode” in *Concepts*.

Object Name Restrictions for Bindery Services

When you create objects to be accessed from a client running a version of NetWare earlier than NetWare 4, the names of the objects must follow bindery naming rules or the pre-NetWare 4 client does not recognize them. Object names in bindery services are interpreted as the following:

- ◆ Spaces in object names are replaced by underscores.
- ◆ Object names longer than 47 characters are cut off after the 47th character.

You cannot use the following characters in an object name that must be accessed from a client running a version of NetWare earlier than NetWare 4:

- / Slash
- \ Backslash
- : Colon
- , Comma
- * Asterisk
- ? Question mark

Naming Restrictions for NetWare Server Objects

The first NetWare Server object for a NetWare 4.1 server must be created with INSTALL. The object is given the same name as the physical server. Rules for naming physical servers are in the <F1> Help of INSTALL.

If you create a NetWare Server object for a non-NetWare 4.1 server, you must use the physical server name as well, because NetWare Directory Services must search for the server on the network to verify its existence.

For example, if you create a Server object for a NetWare 3™ server whose physical name on the network is SURFBOY, you must name the Server object SURFBOY.

Because of these restrictions, you can never rename a NetWare Server object, even if you have the Supervisor object right to it. For more information on NetWare Server objects, see “Object” in *Concepts*.

Creating Searchable Leaf Objects

When you create an object, you enter various types of information about that object into its properties. An object’s properties can include a telephone number, a description, an address, etc.

Many object properties are optional; you are not required to enter information about such properties to create the object. However, information in objects’ properties can help you track and manage those objects.

After you have created objects, you can use the NetWare Administrator, NETADMIN, or NLIST utility to search for and list these objects. You can also search for the values contained in the objects’ properties.

If you enter data into the properties in a consistent format, it is easier to search the Directory database for different types of information when you need it.

For example, you may want to search for all User objects at a certain location, such as building M1. You cannot easily list all objects located in building M1 if you have entered “Bldg. M1,” “BLDG M1,” and “M1” as values in the Location property of multiple User objects.

Standardizing the value for the Location property for all User objects at the site (such as M1, M2, and M3) makes it possible to search for objects located in each building.


Creating Leaf Objects Using NetWare Administrator

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Create object right to the container that will contain the new leaf object

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Select the object that will contain the new leaf object.**

For information on moving around in the browser and selecting objects, press <F1>.



If you are creating User objects, remember that users who are using non-NetWare 4.1 workstations must be created in the container where the bindery services context is set for the server that they need to log in to.

You can create User objects for users who have NetWare 4.1 workstations anywhere in the Directory tree, but the users must know their context in order to log in. You should create User objects in the container where the users will typically log in.

- 3. Choose “Create” from the “Object” menu.**
- 4. From the “New Object” dialog, select the class of object you want to create.**

If the class of object you selected does not appear under “New Object,” you cannot create this object in this container. Select or create another container to hold the object.

- 5. Choose “OK.”**

Each type of leaf object that you create has a different “Create” dialog. For details on each dialog, choose “Help.”

The property fields that are displayed in the “Create” dialog are mandatory. You must enter information in these fields.

The check boxes that are displayed in the “Create” dialog are optional. Usually, you can select only one box, not both.

For example, if you select “Define Additional Properties,” the “Identification” page is displayed immediately after the object is created. If you select “Create Another Object,” another “Create” dialog is displayed immediately after the object is created.

6. Choose “Create.”

If you chose “Define Additional Properties,” the object dialog appears. If you chose “Create Another Object,” another “Create Object” dialog appears.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>
Using the object dialog in NetWare Administrator	“NetWare Administrator” in <i>Utilities Reference</i>


Creating Leaf Objects Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later
- The Create object right to the container object that will contain the new leaf object

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Select the container that will contain the new leaf object.

The objects in the selected container are listed.

To see if you’re in the right context, look at the title bar on the screen.

Press <F1> if you need help.



Note

If you are creating User objects, remember that users who are using non-NetWare 4.1 workstations must be created in the container at which the bindery services context is set for the server that they need to log in to.

You can create User objects for users who have NetWare 4.1 workstations anywhere in the Directory tree, but the users must know their context in order to log in. You should create User objects in the container where the users will typically log in.

4. Press <Insert>.

5. From the “Select an Object Class” screen, select the object class that you want to create.

If the object class you want to create does not appear, you cannot create that object in the selected container. Press <Esc> to return to the browser; then select a different container type.

6. Type the information you are prompted for and press <Enter>.

Each leaf object that you create has a different dialog. For details on each dialog, press <F1> for help.

7. If you want to create another leaf object, choose “Yes.” If you do not, choose “No.”

If you choose “Yes,” you are prompted for information about the next object you want to add. Repeat Step 6, and then continue with Step 8.

If you choose “No”, the leaf object is displayed in the Directory tree. Continue with the next step.

8. Press <F10> to edit this object.

A menu appears from which you can choose to view or edit information about this object.

9. Choose an option from the “Actions” menu and add any necessary information.

10. To exit, press <Esc> until you return to the “NETADMIN Options” menu.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Managing Groups of User Objects

NetWare 4.1 allows you to manage User objects as a group, which is often more efficient than managing them individually. Six objects that can help you manage groups of User objects are described in Table 1-7.

Table 1-7
Objects That Help Manage Users

Object	Description
Organization object	Allows you to assign trustee rights, login scripts, and user defaults to the User objects in the Organization.
Organizational Unit object	Allows you to assign trustee rights, login scripts, and user defaults to the User objects in the Organizational Unit.
Group object	Provides an efficient way for you to manage only one object, the Group object, instead of many individual User objects.
Profile object	Allows you to set up a specific work environment by using a common login script for groups of users who need similar work environments but who are not located in the same container object.
Organizational Role object	<p>Allows you to assign rights to a particular position and set of responsibilities, rather than to a person. The person who occupies that position may change frequently, but the responsibilities of that position do not.</p> <p>The difference between a Group object and an Organizational Role object is that a Group object usually has many members, whereas an Organizational Role object usually has only one or two.</p>
USER_TEMPLATE	Allows you to apply default property values to any user that you create in a container object. You can choose to apply the information in the user template when you create new User objects.

Managing Group Objects

If you want a user to have access to an object, you must make a trustee assignment to that object. Rather than make trustee assignments to many users, you can create a Group object and then, with just one trustee assignment, grant access to all the users who belong to the Group.

If a trustee assignment names a Group object as the trustee, every user in the membership list of the Group object is granted the same access that is granted to the Group object.

After you have created a Group object and added User object names to it, you manage the rights of the Group object rather than the rights of the individual group members.


For example, suppose you have a word processing application on the network that many users need to access. You could create a Group object named WORD PROCESSOR USERS and add the User object names of the users who need access to the application.

Then, rather than granting file trustee rights to each of the individual User objects, you grant the file trustee rights to the Group object WORD PROCESSOR USERS for the application and the working directory.

The users can then use the word processing application just as if you had granted them file trustee rights to the application individually.

When a user is added to the membership list of a Group object, the Group is listed in that user's Security Equal To property. Security Equal To is a property of every User object that lists other objects. The user is granted all rights that any object (User, Group, Printer, etc.) in that list is granted, both object and file rights.

Only User objects can be listed in a Group, and you can add User objects from any part of the Directory tree to a Group.



A Group object is *not* a container. It does not “contain” User objects; users' names are merely assigned to a Group object.

To create a Group object, see “Creating Leaf Objects” on page 35.

You must create User objects before you can add them to the membership list of a Group object. See “Creating Leaf Objects” on page 35 for instructions on creating User objects.

After you have created a Group object, use the following procedures to

- ◆ Add members to a Group object
- ◆ Give a Group object rights to files and directories
- ◆ Delete members from a Group object

You can use either the NetWare Administrator or NETADMIN utility to manage Group objects. Both procedures are documented in this section.

For more information about	Refer to
Groups	“Group object” in <i>Concepts</i>
Object and property rights	“Rights Needed to Create and Manage Objects” on page 17
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Adding Members to a Group Object Using NetWare Administrator

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- At minimum, the Write right to the Members property of the Group object
- At minimum, the Write right to the Security Equal To property of the User object
- At minimum, the Write right to the ACL property of the Group object
- The Group object must already exist, and the User objects you want to add as members of the Group must already exist

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Select the Group object you want to edit.**

For information on moving around in the browser and selecting objects, press <F1>.
- 3. Choose “Details” from the “Object” menu.**
- 4. Select the “Members” button at the right side of the “Object” dialog.**
- 5. Choose the “Add” button to browse the Directory tree for User objects.**
- 6. Browse the Directory tree until the User object you want appears in the “Objects” box.**
- 7. Choose “OK.”**
- 8. Repeat Steps 5 through 7 to add more User objects to the Group object.**

9. When you have finished adding User objects to the Group object, choose "OK" to save your changes and return to the browser.


Adding Members to a Group Object Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Supervisor right to the Group object, or the Write or Supervisor right to the Members property of the Group object
- The Supervisor right to the Group object, or the Write or Supervisor right to the Security Equal To property of the User object
- The Supervisor or Write right to the ACL property of the Group and User objects
- The Group object must already exist, and the User objects you want to add as members of the Group must already exist

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose "Manage Objects" from the "NETADMIN Options" menu.
3. Browse the Directory until the Group object appears on the screen.

Use the instructions at the bottom of the screen to browse the directory. Press <F1> if you need help.

4. When the Group object appears in the "Object" list, select it and press <F10>.

The "Actions" menu appears.

5. Choose “View or Edit Properties of This Object.”
6. Choose Group members from the “View or Edit Group” menu.
7. Press <Insert> at the “Group Members” screen, and then press <Insert> again to browse for the User object you want to add to the Group object.
8. When the User object you want to add appears in the Directory, select it and press <F10>.
9. When the selected User object appears in the “Members” screen, press <Enter>.

To select multiple User objects, use <F5>.
10. Continue to press <Insert> and select User objects until you have added all the users you want as Group members.
11. To save the list of Group members, press <F10>.
12. To exit, press <Esc> until you return to the “NETADMIN Options” menu.

Giving Group Object Rights to Files and Directories Using NetWare Administrator


Prerequisites

Checklist



- A workstation running NetWare Administrator
- The Read object right to the Volume object
- Rights to the file system

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Select the Group object you want to edit.**

For information on moving around in the browser and selecting objects, press <F1>.
- 3. Choose “Details” from the “Object” menu.**
- 4. Select the “Rights to File System” button on the right side of the “Object” dialog.**
- 5. To select a volume, select “Include.”**

A list of volumes appears in the “Select Object” box. Or, you can browse the Directory for a volume.
- 6. From the “Volumes” list, select the volume that contains the directory or file.**
- 7. Choose “Add.”**
- 8. Select the volume that contains the directory or file you want to grant rights to.**
- 9. From the “Files and Directories” dialog, select the directory or file that you want to grant rights to.**

The default rights that make up this object’s trustee assignment to the file or directory appear in the “Rights” area.
- 10. Select the check boxes next to the rights that you want to add.**

You must have the Access Control right to the file or directory to make trustee assignments to the file or directory.
- 11. Choose “OK.”**

The new trustee assignment is now effective for this object.

Giving Group Object Rights to Files and Directories Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Read object right to the Volume object
- Rights to the file system

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Browse the Directory until the Group object appears on the screen.

Use the instructions at the bottom of the screen to browse the directory. Press <F1> if you need help.

4. When the Group object appears in the “Object” list, select it and press <F10>.

The “Actions” menu appears.

5. Choose “View or Edit Rights to Files and Directories.”

6. Select a Volume where you want to make the Group object the trustee of a directory or file.

Press <Insert> to type the Volume object name or press <Insert> twice to browse the Directory.

7. Press <Insert> to type a beginning pathname to the directories in which you want to make trustee assignments, or press <Insert> again to browse for the path.

8. Select “Directories/Files” and press <Enter>.

Choose whether you want to view files, directories, or both when you are selecting one to give a trustee assignment to.

9. Select “Trustee Search Depth” and press <Enter>.

Choose whether you want to view only the files or directories in the current directory, or to search subdirectories.

10. To list the trustee assignments, press <F10>.

The “Trustee Directory Assignments” screen appears.

11. To select a directory or file in which the Group object should be added as a trustee, press <Insert>.

12. To accept the directory you specified earlier, press <Enter>; or, to browse for file system directories, press <Insert>.

13. To add or delete the rights granted, select “Trustee Directory Assignments” and press <Enter>.

The “Trustee Rights Granted” menu appears.

14. To view or add rights that are not yet granted, press <Insert>.

Press <F1> if you help.

15. To save the trustee assignments, press <F10>.

16. Continue selecting directories and files and granting rights until finished.

17. To exit, press <Esc> until you return to the “NETADMIN Options” menu.

Deleting Members from a Group Object Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- The Supervisor right to the Group object, or the Write or Supervisor right to the Members property of the Group object
- The Supervisor right to the Group object, or the Write or Supervisor right to the Security Equal To property of the User object
- The Supervisor or Write right to the ACL property of the Group and User objects

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Select the Group object you want to edit.**

For information on moving around in the browser and selecting objects, press <F1>.
- 3. Choose “Details” from the “Object” menu.**
- 4. Select the “Members” button at the right side of the “Object” dialog.**

The list of User objects for this group appears.
- 5. From the “Members” dialog, select the name you want to delete.**
- 6. Choose “Delete.”**
- 7. If you want to delete other names, continue selecting names and choosing “Delete.”**



You can delete several users at a time by holding down the button on the mouse, dragging the mouse arrow over the names, and choosing “Delete.”

8. When you have finished deleting members, choose “OK” to save your changes and return to the browser.


Deleting Members from a Group Object Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Supervisor right to the Group object, or the Write or Supervisor right to the Members property of the Group object
- The Supervisor right to the Group object, or the Write or Supervisor right to the Security Equal To property of the User object
- The Supervisor or Write right to the ACL property of the Group and User objects

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.
3. Browse the Directory until the Group object appears on the screen.

Use the instructions at the bottom of the screen to browse the directory. Press <F1> if you need help.

4. When the Group object appears in the “Object” list, select it and press <F10>.

The “Actions” menu appears.

5. Choose “View or Edit Properties of This Object.”
6. Select Group members from the “View or Edit Groups” menu.

7. Select the User object you want to delete from the Group object and press <Delete>.

To select multiple User objects, use <F5>.

8. To confirm the deletion, choose “Yes.”

9. To exit, press <Esc> until you return to the “NETADMIN Options” menu.

Managing Profile Objects

Profile objects contain login scripts that are used by groups of users who need similar work environments but who are not located in the same container object.

When a Profile object is named in a User object, the login script contained in the Profile object is executed when the user logs in, after any login script in the Organization or Organizational Unit has executed.

Users can have only one Profile, so only one profile script can execute for any user.

For information about creating a login script, see Chapter 3, “Creating Login Scripts.”

For an example of a login script used in a Profile object, see “Profile Login Script” on page 248.

You can use either the NetWare Administrator or NETADMIN utility to create a Profile object. Both procedures are documented in this section.

For more information about	Refer to
Profile objects	“Profile object” in <i>Concepts</i>
Creating login scripts	Chapter 3, “Creating Login Scripts.”
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Creating Profile Objects Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- The Create object right to the object that will contain the new Profile object

Procedure



1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.

2. Select the object that will contain the new Profile object.

For information on moving around in the browser and choosing objects, press <F1>.



Only Organization and Organizational Unit objects can contain Profile objects.

3. Choose “Create” from the “Object” menu.

4. Select “Profile” from the “New Object” dialog.

The “Create Profile” dialog appears.

If Profile does not appear under “New Object,” you cannot create Profile objects in this container. Select or create another object to contain the Profile object.

5. Choose “OK.”

6. Type the Profile object name in the box provided.

7. (Optional) Select “Define additional properties.”

Select this option if you want to write a Profile login script or supply additional information about the new Profile object. Instructions for creating a Profile script are in Chapter 3, “Creating Login Scripts.”

8. Choose “Create.”

If you selected “Define Additional Properties,” the “Identification” dialog appears.

9. (Optional) Enter information in the fields provided in the “Identification” page of the “Object” dialog.

10. (Optional) Select the “See Also” button at the right side of the object dialog.

The “See Also” page allows you to add information about the Profile object you are creating. For example, you might list the User objects to whom you have assigned this script.

Choose “Help” at any time for information on the current task.

11. (Optional) Choose the “Login Script” page at the right side of the object dialog to add commands to the Profile login script.

Use this page to specify commands that execute when a user logs in, such as a drive mapping.

12. To save the new Profile object and return to the browser, choose “OK.”


Creating Profile Objects Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Create object right to the object that will contain the new Profile object

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Select the object that will contain the new Profile object.

The objects in the selected container are listed. To see if you’re in the right context, look at the title bar on the screen.

Press <F1> if you need help.

4. Press <Insert>.

5. Select “Profile.”

If the Profile object class does not appear, you cannot create that object in the selected container. Press <Esc> to return to the browser, and then select a different container type.

6. Type the new Profile object name and press <Enter>.

7. If you want to create another Profile object, choose “Yes.” If you do not, choose “No.”

If you chose “Yes” you are prompted to type the new Profile object name. Repeat Steps 3 through 7, and then continue with Step 8.

If you chose “No,” then the Profile object is displayed in the Directory tree. Continue with Step 8.

8. To edit this object, press <F10>.

A menu appears from which you can choose to view or edit information about this object.

9. Choose “View or Edit Properties of This Object.”

10. Choose “Login Script.”

11. To enter commands for this Profile login script, choose “No” or, to copy a login script from another Profile object, choose “Yes.”

The commands you place in the Profile login script are executed when users who belong to this Profile object log in.

For information on the commands, press <F1> or see “Login Script Commands and Variables” on page 187.

12. To save your changes, press <F10>.

13. To exit, press <Esc> until you return to the “NETADMIN Options” menu.

Managing Organizational Role Objects

An Organizational Role object allows you to assign rights to a particular position rather than to the person who occupies that position. The people who occupy that position may change frequently, but the responsibilities of the position do not.

The user assigned to an Organizational Role is called the occupant and is granted all rights that are granted to the Organizational Role object.

When a user is added to the occupant list of an Organizational Role object, the Organizational Role is listed in that user’s Security Equal To property.

Security Equal To is a property of every User object that lists other objects. The user is granted all rights that any object (User, Group, Printer, etc.) in that list is granted, both to objects and to files and directories.

You can use the NetWare Administrator or NETADMIN utility to create an Organizational Role object. Both procedures are documented in this section.

For more information about	Refer to
Organizational Role object	“Organizational Role object” in <i>Concepts</i>
Security equivalence	“Security Equal To” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Creating Organizational Role Objects Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- The Create object right to the object that will contain the new Organizational Role object

Procedure



1. **Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
2. **Select the object that will contain the new Organizational Role object.**

For information on moving around in the browser and selecting objects, press <F1>.



Only Organization and Organizational Unit objects can contain Organizational Role objects.

3. **Choose “Create” from the “Object” menu.**
4. **Select “Organizational Role” from the “New Object” dialog.**

If “Organizational Role” does not appear under “New Object,” you cannot create Organizational Role objects in this container. Select or create another object to contain the Organizational Role object.
5. **Choose “OK.”**

The “Create Organizational Role” dialog appears.
6. **Type the Organizational Role object name in the box provided.**
7. **(Optional) “Select Define Additional Properties.”**
8. **Select the “Create” button at the bottom of the window.**
9. **Enter information in the fields provided in the “Identification” dialog.**

The “Identification” page of the “Object” dialog appears.

10. Select the button to the right of “Occupant.”

11. Choose “Add.”

The “Select Object” window appears.

12. Select User objects from the “Directory Context” window until the objects you want are shown in the “Object” window.

13. Select the User object in the left window to occupy the Organizational Role; then choose “OK.”

The object you selected appears in the “Occupant” window.

14. Choose “OK” in the “Occupant” window.

15. When you are finished adding User objects as Occupants, choose “OK” in the “Organizational Role” window.

16. (Optional) Select the “See Also” button at the right side of the object dialog.

The “See Also” page allows you to add information about the Organizational Role object you are creating. For example, you might list the User objects that you have assigned as occupants.

17. To save the new Organizational Role object and return to the browser, choose “OK.”


Creating Organizational Role Objects Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Create object right to the object that will contain the new Organizational Role object

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Select the object that will contain the new Organizational Role object.

The objects in the selected container are listed. To see if you’re in the right context, look at the title bar on the screen.

Press <F1> if you need help.

4. Press <Insert>.

5. Select “Organizational Role.”

If the Organizational Role object class does not appear, you cannot create that object in the selected container. Press <Esc> to return to the browser, and then select a different container type.

6. Type the new Organizational Role object name.

7. Type the Mailbox Location and press <Enter>.

8. If you want to create another Organizational Role object, choose “Yes.” If you do not, choose “No.”

If you chose “Yes,” you are prompted to type the new Organizational Role object name. Repeat Step 6 and then continue with Step 8.

If you chose “No,” then the Organizational Role object is displayed in the Directory tree. Continue with Step 9.

9. To edit this object, press <F10>.

A menu appears from which you can choose to view or edit information about this object.

- 10. Choose “View or Edit Properties of This Object.”**
- 11. Choose “Identification” from the “View or Edit Organizational Role” menu.**
- 12. Specify a User object for the Organizational Role.**
 - 12a. Select the field next to “Occupant” and press <Enter>.**
 - 12b. Press <Insert>.**
 - 12c. Type the complete name of a User object in the space provided, or press <Insert> to browse the Directory and select a User object to be the occupant of the Organizational Role.**

The path from the object to the root of the tree forms the object’s complete name.
- 13. Select additional User objects as needed.**
- 14. To save the list of occupants, press <F10>.**
- 15. Enter information in other fields as needed.**
- 16. To save changes, press <F10>.**
- 17. To exit, press <Esc> to return to the “NETADMIN Options” menu.**

Managing User Templates

A user template contains default information that you can apply to User objects to give them default property values.

You can create a user template in an Organization or Organizational Unit object either when you create the container object or later on.

Then, when you create a User object, you are prompted to use the defaults in the user template. If you do, the property values you entered in the user template, such as login time restrictions, password restrictions, etc., are copied into the User object’s properties.

The user template is actually a User object named USER_TEMPLATE. You enter information in this User object just as you would for any other User object. However, not all properties of a User object can be copied from a user template.

When you create a user template, you can copy information from the parent container's user template. For example, if you create a user template in SALES.NOVELL, you are prompted to copy the user template from NOVELL, if one exists. Using this feature, you can avoid having to reenter similar information for lower-level containers.

User template information is taken from the nearest parent container. If the container object in which you create a User object does not have a user template, you can apply the parent container's user template to the User object.

When working with user templates, remember the following:

- ◆ Changing values in a user template does not change those values in existing User objects. The changes apply only to User objects created after the user template values have been changed.
- ◆ To update information for existing users, you must enter the changes for each User object.
- ◆ You cannot use a user template to grant NDS or file system rights.

You can use the NetWare Administrator or NETADMIN utility to create a user template. Both procedures are documented in this section.

For more information about	Refer to
User defaults	"User template" in <i>Concepts</i>
User objects	"User object" in <i>Concepts</i>
Using the NETADMIN utility	"NETADMIN" in <i>Utilities Reference</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>

Creating and Editing User Templates with NetWare Administrator


Prerequisites



Checklist

- A workstation running NetWare Administrator
- The Create object right to the container object that will contain the user template

Procedure



Procedure

1. **Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
2. **Select the object that will contain the new user template.**

For information on moving around in the browser and selecting objects, press <F1>.



Note

Only Organization and Organizational Unit objects can contain User objects.

3. **Choose “User Defaults” from the “Object” menu.**
4. **If you want the new user template to inherit the properties from the parent container’s user template, choose “Yes.” If not, choose “No.”**

If there is no parent container, you are not given this prompt. Continue with Step 5.

If you chose “Yes,” a USER_TEMPLATE User object is created and the first page of defaults, “Identification,” appears with the same information that is in the parent container’s user template.

If you chose “No,” a USER_TEMPLATE User object is created and the first page of defaults, “Identification,” appears.

5. **(Optional) On the “Identification” page, enter or change the information that you want to apply to new User objects.**

For example, the location for all User objects to which you will apply the user template might be New York. You would then type “New York” in the “Location” field.

6. Select other **USER_TEMPLATE** pages as needed and enter the template information.
7. To save the user template and return to the browser, choose **“OK.”**


Creating and Editing User Templates with NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later
- The Create object right to the container object that will contain the user template

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose **“Manage Objects”** from the **“NETADMIN Options”** menu.

3. **Browse the Directory tree to find and select the container object in which you want to add a new container object.**

The objects in the selected container are listed.

To see if you're in the right context, look at the title bar on the screen. Press <F1> for help.

4. **Press <Insert>.**
5. **From the “Select an Object Class” screen, select the container type that you want to create.**

If the container object class you want to create does not appear, you cannot create that object in the selected container. Press <Esc> to return to the browser, and then select a different container type.

6. Type the new container object name and press <Enter>.

7. If you want to create a user template to be applied to new User objects created in this container, type “Y” and press <Enter>.



You are not prompted to create a user template if you are creating a Country container object.

8. Choose “View or Edit Properties of This Object.”

9. Choose “Edit Template User.”

The “View or Edit User” screen appears.

10. Enter or change the values of the user template as needed.

The help line at the bottom of the screen gives information on each of the options as you highlight them.

For more information, press <F1>.

11. To save the information, press <F10>.

12. To exit, press <Esc> until you return to the “NETADMIN Options” menu.

Searching for Objects

When you want to find information in the Directory database without opening numerous containers to view the objects, you can use the Search feature.

With the Browse object right, you can search for object classes anywhere in the Directory tree. With the Compare property right, you can search for objects with properties that match a particular value.

You can search for objects using the NetWare Administrator, NETADMIN, or NLIST utility. All three procedures are documented in this section.

Searching for Objects Using NetWare Administrator

The NetWare Administrator utility searches each object in the Directory database unless you narrow the search by specifying properties in combination with additional variables such as “less than” or “equal to.”

For example, to find all users in New York, you could search for User objects with State or Province Name “equal to” New York. The Search feature then displays objects that meet the search criteria.

Prerequisites



- A workstation running NetWare Administrator
- The Browse object right to the object you want to see in the search

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Select an object in the Directory tree where the search should start.**
- 3. Choose “Search” from the “Object” menu.**

A “Search” window appears.
- 4. Choose the browser icon (to the right of the “Start From” field) to select an object to start the search from.**
- 5. Specify how much of the Directory tree to search.**
 - ◆ To search everything below the “Start From” object, select “Search Entire Subtree.”
 - ◆ To search only among objects one level below the “Start From” object, do not select “Search Entire Subtree.” Continue with Step 6.

6. Select the down-arrow to the right of the “Search for” field to select an object class to search for.

- ◆ Choose “OK” if you want all the objects of this selected object class to be listed.
- ◆ If you want only those objects having properties that match certain criteria to be listed, continue with Step 7.

7. (Optional) Select the down-arrow to the right of the “Property” field to select which property’s value will be examined.

The properties in this list change depending on which object class you selected in the “Search for” field.

8. (Optional) Select how you want to compare a value to the selected property.

You can select “Equal to,” “Greater than,” “Less than,” or other choices.

9. (Optional) Enter a value in the field to the right of the comparison method that you entered in Step 8.

For example, if you are searching for a Profile object that has the value “Manager’s Profile” in the Other Names property, you might select the property “Other Names,” then use the comparison method “Equal To” and the value of “Manager’s Profile.”

10. Choose “OK.”

Objects that match your selections are displayed in a “Search Results” window.

Additional Information

For more information about	Refer to
Properties	“Property” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Searching for Objects Using NETADMIN

NETADMIN searches each object in the Directory database unless you narrow the search by specifying properties in combination with additional variables such as “less than” or “equal to.”

For example, to find all users in New York, you could search for User objects with State or Province Name “equal to” New York. The Search feature then displays objects that meet the search criteria.


Prerequisites



Checklist

- A workstation running DOS 3.30 or later
- The Browse object right to the object you want to see in the search

Procedure



Procedure

1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage According to Search Pattern” from the “NETADMIN Options” menu.

3. Enter the object name you wish to search on in the “Enter Object Name” field, use * to view all objects, or limit the search by combining characters with the *.

For example, if you want to view all objects whose names begin with “P,” enter “P*” in the “Enter Object Name” field.

4. Select the object class or classes you want to search on in the “Object Class” field. To search on all object types, select “/All classes/.”

4a. Press <Enter> to select the object classes.

4b. If the object classes you want to search on don’t appear in the “Object Classes Included” screen, press <Insert> to select additional object classes.

5. Choose “Yes” in the “Show Alias Class” field if you want to view an Alias object as an alias, and not as the object the alias represents.

6. Press <F10>.

Objects that match your selections are displayed in the browse screen.

Additional Information

For more information about	Refer to
Properties	“Property” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Searching for Objects Using NLIST

NLIST is a workstation command line utility that allows you to

- ◆ List objects and object properties
- ◆ View information about users, groups, volumes, and servers (such as object properties, names, property groups, and login information)
- ◆ Search for objects and object properties (including groups of properties for objects) on NetWare 2 and NetWare 3 servers


For instructions on using NLIST, see “NLIST” in *Utilities Reference*.

Moving Objects in the Directory Tree

While previous versions of NetWare allowed you to move only leaf objects, NetWare 4.1 allows you to move both leaf objects and container objects to other containers in the Directory tree.

You can move a container object only if it is the root of a NetWare Directory Services (NDS) partition that has no subordinate partitions.

When you move a container object, create an alias object that points to the container object you're moving. Then users can continue logging in to the network and finding the container object (and the objects it contains) in its original Directory location.



If you move a container object and do not create an alias, users who are unaware of the object's new location will not easily find the object in the Directory tree, since they will look for it in its original Directory location.

Also, users may not be able to log in if the name context in their configuration file (NET.CFG file) references the moved container.

When you move a leaf or container object, NDS changes all references to the moved object. Although the object's common name remains unchanged, the context name of the object (and of all its subordinates) changes.

Because the context of a container object changes when you move it, users whose name context in their configuration file (NET.CFG file) references the moved container need to update their NET.CFG so that it references the container's new name.

To automatically update users' NET.CFG file with a new name context after you move a container object, use the NCUPDATE utility. For instructions, see "NCUPDATE" in *Utilities Reference*.

You can use NetWare Administrator or NETADMIN to move objects in the Directory tree. Both procedures are documented in this section.

Moving Leaf Objects Using NetWare Administrator

Prerequisites

Checklist



- A workstation running NetWare Administrator
- The Supervisor right to the object you want to move
- The Create object right to the destination container

Procedure

Procedure



- 1. From the MS Windows Program Manager or the OS/2 desktop, choose the “NetWare Administrator” icon.**
- 2. From the browser window, select one or more leaf objects.**

To select multiple leaf objects, press <Ctrl> on the keyboard and, without releasing <Ctrl>, select the objects with the mouse.
- 3. From the “Object” menu, choose “Move.”**
- 4. Select the browser button to the right of the “Destination” box.**

The “Directory Context” box appears in the lower right corner of the screen. Use this field to browse the Directory tree’s containers.

The “Objects” box that appears in the lower left corner shows the containers that appear when you browse the Directory.
- 5. From the “Objects” box, select a container object (an Organization or Organizational Unit) as the place to move the objects to; then choose “OK.”**
- 6. In the “Move” dialog, choose “OK.”**

The listed objects are moved to the destination container.

Additional Information

For more information about	Refer to
Directory tree	"Directory tree" in <i>Concepts</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>
Objects	"Object" in <i>Concepts</i>
Rights	"Rights" in <i>Concepts</i>

Moving Leaf Objects Using NetWare Administrator (Drag-and-Drop)


NetWare Administrator allows you to move an object to a different location in the Directory tree using the drag-and-drop method.

Prerequisites



- A workstation running NetWare Administrator
- The Supervisor right to the object you want to move
- The Create object right to the destination container

Procedure



- 1. From the MS Windows Program Manager, choose the "NetWare Administrator" icon.**
- 2. (Optional) From the "Tools" menu, open another browser window by selecting "Browse."**

If you can see both the object and the destination container in one browser window, it is not necessary to open two browser windows. Go directly to Step 4.

- 3. Reshape the browser windows, using the window borders, so that you can view both windows at once.**

4. From one of the browser windows, select one or more leaf objects.

To select multiple leaf objects, press <Ctrl> on the keyboard and, without releasing <Ctrl>, select the objects with the mouse.

5. Press <Ctrl> on the keyboard and, without releasing <Ctrl>, click and hold down on the object(s) and drag the object(s) to the destination container; then release the mouse button.

6. In the “Move” dialog, choose “OK.”

The listed objects are moved to the destination container.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Directory tree	“Directory tree” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Moving Leaf Objects Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Create object right to the destination container
- The Supervisor right to the object you want to move

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. From the “NETADMIN Options” menu, choose “Manage Objects.”

Your current context appears in the upper left corner.

3. Select the object that you want to move.

- ◆ If the object you want to move appears on the list, select it and press <F10>.
- ◆ If the object is not on the list, browse the Directory by selecting container objects and pressing <Enter> until you see the object you want. Select it and press <F10>.

4. From the “Actions” menu, choose “Move.”

5. Using the Down-arrow key to reach the field, highlight the “New Context” field.

6. Assign a new context to the object you want to move.

- ◆ If you know the new context that you want the object to be in, type the new context in the highlighted field.
- ◆ If you don't know the new context that you want the object to be in, press <Insert> twice to browse the Directory for the destination container; then select the destination container and press <F10>.

7. To accept the new context as the destination container, press <Enter>.

8. To confirm that you want to move the object listed in the “Old Context” field to the container listed in the “New Context” field, press <F10>.

The selected object is moved to the destination container.



Note

You need to wait for processes throughout the Directory to be completed before you can move this object again.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Directory tree	“Directory tree” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Moving Container Objects Using NetWare Administrator


Prerequisites



Checklist

- A workstation running NetWare Administrator
- The Supervisor right to the object you want to move
- The Create object right to the destination container

Procedure



Procedure

- From the MS Windows Program Manager or the OS/2 desktop, choose the “NetWare Administrator” icon.**
- From the “Tools” menu, choose “Partition Manager.”**
- From the “Partition Manager” browser, select the partition that you want to move.**

You can move a container object only if it is the root of a NetWare Directory Services (NDS) partition, and only if it contains no subordinate partitions.

In Partition Manager, the partition icon appears to the left of the object icon. If the container you want to move is not a partition, select the container and choose “Create as New Partition.”

4. From the “Object” menu, choose “Move Partition.”
5. Select the browser button to the right of the “Destination” box.


Use the browser in the “Directory Context” box to view the Directory tree’s containers.

The “Objects” box that appears in the lower left corner shows the containers that you select in the “Directory Context” box.

6. From the “Objects” box, select a container object (an Organization or Organizational Unit) as the place to move the listed objects to; then choose “OK.”

7. Choose “Create Alias in Place of Moved Container.”

The alias object will point to the partition’s new location.



If you move a container object and do not create an alias, users who are unaware of the object’s new location will not easily find the object in the Directory tree, since they will look for it in its original Directory location.

Also, users may not be able to log in if the name context in their configuration file (NET.CFG file) references the moved container.

8. In the “Move” dialog, choose “OK.”

If you chose to create an alias in place of the moved container, NetWare Administrator polls for the creation of the alias object before it moves the selected partition.

If you moved a container and created an alias in its place, you should use the NCUPDATE utility to update the name context of users in the moved container. For instructions, see “NCUPDATE” in *Utilities Reference*.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Alias objects	“Alias object” in <i>Concepts</i>
Directory tree	“Directory tree” in <i>Concepts</i>
NCUPDATE	“NCUPDATE” in <i>Utilities Reference</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>


Moving Container Objects Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Create object right to the destination container

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. From the “NETADMIN Options” menu, choose “Manage Objects.”

Your current context appears in the upper left corner.

3. Select the object that you want to move.

- ◆ If the object you want to move appears on the list, select it and press <F10>.
- ◆ If the object is not on the list, browse the Directory by selecting container objects and pressing <Enter> until you see the object you want. Select it and press <F10>.

4. From the “Actions” menu, choose “Move.”

You can move a container object only if it is the root of a NetWare Directory Services (NDS) partition, and only if it contains no subordinate partitions.

In NETADMIN, when you select a container object that is a partition, the context-sensitive help at the bottom of the screen reads “This is a partition.”

If the container you want to move is not a partition, you must first use a partition management utility (PARTMGR or NetWare Administrator) and create the container as a new partition.

5. Using the Down-arrow key to reach the field, highlight the “New Context” field.

6. Assign a new context to the object you want to move.

- ◆ If you know the new context that you want the object to be in, type the new context in the highlighted field.
- ◆ If you don’t know the new context that you want the object to be in, press <Insert> twice to browse the Directory for the destination container; then select the destination container and press <F10>.

7. To accept the new context as the destination container, press <Enter>.

8. To confirm that you want to move the object listed in the “Old Context” field to the container listed in the “New Context” field, press <F10>.

9. To create an alias in place of the moved container, choose “Yes.”



Important

If you don't create an alias, users who are unaware of the object's new location cannot easily find the object in the Directory tree, since they will look for it in its original Directory location.

Also, users might not be able to log in if the name context in their configuration file (NET.CFG file) references the moved container.

The alias object will point to the partition's new location.

The selected object is moved to the destination container.



Note

You need to wait for processes throughout the Directory to be completed before you can move this object again.

If you moved a container, and created an alias in its place, you should use the NCUPDATE utility to update the name context of users in the moved container. For instructions, see "NCUPDATE" in *Utilities Reference*.


Additional Information

For more information about	Refer to
Objects	"Object" in <i>Concepts</i>
Rights	"Rights" in <i>Concepts</i>
Directory tree	"Directory tree" in <i>Concepts</i>
Using the NETADMIN utility	"NETADMIN" in <i>Utilities Reference</i>

Deleting Objects from the Directory Tree

When you delete leaf objects, NetWare Directory Services (NDS) removes references to the deleted object. For example, if you delete a Profile object, that Profile object is deleted from any User objects that list it. Or, if you delete a User object, all trustee assignments listing that user are deleted.

Some special considerations apply when you delete particular objects, such as a NetWare Server object, a User object, or an Alias object. These cautions are explained in the next three sections.



You cannot undo a Delete Object operation. To get the object back, you must re-create it and reenter all data in its properties.

You can use NetWare Administrator or NETADMIN to delete objects from the Directory tree. Both procedures are documented in this section.

Cautions When Deleting NetWare Server Objects

An NDS NetWare server is one that you have installed in the Directory tree. Any server that is not in the Directory tree is a bindery server.

You can use NetWare Administrator or NETADMIN to delete bindery servers. But consider the following before you delete an NDS Server object:

- ◆ You cannot delete a NetWare server from the Directory tree using the NETADMIN utility.
- ◆ You must use the Partition Manager tool in the NetWare Administrator utility to delete an NDS NetWare server from the Directory tree.
- ◆ If the server you want to delete contains a master replica, you must first change the master to a different type and designate another replica on another server as the master.

Cautions When Deleting User Objects

You must be careful not to delete a trustee object which has the only trustee assignment to a part of the Directory tree. If you did, you could cut off access to that part of the Directory tree.

You must also be careful not to block everyone's rights to an object with an Inherited Rights Filter, leaving no one with access to part of the Directory tree.

Therefore, be aware of the following before you delete a User object:

- ◆ To avoid losing access to any object, check the "Rights to Other Objects" and "Trustees of This Object" attributes of the User object you want to delete. To view these attributes, select the User object, and right-click once. Then select the attribute you want to view from the available options.

If the User object you want to delete has the Supervisor right to another object, transfer that Supervisor right to another User object before you delete the original User object.

Or, give a User object in a higher container object the Supervisor right to objects, and then block other users from deleting those objects.

- ◆ Do not delete user ADMIN until you have given another User object the same Supervisor right that ADMIN has.
- ◆ When you delete a User object, any security equivalences that you have assigned to other trustees are lost.
- ◆ When you delete a User object, the home directories and mail directories will not be deleted.

Cautions When Deleting Alias Objects

Be aware of the following before you delete an Alias object:

- ◆ If you delete an alias object (which appears as the actual object it is pointing to), you only delete the alias object, not what it points to.

When you delete an alias object, it does not affect the object that the alias object points to.

- ◆ If you delete the object that an alias object points to, the alias object is also deleted.


Deleting Objects Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- The Delete object right to the object that you want to delete

Procedure



1. **Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**

2. **Using the browser, select the object you want to delete.**

Only leaf objects and container objects that are empty can be deleted.

3. **Select (check) the “Trustees of This Object” and the “Rights to Other Objects” dialogs.**



If the User object you want to delete has the Supervisor right to another object, give another User object that Supervisor right before you delete the original User object.

4. **Choose “Delete” from the “Object” menu.**



You cannot undo an object deletion. To get the object back, you must re-create it and reenter all data in its properties.

5. **To confirm the deletion, choose “OK.”**

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Deleting Objects Using NETADMIN


Prerequisites



Checklist

- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Delete object right to the object that you want to delete from the Directory tree

Procedure



Procedure

1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Browse the Directory until the object you want to delete appears.

Use the instructions at the bottom of the screen to browse the Directory. Press <F1> if you need help.

4. Select the object and press <Delete>.



Important

If the User object you want to delete has the Supervisor right to another object, give another User object that Supervisor right before you delete the original User object.



You cannot undo an object deletion. To get the object back, you must re-create it and reenter all data in its properties.

5. To confirm the deletion, choose “Yes.”

The object is deleted from the Directory tree.

6. Press <Esc> until you return to the “NETADMIN Options” menu

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>


Renaming Leaf and Container Objects

You may want to rename objects to make their names more descriptive or to reflect your changing environment. For example, you may want to rename a printer from LASER IN BLDG D to LASER IN BLDG A, or an Organizational Unit from SALES to ACCOUNTS.

When you rename an object, NetWare Directory Services changes all references to the renamed object.

Renaming a leaf object changes only the object’s common name, which is the name that is displayed in the Directory tree. It does not change the object’s context. However, renaming a container object changes the object’s common name as well as its context.

When renaming a container object, you should seriously consider creating an alias object that points to the container object you’re renaming. Then users can continue logging in to the network and can see the container object’s original name.



If you rename a container object and do not create an alias, users who are unaware of the object's new name will not easily find the object in the Directory tree, since they will look for its original name.

Also, users may not be able to log in if the name context in their configuration file (NET.CFG file) references the renamed container.

To automatically update users' NET.CFG file with a new name context after you rename a container object, you can place a command in the renamed container's login script that will run the NCUPDATE utility.

Renaming Objects Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- The Rename object right to the object that you want to rename

Procedure



1. **From the MS Windows Program Manager or the OS/2 desktop, choose the "NetWare Administrator" icon.**
2. **Using the browser, select the object that you want to rename.**
3. **From the "Object" menu, choose "Rename."**
4. **Type the name you want to give to the object you selected.**
5. **(Optional) Select "Save Old Name."**

Select this option if you want the old name saved as a value in the "Other Names" field of the "Details" screen.

With the old name saved, users who don't know the object's new name can search for the object under the old name.

6. **(Optional) If you are renaming a container object, select "Create Alias in Place of Renamed Container."**



If you don't create an alias, users who are unaware that the container has been renamed cannot easily find the object in the Directory tree. Also, users whose name context in their NET.CFG file references the renamed container might be unable to log in.

7. To save the changes and return to the browser, choose “OK.”

If you renamed a container object, you should use the NCUPDATE utility to update the name context of users in the renamed container. For instructions, see “NCUPDATE” in *Utilities Reference*.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Alias objects	“Alias object” in <i>Concepts</i>
Directory tree	“Directory tree” in <i>Concepts</i>
NCUPDATE	“NCUPDATE” in <i>Utilities Reference</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Renaming Objects Using NETADMIN


Prerequisites



Checklist

- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Rename object right to the object that you want to rename

Procedure



Procedure

1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. From the “NETADMIN Options” menu, choose “Manage Objects.”

Your current context appears in the upper left corner of the screen.

3. Select the object that you want to rename.

- ◆ If the object appears on the list, select it and press <F10>.
- ◆ If the object is not on the list, browse the Directory by selecting container objects and pressing <Enter> until you see the object you want. Select it and press <F10>.

4. Choose “Rename.”

5. In the “New Name” field, enter a new name for the object.

6. Choose “Yes” to save the old name; choose “No” to discard the old name.

Choose “Yes” if you want the old name saved as a value in the “Other Names” field of the “Details” screen.

With the old name saved, users who do not know the object’s new name can search for the object under the old name.

7. To save your changes, press <F10>.

8. To confirm you want to save the new name, choose “Yes.”

9. (Optional) To create an alias in place of the renamed container, choose “Yes.”



If you don’t create an alias, users who are unaware that the container has been renamed cannot easily find the object in the Directory tree. Also, users whose name context in their NET.CFG file references the renamed container might be unable to log in.

10. Press <Esc> until you return to the “NETADMIN Options” menu.

If you renamed a container object, you should use the NCUPDATE utility to update the name context of users in the renamed container. For instructions, see “NCUPDATE” in *Utilities Reference*.

Additional Information

For more information about	Refer to
Objects	"Object" in <i>Concepts</i>
Rights	"Rights" in <i>Concepts</i>
Alias objects	"Alias object" in <i>Concepts</i>
NCUPDATE	"NCUPDATE" in <i>Utilities Reference</i>
Using the NETADMIN utility	"NETADMIN" in <i>Utilities Reference</i>

Changing Object Property Values

You can use the NetWare Administrator or NETADMIN utility to change object property values.

In this section, the general procedures for changing object property values are followed by tables that describe how to change specific properties of specific objects.

For a complete list of all properties for all objects, see Appendix A, "NDS and Bindery Objects and Properties," in *Utilities Reference*.


Changing Object Property Values Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- The Create object right to the container of the object whose property value will be changed

Procedure



1. Choose the "NetWare Administrator" icon from the MS Windows Program Manager or the OS/2 desktop.
2. From the browser, select the object whose property values you want to change.

For information on moving around in the browser and selecting objects, choose the "Help" button.

3. Choose "Details" from the "Object" menu.

The "Identification" page appears.

4. Make any necessary changes.

See Table 1-8 and Table 1-9 for details on object property values you can change.

Click on "Help" for details on each field and option.

Do not choose "OK" until you have entered all the changes you want to make to every page of the object dialog.

Do not choose "Cancel" unless you want to lose all changes made to every page of the object dialog.

Table 1-8 describes how to change User object property values.

Table 1-8
Changing User Object Property Value with NetWare Administrator

To	Choose this page of the object dialog; then...
Change the User's last name, other names, title, description, department, telephone number, fax number, or email address	"Details"; then enter information in the fields.
Disable the User's account	"Details"; then choose "Login Restrictions," and select (check) "Account Disabled."
Change the User's password	"Details"; then choose "Password Restrictions," choose "Change Password," and then enter information in the fields.

Table 1-8 *continued*

Changing User Object Property Value with NetWare Administrator

To	Choose this page of the object dialog; then...
Set the User's account expiration date or number of concurrent connections	"Details"; then choose "Login Restrictions," check "Account Has Expiration Date" or "Limit Concurrent Connections," and then enter information in the fields.
Change the User's password restrictions (password length, periodic changes, uniqueness, grace logins)	"Details"; then choose "Password Restrictions," and enter information in the fields.
Restrict the User's login times	"Details"; then choose "Login Time Restrictions," and choose the days and hours that the User is restricted from logging in to the network.
Set the User's login address restrictions	"Details"; then choose "Network Address," select the appropriate protocol, choose "Add," and then enter information in the fields.
Change the User's login script	"Details"; then choose "Login Script," and enter login script commands.
Specify the User's Profile login script	"Details"; then choose "Login Script," choose the browser button next to the "Profile" field, and then select a new Profile object from the "Select Object" dialog.
Add the User to an existing Group	"Details"; then choose "Group Membership," choose "Add," and then select a Group from the "Select Object" dialog.
Unlock a User's account	"Details"; then choose "Intruder Lockout," and enter information in fields to reset the User's account.
View, reset, or set up intruder detection on a user's account	"Details"; then choose "Intruder Lockout," and view or reset information in fields to detect intruder attempts.
	To set up intruder detection, you must go to the "Details" page of the container that the users reside in; then choose "Intruder Detection," and enter information in the fields.
Change a User's mailing address	"Postal Address"; then enter information in the fields.
Set up the User's security equivalences	"Security Equivalence"; then choose "Add," and select an object from the "Select Object" dialog.

Table 1-9 describes how to change other object property values.

Table 1-9

Changing Other Object Property Values with NetWare Administrator

To	Select this class of object	Choose this page of the object dialog; then...
Change information about the object's description, location, department, organization, etc.	Any class	"Details"; then enter information in the fields.
Change the object's mailing address	Organization, Organizational Unit	"Details"; then choose "Postal Address," and enter information in the fields.
Change the object's login script	Organization, Organizational Unit, Profile	"Details"; then choose "Login Script," and enter login script commands.
View, reset, or set up intruder detection for the object	Organization, Organizational Unit	"Details"; then choose "Intruder Detection," and enter information in the fields to detect intruder attempts.
Change the list of operators or resources associated with this object	AFP Server, Computer, NetWare Server	"Details"; then choose "Operators," "Supported Services," "Resources," or "User"; click on "Add," and then select other objects from the "Select Object" dialog.
Set the object's network address and protocol	AFP Server, Computer, NetWare Server, Volume	"Details"; then choose "Network Addresses"; select a network protocol, choose "Add," and then complete the fields associated with that protocol.

Table 1-9 *continued*

Changing Other Object Property Values with NetWare Administrator

To	Select this class of object	Choose this page of the object dialog; then...
Change the list of other objects affiliated with this object	Any class	"Details"; then choose "See Also," choose "Add," and then select other objects from the "Select Object" dialog.
Change the list of objects belonging to the Group	Group	"Details"; then choose "Members," choose "Add," and then select other objects from the "Select Object" dialog.
Change the object's volume time or usage information	Volume	"Details"; then choose "Statistics," "Dates and Times," or "User Space Limits"; then enter information in the fields.

5. To save the changes you made in all pages of the dialog, and return to the browser, choose "OK."

Additional Information

For more information about	Refer to
Objects	"Object" in <i>Concepts</i>
Object properties	"Property" in <i>Concepts</i>
Rights	"Rights" in <i>Concepts</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>


Changing Object Property Values Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later
- The Create object right to the container of the object whose property value will be changed

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility.

2. Choose “Manage Objects” from the “NETADMIN Options” menu.

3. Browse the Directory until the object you want appears.

Use the instructions at the bottom of the screen to browse the Directory. Press <F1> if you need help.

4. Select the object and press <F10>.

The “Actions” dialog appears.

5. Choose “View or Edit Properties of This Object.”

6. Select an option and make any necessary changes.

See Table 1-10 and Table 1-11 for details on object property values you can change.

Press <F1> for details on each field and option.

Table 1-10 describes how to modify User object property values.

Table 1-10

Changing User Object Property Values with NETADMIN

To	Choose this option from the "View or Edit Properties of This Object" menu; then...
Change the User's last name, other names, title, description, location, department, telephone number, fax number, or E-mail address	"Identification"; enter information in the fields.
Disable the User's account	"Account restrictions," and then choose "Login restrictions"; enter "Yes" in the "Account disabled" field.
Change the User's password	"Change password"; enter new password as prompted.
Set the User's account expiration date or number of concurrent connections	"Account restrictions," and then choose "Login restrictions"; enter information in the fields.
Change the User's password restrictions (password length, changes, uniqueness, grace logins)	"Account restrictions," and then choose "Password restrictions"; enter information in the fields.
View, reset, or set up intruder detection on a user's account	<p>"Account restrictions," and then choose "Intruder lockout status"; view or reset information in fields to detect intruder attempts.</p> <p>To set up intruder detection, you must go to the "View or Edit Properties of This Object" page of the container that the users reside in; then choose "Intruder Detection," enter information in the fields.</p>
Restrict the User's login times	"Account restrictions," and then choose "Login time restrictions"; enter information in the fields.
Change the User's login script	"Login script"; enter login script commands.
Set the User's login address restrictions	"Account restrictions," and then choose "Net address restrictions"; press <Insert>, select the address type, and then enter information in the fields.
Specify the User's Profile login script	"Groups/Security Equal To/Profile"; choose "Profile," and then enter the name and context of the Profile object (or press <Insert> and select the name from the list).

Table 1-10 *continued*

Changing User Object Property Values with NETADMIN

To	Choose this option from the "View or Edit Properties of This Object" menu; then...
Add the User to an existing Group.	"Groups/Security Equal To/Profile"; select "Groups" and press <Enter>, and then press <Insert> and enter the Group name (or press <Insert> again and select the Group name from the list).
Give the User security equivalence to another User.	"Groups/Security Equal To/Profile"; select "Security Equal To" and press <Enter>, and then press <Insert> and enter the complete name (or press <Insert> again and select the name from the list).
Change the User's mailing address.	"Postal address"; enter information in the fields.
Change the User's account balance, credit, or allow unlimited balance.	"Account restrictions," and then choose "Account balance"; enter information in the fields.
Change the list of other objects affiliated with this User object.	"See also"; enter names (or press <Insert> and select names from the list).

Table 1-11 describes how to modify other object property values.

Table 1-11

Changing Other Object Property Values With NETADMIN

To	Select this class of object	Choose this option from the "View or edit properties of this object" menu; then...
Change the information about an object's description, location, department, organization, etc.	Any class except Print Queue, Print Server, Printer	"Identification"; enter information in the fields.
Change information about the object's mailing address	Organization, Organizational Role, Organizational Unit	"Postal address"; enter information in the fields.
Change the object's login script	Organization, Organizational Unit, Profile	"Login script"; enter login script commands.
Change the object's password	AFP Server	"Change password"; enter a new password.

Table 1-11 *continued*

Changing Other Object Property Values With NETADMIN

To	Select this class of object	Choose this option from the "View or edit properties of this object" menu; then...
Set up the object's accounting information	NetWare Server	"Accounting," and then choose "Submenus"; enter information in the fields.
Set up the object's intruder detection lockout	Organization, Organizational Unit	"Intruder detection"; enter information in the fields.
Set the object's network address and protocol	Computer	"Other properties," and then choose "Network address" and press <Enter>; press <Insert> and select an address type, and then enter information in other fields.
Change the list of other objects affiliated with this object	Any class except Print Queue, Print Server, Printer, User	"See also"; enter names of objects (or press <Insert> and select names from the list).
Change the list of objects belonging to the Group	Group	"Group members"; enter names of objects (or press <Insert> and select names from the list).
View the object's volume time or usage information	Volume	"Volume statistics," "Volume restrictions," "Volume features," or "Volume dates/times"; enter information in the fields.

7. To save your changes, press <F10>.

8. Press <Esc> until you return to the "NETADMIN Options" menu.

Additional Information

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Object properties	“Property” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>



chapter

2

Managing Directories, Files, and Applications

This chapter gives you step-by-step procedures for setting up your network file system and giving users access to the resources they need, while still keeping the files secure.

Planning Directory Structures

How can you design your file system to make network administration easiest? This section gives you some basic information and suggestions about organizing directories and files. See also “Suggestions for Creating Volumes” in Chapter 3 of *Installation*.

System-Created Directories

During installation, the following directories are automatically created.

- ◆ **SYS:DELETED.SAV** contains the files that are deleted before they are purged.
- ◆ **SYS:ETC** contains sample files to help you configure the server.
- ◆ **SYS:LOGIN** contains the programs necessary for users to log in to the network, such as LOGIN.EXE. It has a special subdirectory for OS/2 login programs and a subdirectory called NLS, containing subdirectories for each supported language for login message files.
- ◆ **SYS:MAIL** may or may not contain subdirectories or files.

If you upgrade your server from a previous version of the NetWare[®] operating system, existing users will still have directories here for their login scripts, but their login scripts will become properties of the new User objects.

If you create new users after upgrading, the new users won't have directories in SYS:MAIL.

SYS:SYSTEM contains NetWare operating system files as well as NetWare utilities and programs for the supervisor. SYS:SYSTEM also has an NLS subdirectory, containing subdirectories for each supported language for message files.

SYS:PUBLIC allows general access to the network and contains NetWare utilities and programs for network users. Like SYS:LOGIN, SYS:PUBLIC has a special subdirectory for OS/2 utilities and a subdirectory called NLS, containing the message files for utilities.

SYS:DOC contains electronic versions of the NetWare manuals. This directory is created if you install DynaText.*

High Capacity Storage System (HCSS) Directories

You can create HCSS directories specifically to contain files that migrate to optical disk. Migration begins when allocated space on the hard disk is filled to a set capacity. This feature can be used in place of or in conjunction with file compression and disk suballocation to manage volume space on a server.

HCSS directories can contain DOS program files, applications, batch files, and data files. In HCSS, directories and files are moved back and forth between hard disk and optical disk to optimize the server's storage capacity.

For more information about HCSS directories, see Chapter 6, "Migrating Data Using the High Capacity Storage System." See also "High Capacity Storage System" in *Concepts*.

Workstation Operating System Directories

You may want to put the workstation operating system files on the network to save workstation disk space or to make diskless workstations possible.

Since the workstation operating system files don't normally change, you can keep them on one set of backup diskettes and then skip these directories when you do network backups.

For information about loading operating system software on a network, see "Loading Operating Systems and Applications onto the Network" on page 117.

Application Directories

For ease of management, you should keep application files in a different directory than data files.

Since the application programs don't normally change, you can keep one set of application files on backup diskettes and then skip the application directories when you do network backups.

For more information about loading applications on the network, see "Loading Operating Systems and Applications onto the Network" on page 117.

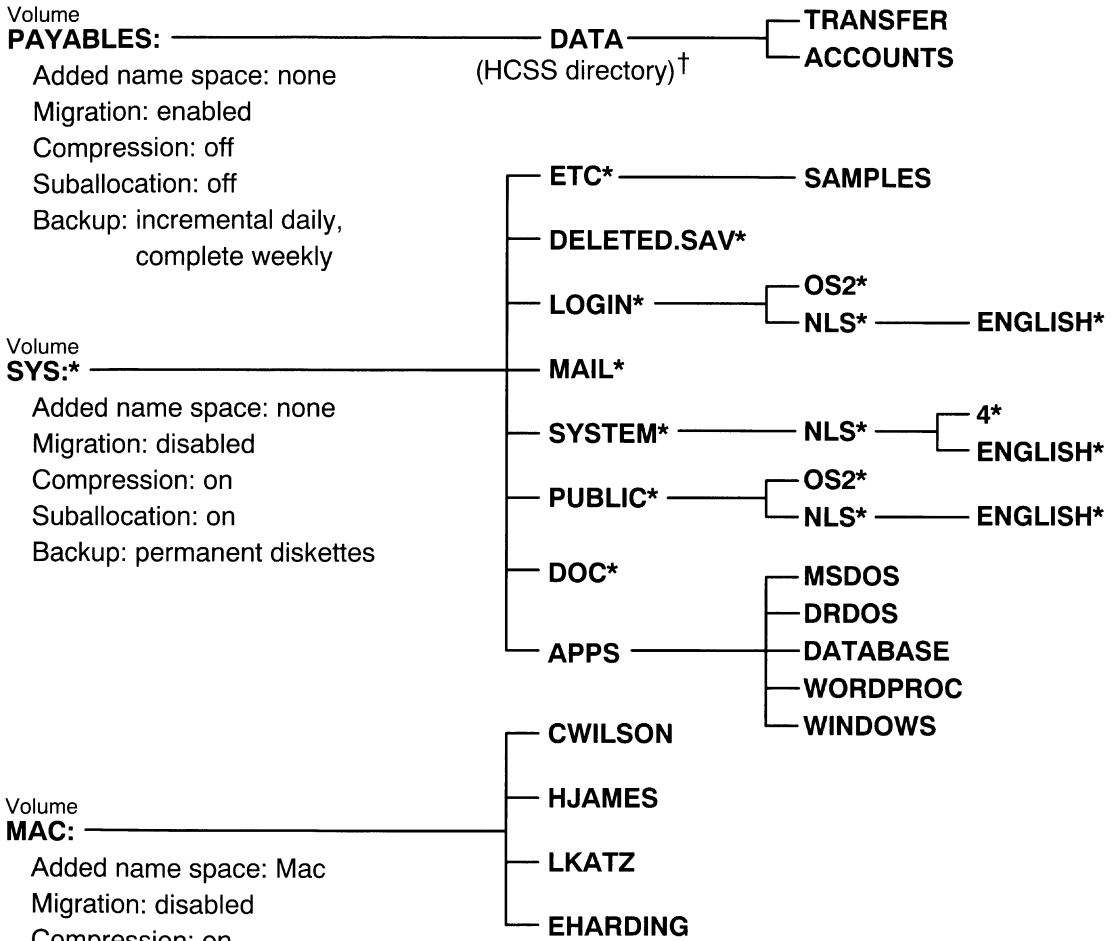
Data Directories

Data directories can include work areas where groups or users keep work files. You can also create a directory to serve as a transfer point for copying files to and from other areas of the network.

Sample Directory Structure

A sample directory structure is shown in Figure 2-1.

Figure 2-1
Sample Network
Directory Structure



* Created during server installation.

† See Chapter 6, "Migrating Data Using the High Capacity Storage System."

Creating Directories and Copying Files

For each NetWare volume, you can create directories and subdirectories to organize data and applications. (If you haven't created the volume, see "Maintaining Volumes" on page 515.)

Use one of several tools to divide a NetWare volume into directories. For example, use the NetWare Administrator graphical utility, the FILER text utility, or the DOS MKDIR (MD) command.

Two methods of creating directories and copying files (NetWare Administrator and FILER) are documented in this section.

Creating Directories Using NetWare Administrator

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Create right to the parent directory of the new directory

Procedure



1. Choose the "NetWare Administrator" icon from the Microsoft Windows (MS Windows) Program Manager or the OS/2 desktop.

2. Using the browser, select the directory to be the parent directory of the one you want to create.

To create a root directory, select the Volume object. For more about moving around in the browser and selecting objects, choose "Help" from the menu bar. For a list of directories and files on a volume, select "View/Set Context" and specify the volume name.

3. Choose "Create" from the "Object" menu.

4. Enter the name for the new directory in the "Directory Name" box.

5. (Optional) Choose “Define Additional Properties” if you want to enter properties for the directory.

Choose this item if you want to assign properties to the new directory now. You can also add or modify properties after the directory is created. Properties include items such as trustees and access rights.

6. (Optional) Choose “Create Another Directory” to create another directory immediately after this one.

7. Choose “Create” to create a directory with the name you entered.

The new directory is created. You are returned to the browser.

Additional Information

For more information about	Refer to
Creating HCSS directories	“Setting Up the HCSS File System” on page 405
Directories for workstation operating systems	“File system” in <i>Concepts</i>
File system rights	“Attributes,” “Effective rights,” “Inherited Rights Filter,” “Rights,” and “Security” in <i>Concepts</i>
Assigning directory and file rights	“Making the File System Secure and Accessible” on page 131
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Creating Directories Using FILER

Prerequisites

Checklist



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Create right to the parent directory of the new directory

Procedure

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen. Make sure you are in the correct location in the Directory tree when you create the new directory.

2. Select “Manage Files and Directories.”

The “Directory Contents” list appears.

3. Find and select the volume, directory, or subdirectory in which you want to create a new directory or subdirectory by completing the following steps:

- ◆ If the item you want appears on the list, select it and press <Insert>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and pressing <Enter> until you see the item you want. Select it and press <Insert>.
- ◆ If you can't find what you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the “Available Options” menu and choosing “Select Current Directory.”

4. When you are prompted, type the name of the new directory and press <Enter>.

The new directory is created. After you create it, you can change default attributes, assign rights, and make trustee assignments for the directory or subdirectory.

5. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Assigning directory and file rights	"Making the File System Secure and Accessible" on page 131
File system rights	"Understanding File and Directory Rights" on page 131 "Attributes," "Effective rights," "Inherited Rights Filter," "Rights," and "Security" in <i>Concepts</i> "Adding a Trustee to a Directory or File" on page 134
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Copying or Moving Files Using NetWare Administrator

If you use the NetWare Administrator to copy a Macintosh file between volumes, you lose extended attribute information on the file unless the two volumes have identical sets of name spaces loaded.

For example, if volume 1 has MAC and DOS name spaces, and volume 2 has MAC, DOS, and OS/2, and you copy a Macintosh file from one volume to another, you lose extended attributes even though both volumes have the MAC name space.

Prerequisites

Checklist



- A 386 or 486 workstation running NetWare Administrator.
- The File Scan right to the source directory and the Create right to the destination directory. In addition, to move files (delete them from the source directory), you need the Erase right.

Procedure

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Using the browser, select the directory, and then select one or more files in that directory that you want to copy or move.**

To select multiple files, press <Ctrl> on the keyboard while clicking on the files.

For information about moving around in the browser and selecting objects, choose “Help” from the menu bar.
- 3. Choose “Move” or “Copy” from the “Object” menu.**

A list appears of the files you selected.
- 4. Select either “Move” or “Copy” from the “Operation” dialog.**
- 5. Choose the browser button next to the “Destination” field.**
- 6. Using the browser, select a directory in a volume for the destination of the directories or files.**
- 7. Choose “OK.”**

The destination directory is listed in the “Destination” field.
- 8. Choose “OK.”**

The listed files are copied or moved to the destination directory

Additional Information

For more information about	Refer to
File system directories and files	"File system" in <i>Concepts</i>
File system rights	"Understanding File and Directory Rights" on page 131 "Attributes," "Effective rights," "Inherited Rights Filter," "Rights," and "Security" in <i>Concepts</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>


Copying or Moving Files Using FILER

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 512 KB of memory available on the workstation.
- The File Scan right to the source directory and the Create right to the destination directory. In addition, to move files (delete them from the source directory), you need the Erase right.

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, volume object, and path are shown in the upper left corner of the screen.

2. Select "Manage Files and Directories."

The "Directory Contents" list appears.

3. Find the volume, directory, or subdirectory that contains the source files.

Browse the directory structure by selecting a directory or subdirectory and pressing <Enter> until you see the directory or files you want to copy.

If you can't find the directory or files you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

Depending on the operation you want to perform next, go to either Step 4 or Step 5.

4. (Optional) If you want to move or copy a directory or subdirectory and all its contents, complete the following steps:

4a. Select the directory or subdirectory and press <F10>.

4b. Select the option you want from the "Subdirectory Options" menu.

4c. At the prompt, type in the complete path to the destination directory, or press <Insert> to browse for the destination directory.

4d. Press <F10> to start the copy or move process.

The progress of the transaction is shown in an information box.

5. (Optional) To copy specific files, complete the following steps:

5a. Select the directory or subdirectory containing the files and press <Enter>.

5b. Press <F5> to mark files you want to copy.

Browse the directory structure by selecting a directory or subdirectory and pressing <Enter> until you see the directory or files you want to copy.

If you can't find the directory or files you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

- 5c. After all of the files are marked, press <F10>.
- 5d. From the “Multiple File Operations” menu, select “Copy Marked Files.”
The “Copy Files to” box appears.
- 5e. Type in the complete path to the destination directory or press <Insert> to browse for the destination directory.
- 5f. Press <F10> to start copying the files.

The progress of the transaction is shown.

- 6. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Creating HCSS directories	“Setting Up the HCSS File System” on page 405
Directories and files	“File system” in <i>Concepts</i>
File system rights	“Understanding File and Directory Rights” on page 131 “Attributes,” “Effective rights,” “Inherited Rights Filter,” “Rights,” and “Security” in <i>Concepts</i>
Using the FILER utility	“FILER” in <i>Utilities Reference</i>

Loading Operating Systems and Applications onto the Network

You may want to load workstation operating system (OS) files onto the network to save workstation disk space or to allow diskless workstations to log in to the network.

Loading DOS onto the Network


Prerequisites



Checklist

- A workstation running DOS 3.30 or later
- The Supervisor or Create right to the directory where you want to load the application

Procedure



Procedure

1. **Create a directory in SYS:PUBLIC for each workstation type and version of DOS you will be using on your network.**

Name your directories according to the following convention:

SYS:PUBLIC/*machine/os_type/os_version*

For each directory, replace *machine* with the six-letter machine name of the workstation (such as IBM_PC or COMPAQ).

Replace *os_type* with the type of DOS you are using (such as MSDOS or DRDOS). Replace *os_version* with the DOS version number.

For example, to install MS-DOS* 5.0 on an IBM PC, go to SYS:PUBLIC to create the DOS subdirectory. To use the DOS MD command, type the commands shown below (in order):

```
MD IBM_PC <Enter>
CD IBM_PC <Enter>
MD MSDOS <Enter>
CD MSDOS <Enter>
MD 50 <Enter>
CD 50 <Enter>
```

2. Load DOS.

Follow the instructions in the DOS documentation to load the DOS software into the directories you created.

3. In the system login script, map the second search drive to the DOS directory.

If all users have the same types of workstations and are using the same version of DOS, you will probably have only one DOS directory. In this case, add a line similar to the following, substituting the correct directory names:

```
MAP S2:=SYS:PUBLIC\IBM_PC\MSDOS\50
```

If your network has more than one DOS directory, use variables to indicate the directory path. These variables are replaced by the correct information from the workstation software when each user logs in. Enter the following command, exactly as shown, in the login script:

```
MAP S2:=SYS:PUBLIC\%MACHINE%\%OS%\%OS_VERSION
```

(The first search drive should be mapped to the PUBLIC directory so that users can access NetWare utilities.) For the %MACHINE variable to work, make sure the long machine type is set in each station's NET.CFG file. For example, a station might have the following line in its NET.CFG file:

```
LONG MACHINE TYPE = IBM_PC
```

4. Add COMSPEC to the system login script.

Following is the proper syntax for the COMSPEC command:

```
COMSPEC=Y:COMMAND.COM
```

This command tells the workstation where to find the command processor

Additional Information

For more information about	Refer to
Loading DOS on the network	The DOS manual that came with your software
Login scripts	Chapter 3, "Creating Login Scripts"
Mapping search drives	"Mapping Search Drives" on page 215
Running DOS on a workstation	<i>NetWare Client for DOS and MS Windows User Guide</i>

Loading MS Windows onto the Network

You can load MS Windows onto a network in several ways:

- ◆ Load all MS Windows files on a user's local hard drive.
- ◆ Load MS Windows program files on the server, and load user files on local hard drives.
- ◆ Load all MS Windows program and user files on the server. In most cases, this installation provides the easiest maintenance and most efficient use of resources.

Instructions for the first two options are included in the documentation that came with your MS Windows software. Instructions for loading program and user files on the server are included in this section.

Following are some advantages of having all MS Windows files installed on the server:

- ◆ Program and configuration files are backed up and secure.
- ◆ No hard disk is required on the user workstation.
- ◆ The .INI and driver files for all users can be updated from one location.
- ◆ Configuration files always match hardware.

The only real disadvantage of having all MS Windows files installed on the server is that it causes more network traffic.

Files Needed for NetWare Functionality in MS Windows

Use the following files to run MS Windows with NetWare functionality. These files are in the Workstation for DOS/Windows kit.

File	Function
NETWARE.DRV	NetWare device driver. Contains executable code for NetWare-related functions.
NETWARE.HLP	Help file for NETWARE.DRV.
NETWARE.INI	Initialization file for NETWARE.DRV and other MS Windows utilities for NetWare. Automatically created by NETWARE.DRV.
NWPOPUP.EXE	Handler for broadcast messages.
VNETWARE.386	Virtual NetWare device driver. Performs virtualization among sessions when MS Windows is in 386 Enhanced mode.
VIPX.386	Virtual IPX device driver. Virtualizes IPX communications among sessions when MS Windows is in 386 Enhanced mode.

Guidelines for Running MS Windows on the Network

- ◆ Use a permanent swap file on a local hard drive if possible; do *not* use network directories for swap files. If a local swap file is not possible, consider increasing RAM to a minimum of 8 MB.
- ◆ RAM plus swap file size should be a minimum of 10 MB.
- ◆ Remove MS Windows search drives from the workstation AUTOEXEC.BAT file.

- ◆ If users are running DR DOS, make sure DR DOS files are dated 07 April 1992 or later.
- ◆ If users are running DR DOS, make sure their CONFIG.SYS files install EMM386.SYS and configure it for use with MS Windows. See your DR DOS manual or DR DOS online help for configuration information.

Loading and Setting Up MS Windows

Prerequisites

Checklist



- A workstation running DOS 3.30 or later
- The Supervisor or Create right to the directory where you are loading MS Windows
- A licensed network copy of MS Windows 3.1 or later
- 16 MB of available disk space
- NetWare Client for DOS and MS Windows

Procedure

Procedure



- 1. Install MS Windows server software, using the SETUP /A option.**

For complete information about the MS Windows SETUP options and the installation procedure, see the documentation that came with your MS Windows software.

- 2. Create a Group object for users who will be running MS Windows.**

For instructions for creating a Group object, see “Managing Group Objects” on page 49.

- 3. Make the MS Windows Group object a trustee of the MS Windows directory.**

For instructions on granting trustee assignments, see “Adding a Trustee to a Directory or File” on page 134.

4. Create a directory for each user to store user-specific MS Windows files.

WIN.COM and files such as .GRP and .INI files are stored here.

5. Add the following information to the system login script.

An example of the syntax for these login script entries appears at the end of this step.

5a. Map a drive to the user-specific directories for the MS Windows group.

5b. Map a search drive to the MS Windows directory for the MS Windows Group object.

5c. Set the MS Windows TEMP directory to a subdirectory of the user directory.

The following example shows the system login script entries you would add to set up MS Windows 3.1 on the network:

```
IF MEMBER OF "WIN31" THEN
MAP INS P:=SYS:USERS\%LOGIN_NAME\WIN31
MAP INS S16:=SYS:APPS\WINAPPS\WIN31
SET TEMP = "P:\USERS\%LOGIN_NAME\WIN31\TEMP"
END
```

6. Set up the workstations by completing the following steps.

6a. Change to the search drive mapped to the MS Windows directory.

Enter the drive letter only.

6b. Modify the user AUTOEXEC.BAT and CONFIG.SYS files by typing:

```
SETUP /N <Enter>
```

Select the "Custom" option to make sure environment variables are correct. When MS Windows prompts for a path during setup, enter the drive letter instead of the path.

6c. Install the NetWare Workstation for MS Windows software.

Follow the instructions in *NetWare Client for DOS and MS Windows User Guide* to install the workstation software and update NetWare-specific files in the MS Windows directory.

Additional Information

For more information about	Refer to
Creating login scripts	Chapter 3, "Creating Login Scripts"
Installing MS Windows on the network	The MS Windows manual that came with your software
Mapping search drives	"Mapping Search Drives" on page 215
Running MS Windows on a workstation	<i>NetWare Client for DOS and MS Windows User Guide</i>

Loading Other Applications onto the Network

You can load various types of network applications, such as word processing or spreadsheet programs, to make them available to users. When loading applications, keep the following in mind:

- ◆ You need the Create right in the directory where you will be loading the application.
- ◆ Follow the instructions in the application's documentation for loading the application onto a network.
- ◆ Make sure the application is designed for network (multiuser) use, and that you observe any licensing restrictions on the number of users who can access the application.
- ◆ To allow users to access network-based applications, map search drives to the directories that contain these applications. To make these search drives permanent, place them in login scripts, which are executed when users log in.

- ◆ If the application requires that it be installed at the root of a volume, but you would rather install it in a subdirectory for security reasons, you can map the directory to a fake root. To map a fake root directory, use the MAP ROOT command as explained in “MAP” on page 212.
- ◆ You can create a Directory Map object that points to an application directory.

Directory Map objects are useful in login scripts. Instead of mapping a drive to a specific directory path, you map a drive to a Directory Map object that points to a directory.

Then, if you change the directory path, you need to change only the Directory Map object’s definition.

Additional Information

For more information about	Refer to
Creating login scripts	Chapter 3, “Creating Login Scripts”
Mapping search drives and fake roots in login scripts	“MAP” on page 212

Assigning Trustee Rights to Operating System and Application Directories

Application programs need a set of file system rights that make them available to users, yet protect them from being corrupted. Keep the following guidelines in mind as you assign file and directory rights to executable, application, and program files.

When you are ready to assign file and directory attributes and rights, see “Making the File System Secure and Accessible” on page 131.

Use the FLAG utility to assign the following rights:

- ◆ Assign the Shareable and Read-Only attributes to application and workstation operating system files.
- ◆ Assign the Execute Only attribute to executable files for which you keep permanent backups.
- ◆ Assign Read and File Scan rights to the User or Group objects that need to use the application.

Additional Information

For more information about	Refer to
Understanding file system rights	“Effective rights” and “Rights” in <i>Concepts</i>
Assigning file and directory rights	“Changing Attributes of a Directory or File” on page 147
Using the FLAG utility	“FLAG” in <i>Utilities Reference</i>

Creating and Using Directory Map Objects

How to Use Directory Map Objects

A Directory Map object represents a particular directory in the file system.

Directory Map objects can be especially useful in login scripts by indicating directories that contain applications or other frequently used files.

For example, if you have a directory that contains DOS 5.0, you will probably map a search drive to that directory in any login scripts you create.

If you should later upgrade to DOS 6.0 and rename the directory, you would have to change the mapping in every login script where that search mapping appears.

By using a Directory Map object, you could avoid having to make changes to the login scripts.

First, you could create a Directory Map object called `CURRENT_DOS` that points to the DOS directory (`SYS:PUBLIC\IBM_PC\MSDOS\5.0`).

Then, in a MAP command in your login scripts, map a search drive to the Directory Map object, rather than to the specific directory:

```
MAP INS S2:=.CURRENT_DOS.SALES.NOVELL_US
```

When users log in, their search drive is mapped to the `CURRENT_DOS` Directory Map object, which points to the directory containing DOS 5.0.

Later, if you upgrade to DOS 6.0 and change the directory's name to `SYS:PUBLIC\IBM_PC\MSDOS\60`, you would change only the Directory Map object to indicate the new path.

You would not have to change the MAP command in the login script because the MAP command still indicates the correct Directory Map object.

Additional Information

For more information about	Refer to
Creating login scripts	Chapter 3, "Creating Login Scripts"
Mapping drives to Directory Map objects in login scripts	"MAP" on page 212
Using the MAP utility	"MAP" in <i>Utilities Reference</i>

Creating a Directory Map Object

You can create a Directory Map object using either the NetWare Administrator graphical utility or the NETADMIN text utility. Both procedures are documented in this section.


Using NetWare Administrator to Create a Directory Map Object

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Create object right to the container where the Directory Map object will be created

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. From the browser, select the Organization or Organizational Unit object that will contain the Directory Map object.**

For information about moving around in the browser and selecting objects, choose “Help” from the menu bar.

- 3. Choose “Create” from the “Object” menu.**
- 4. Choose “Directory Map” from the “New Object” dialog box.**
- 5. Choose “OK.”**
- 6. Enter the name for the Directory Map object in the space provided.**
- 7. In the “Volume” field, enter the name of the volume this Directory Map object will point to.**

You can type in the complete name of the Volume object, or you can choose the browser button to the right of the “Volume” field to browse for the Volume object.

8. In the “Path” field, enter the path of the directory this Directory Map object will point to.

You can type in the path, or you can choose the browser button to the right of the “Path” field to browse for the directory that the Directory Map object will point to.

If the Directory Map object will point to the root of the specified volume, leave the “Path” field blank.

9. (Optional) To define additional properties immediately after creating the Directory Map object, choose “Define Additional Properties.”

10. (Optional) To create another Directory Map object immediately after this one, choose “Create Another Directory Map.”

11. Choose “Create.”

The Directory Map object is created

If you chose “Create Another Directory Map,” the Create dialog appears.

If you chose “Define Additional Properties,” the identification screen appears.

Additional Information

For more information about	Refer to
Directory Map objects	“Directory Map object” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using NETADMIN to Create a Directory Map Object

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Create object right in the container where the Directory Map object will be created

Procedure



1. **At the DOS prompt, type**
NETADMIN <Enter>
2. **Choose “Manage Objects” from the “Netadmin Options” menu.**
3. **Select the Organization or Organizational Unit object that will contain the Directory Map object.**

Browse the Directory tree by selecting objects and pressing <Enter>. When you find the object that you want to contain the Directory Map object, select it and then press <Enter>.

4. **Press <Insert>.**
5. **From the “Select an Object Class” menu, select “Directory Map.”**
6. **Type the name for the Directory Map object and press <Enter>.**
7. **In the “Volume Object Name” field, press <Insert>.**
8. **Enter the name of the Volume object this Directory Map object will point to.**

You can either type the Volume object’s complete name or you can press <Insert> to browse through the Directory tree, select an object, and press <Enter>.

9. **In the “Path on Volume” field, press <Insert>.**

10. Enter the path of the directory this Directory Map object will point to.

You can either type the full path or you can press <Insert> to browse the file system directory structure in the volume defined in Steps 7 and 8; then press <Enter>.

If the Directory Map object will point to the root of the specified volume, leave the "Path" field blank and press <Enter>.

11. From the "Name Spaces" box, select the type of name space for the volume and press <Enter>.

12. To save the changes, press <F10>.

Additional Information

For more information about	Refer to
Directory Map objects	"Directory Map object" in <i>Concepts</i>
Using the NETADMIN utility	"NETADMIN" in <i>Utilities Reference</i>

Making the File System Secure and Accessible

Understanding File and Directory Rights

File system security includes assigning trustee rights and setting file and directory attributes. These two types of security are discussed in the following sections.

Trustee Rights

Trustee rights determine the access users have to directories and files. These rights can be given to User objects, Group objects, or Organizational Role objects.

Trustee rights are explained in Table 2-1.

Table 2-1
Trustee Rights

Right	Allows you to
Access Control	Add and remove trustees and change rights to files and directories.
Create	Create subdirectories and files.
Erase	Delete directories and files.
File Scan	View file and directory names in the file system structure.
Modify	Rename directories and files, and to change file attributes.
Read	Open and read files, and to open, read, and execute applications.
Supervisor	Grant all rights listed in this table.
Write	Open, write to, and modify a file.

Directory and File Attributes

Directory and file attributes assign properties to individual directories or files. Some attributes are meaningful only when applied at the file level, but some apply to both the directory and the file levels.

Be careful when assigning directory and file attributes. The attribute applies to *all* users.

For example, if you assign the Delete Inhibit attribute to a file, no one, including the owner of the file or the system supervisor, can delete the file. But any trustee with the Modify right can change the attribute to allow deletion.

Directory and file attributes are explained in Table 2-2.

Table 2-2

Directory and File Attributes

Attribute	Description	Applies to
A	Archive Needed identifies files that have been modified since the last backup. This attribute is assigned automatically.	Files only
Ci	Copy Inhibit prevents Macintosh users from copying a file. This attribute overrides Read and File Scan trustee rights.	Files only
Dc	Don't Compress keeps data from being compressed. This attribute overrides settings for automatic compression of files not accessed within a specified number of days.	Directories and files
Di	Delete Inhibit means that the file or directory cannot be deleted. This attribute overrides the Erase trustee right.	Directories and files
Dm	Don't Migrate prevents files and directories from being migrated from the server's hard disk to another storage medium.	Directories and files
Ds	Don't Suballocate prevents data from being suballocated.	Files only
H	The Hidden attribute hides files and directories so they can't be listed using the DIR command. A user with File Scan rights can use FILER or the NDIR command to list directories and files with the Hidden attribute.	Directories and files
I	Index allows large files to be accessed quickly by indexing files with more than 64 File Allocation Table (FAT) entries. This attribute is set automatically.	Files only

Table 2-2 continued

Directory and File Attributes

Attribute	Description	Applies to
Ic	Immediate Compress sets data to be compressed as soon as a file is closed. If applied to a directory, every file in the directory is compressed as each file is closed.	Directories and files
N	Normal indicates the Read/Write attribute is assigned and the Shareable attribute is not. This is the default attribute assignment for all new files.	Directories and files
P	Purge flags a file or directory to be erased from the system as soon as it is deleted. Purged files and directories cannot be recovered.	Directories and files
Ri	Rename Inhibit prevents the file or directory name from being modified.	Directories and files
Ro	Read Only prevents a file from being modified. This attribute automatically sets Delete Inhibit and Rename Inhibit.	Files only
Rw	Read/Write allows you to write to a file. All files are created with this attribute.	Files only
Sh	Shareable allows more than one user to access the file at the same time. This attribute is usually used with Read Only.	Files only
Sy	The System attribute hides the file or directory so it can't be seen by using the DIR command. It can be seen if a user with File Scan rights uses FILER or the NDIR command. System is normally used with operating system files, such as DOS system files.	Directories and files
T	Transactional allows a file to be tracked and protected by the Transaction Tracking System (TTS).	Files only
X	The Execute Only attribute prevents the file from being copied, modified, or backed up. It does allow renaming. The only way to remove this attribute is to delete the file. Use the attribute for program files such as .EXE or .COM. Make a copy of a file before you flag it as Execute Only, so you can replace the file if it becomes corrupted.	Files only

Adding a Trustee to a Directory or File

You can add a trustee to a directory or file using either the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.

Using NetWare Administrator to Add a Trustee

Checklist



Prerequisites

- A 386 or 486 workstation running NetWare Administrator
- The Access Control right to the file or directory you want to add the trustee to

Procedure



Procedure

- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Using the browser, select the directory or file that you want to add a trustee to.**

For information about moving around in the browser and selecting objects, choose “Help” from the menu bar.

- 3. From the “Object” menu, choose “Details.”**
- 4. From the “Identification” page, choose “Trustees of this Directory.”**
- 5. From the “Trustees of this Directory” page, choose “Add Trustee.”**
- 6. Select a trustee from the list.**

If the object does not appear in the list, browse the Directory tree to find the object that you want to make a trustee of the file or directory.

- 7. Choose “OK.”**

8. To grant rights to the trustee, mark the appropriate check boxes below the trustee.

9. To return to the browser, choose "OK."

Additional Information

For more information about	Refer to
Trustees	"Trustee" in <i>Concepts</i>
File system rights	"Trustee Rights" on page 131
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>


Using FILER to Add a Trustee

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available in the workstation
- Access Control right to the file or directory you want to add the trustee to

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select "Manage Files and Directories."

The "Directory Contents" list appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and press <F10>.
- ◆ If the item is not in the list, browse a directory or subdirectory by selecting it and pressing <Enter> until you see the item you want. Select it and press <F10>.
- ◆ If you can't find the directory you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

4. Select "View/Set File [or Directory] Information" and press <Enter>.

Information for that file or directory appears.

5. Use the arrow keys to move to the "Trustees" field, and then press <Enter>.

A list of trustees for that file or directory appears.

6. To add a trustee, press <Insert> and locate the trustee's name in the list. Select the name and then press <Enter>.

The new trustee, object type, and default rights appear in the list.

7. (Optional) To add another trustee to this file or directory, press <Esc> until you get to the "File Information" screen. Then repeat Steps 5 and 6.

8. (Optional) Assign rights to the new trustee.

You can assign or modify trustee rights now, or at any time after the trustee has been assigned to the directory or file.

8a. From the trustee list, select the user you want to assign or modify rights for, and then press <Enter>.

The "Trustee Rights" list appears, showing the rights the trustee currently has to this directory or file.

8b. Press <Insert> to see a list of rights you can assign.

- 8c. Select a right you want to give the trustee and press <Enter>. To give the trustee more than one right, press <F5> to mark each right, and then press <Enter>.**

The "Trustee Rights" list reappears with the new rights added.

- 8d. Press <Esc>.**

The new rights appear next to the trustee name.

- 9. Exit FILER by pressing <Esc> until you reach the Exit confirmation box, and then select "Yes."**

Additional Information

For more information about	Refer to
Trustees	"Trustee" in <i>Concepts</i>
File system rights	"Trustee Rights" on page 131
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Deleting a Trustee from a Directory or File

You can delete a trustee from a directory or file using either the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to Delete a Trustee

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Access Control right to the file or directory you want to delete the trustee from

Procedure



1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.
2. Using the browser, select a directory or file that you want to delete a trustee from.

For information on moving around in the browser and selecting objects, choose “Help” from the menu bar.

3. From the “Object” menu, choose “Details.”
4. From the “Identification” page, choose “Trustees of This Directory.”
5. From the “Trustees” list, select a trustee.
6. Choose “Delete Trustee.”
7. To delete that object as a trustee, choose “Yes.”
8. To return to the browser, choose “OK.”

Additional Information

For more information about	Refer to
Trustees	“Trustee” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to Delete a Trustee

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Access Control right to the file or directory you want to delete the trustee from

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select “Manage Files and Directories.”

The “Directory Contents” list appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and pressing <Enter> until you see the item you want. Select it and press <F10>.

If you can't find what you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the “Available Options” menu and choosing “Select Current Directory.”

4. Select “View/Set File [or Directory] Information” and press <Enter>.

Information for that file or directory appears.

5. Use the arrow keys to move to the “Trustees” field, and then press <Enter>.

A list of trustees for that file or directory appear.

6. Select the trustee you want to delete, and then press <Delete>.

You are prompted to delete that trustee from the directory.

7. Select “Yes.”

8. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Trustees	“Trustee” in <i>Concepts</i>
Using the FILER utility	“FILER” in <i>Utilities Reference</i>

Modifying a Trustee’s Rights to a Directory or File

You can modify trustee rights to a directory or file through the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to Modify a Trustee’s Rights

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Access Control right to the file or directory for which you want to change the trustee’s rights

Procedure



1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.
2. Using the browser, select the directory or file for which you want to change trustee rights.

For information on moving around in the browser and selecting objects, choose “Help” from the menu bar.

3. From the “Object” menu, choose “Details.”
4. From the “Identification” page, choose “Trustees of this Directory.”
5. From the “Trustees” list, select a trustee.
6. Grant or revoke rights by marking the check boxes below the trustee.
7. Choose “OK” to save the trustee rights.

Additional Information

For more information about	Refer to
File system rights	“Trustee Rights” on page 131 “Rights” in <i>Concepts</i>
Trustees	“Trustee” in <i>Concepts</i>
Using the NetWare Administrator	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to Modify a Trustee's Rights

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available in the workstation
- The Access Control right to the file or directory for which you want to change the trustee's rights

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select "Manage Files and Directories."

The "Directory Contents" list appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and pressing <Enter> until you see the item you want. Select it and press <F10>.
- ◆ If you cannot find the item you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by backing out to the "Available Options" menu and choosing "Select Current Directory."

4. Select "View/Set File [or Directory] Information" and press <Enter>.

5. **Using the arrow keys, move to the “Trustee” field, and then press <Enter>.**

6. **Select the name of the trustee whose rights you want to modify and press <Enter>.**

A list of the trustee’s current rights appears.

7. **Press <Insert> to see a list of rights you can assign.**

8. **Select a right you want to give the trustee, and then press <Enter>. If you want to assign more than one right, press <F5> to mark each right, and then press <Enter>.**

The “Trustee Rights” list reappears, showing the new list of rights.

9. **To exit, press <Esc>.**

The new rights appear next to the trustee name.

Additional Information

For more information about	Refer to
File system rights	“Trustee Rights” on page 131 “Rights” in <i>Concepts</i>
Trustees	“Trustee” in <i>Concepts</i>
Using the FILER utility	“FILER” in <i>Utilities References</i>

Viewing/Modifying the Inherited Rights Filter for Directories and Files

You can view and modify the Inherited Rights Filter (IRF) for a directory or file using the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to View/Modify an Inherited Rights Filter

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Access Control right to the file or directory for which you want to view or modify the IRF

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Using the browser, select a directory or file.**

For information on moving around in the browser and selecting objects, choose “Help” from the menu bar.
- 3. From the “Object” menu, choose “Details.”**
- 4. From the “Identification” page, choose “Trustees of This Directory.”**
- 5. Under “Inheritance Filter,” select the check boxes for the rights that you want to allow to be inherited for that directory or file.**
- 6. Choose “OK.”**

The “Trustees” dialog reappears.
- 7. To return to the browser, choose “OK.”**

Additional Information

For more information about	Refer to
Inherited rights	“Directory and File Attributes” on page 132 “Attributes,” “Inherited Rights Filter,” and “Rights” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to View/Modify the Inherited Rights Filter

Prerequisites



- A workstation running 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Access Control right to the file or directory for which you want to view or modify the filter

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select “Manage Files and Directories.”

The “Directory Contents” list appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and then pressing <Enter> until you see the item you want. Select it and press <F10>.
- ◆ If you can't find the item you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

4. Select "View/Set File [or Directory] Information" and press <Enter>.

Information for that file or directory appears. The current inherited rights are shown in the "Inherited Rights Filter" field.

5. Use the arrow keys to move to the "Inherited Rights Filter" field and then press <Enter>.

A list of the rights inherited by the file or directory appears.

6. Select a file or directory attribute you want to revoke, and then press <Delete>. To revoke more than one attribute, press <F5> to mark each attribute, and then press <Delete>.

7. Press <Esc>.

The "File [or Directory] Information" screen reappears with a listing of the rights that can be inherited.

8. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Inherited rights	"Directory and File Attributes" on page 132 "Attributes," "Inherited Rights Filter," and "Rights" in <i>Concepts</i>
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Changing Attributes of a Directory or File

You can change the attributes of a directory or file with the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to Change Attributes

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Modify right to the file or directory whose attributes you want to change

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Using the browser, select a directory or file.**

For information on moving around in the browser and selecting objects, choose “Help” from the menu bar.
- 3. From the “Object” menu, choose “Details.”**
- 4. From the “Identification” page, choose “Attributes.”**
- 5. Select the check boxes for the attributes that you want to set or reset for this directory or file.**
- 6. To close the object dialog box and save the new attributes, choose “OK.”**

Additional Information

For more information about	Refer to
File and directory attributes	“Directory and File Attributes” on page 132 “Attributes” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to Change Attributes

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Modify right to the file or directory whose attributes you want to change

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select “Manage Files and Directories.”

The “Directory Contents” list appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and then press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and then pressing <Enter> until you see the item you want. Select it and press <F10>.

- ◆ If you can't find what you want, check the volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

4. Select "View/Set File [or Directory] Information" and press <Enter>.

Information for the file or directory appears.

5. To modify an attribute, use the arrow keys to move to the "Attributes" field, and then press <Enter>.

The attributes for that file or directory appear.

6. Modify the attribute by completing one of the following steps:

- ◆ To delete an attribute, select it and press <Delete>. Select "Yes" when you are prompted to delete the attribute.
- ◆ To add an attribute, press <Insert>. Select the attribute you want to add and press <Enter>.

To assign more than one right, press <F5> to mark each right, and then press <Enter>.

7. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
File and directory attributes	"Directory and File Attributes" on page 132 "Attributes" in <i>Concepts</i>
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Changing the Owner of a Directory or File

You can change the owner of a directory or file using the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to Change the Owner

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Modify right to the file or directory for which you want to change the owner

Procedure



- 1. Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Using the browser, select the directory or file for which you want to change the owner.**

For information on moving around in the browser and selecting objects, choose “Help” from the menu bar.
- 3. From the “Object” menu, choose “Details.”**
- 4. From the “Identification” page, choose “Facts.”**
- 5. To change the owner of this file or directory, click on the browser button to the right of the “Owner” field.**
- 6. Choose the object that you want to make the new owner of this directory or file.**
- 7. When the correct user is displayed in the “Object Name” field, choose “OK.”**

The new owner appears in the “Owner” field of the object dialog box.

- 8. To save any changes, choose “OK.”**

Additional Information

For more information about	Refer to
Objects	Chapter 1, "Managing NetWare Directory Services Objects" "Object" in <i>Concepts</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>

Using FILER to Change the Owner

Prerequisites

Checklist



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Modify right to the file or directory for which you want to change the owner

Procedure

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears. Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select "Manage Files and Directories."

The "Directory Contents" menu appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and then press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and then pressing <Enter> until you see the item you want. Select it and press <F10>.
- ◆ If you can't find what you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

4. Select "View/Set File [or Directory] Information" and press <Enter>.

Information for the file or directory appears. The current owner of the file or directory appears in the "Owner" field.

5. Use the arrow keys to move to the "Owner" field and then press <Enter>.

6. Select the user that you want to be the owner of the file or directory and press <Enter>.



(For directories only) You can apply the change of ownership to either the entire subdirectory structure or to a selected directory.

7. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Objects	Chapter 1, "Managing NetWare Directory Services Objects" "Object" in <i>Concepts</i>
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Viewing Effective Rights and Other Information

Viewing a Trustee's Effective Rights to a Directory or File

You can see the effective rights a trustee has to a directory or file by using either the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to View a Trustee's Effective Rights

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Access Control right to the file or directory, and the Read property right to the trustee object's ACL property

Procedure



1. Choose the "NetWare Administrator" icon from the MS Windows Program Manager or the OS/2 desktop.

2. Using the browser, select a directory or file.

For information on moving around in the browser and selecting objects, choose "Help" from the menu bar.

3. From the "Object" menu, choose "Details."

4. From the "Identification" page, choose "Trustees of this Directory."

5. Select the trustee whose effective rights you want to view.

6. Choose "Effective Rights."

The selected trustee's effective rights are bolded in the "Effective Rights" box.

7. (Optional) To view a different object's rights to this directory or file, choose the browser button to the right of the "Trustee" text box and select an object from the browser.

8. To return to the browser, choose “Cancel.”

Additional Information

For more information about	Refer to
Effective rights	“Effective rights” in <i>Concepts</i>
Modifying trustee rights	“Modifying a Trustee’s Rights to a Directory or File” on page 140
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to View a Trustee’s Effective Rights

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Access Control right to the file or directory, and the Read property right to the trustee object’s ACL property

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select “Manage Files and Directories.”

The “Directory Contents” menu appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and then press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and then pressing <Enter> until you see the item you want. Select it and press <F10>.
- ◆ If you can't find the item on the list, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the "Available Options" menu and choosing "Select Current Directory."

4. Select "View/Set File [or Directory] Information" and press <Enter>.

Information for that file or directory appears.

5. Use the arrow keys to move to the "Trustees" field, and then press <Enter>.

A list of trustees for the file or directory appears, along with the object type of the trustee and the current rights the trustee has to the file or directory.

If the trustee list is empty, no effective rights exist for this file or directory. To assign trustee rights, see "Using FILER to Add a Trustee" on page 135.

Additional Information

For more information about	Refer to
Effective rights	"Effective rights" in <i>Concepts</i>
Modifying trustee rights	"Modifying a Trustee's Rights to a Directory or File" on page 140
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Viewing Other Information about a Directory or File

You can see extended information about a directory or file using either the NetWare Administrator graphical utility or the FILER or NDIR text utilities. You can see file information such as

- ◆ Owner and trustees
- ◆ Attributes, effective rights, and the Inherited Rights Filter (IRF)
- ◆ Name space
- ◆ File size
- ◆ Creation, access, archive, and modify dates

You can see directory information such as

- ◆ Owner and trustees
- ◆ Creation date and time
- ◆ Attributes, effective rights, and the IRF
- ◆ Disk space limitations

Procedures for using the NetWare Administrator, FILER, and NDIR are documented in this section.


Using NetWare Administrator to View Other Information

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Access Control right to the file or directory you want to see information about

Procedure



1. **Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
2. **Using the browser, select the file or directory you want to see information about.**

For help moving around in the browser and selecting objects, choose “Help” from the menu bar.

3. **Choose “Details” from the “Object” menu.**
4. **Select an “Information” page to view for this file or directory.**
5. **If you have made any changes on any pages, choose “OK” to save changes and return to the browser.**

If you haven’t made any changes, choose “Cancel” to return to the browser.

Additional Information

For more information about	Refer to
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to View Other Information

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Access Control right to the file or directory you want to see information about

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select “Manage Files and Directories.”

The “Directory Contents” list appears.

3. Find and select the file, directory, or subdirectory you want.

- ◆ If the item you want appears on the list, select it and then press <F10>.
- ◆ If the item is not on the list, browse a directory or subdirectory by selecting it and then pressing <Enter> until you see the item you want. Select it and press <F10>.
- ◆ If you can't find what you want, check the Volume object name in the upper left corner of the screen. If you are in the wrong volume, you can change it by returning to the “Available Options” menu and choosing “Select Current Directory.”

4. Select “View/Set File [or Directory] Information.”

Information for the file or directory appears.

5. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Using the FILER utility	"FILER" in <i>Utilities Reference</i>

Using NDIR to View Other Information

For online help with the NDIR utility's features, type one of the following commands:

- ◆ For information on the NDIR utility's sorting features, type

NDIR /? SORT <Enter>

- ◆ For information on search filters (restrictions), type

NDIR /? RES <Enter>

- ◆ For information on attribute filters, type

NDIR /? AT <Enter>

- ◆ For information on other options, type

NDIR /? OPT <Enter>

- ◆ For NDIR syntax information, type

NDIR /? SYN <Enter>

- ◆ To view all of the NDIR utility's help screens, type

NDIR /? ALL <Enter>

Prerequisites

Checklist



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Access Control right to the file or directory you want to see information about

Procedure

Procedure



Use the NDIR commands in Table 2-3 to find information about files, directories, volumes, etc. For more information about using NDIR, see “NDIR” in *Utilities Reference*.

Table 2-3
Using NDIR to Get File and Directory Information

To get information about	Use this NDIR command
Files only	NDIR <i>path</i> /FO
Directories only	NDIR <i>path</i> /DO
All subdirectories	NDIR <i>path</i> /S
Volumes	NDIR /VOL
Directory space	NDIR /SPA
File version	NDIR /VER
File details	NDIR <i>filename</i> /D

Salvaging and Purging Deleted Files and Directories

Files deleted from the NetWare server remain on the disk until the deleted files are purged. Deleted files can be salvaged any time before they are purged.

Purging frees the space used to store the deleted files on the server's hard disk. If a disk runs out of free space, NetWare automatically purges first the files that were deleted first.

Salvaging Deleted Files

Files deleted from the NetWare server can be recovered unless they have been purged.

You can salvage files by using either the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.

Using NetWare Administrator to Salvage Files

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Create right to the file that has been deleted. If the file is in the "Deleted Directories" area, you need the Supervisor right to the file

Procedure



1. **Choose the "NetWare Administrator" icon from the MS Windows Program Manager or the OS/2 desktop.**
2. **Select the directory containing the files or directories you want to salvage.**
3. **From the "Tools" menu, choose "Salvage."**

4. Select the source directory (current or deleted) from the “Source Options” drop box.

If you salvage files from deleted directories on the SYS: volume, the list of files takes a long time to appear.

5. Specify the deleted files you want listed by entering a file specification (such as wildcards) in the “Include” field.

For example, to see a list of deleted batch files in the current directory, you would type *.BAT in the “Include” field.

If you leave the field blank or type *.* you get a list of all deleted files in the selected directory.

6. Choose “List” to display the filenames you included.

7. (Optional) Specify how you want the files to be listed by selecting a sort option from the “Sort Options” drop box.

For example, you can sort by deletion date, deleter, etc. The default listing is by filename.

8. Select the files you want to salvage.

9. Choose “Salvage.”

If you salvage files from an existing directory, the files are restored to that directory. If you salvage files from a deleted directory, the files are restored under the root directory.

10. To return to the browser, choose “Close.”

Additional Information

For more information about	Refer to
Salvaging files	“Salvageable files” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to Salvage Files

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Create right to the file that has been deleted. If the File is in the "Deleted Directories" area, you need the Supervisor right to the file

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Your current context, Volume object, and path are shown in the upper left corner of the screen.

2. Select "Salvage Deleted Files."

The "Salvage" menu appears.

3. Salvage files from an existing directory or a deleted directory by completing one of the following steps:

- ◆ To salvage files from an existing directory, select "View/Recover Deleted Files" from the "Salvage" menu.
- ◆ To salvage files from a deleted directory, select "Salvage Deleted Directories" from the "Salvage" menu; then select the volume that contained the deleted directory and press <Enter>.

4. When prompted, type a filename or wildcard to indicate the files you want to salvage, and then press <Enter>.

For example, if you want to recover all files with the .EXE extension, type *.EXE. Or, if you want to recover all files that have RPT in their names, type *RPT*.

If you do not specify a filename or wildcard, all files that have been deleted from the directory will be listed.

A list of deleted the files appears. If you salvage files from deleted directories on the SYS: volume, the list of files takes a long time to appear.

5. (Optional) Specify how you want the filenames to be listed by using the sort option.

5a. Press <F3>.

5b. From the “Salvage Sort Options” menu, select the way you want the files to be listed.

You can sort the files by deletion date, deleter, etc. The default listing is by filename.

6. Select the files to salvage from the list, and press <Enter>. To salvage more than one file, press <F5> to mark each file, and then press <Enter>.

If you salvage files from an existing directory, the files are restored to that directory.

If you salvage files from a deleted directory, the files are restored under the root directory.


7. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Salvaging files	“Salvageable files” in <i>Concepts</i>
Using the FILER utility	“FILER” in <i>Utilities Reference</i>

Purging Files

Purging files frees disk space on the NetWare server's hard disk.



Purged files *cannot* be salvaged.

You can purge files using either the NetWare Administrator graphical utility or the FILER text utility. Both procedures are documented in this section.


Using NetWare Administrator to Purge Files

Prerequisites



- A 386 or 486 workstation running NetWare Administrator
- The Erase right to the deleted file or directory you want to purge. If the file is in the "Deleted Directories" area, you need the Supervisor right to the file

Procedure



- 1. Choose the "NetWare Administrator" icon from the MS Windows Program Manager or the OS/2 desktop.**
- 2. Select the directory containing the files or directories to be purged.**
- 3. From the "Tools" menu, choose "Salvage."**
- 4. Select the source directory (current or deleted) from the "Source" drop box.**

5. Enter a file specification (such as wildcards) in the “Include” field to specify files to be listed.

For example, to see a list of deleted batch files in the current directory, you would type .bat in the Include field.


If you leave the field blank or type *.* you get a list of all deleted files in the current directory.

6. Choose List to display the filenames you specified.

7. (Optional) Specify how you want the filenames to be listed by selecting a sort option from the “Sort Options” drop box.

For example, you can sort by deletion date, deleter, etc. The default listing is by filename.

8. Select the files to be purged.



If no files are selected, then *all* deleted files are purged.

9. Choose “Purge.”

10. To return to the browser, choose “OK.”

Additional Information

For more information about	Refer to
Purging files	“Purge (P) attribute” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Using FILER to Purge Files

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Erase right to the deleted file or directory you want to purge. If the file is in the “Deleted Directories” area, you need the Supervisor right to the file

Procedure



1. At the DOS prompt, type

FILER <Enter>

A list of available options appears.

Purge does *not* allow you to browse directories. Your current context, Volume object, and path are shown in the upper left corner of the screen. Make sure you are in the correct location in the Directory tree.

2. Select “Purge Deleted Files.”

You are prompted for a filename pattern.

3. Type a filename or a wildcard to specify the deleted files that you want to purge, and then press <Enter>.



If you do not specify a filename pattern, *all* deleted files are purged. As soon as you press <Enter> for Step 4, purging starts.

4. From the “Purge Options” menu, specify whether to purge files in the current directory only, or in the entire subdirectory structure of the current directory.

A list of the directories being checked and the files being purged appears.

5. To exit, press <Esc> until the menu you want appears.

Additional Information

For more information about	Refer to
Purging files	"Purge (P) attribute" in <i>Concepts</i>
Using the FILER utility	"FILER" in <i>Utilities Reference</i>



chapter

3

Creating Login Scripts

About Login Scripts

You can use login scripts to automatically set up your users' workstation environments whenever they log in to the network. Login scripts are similar to configurable 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

Login scripts work the same way for DOS, Microsoft* Windows (MS Windows), and OS/2* workstations.

Some login script commands that apply to DOS and MS Windows workstations may not apply to OS/2 workstations. These differences are indicated in the description of each command later in this chapter.

Four Types of Login Scripts

When a user logs in, the LOGIN utility executes the appropriate login scripts. Four types of login scripts are available, and they can be used separately or together to tailor a custom environment for your users. All four types of login scripts are optional.

- ◆ A container login script sets the general environments for all users in that container. The LOGIN utility executes container login scripts first. A user can use only one container login script.



A container login script replaces the system login script from NetWare 3™.

- ◆ A profile login script sets environments for several users at the same time. The LOGIN utility executes a profile login script after the container login script.

A user can be assigned only one profile login script, but can specify other profile login scripts on the command line. Several users can use the same profile login script.

- ◆ A user login script sets environments specific to a single user, such as printing options or a username for electronic mail. The LOGIN utility executes the user login script after any container and profile login scripts have executed.

A user can have only one user login script.

- ◆ The default login script is precoded into the LOGIN.EXE command and is not editable. It executes if a user doesn't have his or her own user login script, even if a container or profile login script exists.

The default login script is executed for all users (including user ADMIN) unless you create a user login script. The default login script contains only essential commands such as drive mappings to the NetWare® utilities.

(To see the commands in the default login script, see "Default Login Script" on page 243.)

If you don't want to create any user login scripts and you don't want the default login script to execute for any users, you can disable the default login script by including the NO_DEFAULT command in the container or profile login script.

To use the login script from an Organization, Organizational Unit, or Profile object, users must have the Browse right to the object and the Read right to the object's Login Script property.



For more information on Browse or Read rights for a file, object, or property, see "Browsing" and "Rights" in *Concepts*.

Deciding Which Login Scripts to Create

Maintaining many user login scripts can be time consuming. Therefore, you should try to include as much customizing information as possible in the container and profile login scripts, which are fewer in number and easier to maintain.

For example, if all users need access to the NetWare utilities in the same volume, put the search drive mapping to that volume in a single container login script rather than in every user login script.

Create profile login scripts if there are several users with identical login script needs.

Finally, in user login scripts, include only those individual items that can't be included in profile or container login scripts.

Since up to three login scripts can execute whenever a user logs in, conflicts can occur. If this happens, the last login script to execute (usually the user login script) overrides any conflicting commands in a previous login script.

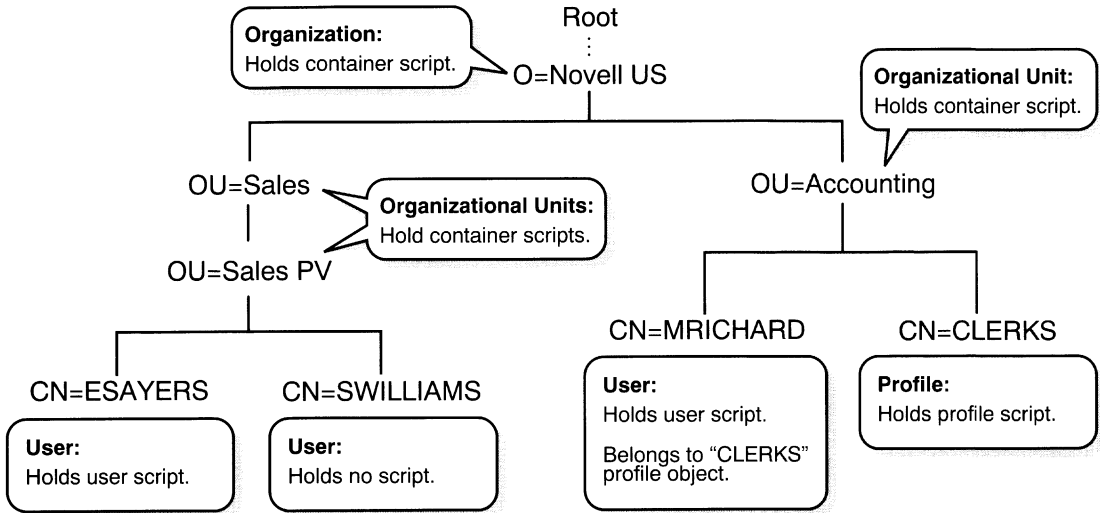
Login scripts are properties of objects. Table 3-1 shows which objects can contain which login scripts.

Table 3-1
Objects That Contain Login Scripts

Object	Type of Login Script
Organization	Container
Organizational Unit	Container
Profile	Profile
User	User

Figure 3-1 shows where the different types of login scripts can reside in a Directory tree.

Figure 3-1
Where Login Scripts
Are Located



In the previous figure, there are three users, ESAYERS, SWILLIAMS, and MRICHARD. The following table shows which login scripts execute when each of these users logs in.

When this user logs in	Login scripts execute in this order
ESAYERS	<ol style="list-style-type: none"> 1. Sales PV's container login script 2. ESAYERS' user login script
SWILLIAMS	<ol style="list-style-type: none"> 1. Sales PV's container login script 2. Default login script

When this user logs in	Login scripts execute in this order
MRICHARD	<ol style="list-style-type: none"> 1. Accounting's container login script 2. CLERKS' profile login script 3. MRICHARD's user login script

Container login scripts only affect users in the Organization or Organizational Unit that contains the login script.

For example, in Figure 3-1, although there are two levels of container objects above users ESAYERS and SWILLIAMS, only the container login script they're in (OU=SALES_PV) executes when they log in.

If the SALES_PV Organizational Unit had no container login script defined, no container login script would execute for ESAYERS and SWILLIAMS, even though a container login script exists at a higher level.

Because user SWILLIAMS has no user login script defined, the default login script executes after the container login script.

Since user MRICHARD belongs to the profile CLERKS, the CLERKS profile login script executes before MRICHARD's user login script. Users can be assigned to only one Profile object, but other profile login scripts can be specified at the command line. For example,

```
LOGIN username /p profile_object
```

You can, however, assign users to more than one Group object. Then use the MEMBER OF "group" identifier variable to specify that different parts of a login script execute, depending on the Group objects that the user belongs to.

For more information about using the MEMBER OF "group" identifier variable in login scripts, see "IF...THEN" on page 204 and "Identifier Variables" on page 236.

Using Login Scripts from Other NetWare Versions

When you use INSTALL.NLM to upgrade a previous version of NetWare to NetWare 4™, the login scripts in the former login directory are automatically added to the Directory database as properties of their respective objects.²¹⁶

In most cases, these earlier login scripts are compatible with NetWare 4 running NetWare Directory Services™. However, we recommend you review all your login scripts to see if they should be updated to take advantage of NetWare 4 functionality.

If you use a previous version of the NetWare MENU utility to execute a menu program from within a login script, see “Converting Old Menu Files” on page 277 for instructions on updating the menu program.

After you convert the menu, change the command in the login script from “#MENU *menu_name*” to “#NMENU *menu_name*.”

In previous versions of NetWare, login scripts for OS/2 and DOS were separate. In NetWare 4, only one login script is used for both OS/2 and DOS environments.


If a NetWare 3.11 user, for example, has both an OS/2 and a DOS login script, both of those login scripts are combined into a single login script when that user is upgraded to NetWare 4.

The commands from the OS/2 login script appear first in the upgraded login script, followed by the commands from the DOS script.

When the user logs in, the LOGIN utility determines whether the user is using OS/2 or DOS. Because of IF statements put in the script during migration/upgrade, LOGIN executes only the section of the upgraded script that contains the corresponding commands.

Creating, Modifying, Copying, and Printing Login Scripts

To create or modify login scripts and to copy one object's login script into another's, you can use either the NetWare Administrator graphical utility or the NETADMIN text utility. Both procedures are documented in this section.



If you are logged in to a server running NetWare 2 or NetWare 3 and that server is in a Directory tree, you must not create or edit a login script using the SYSCON utility. If you do, the changes to that login script will not appear in your NetWare Directory Services login script.

The reason is that your NetWare Directory Services login script is a property of your User object, while your bindery-based login script is a file in your MAIL directory.

The main difference in creating container, profile, and user login scripts is the object you select to contain the login scripts.

- ◆ Container login scripts are assigned to container objects (Organization or Organizational Unit objects).
- ◆ Profile login scripts are assigned to Profile objects. For a User object to use a profile login script, you must select that User object and assign it to the Profile.
- ◆ User login scripts are assigned to User objects.

All four types of login scripts use the same conventions, commands, and variables.

Login Script Hints

The following hints can help you plan effective login scripts. For a description of the commands you can use in a login script, see “Login Script Commands and Variables” on page 187. For login script examples, see “Examples of Login Scripts” on page 243.

Table 3-2
Login Script Conventions

Subject	Convention
Minimum login script	No minimum. All four types of login scripts are optional. Login scripts can have only one line or they can have many. There are no required commands for login scripts.
Case	Either uppercase or lowercase is accepted. Exception: identifier variables enclosed in quotation marks and preceded by a percent sign (%) must be uppercase. See “Identifier Variables” on page 236.
Characters per line	150 characters per line is maximum; 78 characters per line (common screen width) is recommended for readability.
Punctuation and symbols	Type all symbols (#, %, “, _) and punctuation exactly as shown in examples and syntax.
Commands per line	Use only one command per line. Start each command on a new line; press <Enter> to end each command and start a new command. Lines that wrap automatically are considered one command. WRITE command output displays better if WRITE is repeated at the beginning of each wrapped line.
Sequence of commands	Generally, enter commands in the order you want them to execute, with the following restrictions: <ul style="list-style-type: none">◆ ATTACH commands must precede related MAP commands to avoid prompting the user for a username/password during login.◆ If you use “#” to execute an external program, it must follow any necessary MAP commands.◆ If sequence is not important, group similar commands, such as MAP and WRITE commands, together to make the login script easier to read.
Blank lines	Blank lines don’t affect login script execution. Use them to visually separate groups of commands.

Table 3-2 *continued*

Login Script Conventions

Subject	Convention
Remarks (REMARK, REM, asterisks, and semicolons)	Lines beginning with REMARK, REM, an asterisk, or a semicolon are comments, which don't display when the login script executes. Use remarks to record the purpose of each command or group of commands. (For examples, see "REMARK" on page 224.)
Identifier variables	Type identifier variables exactly as shown. For the value of an identifier variable to be displayed on the workstation's screen as part of a WRITE command, you must enclose the identifier in quotation marks and precede it by a percent sign (%). See "Identifier Variables" on page 236.

Creating or Modifying a Login Script Using NetWare Administrator


Use the following instructions to create any of the three user-created types of login scripts (container, profile, or user).

Prerequisites



- A workstation running NetWare Administrator.
- The Write property right to the object that will contain the login script.
- The object to which you are going to assign the login script must already exist (Organization, Organizational Unit, Profile, or User object).

Procedure



1. **Choose the "NetWare Administrator" icon from the MS Windows Program Manager or the OS/2 desktop.**
2. **Using the browser, select the object whose login script you are creating or modifying.**

For information about moving around in the browser and selecting objects, choose Help from the menu bar.

3. **Choose "Details" from the "Object" menu.**

4. Choose the “Login Script” page.

5. Enter the login script commands and information into the login script text box.

For a description of all login script commands, see “Login Script Commands and Variables” on page 187. For login script examples, see “Examples of Login Scripts” on page 243.

6. Choose “OK” to save the login script and close the “Details” dialog box.

If the login script you just created was a container or user login script, you’re finished.

If the login script you just created was for a Profile object, continue with Step 7.

7. (Profile login scripts only) Using the browser, select the User object that needs to use the profile login script.

8. Choose “Details” from the “Object” menu.

9. Choose the “Login Script” page.

10. Enter the name of the Profile object in the “Default Profile” field located under the login script text box.

You can type in the complete name of the Profile object, or you can choose the browser button next to the “Default Profile” field to select the Profile object.

11. To save the Profile object name and close the “Details” dialog box, choose “OK.”

Now you must add the User object as a trustee of the Profile object.

12. Using the browser, select the Profile object.

13. Choose “Trustees of This Object” from the “Object” menu.

14. Choose “Add Trustee.”

15. Enter the name of the User object who is using this Profile object.

You can type in the complete name of the User object, or you can choose the browser button to select the Profile object.

16. Make sure the Browse object right and the Read property right are checked, and then choose “OK” to assign these rights to the User object.

The User object is now a trustee of the Profile object and has the rights necessary to run the profile login script

Additional Information

For more information about	Refer to
Creating a Profile object	“Managing Profile Objects” on page 59
Examples of login scripts	“Examples of Login Scripts” on page 243
Login script commands and variables	“Login Script Commands and Variables” on page 187
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Creating or Modifying a Login Script Using NETADMIN


Use the following instructions to create any of the three user-created types of login script (container, profile, or user).

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Write property right to the object that will contain the login script.
- The object to which you are going to assign the login script must already exist (Organization, Organizational Unit, Profile, or User object).

Procedure



1. At the DOS Prompt, type

NETADMIN <Enter>

2. Choose “Manage Objects” from the “NetAdmin Options” menu.

3. Select the object whose login script you want to create.

- ◆ If the object you want appears in the list, select it and press <F10>.
- ◆ If the object is not in the list, browse the directory by selecting container objects and pressing <Enter> until you see the object you want. Select it and press <F10>.

4. Select “View or Edit Properties of This Object.”

5. Select “Login Script.”

If you are editing an existing login script that already contains some commands, continue with Step 6.

If this login script is empty, a message appears asking if you want to copy a login script from another object.

5a. If you do not want to copy the login script from another object, answer No and continue with Step 6.

5b. To copy a login script from another object, answer Yes and select the name of the object whose script you want to copy. Then continue with Step 6.

The login script is copied into the login script text box.

6. Enter the login script commands and information in the login script text box.

For a description of all login script commands, see “Login Script Commands and Variables” on page 187. For login script examples, see “Examples of Login Scripts” on page 243.

7. To save the login script, press <F10>.

If the login script you just created was a container or a user login script, you’re finished.

If the login script you just created was for a Profile object, continue with Step 8.

8. (Profile login scripts only) Press <Esc> repeatedly until you return to the browser.

9. Select the User object that needs to use the profile login script.

You can either type the object’s complete name and press <F10>, or you can press <Insert> to browse through the Directory tree and choose the name.

10. Select “View or Edit Properties of This Object.”

11. Select “Memberships.”

12. Select the “Profile” field and press <Insert>.

13. Enter the name of the Profile object in the box that appears.

You can either type the object's complete name and press <F10>, or you can press <Insert> to browse through the Directory tree and select the name.

14. Press <F10> to save the changes.

Now you must add the User object as a trustee of the Profile object.

15. Return to the "Manage Objects" menu.

16. Through the browser, select the Profile object.

Press <Insert> to browse through the Directory tree and choose the name.

17. Select "View or Edit the Trustees of This Object."

18. Select "Trustees."

19. To add the User object as a trustee of this Profile object, press <Insert>.

20. Enter the name of the User object who needs to be a trustee of this Profile object.

You can either type the object's complete name and press <F10>, or you can press <Insert> to browse through the Directory tree and choose the name.

21. Select "[All Properties Rights]."

22. To add the user as a trustee and grant the default property right, press <Enter>.

The User object is added as a trustee of the Profile object and is given the Read right to all of the Profile's properties.

Now you must assign the Browse object right to the User object.

23. Enter the name of the User object.

You can either type the object's complete name and press <F10>, or you can press <Insert> to browse through the Directory tree and choose the name.

24. Choose “[Object Right].”

25. To grant the default object right, press <Enter>.

The User object is given the Browse object right. The User object now has all rights necessary to use the Profile object’s login script.

26. To exit NETADMIN, press <Esc> until you get to the confirmation prompt, and then select “Yes.”

Additional Information

For more information about	Refer to
Creating a Profile object	“Managing Profile Objects” on page 59
Examples of login scripts	“Examples of Login Scripts” on page 243
Login script commands and variables	“Login Script Commands and Variables” on page 187
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Copying a Login Script Using NetWare Administrator


Use the following instructions to copy all or part of a login script and paste it into another object’s login script.

Prerequisites



- A workstation running NetWare Administrator.
- The Write property right to the object that will contain the login script.
- The objects whose login scripts you will be working with must already exist (Organization, Organizational Unit, Profile, or User object).

Procedure



1. **Choose the “NetWare Administrator” icon from the MS Windows Program Manager or the OS/2 desktop.**
2. **Using the browser, select the object whose login script you want to copy.**

For information about moving around in the browser and selecting objects, choose Help from the menu bar.
3. **Choose “Details” from the “Object” menu.**
4. **Choose the “Login Script” page.**
5. **In the login script text box, highlight the text you want to copy.**
6. **Press <Ctrl>+<Insert> to copy the highlighted text.**

The highlighted text has been placed in clipboard memory and can be pasted into another login script.
7. **To save the login script and close the “Details” dialog box, choose “OK.”**
8. **Using the browser, select the object whose login script you want to paste the copied text into.**
9. **Choose “Details” from the “Object” menu.**
10. **Choose the “Login Script” page.**
11. **In the login script text box, place the cursor where you want the copied text to appear.**
12. **Press <Shift>+<Insert> to paste the copied text into the login script.**
13. **To save the login script and close the “Details” dialog box, choose “OK.”**

Copying a Login Script Using NETADMIN

Use the following instructions to copy all or part of a login script and paste it into another object's login script.

Prerequisites



- A workstation running DOS 3.30 or later and the NETADMIN utility
- The Write property right to the object that will contain the login script.
- The objects whose login scripts you will be working with must already exist (Organization, Organizational Unit, Profile, or User object).

Procedure



1. **At the DOS prompt, type**
NETADMIN <Enter>
2. **Choose “Manage Objects” from the “NetAdmin Options” menu.**
3. **Select the object whose login script you want to copy.**
 - ◆ If the object you want appears in the list, select it and press <F10>.
 - ◆ If the object is not in the list, browse the Directory by selecting container objects and pressing <Enter> until you see the object you want. Select it and press <F10>.
4. **Select “View or Edit Properties of This Object.”**
5. **Select “Login Script.”**
6. **In the login script text box, place the cursor at the beginning of the text you want to copy and press <F5> to mark the beginning of the text.**

7. Use the arrow keys to move to the end of the text you want to copy.

As you move the cursor, text in the login script is highlighted. This highlighted text will be copied by first deleting it here, then reinserting it here, and then inserting it into the new login script.

8. To delete the text from the login script, press <Delete>.

Although you have deleted the text, the text has been placed in a clipboard memory and can be retrieved.

9. To insert the text from the login script, press <Insert>.

The deleted text has now been restored to your login script. A copy of the text still resides in the clipboard memory, so you can paste it into another object's memory.

10. To exit the login script, press <Esc>, and select "No" when asked if you want to save the changes you made.

11. Return to the browser screen.

12. Select the object whose login script you want to paste the copied text into.

You can either type the object's complete name and press <F10>, or you can press <Insert> to browse through the Directory tree and choose the name.

13. Select "View or Edit Properties of This Object."

14. Select "Login Script."

15. In the login script text box, place the cursor where you want the copied text to appear.

16. To paste the copied text into the login script, press <Insert>.

17. To save the changes, press <F10>.

Printing Login Scripts


Use the following instructions to print a login script.

Prerequisites



- A workstation running DOS 3.30 or later
- The Read and File Scan property rights to the object to be printed.

Procedure



To print a login script from the command line, use the NLIST command and redirect the output to a file or a printer. You must be in an object's parent container to see and print the login script of that object.

To print a user's login script, type:

```
NLIST user <username> show "login script" >LPT1
```

To print a container's login script, type:

```
NLIST "organizational unit" = "ou name" show  
"login script" >LPT1
```

Any parameter of the NLIST command that includes a space in its name must be enclosed in quotes.

Login Script Commands and Variables

This section describes the commands you can use in a login script. The commands are presented in alphabetical order.

Syntax conventions for login script commands, as shown in Table 3-3, are the same as those for workstation text utilities, with one exception: Some login script commands must be preceded by the # symbol.

Following is an example of the syntax for the MAP login script command:

```
MAP [option] drive:=drive:|path <Enter>
```

The command syntax is described in the table.

Table 3-3
Command Syntax Conventions

Convention	Explanation
MAP	Words in uppercase letters are keywords that must be included in the command and spelled exactly as shown. However, it doesn't matter if you type them in uppercase or lowercase letters.
[]	Square brackets indicate that the enclosed item is optional.
	A vertical bar means you can use either the item to the left of the bar or the item to the right, but not both. In the previous MAP example, you can enter either the drive letter or a complete path.
<i>drive</i>	Words in italics are variables. Replace variables with information specific to your task.
[<i>option</i>]	Options or parameters for each command are listed with the command. Options and parameters can often be abbreviated.
<Enter>	Angle brackets indicate a key you should press.
[[]]	Nested square brackets indicate that all enclosed items are optional. However, if you use the items within the innermost brackets, you must also use the items within the outer brackets.

(Execute External Program)

Use this command (the # symbol) to execute a program that is external to the login script.

Command Format

```
# [path] filename [parameter]
```

Replace *path* with a drive letter, or, if you have specified NOSWAP on the command line or in the login script, you may replace *path* with a full directory path beginning with the NetWare volume name.

Replace *filename* with an executable file (files that end in .EXE, .COM, or .BAT, for example). Do not include the extension.

Replace *parameter* with any parameters that must accompany the executable file.

Using

If you want the LOGIN utility to execute a program that is external to the login script, enter the # command (symbol) followed by the name of the file you want to execute.

This command fails when

- ◆ The given directory is invalid
- ◆ Proper security rights are lacking
- ◆ The executable file cannot be found
- ◆ Insufficient workstation memory is available to load the file



LOGIN swaps to extended or expanded memory or to disk unless NOSWAP is specified on the command line or in the login script.

NOSWAP prevents LOGIN from being swapped out of conventional memory. Then, if the station does not have enough memory to handle both LOGIN and the # command, the # command fails but the rest of the login script executes normally.

For more information, see "SWAP" on page 231.

Examples

When you want to define a default print queue and printer, you can make the login script execute the NetWare CAPTURE utility. This allows you to send print jobs to a network print queue (named QUEUE_FOR_LASERJET in this example).

If you have a search drive mapped to SYS:PUBLIC where the NetWare utilities are stored, you could enter the following command in the login script:

```
#CAPTURE Q=QUEUE_FOR_LASERJET NB TI=10 NFF
```

You do not need to enter a path in this case because CAPTURE is located in a search drive.

If you do not have a search drive mapped to a directory, include the path to that directory in the command. For example, to run a batch file named BATCH.BAT in the ACCOUNTS directory, use the following command:

```
#Z : \ACCOUNTS\BATCH
```

Additional Information

For more information about	Refer to
Using the NOSWAP command	"NOSWAP" on page 221
Using the SWAP command	"SWAP" on page 231

ATTACH

Use ATTACH to connect to a bindery-based NetWare server (NetWare 2 or NetWare 3), or to a NetWare 4 server using bindery services, while the login script is running.

Command Format

```
ATTACH [server[/username[;password]]]
```

Using ATTACH

Replace *server* with the name of the NetWare server to which you want to attach.

Replace *username* with the login name. If you do not include the username, the user is prompted for a login name when the ATTACH command is executed from the login script.

You can replace *password* with the correct password for that user and server. If the username and password are the same as the primary login username and password, you can omit the password and not be prompted for it.

Use caution when including passwords in a login script, however. It is more secure to eliminate the password. Then, at the point in the login script where the ATTACH command is executed, the user is prompted for the password.

Example

To attach user MRICHARD (whose password is "GOLFING") to a server named REPORTS (which is a bindery-based server running NetWare 3), add the following line to her login script:

```
ATTACH REPORTS/MRICHARD;GOLFING
```

BREAK

Use **BREAK ON** to allow the user to terminate execution of the login script. The default is **BREAK OFF**.

Command Format

BREAK ON|OFF

Using BREAK

If **BREAK ON** is included in a login script, you can press **<Ctrl>+<C>** or **<Ctrl>+<Break>** to abort the normal execution of your login script.

Including **BREAK ON** in a login script does not affect the DOS **<Ctrl>+<Break>** check. For more details, see “DOS BREAK” on page 196.

When the **BREAK** option is **ON**, type-ahead keyboard input is not saved in the buffer.

CLS

Use **CLS** to clear the display from the workstation’s screen during the login process.

Command Format

CLS

Using CLS

When a user logs in, a login script may display messages on the user’s workstation screen.

If the **CLS** command is added to the login script, any messages generated by commands earlier in the login script are cleared from the screen.

COMSPEC

To execute DOS commands from the network, use COMSPEC in the login script. Specify the directory where DOS and the DOS command processor (COMMAND.COM) are loaded.

COMSPEC is originally set when DOS is booted. It must be reset after you log in to change the location that COMMAND.COM loads from while you're in the network.



Users running virtual DOS sessions under OS/2 should not use this command in the login script. For more information, see "Using COMSPEC with OS/2" on page 194.

Command Format

COMSPEC=[*path*]COMMAND.COM

Replace *path* with either a drive letter or a full directory path beginning with the NetWare volume name.

Using COMSPEC

- ◆ If users are running DOS from a network directory, first map a search drive in the login script to that directory and then add the COMSPEC command to the login script.

(You may want to map a fake root to the DOS directory. For information about mapping a fake root, see "MAP" on page 212.)
- ◆ If all users use the same version of DOS from the network, you can add the COMSPEC command to the container login script.
- ◆ If more than one version of DOS is available on your network, a network directory should exist for each DOS version. In this case, you can put COMSPEC commands in either profile or user login scripts, to make sure each workstation accesses the version of DOS it needs.

- ◆ If users are running DOS from their local drives, do not add COMSPEC to login scripts.
- ◆ To use an environment variable as the value in a COMSPEC command, precede it with the percent sign (%), as follows:

COMSPEC=%environment variable

Using COMSPEC with OS/2

For information about accessing the correct version of DOS from OS/2, see *NetWare Client for OS/2 User Guide*.

CONTEXT

Use CONTEXT to set a user's current context in the Directory tree.

Command Format

CONTEXT *context*

Using CONTEXT

To change the current Directory tree context, replace *context* with the context you want the user to see after login.

Similarly to the workstation CX utility, you can enter a complete name to move down through the context, or you can use periods to move up toward the root of the tree.



CONTEXT does not support all options that the CX workstation utility does. It only sets the context.

Example

To change the context to the Organizational Unit SALES, under the Organization NOVELL_US, add the following line to the login script:

CONTEXT .SALES.NOVELL_US

You can type a single period instead of a container name to indicate that you want to move up one level.

For example, if you are in the context SALES.NOVELL_US and you want to move up one level to the context NOVELL_US, add the following line to the login script.

CONTEXT .

To move up two levels, enter two periods, and so on.

Additional Information

For more information about	Refer to
Context	"Context" in <i>Concepts</i>
Using the CX utility	"CX" in <i>Utilities Reference</i>

DISPLAY

Use DISPLAY to show the contents of a text file on a workstation's screen when the user logs in.

This command works best with an ASCII file. If you use DISPLAY with a word-processing file, printer and word-processing codes are displayed with the text.

Command Format

DISPLAY [*path*] *filename*

Replace *path* with either a drive letter or a full directory path beginning with the NetWare volume name.

Replace *filename* with the complete name (including the extension) of the file you want to display.

Using DISPLAY

When you use DISPLAY to display the contents of a file on the screen, the exact characters in the file, including any printer and word-processing codes, appear on the workstation screen. (To display only the text and suppress codes, use FDISPLAY. See "Using FDISPLAY" on page 201)

If the given directory does not exist or if the file is not found, no error message appears on the screen when the user logs in.

Example

Suppose you put messages in a file called SYSNEWS.TXT, in the directory SYS:PUBLIC\MESSAGES, and you want your users to see this file on their screens when they log in on Mondays. Add the following lines to the container login script:

```
IF DAY_OF_WEEK="Monday" THEN  
    DISPLAY SYS:PUBLIC\MESSAGES\SYSNEWS.TXT  
END
```

DOS BREAK

Use DOS BREAK to set the <Ctrl>+<Break> checking level for DOS.

If DOS BREAK is set to ON, you can terminate a program (other than the login script) by pressing <Ctrl>+<Break>. This command does not apply to OS/2 workstations.



This command is different from the BREAK command that terminates a login script. For more details, see "BREAK" on page 192.

Command Format

```
DOS BREAK [ON|OFF]
```

Using DOS BREAK

Enter the following command in the login script:

```
DOS BREAK ON
```

The default is DOS BREAK OFF.

Additional Information

For more information about	Refer to
Using the DOS BREAK command	Your DOS manual
Using the BREAK login script command	"BREAK" on page 192

DOS SET

See "SET" on page 226.

DOS VERIFY

Use DOS VERIFY to verify that data written to a local drive is not written to a bad sector and can be read without an error.

Command Format

DOS VERIFY [ON|OFF]

Using DOS VERIFY

The DOS COPY and NCOPY commands do not automatically verify that data copied to a local drive can be read after the copy.

To assure verification of each copy operation after login, add the VERIFY ON and DOS VERIFY ON commands (for network and DOS copies respectively) to the login script.


Another option, since VERIFY ON can affect performance by slowing down write operations, is to use the /V option at the command line with each COPY or NCOPY operation.

The default in the login script is DOS VERIFY OFF.

These commands may not work with some software that is copy protected.

DRIVE

Use **DRIVE** to change the default drive while the login script is executing.



On OS/2 workstations, **DRIVE** is effective only during execution of the login script. When the login script completes, you're returned to the drive in effect when the **LOGIN** command was issued..

Command Format

DRIVE [*drive*:|**n*:]

Replace *drive* with a local or network drive letter, or replace *n* with a drive number. Use of either is dependant on their already being assigned within the login script.

Using DRIVE

Unless this command is in your login script, the default drive is set to the first network drive, which is often assigned to your home directory when you log in.

If you don't want the default drive to be the first network drive, map a drive in the login script to the directory you want to be the default; then, use the **DRIVE** command to change the default drive.

Instead of specifying a drive letter such as F: or G:, you can use an asterisk followed by a number *n* to represent the *n*th network drive (for example, *3). This allows drive letters to reorder themselves automatically if previous drive mappings are deleted or added.

Example


Suppose you expect to work on only one project for several days and the files for that project are located on drive S:. You can use the **DRIVE** command to set your default drive to S: so you won't have to change your default drive manually every time you log in.

First, make sure you've mapped drive S: to the correct directory in your login script. Then enter the following command in the login script:

DRIVE S:

EXIT

Use EXIT to terminate execution of the login script and execute an external program.



This command does not apply to OS/2 workstations.

Command Format

```
EXIT ["filename [parameters]"]
```

Using EXIT

The length of information between quotes can't exceed your keyboard buffer length minus 1 (commonly 15 - 1 = 14 characters).

You can use the EXIT command in a login script to stop the login script and execute a program, such as a word processing or menu program.

Because EXIT stops the login script, make sure you put this command at the end of the login script.

You can also use EXIT in an IF...THEN statement, so that the login script stops and exits to an external program only if a certain condition exists. If the condition doesn't exist, the login script skips the EXIT command and continues executing.

If the program you are executing with the EXIT command requires any DOS paths or NetWare search drives to be set, make sure they are specified in the login script ahead of the EXIT command.

If you add EXIT to a container login script, it prevents other profile or user login scripts from running. If you put EXIT in a profile login script, it prevents the user login script from running.

The EXIT command works only on IBM-compatible workstations running DOS. Therefore, if your DOS workstation has a machine name different from IBM_PC specified in its NET.CFG file, you must add the PCCOMPATIBLE login script command to the login script.

For more information about the PCCOMPATIBLE command, see "PCCOMPATIBLE" on page 222.

Examples

- ◆ Suppose the workstation's long machine name is IBM_PC. To execute a menu program called TRAINING when the login script is finished, add the following line at the end of the login script:

```
EXIT "NMENU TRAINING"
```

- ◆ If you are using a Hewlett Packard computer and you have changed the long machine name to HE_PAC in the NET.CFG file, add the following lines at the end of the login script:

```
PCCOMPATIBLE
```

```
EXIT "NMENU TRAINING"
```

- ◆ Suppose you want the login script to exit to a word processing program when the user logs in on Mondays, but not on other days. You could add the following IF...THEN statement to the login script:

```
IF DAY_OF_WEEK="MONDAY" THEN EXIT "WP"
```

Additional Information

For more information about	Refer to
Changing the machine name in NET.CFG	<i>NetWare Client for DOS and MS Windows User Guide</i>
Creating a menu	"Creating a Menu File" on page 272 in <i>Creating Menus</i>

FDISPLAY

Use **FDISPLAY** to show the text of a word processing file on a workstation's screen when the user logs in.

To display both the text and the printer and word processing codes of a file, or to display an ASCII file, see "DISPLAY" on page 195.

Command Format

```
FDISPLAY [path] filename
```

Replace *path* with either a drive letter or a full directory path beginning with the NetWare volume name.

Replace *filename* with the complete name (including the extension) of the file you want to display.

Using FDISPLAY

When you use **FDISPLAY** to display the contents of a word processing file on the screen, the text in the file is filtered and formatted so that only the text itself is displayed. **FDISPLAY** will not display tabs.

If the given directory does not exist or if the file is not found, no error message appears on the screen when the user logs in.

Example

Suppose you put messages in a file called **SYSNEWS.TXT**, in the directory **SYS:PUBLIC\MESSAGES**, and you want your users to see this file on their screens when they log in on Mondays.

Add the following lines to the container login script:

```
IF DAY_OF_WEEK="Monday" THEN  
  FDISPLAY SYS:PUBLIC\MESSAGES\SYSNEWS.TXT  
END
```

FIRE PHASERS

FIRE PHASERS signals the workstation to emit a phaser sound.

Command Format

FIRE *n*

Replace *n* with the number of times you want this sound to occur.

Using FIRE PHASERS

Use this command by itself to generate the phaser sound whenever a user logs in. Use FIRE PHASERS with the IF...THEN command to make the sound execute a different number of times depending on the circumstances of the login.

Example

The following line executes the phaser sound four times upon login:

FIRE 4

To use an environment variable as the number of times to fire, use % before the variable, as follows:

FIRE %*environment variable*

Either of the following lines fires the phaser five times on Thursdays:

IF DAY_OF_WEEK="Thursday" THEN FIRE 5

or

FIRE %NDAY_OF_WEEK

The identifier variable %NDAY_OF_WEEK indicates a number that corresponds to the day of the week. Since Thursday is the fifth day of the week, phasers fire five times on Thursdays.

For more information about using identifier variables, see "Identifier Variables" on page 236.

GOTO

Use GOTO to execute a portion of the login script out of the regular sequence.

Command Format

GOTO *label*

Use *label* to indicate where the login script should continue executing.

Using GOTO

Set BREAK ON in your login script before experimenting with GOTO loops so that you can break out of a login script if necessary.

For more information about the BREAK login script command, see "BREAK" on page 192.



Do not use GOTO to enter or exit a nested IF...THEN statement. This usage confuses the program.

Example

To execute a loop of commands, you could include the following lines in your login script. In this case, the commands to be executed are labeled AGAIN (as indicated in the second line).

```
SET X="1"  
AGAIN:  
SET X=<X> + "1"  
;see compound strings for this  
WRITE <X>  
IF <X> < "9" THEN GOTO AGAIN
```

The GOTO command looks at the value of <X> (a DOS environment variable). If the value of <X> is less than 9, then <X> increments by 1 and GOTO loops back to the AGAIN label. When <X> gains the value of 9, the IF...THEN test becomes false, the GOTO is ignored, and the script continues normally. See the IF...THEN command below.

IF...THEN

Use IF...THEN when you want the login script to perform an action only under certain conditions.

Command Format

```
IF conditional [AND|OR [conditional]] THEN  
    commands  
[ELSE  
    command]  
[END]
```

Replace *conditional* with identifier variables. For more information about identifier variables, see "Identifier Variables" on page 236.

Replace *commands* with any login script commands that you want to be executed if the specified condition is true.

Using IF...THEN

- ◆ Use IF...THEN statements to execute commands only under certain conditions.

An example of a conditional statement is

```
IF MEMBER OF "CLERKS"
```

In this statement, some action is performed if the user who logged in belongs to the Group object named CLERKS.

The following is a different type of conditional statement:

```
IF DAY_OF_WEEK="MONDAY"
```

In this statement, the equal sign (=) indicates the relationship between the variable (DAY_OF_WEEK) and its value (Monday). Note that the value (Monday) is inside quotation marks.

- ◆ When using IF...THEN statements, be aware of the following syntax rules:
 - ◆ Use AND or OR to include two or more conditionals in an IF...THEN statement.
 - ◆ Values of conditional statements must be enclosed in quotation marks.
 - ◆ The ELSE statement is optional.
 - ◆ IF, ELSE, and END must be on separate lines. THEN does not need to be on a separate line.
 - ◆ If you include a WRITE command as part of the IF...THEN command, the WRITE command must be on a separate line.
 - ◆ IF...THEN statements can be nested (up to 10 levels). However, GOTO should not be used in a nested IF...THEN statement to enter or exit from the body of an IF...THEN statement.
 - ◆ If your IF...THEN statement consists of only one line, even if that line wraps, you do not need to include END. If your IF...THEN statement must be on more than one line (for example, you used ELSE or WRITE, which must be on separate lines), you must include END.

Six relationships are possible between the elements of an IF...THEN statement. Represent these relationships with the following symbols:

Symbol	Definition
=	Equals
<>	Does not equal
>	Is greater than
>=	Is greater than or equal to
<	Is less than
<=	Is less than or equal to

Examples

- ◆ If you place the following command in a login script, the message "Status report is due today" appears when the user logs in on Monday and "Have a nice day!" on other days.

```
IF DAY_OF_WEEK="MONDAY" THEN
    WRITE "Status report is due today"
ELSE
    WRITE "Have a nice day!"
END
```

- ◆ The following line means "If the hour (on a 24-hour scale) is greater than or equal to 12, then write 'afternoon'."

```
IF HOUR24>="12" THEN
    WRITE "afternoon"
END
```

- ◆ The following command executes the CAPTURE utility on the fourth day of the week (Wednesday):

```
IF NDAY_OF_WEEK="4" THEN
    #CAPTURE Q=FAST_Q NB TI=10 NFF
END
```

- ◆ The following example shows nested IF...THEN statements. Notice that there are two IF statements, so each one must have its own END statement.

```
IF DAY_OF_WEEK="MONDAY" THEN
    MAP *6:=VOL1:APPL\WP
    IF MEMBER OF CLERKS THEN
        WRITE "Your report is due immediately!"
    END
END
```

- ◆ Conditionals can be joined with commas, the word AND, or the word OR to form compound conditionals.

The first line of the following IF...THEN statement is a compound conditional that means "If it is the evening of the first day of the month":

```
IF GREETING_TIME="EVENING" AND DAY="01" THEN  
    WRITE "The system will be backed up tonight."  
END
```

The following line is a compound conditional that means "If it is 11:59:59 p.m.":

```
IF HOUR24="23" AND MINUTE="59" AND SECOND="59"
```

- ◆ An IF...THEN statement can include several commands that must be executed if the conditional is true.

The following example shows two commands that are executed on Tuesdays: a WRITE command that displays a message about a staff meeting, and an INCLUDE command that tells the login script to process any commands or messages contained in the file SYS:PUBLIC\UPDATE.

```
IF DAY_OF_WEEK="TUESDAY" THEN  
    WRITE "Staff meeting today at 10 a.m."  
    INCLUDE SYS:PUBLIC\UPDATE  
END
```

Additional Information

For more information about	Refer to
Using identifier variables	"Identifier Variables" on page 236
Using the WRITE login script command to display messages	"WRITE" on page 233

INCLUDE

Use INCLUDE to execute independent files or another object's login script as a part of the login script currently being processed.

These sub-scripts can be text files that contain valid login script commands (any of the commands explained in this section) or login scripts that belongs to a different object you have rights to.

Command Format

INCLUDE [*path*] *filename*

or

INCLUDE *object_name*

To use a text file as a sub-script, replace *path* with either a drive letter or a full directory path beginning with the NetWare volume name.

Replace *filename* with the complete name (including the extension) of the text file.

To execute another object's login script as part of a login script, replace *object_name* with the name of the object whose login script you want to use.

Using INCLUDE

Text files that contain login script commands and other object's login scripts can be used as sub-scripts. Use these sub-scripts to supplement the main login script.

You can create and edit text file sub-scripts using any text editor. Sub-scripts do not have to have any particular filenames or extensions.

The INCLUDE command executes the login script commands contained in the sub-script. It does not display the text of the sub-scripts.

INCLUDE nesting is limited only by available memory. This means that one sub-script file can include another sub-script file, which can include yet another sub-script file, and so on.

If the sub-script is a text file, users must have at least File Scan and Read rights to the directory containing the sub-script.

If you are using another object's login script as a sub-script, users must have the Browse right to the object whose script you are including, and the Read right to the object's Login Script property.

Examples

- ◆ To execute a text file called SCRIPT.NEW (located in the VOL1: volume) as a sub-script, add the following line to your main login script:

```
INCLUDE VOL1:ADMIN\USERS\SCRIPT.NEW
```

- ◆ Suppose you are creating a container login script for all users under the Organizational Unit object SALES_LA. You recently created a container login script for users under the Organizational Unit object SALES_PV.

Now you've decided that the login scripts for the two different groups of users are very similar. In fact, you decide that the SALES_LA users could use the same login script as the SALES_PV users, but with a few more drive mappings.

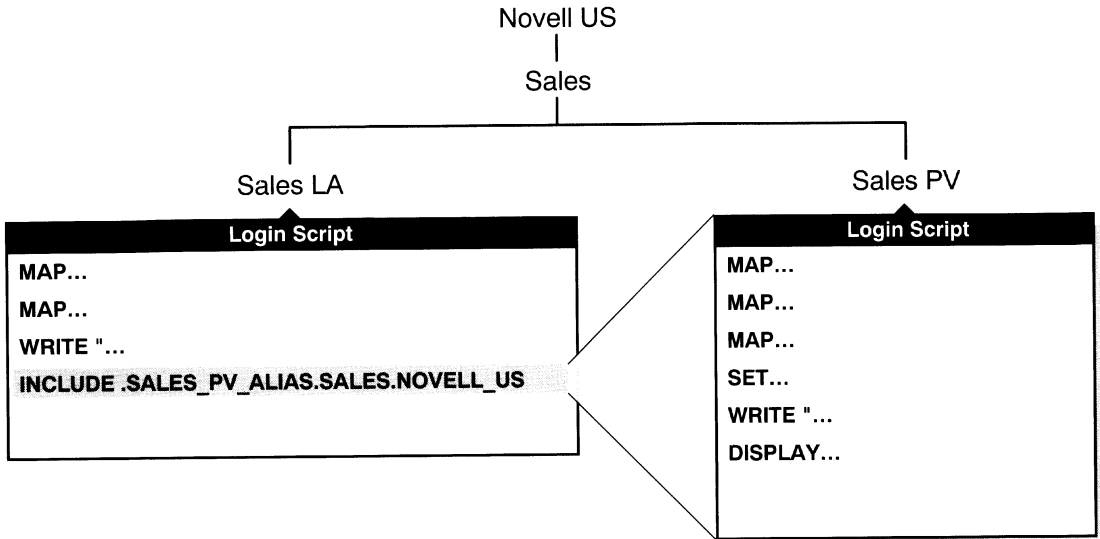
In the SALES_LA login script, you could add the additional drive mappings, and then use the INCLUDE command to execute the entire SALES_PV login script as a part of the SALES_LA login script, as follows:

- ◆ Create an alias for the SALES_PV Organizational Unit in the SALES_LA Organizational Unit.
- ◆ Add this line to the SALES_LA Organizational Unit's login script.

```
INCLUDE .SALES_PV_ALIAS.SALES.NOVELL_US
```

Figure 3-2 illustrates how the INCLUDE command executes the SALES_PV login script as part of the SALES_LA login script.

Figure 3-2
Using INCLUDE in a Login Script



LASTLOGINTIME

Use LASTLOGINTIME to display the last time the user logged in.

Command Format

LASTLOGINTIME

Using LASTLOGINTIME

If you include this command in your login script, the time of the last login is displayed on the user's workstation screen.

MACHINE

Use MACHINE to set the DOS machine name (such as IBM or EDIT_ROOM) of the workstation. The MACHINE command is necessary for some programs (such as NETBIOS) written to run under PC DOS.

The MACHINE command does not apply to OS/2 workstations.



Do *not* confuse the MACHINE command with the identifier of the same name. The identifier is used with a preceding percent (%) sign in MAP and WRITE statements. The identifier reads its value from the NET.CFG file.

It is unlikely that you will need to use this command. However, it is provided in case you come upon a program that requires it.

Command Format

MACHINE=*name*

Using MACHINE

The machine name can be up to 15 characters. (Longer machine names are truncated to 15 characters.)

For example, to specify that the machine name is IBM_PS2, add the following line to the login script:

MACHINE=IBM_PS2

Additional Information

For more information about	Refer to
Setting machine names in the NET.CFG file	<i>NetWare Client for DOS and MS Windows User Guide</i>
Using identifier variables	"Identifier Variables" on page 236

MAP

Use MAP to map drives and search drives to network directories.

Command Format

```
MAP [option] [drive:=path]
```

Replace *drive* with any valid network drive letter, local drive letter, or search drive number.

Replace *path* with either a drive letter, a full directory path, or a Directory Map object.

More than one command can be on the map line if the commands are separated by a semicolon (;), as shown in the following example:

```
MAP *1:=SYS:PUBLIC;*2:=SYS:PUBLIC\DOS
```

When mapping a drive to a directory on a NetWare Directory Services (NDS) server, begin the path with either the Volume object name or server/volume.

When mapping to a directory on a bindery-based server or to an NDS server that isn't your current server, begin the path with the server's name.

Replace *option* with one of the following:

- ◆ **DISPLAY ON/OFF:** Determines whether drive mappings are displayed on the screen when the user logs in. The default setting is ON. This option is valid only in login scripts.
- ◆ **ERRORS ON/OFF:** Determines whether MAP error messages are displayed when the user logs in. MAP ERROR OFF must be placed before MAP commands in the login script. The default setting is ON. This option is valid only in login scripts.
- ◆ **INS:** Inserts a drive mapping between existing search mappings. This option is valid in login scripts and at the DOS command line.

- ◆ **DEL:** Deletes a drive mapping, making that drive letter available for other mapping assignments. This option is valid in login scripts and at the command line.
- ◆ **ROOT:** Maps a fake root. OS/2 is always mapped to the root. Some applications require their executable files to be located in a root directory.

Since you may not want users to have rights at the root directory, you can map a fake root to a subdirectory instead. This option is valid in login scripts and at the command line.

- ◆ **C (CHANGE):** Changes a search drive mapping to a regular mapping, and a regular mapping to a search drive mapping. This option is valid in login scripts and at the command line.
- ◆ **NP (No prompt):** When a MAP command conflicts with an existing drive mapping, MAP NP eliminates the prompt that asks the user if the new drive mapping should overwrite the old mapping.

This option is valid only at the command line.

- ◆ **P (Physical):** Maps a drive to the physical volume of a server, rather than to the Volume object's name.

It is possible to have a Volume object name that conflicts with a physical volume name. (For example, object ACCT is an Accounting volume, but there is also an ACCT which is a physical volume.)

Therefore, if you prefer to map a drive to the physical volume name, use MAP P. This option is valid in login scripts and at the command line.

- ◆ **N (Next):** When used without specifying a drive number or letter, maps the next available drive. This option is valid in login scripts and at the command line.

Using MAP

- ◆ If you use MAP to automate drive map assignments during execution of the login script, users don't have to manually map drives every time they log in.
- ◆ Specify drive mappings in a login script by entering the same commands that you would enter if you were using MAP at the command line.
- ◆ To avoid having the result of each mapping displayed as it is executed, you can put the MAP DISPLAY OFF command at the beginning of your login script. When all drive map assignments have been completed, add the line MAP DISPLAY ON and MAP to your login script. This sequence provides a cleaner display for the users as they log in.
- ◆ Instead of specifying drive letters such as F: or G:, you could use an asterisk followed by a number *n* to represent the *n*th network drive. For example, if your first network drive is F: then using MAP *3:= would assign H: {F G H—1 2 3}, or if your first network drive is D: then using MAP *4:= would assign G: {D E F G—1 2 3 4}.

This allows drive letters to reorder themselves automatically when local drives are removed or added or when the first network drive is changed.

This also allows users to log in from workstations with a different number of local drives than their regular workstation.

- ◆ You can map a local drive (usually A: through C:) to a network directory, but you cannot access the local drive until you remove the network drive mapping.
- ◆ You must not map a redirected drive, such as a CD-ROM drive, to a network drive.

- ◆ In many cases, you might find it useful to map network drives in the following order:

- ◆ Map the first network drive to the user's home directory.
- ◆ If your users are running OS/2, map network drive P: to the SYS:PUBLIC directory.

This drive mapping is used by OS/2 workstations, as explained in *NetWare Client for OS/2 User Guide*.

- ◆ If users are running MS Windows from the network, map a drive to each user's directory that contains user-specific files.
- ◆ Map remaining drives to any directories that users work in frequently, such as project directories.

Mapping Search Drives



OS/2 workstations do not use search drives. Instead, OS/2 users use the OS/2 commands PATH, DPATH, and LIBPATH in their CONFIG.SYS files.

For DOS and MS Windows workstations, you can map search drives to directories that contain applications, executable files, and so forth. Then users can execute those applications regardless of the directory they are currently working in.

A maximum of 16 NetWare search drives are allowed.



Any object names in the login script should either be in the user's context or should have an alias point to the real object in another context. Any object referenced by a name outside the user's context will break when that object is moved or renamed, or the context is renamed. See "Alias object" in *Concepts*.

When you map a search drive, use a search drive number (an S followed by a number). This search drive number assigns the next available drive letter to the mapping, starting with Z and working backwards through the English alphabet.

The letter assigned to the search drive is put into the DOS path statement. If you already have search drives in the path statement, the command MAP S1: will overwrite the first one in the path. To prevent search drive assignments from overriding existing DOS PATH letters, use the INSERT option when assigning search drives. For example, type

```
MAP INS S16:=path
```

To ensure that users can access NetWare utilities, DOS directories, and applications, we recommend you map search drives to these directories in the following order:

- ◆ Map the first search drive (S1:) to the SYS:PUBLIC directory, which contains the NetWare utilities for DOS and MS Windows workstations.
- ◆ Map the second search drive (S2:) to the DOS directory if users access DOS from the network.
- ◆ Map the third and subsequent search drives (S3:, S4:, etc.) to directories containing applications and the electronic NetWare documentation.
- ◆ If your users are running MS Windows from the network, map a search drive to the MS Windows directory for the MS Windows group.

To avoid inadvertently changing the order of any search drives that must be mapped to a specific drive letter, you can map all remaining search drives with the number S16:, which assigns the next lowest search number each time it is used.

This command assigns the next available drive letter to the search drive, without displacing the previous search drives.

If you have an application that requires a particular drive letter, you can use the following command to map the search drive, replacing *drive* with the drive letter:

MAP S16:=*drive*:=*path*

If you map a search drive using a number already assigned to a search drive, NetWare makes the old search drive a network drive. The letter assigned to the old search drive remains assigned as the converted drive mapping. The new search drive takes the next unused letter of the alphabet.

Mapping Drives to Directory Map Objects

Another way to map a drive to a directory is to create a Directory Map object that points to the directory. Then, if you move the directory, you only need to change the Directory Map object rather than all of the login scripts that may include that mapping.



It is best to use map objects in the user's current context. Don't use complete names that point to other contexts. If the map is in another context, you should create an alias that points to the real directory map object. See "Alias object" in *Concepts*.

For example, to map a search drive to a Directory Map object whose complete name is APPL.SALES_LA.NOVELL_US, add the following line to the login script:

```
MAP S2:=.APPL.SALES_LA.NOVELL_US
```

In the above example, the Directory Map object's name begins with a period, which indicates that the drive is mapped to the drive root.

If the user whose login script this line appears in is also located in the SALES_LA.NOVELL_US context, the MAP command does not have to specify the Directory Map object's complete name. Instead, the line would be

```
MAP S2:=APPL
```

For more information about using Directory Map objects, see "Loading Operating Systems and Applications onto the Network" on page 117.

Examples

- ◆ To map the first search drive to the SYS:PUBLIC directory (which contains the NetWare utilities for DOS and MS Windows workstations) add the following line to the login script:

```
MAP S1:=SYS:PUBLIC
```

The second search drive should be mapped to the DOS directory if users run DOS from the network.

- ◆ If your network has more than one DOS directory, use variables to indicate the directory path. These variables are replaced by the correct information from the workstation software when each user logs in.

To use variables for the DOS directory path, enter the following command in the login script:

```
MAP S2:=SYS:PUBLIC\%MACHINE%\%OS%\%OS_VERSION
```

- ◆ If all users have the same types of computers and are using the same version of DOS, you probably only have one DOS directory. In this case, add a line similar to the following, substituting the correct directory names:

```
MAP S2:=SYS:PUBLIC\IBM_PC\MSDOS\50
```

You can also create a Directory Map object that points to the DOS directory, then map the search drive to the Directory Map object.

For more information about creating DOS directories, see "Loading DOS onto the Network" on page 117.

- ◆ To map the next available search drive to the SYS:APPL\WORDPROC directory, add the following line to the login script:

```
MAP S16:=SYS:APPL\WORDPROC
```

If you have mapped a Directory Map object to this directory, you can substitute the Directory Map object's name for the directory path.

For example, suppose you created a Directory Map object called WPROC, which is located in the context SALES.NOVELL_US, and mapped that object to the SYS:APPL\WORDPROC directory.

Following the rule of no complete names in a login script, you would create an alias in your own context to the object in .SALES.NOVELL_US. Then use the following line in your login script:

```
MAP S16:=WPROC
```

- ◆ Suppose an application in the FORM directory requires that it reside in the root directory of drive P:, but you don't want to put the application in the root directory for security reasons.

You can map a fake root to the directory and map a search drive to it at the same time by adding the following line to the login script:

```
MAP ROOT S16:=P:=SYS:APPL\FORM
```

- ◆ To map Richard's first four regular drive mappings to his home directory, the SYS:PUBLIC\OS2 directory (which contains the NetWare utilities for OS/2), the REPORTS directory, and the PROJECTS directory, add the following four lines to Richard's user login script:

```
MAP *1:=VOL1:HOME\RICHARD
```

```
MAP *2:=SYS:PUBLIC\OS2
```

```
MAP *3:=VOL1:ACCOUNTS\REPORTS
```

```
MAP *4:=VOL1:UPDATES\PROJECTS
```

To map Richard's fifth drive to the PUBLIC directory on a NetWare 3.11 server named FS1, you need to include the server name in the MAP command. Use the following line in your login script:

```
MAP *5:=FS1\SYS:PUBLIC
```

- ◆ If you are mapping a drive to a directory that is located on a volume within your current context, include the volume's common name in the MAP command, as demonstrated in the previous examples.

If you are mapping a drive to a volume that is not in your current context, first create an alias to that volume; then include the volume's common name in the MAP command.

For example, if the complete name of a volume not in your current context is VOL1.SALES.NOVELL_US, then create an alias named VOL1: to that volume. The MAP command would include only this common name.

To map a drive to the APPL directory in this volume, the line in the login script would be

```
MAP *6 :=VOL1:APPL
```

Additional Information

For more information about	Refer to
Drive mappings	"Drive mapping" in <i>Concepts</i>
Enabling users to run DOS from the network	"Loading DOS onto the Network" on page 117
Setting paths for OS/2 workstations	<i>NetWare Client for OS/2 User Guide</i>
Using Directory Map objects	"Loading Operating Systems and Applications onto the Network" on page 117

NO_DEFAULT

Use NO_DEFAULT in a container or profile login script if you do not want the default user login script to run.

Command Format

```
NO_DEFAULT
```

Using NO_DEFAULT

If you do not want to create any user login scripts, and you do not want the default user login script to run, add this command to either the container or the profile login script.

If you have created a user login script for someone, that login script executes whether or not the NO_DEFAULT command is in the container or profile login script.

NOSWAP

Use NOSWAP to prevent the LOGIN utility from being moved out of conventional memory into higher memory (if available) or onto the disk to execute a # command and LOGIN at the same time.

Command Format

NOSWAP

Using NOSWAP

LOGIN always swaps to extended or expanded memory unless NOSWAP is specified on the command line or in the login script.

If you do not want LOGIN to be temporarily stored in higher memory or on the workstation's disk, use the NOSWAP command. NOSWAP prevents LOGIN from being swapped out of conventional memory.

Then, if the workstation does not have enough memory to handle both LOGIN and the # command, the # command fails but the rest of the login script executes as usual.

If you want LOGIN to be swapped out of conventional memory immediately every time a # command is executed, place the SWAP command in the login script, before the # command.

Additional Information

For more information about	Refer to
Using the # command	"# (Execute External Program)" on page 189
Using the SWAP command	"SWAP" on page 231

PAUSE

Use PAUSE to create a pause in the execution of the login script.

Command Format

PAUSE

Using PAUSE

Enter this command in your login script at any point you want a pause to occur.

You can add PAUSE to the login script following a message so that the user has time to read the message before it scrolls off the screen.

If you include PAUSE, the message "Strike any key when ready..." appears on the workstation screen. The LOGIN utility then waits for a key to be pressed before it executes the rest of the login script.

PCCOMPATIBLE

Use PCCOMPATIBLE to enable the EXIT "command" login script command to work if your workstation's LONG MACHINE NAME is not IBM_PC.

This command does not apply to OS/2 workstations.

Command Format

PCCOMPATIBLE

Using PCCOMPATIBLE

If your computer is an IBM PC compatible machine and not an IBM PC, use PCCOMPATIBLE in your login script to inform the LOGIN utility that your machine's long name is something other than IBM_PC. The LONG MACHINE NAME (AT&T, COMPAQ, or others) must be included in the NET.CFG file.

Place the following anywhere before EXIT in the login script:

PCCOMPATIBLE

Example

If you have a Hewlett Packard computer and you have changed the LONG MACHINE NAME to HE_PAC in the NET.CFG file, and you want to exit to NETADMIN from within your login script, put the following commands in your login script:

PCCOMPATIBLE

EXIT "NETADMIN"

Additional Information

For more information about	Refer to
Setting the machine name in NET.CFG	<i>NetWare Client for DOS and MS Windows User Guide</i>
Using the EXIT command	"EXIT" on page 199

PROFILE

Use PROFILE in a container script to set or override a user's assigned or command-line-specified profile script. It is useful when defining a group profile.

Command Format

PROFILE *profile_object_name*

Example

To override the profile script assigned to a user or specified at the command line, and cause the user to execute a PROFILE script called *team_profile*, use the following command:

PROFILE *team_profile*

REMARK

Use REMARK to insert explanatory text into your login script. This text does not display on the screen.

Command Format

REM[ARK] [*text*]

or

***** [*text*]

or

; [*text*]

Replace *text* with the comment you want to include in the login script.

Using REMARK

To include explanatory text in your login script, begin a line with REMARK, REM, an asterisk (*), or a semicolon (;).

Any text that follows these symbols is ignored when the LOGIN command executes your login script. Remarks do not appear on the screen.

Using remarks in your login script can make the script much easier for you or other users to read and understand.

The REMARK command and its associated text must be the only entry on a line. Placing remarks on the same line as other login script commands can cause errors.

If a remark is several lines long, begin each line with the remark keyword (REMARK, REM, an asterisk, or a semicolon).

Example

The following are examples of explanatory text that you might use with the REMARK command and its variants:

```
* This is Richard's login script
; Mapped network drives follow:
REM The next mapping is a fake root.
REMARK This login script is for new users.
```

SCRIPT_SERVER

NetWare 2 and 3 users can use SCRIPT_SERVER to set a home server from where the bindery login script is read.



SCRIPT_SERVER has no effect on NetWare 4 users.

Command Format

```
SCRIPT_SERVER server_name
```

Using SCRIPT_SERVER

This command has no effect on NetWare Directory Services.

SET

Use SET to set a DOS or OS/2 environment variable to a specified value.

For OS/2 workstations, SET commands affect the environment only while the login script is running; the settings disappear when LOGIN terminates.

Command Format

```
[TEMP] SET name="value"
```

Replace *name* with an environment parameter that identifies the environment you want to change.

Replace *value* with identifier variable substitutions. Values must be enclosed in quotation marks.

To change the environment for the login script, but not for the workstation itself after the login script has finished executing, use the optional keyword TEMP.



Variables set in the login script for an OS/2 workstation affect the environment only while the login script is running. Therefore, the TEMP option does not work with OS/2 workstations.

Using SET

Use the SET login script command the same way you use the DOS command called SET. However, when you use SET in a login script, you must enter quotation marks (" ") around values.



If a variable is set to a path that ends in a \", these two characters are interpreted as an embedded quote preceded by an escape character. To avoid this problem, use two backslashes before the ending double quotes (\\").

SET commands do not have to be included in login scripts.

For example, you may decide that it is easier to put some SET commands in the workstation's AUTOEXEC.BAT file. Where you use SET commands depends upon your individual needs.

For information about values you can set, see the SET command in your DOS or OS/2 documentation.

This command does not work in a login script if the DOS workstation's environment is too small. In this case, you should set the environment size in the CONFIG.SYS file.

See the SHELL command in your DOS manual for more information about the environment size. (OS/2 workstations do not have this limitation.)

After you use the SET command to set a value for an environment variable, you can use that variable in other login script commands.

To include an environment variable as an identifier variable in a command, enclose the name of the variable in angle brackets; for example, <emailuser>.

Examples

- ◆ You can use SET to make a prompt display the current directory path, such as F:\HOME\MARY>, rather than just the drive letter. To do this, add the following line to the login script:

```
SET PROMPT="$P$G"
```

"\$P" lists the current directory path; "\$G" displays a ">" (greater than) character. See your DOS or OS/2 manual for more information.

- ◆ To set a path for a program called DAILY, which is in the REPORTS subdirectory beneath drive G:, you would add the following line:

```
SET PATH="G:\REPORTS\DAILY"
```

This sets the variable PATH to G:\REPORTS\DAILY.



Setting the variable PATH in the login script removes any search drives previously assigned. Use SET PATH only before you map search drives. SET PATH also overwrites any paths set in the user's AUTOEXEC.BAT file.

To display this path, you can include PATH as an identifier variable in a WRITE command by enclosing the variable (not the value) in angle brackets. For example, the following line displays "My path is G:\REPORTS\DAILY"

```
WRITE "My path is "%<path>"
```

- ◆ To include an environment variable in a MAP command, precede the variable with a percent sign (%).

For example, you could include the following lines in a login script to set and map a drive to the variable NWS:

```
SET NWS="C:\XYZ"  
MAP S16:=%<NWS>
```

Additional Information

For more information about	Refer to
Using environment variables as identifier variables in other login script commands	"Identifier Variables" on page 236
Using the SET command	Your DOS or OS/2 manual

SET_TIME

Use SET_TIME to set the workstation time equal to the time on the NetWare server that the workstation first connects to.

Command Format

```
SET_TIME ON|OFF
```

Using SET TIME

The default value is SET_TIME ON, which means the workstation time is set to the NetWare server time whenever the user logs in. If you include SET_TIME OFF in the login script, the workstation time does not update to the server's time.

SHIFT

Use SHIFT to change the order in which %*n* identifier variables are interpreted in the login script. SHIFT allows users to enter LOGIN parameters in any order.

Command Format

SHIFT [*n*]

Replace *n* with the number of places you want the variable to shift. The default is SHIFT 1.

Using SHIFT

You can shift up to 10 arguments.

When users execute LOGIN, they can include additional parameters. Each of these parameters is assigned a %*n* variable; in this way, the parameter's real value can be substituted for the %*n* variable that appears in the login script.

In the login script, you can add SHIFT with a positive or negative number to move the variables in either direction. For example, SHIFT -3 moves each %*n* variable three positions to the left.

Example

When Mary logs in, she wants to access her word processing program, change the way it is set up, and map a drive to her work directory called ACCNTS.

Mary also has a command in her login script to map a drive to her LOTUS directory, but she does not need it today. The commands in Mary's login script are shown here.

```
LOOP  
IF "%2"="WP" THEN  
    SET WP="\U-CML\B-10\D"  
    MAP S16:=SYS:APPL\WP\SETUP  
IF "%2"="ACCNTS" THEN  
    MAP G:=SYS:ACCNTS  
IF "%2"="LOTUS" THEN  
    MAP S16:=SYS:APPL\LOTUS  
SHIFT  
IF "%2"<>" THEN GOTO LOOP
```

(In the last line, "IF "%2" <>" is followed by closed quotation marks, which means "If %2 isn't blank".)

With these commands in her login script, Mary can log in to the primary file server (named FS1) using her username, MARY, as follows:

```
LOGIN FS1\MARY WP ACCNTS
```

The parameters in Mary's LOGIN command are given the following values:

```
%0=FS1  
%1=MARY  
%2=WP  
%3=ACCNTS
```

Mary's login script looks for %2, which is WP, and sets the word processing environment. Then the login script shifts the variables one to the right so that %2 now becomes ACCNTS. Upon executing the loop, the login script maps a drive to the ACCNTS directory.

Mary could also change the order of her LOGIN command without affecting the way her work environment is set up, as follows:

```
LOGIN MARY ACCNTS WP
```

The parameters in this LOGIN command are given the following values:

%0=FS1
%1=MARY
%2=ACCNTS
%3=WP

In this case, Mary's login script looks for %2, which is now ACCNTS.

The login script maps a drive to the ACCNTS directory. Then the login script shifts the variables to the right so that %2 now becomes WP.

Upon executing the loop, the login script sets the word processing environment.

Additional Information

For more information about	Refer to
Using %n variables in login scripts	Table 3-4 on page 237

SWAP

Use SWAP to move the LOGIN utility out of conventional memory into higher memory (if available) or onto the disk. This allows execution of a # command and LOGIN at the same time.

Command Format

SWAP [*path*]

You can replace *path* with either a drive letter or a full directory path beginning with the NetWare volume name.

Using SWAP

By default, the LOGIN utility always swaps to extended or expanded memory, unless NOSWAP is specified on the command line or in the login script.



Note

The SWAP option doesn't work with the DR DOS 6.0 EMM386 Memory Manager option unless upper memory is disabled.

If you specify a path in the SWAP command, LOGIN swaps into the directory you specified. If the directory specified in that path does not exist or if you do not have rights there, LOGIN prompts you for another path.

If you don't specify a path, LOGIN swaps either into higher memory (if available) or to the current drive. If LOGIN tries to swap to the current drive and you don't have rights to the current drive, LOGIN prompts you for a path to use. If you specify a valid path, LOGIN always swaps to the specified path.

Then, if the workstation doesn't have enough memory to handle both LOGIN and the # command, the # command fails but the rest of the login script executes as usual.

Additional Information

For more information about	Refer to
Using the # command	"# (Execute External Program)" on page 189
Using the NOSWAP command	"NOSWAP" on page 221

TEMP SET

See "SET" on page 226.

WRITE

Use WRITE to display messages on the workstation screen when a user logs in.

Command Format

```
WRITE "[text][%identifier]" [;][identifier]
```

Replace *text* with the words you want to display on the screen.

Replace *identifier* with a variable you want to display, such as a user's login name. (See "Using Identifier Variables" on page 240 for a complete list of variables.)

Using WRITE

Text you want to display must be enclosed in quotation marks ("").

There are several ways to display variables in the text message. The way you enter the variable in the WRITE command determines the display format, as follows:

- ◆ If you type the identifier variable exactly as shown, with no special punctuation, only the variable is displayed on the screen. (See "Login Script Identifier Variables" on page 237)
- ◆ If you enclose the identifier variable inside quotation marks, precede the variable with a percent sign (%), and type it in uppercase letters.

This method is often used to combine regular text with an identifier variable because both the text and the variable can be enclosed in the same quotation marks.

To join several text strings and identifier variables into a single display without enclosing the variables in quotation marks, use a semicolon between the text and the variables.

If you have several WRITE commands, each one appears on a separate line on your workstation. However, if you put a semicolon at the end of all but the last WRITE commands, the displays all appear as one continuous sentence or paragraph (although they may wrap onto additional lines on the workstation's screen).

Text strings can include the following special characters:

Character	Meaning
\r	Makes a carriage return occur
\n	Starts a new line of the text
\"	Displays a quotation mark on the screen
\7	Makes a beep sound

In addition to the semicolon, there are additional operators you can use to form compound strings (in other words, to join text and identifier variables into one command). These operators are listed in the following table, in order of precedence.

Operator	Meaning
* / %	Multiply, divide, modulus
+ -	Add, subtract
>> <<	Shift left or right (1000 >> 3 becomes 1)

Examples

- ◆ To display the message “Hello,” add the following line to the login script:

```
WRITE "Hello"
```

- ◆ To display the user’s last name (surname) along with a greeting, add the identifier `LAST_NAME` to the command. To do this, either join the text and the identifier with a semicolon or include the variable in the quotation marks with the text.

Either of the following lines displays “Hello, Smith” when user Bob Smith logs in:

```
WRITE "Hello, " ; %LAST_NAME
```

```
WRITE "Hello, %LAST_NAME"
```

- ◆ To make a beep sound occur while the phrase “Good morning” appears on the screen, add the following line to the login script:

```
WRITE "Good %GREETING_TIME \7"
```

Identifier Variables

With many login script commands, you can take advantage of identifier variables to make your login script more efficient and flexible.

Identifier variables allow you to enter a variable (such as `LAST_NAME`), rather than a specific name (such as `Smith`) in a login script command. When the login script executes, it substitutes real values for the identifier variables.

By using the variable, you can make the login script command applicable to many users instead of limiting it to one user.

For example, the command

```
WRITE "Hello, "%LAST_NAME
```

displays the following message on Bob's Smith's workstation screen when he logs in:

```
Hello, SMITH
```

Similarly, when Mary Jones logs in, the message she sees is

```
Hello, JONES
```

In the previous example, when the user logged in, the user's actual last name was substituted for the `LAST_NAME` variable in the command.

Table 3-4 lists all the available identifier variables.

Table 3-4
Login Script Identifier Variables

Category	Identifier Variable	Function
Date	DAY	Day number (01 through 31)
	DAY_OF_WEEK	Day of week (Monday, Tuesday, etc.)
	MONTH	Month number (01 through 12)
	MONTH_NAME	Month name (January, February, etc.)
	NDAY_OF_WEEK	Weekday number (1 through 7; 1=Sunday)
	SHORT_YEAR	Last two digits of year (94, 95, 96, etc.)
	YEAR	All four digits of year (1994, 1995, 1996, etc.)
Time	AM_PM	Day or night (am or pm)
	GREETING_TIME	Time of day (morning, afternoon, or evening)
	HOUR	Hour (12-hour scale; 1 through 12)
	HOUR24	Hour (24-hour scale; 00 through 23; 00=midnight)
	MINUTE	Minute (00 through 59)
	SECOND	Second (00 through 59)
User	%CN	User's full login name as it exists in NDS
	ALIAS_CONTEXT	"Y" IF REQUESTER_CONTEXT is an Alias
	FULL_NAME	User's unique username. It is the value of the FULL_NAME property for both NDS and bindery-based NetWare. Spaces are replaced with underscores
	LAST_NAME	User's last name (surname) in NetWare Directory Services, or full login name in bindery-based NetWare
	LOGIN_CONTEXT	Context where user exists
	LOGIN_NAME	User's unique login name (long names are truncated to eight characters)

Table 3-4 *continued***Login Script Identifier Variables**

Category	Identifier Variable	Function
	MEMBER OF "group"	Group object that the user is assigned to
	NOT MEMBER OF "group"	Group object that the user is not assigned to
	PASSWORD_EXPIRES	Number of days before password expires
	REQUESTER_CONTEXT	Context when login started
	USER_ID	Number assigned to each user
Network	FILE_SERVER	NetWare server name
	NETWORK_ADDRESS	IPX external network number of the cabling system (8-digit hexadecimal number)
Workstation	MACHINE	Type of computer (IBM_PC, etc.)
	NETWARE_REQUESTER	Version of the NetWare Requester for OS/2 or VLM users
	OS	Type of operating system on the workstation (MSDOS, OS2, etc.)
	OS_VERSION	Operating system version on the workstation (3.30, etc.)
	P_STATION	Workstation's node number (12-digit hexadecimal)
	SHELL_TYPE	Version of the workstation's DOS shell (1.02, etc.); supports NetWare 2 and 3 shells and NetWare 4 Requester for DOS.
	SMACHINE	Short machine name (IBM, etc.)
	STATION	Workstation's connection number
DOS environment	<variable>	Any DOS environment variable can be used in angle brackets (<path>, etc.). To use a DOS environment variable in MAP, COMSPEC, and FIRE PHASERS commands add a percent sign (%) in front of the variable. For example, MAP S16:=%<path>

Table 3-4 continued

Login Script Identifier Variables

Category	Identifier Variable	Function
Miscellaneous	ACCESS_SERVER	Shows whether the access server is functional (TRUE=functional, FALSE=not functional)
	ERROR_LEVEL	An error number (0=no errors)
	%n	Replaced by parameters the user enters at the command line with the LOGIN utility. For more information, see Table 3-4 on page 237.
Object properties	property name	<p>You can use property values of NDS objects as variables. Use the property values just as you do any other identifier variable. If the property value includes a space, enclose the name in quotation marks.</p> <p>To use a property name with a space, within a WRITE statement, you must place it at the end of the quoted string:</p> <pre>WRITE "Given name=%GIVEN_NAME" IF "%MESSAGE SERVER"="MS1" THEN MAP INS S16:=MS1\SYS:EMAIL</pre> <p>To see a list of object properties, see Appendix A, "NDS and Bindery Objects and Properties," of <i>Utilities Reference</i>. Not all properties are supported.</p>

Using Identifier Variables

When using identifier variables in login script commands, observe the following conventions.

- ◆ Identifier variables are used most often with commands such as IF...THEN, MAP, and WRITE. They can also be used with commands for which you can specify a path, such as COMSPEC.
- ◆ Type the variable exactly as shown.
- ◆ To use DOS environment variables as identifiers, enclose them in angle brackets.
- ◆ Identifier variables can be placed within literal text strings in a WRITE statement. However, the identifier variable must be in uppercase letters and preceded by a percent sign. (Literal text is the text that is displayed on the screen, such as "Sales report is due today." Literal text must be enclosed in quotation marks.)

Examples:

- ◆ If user Smith logs in during the morning, both of the following lines display the same message on his screen ("Good morning, SMITH"):

```
WRITE "Good "; GREETING_TIME; ", "; LAST_NAME  
WRITE "Good %GREETING_TIME, %LAST_NAME"
```

- ◆ To use DOS environment variables as identifiers, enclose them in angle brackets. The following example uses the DOS environment variable "path:"

```
WRITE "my path is "%<path>
```

The text displayed on the screen is similar to

```
my path is z:.;y:.;c:\WINDOWS
```

Using LOGIN Parameters with %n Variables

Some variables in a login script can be indicated by a percent sign (%) followed by a number from 0 to 9.

When a user logs in, he or she can type additional parameters that the LOGIN utility passes to the login script. The login script then substitutes these parameters for any %n variables in the login script. These variables are replaced in order by the parameters the user typed when executing the LOGIN utility.

The %0 variable is replaced by the name of the NetWare server the user typed at the command line, and %1 is replaced by the user's login name. The remaining variables change, depending on what the user types when executing LOGIN. The %n variables must precede all command line options.

The SHIFT login script command allows you to change the order in which these %n variables are substituted. For more information about the SHIFT command, see "SHIFT" on page 229.

The %n variables can be used in WRITE statements if included within the quotes:

```
WRITE "My login name is %1."
```

Example

Suppose a login script contains the following commands:

```
IF "%2"="SALES" THEN  
  WRITE "Meeting today"  
END  
IF "%3"="LEGAL" THEN  
  WRITE "Report is due tomorrow"  
END
```

If user RON logged in by typing the following command:

LOGIN COUNT\RON SALES MARKETING

then the login script would substitute the values Ron entered at the keyboard for the %n variables in the login script, as shown here:

%0=COUNT

%1=RON

%2=SALES

%3=MARKETING

Since %2 is replaced by "SALES," the message "Meeting today" is displayed on Ron's screen. However, since %3 is replaced by "Marketing," Ron doesn't see "Report is due tomorrow."

Examples of Login Scripts

The following examples of login scripts may help you plan your own container, profile, and user login scripts. Each example login script is shown in a table.

The left column of each table shows the commands in the login script. The right column explains the command's purpose.

For instructions on creating login scripts, see "Creating or Modifying a Login Script Using NetWare Administrator" on page 177 or "Creating or Modifying a Login Script Using NETADMIN" on page 180.

Default Login Script

The default login script executes the first time User object ADMIN logs in. It also executes for any users who do not have user login scripts.

You can't modify the default login script because it is coded into the LOGIN utility. Instead, you can create container, profile, or user login scripts.

The following table lists the content of the default login script.

Table 3-5
Default Login Script

Login Script Command	Purpose
MAP DISPLAY OFF	Prevents map commands from displaying on the screen.
MAP ERRORS OFF	Prevents mapping errors from displaying on the screen.
MAP *1:=SYS:	Maps the first drive to volume SYS:.
MAP *1:=SYS:%LOGIN_NAME	Maps the first drive to the user's home directory, if LOGIN_NAME is the same as the user's home directory. If the user has no home directory, the first drive is still mapped to SYS:.
IF "%1"="ADMIN" THEN MAP *1:=SYS:SYSTEM	If the login name is ADMIN, the first drive is mapped to SYS:SYSTEM instead of to the user's home directory.

Table 3-5 *continued*

Default Login Script

Login Script Command	Purpose
MAP P:=SYS:PUBLIC	<p>If the user logs in from an OS/2 workstation, drive P: is mapped to SYS:PUBLIC.</p> <p>This mapping is used by OS/2 workstations, as explained in <i>NetWare Client for OS/2 User Guide</i>.</p> <p>If the user is not using an OS/2 workstation, this drive mapping is not included in the default login script.</p>
MAP INS S1:=SYS:PUBLIC MAP INS S2:=SYS:PUBLIC\%MACHINE%\OS\ %OS_VERSION	<p>If the user is using a DOS or MS Windows workstation, the first search drive is mapped to SYS:PUBLIC, where DOS-based NetWare utilities are stored.</p> <p>Then the second search drive is mapped to the directory where DOS is stored. (The two MAP commands are joined by a semicolon.)</p> <p>If the user logs in from an OS/2 workstation, these drive mappings are not included in the default login script.</p>
MAP DISPLAY ON	Allows MAP commands to display.
MAP	Displays a list of all drive mappings on the user's screen.

Container Login Script

The container login script should contain as much information as possible that applies to all users. An example is shown in the following table.

Table 3-6
Sample Container Login Script

Login script command	Purpose
MAP DISPLAY OFF	Prevents MAP commands from displaying on the screen as they are assigned.
MAP ERRORS OFF	Prevents mapping errors from displaying on the screen.
MAP *1:=SYS:	Maps the first drive to volume SYS:.
MAP *1:=SYS:%LOGIN_NAME	Maps the first drive to the user's home directory, if LOGIN_NAME is the same as the user's home directory. If the user has no home directory, the first drive is still mapped to SYS:.
IF "%1"="ADMIN" THEN MAP *1:=SYS:SYSTEM	If the login name is ADMIN, the first drive is mapped to SYS:SYSTEM instead of to the user's home directory.

Table 3-6 *continued*

Sample Container Login Script

Login script command	Purpose
<pre>IF OS="OS2" THEN MAP P:=SYS:PUBLIC ELSE MAP INS S1:=SYS:PUBLIC MAP INS S2:=SYS:PUBLIC\%MACHINE%\OS\ %OS_VERSION END</pre>	<p>If the user logs in from an OS/2 workstation, drive P: is mapped to SYS:PUBLIC.</p> <p>This mapping is used by OS/2 workstations, as explained in <i>NetWare Client for OS/2 User Guide</i>.</p> <p>If the user is not using an OS/2 workstation, the first search drive is mapped to SYS:PUBLIC, where DOS-based NetWare utilities are stored. The second search drive is mapped to the directory where DOS is stored.</p> <p>If all workstations use DOS, you can use the following two commands instead of the IF...THEN command:</p> <pre>MAP INS S1:=SYS:PUBLIC MAP INS S2:=SYS:PUBLIC\ %MACHINE%\OS%\OS_VERSION</pre>
<pre>IF MEMBER OF "WIN31" THEN MAP INS *2:=SYS:USERS%\%LOGIN_NAME\WIN31 MAP INS S16:=SYS:APPS\WINAPPS\WIN31 SET TEMP="P:\USERS%\%LOGIN_NAME\ WIN31\TEMP" END</pre>	<p>If the user who logs in is a member of the Group object WIN31, the next available drive is mapped to that user's MS Windows directory. Then the next available search drive is mapped to the MS Windows directory for the WIN31 group. Finally, the MS Windows TEMP directory is set to a subdirectory of the user's MS Windows directory.</p>
<pre>MAP INS S16:=VOL1:APPL\WP</pre>	<p>Maps the next available search drive to the directory that contains WordPerfect.</p>
<pre>MAP INS S16:=VOL1:APPL\LOTUS</pre>	<p>Maps the next available search drive to the directory that contains Lotus.</p>
<pre>MAP INS S16:=SYS:EMAIL</pre>	<p>Maps the next available search drive to the EMAIL directory.</p>

Table 3-6 *continued*

Sample Container Login Script

Login script command	Purpose
MAP O:=SYS:DOC	Maps drive O: to a directory necessary for running the electronic NetWare documentation (Novell DynaText*).
IF MEMBER OF "MANAGERS" THEN MAP *3:=VOL1:PROJECTS\REPORTS END	If the user belongs to the Group object MANAGERS, the login script maps the third network drive to the REPORTS directory.
IF MEMBER OF "TESTERS" THEN MAP *4:=INPUT:STATUS\UPDATES END	If the user belongs to the Group object TESTERS, the login script maps the fourth network drive to the UPDATES directory.
COMSPEC=S2:COMMAND.COM	Sets COMSPEC to the DOS command processor, located in the DOS directory (in the second search drive).
SET PROMPT="\$P\$G"	Sets the prompt to display the user's current directory path, followed by the ">" symbol.
MAP DISPLAY ON	Allows MAP commands to display.
MAP	Displays a list of all drive mappings.
WRITE	Displays a blank line between the list of mappings and following lines.
WRITE "Good %GREETING_TIME, %LAST_NAME."	Displays a greeting to the user. Example: "Good morning, SMITH."
WRITE "Your password expires in %PASSWORD_EXPIRES days."	Displays a message indicating the number of days before the user's password expires.
FIRE 3	Makes the phaser sound occur three times to tell the user that the login process is complete.

Profile Login Script

If you have groups of users with identical login script needs, you can create a Profile object, then create a login script for it. Then you can assign each user to be a member of that Profile object.

The following table shows an example of a profile login script you might create for users in the Profile object ACCOUNTING. This profile login script would execute after the container login script had executed.

Table 3-7

Sample Profile Login Script

Login script command	Purpose
MAP DISPLAY OFF	Prevents MAP commands from displaying on the screen as they are assigned.
MAP ERRORS OFF	Prevents mapping errors from displaying on the screen.
MAP INS S16:=VOL1:APPL\DB	Maps the first available search drive (after those assigned in the container login script) to the directory that contains the database program.
MAP *5:=VOL1:ACCOUNTS\NEW	Maps the fifth network drive (after those assigned in the container login script) to the NEW subdirectory.
MAP *6:=VOL1:ACCOUNTS\RECORDS	Maps the sixth network drive (after those assigned in the container login script) to the RECORDS subdirectory.
#WSUPDATE S1:IPXODI.COM /LOCAL	Executes WSUPDATE, which updates the IPXODI.COM file on the user's workstation with a new version of the file located in the first search drive.
MAP DISPLAY ON	Allows MAP commands to display.
MAP	Displays a list of all drive mappings.
WRITE	Displays a blank line between the list of mappings and following lines.

Table 3-7 *continued*

Sample Profile Login Script

Login script command	Purpose
<pre>IF DAY_OF_WEEK="FRIDAY" THEN WRITE "Weekly progress report is due today" FIRE 2 END</pre>	<p>On Fridays, the phaser sound occurs twice to alert the user while the message "Weekly progress report is due today" displays on the screen.</p>
<pre>PCCOMPATIBLE EXIT "NMENU WORK"</pre>	<p>Stops the profile login script and sends the user into a menu program called WORK.</p> <p>EXIT also prevents any user login scripts from executing.</p> <p>If you want a user login script to execute after the profile login script, put these lines at the end of the user login script instead.</p> <p>(DOS workstations with the machine name IBM_PC do not need the PCCOMPATIBLE line.)</p> <p>The EXIT command does not work on OS/2 workstations.</p>

User Login Script

The following table shows an example of a user login script for user Mary. The user login script executes after the container and profile login scripts have executed.

Table 3-8
Sample User Login Script

Login script command	Purpose
MAP DISPLAY OFF	Prevents MAP commands from displaying on the screen as they are assigned.
MAP ERRORS OFF	Prevents mapping errors from displaying on the screen.
MAP *7:=VOL1:MARY\PROJECTS\RESEARCH	Maps Mary's seventh network drive (after those assigned in the container and profile login scripts) to the RESEARCH subdirectory in her home directory.
MAP *8:=VOL1:FORMS	Maps Mary's eighth network drive (after those assigned in the container and profile login scripts) to the FORMS directory.
REM Mary needs access to FORMS while she's on the REM troubleshooting team. Remove this drive mapping REM when she's reassigned.	This remark is intended as a reminder to the person who created the login script. This remark isn't displayed on the user's screen. (Because the remark is several lines long, each line starts with REM.)
SET WP="/u-mjr/b-5"	Sets Mary's environment variables for her word processing application.
SET USR="mrichard"	Sets Mary's username (mrichard) for the electronic mail program.
#CAPTURE Q=FAST_Q NB TI=10 NFF	Executes the CAPTURE utility so Mary can print from nonnetwork applications.
PCCOMPATIBLE EXIT "NMENU TRAINING"	Stops the user login script and sends the user into a menu program called TRAINING. (DOS stations with the machine name IBM_PC don't need the PCCOMPATIBLE line.) EXIT doesn't work on OS/2 workstations.



chapter

4

Creating Menus

Creating, Converting, and Modifying Menu Files

Menus create a simple front end to the users' working environment. They make it easy to access network resources by presenting a list of options instead of requiring the user to enter cryptic DOS commands.

Menus can be shared by multiple users, requiring the creation of fewer menus. Or, for the unique requirements of some users, custom menus can be easily created as needed.

In the NetWare 4™ operating system, NMENU.BAT replaces the MENU.EXE utility included in earlier versions of NetWare. NMENU files are easy to script, require less memory, and run faster than did the old MENU utility files.

Existing menu files from earlier NetWare versions can be easily converted for use by the newer NMENU program (see "Converting Old Menu Files" on page 277). NMENU works in all versions of NetWare® 2, NetWare 3, and NetWare 4.

Because NMENU relies on special files to execute, and those files are dependent upon a few special commands, this chapter focuses on use of the scripting commands that make up a menu file.

Guidelines for planning your menus are located in "Planning Your Menus" on page 259. Commands and their usage are explained in "Using the NMENU Commands" on page 260. Several example programs are given in "Creating a Menu File" on page 272.

Making menus available to users (see "Setting Up the User Environment" on page 275) and converting old menu files (see "Converting Old Menu Files" on page 277) are the final two topics of this chapter.

Getting Acquainted with NMENU

Since the sole function of NMENU is to cause a scripted file to execute, there is little to learn about NMENU.BAT except its syntax. The command is followed by the menu filename.

NMENU Syntax

```
nmenu filename
```

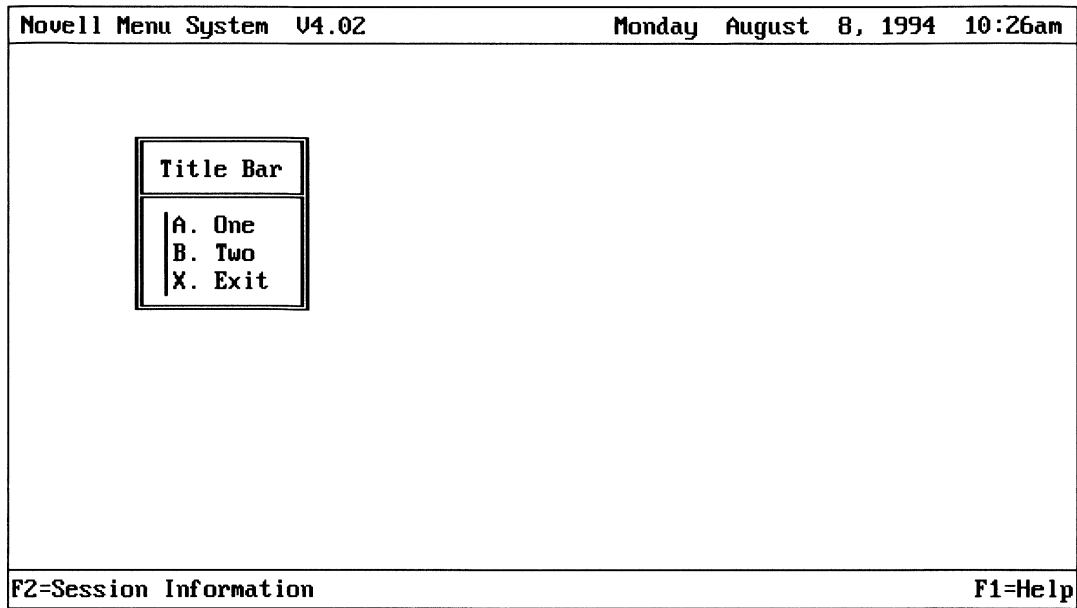
Replace *filename* with the menu filename.

What Menus Look Like

Figure 4-1 shows the display created by a single-window menu with only three options. Whether a menu displays one or more windows, the top and bottom bars of the screen remain the same.

The top bar always displays the NMENU version number on the left and the day, date, and time on the right. The bottom bar displays the available options.

Figure 4-1
A Single-Window
Menu



When you create menus with multiple windows, the windows cascade from left to right (see "A Multiple-Window Menu" on page 255). Window sizing is automatically determined by the content of each window.

What Makes Menus Work

You create menu files with a text editor and save the file with a .SRC extension. Then you use the MENU MAKE program to compile the file. It is given a .DAT extension and, as a compiled file, is no longer editable. Any edits must be made to the .SRC file and then the .SRC file must be recompiled.

Elements of a Menu

There are three primary elements to every menu: MENU, ITEM, and EXEC. As these elements are expanded and repeated, controlled options are displayed within each window.

Figure 4-2 shows the .SRC file used to create the menu in Figure 4-1. This file shows how these three elements are used. Details about these elements and their options are covered later in the chapter.

Figure 4-2
Primary Elements of a Menu

```
Menu 1,Title Bar
  Item One {pause}
    Exec dir
  Item Two {show pause}
    Exec ver
  Item ^XExit
    Exec EXIT
```

◆ Element 1: MENU

MENU specifies a window within the menu file. It includes a menu number and a menu name. The menu name is the title bar for the menu.

◆ Element 2: ITEM

ITEM includes the text the user will see. It is given an indicator letter by default, or you can determine the character to precede it. Each option to be displayed in the window must be preceded by the word ITEM. Some options are available for the ITEM line; these options are discussed in "ITEM" on page 262.

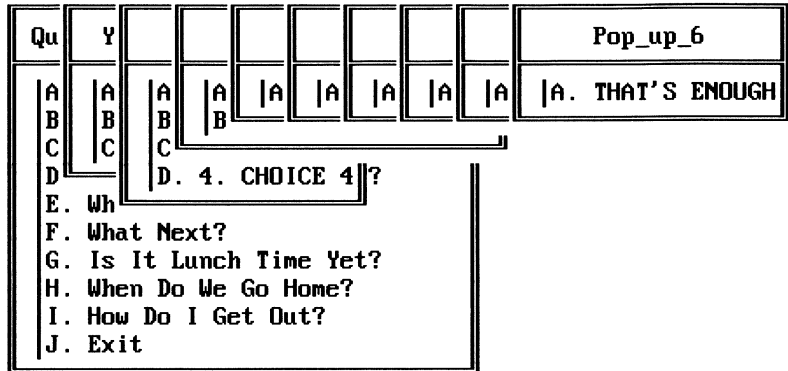
◆ Element 3: EXEC

EXEC is the primary command for the third element. In the menu file shown in Figure 4-2, EXEC issues three separate commands: a directory listing, a display of the current version of DOS running, and a command to exit the menu.

Combining the Elements

Figure 4-3 illustrates a menu with 10 windows, created by including 10 MENU commands in the same file. Each window is automatically sized and cascaded across the screen.

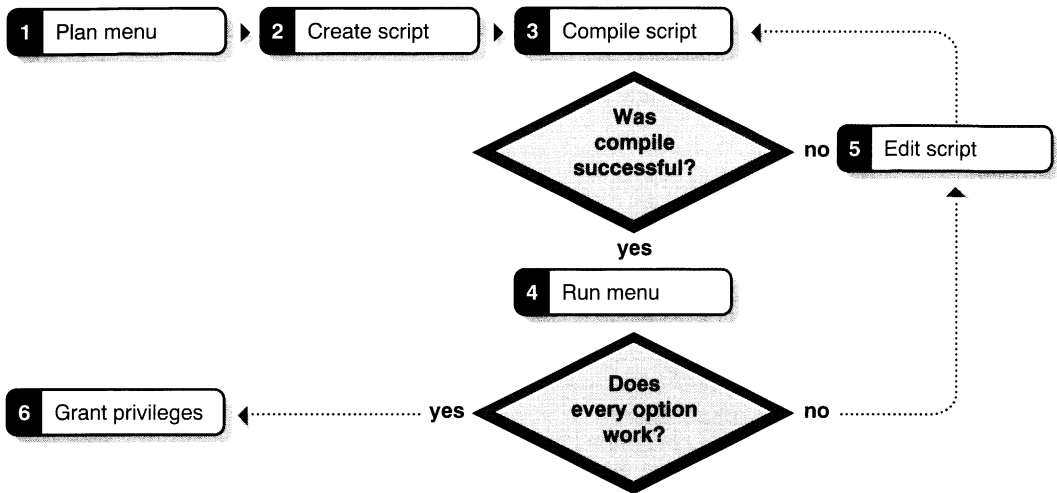
**Figure 4-3
A Multiple-Window
Menu**



Menu Creation Steps


There are generally six steps to follow when considering a new menu. The relationship between steps is shown in Figure 4-4.

Figure 4-4
Charting the Steps
in Creating a Menu



Procedure

Each step in the flowchart is explained below.



1. Plan and design your menu. Before starting, answer the following questions:

- ◆ Who is it for?
- ◆ What do the users need access to?
- ◆ Should they have access to a NetWare prompt?
- ◆ Should they be forced to log out when exiting?
- ◆ How complex is the menu? Will it need to be in multiple files?

2. Use a text editor to create your menu with a .SRC extension.

3. Compile the .SRC file with the MENU MAKE program. This will create a .DAT version of the file.

The results of any errors occurring during the compile process, including detailed error messages per line of script, are displayed on your screen. Fix the errors in the .SRC version of the file according to the error messages, and then recompile. When you have eliminated all of the errors, continue with Step 4.

- 4. Run NMENU filename.dat to verify that it does what you expect.**
- 5. Change the program if needed, as described in Step 3, and then repeat Step 4. When Step 4 is successful, go on to Step 6.**
- 6. Provide access to the intended users by placing the .DAT file in an appropriate directory and granting sufficient rights to the users.**



Note

Information on placement of files and required rights is in “Making Menus Work” on page 275.

Following these six steps can help you learn how to use the scripting language and how to manage the menu files you create.

Use of the NMENU program is almost unlimited. When considering the uses for menus, remember that they can be as simple as presenting available application programs (see Figure 4-5).

**Figure 4-5
Example of a Simple
Menu**

AV	Available Word Processors
A	A. WordPerfect 5.1
B	B. WordPerfect 6
X	C. Word

They also can be more complex, such as guiding data entry for cataloging (see Figure 4-6).

Figure 4-6
Example of a More
Complex Menu

Data Entry
A. Add Entry to Catalogue
User Input Requested
Enter Title Enter Author's Name Category: 1=Apps, 2=Prod, 3=Self

The text used to create the four menus shown in this section (Figure 4-1, Figure 4-3, Figure 4-5, and Figure 4-6) is included in "Example Menu Programs" on page 272.

Before you look at how these menus were created, however, you should be familiar with the guidelines for planning menus and with the scripting language rules.

Planning Your Menus

Several considerations are important when planning a menu. Details are presented throughout this chapter, but some general rules include the following:

- ◆ You must include an EXIT or LOGOUT command in the menu to be able to leave the menu. For security reasons, the Esc key doesn't work.
- ◆ The maximum number of windows per menu is 11 (1 main window and 10 subwindows).
- ◆ The maximum number of menus you can define in a single file is 255.
- ◆ The maximum "MENU" name length (title bar) is 40 characters.
- ◆ The maximum number of "ITEM" lines per window is 12.
- ◆ The maximum width of each window "ITEM" is 40 characters.
- ◆ The maximum text file width is 78 characters.
- ◆ The main menu must be at the beginning of the .SRC file.
- ◆ Submenus defined in the same file are called with the SHOW command by menu number.
- ◆ Menus in separate files are called as submenus, by menu name, with the LOAD command.
- ◆ If a command wraps to another line, type a plus sign (+) at the end of the line and continue the command on the next line.
- ◆ Menu colors are determined by the settings in the COLORPAL utility. See "COLORPAL" in *Utilities Reference*.

Using the NMENU Commands

Menu programs use two types of commands: organizational and control. The following sections describe these commands and their options.

NMENU Organizational Commands and Options

Organizational commands establish the content and organization of the menus the user sees on the screen. Use these commands to determine what the menus look like.

Table 4-1 gives an overview of the organizational commands. Detailed information about each command follows the table.

Table 4-1

NMENU Organizational Commands

Command	Explanation
MENU	Marks the beginning of a new menu or submenu screen in the text.
ITEM	Identifies an item to be listed on the menu. Specifies execution parameters.

MENU

Indicates a new window definition within your menu file.

Command Format

MENU *menu_number,menu_name*

Replace *menu_number* with the number you want to assign to this menu. Menus are called and displayed by their number.

A menu number can be any number from 1 through 255. Each menu within a source file must be assigned a unique number, but the number sequence doesn't matter.

Replace *menu_name* with the title you want to appear at the top of the menu. Menu names can be a maximum of 40 characters long.

Using MENU

The first menu defined in the source file is always the first menu displayed, no matter what number is assigned to it. Subsequent menus are referenced by their numbers, no matter what order they appear in the .SRC file.

For example, if you have defined three menus and assigned them numbers 1, 2, and 5, and menu 5 is the first menu in the source file, menu 5 is displayed first. Menus 1 and 2 are referenced in control commands and displayed depending upon user selection.

Example

To define a main menu for a program to be used by accountants in your company, type a line similar to the following:

MENU 1,Accounting Main Menu

Since this is the first menu you want to appear, place it at the beginning of your menu program.

ITEM

Indicates an option in the menu.

Command Format

```
ITEM item_name {[option...]}
```

Replace *item_name* with the list item you want to appear in the menu. The maximum length for an item is 40 characters.

Replace *option* with one or more of the options shown in the following table. Separate multiple options with a space. Enclose all options for a single ITEM in a single set of braces.

These options provide you with better control of menu execution. You determine how memory is used, which directory the user remains in, and what information is presented to those users.

Table 4-2

ITEM Command Options

ITEM Option	Explanation
BATCH	<p>Removes the menu program from memory before executing the item.</p> <p>Without this option, a portion of the memory stays resident, requiring approximately 32 KB of available memory, even though an application may be running.</p> <p>For example, if you enter</p> <pre>ITEM Word 5.0 {BATCH}</pre> <p>the menu program is removed from memory when Word 5.0 is executed.</p> <p>Setting this option automatically sets the CHDIR option. Don't use this command with the EXEC DOS command; use EXEC CALL.</p> <p>For more information, see "EXEC" on page 266.</p>

Table 4-2 *continued*
ITEM Command Options

ITEM Option	Explanation
CHDIR	<p>Changes back to the drive and directory that were in effect before an ITEM was executed.</p> <p>For example, to change back to the drive and directory the user was in before executing a word processing application in another directory, type]</p> <p>ITEM WordPerfect 5.0 {CHDIR}.</p> <p>When the word processing application is closed, the directory is changed to the original directory for the menu.</p>
PAUSE	<p>Allows users to read messages associated with a command being executed from the menu, by pausing the screen display.</p> <p>The message "Press any key to continue" is also displayed, and the screen does not change until a user presses a key.</p> <p>For example, to display a message when a user selects the item "DOS Copy" from the menu, include the line:</p> <p>ITEM DOS Copy {PAUSE}</p> <p>When the copy function is complete, the display waits at the "Press any key..." prompt for a key to be pressed before returning to the menu.</p>
SHOW	<p>Displays the command name, such as COPY or DIR, in the upper left corner of the screen.</p> <p>For example, if you enter</p> <p>ITEM Copy Files {SHOW}</p> <p>the DOS command COPY is displayed when the item is executed.</p>

Using ITEM

List items appear on the menu in the order in which they appear in the source file. They are not displayed in alphabetical order.

Each item is automatically assigned an alphabetic selection character.

If you want to assign a different character, place a carat (^) and the character you want in front of the item name (no spaces).



Forcing the selection character shortens the maximum line length to 38 characters.

If you assign selection letters, you should assign them to all menu list items. Otherwise another item may be automatically assigned a character you have previously assigned.

Example

For example, if the first menu item is "Word Processing" and you want to assign it the letter "W" instead of the automatic letter designation of "A," type

ITEM ^WWord Processing

NMENU Control Commands and Options

Control commands tell NMENU how to perform an action, such as displaying a submenu or user prompt, performing a DOS function, or starting an application. You also use these commands to tell the menu program how to process information and execute commands.

Table 4-3 gives an overview of the control commands. Detailed information about each command follows the table.

Table 4-3
NMENU Control Commands

Command	Explanation
EXEC	Executes a DOS or NetWare command following an ITEM statement.
SHOW	Displays a submenu from the same .DAT file.
LOAD	Calls and displays a menu as a submenu from a different .DAT file than the one you are running.
GETO	Requests information from the user before a menu item is executed. User input is <i>optional</i> . (See “GETx (GETO, GETP, GETR)” on page 269.)
GETP	Requests information from the user before a menu item is executed. User input is <i>required</i> for the program to proceed. Assigns a variable (%n) to the information so it can be used again. (See “GETx (GETO, GETP, GETR)” on page 269.)
GETR	Requests information from the user before a menu item is executed. User input is <i>required</i> for the program to proceed. (See “GETx (GETO, GETP, GETR)” on page 269.)

EXEC

Instructs NMENU to perform the command that follows EXEC.

Command Format

EXEC *command*

Replace *command* with the command required to execute the ITEM. This could be the name of an executable file, a DOS or NetWare command, or one of the options associated with EXEC. The EXEC options are described in Table 4-4.

Table 4-4
EXEC Command Options

EXEC Option	Explanation
EXEC CALL	Runs a batch file, and returns to NMENU. If you want to return to NMENU after a batch file executes, use this command to call batch files.
EXEC DOS	Runs the DOS command processor. If this command is used, the menu user must type "EXIT" when DOS processing is completed in order to return to NMENU.
EXEC EXIT	Exits the user from NMENU, but leaves the user logged in to NetWare. For security reasons, users cannot access the NetWare prompt unless this command is included in the menu.
EXEC LOGOUT	Exits the user from NMENU and logs the user out of the network, leaving the user at the DOS prompt. (See "Setting Up the User Environment" on page 275.)

Using EXEC

EXEC must follow the ITEM it applies to. It must also follow other control commands needed by ITEM.

Example

To load Word 5.0 from a menu, include the two following lines in the menu file:

```
ITEM Word 5.0  
EXEC word5
```

LOAD

Instructs NMENU to execute a separate menu file.

Command Format

```
LOAD filename
```

Replace *filename* with the name of another NMENU program.

If you have several menu programs created with the .DAT extension, use this command to call another menu program from the active menu program. LOAD always calls menus by their filename, not number.

Using LOAD

Use this command when you have multiple menu programs defined. Although you can define up to 255 menus per source file, smaller, separate files are easier to manage and update.

The NMENU file being loaded must be in the current directory, or you must have a search path to the file defined.

Example

If you are writing a menu program for Sales and you want to call up the Accounting menu program, type

```
MENU 1,Sales Main Menu  
ITEM Accounting Menu  
LOAD ACCOUNT
```

SHOW

Instructs NMENU to execute a submenu defined within the same file.

Command Format

SHOW *menu_number*

Replace *menu_number* with the number of the submenu to be displayed.

Using SHOW

Include SHOW commands to identify the submenu to be displayed when an item is selected from a menu.

SHOW commands always use menu numbers, not titles.

Example

If you have a menu with items listing categories of applications such as "Word Processing" and "Spreadsheets," each ITEM has the following SHOW commands to call up the appropriate submenus for those categories:

ITEM Word Processing

SHOW 3

ITEM Spreadsheets

SHOW 5

To continue this example, if menu "3,Word Processing" is a list of available word processing programs needed by both Sales and Accounting, and each group has its own menus defined, the programs can "share" that screen and the screen's calls to execute the applications.

GETx (GETO, GETP, GETR)

The GETx commands provide access to user input. You can request or require user input. You can even store user input for future use. Each variation of the GETx command uses the same parameters to control what the users sees and does. Pay close attention to the use of spaces, commas, and braces { }.

Command Format

GETx *prompt* {*prepend*} *length*, [*prefill*], {*append*}

Replace *x* with the letter "O," "R," or "P."

- ◆ If you want the user to decide whether to enter information, use GETO.
- ◆ If you want to require the user to enter information (such as a password), use GETR.
- ◆ If you want user input to be used by more than one EXEC command after it is entered, use GETP. The information is assigned a variable (such as %1, %2, etc.).

Replace *prompt* with the text (instructions) to be displayed to the user (40 characters maximum). For example: "Enter your Password."

Prepend is a value added to the beginning of the user response. A space is usually required in the first GETx command, to separate it from the command issued in its associated EXEC command.

Replace *length* with the maximum number of characters the user can enter. This parameter is required. For example, if you use GETR for a phone number, limit the *length* to 11 characters to match the field length for the phone number. Or, if you want a state or country abbreviation entered, limit the *length* to two characters.

Replace *prefill* with a default response displayed with the *prompt*. The user can accept the default, change the response by typing over the default, or cancel selection of the item. *Prefill* cannot be longer than the specified *length*.

The *prefill* parameter is optional. If you do not want to include a default response, enter the two commas together. (See the GETx examples on the following page.)

Append is the value added to the end of the user response. If no value is needed, enter braces with a blank space between them.



Append doesn't work with GETP.

Using GETx Interactive Commands

Following are some general guidelines on the use of the GETx commands.

- ◆ You must define commands for each menu ITEM separately.
- ◆ Enter commands between the ITEM and the EXEC command(s) associated with the ITEM.
- ◆ Enter the commands in either uppercase or lowercase.
- ◆ Limit each prompt to one line.

GETx Command Examples

- ◆ Assume that you want to require users to enter a project code to keep track of work performed for that project. Type the following:

```
GETR Please enter the project code: { } 08,, { }  
GETR Please enter your password: { } 08,, { }  
GETO Load default macros? { } 01,, { }
```

- ◆ To get input for a program that calculates mortgages, using GETR, type

```
GETR Enter the loan amount: { } 7,,{ }  
GETR Enter the interest rate: { } 5,8.5, { }  
GETO Enter the period (/m=months or /y=years):+  
  { } 7,/y=30, { }  
EXEC mortgage
```

The iFnpnt values will be appended to the "EXEC mortgage" command before it is executed.

- ◆ To create the previous program using GETP in place of other GET commands, type

```
GETP Enter the loan amount: { } 7,,{ }  
GETP Enter the interest rate: { } 5,8.5, { }  
GETP Optional-Enter the period (/m=months or+  
  /y=years): { } 7,/y=30, { }  
EXEC echo %1  
EXEC echo %2  
EXEC echo %3  
EXEC mortgage %1 %2 %3  
EXEC pause
```

Creating a Menu File

Prerequisites



Checklist

- A workstation running DOS 3.30 or later
- The Read, File Scan, Write, Create, and Erase file system rights to the directory where you will create and edit the menu
- A DOS text editor

Example Menu Programs

Study the following examples by creating and trying them. By entering, executing, and debugging each example, you will gain the depth of understanding necessary to use the scripting language effectively.



Note

In the following examples, indents are used to help readability; they are not required.

The examples include those script files used for the menus in Figure 4-1, Figure 4-3, Figure 4-5, and Figure 4-6. Subsequent examples show other ways of using the NMENU scripting language.

Example 1: Single-Window Menu

This example in the following figure is the text used to create the menu in Figure 4-1 on page 253.

Figure 4-7
Single-Window
Menu Script

```
Menu 1,Title Bar
  Item One {pause}
    Exec dir
  Item Two {show pause}
    Exec ver
  Item ^XExit
    Exec EXIT
```

Example 2: A Simple Menu

Text from the .SRC file used for the menu in Figure 4-5 on page 257 is shown in the following figure. This menu demonstrates how submenus can be used.

Figure 4-8
A Simple Menu
Script

```
Menu 5, Available Applications
  Item Word Processors { }
    Show 10
  Item SpreadSheets { }
    Show 15
  Item ^XExit Menu { }
    Exec EXIT
Menu 10, Available Word Processors
  Item WordPerfect 5.1
    Exec wp51
  Item WordPerfect 6
    Exec wp6
  Item MSWord
    Exec Word
Menu 15, Available SpreadSheets
  Item Quattro Pro
    Exec Q
  Item Lotus 123
    Exec 123
```

Example 3: A More Complex Menu

The menu in Figure 4-6 on page 258 was created from the text file in the following figure. This menu incorporates the GETR command with simple DOS functionality.

Figure 4-9
A More Complex
Menu Script

```
menu 22, Data Entry
  item ^AAdd Entry to Catalogue {show}
    getr Enter Title { } 40,, {, }
    getr Enter Author's Name { } 40,, {, }
    getr Category: 1=Apps, 2=Prod, 3=Self {} 1,, {,}
    exec echo >>datalist.cat
  item ^View Catalogue {pause}
    exec sort <datalist.cat | more
  item ^Instructions {pause}
    exec type instruct.cat | more
  item ^XExit the menu
    exec EXIT
```

As this menu script shows, by combining regular DOS commands with special NMENU conventions you have a simple means of organizing data. You don't need a database program just to catalog some information.

Example 4: Combining DOS and NetWare Commands

Users can accomplish more work when access to DOS and NetWare functionality is combined in the same menu.

Figure 4-10
Combining
Commands in a
Script

```
Menu 01,User Options
  Item Utilities
    Show 10
  Item DOS Prompt
    Exec DOS
  Item Log out of the network
    Exec Logout
Menu 10,Utilities
  Item ^1NetWare Menu Utilities
    Show 12
  Item ^2NetWare Command Line Utilities
    Show 14
Menu 12,NetWare Menu Utilities
  Item NETADMIN {Batch}
    Exec netadmin
  Item FILER {Batch}
    Exec filer
  Item NETUSER {Batch}
    Exec netuser
Menu 14,NetWare Command Line Utilities
  Item NLIST {show}
    Geto Class & Option: { } 25,user /a, {}
    Exec nlist
  Item COPY Files {pause}
    Getp Enter Source { } 25,, {}
    Getp Enter Destination { } 25,, {}
    Exec ncopy %1 %2
    Exec dir %2 /w
  Item Display a MAP listing {show pause}
    Exec map
```


Making Menus Work

Now that you have working menus, you must make them available to your users. Following is a list of rules for usage.

- ◆ You must create a search drive to the directory where the menu files exist. Search mappings can be created in container, profile, or user login scripts.
- ◆ Users must have at least Read and File Scan rights to the directory where the menu files exist.
- ◆ NMENU uses temporary files. You must create a permanent drive mapping assignment to a directory where the temporary files will be used.
- ◆ Users must have at least Read, File Scan, Write, Create, and Erase rights to the directory where the temporary files will be stored.

You can simplify management of the menu system by keeping all menu files in one place, such as SYS:MENUS. It becomes even easier to manage if the temporary files are kept in one place, such as SYS:MENUS\TEMP.

Setting Up the User Environment

If you are storing users' temporary files in a network directory, place the following commands in the login script:

```
SET S_FILEDIR=  
SET S_FILE=.
```

For example, if you have created a subdirectory called TEMP under a MENUS directory, you would type

```
SET S_FILEDIR="Z:\\MENUS\\TEMP\\"  
SET S_FILE="%STATION"
```

These commands point to the directory where temporary files are stored and create unique files in the temporary directory for each workstation ID number. The trailing backslash on the S_FILEDIR path is required.



If you choose not to use the %STATION identifier variable, you must manually create a file with a maximum length of seven characters. The NMENU program automatically prepends the # symbol to the beginning of the S_FILE filename.

If the user will be using a menu with the logout option, set the S_FILEDIR environment variable to a path on the user's local drive. Set the S_FILE to "%STATION."

A copy of the MENU-X.BAT file needs to be in the SYS:LOGIN directory if the NMENU LOGOUT option is going to be used.

EXEC LOGOUT gives the message "Batch file not found" unless the temporary directory is on a local drive, and MENU_X.BAT is in the LOGIN directory.

Starting NMENU from a Login Script


If you want the menu to execute from within a login script, add the EXIT login script command. For example, to execute the ACCOUNT menu from a user's login script, add the following line to the login script:

```
EXIT "NMENU ACCOUNT"
```

Converting Old Menu Files

Use this procedure to convert an existing MENU file (created with an earlier version of NetWare) to an NMENU file. These older menu programs usually have the extension .MNU.

Procedure



1. Create a directory for the new menu files and temporary files.

If you want to use the directories you already have set up for menu programs, you do not need to create new ones.

2. Change the old .MNU file to the new .SRC source file format with the MENU CNVT command. Complete the following steps.

2a. At the workstation command line, type

```
MENUCNVT menu_name.mnu <Enter>
```

Replace *menu_name* with the name of the old menu you want to convert. If the menu is not in your current directory, enter the complete path to the menu name.

The program creates a new .SRC file, leaving the original .MNU file unchanged.

2b. Edit the new menu source file.

You will probably need to edit the new file to eliminate incorrect commands. For example, if your original MENU file used the SYSCON command, you must remove it from your new NMENU file and replace it with an appropriate NetWare 4 command (such as NETADMIN or NETUSER). NetWare 4 does not use SYSCON.

There can be substantial differences between the original .MNU file and the converted .SRC file. MENU CNVT takes care of most differences, but not all of them. Some typical entries requiring changes by hand include the following:

- ◆ Invalid commands such as SYSCON
- ◆ Preceding letters or numbers from the old menu
- ◆ Conversion of former @1, @2, etc., variables into the newer GETP variables

3. Follow the same six steps discussed in “Menu Creation Steps” on page 256.

Example

To help you see the differences between MENU and NMENU files, Figure 4-11 shows a part of the file for the menu in Figure 4-3 on page 255.

The figure shows both the .MNU and the unedited .SRC file formats. They are shown side-by-side to help you make the comparison.

The “After Conversion (.SRC)” side has not yet been edited. Try to find changes to make the file execute properly and efficiently. Look for commands that are no longer supported in NetWare 4 and replace them with appropriate commands. One example is SLIST, which must be replaced with NLIST SERVER.

If you create your own file to test the .SRC file, remember to create some text file(s) to be displayed by the TYPE command.

**Figure 4-11
Conversion Scripts**

Before Conversion (.MNU)

```
%Questions Questions Questions!,12,20,2
1. Who Am I?
   whoami
   pause
2. Where Am I?
   type where.txt
   pause
3. Am I Alone?
   userlist
   pause
4. What Am I Doing Here?
   type what.txt
   pause
5. What Can I Do?
   %Your Choice
6. What Next?
   %What Next
```

After Conversion (.SRC)

```
MENU 01,Questions Questions Questions!
ITEM Who Am I? { }
EXEC whoami
EXEC pause
ITEM Where Am I? { }
EXEC type where.txt
EXEC pause
ITEM Am I Alone? { }
EXEC userlist
EXEC pause
ITEM What Am I Doing Here? { }
EXEC type what.txt
EXEC pause
ITEM What Can I Do? { }
SHOW 02
ITEM What Next? { }
SHOW 03
```

(continued)

Before Conversion (.MNU)

%Your Choice,5,50,4

- A. Quit
 leave.bat
- B. Try Again
 %Utilities
- C. Experiment
 %Demo

%Demo,12,40,1

- 1. CHOICE 1
 %Utility
- 2. CHOICE 2
 %Variety
- 3. CHOICE 3
 %TalkToMe
- 4. CHOICE 4
 %Beyond

%Beyond,20,75,5

MOVE ABOUT
 %Pop_up_1
THE OUTSIDE WORLD
 slist
 pause

%Pop_up_1,24,10,6

MOVE ABOUT
 %Pop_up_2

%Pop_up_2,19,30,1

MOVE ABOUT
 %Pop_up_3

%Pop_up_3,1,1,2

MOVE ABOUT
 %Pop_up_4

%Pop_up_4,8,60,3

MOVE ABOUT
 %Pop_up_5

%Pop_up_5,12,4,4

MOVE ABOUT
 %Pop_up_6

%Pop_up_6,12,78,5

MOVE ABOUT
 type theend.txt
 pause

After Conversion (.SRC)

MENU 02,Your Choice

- ITEM Quit { }
 EXEC leave.bat
- ITEM Try Again { }
 SHOW 06
- ITEM Experiment { }
 SHOW 07

MENU 07,Demo

- ITEM CHOICE 1 { }
 SHOW 08
- ITEM CHOICE 2 { }
 SHOW 09
- ITEM CHOICE 3 { }
 SHOW 10
- ITEM CHOICE 4 { }
 SHOW 11

MENU 11,BEYOND

- ITEM MOVE ABOUT { }
 SHOW 12
- ITEM THE OUTSIDE WORLD { }
 EXEC SLIST
 EXEC PAUSE

MENU 12,Pop_up_1

- ITEM MOVE ABOUT { }
 SHOW 13

MENU 13,Pop_up_2

- ITEM MOVE ABOUT { }
 SHOW 14

MENU 14,Pop_up_3

- ITEM MOVE ABOUT { }
 SHOW 15

MENU 15,Pop_up_4

- ITEM MOVE ABOUT { }
 SHOW 16

MENU 16,Pop_up_5

- ITEM MOVE ABOUT { }
 SHOW 13

MENU 12,Pop_up_6

- ITEM MOVE ABOUT { }
 EXEC TYPE THEEND.TXT
 EXEC PAUSE



chapter

5

Managing the NetWare Directory Tree

About NetWare Directory Services

NetWare® Directory Services™ (NDS) is a distributed name service that provides global access to all network resources regardless of where they are physically located. Users log into a multiserver network and view the entire network as a *single information system*.

This single information system is the basis for increased productivity and reduced administrative costs. You can manage NDS™ by using the utilities and programs described in this chapter.

The following utilities and programs help you maintain the NDS database.

- ◆ **Partition Manager** utilities allow you to distribute your NDS database by placing copies (replicas) of parts (partitions) of the database on different servers.

A text utility (PARTMGR) and a graphical utility (NetWare Administrator) help you manage the database partitions and replicas. You can also remove NDS from a server using the INSTALL NetWare Loadable Module (NLM) program.

- ◆ **UIMPORT** allows you to import information from an existing database (any database capable of converting records to an ASCII file will work) to the NDS database. Use this utility and your database records to create, delete, and update User objects and their properties.
- ◆ **DSREPAIR** checks and repairs local portions of the database, similar to the way VREPAIR fixes volumes on a server.

- ◆ **DSMERGE** allows you to merge one Directory tree into another at the [Root], resulting in a single tree. This merging process would be used to create a single tree structure when previously separate organizations are joined together.
- ◆ **DSTRACE** allows you to determine whether NDS synchronization processes are complete and diagnose NDS errors.

Creating and Managing Directory Services Partitions

About Partitions and Replicas

A partition is a part of the total Directory and contains one or more containers and its associated leaf objects.

When a partition is subordinate to another in the Directory tree, it is referred to as a *child partition*. The partition above it is referred to as the *parent partition*.

You can make copies of a partition, called *replicas*, and store them on different servers in your network. Distributing replicas reduces network traffic by making information accessible locally and enabling users to utilize the network services even when a particular server is down.

Replicas also provide fault tolerance by ensuring that more than one copy of the partition information is available. For example, if a replica of the partition becomes unavailable, you can use a different replica to re-create it.



Partitions contain only NDS database information, *not* file and directory information.

Planning and Using Partitions and Replicas

What Happens During Installation?

By default, the installation utility adds a replica of the partition that contains the server's context only if the total of existing replicas is less than three.

However, if the server is not a NetWare 4 server and contains bindery files (SYS:SYSTEM\NET\$.SYS) a replica will be added, regardless of the number of replicas.

These default settings ensure that bindery services will work correctly for networks running both NetWare 3™ and 4 software.

For more information about bindery services, see Chapter 3, "Understanding Bindery Services," of *Introduction to NetWare Directory Services* and "Bindery services" in *Concepts*.

You can create additional partitions and replicas using the NetWare Administrator graphical utility or the PARTMGR text utility. Instructions for using these utilities are included in this section.

For more information on Directory partitions, see Chapter 2, "Understanding Management Features" of *Introduction to NetWare Directory Services* and "NetWare Directory partition" in *Concepts*.

Guidelines for Managing Partitions and Replicas

- ◆ Make sure all the servers that contain partition replicas are up and running before you attempt any partition operation.

If you attempt a partition operation while a server (that contains a replica) is down, NetWare Directory Services (NDS) won't be able to synchronize, since it won't be able to communicate with the replica on the downed server.

- ◆ You can create partitions only at the container level of the Directory tree.
- ◆ NDS does not require you to have a replica of the complete database in any location; the database is designed to be divided into usable pieces.

- ◆ Operations to manage partitions or replicas take place in the background and take time to complete, since the replicas need to be synchronized with new information.
- ◆ You cannot store more than one replica of the same partition on a server. It is not necessary for all servers in your tree to store replicas, unless you want to be able to log in to every server in your tree using bindery services.
- ◆ You can create as many replicas of a partition as you need; however, we recommend that you create between three and six replicas of each partition in your Directory tree.

Having too many replicas may slow down your network because it will be busy synchronizing.

- ◆ You can have only one master replica. Replica types are explained in the following table.

Replica	Description
Master	<p>Partition information can be read from and written to this replica. The master replica can also be used to change the logical structure of the Directory (by creating a new partition, for example).</p> <p>When you create a new partition, a master replica is created and stored on the same server as the parent partition.</p>
Read/write	<p>Partition information can be read from and written to this replica. For example, it can be used to log in, add or delete objects, and view directory information.</p>
Read-only	<p>Partition information can be read from this replica, but it cannot be written to by anything other than a read/write or master replica.</p> <p>Users cannot authenticate the network through a read-only replica because the login process changes the NDS database.</p> <p>You cannot set a bindery context when you log in to a read-only replica.</p>

Replica	Description
Subordinate	<p>This replica cannot be modified by any user. It is automatically placed on a server by NDS if the parent partition has either a master, read/write, or read-only replica on the server and the child partition does not.</p> <p>If you add a read/write or read-only replica of the child partition to the server, the subordinate replica is removed.</p>

Because partition and replica management is such a vital part of managing your network, read Chapter 7, “Managing NetWare Directory Services,” of *Introduction to NetWare Directory Services* before you use the Partition Manager utilities.

For more information about	Refer to
Bindery services	“Bindery services” in <i>Concepts</i>
Partition Manager utilities	“PARTMGR” and “NetWare Administrator” in <i>Utilities Reference</i>
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Planning and using partitions	Chapter 5, “Planning NetWare Directory Services Implementation,” of <i>Introduction to NetWare Directory Services</i>

Restricting Access to Partition Manager

As administrator, you may want to limit which users in your network can access the Partition Manager that can be launched from NetWare Administrator.

There are two ways you can limit access to Partition Manager:

- ◆ Restrict rights to NWPART.DLL (found in SYS:PUBLIC) by giving Read and File Scan rights only to those who you want to have access to Partition Manager.
- ◆ Remove NWPART.DLL from SYS:PUBLIC and place it in your path or in a directory to where you have a search drive mapped.

This allows only you to see Partition Manager under the “Tools” menu in NetWare Administrator.



You can also limit rights to PARTMGR (in DOS) by restricting rights to PARTMGR.EXE or by removing it from SYS:PUBLIC and placing it in your path or in a directory to where you have a search drive mapped.

Creating a New Partition

A partition can consist of one or more container objects and their associated leaf objects. A partition cannot contain only leaf objects. The container that is the first object in the partition is called the “root” of the partition.

When you create a new partition, you split the parent partition and end up with two partitions. The new partition becomes a child partition.

For example, if you select an Organizational Unit and choose to create it as a new partition, you are choosing to split the Organizational Unit from its parent partition ([Root], for example, which is always a partition). The Organizational Unit you selected becomes the partition root of a new partition.

The replicas of the parent partition will remain on the same servers, and information for the new partition will migrate from the parent partition’s replicas to the new partition’s replicas.

The master replica of the new partition will be stored on the same server as the master replica of the parent partition.

Creating a partition may take some time, since all of the replicas need to be synchronized with the new partition information.

When you create a new partition, the utility you use will inform you that the partition is created successfully, but the actual creating is still completing on the servers. You will need to wait a while before being able to perform another partition operation.

Partitions can be created using the PARTMGR text utility or the NetWare Administrator graphical utility. Both procedures are documented in this section.

Creating a Partition Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the container object you are partitioning

Procedure



- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the “Tools” menu, choose “Partition Manager.”**

Partition Manager will launch from the container you selected in NetWare Administrator.

- 3. To locate the container object that you want to create as a partition, browse the “Partition Manager” screen.**

If the container object does not appear in the window, browse the tree by choosing an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. **Select a container object.**
5. **Choose “Create as New Partition.”**
6. **To confirm the creation of a new partition, choose “Yes.”**

The new partition is created when the process is completed on the servers. To see the partition icon that signifies that a container is a partition, you need to refresh the screen by choosing the up arrow and then expanding the parent container again.

A master replica is stored on the server where the parent partition’s master replica resides. An icon appears next to the Organization (O) or Organizational Unit (OU) to show that the container is the root of a partition.

To see where the master replica is stored, select the Organization (O) or Organizational Unit (OU) that you just partitioned, and then choose “Replicas.” The server name appears in the “Servers” column and “Master” appears next to it in the “Type” column.

If you want to make additional replicas of this partition, see “Creating a Replica” on page 302.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Creating a Partition Using PARTMGR

Prerequisites



- A workstation logged in to the network, running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the container object you are partitioning

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner of the screen.

3. Select the container object to partition.

- ◆ If the container you want to partition appears on the list, select it and press <F10>.
- ◆ If the container is not on the list, browse the directory by selecting containers and pressing <Enter> until you see the container you want. Select it and press <F10>.

4. To create the new partition, choose “Yes.”

The new partition is created when the process is completed on the servers. To see the partition icon that signifies that a container is a partition, you need to refresh the screen by choosing the up arrow and then expanding the parent container again.

A master replica is stored on the server where the parent partition’s master replica resides. An icon appears next to the Organization (O) or Organizational Unit (OU) to show that the container is the root of a partition.

To see where the master replica is stored, select the Organization (O) or Organizational Unit (OU) that you just partitioned, and then choose "View/Edit Replicas." The server name appears in the "Replicas Stored on Server" column, and "Master" appears next to it in the "Type" column.

If you want to make additional replicas of this partition, see "Creating a Replica" on page 302.

For more information about	Refer to
Partitions	"NetWare Directory partition," "Partition management," and "Replica" in <i>Concepts</i>
Using the PARTMGR utility	"PARTMGR" in <i>Utilities Reference</i>

Merging Partitions

You might want to merge two partitions if the Directory information in the two partitions is closely related.

In other words, you would merge a partition with its parent partition when you want to delete the partition without deleting the objects in the partition.

Consider keeping partitions separate if the partitions are large, because large partitions slow down response time.

The partition is merged when the process is completed on the servers. To see that the icon of the partition you merged is gone (which signifies that the merge is complete), you need to refresh the screen by choosing the up arrow and then expanding the container again.

Merging a partition with its parent partition might take some time, since the replicas need to be deleted and the parent replicas updated with the merging partition information.

You can merge a subordinate partition with its parent partition using either the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.


Merging Partitions Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object and its parent partition

Procedure



- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the “Tools” menu, choose “Partition Manager.”**
- 3. To locate the container object that you want to merge with the parent partition, browse the “Partition Manager” screen.**

If the partition does not appear in the window, browse the tree by selecting an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

- 4. Select the partition you want to merge.**
- 5. Choose “Merge Partition.”**
- 6. To merge the partition with its parent, choose “Yes.”**

The partition is merged when the process is completed on the servers. To see that the icon of the partition you merged is gone (which signifies that the merge is complete), you need to refresh the screen by choosing the up arrow and then expanding the container again.

Merging a partition with its parent partition might take some time, since the replicas need to be deleted and the parent replicas updated with the merging partition information.

The replicas of the partitions are also merged.

A partition with no parent partition will merge with the [Root] partition.

For more information about	Refer to
Partitions	"NetWare Directory partition," "Partition management," and "Replica" in <i>Concepts</i>
Using the NetWare Administrator utilities	"NetWare Administrator" in <i>Utilities Reference</i>


Merging Partitions Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the object at the root of the parent partition

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the "Partition Administration" menu, choose "Manage Partitions."

Your current context appears in the upper-left corner of the screen.

3. Select the partition to merge.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the directory by selecting objects and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose "Merge with the Parent Partition."

5. To merge the selected partition with its parent partition, choose “Yes.”

The partition is merged when the process is completed on the servers. To see that the icon of the partition you merged is gone (which signifies that the merge is complete), you need to refresh the screen by choosing the up arrow and then expanding the container again.

Merging a partition with its parent partition might take some time, since the replicas need to be deleted and the parent replicas updated with the merging partition information.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Moving a Partition

You can move a container object only if it is the root of a Directory partition that has no subordinate partitions. So, moving a container is really moving a partition.

When you move a container object, NetWare Directory Services (NDS) changes all references to the container. Although the object’s common name remains unchanged, the context name of the container (and of all its subordinates) changes.

When you move a partition, create an alias object that points to the partition you’re moving. Doing so allows users to continue logging into the network and finding objects in their original Directory location.



Note

If you move a partition and do not create an alias, users who are unaware of the partition's new location will not easily find objects in the Directory tree, since they will look for them in their original Directory location.

This might also cause client workstations to fail at login if the NAME CONTEXT parameter in the NET.CFG file is set to the original location in the Directory tree.

Because the context of an object changes when you move it, users whose name context in their configuration file (NET.CFG file) references the moved object need to update their NET.CFG so that it references the object's new name.

To automatically update users' NET.CFG file with a new name context after moving an object, use the NCUPDATE utility. For instructions, see "NCUPDATE" in *Utilities Reference*.

You can use NetWare Administrator or NETADMIN to move partitions in the Directory tree. Both procedures are documented in this section.

Moving a Partition Using NetWare Administrator

Prerequisites



Checklist

- A workstation running MS Windows 3.1 and NetWare Administrator
- The Supervisor right to the object you want to move
- The Create object right to the destination container

Procedure



Procedure

1. **From the MS Windows Program Manager, choose the "NetWare Administrator" icon.**
2. **From the "Tools" menu, choose "Partition Manager."**
3. **From the "Partition Manager" browser, select the partition that you want to move.**



Note

You can move a container object only if it is the root of a partition, and only if it contains no subordinate partitions.

In Partition Manager, the partition icon appears to the left of the object icon. If the container you want to move is not a partition, select the container and choose "Create as New Partition." Then refresh the screen so that the partition icon appears and go to the next step.

If there are subordinate partitions in the container you want to move, you can merge those partitions with their parent partition.

4. **From the "Object" menu, choose "Move Partition."**
5. **Select the browser button to the right of the "Destination" box.**

Use the browser in the "Directory Context" box to view the Directory tree's containers.

The "Objects" box that appears in the lower-left corner shows the containers that you select in the "Directory Context" box.

6. **From the "Objects" box, select a container object (an Organization or Organizational Unit) as the place to move the listed objects to; then choose "OK."**
7. **Choose "Create Alias in Place of Moved Container."**

The alias object will point to the partition's new location.



If you move an object do not create an alias, users who are unaware of the object's new location will not easily find objects in the Directory tree, since they will look for them in their original Directory location.

This might also cause client workstations to fail at login if the NAME CONTEXT parameter in the NET.CFG file is set to the original location in the Directory tree.

Because the context of an object changes when you move it, users whose name context in their configuration file (NET.CFG file) references the moved object need to update their NET.CFG so that it references the object's new name.

To automatically update users' NET.CFG file with a new name context after moving an object, use the NCUPDATE utility. For instructions, see "NCUPDATE" in *Utilities Reference*.

8. In the “Move” dialog box, choose “OK.”

If you chose to create an alias in place of the moved container, NetWare Administrator polls for the creation of the alias object before it moves the selected partition.

Unless you want the partition you just moved to remain a partition, you should merge it with its parent partition to avoid having an unnecessary partition in the Directory tree. See “Merging Partitions” on page 290.



You will need to wait for processes throughout the Directory to be complete before you can perform a partition operation with this object again.

For more information about	Refer to
Merging partitions	“Merging Partitions” on page 290
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Alias objects	“Alias object” in <i>Concepts</i>
Directory tree	“Directory tree” in <i>Concepts</i>
NCUPDATE	“NCUPDATE” in <i>Utilities Reference</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>


Moving a Partition Using NETADMIN

Prerequisites



- A DOS workstation running DOS 3.30 or later and the NETADMIN utility
- The Create object right to the destination container

Procedure



1. At the DOS prompt, type

NETADMIN <Enter>

For information on moving around in NETADMIN and selecting objects, press <F1> after starting the utility. To see which container objects in the Directory tree are partitions, exit NETADMIN and type "PARTMGR" at the command line. Then browse the tree.

2. From the NETADMIN options menu, choose "Manage Objects."

Your current context appears in the upper-left corner.

3. Select the object that you want to move.

- ◆ If the object you want to move appears on the list, select it and press <F10>.
- ◆ If the object is not on the list, browse the directory by selecting container objects and pressing <Enter> until you see the object you want. Select it and press <F10>.

4. From the "Actions" menu, choose "Move."

You can move a container object only if it is the root of a Directory partition, and only if it contains no subordinate partitions.

In NETADMIN, when you select a container object that is a partition, the context-sensitive help at the bottom of the screen reads "This is a partition." Also, an asterisk (*) is displayed in front of the object name.

If the container you want to move is not a partition, you must first use a partition management utility (PARTMGR or NetWare Administrator) and create the container as a new partition.

5. **Use the Down-arrow key and highlight the “New Context” field.**
6. **Assign a new context to the object you want to move.**
 - ◆ If you know the new context that you want the object to be in, type the new context in the highlighted field.
 - ◆ If you don’t know the new context that you want the object to be in, press <Insert> twice to browse the Directory for the destination container; then select the destination container and press <F10>.
7. **To accept the new context as the destination container, press <Enter>.**
8. **To confirm that you want to move the object listed in the “Old Context” field to the container listed in the “New Context” field, press <F10>.**
9. **To create an alias in place of the moved container, choose “Yes.”**

The alias object will point to the partition’s new location, and the selected object is moved to the destination container.



If you move an object do not create an alias, users who are unaware of the object’s new location will not easily find objects in the Directory tree, since they will look for them in their original Directory location.

This might also cause client workstations to fail at login if the NAME CONTEXT parameter in the NET.CFG file is set to the original location in the Directory tree.

Because the context of an object changes when you move it, users whose name context in their configuration file (NET.CFG file) references the moved object need to update their NET.CFG so that it references the object’s new name.

To automatically update users’ NET.CFG file with a new name context after moving an object, use the NCUPDATE utility. For instructions, see “NCUPDATE” in *Utilities Reference*.



You will need to wait for processes throughout the Directory to be complete before you can perform a partition operation with this object again.

Unless you want the partition you just moved to remain a partition, you should merge it with its parent partition to avoid having an unnecessary partition in the Directory tree. See “Merging Partitions” on page 290.

For more information about	Refer to
Objects	“Object” in <i>Concepts</i>
Rights	“Rights” in <i>Concepts</i>
Directory tree	“Directory tree” in <i>Concepts</i>
Using the NETADMIN utility	“NETADMIN” in <i>Utilities Reference</i>

Aborting a Partition Operation

If you have begun the process of creating, merging, or moving a partition, or changing a replica type, you can still abort the process since partition operations take time. You can abort a partition operation only before the operation is in its final stages.

You would want to use this feature if you begin a partition operation and find that your database will not synchronize.

If NetWare Directory Services (NDS) cannot synchronize replica information in your database because it is corrupted, or because a server in your Directory tree is down, you probably want to abort any partition operation in progress.

If you choose to abort a partition operation when there is no operation in progress, no partitions will be affected.

You can abort a partition operation using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.

Aborting a Partition Operation Using NetWare Administrator

Prerequisites

Checklist



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition

Procedure

Procedure



1. **From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
2. **From the “Tools” menu, choose “Partition Manager.”**
3. **To locate the partition that is executing a partition operation, browse the “Partition Manager” screen.**

If the container object does not appear in the window, browse the tree by selecting an object to see its subordinates, or by selecting the arrow key to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. **Select the partition that is executing a partition operation and choose “Abort Partition Operation.”**

As long as at least one of the replicas has not yet completed the operation (as shown in the “State” box), you can still abort the operation.

5. To abort the partition operation, choose “Abort” and then choose “Yes.”

The partition operation is aborted, and any replicas that have been merged, created, or moved (depending on the operation you had begun) are returned to their original Directory location.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Aborting a Partition Operation Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the parent partition that is involved in a partition operation.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the directory by selecting objects and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “Abort Partition Operation.”

A list of the replicas of the selected partition appears. Each replica’s type and state are also displayed.

As long as at least one of the replicas has not yet finished the operation you had begun, you can still abort the operation. If the state of the replica is “On,” all operations are complete and cannot be aborted.

5. To abort the partition operation, press <F10> and choose “Yes.”

The partition operation is aborted, and any replicas that began to be merged or created (depending on the operation you had begun) are returned to their previous state.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Creating a Replica

When you create a partition, a master replica is automatically created and stored on the server where the parent partition’s master replica resides. You can create additional replicas of the partition, within these guidelines:

- ◆ You can have only one master replica. Additional replicas must be read/write or read-only. For a description of replica types, see “Planning and Using Partitions and Replicas” on page 283.
- ◆ You can store only one replica of a partition on a server.

Replicas can be created using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.


Creating a Replica Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition

Procedure



1. **From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
2. **From the “Tools” menu, choose “Partition Manager.”**
3. **Browse the “Partition Manager” screen to locate the partition that you want to create a replica for.**

If the container object does not appear in the window, browse the tree by selecting an object to see its subordinates, or by selecting the arrow key to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. **Select the container object and choose “Replicas.”**
5. **Choose “Add Replica.”**
6. **Browse the objects in the “Server” box to locate the server you want to store the replica on.**

The partition for which you are creating the replica appears at the top of the screen. The context and server information changes as you move through the Directory tree structure.

7. **Select the server you want the new replica to be stored on.**
8. **Choose the replica type you want to create.**

9. To create the replica, choose “OK.”

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Creating a Replica Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the partition to replicate.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the tree by selecting objects and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “View/Edit Replicas.”

A list of the replicas of the selected partition appears. Each replica’s type is also displayed.

5. To add a replica to the server, press <Insert>.

6. Enter the replica information.

6a. At the “Replica Type” field, press <Enter>.

6b. Choose the type of replica you want to create from the “Replica Type” menu at the top of the screen.

6c. At the “Store on Server” field, press <Enter>.

6d. Type the name of the server you want to add the replica to, or press <Insert> to select a server from the browser.

6e. To create the replica, press <Esc> or <F10> and choose “Yes.”

The new replica appears on the list of replicas for the server.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Deleting a Replica

Replicas can be deleted using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.

When you delete replicas, keep the following guidelines in mind:

- ◆ If a master replica becomes corrupted (for example, if the server will be down for a long time or the master replica becomes corrupted), change a read/write or read-only replica on another server to the master replica. This automatically changes the old master replica to a read/write replica.
- ◆ For fault tolerance, you should maintain at least one replica of the master partition on different servers.

Deleting a Replica Using NetWare Administrator

Prerequisites

Checklist



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replica you want to delete

Procedure

Procedure



- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the “Tools” menu, choose “Partition Manager.”**
- 3. Browse the “Partition Manager” screen to locate the partition whose replica you want to delete.**

If the object where the container is stored does not appear in the window, browse the tree by selecting an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. **Select the partition and choose “Replicas.”**
5. **Select the server you want to delete the replica from and choose “Delete Replica.”**

You cannot delete a master replica. If the replica you want to delete is a master, go to a server with another replica of the master and make it the new master replica. This automatically changes the old master replica to a read/write replica, which you can then delete.

For instructions, see “Viewing a List of Partitions in a Directory Tree” on page 320.

6. **Choose “OK.”**

For more information about	Refer to
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>


Deleting a Replica Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replica you want to delete

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the partition that has a replica you want to delete.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the directory by selecting objects and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “View/Edit Replicas.”

A list of the replicas of the selected partition appears. Each replica’s type is also displayed.

5. Select the replica to delete and press <Delete>.

You cannot delete a master replica. If the replica you want to delete is a master, go to a server with another replica of the master and make it the new master replica. This automatically changes the old master replica to a read/write replica, which you can then delete.

For instructions, see “Viewing a List of Partitions in a Directory Tree” on page 320.

6. To delete the replica, choose “Yes.”

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Sending Updates to Other Replicas

Although NetWare Directory Services automatically synchronizes the directory data of replicas (so that each replica contains the most recently updated data), you can manually start the process to synchronize (update) the directory data of replicas, if necessary.

You should use the DSREPAIR utility to discover if the data in some replicas is out of sync with the master replica or not. If the data is out of sync, you would want to manually send updates to those replicas. The data in the other replicas would be updated and the directory data of each replica would be synchronized.

You can send updates to other replicas using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.

Sending Updates Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replicas you want to send updates to

Procedure



- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the “Tools” menu, choose “Partition Manager.”**
- 3. Browse the “Partition Manager” screen to locate the partition whose replicas you want to update.**

If the container object does not appear in the window, browse the tree by selecting an object to see its subordinates, or by selecting the arrow key to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. Select the partition and choose “Replicas.”

5. From the “Replicas” screen, select which replica you want to send updates from (the one that contains updated data) and choose “Send Updates to Other Replicas.”

The updated data is sent to all existing replicas of the same partition.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Sending Updates Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replicas you want to send updates to

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the partition whose replicas you want to update.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the tree by selecting objects and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “View/Edit Replicas.”

A list of the replicas of the selected partition appears. Each replica’s type is also displayed.

5. Select the replica you want to send updates from and press <F10>.

6. Choose “Send Updates to Other Replicas.”

7. To send updates to all the other replicas, choose “Yes.”

The replica information is sent to all other replicas of the partition (including the master replica).

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Receiving Updates from Other Replicas

Although NetWare Directory Services (NDS) automatically synchronizes the directory data of replicas (so that each replica contains the most recently updated data) you can manually synchronize (update) the directory data of replicas if they ever get out of sync.

You should use the DSREPAIR utility to discover if the data in some replicas is out of sync with the master replica or not. If the data is out of sync, you would want to manually receive updates from the Master replica.

You cannot choose “Receive Updates From Other Replicas” from a master replica. The master is assumed to be the most current and accurate copy of the partition. If it’s not, you should assign one of the other replicas to be the master using the PARTMGR utility/Replica.

If you choose “Receive Updates From Other Replica” from any replica, that replica will receive NDS information from the master.

You can receive updates from other replicas using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.

Receiving Updates Using NetWare Administrator


Prerequisites

Checklist



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replicas you want to send updates to

Procedure



1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.
2. From the “Tools” menu, choose “Partition Manager.”
3. Browse the “Partition Manager” screen to locate the partition whose replicas you want to update.

If the container object does not appear in the window, browse the tree by selecting an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. Select the partition and choose “Replicas.”
5. From the “Partition Replicas” screen, select which replica you want to update and choose “Receive Updates.”

The replica you have chosen receives NDS information from the master replica.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Receiving Updates Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replicas you want to update

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the partition whose replicas you want to update.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the tree by selecting objects and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “View/Edit Replicas.”

A list of the replicas of the selected partition appears. Each replica’s type is also displayed.

5. Select the replica you want to update and press <F10>.

6. Choose “Receive Updates from Other Replicas.”

7. To receive updates from all the other replicas, choose “Yes.”

The selected replica receives NDS information from the master replica.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Viewing a List of Partitions Stored on a NetWare Server

When you view a list of partitions stored on a NetWare server, you are seeing all the partitions that have a replica stored on the selected server and the type of each replica.

You can see a list of partitions stored on a NetWare Server object using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.


Listing Partitions Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Write property right to the ACL property of the NetWare Server object whose partitions you want to view

Procedure



1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.
2. From the “Tools” menu, choose “Partition Manager.”
3. Browse the “Partition Manager” to find the NetWare Server object whose partition list you want to view.

If the NetWare Server object does not appear in the window, browse the tree by choosing an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

4. Select the NetWare Server object and choose “Server Partitions.”

A list of the partitions that have a replica stored on the server is displayed.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Listing Partitions Using PARTMGR

Prerequisites

Checklist



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Write property right to the ACL property of the NetWare Server object whose partitions you want to view

Procedure

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the server whose partitions you want to list.

- ◆ If the server appears on the list, select it and press <F10>.
- ◆ If the server is not on the list, browse the tree by selecting containers and pressing <Enter> until you see the server you want. Select it and press <F10>.

A list of partitions stored on the server appears.

For more information about

Refer to

Partitions

“NetWare Directory partition,” “Partition management,” and “Replica” in *Concepts*

Using the PARTMGR utility

“PARTMGR” in *Utilities Reference*

Viewing a List of Replicas of a Partition

You can see a list of a partition's replicas, the servers where the replicas are stored, and whether a replica is a master, read/write, read-only, or subordinate replica type.

You can view a replica list using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.


Listing Replicas Using NetWare Administrator

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Write property right to the ACL property of the NetWare Server object whose replicas you want to view

Procedure



- 1. From the MS Windows Program Manager, choose the "NetWare Administrator" icon.**
- 2. From the "Tools" menu, choose "Partition Manager."**
- 3. Browse the "Partition Manager" screen to locate the partition whose replica list you want to view.**

If the partition does not appear in the window, browse the tree by choosing an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the "Partition Manager" screen and selecting objects, choose the "Help" button.

- 4. Select the partition and choose "Replicas."**

A list of replicas appears, including the server where each replica is stored, the replica type, and the time and date of the last time the replica was synchronized with the other replicas.

For more information about	Refer to
Partitions	"NetWare Directory partition," "Partition management," and "Replica" in <i>Concepts</i>
Using the NetWare Administrator utility	"NetWare Administrator" in <i>Utilities Reference</i>


Listing Replicas Using PARTMGR

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 512 KB of memory available on the workstation
- The Write property right to the ACL property of the Server object whose replicas you want to view

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the "Partition Administration" menu, choose "Manage Partitions."

Your current context appears in the upper-left corner.

3. Select the partition whose replicas you want to list.

- ◆ If the partition appears on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the tree by selecting containers and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “View/Edit Replicas.”

The “Replicas Stored on Server” screen appears, which lists the replicas of the selected partition, the server they reside on, and the type of each replica.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Viewing a List of Partitions in a Directory Tree


You can see a list of all partitions to which you have Browse rights in a NetWare Directory Services (NDS) tree using the NetWare Administrator graphical utility.

Prerequisites



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation

Procedure



1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.
2. From the “Object” menu, choose “Search.”
3. Choose the drop-down arrow button at the right of the “Search For” field.
4. Scroll the drop-down list until you find “Partition.”

5. Select “Partition” and choose “OK.”

A list of the partitions in the NDS tree appears. If you do not have the Browse right to the object at the root of the partition, the partition does not appear in the list.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Changing a Replica’s Type

You can change replica types according to the following guidelines:

- ◆ You can have only one master replica of a partition. Creating a new master replica automatically changes the old master replica to a read/write replica.
- ◆ You can change read/write replicas to read-only, and vice versa, without affecting other replicas of the same partition.

You can change a replica’s type using the NetWare Administrator graphical utility or the PARTMGR text utility. Both procedures are documented in this section.

Changing a Replica’s Type Using NetWare Administrator


Prerequisites

Checklist



- A workstation running NetWare Administrator
- A minimum of 6 MB of memory available on the workstation
- The Supervisor object right to the object at the root of the partition whose replica you want to change

Procedure



- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the “Tools” menu, choose “Partition Manager.”**
- 3. Browse the “Partition Manager” screen to locate the partition that has a replica whose type you want to change.**

If the NetWare Server object does not appear in the window, browse the tree by choosing an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

- 4. Select the partition whose replica you want to change and choose “Replicas.”**
- 5. From the list of replicas, select the replica you want to change.**

Be careful to select the correct replica by noting the server’s name as well as its type.
- 6. Choose “Change Type.”**
- 7. Choose the type you want the replica to become.**

8. Choose “OK.”

The new replica type appears on the “Partition Replicas” screen.

If you changed a replica type to a master and a master replica already existed, the replica you just changed to master is now the master replica, and the old master replica has been changed automatically to a read/write replica.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>

Changing a Replica’s Type Using PARTMGR

Prerequisites



- A workstation running DOS v3.30 or later
- The Supervisor object right to the object at the root of the partition

Procedure



1. At the DOS prompt, type

PARTMGR <Enter>

2. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

3. Select the partition whose replica you want to change.

- ◆ If the partition is on the list, select it and press <F10>.
- ◆ If the partition is not on the list, browse the tree by selecting partitions and pressing <Enter> until you see the partition you want. Select it and press <F10>.

4. Choose “View/Edit Replicas.”

A list of the replicas of the selected partition appears. Each replica’s type is also displayed.

5. Select the replica you want to change and press <F10>.

6. Choose “Change Replica Type.”

7. Press <Enter> and choose the replica type you want the replica to become.

8. To save the replica as the new type, press <F10>.

The new replica type appears on the list. The new replica type appears on the “Partition Replicas” screen.

If you changed a replica type to a master and a master replica already existed, the replica you just changed to master is now the master replica and the old master replica has been changed automatically to a read/write replica.

For more information about	Refer to
Partitions	“NetWare Directory partition,” “Partition management,” and “Replica” in <i>Concepts</i>
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Removing NetWare Directory Services from a Server

Considerations Before Removing NDS

Removing NetWare Directory Services (NDS) from a server may corrupt your NDS database. In most cases, this procedure is not necessary to correct a NetWare 4 problem. Before removing NDS, ask

- ◆ Do other servers get Directory information (such as partition replicas) from this server?
- ◆ Could a bad LAN driver (one that doesn't communicate properly with NDS) be causing the problem?
- ◆ Could routers, network boards, or other hardware components be causing the problem?

If you answered "yes" to any of these questions, only remove NDS as a last resort, after resolving all other problem areas.

By removing NDS from a server, you remove the NetWare Server object from the Directory tree and downgrade the server's volumes to bindery volumes. Trustee assignments and all NDS information, such as links to subordinate partitions, are lost.

If you installed this server into the wrong container, it isn't necessary to remove NDS. Use the Partition Manager tool in the NetWare Administrator utility to move the Server object to the correct container.



Removing NDS may corrupt your Directory database. You must remove NDS if you are reformatting this server's hard drive. Do not remove NDS unless Novell Technical Support advises you to remove NDS from this server.

Removing NDS from a Server

Prerequisites

Checklist



- Access to the server console.
- The Supervisor object right to the NetWare Server object and its associated Volume objects.
- Make sure all of the server's volumes are mounted. Unmounted volumes won't get their Volume objects removed from the Directory database.

Procedure

Procedure



1. **At the server console, type**

LOAD INSTALL <Enter>
2. **From the "Installation Options" menu, select "Directory Options (Install NetWare Directory Services" and press <Enter>.**
3. **From the "Directory Services Options" menu, select "Remove Directory Services from this Server" and press <Enter>.**
4. **At the confirmation prompt, choose "Yes."**

The "Directory Services Login/ Authentication" screen appears.
5. **Type the password for User ADMIN in the "Password" field and press <Enter>.**

After you enter the password, the server's mounted volumes are downgraded and a message displays the number of volumes affected. Then the system checks for Directory connections to other servers.

6. (Conditional) If this server contains a master replica, you must designate another server to hold it. Choose one of the following and press <Enter>:

- ◆ Do it automatically.

INSTALL finds the first server with a replica of this partition and changes its type to “master.”

- ◆ Designate the server yourself.

Highlight the server where you want the master replica located.

Press <F10> to save and continue.

A message notifies you that the master replica has moved.

7. To continue, press <Enter>.

INSTALL removes NDS and deletes the Server object and Volume objects associated with it from the Directory database.



If a Directory link is down (such as another server containing objects from this server), you must use Partition Manager to delete the Server and Volume objects from the Directory database.

8. To exit INSTALL, press <Esc>.

Reinstalling NDS on a Server


This procedure reinstalls Directory Services on a server. It doesn't preserve the trustee assignments from the first time NDS was installed on the server; you must make those assignments again.

Prerequisites



- Access to the server console
- The Supervisor object right to the container you want to install the server in

Procedure



1. **At the server console, type**

LOAD INSTALL <Enter>
2. **From the “Installation Options” menu, select “Directory Options (Install NetWare Directory Services)” and press <Enter>.**
3. **From the “Directory Services Options” menu, select “Install Directory Services onto this Server” and press <Enter>.**
4. **Install Directory Services on the server.**

To select a Directory tree and install Directory Services on the server, follow the instructions in Chapter 2, “Simple Installation,” or Chapter 3, “Custom Installation,” of *Installation*.
5. **To exit INSTALL, press <Esc>.**

Deleting a NetWare Server Object from the NDS Database

Considerations Before Deleting NetWare Server Objects



Deleting a NetWare Server object permanently removes it from your network! It also permanently removes its data and resources from the network.

Deleting a NetWare Server object may corrupt your NDS database, especially if the Server object provides NDS database services (such as storing partition replicas). Before deleting the Server object with the Partition Manager tool, consider these alternatives:

- ◆ Change the master replicas stored on this server to read/write replicas and then delete all replicas on the server. Once all the processes are complete, you can delete the server.
- ◆ Remove NDS from the server with the INSTALL NLM (described in “Removing NDS from a Server” on page 326). This procedure protects your NDS database from lost services.
- ◆ To move the NetWare Server object to another context, use the “Move Server Object” option in the Partition Manager tool . (You do *not* need to delete the Server object and re-create it in another context.)
- ◆ If the Server object you are deleting doesn’t contain any NDS database partitions or replicas, follow the procedures in “Moving Container Objects Using NETADMIN” on page 83.

Deleting a NetWare Server Object Using NetWare Administrator

Prerequisites

Checklist



- A workstation running NetWare Administrator.
- A minimum of 6 MB of memory available on the workstation.
- The Supervisor object right to the container of the Server object you want to delete.
- Create a new master replica (if a master replica is stored on the NetWare Server object). See “Viewing a List of Partitions in a Directory Tree” on page 320.

Procedure

Procedure



- 1. At the server console, bring down the server by typing**

DOWN <Enter>

For more information, see “Bringing Down a Server” on page 446.

- 2. From the MS Windows Program Manager at your workstation, choose the “NetWare Administrator” icon.**
- 3. From the “Tools” menu, choose “Partition Manager.”**
- 4. Browse the “Partition Manager” screen to locate the NetWare Server object that you want to delete.**

If the NetWare Server object does not appear in the window, browse the tree by choosing an object to see its subordinates, or by choosing the arrow to move toward the [Root]. Your current context appears in the upper-left corner of the screen.

For information about moving around in the “Partition Manager” screen and selecting objects, choose the “Help” button.

5. Select the NetWare Server object and choose “Delete Server.”

The NetWare Server object is deleted from the NDS database.

For more information about	Refer to
Deleting objects from the NDS database	“Moving Container Objects Using NETADMIN” on page 83
Using the NetWare Administrator utility	“NetWare Administrator” in <i>Utilities Reference</i>


Deleting a NetWare Server Object Using PARTMGR

Prerequisites



- A workstation running DOS v3.30 or later.
- A minimum of 512 KB of memory available on the workstation.
- The Supervisor object right to the container of the Server object you want to delete.
- Create a new master replica (if a master replica is stored on the NetWare Server object). See “Viewing a List of Partitions in a Directory Tree” on page 320.

Procedure



1. At the server console, bring down the server by typing

DOWN <Enter>

For more information, see “Bringing Down a Server” on page 446.

2. At your workstation’s DOS prompt, type

PARTMGR <Enter>

3. From the “Partition Administration” menu, choose “Manage Partitions.”

Your current context appears in the upper-left corner.

4. Select the NetWare Server object you want to delete.

- ◆ If the NetWare Server object appears on the list, select it and press <Delete>.
- ◆ If the NetWare Server object is not on the list, browse the tree by selecting objects and pressing <Enter> until you see the NetWare Server object you want. Select it and press <Delete>.

The NetWare Server object is deleted from the NDS database.



It may take a long time for all the other servers to know that a server is down. You should wait awhile after downing a server before deleting it.

For more information about	Refer to
Deleting objects from the NDS database	“Moving Container Objects Using NETADMIN” on page 83
Using the PARTMGR utility	“PARTMGR” in <i>Utilities Reference</i>

Importing User Information into the Directory Services Database

The UIMPORT utility allows you to import data from an existing database into the NetWare Directory Services (NDS) database. This utility is particularly valuable if you have hundreds, or thousands, of user records that you want to record in NDS without having to manually re-create each user.

Any application capable of converting records to a comma-separated ASCII file will work with UIMPORT.

You can use UIMPORT to automate the maintenance of your NDS database when you want to

- ◆ Create User objects in the NDS database using records from another database.
- ◆ Update User properties in the NDS database when records are changed in your original database program.
- ◆ Delete User objects when their accounts on the network are no longer needed.

Understanding the Import Process

The process for importing records from your database into the NDS database is as follows:

1. **Generate a data file**—a delimited ASCII text file—from your existing database records.

Most applications let you save data as a comma-separated ASCII text file with quotation marks as delimiters.

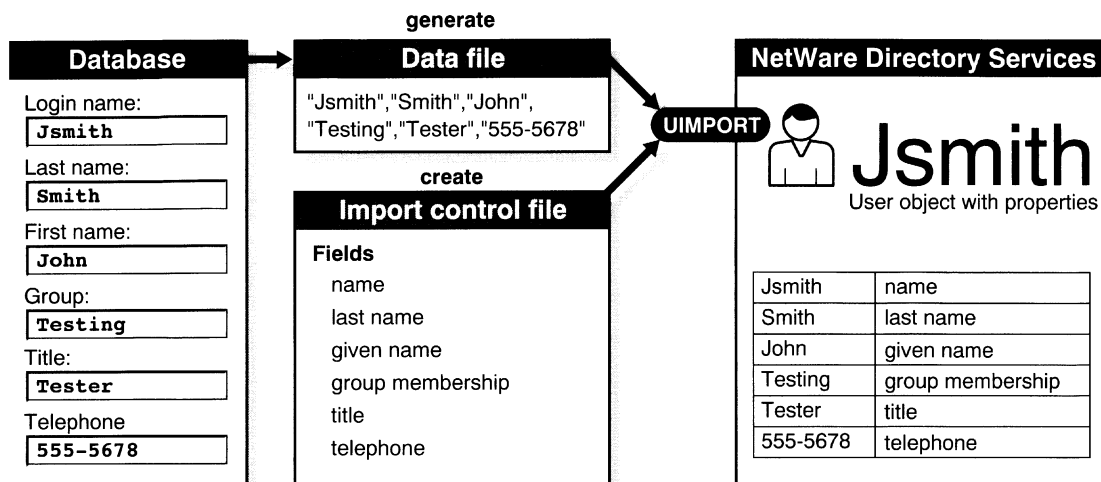
2. **Create an import control file** to interpret the data file.

You can create the import control file using any text editor under DOS or OS/2. This file defines how records in the data file are imported into NDS and which fields (properties) the data will be placed in.

3. **Create User objects** using UIMPORT to transfer data from your database to the NDS database.

The following figure provides a general overview of the import process.

Figure 5-1
Import Process



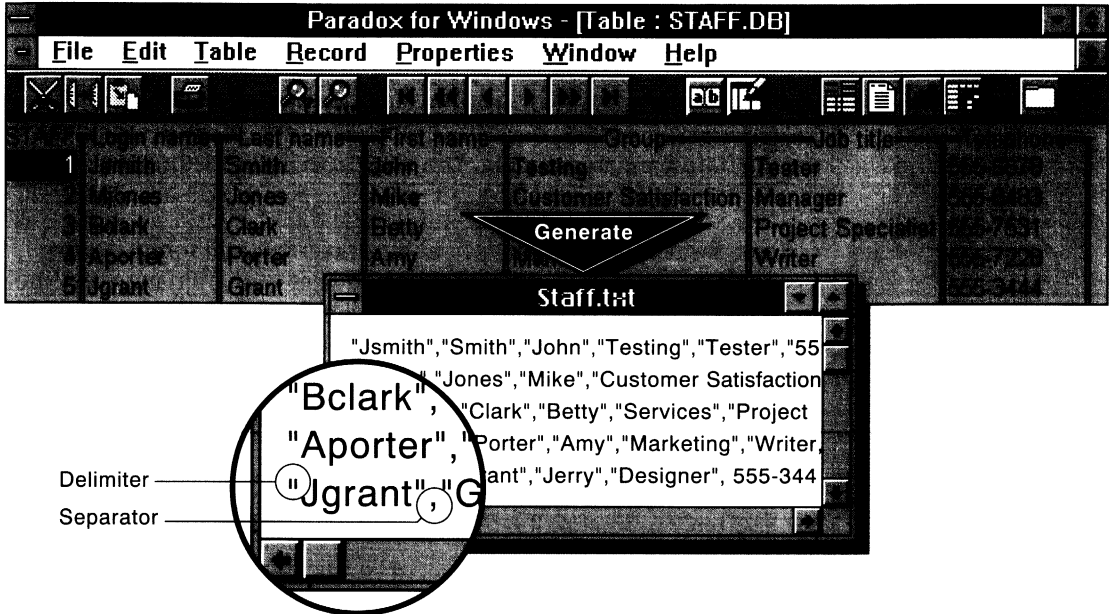
Preparing to Import User Information into the Directory Services Database

Generating the Data File

You generate a data file from within your database application, or any other application, by saving the data as a delimited ASCII text file. Most applications use commas as the default field separator and quotation marks as the default delimiter.


The following figure shows a sample database table created using Borland's Paradox for Windows, and the data file generated from it.

Figure 5-2
Generate a Data File



If you need instructions on how to save your data as a delimited ASCII file, see the documentation accompanying your application.

Before you generate a data file from within your database application, look for embedded punctuation in the data fields.



If the fields in your database have embedded punctuation (for example, quotation marks or commas), you should not use the same punctuation as field separators or delimiters when you export the data.

If you do, UIMPORT won't be able to distinguish between new fields and embedded punctuation.

Since embedded quotation marks and commas are the most commonly found punctuation in database fields, here are some guidelines to help you avoid problems in generating a data file.

They are grouped according to four situations you might encounter:

- ◆ You have both embedded quotation marks and embedded commas that you don't want to delete:

We recommend that you export the data using an alternate character for the comma separator, and that you *not* use quotation marks as delimiters.

For example, you could export data with carets (^) in place of comma separators. As a result, carets would separate each field in the data file you generate. You must specify in the import control file the replacement character you use. (See Table 5-1 on page 340.)

You should also export data without using quotation marks to delimit the fields. If your application uses quotation marks as the default delimiter, you must change the delimiter to an alternate character or choose to use no delimiter.

If you use an alternate character, you must specify in the import control file the replacement character you use. (See Table 5-1 on page 340.)

- ◆ You have neither embedded quotation marks nor embedded commas:

We recommend that you export the data using commas to separate the fields.

If your application inserts quotation marks as delimiters around each field, you can accept this default, or you can choose to use no delimiter.

- ◆ You have embedded commas, but no embedded quotation marks:

We recommend that you export the data using an alternate character for the comma separator.

For example, you could export data with carets (^) in place of comma separators. As a result, carets would separate each field in the data file you generate. You must specify in the import control file the replacement character you use. (See Table 5-1 on page 340.)

If your application inserts quotation marks as delimiters around each field, you can accept this default, or you can choose to use no delimiter.

- ◆ You have embedded quotation marks, but no embedded commas:

We recommend that you export data using commas to separate the fields and without using quotation marks to delimit the fields.

However, if your application uses quotation marks as the default delimiter, you must change the delimiter to an alternate character or tell your application to use no delimiter.

If you use an alternate character, you must specify in the import control file the replacement character you use. (See Table 5-1 on page 340.)

One of the fields required to create new user objects in the NDS database is *name*. If you don't have a field in your database that records each individual's login name, you have a few options:

- ◆ You can create a login name field within your application (and assign each individual a unique eight-character login name) before generating the data file. This option is recommended.
- ◆ You can generate the data file and then add a unique login name field to each string of user information in the data file.
- ◆ You can use an existing field in your application to represent the login name field. For example, you may have a field of employee or ID numbers. You can generate the data file and import the unique numbers into NDS as user login names. (Keep in mind that numbers are not as easy to recognize or manage as names.)

Creating the Import Control File

You can create the import control file using any text editor under DOS or OS/2.

This file controls how the information in the data file will be imported into the Directory Services database, and determines which fields (called properties in NDS) the data will be placed in. The import control file contains two types of information:

- ◆ **Control parameters** define characters used in the data file and specify how data is updated.
- ◆ **Field definitions** determine where data is put in the NDS database (which properties the fields are imported into).

Enter control parameters first, followed by field definitions. (Control parameters are explained in Table 5-1 on page 340. Field definitions are explained in Table 5-2 on page 345.) For example:

```
Import control
  Separator=^
  User template=y
Fields
  Last name
  Name
  Telephone
  Title
```

Control and field parameters are not case-sensitive, though the headings "Import Control" and "Fields" must be left aligned (you don't have to specify import control parameters if you use the defaults). Entries under the headings must be preceded by at least one space or a tab.

To create a new user, two field definitions are required: "Name" (the user's login name) and "Last name." To update records for users who have already been created, only the "name" field is required. These fields are explained in Table 5-2 on page 345.

All fields in the data file must have a corresponding field definition in the import control file. To see a sample data file and its corresponding import control file, see “Data File Created from Exported Database Fields” on page 356.

Although you don’t have to save the import control file in the same directory in which you saved the data file, doing so makes locating the files and running UIMPORT easier. If you don’t, you will have to specify the path to each of the files when you run UIMPORT.

After generating the data file and creating the import control file, see “Creating User Objects with the UIMPORT Utility” on page 359 to start the import process.

For more information about	Refer to
Creating a comma-separated ASCII file	Your database application documentation
Creating and managing User objects	Chapter 1, “Managing NetWare Directory Services Objects”
Starting the import process	“Creating User Objects with the UIMPORT Utility” on page 359
NetWare Directory Services	“NetWare Directory Services” in <i>Concepts</i> , and Chapter 1, “Understanding NetWare Directory Services,” of <i>Introduction to NetWare Directory Services</i>

Control Parameters

Use the applicable control parameters in the following table to create the "Import control" section of the import control file. These can be changed by adding the control parameter in the data file. When doing so, add an exclamation point (!) before the control parameter in the data file. (See "Data File Edited to Update User Objects" on page 357.)

Table 5-1

UIMPORT Control Parameters

Control Parameter	Explanation
Create home directory	<p>Allows you to create a home directory for User objects.</p> <p>If you create a home directory, users are automatically given file system rights to work in that directory. This option <i>will not</i> work unless you set the "Home directory path" and "Home directory volume" control parameters.</p> <p>(Also, this option will not work if you define a data field for "Home Directory," because that field will override Create home directory.)</p> <p>For example, if you want the users you are importing to have a home directory, type <code>CREATE HOME DIRECTORY=Y</code>.</p> <p>The default setting is "N", so if you don't want to create home directories, no parameter is needed.</p> <p>Note: If you don't set this parameter to Y, you can still import Home Directory values through the control section of the import control file or through the data file. These values will be put into the object's Home Directory property, but the directory on the file system will not be created.</p>
Delete Mailbox Dirs	<p>Allows you to delete the mailbox directories when moving a user's mailbox from one messaging server to another or changing the mailbox ID of a user.</p> <p>To delete the mailbox directories, type <code>DELETE MAILBOX DIRS=Y</code>.</p> <p>The default is "N"; so if you don't want to delete the mailbox directories, no parameter is needed.</p>

Table 5-1 continued

UIMPORT Control Parameters

Control Parameter	Explanation
Delete property	<p>Allows you to delete values for a property of a User object.</p> <p>For example, to delete all titles from a User object's Title property, type DELETE PROPERTY=#DEL, and then edit the data file by putting "#DEL" in the Title field.</p> <p>If you delete a group membership, you will also delete the security equals for the group.</p> <p>You cannot delete the following properties with the Delete property parameter:</p> <p><i>Volume restrictions.</i> To delete volume restrictions, enter the volume name and a negative integer such as -1 for the restrictions. For example: VOL1:-1</p> <p><i>Password.</i> To delete the password, enter "" (a null field) for the password field.</p> <p><i>Home directory.</i> You can delete the Home directory property, but the home directory path will not be deleted by UIMPORT. You will need to manually go to this directory and delete it, if desired.</p> <p>There is no default for this property, so if you don't want to delete properties, no parameter is needed.</p> <p>For more information, see "Creating User Objects with the UIMPORT Utility" on page 359.</p>
Home directory path	<p>If you create a home directory for User objects, you <i>must</i> specify a path in the file system where you want the directories to be created. This and the Home directory volume must be entered as a pair. If you specify a Home directory volume, the path will be assumed to be null unless this field is also specified.</p> <p>If you want home directories created in the Users directory, type HOME DIRECTORY PATH="USERS". Do not include the volume name in the path.</p>

Table 5-1 *continued*

UIMPORT Control Parameters

Control Parameter	Explanation
Home directory volume	<p>If you create a home directory for User objects, you <i>must</i> specify the name and context of the Volume object associated with the file system where the home directories are created.</p> <p>For example, if the file specified in “Home directory path” is on volume SYS:, enter the Volume object’s complete DS name. For example: HOME DIRECTORY VOLUME=“SYS.STUDENT RECORDS.UNIVERSITY”.</p> <p>When creating users, the Home directory volume is set in the following order:</p> <ol style="list-style-type: none"> 1. Information in the data file. If the data file has a home directory specified, its value will be used. 2. If there is no home directory specified in the data file, the home directory from the User Template object will be used. This happens regardless of whether or not you select to copy the User Template properties when creating users. 3. If the home directory is not found in the data file or the User Template, this control parameter (Home directory volume) will be used.
Import mode	<p>Controls how the User objects will be created or updated. Options are C (create new objects only), U (update data for existing objects only), B (both create and update), and R (remove objects).</p> <p>The default is “B”, so if you want new objects to be created <i>and</i> existing users updated, no parameter is needed.</p> <p>Note: You can use the same import control file for adding and deleting users by just changing the Import Mode to R (Remove objects). This can be especially useful if you have just created some users with incorrect information and need to restart the entire process. Just change the Import mode to R, run UIMPORT to delete the users, then change the Import Mode back to C or B and rerun UIMPORT.</p>

Table 5-1 *continued*

UIMPORT Control Parameters

Control Parameter	Explanation
Maximum directory retries	<p>Allows you to specify how many times UIMPORT should attempt to get the object ID of the user you have just added in order to create home directories and mailbox directories.</p> <p>For example, type Maximum directory retries=5.</p> <p>If Directory Services creates your user on one server's replica and you create the home or mailbox directory on another server, the second server may not know about the user immediately. This option allows time for the network to catch up with the request to add the user on the other server.</p> <p>Initially this option is set to 5; but if you get the message "991: An error occurred in NWDSMapNameToID..." you should increase this number.</p>
Name context	<p>Specifies the Directory Services context where user objects will be created. You should always use this parameter.</p> <p>For example, if you are in the Organization Development.ACME and you want to create user objects in a different Organizational Unit (Engineering), type NAME CONTEXT=.Engineering.ACME</p> <p>The default is your current context (the context displayed when you type CX from the command line).</p>
Quote	<p>Specifies the character used to delimit fields when exporting data to the data file. You might want to change the delimiter from quotes to different characters if your data has embedded quotation marks and you don't want to delete them.</p> <p>For example, if you have any embedded quotation marks in fields in your database, you could export data using carets (^) as the delimiter. You would enter QUOTE=^ as the quote parameter.</p> <p>The default character for delimiters are quotes (""). If you don't have embedded quotation marks in fields in your database, and you generate the data file with quote delimiters, no parameter is needed.</p>
Replace value	<p>Allows you to overwrite or add data to multivalued properties.</p> <p>For example, to import new telephone numbers to User properties without saving the existing numbers, type REPLACE VALUE=Y.</p> <p>The default is "N," so if you don't want data to be added to multivalued properties, no parameter is needed. You do not need to set this parameter for single-value fields, because the new value automatically overwrites the existing value.</p>

Table 5-1 *continued*

UIMPORT Control Parameters

Control Parameter	Explanation
Separator	<p>Specifies the character used to separate fields when exporting data to the data file.</p> <p>If the fields in your database have embedded punctuation, (such as commas), you should not use the same punctuation as field separators when you export the data to the data file. If you do, UIMPORT won't be able to distinguish between new fields and embedded punctuation. Instead, use a different character as the separator.</p> <p>For example, if you have any embedded commas in fields in your database, you could export the data using a semicolon (;) as the field separator. You would enter SEPARATOR=; as a control parameter.</p> <p>The default separator is a comma. If you have no embedded commas in fields in your database, and you generate the data file with comma separators, no parameter is needed.</p>
User template	<p>Specifies whether you want user template defaults to be applied to the User objects being created. (User template defaults are explained in "Managing User Templates" on page 67.)</p> <p>To apply template properties to User objects, type USER TEMPLATE=Y. Template properties are applied first, and then properties from the data file are imported.</p> <p>The default is "N", so if you do not want the user template defaults applied to the new User objects, no parameter is needed.</p> <p>The following fields are copied from the user Template:</p> <p>Account balance, Account has expiration date, Allow unlimited credit, Allow user to change password, City, Days between forced changes, Default server, Department, Description, Fax number, Foreign Email address, Foreign Email alias, Full name, Generational qualifier, Given name, Grace logins allowed, Group membership, Home directory, Language, Location, Login allowed time, Login script, Low balance limit, Mailbox location, Mailing label information, Maximum connections, Minimum password length, Network address restriction, Postal (ZIP) code, Postal office box, Profile, Remaining grace logins, Require a password, Require unique passwords, Security equal to, See also, State or province, Street address, Telephone, Title</p>

Field Definitions

Use the applicable definitions in the following table to create the import control file.

A single-value property can contain only one entry (property). If you put more than one entry for a single-value property in the control file, only the last one will be saved.

A multivalue property can contain more than one entry. If you want to replace existing values in existing users, make sure to set the Replace Value control parameter to Y. By default, all values will be added to those already in the user object.

Table 5-2
UIMPORT Field Definitions

Property	Property description
Account balance	Single-value property. In the data file, enter the beginning account balance for the user. If accounting is not turned on, this field will have no use.
Account disabled	Single-value property. In the data file, enter Y if you want to disable this account. If you enter Y, you can't modify Password or Login script. The default is N; so if you don't want the account disabled, no field is required.
Account has expiration date	Single-value property. In the data file, enter the date in the form MM/DD/YY. It will expire at 12:01 am on that date.
Allow unlimited credit	Single-value property. In the data file, enter Y to allow unlimited credit. Enter N to not allow unlimited credit. If accounting is not turned on, this field will have no use. The default value for this property is N.
Allow user to change password	Single-value property. If you don't want the user to be able to change his/her password, enter N in the data file. The default is "Y"; so if you want to allow the user to change his/her password, no parameter is needed.
City	Single-value property. In the data file, enter the name of the city.

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Date password expires	Single-value property. In the data file, enter the date the password expires in the form MM/DD/YY. It will expire at 12:01 am on that date.
Days between forced changes	Single-value property. In the data file, enter the number of days before requiring the user to change his/her password. The default for this value is 40. The possible values for this field are numbers 1–365. Do not set this value to 0.
Default server	Single-value property. In the data file, enter the name of the server from which the user gets his/her messages when they are sent with the SEND utility. In the data file, enter the Directory Services name of the server, including the context if the server is not in the same context as the user.
Department	Multivalued property. In the data file, enter the department name, code, number, or other type of information.
Description	Single-value property. In the data file, enter any type of information about the user that you want to import. The field must be enclosed in quotes (or the delimiter defined in the control file) if it contains commas or new-line characters. This field is a string; all data from the opening quote to the closing quote is considered part of the field.
Fax number	Multivalued property. In the data file, enter the fax number for each user you want to import.
Foreign Email address	Single-value property. In the data file, enter the foreign Email address of the user, which specifies a mailbox that resides in a foreign Email system. This information is used by NetWare MHS Services. The format for this value is Type:Value, where Type is the type of messaging protocol and Value is the the user's address (in the format required by the foreign Email system). For example, in your data file you could enter a field such as SMTP:Jjones@Acme.com

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Foreign Email alias	<p>Multivalued property. In the data file, enter the Foreign Email aliases of the user, which specify an object's aliases as known in a foreign messaging system. This information is used by NetWare MHS Services.</p> <p>For example, an MHS user (a user whose mailbox is located on an MHS Messaging Server) can have an X.400 alias—with an X.400 native name—so that X.400 users can use this alias to send mail to the MHS user.</p> <p>The format for this value is Type:Value, where Type is the type of messaging protocol and Value is the user's name (in the format required by the foreign Email system).</p> <p>For example, in your data file you could enter a field such as X400:g=joe;s=jones;ou=sales;o=acme;p=acmemd;a=att;c=us</p>
Full name	<p>Single-value property. In the data file, enter the user's full name.</p> <p>For example: Roland D Bruns</p>
Generational qualifier	<p>Single-value property. In the data file, enter the generational qualifier for the user's name. Usually this is a Jr., Sr., II or III.</p> <p>If the user's name is Bob Wilson III, enter III in the data file.</p>
Given name	<p>Single-value property. In the data file, enter the given name of the user. For example, for John Doe, enter John.</p>
Grace logins allowed	<p>Single-value property. In the data file, enter the number of grace logins allowed before the account is locked. 0 implies no limit.</p>
Group membership	<p>Multivalued property. This field allows you to organize groups of users who need similar rights and access to network resources. In the data file, enter the names of the groups to which this user belongs. If the group object is in a different context than the user, enter the complete Directory Services name of the group, including its context.</p>

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Home directory	<p>Single-value property. In the data file, enter the volume and path of the home directory in the following format:</p> <p>Volume Name:path</p> <p>If the volume is in a different context than the user, enter the complete Directory Services name. Separate the volume name and the path by a colon (:). Include the full path, including the final directory.</p> <p>For example, for user ABC1, create the home directory on volume object VOL1. VOL1 is in the same context as the user.</p> <p>VOL1:USERS\ABC1</p> <p>This field in the data file will override the Home Directory field in the control section.</p>
Initials	<p>Single-value property. In the data file, enter the middle initial of the user. For John I. Doe, enter I here.</p>
Language	<p>Multivalued property. In the data file, enter the language directories to be searched to find the message files for NetWare utilities for this user.</p> <p>If more than one language directory should be searched, enter multiple language fields. The order is important. The first entry will be the first language directory to be searched, the second is the second language directory to be searched, etc.</p>
Last name	<p>Single-value property. In the data file, enter the user's last name. A value is required for this field when you are creating new User objects.</p>
Location	<p>Multivalued property. In the data file, enter any information about a user's work, department, or division location.</p>

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Login script	<p>The value of the login script is the name of a file containing the ASCII text for the login script. You can reuse the same file for all users or have a separate file for each user. In the data file, enter the path to the file as a DOS path. For example:</p> <p style="text-align: center;">C:\DIR1\MYLOGIN.SCR</p> <p>The maximum size of the login script is 16,384 bytes. If the file is larger than this, only the first 16,384 bytes will be copied to the login script.</p>
Low balance limit	<p>Single-value property. If unlimited credit is not allowed, in the data file, enter the lowest balance the user can have and still receive services that are chargeable.</p> <p>If accounting is not turned on, this field has no effect. Also, if Allow unlimited credit is set to "Yes," this field has no effect.</p>
Mailbox ID	<p>Single-value property. This is a unique name that specifies the directory in which all of the object's inbound mail will be placed. This information is used by NetWare MHS Services.</p> <p>In the data file, assign the user a mailbox ID using the user's login name. If the user's name has spaces in it, or if it uses non-DOS characters, assign the user a mailbox ID using the user's login name, but eliminate the spaces and other illegal characters to form a legal DOS name.</p> <p>If you do not assign a user a mailbox ID, UIMPORT will do so automatically.</p>
Mailbox location	<p>Single-value property. In the data file, enter the name of the messaging server on which the user's mailbox is located. If the messaging server is not in the same context as the user, enter the complete Directory Services name of the messaging server.</p> <p>For example, if the user's mailbox is located on a messaging server named SERVMAN_MSG and the context of SERVMAN_MSG is Publications.Novell, enter SERVMAN_MSG.Publications.Novell.</p>
Maximum connections	<p>Single-value property. In the data file, enter the number of workstations the user can login from. For no limit, enter 0.</p>

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Minimum password length	Single-value property. In the data file, enter the minimum password length.
Name	Single-value property. A value is required for this field. In the data file, enter any unique user name. For example, you can use a student or employee identification number as the login name. This field will be the user's login name in NDS.
Other names	Multivalued property. For example, if a user has two additional names that you want to import, enter both them in the data file and enter two field definitions in the import control file called "Other names".
Password	Single-value property. In the data file, enter a unique password for the user. Numbers, letters, and special characters may be used.

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Mailing label information	<p>Multivalue property. This is a single property with up to six lines of information. If you want a user's entire address imported as a single property, use "Mailing label information" for each field that contains a part of the address. The first field will be the first line of the postal address, the second field will be the second line, etc.</p> <p>If you are going to define only four lines of the postal address in the data file, you only need to include four Mailing label information fields. Each line of the postal address in the data file must be a separate data string (i.e. separated by a comma or other separator).</p> <p>To import four lines of the postal address, you would type the following in the Fields section in the Control file:</p> <p>Fields</p> <p style="padding-left: 40px;">...(other fields) Mailing label information Mailing label information Mailing label information ...(other fields)</p> <p>In the data file, you might enter a line similar to the following:</p> <p>...,"1234 Any Street","Torrance","CA","98550",...</p> <p>You can use the Mailing label information field in place of Postal code, Post office box, State or province, and City.</p>
Postal (ZIP) code	<p>Single-value property. Enter in the data file the five- or ten-digit zip code. (84111 or 84111-1111). For a Canadian address this is the postal code.</p>
Post office box	<p>Single-value property. In the data file, enter the user's post office box.</p>
Profile	<p>Single-value property. In the data file, enter the name of the profile object to which this user belongs. This must include the complete Directory Services name of the profile object if the profile is in a different context than the user.</p>

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Remaining grace logins	<p data-bbox="539 291 1208 434">Single-value property. Normally this field is not imported. This is the number of grace logins still remaining. It is set to the same value as in "Grace logins allowed." If you want this number to be different than the "Grace logins allowed" number, you will have to use the modify option in UIMPORT.</p> <p data-bbox="539 462 1208 548">For those users whose Remaining grace logins should be the same as the Limit grace logins, either don't import this field or set it to the same number you entered for Limit grace logins.</p>
Require a password	Single-value property. In the data file, enter Y if a password is required. If no password is required for this user, enter N.
Require unique passwords	Single-value property. If unique passwords are required, enter Y in the data file. If unique passwords are not required, in other words the user can reuse passwords, enter N.
See also	Multivalued property. Enter in the data file any other related objects. Enter the complete Directory Services name of the object if it is not in the same context as the User object.

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Skip	<p>This is a unique processing option for UIMPORT. The values in the data file that correspond to Skip fields are ignored by UIMPORT.</p> <p>For example, if you exported into the data file users' birthdays and you do not want this information imported into NDS, enter "Skip" as the field definition. So if you export data that you don't want to import, rather than reexporting the entire data file, you can simply specify "Skip" for the column that contains the information.</p> <p>Suppose you export to the data file a work and a home phone number, but you don't want to import the home phone number. Your data file may have a string like this:</p> <p>...,"801-555-4141","801-555-8677",...</p> <p>Your import control file would correspond like this:</p> <p>Fields</p> <p>...</p> <p> Telephone</p> <p> Skip</p> <p> ...</p> <p>The 801-555-8677 field in the data file will be ignored.</p>
Security equal to	<p>Multivalued property. In the data file, enter the name of the objects to which this user is security equivalent. This name must include the complete Directory Services name of the object if the object is in a different context than the user.</p>
State or province	<p>Single-valued property. In the data file, enter the State or province of the user.</p>
Telephone	<p>Multivalued property. For example, if a user has three telephone numbers that you want to import, enter in the data file each of the three numbers. In the control file, you would need three "Telephone" field entries.</p>
Title	<p>Multivalued property. In the data file, enter the user's title.</p>

Table 5-2 *continued*

UIMPORT Field Definitions

Property	Property description
Volume restrictions	<p>Single-value property. In the data file, enter the name of the volume object and the space restrictions on that volume in the following format:</p> <p>Volume object name:Restriction amount</p> <p>To remove volume restrictions, enter -1 (or any other negative number) for the restriction amount.</p>

UIMPORT Field Name Changes

UIMPORT data field names have changed in NetWare 4.1. Field names used in NetWare 4.02 are listed in the left column of the following table and new NetWare 4.1 names are listed in the right column.

4.02 UIMPORT Field Name	4.1 UIMPORT Field Name
Minimum account balance	Low balance limit
Login expiration time	Account has expiration date
Login grace limit	Grace logins allowed
Login maximum simultaneous	Maximum connections
Login disabled	Account disabled
Password expiration time	Date password expires
Password expiration interval	Days between forced changes
Password minimum length	Minimum password length
Password required	Require a password
Password unique required	Require a unique password
Password allow change	Allow user to change password
Postal address	Mailing label information
Security equals	Security equal to

4.02 UIMPORT Field Name	4.1 UIMPORT Field Name
Facsimile telephone number	FAX number
Telephone number	Telephone
New field	Foreign Email address
New field	Foreign Email alias
New field	Full name
New field	Generational qualifier
New field	Initials
New field	Mailbox ID
New field	Mailbox location

Examples of Files Used by UIMPORT

The following examples are based on records taken from a database sample that uses these fields:

Last name:

First name:

Middle initial:

Local address

Street:

City:

State or province:

Zip code:

Student number:

Year:

Major:

Cumulative grade point average:

Department:

Data File Created from Exported Database Fields

When you export your database records to a comma-separated ASCII file, the records appear in the data file as shown here:

```
"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"
```

Corresponding Import Control File

Using the data file as a guide, you could set up the import control file to import the data file fields as shown here. This example assumes that you are creating new users in the NDS database:

Import control

```
Name context=.administration.student_accts  
User template=y  
Create home directory=y  
Home directory path="Students/Home"  
Home directory volume=".SYSVOL.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
```

Note how the Name field corresponds to a student ID number. When the ID number is imported into NDS, it will become the user's login name—which is the User object name.

Managing numbers is difficult because they don't uniquely identify a user. To avoid having to manage numbers, you can create a field of unique login names in your data file.

Data File Edited to Update User Objects

Suppose that later you wanted to update User object information and add properties that were not imported when the users were created.

For example, if you wanted to delete values from the "Middle initial" field, import new values from the "Year," "Major," and "Grade point average" fields, and change the separator used for all user information after Adam Jones, you would edit the data file as shown here:

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

```
!Separator=/  

```

```
"Smith"/"John"/"#DEL"/"222 North Cerillos"/"Los  
Angeles"/"California"/  
"96000"/"2875"/"Senior"/"Accounting"/"3.0"/"Business  
Administration"
```



Note Import control parameters are initially set in the control file, but can be included in the data file. If you include control parameters in the data file, the parameter should be preceded by an exclamation point (!).

For example, to place the users in the Engineering.Novell context, and to change the separator to a semicolon (;) you would add the following lines to the data file:

```
!Name context=.Engineering.Novell.  
!Separator=;
```

Corresponding Changes to the Import Control File

After editing the data file, you must change the import control file to reflect the changes you made.

In the following example, the “delete property” and “replace value” control parameters have been added to delete the marked properties and specify that you want properties in both single and multivalued fields to be replaced by new properties. Fields marked “Skip” are replaced by “See also,” so the properties will be imported to the “See also” field for the users.

Import control

Name context=.administration.student_accts

User template=y

Create home directory=y

Home directory path=students/home

Home directory volume=.SYSVOL.Student Records

Delete property=#DEL

Replace value=y

Fields

Last name

Given name

Middle initial

Postal address

Postal address

Postal address

Postal address

Name

See also

See also

See also

Department

Creating User Objects with the UIMPORT Utility

Prerequisites



- A workstation logged in to the network, running DOS 3.30 or above or OS/2 v2.x, and using NetWare 4.1 client software
- A minimum of 512 KB of memory available on the workstation
- The Create or Supervisor object right to the container object where the User objects will be created
- The Create or Supervisor file system right to the directory where user home directories will be created (if applicable)

Procedure



1. Generate a data file from within your database application.

For general instructions, see “Generating the Data File” on page 334. For specific instructions, see the documentation accompanying your application.

2. Create an import control file.

If you have enough memory, save this file in the same directory in which you saved the data file. Otherwise, you will need to specify the path to both the data file and the import control file when you run UIMPORT.

3. Run UIMPORT from the directory in which you saved the data file and the import control file.


At the DOS command line, type

```
UIMPORT [control_file] [data_file] [/C] <Enter>
```

Replace *control_file* with the name of the import control file you created. Replace *data_file* with the name of the data file you generated from your database application. Use /C to run UIMPORT with continuous output.



If you don't run UIMPORT from a directory that contains both the import control file and the data file, you must specify the path to each file.



Important

Importing hundreds or thousands of users may take several hours. Therefore, you might want to run UIMPORT at night.

4. (Optional) Route import errors to a file.

If errors occur during the import process, messages appear on your workstation screen. If you want the messages sent to a file, add the DOS ">filename" option to the command.

For example, to route the messages to a file called uimport.log in your home directory, use a command similar to the following:

```
UIMPORT [control_file] [data_file]  
          >home\paul\uimport.log <Enter>
```



Note

When importing large groups of users, the performance of your server might begin to slow down. This reduction in performance is related to the amount of memory available for Directory Services and TTS to use with UIMPORT's numerous transactions. This reduction is related to the size of UIMPORT batches available server memory.

If you will be importing large groups of users or experience decreased server performance while running UIMPORT, you can add server memory, process smaller batches of users, or process UIMPORT batches during periods of low server utilization.

For more information about	Refer to
Control parameters	Table 5-1 on page 340
Creating import files	"Repairing the NetWare Directory Database" on page 361
Field definitions	Table 5-2 on page 345

Repairing the NetWare Directory Database

The DSREPAIR utility is provided with NetWare 4™ software to repair problems with NetWare Directory Services on a single-server basis. It does not correct problems on other servers from a single, centralized location. It must be run on each server that you want to correct Directory database errors on.

DSREPAIR Overview

The DSREPAIR utility allows you to maintain and repair the local NetWare Directory database of a tree. This utility performs the following operations:

- ◆ Repairs the local database

The repair is performed for the DS.NLM file stored on the current server you are running the DSREPAIR utility on.

- ◆ Repairs local replicas

You are provided with tools to repair replicas, replica rings, and server objects. You can also view the purge time of each replica to ensure that data in each replica is the same.

- ◆ Analyzes each server in each local partition for synchronization errors

You can view errors and list the partition name, server name, synchronization time, and error code for each error.

- ◆ Writes replica information to a log file

The log file contains detailed information about local partitions and servers. This information helps you diagnose damage to the database.

- ◆ Creates a dump file of a damaged database

The dump file is saved in a compressed format. You can use DSREPAIR to diagnose and repair the damaged database.


- ◆ Checks the remote server ID list

You can access a list of the identification numbers for remote servers. You can use this list to verify the identification numbers of remote servers and modify the numbers if you need to.

- ◆ Searches for local database objects

A browser allows you to locate and synchronize objects in the local database.

Some NDS database problems are not fatal, and Directory Services continues to operate. But if the database becomes corrupted, you get a message on the console that the server could not open the local database. In this case, run DSREPAIR or reinstall the NDS database to fix the problem so that the database can be opened.



DSREPAIR affects only the parts of the database stored on the server where you run it. To fix the entire database, you must run the utility on each server which contains a part of the database.



DSREPAIR changes inconsistent objects to Unknown objects when they do not have mandatory properties or are invalid in other respects (their properties don't meet minimum requirements for an object type). Unknown objects can be deleted but cannot be changed back to their original object type.

In NetWare 4.1, Unknown objects are represented by question mark icons in the NetWare Administrator utility.

DSREPAIR Options

After you load DSREPAIR, you can use the following options:

Option	Use to
Unattended Full Repair	Automatically perform all possible repair operations to the Directory database that do not require operator assistance.
Time Synchronization	Contact all servers within this server's local database to request information about Directory Services and time synchronization. If a replica of the root partition is contained on this server, then all servers in the Directory tree are contacted.
Replica Synchronization	Determine the status of synchronization for every replica in the replica table for the Directory tree. The status of synchronization can inform you of the current condition of the Directory tree.
View/Edit Repair Log File	Track all operations of the DSREPAIR utility to a single file. The default log file is <code>SYS:SYSTEM\DSREPAIR.LOG</code> . You can configure options for the log file by accessing "Log File And Login Configuration" in the "Advanced Options" menu.
Advanced Options Menu	The "Advanced Options" menu allows you to manually perform individual or global repair operations on the Directory tree. You can also access diagnostic information about the Directory tree database to analyze the status of the tree. See "Using the Advanced Options" on page 371 for more information.

Running an Unattended Full Repair

Use this procedure to perform an automatic repair of the Directory database. Any operation requiring an operator's assistance, such as managing partition replicas or editing remote server ID numbers, is not performed.


If you want to manually repair the Directory database, use options available in the "Advanced Options" menu. See "Using the Advanced Options" on page 371 for more information.

Prerequisites



- Access to the server console or an established RCONSOLE session with the server
- All servers in a tree using the same time source
- The Supervisor object right to the [Root] object of the Directory tree

Procedure



1. At the server console prompt, load the module by typing

LOAD DSREPAIR <Enter>

If you have placed DSREPAIR in a directory other than SYS:SYSTEM, you must enter the full path to the utility.

The utility locks the database.

2. Choose "Unattended Full Repair."

The "Repairing Directory on Server *servername*" screen appears. A window is available to observe the repairs in process.

Following the automatic repair, a message window appears that informs you of the repair status, the total number of errors corrected, and the amount of time used to complete the repair operation.

3. Press <Enter> to display the error log file.

The error log is displayed within a full text editor. You can annotate or modify the DSREPAIR log if you want to.

Checking Time Synchronization (DSREPAIR)

Use this procedure on the server before or after performing a repair. It allows you to contact all servers within this server's local database to request information about Directory Services and time synchronization.

If a replica of the root partition is contained on this server, then all servers in the Directory tree are contacted.




Time synchronization is very important to Directory Services. All servers in a tree must be synchronized to the same time source. If they are not, collision will occur when synchronizing objects in replicas.

Prerequisites



- Access to the server console or an established RCONSOLE session with the server
- All servers in a tree using the same time source

Procedure



1. At the server console prompt, load the module by typing

LOAD DSREPAIR <Enter>

If you have placed DSREPAIR in a directory other than SYS:SYSTEM, you must enter the full path to the utility.

The utility locks the database.

2. Choose “Time Synchronization.”

The “Collecting Time Synchronization and Server Status” screen appears. A window is available to observe the operations in process.

Following the operation, the DSREPAIR log is displayed within a full text editor. You can annotate or modify the DSREPAIR log if you want to.

The DSREPAIR log contains the following:

Field	Indicates
Server Name	All of the server names known to the local Directory database. If this server contains a replica of the root partition, then this list contains all the servers in the tree.
Version	The version of DS.NLM running on the server.
Replica Depth	The replica depth field reports “-1” if no replicas are stored on the server or “0” if the server contains a replica of the root partition. A positive integer indicates how many objects deep from the root the first replica is on that server.

Field	Indicates
Time Source	<p>The time server type. This information helps you determine if time synchronization for all the trees in the server is configured properly.</p> <p>All servers in a tree must be using the same time source.</p> <p>For example, if there are two Single time servers then you know that all servers in the tree cannot be polling the same time source and there is a configuration problem.</p>
Time in Sync	<p>The local time synchronization status on each server. The status should be "Yes." If the status is "No," then the server is not able to contact its time source.</p>
Time +/-	<p>The difference in time between the local server and the selected server in the list. All servers should be within one second of each other; if they are not, they have not been configured properly.</p> <p>This field reports up to 999 minutes and 59 seconds (which is approximately 16 hours and 30 minutes) in the form <i>minutes:seconds</i>. If the time difference is greater than 16 hours and 30 minutes, then the maximum value is displayed as: -999:59.</p> <p>If the difference in time is more than a few minutes, it might indicate that the servers are using different time source servers.</p>

Checking Replica Synchronization


Use this procedure on the server to determine the status of synchronization for every replica in the replica table for the Directory tree.

The synchronization status can inform you of the current condition of the Directory tree.


If you want to manually manage the Directory partitions and replicas, use the “Replica And Partition Operations” option in the “Advanced Options” menu. See “Using the Advanced Options” on page 371 for more information.



Normal partitioning operations should be done with one of the client utilities (NWADMIN or PARTMGR). The “Replica Synchronization” option should only be used when the master replica of a partition has been lost because of server or hardware failure and the server will be reinstalled into the NetWare Directory tree.



You need to synchronize damaged replicas manually, using the PARTMGR utility or Partition Manager in the NetWare Administrator utility.




Be careful when changing a secondary replica to a master. If a server containing a secondary is changed to a master and the server containing the original master is brought back up, there is a possibility that the partition will have two master replicas.

Prerequisites



- Access to the server console or an established RCONSOLE session with the server
- The Supervisor object right to the [Root] object of the Directory tree

Procedure



1. At the server console prompt, load the module by typing

LOAD DSREPAIR <Enter>

If you have placed DSREPAIR in a directory other than SYS:SYSTEM, you must enter the full path to the utility.

The utility locks the database.

2. Choose “Replica Synchronization.”

The “Log In as the Admin” screen appears.

3. Enter the administrator name and password to log in to the tree.

Log in as a user who has the Supervisor object right to the [Root] object of the tree. Enter your complete name (Distinguished Name), such as ADMIN.NOVELL or CN=ADMIN.O=NOVELL. Entering only ADMIN is invalid because it is not a complete name.

4. Press <Enter>.

The “Collecting Replica Synchronization And Server Status” screen appears. A window is available to observe the operations in process.

Following this operation, the DSREPAIR log is displayed within a full text editor. You can annotate or modify the DSREPAIR log if you want to.

The error log contains the following fields:

Field	Indicates
Servers That Contain Replicas	The servers in the tree that contain replicas.
Replica Type	<p>The type of replica for each server that contains replicas of partitions.</p> <p>Following is a list of the possible types of replicas.</p> <p>Master. This is a writable replica that can also handle partition operations. There is only one master replica per partition. Only one partition operation is valid at any time for a partition, and the master enforces that requirement. For all nonpartition operations, this replica is equivalent to a read/write replica.</p> <p>Read/Write. This is a writable replica that can be updated from a client workstation. Both read/write and master replicas are valid for login and authentication requests.</p> <p>Read-Only. This is a replica that cannot be changed from a client workstation. It is updated with the changed data in the replica from another read/write or master. This replica cannot be used for bindery services because bindery services requires a writable replica to be contained on the server for bindery users.</p> <p>Subordinate Reference. This is a replica of the root partition which includes the replica list (ring). As a child partition, it resides on every server that holds a copy of it's parent partition.</p> <p>This replica is used to facilitate tree connectivity.</p>
Status	The synchronization status of partition replicas for servers in the tree. The server running DSREPAIR does not synchronize to itself, so the status for the server's own replica is displayed as "Host".

Viewing and Editing the Repair Log File

Use this procedure on the server to view and edit repair information for the local Directory tree.

The default log file is SYS:SYSTEM\DSREPAIR.LOG. You can change the filename with the “Log File and Login Configuration” option in the “Advanced Options” menu. See “Using the Advanced Options” on page 371 for more information.

Prerequisites



- Access to the server console or an established RCONSOLE session with the server
- The Supervisor object right to Directory objects in the local tree

Procedure



1. **At the server console prompt, load the module by typing**

LOAD DSREPAIR <Enter>

If you have placed DSREPAIR in a directory other than SYS:SYSTEM, you must enter the full path to the utility.

The utility locks the database.

Using the Advanced Options

Use this procedure to access the “Advanced Options” for manual repair of the Directory database.

If you want the Directory database to be repaired automatically, use the options available from the “Unattended Full Repair” option in the main menu. See “Running an Unattended Full Repair” on page 364 for more information.

The “Advanced Options” menu includes the following options:


- “Log File Login Configuration”
- “Repair Local DS Database”
- “Servers Known to this Database”
- “View Remote Server ID List”
- “Replica and Partition Operations”
- “Security Equivalence Synchronization”
- “Global Schema Update”
- “View/Edit Repair Log File”
- “Create a Database Dump File”

Prerequisites



- Access to the server console or an established RCONSOLE session with the server
- The Supervisor object right to the [Root] object in the tree

Procedure



1. **At the server console prompt, load the module by typing**

LOAD DSREPAIR <Enter>

If you have placed DSREPAIR in a directory other than SYS:SYSTEM, you must enter the full path to the utility.

The utility locks the database.

2. **Choose “Advanced options menu.”**
3. **From the menu, select the repair operation you want to perform.**

Refer to the online help in DSREPAIR for more information and instructions on how to use these options in the DSREPAIR utility.

For more information about	Refer to
Directory Services database	"NetWare Directory Services" in <i>Concepts</i>
DSREPAIR utility	"DSREPAIR" in <i>Utilities Reference</i>

Merging NDS Trees

The DSMERGE utility allows you to merge the [Root] of two separate NDS trees. Only the [Root] objects are merged; container objects and their leaf objects maintain separate identities within the newly merged [Root].

DSMERGE Overview

The two trees you merge are called the local source tree and the target tree. To merge two trees, you load DSMERGE on a server in the local tree.

DSMERGE does not change Directory names or contexts within the containers. Object and property rights for the merged objects are retained.



Note

You can't merge container or leaf objects with DSMERGE. To move leaf objects, use the NetWare Administrator graphical utility or the NETADMIN text utility. To merge partitions, use the PARTMGR utility.

Merging the Source into the Target Tree

When you merge the trees, the servers in the source tree become part of the target tree. (The target tree is the tree that the local tree will be merged into.)

The target [Root] object becomes the new [Root] for objects in the source tree, and the tree name of all servers in the source tree is changed to the target tree name.

After the merge, the target tree name is retained.

The objects that were subordinate to the local [Root] object become subordinate to the target [Root] object.

Partition Changes

During the merge, all replicas of the root partition are removed from servers in the local tree. The server that contained the master replica of the local tree receives a replica of the target tree's root partition.

During the merge, DSMERGE splits the objects below the local [Root] object into separate partitions.

The following two figures illustrate the effect on partitions when you merge two trees.

Figure 5-3
Directory Trees
Before a Merge

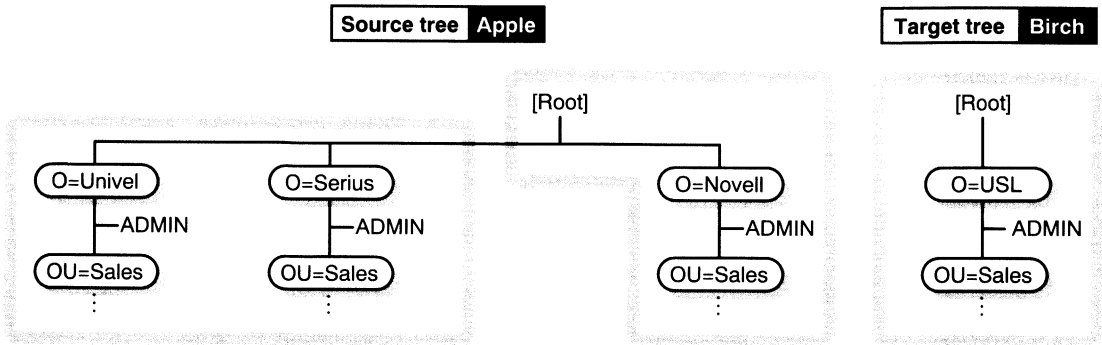



Figure 5-4
Directory Tree After
a Merge



Preparing the Source and Target Trees

Before performing a merge operation, ensure that the state of synchronization for all servers affected by the operation is stable. The following table provides recommendations for preparing source and target trees for merging.

Necessary Precondition	Required Action
No aliases or leaf objects can exist at the [Root] of the source tree.	Delete any aliases or leaf objects at the [Root] of the source tree.
No similar names can exist between the source and target trees.	Rename objects on the source and target trees if similar names exist. Move objects from one of the containers to a different container in its tree if you don't want to rename objects. Then delete the empty container before running DSMERGE. (See "Moving Objects in the Directory Tree" on page 76 and "Moving Container Objects Using NETADMIN" on page 83.) You can have identical container objects in both trees if they are not immediately subordinate to the [Root]. They are uniquely identified by their immediate container object.
No login connections can exist on either the source or target trees.	Close all connections on the source and target trees.
The NDS version must be the same on both source and target tree.	Upgrade all non-4.1 servers that have a replica of the [Root] object.

Necessary Precondition	Required Action
Any server that contains a replica of the [Root] object on both source and target trees must be up and running.	Ensure that all servers containing a replica of the [Root] partition on both source and target trees are up and running. Ensure that any WAN links affected are stable.
Schema on both source and target trees must be the same.	Ensure that both source and target trees are time synchronized within 2 seconds of each other.

Because the merge operation is one single transaction, it is not subject to catastrophic failure caused by power outages or hardware failure. However, you should perform a regular backup of the Directory database before using DSMERGE. See Chapter 9, “Backing Up and Restoring Data.”

Time Synchronization before the Merge



Proper configuration of time synchronization is a very involved process. Make sure you allow enough time to synchronize both trees before you merge the trees.

Directory Services will not work properly if two external time sources are used, or if all servers in a tree are not synchronized.

Before you do the merge, make sure that all servers in both trees are synchronized, and use only one time source.

There should be a maximum of one Reference or one Single Reference time server in a tree. Likewise, after the merge, the tree should contain only one Reference or one Single Reference time server.

If each of the trees you are merging has either a Reference or one Single Reference time server, you must reassign one of them to refer to the Reference or Single Reference time server in the other tree so that the final tree contains only one Reference or Single Reference time server (if any).

For more information on time synchronization, see “Managing Network Time Synchronization” on page 593.

To view time synchronization information, see “Checking Time Synchronization (DSMERGE)” on page 381.

DSMERGE Options

After you load DSMERGE, you can use the following options:

Option	Use to
Check Servers in This Tree	<p>Contact all servers in the local tree to verify that each server has the correct version, status, and tree name.</p> <p>The server you are on must have a replica of the [Root] partition. It does not require the master replica.</p>
Check Time Synchronization	<p>Display a list of all servers in this tree, along with information about time sources and time synchronization.</p> <p>The server you are on must have a replica of the [Root] partition. It does not require the master replica.</p>
Merge Two Trees	<p>Merge the [Root] of the local (source) tree to the [Root] of the target tree.</p> <p>The server you are on must have the master replica of the local tree's [Root] partition.</p>
Rename Tree	<p>Rename the local tree. Use this option if you are merging two roots with the same name.</p> <p>You can rename only the local tree name. To rename the target tree name, load DSMERGE on a server in the target tree. Then load DSMERGE on the source tree to perform the merge.</p> <p>This option requires that the server you are on has the master replica of the [Root] partition.</p>

Checking Servers in the Tree

Before you rename or merge trees, use this option to contact all servers in the tree and verify that all servers have the same tree name.

After you rename or merge trees, use this option to verify that all servers have the new tree name.

Prerequisites



- Access to a server console on the local tree or an established RCONSOLE session with the server
- The Supervisor object right to the [Root] object of the local tree

Procedure



1. **Load DSMERGE on the server where a replica of the [Root] partition of the local tree is stored.**

At the server console prompt, type

LOAD DSMERGE <Enter>

2. **Choose “Check servers in this tree.”**

Each server in the tree is listed in the “Status of Servers in the Tree” screen., with their corresponding status information. Any servers that have existing problems are flagged and then listed at the top of the server list.

You should confirm that each server’s status is marked as “Verified” before completing a merge of two trees.

The following table describes the information provided in the “Status of Servers in the Tree” screen.

Field	Operation
Server Name	Lists the names of all servers DSMERGE contacted, and shows their context within the tree.
Version	Indicates the version of NetWare running on the server.
Status	<ul style="list-style-type: none">◆ UP Indicates that the server is in the right tree.◆ Error <i>number</i> All Directory Services errors are numbered between –600 and –699 in decimal notation. See <i>System Messages</i> for an explanation of specific Directory Services error messages.◆ Unknown Indicates the server is not responding. This is usually the cause of a downed server or of communication problems.◆ Wrong Tree Indicates that this server does not belong to this Directory tree. This status might occur if the tree was recently merged or renamed because the server might take a few minutes to recognize the change. Or, this status can occur if the server was reinstalled in another tree, but not properly removed from this tree. If so, delete this server's object from this tree.

Checking Time Synchronization (DSMERGE)


Use this procedure on both of the trees before merging them.

Prerequisites



- Access to a server console on both the source and target trees or an established RCONSOLE session with the servers
- The Supervisor object right to the [Root] object of both the source and target trees

Procedure



1. **Load DSMERGE on the server where a replica of the [Root] partition of the local tree is stored that has the master replica.**

At the server console prompt, type

LOAD DSMERGE <Enter>

If you don't know where the master replica is, load DSMERGE on any server. You will be prompted with the name of the server that contains the master replica when it is required.

2. **Choose "Check time synchronization."**

The "Time Synchronization Information for Tree *treename*" screen appears.

This option displays a list of all servers in the tree, along with information about their time sources and the server time.

Verify that all servers in the tree are synchronized and that they are using the same time source.

The following table describes the information provided in the “Time Synchronization Information for Tree *treename*” screen.

Field	Operation
Server Name	Lists the name of each server recorded in the tree's [Root] partition. If the server could not be contacted, it is listed as “Unknown.”
Type	Indicates the kind of time server the server is using.
In Sync	<ul style="list-style-type: none">◆ YES Indicates that the server is in sync with the time server.
	<ul style="list-style-type: none">◆ NO Indicates that the server is in not sync with the time server.
	This option does not indicate if the local server is in sync with the server you have selected. Check the type of time server each server is using to determine if they are using different time servers.
Time Delta	Displays the difference in time between the local server and the selected server in the list.
	If the difference in time is more than a few minutes, it may indicate that the servers are using different time servers.

Merging Two Trees

For complete functionality of all menu options in DSMERGE, load DSMERGE on a server that contains the master replica.

If you don't know where the master replica is stored, you will be prompted with the correct server name when you attempt an operation that requires the master replica.

To perform a merge operation, you must load DSMERGE on the local tree.

When merging large trees, it's significantly faster to designate the source tree as the tree with fewer objects immediately subordinate to the [Root]. By doing this, you create fewer partition splits during the merge, since all target tree objects result in new partitions.



Because the local tree name no longer exists after the merge, you must update the "PREFERRED TREE" statement in the NET.CFG files of all local client workstations.

To minimize the number of client workstations you need to update, designate the tree with the most client workstations as the target tree, because the final tree retains the name of the target tree.


Or, rename the tree after the merge operation so that the final tree name corresponds to the tree with the greater number of client workstations attaching to it. See "Renaming the Tree" on page 385.

Prerequisites



- Access to the server console on the local tree or an established RCONSOLE session with that server
- The Supervisor object right to the [Root] object of both trees you want to merge
- A regular backup of the Directory Services database for the two trees
- All servers in both trees are synchronized and using the same time source
- (Optional) All servers in the tree are verified

Procedure



- 1. Load DSMERGE on the server where the master replica on the local tree is stored.**

At the server console prompt, type

```
LOAD DSMERGE <Enter>
```

If you don't know where the master replica is stored, you will be prompted with the correct server name when you attempt to merge the trees.

- 2. Choose "Merge two trees."**

The "Merge Tree Information" screen appears.

- 3. Enter the administrator name and password to log in to the local (source) tree.**

Log in as a user who has the Supervisor object right to the [Root] object on the local tree. Enter the typeless or typeful Distinguished Name, such as ADMIN.NOVELL or CN=ADMIN.O=NOVELL. Entering only ADMIN is invalid because it is not the complete name of the User object.

- 4. Choose "Target Tree" and select a target tree from the list servers in the "Available Trees" window.**

If the tree you want is not in the list, press <Insert> and enter the target tree's network address.

- 5. Enter the administrator name and password to log in to the target tree.**

- 6. Press <F10> to perform the merge.**

A message stating that the trees have been merged successfully is displayed.

Renaming the Tree

You must rename a tree if the two trees you want to merge have the same name.

You can rename only the local (source) tree name. To rename the target tree, run DSMERGE from a server on the target tree.

After you change a tree's name, you must update the clients' "PREFERRED TREE" statement in their NET.CFG files.

When you merge two trees, to minimize the number of client workstations that need to be updated, designate the tree with the most client workstations as the target tree because the final tree retains the name of the target tree.

Or, rename the tree after the merge so that the final tree name corresponds to the tree name with the majority of client workstations.


Another option is to rename the merged tree to the name of the original source tree. If you choose this option, then you must update the NET.CFG files on the target tree client workstations.

Prerequisites



- Access to a server console on the local tree or an established RCONSOLE session with the server
- The Supervisor object right to the [Root] object of the local tree
- (Optional) All servers in the tree are verified

Procedure



1. **Load DSMERGE on the server where a master replica of the [Root] partition is stored.**

At the server console prompt, type

LOAD DSMERGE <Enter>

If you don't know where the master replica is, load DSMERGE on any server in the local tree. Then you will be prompted with the correct server name when you attempt to rename a tree.

2. **Choose "Rename this tree."**
3. **Enter the administrator name and password to log in to the local (source) tree.**

Log in as a user who has the Supervisor object right to the [Root] object on the local tree. Enter your complete name, such as ADMIN.NOVELL or CN=ADMIN.O=NOVELL. Entering only ADMIN is invalid since it is not a complete name.

4. **Enter the new tree name.**
5. **Press <F10> to perform the rename.**

Completing the Tree Merge

Following the merging of two trees, it might be necessary to complete the following tasks:

1. (Optional) Choose “Checking Servers in the Tree” in the DSMERGE main menu to confirm that all tree names were changed correctly. See “Checking Servers in the Tree” on page 379.
2. Check the new partitions that the merge operation created. If you have many small partitions in the new tree, or if you have partitions that contain related information, you might want to merge them. See “Merging Partitions” on page 290.
3. Copy a new replica to any non-NetWare 4.1 servers after the merge is complete, if you did not upgrade before running DSMERGE.
4. Re-create any leaf objects or aliases at the [Root] that were deleted before you run DSMERGE.
5. Evaluate partitioning of the Directory tree.

Merging trees might significantly change replica placement on the source tree. You should carefully evaluate and change the partitioning as needed.

6. Update client workstations’ “PREFERRED TREE” statements in their NET.CFG files or rename the target tree.

Only the target tree name is retained after the merge.




To minimize the number of NET.CFG files you need to update, designate the tree with the most client workstations as the target tree because the final tree retains the name of the target tree.

Or, rename the tree after the merge operation so that the final tree name corresponds to the majority of the client workstations’ NET.CFG files. See “Renaming the Tree” on page 385.

The target [Root] object becomes the new [Root] for servers moved from the source tree. The Access Control List (ACL) for the [Root] object of the source tree are preserved. Therefore, the rights of the source tree's user ADMIN to the [Root] object are still valid.

After the merge is complete, both ADMIN users still exist and are uniquely identified by different container objects.



For security reasons, you might want to delete one of the two ADMIN User objects or restrict the rights of the two objects.

For more information about	Refer to
Object and property rights	"Object Rights" on page 18 and "Property Rights" on page 20
Partitions and replicas	"Creating and Managing Directory Services Partitions" on page 282
[Root] object	"The Directory Tree" and "[Root] Object" in Chapter 1 of <i>Introduction to NetWare Directory Services</i> "Root object," "Directory tree," and "Objects" in <i>Concepts</i>
Time synchronization	"Managing Network Time Synchronization" on page 593
DSMERGE	"DSMERGE" in <i>Utilities Reference</i>

Viewing and Managing NDS Synchronization Status

This section explains the NDS Trace feature that you can use from the server console. You should use this feature:

- ◆ To determine whether NDS synchronization processes are complete.
- ◆ To diagnose NDS errors. These errors may appear when you are manipulating NDS objects with the administration utilities. NDS errors also show up on the NDS Trace screen.

You can identify NDS-related system messages by their numbering: -601 through -699 and F966 through F9FE.




As with all NetWare system messages, an NDS system message does not necessarily indicate an error condition, but may simply indicate general NDS status. For more information, see the specific message in *System Messages*.

Prerequisites



- Access to the server console, or an established RCONSOLE session with the server

Procedure



1. Turn on the NDS Trace screen by typing the following at the server console:

```
SET NDS TRACE TO SCREEN = ON <Enter>
```

2. (Optional) To copy NDS Trace messages to a file, type

```
SET NDS TRACE TO FILE = ON <Enter>
```

By default, the messages are copied to the file DSTRACE.DBG in the SYS:SYSTEM directory.

To copy the messages to a different directory and/or file on volume SYS:, type

```
SET NDS TRACE FILENAME = path\filename <Enter>
```



Note

Copying NDS Trace messages to a file may be helpful if you need to ask someone else for diagnostic help.

3. To view NDS Trace screen messages, press <Alt>+<Esc> until you see the “Directory Services” screen.

Two types of messages are of particular interest:

- ◆ “All processed = YES” indicates that all pending NDS synchronization actions have been processed.
- ◆ Messages numbered –601 through –699 and F966 through F9FE indicate NDS status or errors. For explanations and suggested actions, see *System Messages*.

4. To disable tracing, type

SET NDS TRACE TO SCREEN = OFF <Enter>

For more information about	Refer to
NDS Trace options	“Directory Services Parameters” under SET in <i>Utilities Reference</i>
NDS system messages	System Messages



chapter

6

Migrating Data Using the High Capacity Storage System

What is the High Capacity Storage System?

The High Capacity Storage System (HCSS) is a utility that extends the storage capacity of a NetWare® server by integrating an optical disk library, or jukebox, into the NetWare file system.

HCSS moves files between the server's faster, but limited-capacity, storage devices (hard disks) and the slower, high-capacity storage devices (magneto-optical disks) in a jukebox.

This process, called data migration and demigration, is mostly hidden from the end user.

HCSS uses a NetWare volume to cache frequently used files, and manages the migration of less frequently used files to media in a jukebox.

Once a jukebox and the HCSS directories are set up, users and applications can access files stored on media by using the same NetWare commands and function calls with which they access files stored on the volume.

After the initial configuration, you can use HCSS in the NetWare Administrator graphical utility to

- ◆ Import and export media
- ◆ Format and reformat media
- ◆ Delete HCSS directories
- ◆ Change HCSS file storage parameters

HCSS also gives you the flexibility to physically transport directory and file data. You can move media to another jukebox, or you can store media on a shelf for archiving.

HCSS Concepts

To understand the information in this chapter, you should be familiar with the following entries in *Concepts*:

“Data migration”

“High Capacity Storage System”

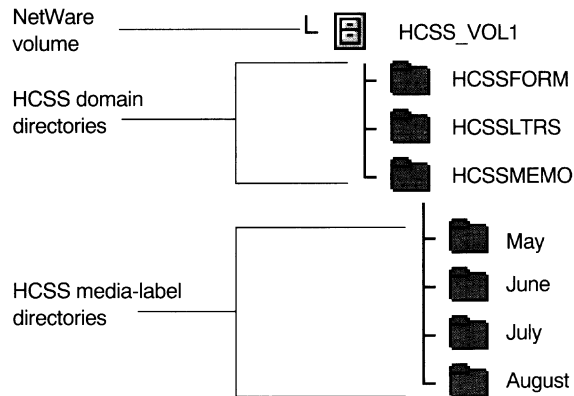
“Jukebox”

How the HCSS File System Is Structured

An HCSS directory looks like other directories in the NetWare file system. However, certain levels in an HCSS directory are created and handled differently.

Figure 6-1 illustrates these levels in a NetWare file system structure.

Figure 6-1
Sample HCSS File
System Structure



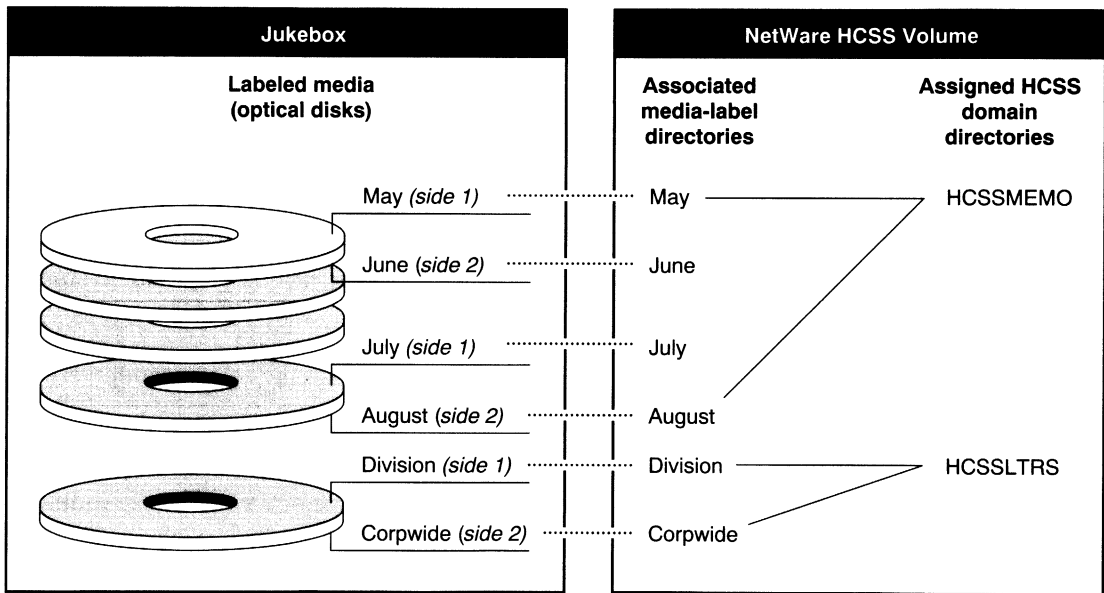
You create HCSS domain directories off the root of an HCSS volume. Within these directories, you create pairs of media-label directories. Each media-label directory corresponds to one side of media (that is, to one side of magneto-optical media that is labeled “Formatted Rewriteable Optical Disk”).



Creating HCSS directories and non-HCSS directories on the same volume can cause problems for HCSS users. Non-HCSS directories can fill the volume to its migration threshold and trigger migration inefficiently.

Figure 6-2 illustrates the association between media in the jukebox and media-label directories in the HCSS volume on the server.

Figure 6-2
Labeled Media and
Related Directories



An HCSS volume can contain an unlimited number of HCSS domain directories, and each domain directory can have several media-label directories (in multiples of two only, one media-label directory for every side of media assigned to that domain directory).

Each media-label directory can have multiple levels of subdirectories and files, like any other NetWare directory.

You are limited, however, to only one HCSS volume per NetWare server and to two media-label directories for each piece of media (that is, two for each optical disk) your jukebox will hold.



Although it is possible for HCSS domain directories and non-HCSS directories (regular NetWare directories) to coexist on the same volume, for management simplicity, we recommend you create and dedicate one volume exclusively for HCSS domain directories and their contents.

What Happens When Data Migrates

A file is considered migrated when the file data (everything inside of the file) has been moved from the volume to the corresponding labeled side of media in the jukebox.

File information, such as the file name and file size, remains on the volume so you can view the complete directory structure and access a file at any time.

The space the file data occupied on the NetWare volume is available after the file data has migrated, thus allowing more data to be stored in the volume. A 100MB volume, for example, has the potential to store many gigabytes as some of its files migrate to the jukebox.

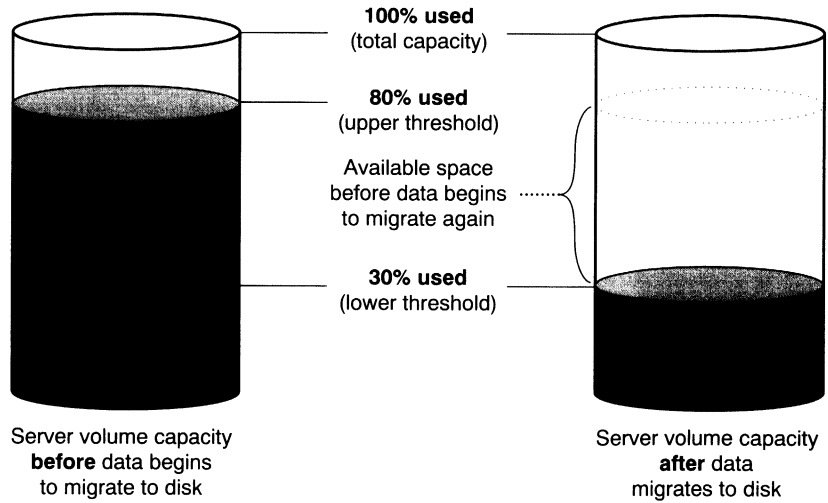
The triggers that start and stop migration are the upper and lower threshold parameters, which are set with HCSS commands.

The least-recently-used file data migrates when the upper volume-capacity threshold is reached or at a set time of day, and migration continues until the lower volume-capacity threshold is reached.

Figure 6-3 illustrates the upper and lower threshold concept. Assume that this is a 100MB volume and that its upper threshold is set to 80 percent and its lower threshold to 30 percent.

When 80 MB of the volume contain data, the oldest file data begins to migrate to the jukebox. The data continues to migrate until only 30 MB of data remain on the volume, making 50 MB of space on the volume available for more file data.

Figure 6-3
Volume Capacity
Before and After
Data Migration



This migration process might happen several times a day, depending on the size of the volume.

Migration occurs on a file-by-file basis according to the last time a file was accessed (least-recently-used files migrate first).

The migration process doesn't allow file data to flow from one side of media to the other. If one side of media is nearly full, you are notified so you can decide whether to delete or move files.



Note

For information on how to determine if a file has been migrated, see "NDIR" or "FLAG" in *Utilities Reference*.

Migrated file data stays on the media until the file is accessed; then the file data is automatically demigrated to the volume. Migrated files are always available, but file access may be slower.

Media Access Control

HCSS has options you can adjust to provide fair and efficient access to media in the jukebox. These options allow you to specify the amount of time a side of media is loaded in the jukebox drive to fill requests before the changer switches to another side of media.

The way HCSS controls media access is called elevator queuing. Elevator queuing for a jukebox works like an elevator in a tall building. Requests for any given floor come in sporadically. Rather than filling the requests in a first-come-first-served order, the elevator system determines the most efficient way to fill every request.

The number of people who can get on or off the elevator on any floor is determined by how long the doors remain open on a floor. If the elevator doors are closing and a person in the hallway pushes the elevator button, the doors will open again for a given period.

Likewise with HCSS, all requests for a side of media are queued by the system. Once the side of media is loaded into the drive, the system tries to fill as many requests as possible within a given time limit.

Media Export and Import Operations

HCSS can export (remove) media from the jukebox and later import (load) it into the same jukebox or into a different jukebox on another server with HCSS installed.

Before exporting a piece of media, HCSS attempts to migrate all file data under the two affected media-label directories to the media.

If all file data is migrated, the media, complete with the two media-label directories and files, is ejected from the jukebox while all references to the two directories are deleted from the volume.

If the export operation fails and the media is not ejected from the jukebox, it may be because the contents of the media-label directories can't fit on the media. You may need to move or delete files and try to export again.


When media is imported and assigned to an HCSS domain directory, the two media-label directories (including all subdirectories, files, and file attributes) are re-created in the domain directory just as they were before the media was exported.

Restrictions

The following restrictions apply when you use the HCSS utility.

- ◆ Supports only one volume per server.

Only one NetWare volume on the server where HCSS is installed can have the migration attribute set.



Do not have more than one volume with the migration attribute set on the server where HCSS is installed. This would compromise migration and may cause migration to fail.

- ◆ Restoring an HCSS volume requires completing a mixture of procedures using both the HCSS utility and the SBACKUP utility (or similar backup engine).

For more information, see “Restoring Migrated Files” in Chapter 9.

- ◆ Supports only one jukebox per server.
- ◆ Supports only the DOS name space.
- ◆ Requires an MS Windows client.
- ◆ Requires use of NetWare Administrator to delete media-label and domain directories.

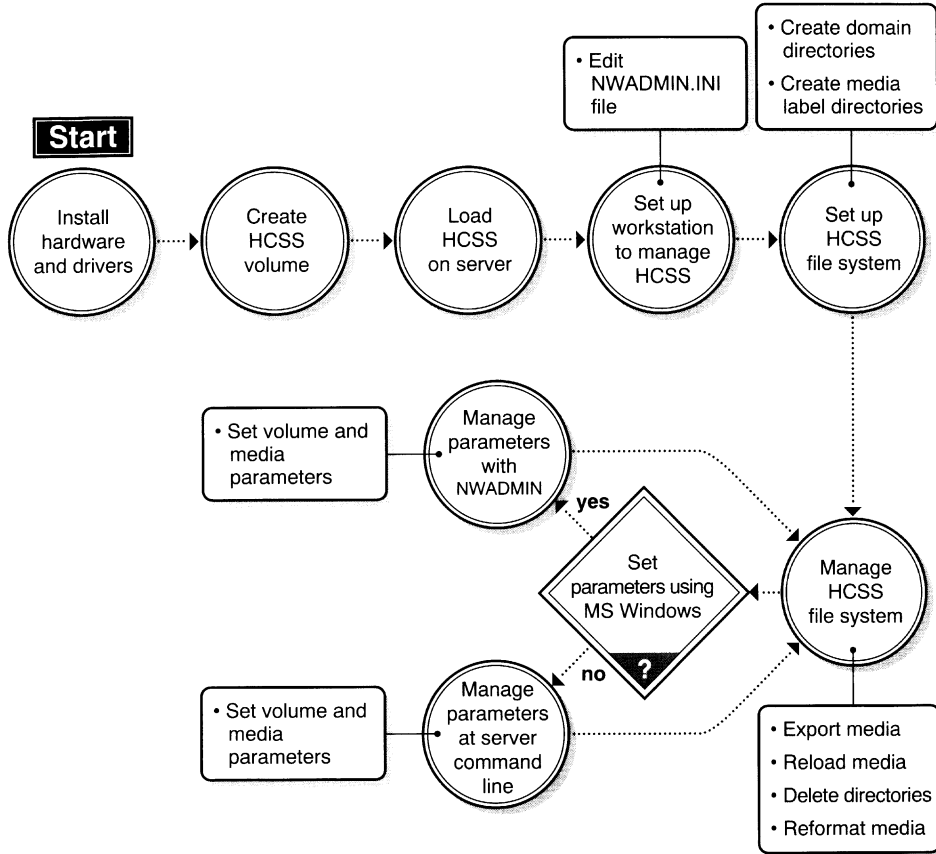
You can delete HCSS domain directories and media-label directories only by using the HCSS tools in the NetWare Administrator graphical utility.

- ◆ Doesn't support modifying directory attributes.
You shouldn't modify the directory attributes of an HCSS domain directory or any subdirectory under it.
- ◆ Doesn't support renaming of any subdirectory under a media-label directory.

Overview of HCSS Setup and Management

Figure 6-4 illustrates the tasks required to integrate HCSS with your regular NetWare® file system.

Figure 6-4
HCSS Tasks



How the HCSS File System Is Managed

After you install HCSS on the server, you can use the NetWare Administrator graphical utility (NWADMIN) at a workstation to perform all routine file system management tasks.

Table 6-1 summarizes how you must create, set up, and manage the levels of the HCSS directory structure.



Do not use other NetWare utilities to manage HCSS domain directories or media-label directories. Doing so could cause your file system to become corrupted.

Table 6-1
Managing the HCSS File System

File System Level	Management Method
HCSS volume	<p>Use the NetWare INSTALL module to create and modify volumes.</p> <p>Use HCSS commands either in NetWare Administrator or at the server console prompt (MSEngine prompt if you are running SFT III) to set data migration parameters for the volume.</p> <p>Requires Supervisor right, or equivalent, in the file system.</p>
HCSS domain directory on the volume	<p>Use HCSS commands in NetWare Administrator to create and manage HCSS domain directories.</p> <p>Requires Supervisor right, or equivalent, in the file system.</p>
HCSS media-label directory	<p>Use HCSS commands in NetWare Administrator to create and modify media-label directories.</p> <p>Requires Supervisor right, or equivalent, in the file system.</p>
Subdirectories below media-label directory	<p>Create and modify by using standard NetWare and DOS utilities and commands.</p> <p>Users with appropriate rights can work at this level the same as they do in the regular NetWare file system.</p>

Additional Information

For more information about	Refer to
Creating HCSS volumes	“Creating an HCSS Volume” on page 402 “Manage NetWare Volumes” in Chapter 3 of <i>Installation</i>
HCSS functionality	“High Capacity Storage System” and “Data migration” in <i>Concepts</i>
HCSS system messages	<i>System Messages</i>

Installing and Setting Up HCSS

Installing Hardware

The NetWare 4.1 release of HCSS requires the following:

- ◆ Hewlett-Packard 1718, 1715, or 1710 SCSI jukeboxes or compatible jukeboxes using the HPCHGR.DSK driver (which supports 512- or 1024- bytes/sector media)
- ◆ SCSI host bus adapters and drivers that support or emulate ASPITRAN
- ◆ Magneto-optical media that have been low-level formatted (labeled “Formatted Rewriteable Optical Disk” by the manufacturer)
- ◆ A workstation running MS Windows

Install the jukebox and its adapter board according to the documentation that came with the products. We recommend that the jukebox be the only device connected to the SCSI adapter board.


For more information on supported jukeboxes, call FAXBACK at 1-800-414-LABS.

Loading HCSS Drivers



If you are using NetWare 4.1 SFT III™ on your network, load drivers using the procedures under “Loading HCSS Drivers with SFT III” on page 438.

Procedure



1. Ensure that the appropriate hardware has been installed as recommended by the hardware manufacturer.

2. Load the applicable drivers.

Use the following syntax:

```
LOAD [path]controller_board_driver [option] <Enter>  
LOAD HPCHGR <Enter>
```

HPCHGR is Novell’s driver that ships with NetWare to operate the jukebox.

For example, if you are using an Adaptec 1740 board, type the following at the server console prompt, in the order given:

```
LOAD C:AHA1740 DEV_ENABLE=0 <Enter>  
LOAD HPCHGR <Enter>
```

The AHA driver autoloads ASPITRAN.DSK.



Some drivers require specifying an option (such as “DEV_ENABLE=0” in the previous example) to prevent the driver from taking control of the jukebox.

For more information about driver options, see the hardware manufacturer’s documentation.

For more information about loading drivers, see “Loading Disk Drivers” on page 559.

3. Type the following at the server console prompt:

```
SCAN FOR NEW DEVICES <Enter>
```

The SCAN FOR NEW DEVICES command checks for peripheral devices added since the server was last booted, and then registers the added devices with the operating system.

4. (Optional) Verify that the device was successfully loaded.

At the server console prompt, type

LIST DEVICES <Enter>

A list is displayed of all active peripheral devices on your server, including the magneto-optical drives in the jukebox.

Creating an HCSS Volume

An HCSS volume *must* be in the NetWare Directory Services™ database, and the HCSS volume *must not* be in the SYS: volume.



To simplify management of your HCSS system, we recommend that you create a dedicated HCSS volume for your domain directories and their contents.

Make the HCSS volume capacity at least 10 percent of the full capacity of the jukebox. Consider adding hard disk space or moving existing directories if necessary. You should provide enough room for frequently accessed files, or access time may be slower due to file migration and demigration.

An HCSS volume is created like any other NetWare volume, but it requires some specific settings.

Create an HCSS volume using the procedure detailed in “Creating Volumes” on page 515, but, before mounting it, do the following:

- ◆ Choose the volume and set the following parameters:

File Compression: OFF or ON (as appropriate)

Block Suballocation: OFF or ON (as appropriate)

Data Migration: ON

To change any of these parameters, select the parameter and press <Enter> to toggle between off and on.

- ◆ Choose the volume and set the volume block size, using one of these supported settings, to match your average file size:

Settings: 4 KB, 8 KB, 16 KB, 32 KB, or 64 KB

If Block Suballocation is ON, set “Volume Block Size” to 64 KB.

- ◆ Save all changes

Loading HCSS on the Server



Note If you are using SFT III on your network, load HCSS using the procedures under “Loading HCSS on the MEngine” on page 440.

Prerequisites



Checklist

- Jukebox and adapter installed according to the hardware manufacturer’s instructions
- HCSS volume created and mounted
- HCSS drivers loaded, and scan for new devices done (see “Loading HCSS Drivers” on page 401)

Procedure



Procedure

1. At the server console prompt, type

LOAD HCSS <Enter>

The HCSS module, as well as several other loadable modules, autoloads RTDM.NLM (a module that enables migration).

When HCSS is loaded, it takes inventory of the media in the jukebox and verifies a correlation between the media-label directories (directories created when importing media) and the media. Any media without corresponding media-label directories begins ejecting.

When HCSS has completed the inventory and verification, the message “HCSS is up and running” appears on the server console and is broadcast to user ADMIN. (The message doesn’t appear when HCSS is loaded for the first time.)



Note

When HCSS is loaded for the first time, it will not be up and running until you create the volume, domain directory, and media-label directories, in that order. (For information on the required flow of tasks, see “Overview of HCSS Setup and Management” on page 398.)

2. (Conditional) If the jukebox ejects media when you load HCSS, remove the media and, at the server console prompt, type the following:

MEDIA REMOVED <Enter>



Note


Enter this command for each piece of media that is ejected.

You can prevent more media from being ejected by typing the following command at the server console: HCSS Eject Media Override=On. See "Setting Parameters at the Server Console" on page 434.

3. (Optional) Add the commands to load the drivers and loadable modules to the server's AUTOEXEC.NCF file.

You can load the drivers and modules automatically every time the server is booted by adding the commands to your AUTOEXEC.NCF file.

For more information about adding commands to the .NCF file, see "Creating or Editing a Server Batch (.NCF) File" on page 459.



Important

Two .NCF files are included in the HCSS module for unloading and reloading HCSS when needed. Run HUNLOAD.NCF if it becomes necessary to unload HCSS. Run HRELOAD.NCF to reload (reinitialize and set up) HCSS after a fatal error has occurred.

4. At the workstation you plan to use to manage HCSS, edit the NWADMIN.INI file as described in the following section, "Setting Up a Workstation to Manage HCSS."

Setting Up a Workstation to Manage HCSS

Many HCSS management tasks can be done only through NetWare Administrator, which runs under MS Windows.

In order for HCSS to appear in the "Tools" pull-down menu in NetWare Administrator, you must modify the NWADMIN.INI file on the workstation where you will do these tasks.


Prerequisites



Checklist

- HCSS hardware and server software installed
- A volume created and set up for the HCSS file system
- A workstation running NetWare Administrator under MS Windows

Procedure



1. **At the workstation you plan to use to manage HCSS, log in as ADMIN to the server where HCSS is loaded.**
2. **At the same workstation, change to the MS Windows directory where the NWADMIN.INI file is stored.**

If you can't locate the NWADMIN.INI file, verify that you have successfully loaded NetWare Administrator in your MS Windows directory.

For more information, see *NetWare Client for DOS and MS Windows User Guide*.

3. **Using a text editor, open the NWADMIN.INI file and add the following lines under the [Snapin Object DLLs] section heading:**

```
Vis HCSSDecider=hdecsnap.dll
Vis HCSSObject=hobjsnap.dll
Vis HCSSMedia=hmedsnap.dll
```

4. **Save your changes and exit the NWADMIN.INI file.**
5. **Create HCSS directories, using the procedures in the following section.**

Setting Up the HCSS File System

To set up your HCSS file system:

1. Create one or more HCSS domain directories off the root of the HCSS volume. (See "Creating an HCSS Domain Directory" on page 406.)
2. Create media-label directories immediately below each domain directory. (See "Creating HCSS Media-Label Directories" on page 408.)



Because the domain directory and the media-label directory must be managed with HCSS commands, you might want to give them distinctive names to distinguish them from other NetWare directories.


Creating an HCSS Domain Directory

Prerequisites



- A workstation running NetWare Administrator with the NWADMIN.INI file modified for HCSS
- HCSS hardware and server software installed
- A volume created and set up for the HCSS file system
- You are logged in as ADMIN to the server on which HCSS is loaded

Procedure



- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the browser, select the HCSS volume where you want to create the HCSS domain directory.**

For information about moving around in the browser and selecting objects, choose “Help” from the menu bar and look up “Navigation.”

- 3. Verify that HCSS options appear in the “Tools” menu.**

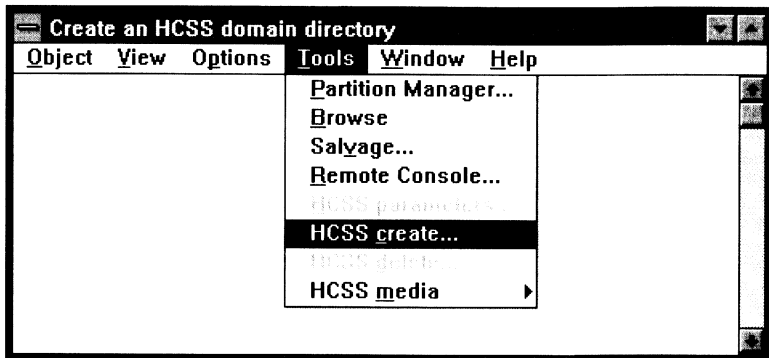
The options include “HCSS Parameters” and “HCSS Create.”

If the HCSS options do not appear in the “Tools” menu, verify that you have met the prerequisites listed for this procedure.

- 4. From the “Tools” menu, select “HCSS Create.”**

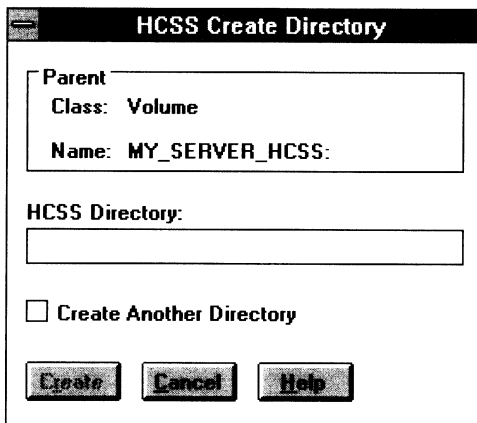
(See Figure 6-5.)

Figure 6-5
 Selecting the
 “HCSS Create”
 Option



The “HCSS Create Directory” dialog box appears. (See the example in Figure 6-6).

Figure 6-6
 “HCSS Create
 Directory” Dialog
 Box



5. Complete the “HCSS Create Directory” dialog.
 - 5a. In the “HCSS Directory” text box, type a unique name for the HCSS domain directory.

The name can be no longer than eight characters. Allowable characters are: uppercase A–Z and 0–9.
 - 5b. (Optional) If you want to create more than one HCSS domain directory, select “Create Another Directory.”

5c. Choose “Create.”

If you selected “Create Another Directory,” the dialog box remains on the screen so you can type in your next directory name.

If you did not select (or have unselected) “Create Another Directory,” you are returned to the browser, where the name or names of the new HCSS directories appear under the HCSS volume.

You can now create HCSS media-label directories, as described in the following section.

Creating HCSS Media-Label Directories

Prerequisites



- A workstation running NetWare Administrator with the NWADMIN.INI file modified for HCSS
- HCSS hardware and server software installed
- A volume created and set up for the HCSS file system
- You are logged in as ADMIN to the server on which HCSS is loaded

You create media-label directories when you assign labeled media to a domain directory.

To label or assign new or existing media, use the following table and determine which procedures to follow.

If you need to	Complete these procedures
Load new media (or media that you intend to overwrite) into the jukebox and file system	<ol style="list-style-type: none"> 1. "Importing Media" on page 409 2. "Formatting Media" on page 411 3. "Assigning HCSS Media to Domain Directories" on page 413
Import previously formatted media, with files and directories intact, into the jukebox and file system	<ol style="list-style-type: none"> 1. "Importing Media" on page 409 2. "Assigning HCSS Media to Domain Directories" on page 413


For an example of assigned media, see Figure 6-2.



Always use the "HCSS Import" option to load media; don't load it into the jukebox manually.

Importing Media

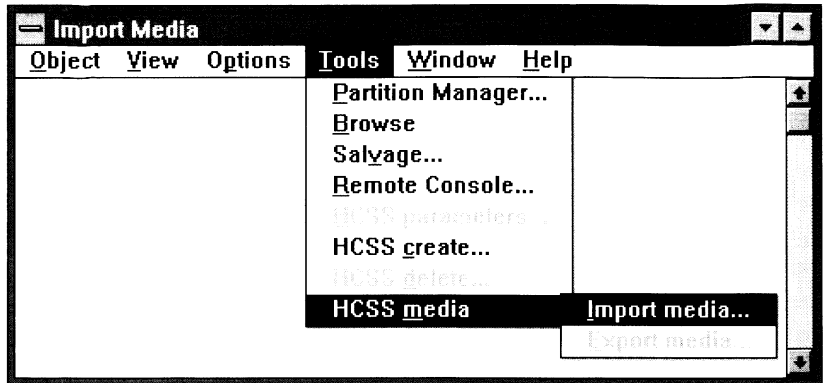
Procedure



1. **From the MS Windows Program Manager, choose the "NetWare Administrator" icon.**
2. **From the browser, choose an HCSS domain directory.**
3. **From the "Tools" menu, select "HCSS Media."**

(See Figure 6-7.)

Figure 6-7
**Selecting the
 “Import Media”
 Option**



4. From the “HCSS Media” menu, select “Import Media.”


The “Import” dialog box appears, prompting you to insert the media.



Suggestion

Before inserting media, label the cartridges with the names of the two media-label directories you plan to create.

5. Insert the media into the jukebox's mail slot.



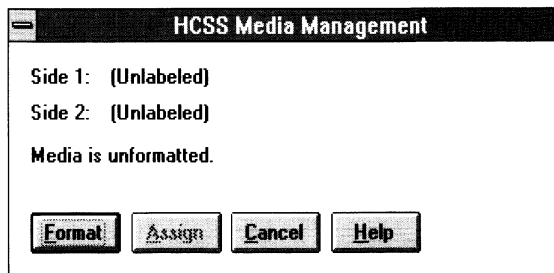
Important


Use only media labeled “Formatted Rewriteable Optical Disk.”

6. Choose “OK.”

The “HCSS Media Management” dialog box (Figure 6-8) appears, displaying the state of the media that was loaded (either “Media is unformatted” or “Media is formatted and unassigned”).

Figure 6-8
**“HCSS Media
 Management”
 Dialog Box**





Important

If the media is unformatted, you must format it before proceeding.

7. Choose a command from the dialog box and then go to the appropriate procedures, as shown in the following table:

If you need to	Choose	Then go to
Format or reformat media (when media is new or you plan to overwrite existing contents)	Format	“Formatting Media” on page 411
Create HCSS media-label directories by assigning media to an HCSS domain directory	Assign	“Assigning HCSS Media to Domain Directories” on page 413
Eject media from the jukebox	Cancel	Step 5 if you want to import media, or continue to select “Cancel” to return to the browser

Formatting Media

Before files can be stored on media, the media must be high-level formatted by HCSS.


Prerequisites



Checklist

- HCSS hardware and server software installed
- A volume created and set up for the HCSS file system
- You are logged in as ADMIN to the server on which HCSS is loaded
- A workstation running NetWare Administrator with the NWADMIN.INI file modified for HCSS
- Magneto-optical media imported into the jukebox (see “Importing Media” on page 409)

Procedure



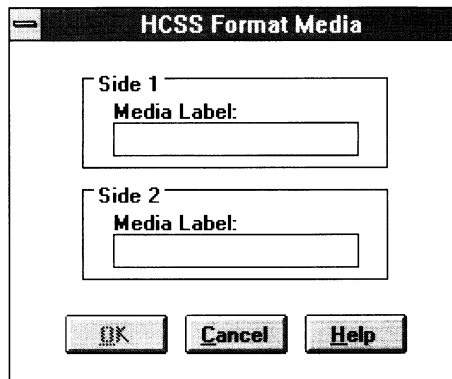
1. Ensure that you meet the prerequisites.
2. From the “HCSS Media Management” dialog box, choose “Format.”

A message appears, warning you about erasing data and asking if you are sure you want to format the media. When you select “Yes,” the dialog box in Figure 6-9 appears.



If you are reformatting media, the previous media labels appear as defaults in the “Media Label” boxes.

Figure 6-9
“HCSS Format
Media” Dialog Box




3. Enter the media label names.



Since labels should not be changed except by reformatting the media, you should plan your labels carefully.

Type a unique label for each side of media in the “Media Label” boxes for side 1 and side 2.



The labels must be different from all other labels in the HCSS volume.

Labels can be no longer than eight characters. The following character sets are allowed: uppercase A–Z and 0–9.

The media labels you enter here become the HCSS media-label directory names.

4. Select one of the options from the following table:

If you	Choose	Then go to
Want to accept the format, label the media, and assign the media-label directories	OK	“Assigning HCSS Media to Domain Directories” on page 413
Do not want to format or label the media	Cancel	Step 7 in “Importing Media” on page 409 (where you can exit through the “HCSS Media Management” dialog box)

Assigning HCSS Media to Domain Directories

The existence of HCSS media-label directories implies the existence of media labeled with the same name.

You can assign one or more pieces of media (and therefore media-label directories) under an HCSS domain directory. You are limited only to the number of media your jukebox will hold.

Procedure




1. Ensure that the “HCSS Assign Media” dialog box is on the screen.

Figure 6-10 shows an example of this dialog box.

If you have just imported or formatted media, the box should already be displayed. (If the box isn’t displayed, refer to the procedures in “Importing Media” on page 409 and make the appropriate selections until the box appears.)

Figure 6-10
“HCSS Assign
Media” Dialog Box



The two media labels you entered when you formatted the media appear in the "Media Label" list box.

All domain directories assigned to the HCSS volume are listed in the "HCSS Directory" drop box on the right side of the screen.

2. From the "HCSS Directory" drop box, select an HCSS directory.

3. Choose "OK."

The HCSS media-label directories are assigned to the HCSS domain directory you selected.

4. (Optional) To return to the "HCSS Media Management" dialog box, select "Cancel."

5. (Optional) To return to the browser, select "Cancel" one or more times.

As you return to the browser, the unassigned piece of media is ejected.


6. (Optional) To label media and assign other media-label directories to a domain directory, return to Step 3 in "Importing Media" on page 409.

Your HCSS directories and files should now be ready for migration.

Managing the HCSS File System

You might need to perform some or all of the following tasks to manage or maintain your HCSS file system:

- ◆ Export media
- ◆ Reload (import) media
- ◆ Delete an existing HCSS directory
- ◆ Reformat media



Always use the HCSS options in the NetWare Administrator utility to change HCSS domain directories and media-label directories.

Always use the “HCSS Import Media” and “HCSS Export Media” options in NetWare Administrator to load and unload media; do not manually load media into the jukebox until prompted by the software.

Never rename subdirectories once they are created under media-label directories.

Exporting Media

You must use HCSS commands to export (remove) media from the jukebox.

Exporting media removes the two associated HCSS media-label directories and their contents from the NetWare® file system. All files that were cached on the volume are first migrated to media, and then the media is ejected.

You might want to export media to

- ◆ Store media with its directories and files intact
- ◆ Assign specific media-label directories and data to a different HCSS domain directory
- ◆ Reformat media or move it to a jukebox on a different server


If the export operation fails, you get a message that HCSS cannot export the media.

The export operation will fail if:

- ◆ Files are open
- ◆ Files are not archived and the Migrate Unarchived Files parameter is turned off
- ◆ Files are flagged with the Execute Only file attribute
- ◆ Subdirectories under the media-label directory have the Don't Migrate attribute set
- ◆ There is not enough room on the media to hold all of the unmigrated files that are assigned to it

To make sure files are not in use before exporting media, use the SEND utility at the server console (or at the MSEngine prompt if you are running SFT III) to notify users that you are exporting media.

Procedure



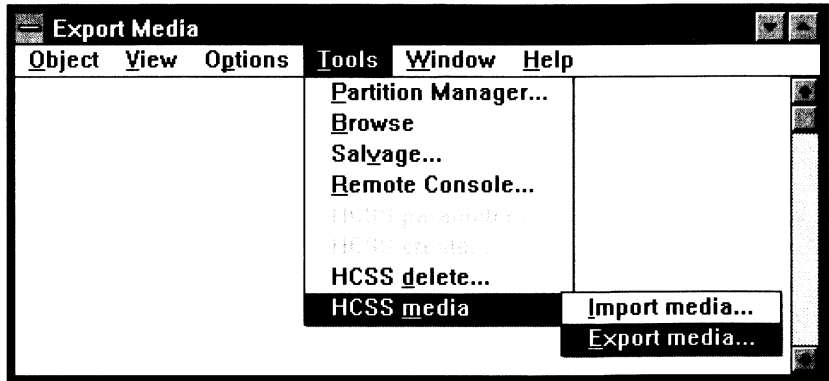
- 1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.**
- 2. From the browser, select the HCSS volume; then select the HCSS domain directory to which the media is assigned.**

For information about moving around in the browser and selecting objects, choose “Help” from the menu bar.

- 3. From the “Tools” menu, select “HCSS Media.”**
- 4. From the “HCSS Media” menu, select “Export Media.”**

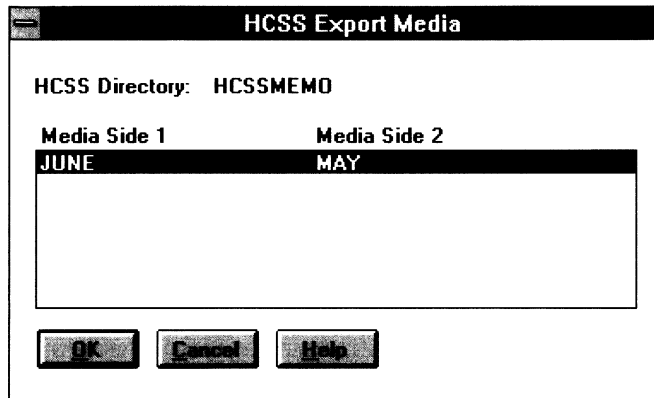
(See Figure 6-11.)

Figure 6-11
Selecting the
"Export Media"
Option



A dialog box similar to Figure 6-12 appears.

Figure 6-12
"HCSS Export
Media" Dialog Box



5. Select the pair of HCSS media-label directories for the media you want to export.
6. Choose "OK."

The names of the HCSS directories are removed from the browser. If the HCSS files are currently located on the server's hard disk, you must wait while they are migrated back to the media in the jukebox.

7. When the “Remove Media” prompt is displayed, remove the media from the jukebox.

Both the directory structure and the file data are removed from the file system.

8. Choose “OK.”

The dialog box disappears and you are returned to the browser.


Reloading Media

You can reload the files on exported media, and add the corresponding HCSS media-label directories and contents to your file system, by importing the media with the NetWare Administrator utility.

To import previously formatted media into the jukebox with files and directories intact, follow the procedures in “Creating HCSS Media-Label Directories” on page 408.


Deleting an HCSS Domain Directory

Deleting an HCSS domain directory removes it from the NetWare file system.



You must use NetWare Administrator to delete an HCSS domain directory. Using a command line delete or some other delete utility will confuse the jukebox inventory and might cause media to begin ejecting from the jukebox.

Procedure



- 1. Export all media whose labels appear as HCSS media-label directories under the HCSS domain directory you are deleting.**

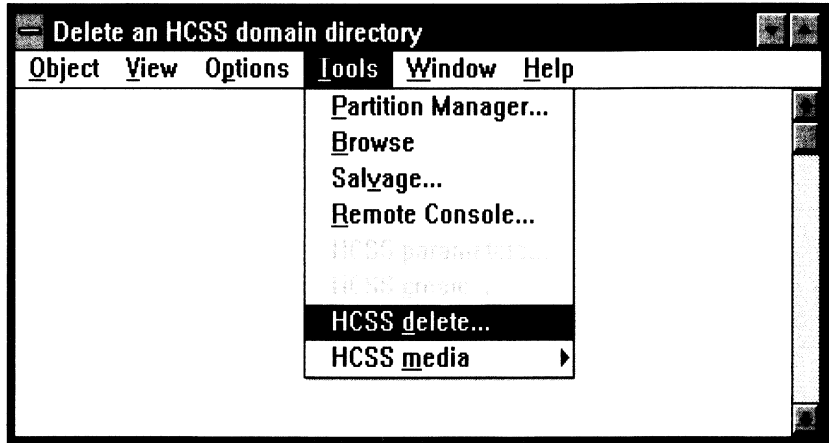
For instructions, see “Exporting Media” on page 415.

- 2. Delete or move any non-HCSS files and directories from the HCSS domain directory, using standard NetWare commands.**

If you had both HCSS and non-HCSS files in the same domain directory, the non-HCSS files are probably those remaining after you have exported all media in the domain directory.


3. From the browser, select the name of the HCSS domain directory you want to delete.
4. From the “Tools” menu, select “HCSS Delete” (Figure 6-13).

Figure 6-13
Selecting the
“HCSS Delete”
Option



The “HCSS Delete Directory” dialog box appears. (See the example in Figure 6-14.)


Figure 6-14
“HCSS Delete
Directory” Dialog
Box



5. Choose “Delete.”

The HCSS domain directory name is removed from the browser and the file system.

Reformatting Media



Warning Reformatting media erases all files and directories stored on the media.

You might want to reformat media to

- ◆ Erase all files and directories stored on the media
- ◆ Create new media labels (HCSS media-label directory names)
- ◆ Assign media to an HCSS directory or eject media and reserve it for future use

To reformat media:

1. Perform the procedures under “Importing Media” on page 409.
2. Perform the procedures under “Formatting Media” on page 411.
3. (Optional) If you want to assign media to domain directories at this time, perform the procedures under “Assigning HCSS Media to Domain Directories” on page 413.

Checking Media Status

HCSS provides a command you can use to check media status and verify that migration has occurred.

To check HCSS media status, from the server console prompt (MSEngine prompt if you are running SFT III), type the following:

```
HCSS Media Status = [path]
```

The path should include the volume name.

For example:

```
HCSS Media Status = HCSSVOL:\domain_dir\media-  
label_dir
```

After you enter the command, the following information about the specified media appears:

- ◆ **Total Bytes:** The total capacity of media, in bytes, for a media-label directory
- ◆ **Used Bytes:** How much space is currently being used in bytes
- ◆ **Migrated Files:** The number of files that have been migrated to this side of media

Managing HCSS Parameters

You can customize your HCSS volume and define how the volume and jukebox interact by setting various HCSS parameters. You can set or change these parameters either through MS Windows, using NetWare Administrator, or at the server console prompt (MSEngine prompt if you are running a NetWare 4.1 SFT III™ system).



Note

All HCSS parameters have default settings. You don't need to change these settings unless you want to customize HCSS operations.

(For information on the default settings, see “436Setting Parameters Using NetWare Administrator” on page 427 or “Setting Parameters at the Server Console” on page 434.)

What the Parameters Do

Several parameters work together and allow you to specify

- ◆ If and when data migration occurs (see “Migration If/Then Parameters” on page 422)
- ◆ How media requests are handled (see “Media and Media Request Parameters” on page 423)
- ◆ Which decision-making features are used, either by you or by HCSS (see “Decision-Making Parameters” on page 425)

Migration If/Then Parameters

The following HCSS parameters let you specify if and exactly when data will migrate to the jukebox.

Disable Migration

Turn on this parameter when you need to disable migration for any reason—such as reducing demands on the processor or jukebox.

Migrate Unarchived Files

Turn on this parameter to allow migration of files that have not been backed up. By default this parameter is turned off, which means that files won't migrate until they have been backed up.



Note

If you attempt to export media with the Migrate Unarchived Files parameter turned off, and the associated files are unarchived, the media will not be exported.

Migrate Compressed Files Only

If this parameter is set, HCSS migrates only those files that have been compressed (files flagged with "Co") or files that can't be compressed (files flagged with "Cc"). Files that can't be compressed are those that would n't save any space if they were compressed.

For more information on file flags, see "FLAG" in *Utilities Reference*.

Upper Threshold

This parameter specifies how full the volume can become before file data begins migrating to the jukebox. This value can be set from 1% to 100%, but it must be higher than the lower threshold.

The upper threshold should be set so that the most frequently used files remain cached on the volume.

Lower Threshold

This parameter specifies how full the volume can end up after migration ends. The value is a percentage of full, and it must be lower than the upper threshold percentage.

Migrate to Lower Threshold Hour/Minute

This parameter specifies a time when data will migrate, regardless of whether the upper threshold is reached.

You can disable this parameter by setting the hour and minute to exactly midnight. Migration will still be initiated when the volume capacity reaches the upper threshold.

Media and Media Request Parameters

Eject Media Override

Turn on this parameter at the server console prompt to send a message to HCSS to stop media from ejecting.

This parameter can be set only from the server console prompt.

When HCSS is loaded, it performs a jukebox inventory that detects any inconsistencies between the media-label directories and the labeled media in the jukebox. HCSS will attempt to eject any media without corresponding media-label directories.

If the jukebox begins ejecting media when you load HCSS, and you believe the media belongs in the jukebox

1. Turn this parameter on, allow HCSS to finish loading, and then manually place the media back in the jukebox.
2. Unload HCSS and the jukebox driver (HPCHGR) and determine the status of the volume and media-label directories.
3. Correct the problem (for instance, if your volume wasn't mounted, mount it now) and reload the jukebox driver and then HCSS.

Delete Through

Setting this parameter turns off asynchronous delete requests (where response to the request is not synchronized with the actual deletion of files from media), and files are deleted from media before the delete request is completed.

Normally, when you delete a file, the file is placed in a delete queue and the delete request is completed. It is not actually deleted from the media until other, higher priority requests have been completed (such as demigrations and renaming migrated files).

Use this parameter if, for security purposes, you want to delete all the way through to the media and not be placed in a delete queue, even if the delete request takes longer.

Minimum Time In Drive

This parameter works in conjunction with the Maximum Time in Drive parameter and the Request Idle Time parameter, and it specifies the minimum amount of time that the media remains in the drive for servicing (or waiting to service) requests for the loaded side of media.

If the media is not already in the drive, part of the Minimum Time In Drive time is spent moving it into place.

During the Minimum Time In Drive time, the media remains in the drive even if there are no additional requests for that side of media.

Maximum Time In Drive

This parameter works in conjunction with the Minimum Time in Drive parameter and the Request Idle Time parameter, and it specifies the maximum amount of time media remains in the drive before being removed to allow a request on another side of media.

However, the media will remain in the drive indefinitely if no other sides of media are needed.

The Maximum Time In Drive time must be greater than the Minimum Time In Drive time.

Request Idle Time

This parameter works in conjunction with the Maximum Time in Drive parameter and the Minimum Time In Drive parameter, and it defines a short grace period to wait for additional requests for files on the media loaded in the drive.

This grace period begins after the Minimum Time In Drive time has expired.

If the grace period expires before a request is received, the media may be removed to service a request on another side of media.

If a request is received during this grace period, the grace period begins again. This process continues until the Maximum Time In Drive time is reached.

The Request Idle Time must be less than the difference of the Maximum Time In Drive and Minimum Time In Drive settings.

Decision-Making Parameters

HCSS periodically scans the HCSS domain directories and their subdirectories for migratable files, sorts the files by the least-recently-used (LRU) file, and builds an LRU list. The first file on the LRU list (the oldest file on the list) is the first file migrated when either the upper threshold or the preset migration time is reached.

If you have a large number of files in the HCSS volume and HCSS is dominating your hard drive for an unacceptable period while building the LRU list, you can reduce the impact by doing one or both of the following.


- ◆ Setting or decreasing the Marked Files Limit parameter (default = no limit)
- ◆ Increasing the Marking Frequency parameter

Marked Files Limit

This parameter is used to limit the number of files on the LRU list. When the limit is reached, HCSS stops scanning the HCSS volume for migratable files.

If, for example, a file system had 10,000 migratable files with no Marked Files Limit (the setting was 0), the LRU list would have 10,000 files as candidates for migration. If the same system had a Marked Files Limit of 5000, then only the first 5000 migratable files found would go on the list. After the first 5000 files had migrated, the next list generated should include the remaining 5000 migratable files.

Placing a limit on the LRU list means that the first file migrated is the least recently used of those on the list and *not* necessarily the least recently used file on the volume.



Set this number to half of the number of existing migratable files and try it in a working environment. If your hard disk is tied up for an unacceptable period, reduce the limit.

Marking Frequency

Use this to specify how often HCSS builds an LRU file list. The larger this number, the longer HCSS will rest between building LRU lists.

Polling Frequency

This parameter specifies how often you want HCSS to check if the upper threshold has been reached, indicating that it is time to migrate.

Remaining Capacity Before Warning

This parameter specifies at what capacity you want to be notified that a side of media is almost full. For example, if you set the value to 20%, you will be notified when the media is 80% full.

Warning Frequency

This parameter specifies how often you want to be warned that sides of media are almost full. The warning is a broadcast message to user ADMIN.

436 Setting Parameters Using NetWare Administrator

For more information on HCSS parameters, see “What the Parameters Do” on page 421.


Prerequisites



Checklist

- HCSS hardware and server software installed
- A volume created and set up for the HCSS file system
- You are logged in as ADMIN to the server on which HCSS is loaded
- A workstation running NetWare Administrator with the NWADMIN.INI file modified for HCSS

Procedure



Procedure

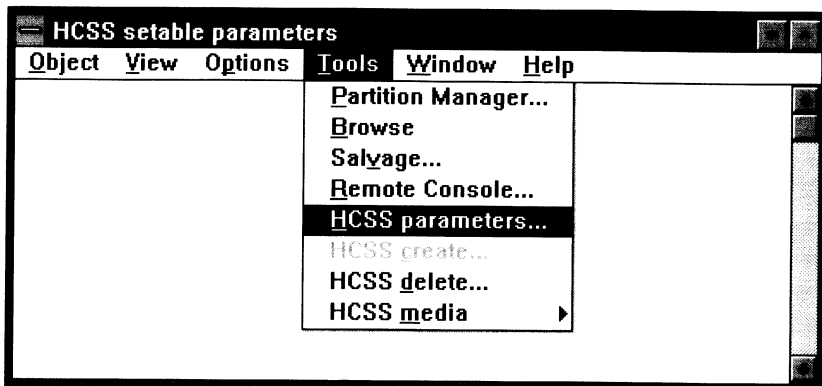
1. From the MS Windows Program Manager, choose the “NetWare Administrator” icon.
2. From the browser, select the HCSS volume.

For information about moving around in the browser and selecting objects, choose “Help” from the menu bar.

3. From the “Tools” menu, select “HCSS Parameters.”

(See Figure 6-15.)

Figure 6-15
Selecting the
“HCSS
Parameters” Option



The “HCSS Parameters” dialog box appears. (See Figure 6-16.)

Figure 6-16
The “HCSS
Parameters” Dialog
Box

HCSS Parameters

Volume Name: GIDDY_HCSS:

Disable Migration:

Migrate Unarchived Files

Migrate Compressed Files Only

Capacity Thresholds

Upper Threshold: 80 % Full

Lower Threshold: 50 % Full

Migrate to Lower Threshold Time: 3:00 AM Off

Media

Delete Through to Media

Maximum Time in Drive: 30 Seconds

Minimum Time in Drive: 20 Seconds

Request Idle Time: 2 Seconds

Advanced

Remaining Capacity Before Warning: 20 % Remaining

Polling Frequency: 1 Minutes

Marking Frequency: 30 Minutes

Warning Frequency: 2 Minutes

Marked Files Limit: 0

OK Cancel Help

4. In the “HCSS Parameters” dialog box, make changes to some or all of the parameters, as needed.

See Table 6-2 for information on how to change individual parameter settings.



Note

Parameter settings marked with an asterisk (*) in the table are persistent (that is, the setting stays the same unless you change it). The nonpersistent parameter settings reset to the default setting every time HCSS is loaded.

You can create an .NCF file to load and set the nonpersistent parameters to your desired settings. (For more information about creating an .NCF file, see “Creating or Editing a Server Batch (.NCF) File” on page 459.)

5. When you have finished setting the parameters, choose “OK.”

The new settings are activated. (Choosing “Cancel” returns you to the browser and does not save the new settings.)

Table 6-2

How to Set HCSS Parameters Using NetWare Administrator

Parameter	Purpose	Parameter Setting Information
Disable Migration	Turns migration on or off. (For more information, see “Disable Migration” on page 422.)	Select (place an X in the box) to turn off migration. Deselect (remove X from box) to reenable migration. Default setting: disabled (no X).
Migrate Unarchived Files	Allows migration of files that have not been backed up. (For more information, see “Migrate Unarchived Files” on page 422.)	Select (place an X in the box) to allow migration of unarchived files. Deselect (remove X from box) to prevent migration of unarchived files. Default setting: disabled (no X).
Migrate Compressed Files Only	Allows only compressed files to be migrated. (For more information, see “Migrate Compressed Files Only” on page 422.)	Select (place an X in the box) to enable the parameter. Only compressed files, or files that can’t compress, will be migrated. Deselect (remove X from box) to disable the parameter. Default setting: disabled (no X).

Table 6-2 *continued*

How to Set HCSS Parameters Using NetWare Administrator

Parameter	Purpose	Parameter Setting Information
Upper Threshold*	Sets the upper threshold of volume space that can be used before files are migrated to media. (For more information, see "Upper Threshold" on page 422.)	Click on the up or down arrow until you find the setting you want. Set anywhere from 1 to 100%, but higher than the lower threshold. Default setting: 80 (migration begins when the volume has 20% free volume space remaining).
Lower Threshold*	Sets the lower threshold of volume space at which files stop migrating to media. (For more information, see "Lower Threshold" on page 422.)	Click on the up or down arrow until you find the setting you want. Set anywhere from 0 to 99%, but lower than the upper threshold. Default setting: 50 (migration stops when the volume has 50% free volume space remaining).
Migrate to Lower Threshold Time*	Sets the time when you want files to begin migrating automatically. (For more information, see "Migrate to Lower Threshold Hour/Minute" on page 423.)	Click on the up or down arrow until you find the setting you want or select the "Off" box to disable this parameter. Select a time when file access is least likely to occur, but not to exactly midnight, unless you want to disable this parameter. Set the migration time to exactly 12:00 a.m. to disable the parameter. Default setting: 3:00 a.m. (migration begins at that time).

Table 6-2 *continued*

How to Set HCSS Parameters Using NetWare Administrator

Parameter	Purpose	Parameter Setting Information
Delete Through to Media	<p>Turns off asynchronous deletes, allowing deletion of data from the media to synchronize with the response to a requested delete</p> <p>(For more information, see "Delete Through" on page 423.)</p>	<p>Select (place an X in the box) to enable the parameter, which allows data to be deleted from media before the delete request completes.</p> <p>Deselect (remove X from box) to disable the parameter.</p> <p>Default setting: disabled (no X).</p>
Maximum Time In Drive	<p>Sets the maximum amount of time (in seconds) a side of media will remain active before changing to another side of media</p> <p>(For more information, see "Maximum Time In Drive" on page 424.)</p>	<p>Select this parameter and type in the time (in seconds).</p> <p>Supported settings: 0 to 3600, but must be greater than the Minimum Time In Drive time.</p> <p>Can be set for 1 hour if converted into seconds (for example, 1 hour=3600.)</p> <p>Default setting: 30 seconds.</p>
Minimum Time In Drive	<p>Sets the minimum amount of time (in seconds) a side of media will remain active before changing to another side of media.</p> <p>(For more information, see "Minimum Time In Drive" on page 424.)</p>	<p>Select this parameter and type in the time (in seconds).</p> <p>Supported settings: 0 to 3600, but must be less than Maximum Time In Drive.</p> <p>Default setting: 20 seconds.</p>

Table 6-2 *continued*

How to Set HCSS Parameters Using NetWare Administrator

Parameter	Purpose	Parameter Setting Information
Request Idle Time	<p>Sets the grace period during which additional requests for the media loaded in the drive will be accepted.</p> <p>(For more information, see “Request Idle Time” on page 424.)</p>	<p>Select this parameter and type in the time (in seconds).</p> <p>Supported settings: Any fraction of the difference between the Maximum Time In Drive and Minimum Time In Drive times.</p> <p>Default setting: 2 seconds.</p>
Remaining Capacity Before Warning	<p>Specifies at what media capacity you want to be notified that a side of media is filling up.</p> <p>This works in conjunction with the Warning Frequency parameter to determine when you are notified.</p> <p>(For more information, see “Remaining Capacity Before Warning” on page 426.)</p>	<p>Click on the up or down arrow until you find the setting you want.</p> <p>Supported settings: 0 to 100 percent.</p> <p>Default setting: 20.</p>
Polling Frequency*	<p>Specifies how often (in minutes) you want HCSS to check the thresholds you have set.</p> <p>(For more information, see “Polling Frequency” on page 426.)</p>	<p>Set this anywhere from 1 minute to 24 days (days converted into minutes).</p> <p>Supported settings: 1 to 34560.</p> <p>Default setting: 1.</p>
Marking Frequency*	<p>Specifies how often (in minutes) HCSS builds a least-recently used (LRU) file list.</p> <p>(For more information, see “Marking Frequency” on page 426.)</p>	<p>Set this anywhere from 1 minute to 24 days (days converted into minutes).</p> <p>Supported settings: 1 to 34560.</p> <p>Default setting: 30.</p>

Table 6-2 continued

How to Set HCSS Parameters Using NetWare Administrator

Parameter	Purpose	Parameter Setting Information
Warning Frequency*	<p>Sets how often (in minutes) you want to be warned that the volume has exceeded its upper capacity threshold.</p> <p>(For more information, see "Warning Frequency" on page 426.)</p>	<p>Set this anywhere from 1 minute to 24 days (days converted into minutes).</p> <p>Supported settings: 1 to 34560.</p> <p>Default setting: 2.</p>
Marked Files Limit	<p>Limits the number of migratable files that are placed on the LRU list.</p> <p>(For more information, see "Marked Files Limit" on page 425.)</p>	<p>Supported settings: Any number smaller than your total number of migratable files.</p> <p>Default setting: no limit (0).</p>

Setting Parameters at the Server Console


For an overview of all the HCSS parameters, see “What the Parameters Do” on page 421.

Prerequisites



- HCSS hardware and server software installed
- A volume created and set up for the HCSS file system
- You are logged in as ADMIN to the server on which HCSS is loaded
- If you are running an SFT III™ system, you are logged in as ADMIN on the MEngine. (See “Using HCSS with SFT III” on page 438.)

Procedure



1. **(Optional) View the list of HCSS parameters and their current settings.**

At the server console prompt, type

```
HCSS <Enter>
```

2. **(Optional) Change any HCSS parameter, as needed, using the following syntax:**

```
HCSS [parameter] = [setting]
```

Replace *parameter* and *setting* with a parameter name and the applicable setting found on Table 6-3.



Note

Parameters marked with an asterisk (*) in the table are persistent (that is, the setting stays the same unless you change it). The nonpersistent parameters reset to the default setting every time HCSS is loaded.

You can create an .NCF file to load and set the nonpersistent parameters to your desired settings. (For more information about creating an .NCF file, see "Creating or Editing a Server Batch (.NCF) File" on page 459.)

Table 6-3
How to Set HCSS Parameters at the Server Console

Parameter Syntax	Supported Settings	Purpose of Parameter
(no parameter)	(no setting)	Displays a list of HCSS parameters and current settings.
Eject Media Override = <i>setting</i>	ON, OFF. Default: OFF.	Stops media in the jukebox from being ejected. (For more information, see "Eject Media Override" on page 423.)
Delete Through = <i>setting</i>	ON, OFF. Default: OFF.	When turned off, expedites requests for deletes. (For more information, see "Delete Through" on page 423.)
Migrate Unarchived Files = <i>setting</i>	ON, OFF. Default: OFF.	Allows migration of files that have not been backed up. (For more information, see "Migrate Unarchived Files" on page 422.)
Migrate Compressed Files Only = <i>setting</i>	ON, OFF. Default: OFF.	Allows only files that are compressed (or that can't compress) to be migrated. (For more information, see "Migrate Compressed Files Only" on page 422.)
Maximum Time In Drive = <i>setting</i>	0 to 3600, but must be greater than the Minimum Time In Drive setting. Default: 30 (seconds).	Sets the maximum amount of time (in seconds) a side of media will remain active before changing to another side of media. (For more information, see "Maximum Time In Drive" on page 424.)

Table 6-3 *continued***How to Set HCSS Parameters at the Server Console**

Parameter Syntax	Supported Settings	Purpose of Parameter
Minimum Time In Drive = <i>setting</i>	0 to 3600, but must be less than the Maximum Time In Drive setting. Default: 20 (seconds).	Sets the minimum amount of time (in seconds) a side of media will remain active before changing to another side of media. (For more information, see “Minimum Time In Drive” on page 424.)
Request Idle Time = <i>setting</i>	Set this to a fraction of the difference between the Maximum Time In Drive and Minimum Time In Drive settings. Default: 2 (seconds).	Sets the grace period during which additional requests for the media loaded in the drive will be accepted. (For more information, see “Request Idle Time” on page 424.)
Migration = <i>setting</i>	ON, OFF. Default: ON	Turns migration on or off. (For more information, see “Disable Migration” on page 422.)
Upper Threshold = <i>setting</i>	1 to 100.* Default: 80 (percent).	Sets the upper threshold of volume space that can be used before files are migrated to media. (For more information, see “Upper Threshold” on page 422.)
Lower Threshold = <i>setting</i>	0 to 99.* Default: 50 (percent).	Sets the lower threshold of volume space at which files stop migrating to media. (For more information, see “Lower Threshold” on page 422.)
Migrate to Lower Threshold Hour = <i>setting</i>	0 to 23.* Default: 3 (3:00 a.m.).	Sets the hour when files begin migrating automatically. When set to exactly 0 (the equivalent of 12:00 a.m.), the parameter is disabled. (For more information, see “Migrate to Lower Threshold Hour/Minute” on page 423.)

Table 6-3 *continued*

How to Set HCSS Parameters at the Server Console

Parameter Syntax	Supported Settings	Purpose of Parameter
Migrate to Lower Threshold Minute = <i>setting</i>	0 to 59.* Default: 0 (12:00 a.m.).	Sets the minute of the hour when files begin migrating automatically. When set to exactly 0 (the equivalent of 12:00 a.m.), the parameter is disabled. (For more information, see “Migrate to Lower Threshold Hour/Minute” on page 423.)
Remaining Capacity Before Warning = <i>setting</i>	10 to 100. Default: 20 (percent).	Specifies when you want to be alerted about the volume’s remaining capacity. (For more information, see “Remaining Capacity Before Warning” on page 426.)
Polling Frequency = <i>setting</i>	1 to 34560.* Default: 1 (minute).	Specifies how often (in minutes) you want HCSS to check the upper and lower thresholds. (For more information, see “Polling Frequency” on page 426.)
Marking Frequency = <i>setting</i>	1 to 34560.* Default: 30 (minutes).	Specifies how often (in minutes) a list of the least recently used directories that are managed by HCSS will be built. (For more information, see “Marking Frequency” on page 426.)
Warning Frequency = <i>setting</i>	1 to 34560.* Default: 2 (minutes).	Specifies how often (in minutes) you will be warned that the volume has exceeded its upper capacity threshold. (For more information, see “Warning Frequency” on page 426.)
Marked Files Limit = <i>setting</i>	The total number of files flagged for migration, or any number smaller than the total. Default: 0 (no limit).	Limits the number of migratable files that are scanned for the LRU list with any setting other than zero (0). (For more information, see “Marked Files Limit” on page 425.)

Using HCSS with SFT III


A NetWare 4.1 SFT III™ system uses two mirrored NetWare servers to provide uninterrupted network services to clients. If one server fails, its mirrored partner automatically takes over. (For more information, see “System Fault Tolerance” in *Concepts*.)

If you are using HCSS with an SFT III system

- ◆ Use the following two sections—“Loading HCSS Drivers with SFT III” and “Loading HCSS on the MEngine”—in place of the sections titled “Loading HCSS Drivers” and “Loading HCSS on the Server” in this chapter.
- ◆ Substitute “MEngine prompt” for “server console prompt” where noted in this chapter.

Loading HCSS Drivers with SFT III

Procedure



1. **Ensure that the appropriate hardware has been installed as recommended by the hardware manufacturer.**
2. **At the IOEngine console, load the applicable drivers.**

Use the following syntax:

```
LOAD [path]controller_board_driver [option]  
    <Enter>  
LOAD HPCHGR <Enter>
```

HPCHGR is the driver that ships with NetWare to operate the jukebox.

For example, if you are using an Adaptec 1740 board, type the following at the MEngine prompt, in the order given:

```
LOAD C:AHA1740 DEV_ENABLE=0 <Enter>  
LOAD HPCHGR <Enter>
```

The AHA driver autoloads ASPITRAN.DSK.



Some drivers require specifying an option (such as “DEV_ENABLE=0” in the previous example) to prevent the driver from taking control of the jukebox.

For more information about driver options, see the hardware manufacturer’s documentation.

For more information about loading drivers, see “Loading Disk Drivers” on page 559.

3. Type the following at the MEngine prompt:

SCAN FOR NEW DEVICES <Enter>

The SCAN FOR NEW DEVICES command checks for peripheral devices added since the server was last booted, and then registers the added devices with the operating system.

4. (Optional) Verify that the device was successfully loaded.

At the MEngine prompt, type

LIST DEVICES <Enter>

A list is displayed of all active peripheral devices on your server, including the magneto-optical drives in the jukebox.

5. Return to the section titled “Creating an HCSS Volume” on page 402.

Loading HCSS on the MEngine


Load the HCSS loadable modules.

Prerequisites



- HCSS volume created and mounted
- Applicable drivers loaded at the IOEngine console

Procedure



1. At the MEngine prompt, type

LOAD HCSS <Enter>

The HCSS module (as well as several other loadable modules) autoloads RTDM.NLM (a module that enables migration).

When HCSS is loaded, it takes inventory of the media in the jukebox and verifies a correlation between the media-label directories (directories created when importing media) and the media. Any noncorresponding media begins ejecting.

When HCSS has completed the inventory and verification, the message “HCSS is up and running” appears on the MEngine console and is broadcast to user ADMIN. (The message doesn’t appear when HCSS is loaded for the first time.)



When HCSS is loaded for the first time, it will not be up and running until you create the volume, domain directory, and media-label directories in that order. (For information on the required flow of tasks, see “Overview of HCSS Setup and Management” on page 398.)

2. (Conditional) If the jukebox ejected media when you loaded HCSS, remove the media and, at the MEngine prompt, type the following:

MEDIA REMOVED <Enter>

Enter this command for each piece of media that is ejected.



You can prevent more media from being ejected by typing the following command at the MEngine prompt: HCSS Eject Media Override=On. See “Setting Parameters at the Server Console” on page 434.

3. (Optional) Add the commands to load the drivers and loadable modules to the MEngine's MSAUTO.NCF file.

You can load the drivers and modules automatically every time the MEngine is booted by adding the commands to your MSAUTO.NCF file.

For more information about adding commands to the .NCF file, see "Creating or Editing a Server Batch (.NCF) File" on page 459.



Two .NCF files are included in the HCSS module for unloading and reloading HCSS when needed. Run HUNLOAD.NCF if it becomes necessary to unload HCSS. Run HRELOAD.NCF to reload (reinitialize and set up) HCSS after a fatal error has occurred.

If you are running SFT III and the server to which the jukebox is attached fails, you must run HRELOAD after the mirrored servers resynchronize.

4. At the workstation you plan to use to manage HCSS, edit the NWADMIN.INI file as described in "Setting Up a Workstation to Manage HCSS" on page 404.



chapter

7

Maintaining the NetWare Server

This chapter contains procedures designed to help you manage, optimize, and protect NetWare[®] servers. Procedures are grouped into the following sections:

- ◆ Common Management Tasks—including basic procedures such as downing a server and creating server batch files
- ◆ Managing Server Connections—including procedures to monitor connection problems and manage communications buffers
- ◆ Monitoring and Optimizing the Server—including ways to manage memory, improve performance, and monitor errors
- ◆ Maintaining Volumes—including mounting and dismounting, creating, deleting, and renaming volumes
- ◆ Managing Server Hard Disks—including instructions for file compression, adding new disks, and mirroring or duplexing hard disks
- ◆ Disaster Prevention and Recovery—including procedures to prevent packet forgery, establish UPS monitoring, and protect the system's memory
- ◆ Managing Network Time Synchronization—including how to configure time synchronization for large networks
- ◆ Using Remote Console to Manage a Server—including procedures for executing remote console sessions over both a direct connection and a modem
- ◆ Administering Accounting—including instructions for setting up accounting, calculating charge rates, and viewing accounting totals

Common Management Tasks

This section describes typical network supervisor tasks and explains ways you can change the server environment you created during installation of the NetWare 4™ operating system.

Sending Console Messages to Workstations

To send a message from the server console to all workstations or to a single user, use the SEND command.

You can still use the BROADCAST command, as you may have done with previous versions of NetWare. The functionality is exactly the same as that of the SEND command.



Unlike the client SEND command, *groupname* is not a valid parameter at the server.

Sending a Message to a Specific User

Procedure



- 1. Make sure you are at the server console prompt (:).**

If another screen is displayed (such as MONITOR), press <Alt>+<Esc> to return to the console prompt.

- 2. Type a message, or a series of messages, using the SEND command with the “TO” option.**

Messages can be up to 55 characters.

For information on SEND command syntax, see “SEND” in *Utilities Reference*.

- 3. Enter the username or connection number in the SEND command.**

To determine the username or connection number of a user you want to send a message to, use the “Connection Information” screen in MONITOR.

The form of the user's name depends on whether the user is logged in as a NetWare Directory Services™ object or as a bindery object, and on where the master replica database is stored.

For example, there are three ways you could send a message to JSMITH, depending on the way the way the name appears in "Connection Information" in MONITOR.

- ◆ If the name appears as "JSMITH," type the following:

```
SEND "REMEMBER THE MEETING." TO JSMITH <Enter>
```

- ◆ If the name appears as "CN=JSMITH.OU=SALES.O=ABC_INC," type the following:

```
SEND "REMEMBER THE MEETING." TO  
CN=JSMITH.OU=SALES.O=ABC_INC <Enter>
```

- ◆ If JSMITH's connection number in the "Connection Information" display is "2," type the following:

```
SEND "REMEMBER THE MEETING." TO 2 <Enter>
```

Sending a Message to All Workstations

Procedure



1. Make sure you are at the server console prompt (:).
2. Type a message, or a series of messages, using the SEND command.

Messages can be up to 55 characters.

For example, to send a message to all workstations, type:

```
SEND "SERVER GOING DOWN" <Enter>
```

The message is sent to all workstations.

Additional Information

For more information about	Refer to
Using the SEND utility	"SEND" in <i>Utilities Reference</i>
Receiving and clearing messages at the workstation	<i>NetWare Client for DOS and MS Windows User Guide</i> <i>NetWare Client for OS/2 User Guide</i>

Bringing Down a Server

Use the DOWN utility at the server console to ensure data integrity before turning off power to the server.

DOWN ensures data integrity by writing all cache buffers to disk, closing all files, and updating the appropriate Directory Entry Tables and File Allocation Tables.




Note

Bringing down a server during a remote console session terminates the session.

If you bring down the server but do not exit to DOS, the server is still connected to the network, and it still receives packets. You can continue to execute console commands, such as TRACK ON and TRACK OFF, that deal with packets.

Procedure



Procedure

- 1. At the server console prompt, use SEND to inform all users that they should close all files and log out.**

For example, type

```
SEND "Please close files and log out from server  
servername." <Enter>
```

```
SEND "Going down in 5 minutes." <Enter>
```

- 2. Bring down the server by typing**

```
DOWN <Enter>
```

If files are still open, the console displays the files that are open, the user who opened them, and the workstation connection number.

After executing the DOWN command, you may need to go to the workstation with open files and close them for the user if the workstation is not attended.

You may now either return to DOS or restart the server.

3. (Optional) To return to DOS, type

EXIT <Enter>

4. (Optional) To restart the server without returning to DOS, type

RESTART SERVER <Enter>

Booting the Server with the EXIT Command

The following procedure explains how to use the EXIT command to reboot the server. This procedure allows users who are using RCONSOLE to reboot the NetWare server remotely.

Prerequisites



- Commands in the server AUTOEXEC.BAT to execute SERVER.EXE. (If these commands are not in the AUTOEXEC.BAT file, this procedure reboots the computer and displays the DOS prompt, but does not load the NetWare server software.)

Procedure



- 1. At the server console prompt, remove DOS from memory by typing**

REMOVE DOS <Enter>

- 2. Bring down the server by typing**

DOWN <Enter>

- 3. Reboot the server by typing**

EXIT <Enter>

This reboots the computer, executes the SERVER.EXE file, and displays the server console prompt.

Loading a NetWare Loadable Module

This procedure explains how to load a NetWare Loadable Module™ program (NLM) from the server console.

Prerequisites




- Determine the memory protection domain you want the NLM™ to run in. For a list of NetWare NLM programs and the domains they can run in, see Appendix C, “Memory Domains for NLMs and Autoloaded NLMs,” of *Utilities Reference*.

You do not need to set up a memory domain if you want the NLM to run in the OS domain. If you want the NLM to run in a memory domain outside of the OS domain, see “Protecting the Operating System’s Memory” on page 581.

- If you are running a message file for the NLM in a language other than English, make sure the language is specified properly. See “Specifying a Language for an NLM” on page 466.
- If the NLM is not in the SYS:SYSTEM directory, either move the NLM to the SYS:SYSTEM directory or add a server search path to the NLM. Otherwise, you must enter the path for the NLM when you execute the LOAD command. See “Viewing and Adding Server Search Paths” on page 454.

Procedure



1. At the server console prompt, enter the LOAD command.

- ◆ Use this format to load an NLM residing in SYS:SYSTEM and running in the OS domain:

```
LOAD NLM <Enter>
```

- ◆ Use this format to load an NLM with a specified memory domain:

```
LOAD NLM DOMAIN=domain_name <Enter>
```

- ◆ Use this format to load an NLM located on the A: drive:

```
LOAD A:NLM <Enter>
```

2. If you want the NLM to be loaded whenever the server boots, add the LOAD command to the AUTOEXEC.NCF file.

Additional Information

For more information about	Refer to
Using the LOAD utility	“LOAD” in <i>Utilities Reference</i>
Loading NLM programs in different domains	“Protecting the Operating System’s Memory” on page 581
Memory protection	“Memory protection” in <i>Concepts</i>

Loading and Binding LAN Drivers

After you add a network board to your NetWare server, you must load and bind the corresponding LAN driver. Loading a LAN driver establishes a network connection (if the server is physically connected to the network cabling). Binding a LAN driver links a network protocol to the driver and the network board.

The LAN driver you choose depends on the cabling system and the network board you are using. Refer to Table 7-1 to determine which driver to load.

Table 7-1
NetWare LAN Drivers

Cabling System	Network Board	LAN Driver
ARCnet*	RX-Net™	TRXNET.LAN
	RX-Net II	
	RX-Net/2™	
Ethernet	NE/2™	NE2.LAN
	NE/2T	
	NE/2-32™	NE2_32.LAN
	NE1000™ - ASSY 950-054401	NE1000.LAN
	NE1000 - ASSY 810-160-001	
	NE2000™ - ASSY 810-149	NE2000.LAN
	NE2000T - ASSY 810-000220	
	NE2100™ - ASSY 810-000209	NE2100.LAN
	NE2100 - ASSY 810-000214 (twisted-pair version of NE2100)	NE1500T.LAN
	NE3200™	NE3200.LAN
NE32HUB™	NE32HUB.LAN	
IBM* PC network	PCN2 broadband	PCN2L.LAN
	PCN2/A broadband	
	PCN baseband	
	PCNA baseband	
Token ring	NTR2000	NTR2000.LAN

To load and bind LAN drivers, you can use

◆ The INETCFG NLM

Use INETCFG to load LAN drivers and bind them to any supported protocol. For more information, see “INETCFG” in *Utilities Reference*.

◆ The LOAD and BIND commands


If you know the parameters required by the communication protocol, you can use the LOAD and BIND commands to load and bind LAN drivers at the command line. For more information, see “LOAD” and “BIND” in *Utilities Reference*.

◆ The INSTALL NLM

You can use INSTALL to load LAN drivers and bind them to the default protocol IPX™.

The following procedure explains how to use INSTALL to load a LAN driver and bind it to IPX. To bind a LAN driver to another protocol, use INETCFG.

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Choose “Driver Options” from the “Installation Options” menu.

3. Select “Configure Network Drivers” from the “Driver Options” menu.

4. Choose “Select an Additional Driver” from the “Additional Driver Actions” menu.

The system displays a list of available drivers.

5. Choose the driver you want to load or, if the driver is not on the list, press <Insert>.

If you press <Insert> to load an unlisted driver, follow the screen prompts. If you select a listed driver, continue with the following steps.



For some drivers, a message may appear indicating that the driver must be loaded manually (at the console prompt). To load these drivers, follow the screen prompts or press <F1> for more information.

6. Choose “Select/Modify Driver Parameters.”

The system displays a window where you can set values for driver parameters. If you selected an NE2000 driver, the system also displays a window containing protocol options. The cursor is active in the protocol options window, if it is displayed.

7. If the window containing protocol choices is displayed, accept the default IPX protocol, by pressing the downarrow until the cursor moves to the parameter window.



The protocol choices window also lists TCP/IP and AppleTalk. These are nonrouting protocols which you may select instead of IPX. To configure standard TCP/IP, AppleTalk, or other non-IPX protocols, use INETCFG.

8. Enter parameter values in the fields on the parameter window.

Press <F1> for help if necessary.

In some cases, the system displays a pop-up list of values for the field from which you select the desired value. In other cases, you must type in a value and press <Enter> to move to the next field.

9. (Optional) To specify a frame type other than 802.2 for an Ethernet driver, press <F3> to display a selection list of frame types. Select the frame type from the list and press <Enter>.

Ethernet drivers for NetWare 4 default to the Ethernet 802.2 frame type. If you want compatibility with earlier versions of NetWare, load the LAN driver with frame type 802.3, as well.

10. When finished, press <F10> to save the values and exit the window.

The system loads the LAN driver and then displays a confirmation window containing the command line to bind IPX to the LAN driver with the specified frame type.

At this point, you can either confirm binding of the protocol with the specified frame type or change the frame type.

11. To bind the protocol, press <Enter>. To display the command line with a different frame type, press <F3>.

Each time you press <F3>, the system displays the command line with another of the available frame types. Press <F3> until the desired frame type is displayed.

11a. When the frame type you want is displayed, press <Enter> to bind the LAN driver.

INSTALL automatically places the LOAD and BIND commands in the AUTOEXEC.NCF file.


Additional Information.

For more information about	Refer to
AppleTalk protocol	"AppleTalk protocols" in <i>Concepts</i>
Frame types	"Ethernet configuration" in <i>Concepts</i>
IPX protocol	"IPX" in <i>Concepts</i>
TCP/IP protocol	"TCP/IP" in <i>Concepts</i>

Viewing and Adding Server Search Paths

Adding a server search path allows you to run server utilities and NetWare Loadable Module (NLM) programs and to access server batch files (.NCF files) without specifying the full path to them. Whenever you run a utility or NLM from the server, the server searches for the utility in the current directory and also in the path you specified as a server search path.

Procedure



1. To view current search paths, type the following at the server console prompt:

```
SEARCH <Enter>
```

The system displays the current search paths in a numbered list. For example,

```
Search 1: [Server path] SYS:SYSTEM
```

```
Search 2: [Server path] A:
```

```
Search 3: [Server path] VOL1:NCF
```

2. To add a search path, use one of the following formats:

- ◆ Add a directory or drive as a search path by typing

```
SEARCH ADD path <Enter>
```

For example, this command specifies a search path to VOL1:NCF

```
SEARCH ADD VOL1:NCF <Enter>
```

This command specifies a search path to drive A:

```
SEARCH ADD A: <Enter>
```

- ◆ Add a search path in a particular place in the numbered list of search paths by typing

```
SEARCH ADD number path <Enter>
```

For example, the following command both creates the search path to drive A: and places the path in the second position in the numbered list of search paths.

SEARCH ADD 2 A: <Enter>

The search path that was previously number 2 becomes number 3, and so on.

3. To delete a search path, use the following format:

SEARCH DEL *number* <Enter>

Replace *number* with the number of the search path you want to delete.

For more information, see "SEARCH" in *Utilities Reference*.


Installing, Uninstalling, and Configuring a Server Product

The INSTALL NLM allows you to install, uninstall, and configure other NetWare products such as name space modules and Storage Management Services™ (SMS).

If you have questions during the installation or configuration of a product, refer to the documentation that came with the product.

Some products may not have configuration or uninstall options. A message explaining this is displayed if you try to perform unsupported operations.

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Choose "Product Options" from the "Installation Options" menu.

3. Choose “View/Configure/Remove an Installed Product” from the “Other Installation Actions” menu.

The “Currently Installed Products” list appears. If no products are currently installed, the list is empty.

4. Do one of the following:

- ◆ To install a product, press <Insert> and follow the prompts.
- ◆ To uninstall a product, select the product from the list and press <Delete>.
- ◆ To set configuration options for a product, select the product from the list and follow the prompts.

Copying NetWare Files


When you do a simplified NetWare install, the system installs some, but not all, of the NetWare files. If desired, you can copy the rest of the installation files later by using the “Copy Files” option of the INSTALL NLM. Follow this procedure.

Prerequisites



- NetWare 4 installation CD-ROM

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Select “Copy Files Option” from the “Installation Options” menu.

The system displays a box containing the default source path from which the files will be copied.

3. Press <Enter> to accept the default path, or press <F3> to specify a different path.

If you press <F3>, the system displays a prompt with the current source path. Backspace to delete the path and type in the new path. Press <Enter>.

The system displays the prompt "Specify a server boot path (where SERVER.EXE will be):" This is the destination directory—the location to which the files will be copied. This directory must be the boot directory—the directory where SERVER.EXE is located. The default path is C:\NWSERVER.

4. Accept the default or type in a new path if your bootstrap directory is not C:\NWSERVER.

The system displays a list of file groups. An X next to each group means the group is preselected.

5. Delete the X next to any group(s) you do not want to copy.

To delete the X, use the arrow keys to move the cursor to the desired group and press <Enter>. The X next to the box disappears.

To reselect a group, move the cursor to the group and press <Enter> again.

6. When finished selecting groups, press <F10>.

The system copies the files.

Extracting NetWare Files from the Installation CD-ROM

If a NetWare system file is accidentally deleted and purged from the network, you can use the NWXTRACT utility to extract the file from the NetWare 4.1 installation CD-ROM or diskettes.

NWXTRACT copies selected files from the installation CD-ROM or diskettes to their default locations on the network or to a user-specified destination.


Prerequisites



- A workstation running DOS 3.30 or later
- Rights to copy files to the destination directory

- The NetWare 4.1 installation CD-ROM or diskettes
- Name space support, if extracting Macintosh*, UNIX*, or OS/2* files

Procedure



1. **If using diskettes, load the diskette containing the file or files you want to extract.**
2. **To extract a file or group of files, type**

```
NWTRACT path filename | groupname [destination]
           [option] <Enter>
```

Replace *drive* with the drive letter for the installation CD-ROM or the diskette.

Replace *filename* with the name of the file you want to extract from the CD-ROM or diskette. Or, replace *groupname* with the name of the group of files you want to extract.

Replace *destination* with the destination path for the extracted files, if different from the default destination.

Replace *option* with any of the following:

Option	Description
<i>/S=server</i> <i>server object</i>	Copies extracted files to their default location on the specified server.
<i>/T=type</i>	Specifies the type of file to be extracted. Valid types are DOS, MAC, OS2, SER (server), UNIX, and WIN.
<i>/?</i>	Displays online help.

3. **To make sure the extracted system files are flagged Read Only, change to the destination directory and type**

```
FLAG filename RO <Enter>
```

See "FLAG" in *Utilities Reference* for information on flagging files.


Creating or Editing a Server Batch (.NCF) File

When a server is booted, SERVER.EXE executes two files:

- ◆ **STARTUP.NCF**—This file loads the server’s disk drivers and some SET parameters.
- ◆ **AUTOEXEC.NCF**—This file stores the server name and IPX internal network number, loads the LAN drivers and settings for the network boards, binds the protocol to the installed drivers, and loads other NLM programs.

The following procedure explains how to edit either of these files.

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Choose “NCF Files Options” from the “Installation Options” menu.

The system displays a menu containing options to create or edit both the STARTUP.NCF and AUTOEXEC.NCF files.

3. Select the desired option and press <Enter>.

If you select an option to edit one of the files, the selected file appears in a window.

If you select an option to create one of the files, the system displays the existing file in one window and a new file with default entries in another window, so you can compare the two.


The default entries include only essential contents of the file and reflect the current system setup. For example, the default entries include time zone information and load and bind commands for currently installed LAN drivers. The new file does not include custom configuration settings you may have in the original file.

The cursor is active in the new file. Press <Tab> to toggle between the two files.

4. Edit the new or existing file as necessary.

Editing and navigation keystrokes are listed at the bottom of the screen. For additional help, press <F1>.

5. When finished, save the file by pressing <F10> and selecting "Yes" when prompted. To exit without saving, press <Alt>+<F10>.



Each time you edit the AUTOEXEC.NCF file or the STARTUP.NCF file, NetWare saves the previous version as AUTONCF.OLD or STARTNCF.OLD.



Changes to the file take effect once you reboot the server.

Editing Text Files from the Server Console


Although you usually edit server files from a workstation, sometimes it is more convenient to edit text files from the server console.

The following procedure explains how to use the EDIT NLM to edit files. You can use EDIT to edit DOS text files on either DOS or NetWare partitions.



EDIT allows you to edit only DOS text files. Each time you save the file, you can increase the file size by up to 4 KB.

Procedure



1. At the server console prompt, type

LOAD EDIT *pathname* <Enter>

If you do not specify a file pathname, the system prompts you to enter the pathname. Type the pathname of the file to be edited and press <Enter>.

The file must be located on a mounted volume or a local drive of the server.

If you enter only the filename, the system assumes a default directory location of SYS:SYSTEM.

If the file exists, it is displayed on the screen, ready for editing.

If the file does not exist, EDIT displays a prompt that asks whether the file should be created. Respond as described in Step 2.

2. Choose “Yes” to create the new file or “No” to redisplay the pathname prompt.

A blank screen appears if this is a new file.

3. Edit the file as needed.

The keys used to edit text are described in Table 7-2.

Table 7-2
Editing Keys Available in the EDIT NLM

Key	Function
Arrow keys	Move up/down one line or left/right one character.
Backspace	Delete character to left of cursor.
<Ctrl>+<PageUp>	Move to beginning of document.
<Ctrl>+<PageDown>	Move to end of document.
<Delete>	Delete the character at the cursor location or delete blocked text.
<End>	Move to end of line.
<Esc>	Exit the file. You are prompted to save any changes.
<F5>	Highlight (block) text for copying or deleting. Toggles highlighting on and off.
<F6>	Copy highlighted (blocked) text to a buffer.
<Home>	Move cursor to beginning of line.
<Insert>	Insert text from buffer.
<PageUp>	Move up one page
<PageDown>	Move down one page.
<Tab>	Move cursor four spaces to the right.

4. **When finished editing, press <Esc>. Select “Yes” to save the file or “No” to exit the file without saving changes.**

The system redisplay the pathname prompt.

5. **To edit another file, enter the file pathname. To redisplay the console prompt, press <Esc>.**

Setting a Server’s Language

The NetWare 4.1 operating system and its NLM programs and utilities use the language the server was installed in, unless you specify otherwise. The server languages supported for the initial release of NetWare 4.1 are French, Italian, German, Spanish, and English.

This section explains how to specify a new server language. The first step is to copy the new language files to the server. The second step is to set the server to the new language.

Subsequent sections explain how to change the language for the NLM programs you load on the server and how to change the keyboard type.

Copying the Language Files to the Server


Before you can change the server language, you must copy the language files to the server from the installation CD-ROM or another server. Follow this procedure to copy the files.

Prerequisites



- Installation CD-ROM or access to another NetWare server containing the language files

Procedure



1. **At the server console prompt, type**

LOAD INSTALL <Enter>

2. **From the “Installation Options” menu, select “Product Options.”**

- 3. From the “Other Installation Actions” menu, select “Choose an Item or Product Listed Above.”**
- 4. From the “Other Installation Items/Products” menu, select “Install an Additional Server Language.”**

The system displays a box containing the path from which you last installed files. In most cases, this is the path to the language files for the *current* language. For example, if you last installed NetWare files in English, this is the path to the English files.

You should accept the path to the current language files so that help and system messages are displayed in the current language during the install procedure.

If, for some reason, the path is not to the current language files, then you should enter the path to the current language files. See the next step.

- 5. Press <Enter> to accept the displayed path, or press <F3> and type in the correct path to the current language files.**

The path to the current language files has this syntax:

```
SERVER | CD:\NW410\INSTALL\current_language
```

If you are installing the files from another server, the system may prompt you for a username and password for that server. Follow the prompts.

The system displays a window listing groups of language files.

- 6. Use the arrow keys to move the cursor to the language you want to install and press <Enter>.**

The system marks the selected language with an “X.” You can specify several languages, if desired.

To deselect a language, move the cursor to the language and press <Enter> again.

- 7. When you have finished selecting languages, press <F10>.**

The system copies the selected language files.

After the files are copied, you can change the server language at any time. See the next section, “Changing the Server’s Language.”

Changing the Server's Language

Follow this procedure to change the server language.



Note

This procedure requires you to reboot the NetWare server because you set the server language when executing SERVER.EXE. You may want to perform this procedure after business hours.

Prerequisites



Checklist

- Language files copied to the server. (See “Copying the Language Files to the Server” on page 462.)

Procedure



Procedure

1. **At the server console prompt, type**

LOAD INSTALL <Enter>
2. **From the “Installation Options” menu, select “Product Options.”**
3. **From the “Other Installation Actions” menu, select “Choose an Item or Product Listed Above.”**
4. **From the “Other Installation Items/Products” menu, select “Change Server Language.”**

The system displays a box containing the path from which you last installed files. In most cases, this is the path to the language files for the *current* language. For example, if you last installed NetWare files in English, this is the path to the English files.

You should accept the path to the current language files, so that help and system messages are displayed in the current language until you have completed the change to the new language.

If, for some reason, the path is not to the current language files, then you should enter the path to the current files. See the next step.

5. Press <Enter> to accept the displayed path, or press <F3> and type in the correct path to the current language files.

The path to the current language files has this syntax:

SERVER | CD:\NW410\INSTALL\current_language

If the current language files are on another server, the system may prompt you for a username and password for that server. Follow the prompts.

The system displays a menu listing the available languages.

6. Choose the desired language from the menu.

You must choose a language that you installed earlier with the “Install an Additional Server Language” option of INSTALL. (See “Copying the Language Files to the Server” on page 462.)

If this is the first time you have changed the language, the system displays a prompt specifying the default destination where the language files are to be placed. The default location is the bootup directory, C:\NWSERVER.

The language files must be placed in your bootup directory.

7. Accept the default or type a new path if your bootup directory is not C:\NWSERVER.

Once you specify the bootup directory, the system loads the language files.

8. Reboot the server to change the server language.


Additional Information

For more information about	Refer to
Setting language for a workstation	<i>NetWare Client for DOS and MS Windows User Guide</i> <i>NetWare Client for OS/2 User Guide</i>
Languages for NetWare	“International use of NetWare 4.1” in <i>Concepts</i>

Specifying a Language for an NLM

Use the LANGUAGE command to change the language for subsequently loaded NLM programs. The LANGUAGE command does not change the language for currently loaded modules. .

For consistent language display, you should also change the server language as explained in the previous section, "Setting a Server's Language," on page 462.



If you change the NLM language without changing the server language, the NLM still displays some character strings and screen titles in the server language.

To see a list of languages, type

LANGUAGE LIST <Enter>

The system displays a list of languages and their ID numbers. (Not all languages in the list are supported. Additional languages may be available in the future.)

To determine the currently specified NLM language, type


LANGUAGE <Enter>

Prerequisites



- The name or number of the language you want to use. Use the LANGUAGE LIST command to determine the correct name and number. English, French, Italian, German, and Spanish - Latin America are supported with the initial release of NetWare 4.1.

Procedure



1. To change the language for subsequently loaded NLM programs, type

```
LANGUAGE language_number | language_name <Enter>
```

For example, to change the language to German, type

```
LANGUAGE 7 <Enter>
```

or

```
LANGUAGE GERMAN <Enter>
```


Additional Information

For more information about	Refer to
Languages for NetWare	"LANGUAGE" in <i>Utilities Reference</i> "International use of NetWare 4.1" in <i>Concepts</i>

Changing the Server Keyboard Type

NetWare 4 allows you to use keyboard types other than U.S. English by loading the KEYB.NLM program. The language you specify with KEYB.NLM must match the language of your keyboard, not the server language.

The keyboard type can be changed without rebooting the server.



Changing the keyboard type causes some of the keys on your keyboard to represent different characters. Do not use this command unless you have the appropriate keyboard for the language you are specifying and you are familiar with the keyboard's use. Otherwise, you may not know which keys to use to change the keyboard back to its original language.

Procedure



1. To view a list of valid keyboard types, type

```
LOAD KEYB <Enter>
```

2. From the list of keyboard types, find the keyboard type that matches the one you are using.

3. To change the keyboard type to match your keyboard, type

```
LOAD KEYB <keyboard_type> <Enter>
```

For example,

```
LOAD KEYB GERMANY <Enter>
```



Keyboard types are dependent on the code page set by DOS. If you must change the DOS code page setting, see your DOS manual for more information.

Additional Information

For more information about	Refer to
Languages for NetWare	"International use of NetWare 4.1" in <i>Concepts</i>
Languages supported by KEYB	"KEYB" in <i>Utilities Reference</i>

Viewing and Setting Server Time and Time Zone

Three server commands allow you to view and set server time and time zone: `TIME`, `SET TIME`, and `SET TIME ZONE`.

Viewing Server Time

To display the date and time kept by the NetWare server's clock, type the following at the server console:

```
TIME <Enter>
```

In addition to the date and time, the `TIME` command displays the server's daylight-saving-time status and time synchronization status. For more information, see "TIME" in *Utilities Reference* and "Time synchronization" in *Concepts*.

Setting Server Time



Because time synchronization between servers is critical to NetWare Directory Services, make sure you understand time synchronization before you change the time or time zone on a server.

If you change the time on a Primary, Reference, or Single Reference server, you affect the time on other servers that use it as a time source.

Setting the time incorrectly can adversely affect the performance of NetWare Directory Services. For more information, see "Managing Network Time Synchronization" on page 593 and "SET TIME" in *Utilities Reference*.

To set or change the date and time kept by the NetWare server's clock, type the following at the server console prompt:

```
SET TIME [mo/day/yr] [hr:min:sec] <Enter>
```

If you don't specify a.m. or p.m., NetWare 4.1 uses a 24-hour clock. For example, if you enter 1:00, the time is set to 1:00 a.m. If you enter 13:00, the time is set to 1:00 p.m.

Changing the Server Time Zone

Important



Because time synchronization between servers is critical to NetWare Directory Services, make sure you understand time synchronization before you change the time or time zone on a server. Setting the time incorrectly can adversely affect the performance of NetWare Directory Services.

The SET TIME ZONE command specifies the abbreviation for the local time zone and for daylight saving time. It also specifies the difference in hours between the local time zone and Coordinated Universal Time. Coordinated Universal Time (UTC) has historically been known as Greenwich Mean Time.

The SET TIME ZONE command does not set the time or turn daylight saving time on and off. You turn on daylight saving time and set other related parameters by using the SET command. For more information, see "SET" in *Utilities Reference*.

To display the current time zone setting, type the following at the server console prompt:

```
SET TIME ZONE <Enter>
```

Procedure

To change the time zone information for your NetWare server, follow these steps.

Procedure



1. At the server console prompt, type

```
SET TIME ZONE zone [+|-] hr:min:sec[daylight]  
<Enter>
```

Where

- ◆ *zone* is a standard three-letter abbreviation for the time zone, such as PST for Pacific Standard Time.

[+ | -] is the number of hours east or west of the Coordinated Universal Time (UTC) meridian. If you do not specify + or -, the system defaults to +.
- ◆ *hr:min:sec* indicates the time difference between UTC and the local time zone. You can specify this difference simply in hours. Or, for extreme accuracy, specify hours, minutes, and seconds.

- ◆ *daylight* represents a standard three-letter abbreviation for daylight saving time, such as PDT for Pacific Daylight Time.

For example, the following command line specifies PST as the standard time abbreviation, 8 as the number of hours to add to Coordinated Universal Time, and PDT as the abbreviation for daylight time:

```
SET TIME ZONE PST8PDT <Enter>
```

2. Change the SET TIME ZONE command in the AUTOEXEC.NCF file to the new time zone.

For more information, see “SET TIME ZONE” in *Utilities Reference*.

Changing a Server’s Name or IPX Internal Network Number

The server’s name and IPX internal network number are stored in the AUTOEXEC.NCF file. You can edit this file by using INSTALL or EDIT at the server console, or by using an ASCII text editor at the workstation.




Note

This procedure requires you to reboot the NetWare server. You may want to do it after business hours or when no one is accessing the server.

To use INSTALL to change the settings, follow the procedure in this section. To use EDIT, see “Editing Text Files from the Server Console” on page 460.

Procedure



Procedure

1. At the server console prompt, type

```
LOAD INSTALL <Enter>
```

2. Choose “NCF Files Options” from the “Installation Options” menu.

3. Choose “Edit AUTOEXEC.NCF File” from the “NCF Files Options” menu.

4. Change the following lines in the file:


File Server Name *server_name*
IPX Internal Net *net_number*

5. To exit and save the file, press <F10>, and then select “Yes” from the “Save file AUTOEXEC.NCF?” box.

6. To exit INSTALL, press <Alt>+<F10>.

7. Reboot the server to make the changes take effect.

See “Bringing Down a Server” on page 446 if you need help rebooting the server.



Using INSTALL to rename a server does not affect NDS™ server objects. You must use NETADMIN or NetWare Administrator to rename NetWare Server objects.

Likewise, using INSTALL to rename a volume does not affect NDS volume objects. For more information, see “Renaming Volumes” on page 521.

Be sure to update any references to the former server name in the Directory tree and in the TIMESYNC.CFG file.

Managing Server Connections

This section explains how to enable and disable access to the server and how to monitor and solve common connection problems.

Disabling and Enabling Logins

If you need to make repairs to the server, use the `DISABLE LOGIN` command to prevent users from logging in.



This does not affect users who are already logged in to the network.

To reenable logins, or to enable the `ADMIN` account when it has been locked by the intruder detection function, use `ENABLE LOGIN`.

For more information about these commands, see “`DISABLE LOGIN`” and “`ENABLE LOGIN`” in *Utilities Reference*.

For more information about the intruder detection function, see Table 1-8 and Table 1-10 in Chapter 1.

Clearing a Workstation Connection

You can use either the `MONITOR` utility or the `CLEAR STATION` command to clear a connection, as described in the following sections.

Use either of these procedures to clear a workstation connection when the workstation has crashed and left open files on the server.



If you clear the connection while the workstation is in the middle of a transaction or a file update, the files may be saved with incorrect data.

Using MONITOR to Clear the Connection

Procedure

Procedure



1. **At the server console prompt, type**

LOAD MONITOR <Enter>

2. **Select “Connection Information” from the “Available Options” menu.**

A list of active connections appears.

3. **Select the connection you want to clear.**

4. **(Optional) Press <Enter> to view files that are open for the workstation.**

After viewing files, press <Esc> to go back to the “Connection Information” screen.

5. **To clear the workstation connection, press <Delete>.**

6. **At the “Clear Connection?” box, select “Yes.”**

The connection is no longer listed on the “Connection Information” screen.

Using the CLEAR STATION Command to Clear the Connection


Prerequisites

Checklist



- The connection number for the workstation. You can determine the connection number by selecting “Connection Information” from the “Available Options” menu in MONITOR.

Procedure



1. At the server console prompt, type

CLEAR STATION *n* <Enter>

Replace *n* with the connection number of the workstation.

The system displays a message stating that the connection is cleared.

Additional Information

For more information about	Refer to
Connection numbers	“Connection number” in <i>Concepts</i>
Using the CLEAR STATION command	“CLEAR STATION” in <i>Utilities Reference</i>
Using the MONITOR utility	“MONITOR” in <i>Utilities Reference</i>

Monitoring Workstation Connections

Use this procedure for early warning that a workstation connection has become inactive. This procedure changes the SET watchdog parameters, which a server sends to a workstation to determine if the workstation connection is still active.

When the parameters are set to the values specified in the following procedure, the server sends 20 watchdog packets (twice the default value) but sends them in the shortest time allowed. If there is no response from the workstation, the server disconnects the workstation and displays an alert on the console.

Procedure



1. **At the server console prompt, type**

```
LOAD SERVMAN <Enter>
```

2. **Choose “Server Parameters” from the “Available Options” menu.**
3. **Choose “Communications” from the “Categories” menu.**
4. **Change the following watchdog parameters to the indicated values:**

```
Number of Watchdog Packets=20  
Delay Between Watchdog Packets=10  
Delay Before First Watchdog Packet=16  
Console Display Watchdog Logouts=ON
```

5. **Press <Esc> twice to reach the “Update Options” menu.**
6. **Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

A window appears, indicating the path to the AUTOEXEC.NCF file.

7. **If desired, press <Enter> to update the file.**

The system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file.

If you do not update the AUTOEXEC.NCF file, the parameter changes last only until the server is rebooted.

Additional Information

For more information about	Refer to
Watchdog packets	“Watchdog” in <i>Concepts</i>
SET parameters	“SET” in <i>Utilities Reference</i>
Using SERVMAN to set parameters	“SERVMAN” in <i>Utilities Reference</i>

Increasing the Maximum Number of Packet Receive Buffers

Packet receive buffers store incoming data packets until they can be processed by the server.

You should increase the maximum number of packet receive buffers only if the server is running out of buffers. The server is running out of buffers if it is

- ◆ Slowing down and losing workstation connections
- ◆ Receiving “No ECB available count” errors

The “General Information” window of MONITOR displays the total number of currently allocated packet receive buffers.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

- 2. Choose “Server Parameters” from the “Available Options” menu.**
- 3. Choose “Communications” from the “Categories” menu.**
- 4. Choose “Maximum Packet Receive Buffers.”**

5. Increase the value of this parameter and press <Enter>.

A good rule of thumb is to set this value to twice the size of the Minimum Packet Receive Buffer value.

For other suggestions, see the discussion of the Maximum Packet Receive Buffer parameter under “SET” in *Utilities Reference*.

6. Press <Esc> twice to reach the “Update Options” menu.

7. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

You are prompted to update the STARTUP.NCF file.

8. If the SET parameter is in the STARTUP.NCF file, press <Enter> to update the value in the file. Otherwise press <Esc>.

You are prompted to update the AUTOEXEC.NCF file.

9. If desired, press <Enter> to update the value in the file. Otherwise, press <Esc>.

If you press <Enter> to update the file, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file.

Increasing the Minimum Number of Packet Receive Buffers


The operating system allocates a minimum number of packet receive buffers as soon as the server boots. The minimum number is specified by the Minimum Packet Receive Buffer SET parameter.

Refer to the Packet Receive Buffer value in the “General Information” window of MONITOR to determine how many buffers the server is currently allocating.

Use the following procedure to increase the minimum number of packet receive buffers if

- ◆ The allocated number is higher than 10 and the server doesn’t respond immediately after booting
- ◆ The “No ECB available” count in MONITOR grows continuously.

Procedure



1. **At the server console prompt, type**

LOAD SERVMAN <Enter>

2. **Choose “Server Parameters” from the “Available Options” menu.**
3. **Choose “Communications” from the “Categories” menu.**
4. **Choose “Minimum Packet Receive Buffers.”**
5. **Increase the value of this parameter.**

As a rule of thumb, allocate at least two packet receive buffers for each workstation connection.

For other suggested settings, see the discussion of the Minimum Packet Receive Buffer parameter under “SET” in *Utilities Reference*.



The Minimum Packet Receive Buffers value should be smaller than the Maximum Packet Receive Buffers value. If it is greater than the maximum value, the system changes the maximum value to match the minimum value.

6. **Press <Esc> twice to reach the “Update Options” menu.**
7. **Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

A window appears, indicating the path to the STARTUP.NCF file.

8. **Press <Enter> to update the file.**

The system places the parameter in the STARTUP.NCF file the next time the server boots. The parameter can be set only in the STARTUP.NCF file.

9. **When you want the changes to take effect, reboot the server.**


Avoiding Bottlenecks During Peak Usage

Use this procedure if your server is slow to respond during periods of heavy use.

This procedure decreases the value of the New Packet Receive Buffer Wait Time SET parameter, which specifies how long the server must wait before granting a new packet receive buffer. By decreasing the value of this parameter you cause the system to grant buffers more quickly.

You may also want to increase the value of the Minimum Packet Receive Buffers parameter. "Increasing the Minimum Number of Packet Receive Buffers" on page 478.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose "Server Parameters" from the "Available Options" menu.

3. Choose "Communications" from the "Categories" menu.

4. Choose "New Packet Receive Buffer Wait Time."

5. Decrease the value of this parameter and press <Enter>.

6. Press <Esc> twice to reach the "Update Options" menu.

7. Choose "Update AUTOEXEC.NCF and STARTUP.NCF Now" from the "Update Options" menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

8. If desired, press <Enter> to update the file.

If you press <Enter>, the system writes the parameter to the AUTOEXEC.NCF file or updates the parameter if it is already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter change lasts only until you reboot the server.

Monitoring and Optimizing the Server

This section describes how to prevent server problems that involve memory management, CPU utilization, process scheduling, and LAN traffic. There are also procedures to help improve server performance.

Assessing Server RAM

One of the most common causes of a slow network is insufficient Random Access Memory (RAM) in the NetWare server. This procedure helps you determine if your server has enough RAM.

Procedure



1. At the server console prompt, type

LOAD MONITOR <Enter>

2. Choose “Cache Utilization” from the “Available Options” menu.

3. Write down the percentage of Long Term Cache Hits shown in the lower window of the screen.

Tracking cache hits helps you predict when you need more RAM before you experience serious performance loss.

If the percentage of Long Term Cache Hits falls below 90 percent, add more RAM to the server.

4. (Optional) To temporarily free up RAM, do either or both of the following:

- ◆ Use the REMOVE DOS or SECURE CONSOLE command to make the memory used by DOS available to cache.
- ◆ Unload any NLM programs that are not critical.

5. To exit MONITOR, press <Esc>, or to return to the console prompt without exiting MONITOR, press <Alt>+<Esc>.

Additional Information


For more information about	Refer to
Cache buffers	"Cache buffer" in <i>Concepts</i>
MONITOR utility	"MONITOR" in <i>Utilities Reference</i>
REMOVE DOS command	"REMOVE DOS" in <i>Utilities Reference</i>
SECURE CONSOLE command	"SECURE CONSOLE" in <i>Utilities Reference</i>

Controlling Memory Allocation

This procedure explains how to limit the amount of memory available for file and directory caching, thus making more memory available for NetWare Loadable Module™ (NLM) programs and server processes. Use this procedure if you receive "insufficient memory" alerts when trying to load NLM™ programs.

To solve memory problems without limiting memory for file and directory caching, add more memory.

Procedure



1. **At the server console prompt, type**

LOAD SERVMAN <Enter>

2. **Choose “Server Parameters” from the “Available Options” menu.**
3. **Choose “File Caching” from the “Categories” menu.**
4. **Choose “Minimum File Cache Buffers.”**

This parameter specifies the minimum amount of memory the system will keep for file caching. Memory above this minimum amount is available for other processes.

5. **Reduce the value of this parameter.**
6. **Press <Esc> twice to reach the “Categories” menu.**
7. **Choose “Directory Caching” from the “Categories” menu.**
8. **Choose “Maximum Directory Cache Buffers.”**

This parameter specifies the maximum number of buffers that can be allocated for directory caching. Once allocated, the buffers remain allocated, even if they are not being used. Thus, they are not available for other processes.

By reducing this parameter, you prevent the system from allocating too many directory cache buffers and you ensure that more memory is available for other server processes.



Reducing the maximum number of directory cache buffers may impede system performance.

9. **Reduce the value of the “Maximum Directory Cache Buffers” parameter.**
10. **Press <Esc> twice to reach the “Update Options” menu.**


11. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

12. Press <Enter> to update the file.

When you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file.

13. Reboot the server to free the memory.



Other temporary measures to free memory include using the REMOVE DOS or SECURE CONSOLE command to make the DOS partition memory available, and unloading noncritical NLM programs. See “REMOVE DOS” and “SECURE CONSOLE” in *Utilities Reference*.


Speeding Directory Searches

Use this procedure if the server responds slowly to directory searches.



This procedure requires you to increase the number of directory cache buffers that can be allocated by the system. Memory allocated for directory cache buffers is no longer available for use by other processes.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose “Server Parameters” from the “Available Options” menu.

3. Choose “Directory Caching” from the “Categories” menu.

4. Select “Directory Cache Allocation Wait Time.”

This parameter specifies how long the system must wait after allocating one directory cache buffer before it can allocate another buffer. If the value is too high, directory searches are slow.

5. Decrease the value of this parameter.

6. Select “Maximum Directory Cache Buffers.”

This parameter specifies the maximum number of cache buffers the system can allocate for directory caching. If the number is too low, directory searches are slow.

7. Increase the value of this parameter.

8. Select “Minimum Directory Cache Buffers.”

This parameter specifies the minimum number of directory cache buffers the system allocates. If the number is too low, directory searches are slow.

9. Increase the value of this parameter.

10. Press <Esc> twice to reach the “Update Options” menu.

11. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

12. If desired, press <Enter> to update the file.

If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.

Prioritizing Server Processes

You can use SCHDELAY to prioritize server processes. SCHDELAY allows you to delay the execution of specific processes, thereby giving more CPU time—and higher priority—to other processes.

The command syntax is

```
LOAD SCHDELAY process_name = number [process_name =  
number] <Enter>
```

The *process name* value is the name of a process or thread that accesses the server CPU. You can specify the process value for multiple processes at one time.

To view the list of processes that can be prioritized, along with their current SCHDELAY values, type

```
LOAD SCHDELAY <Enter>
```

The *number* value determines how many times the process delays execution. A value of 0 causes no delay, while a value of 2 skips every other execution of the process.

To delay a process, you can enter a value from 2 to 10000—the higher the value, the lower the priority of the process.



The SCHDELAY values only have an effect if the system is busy and two or more processes are vying for the CPU's time.

For example, if you want to assign a database NLM called DBNLM a lower priority than other processes so it won't monopolize the NetWare server's CPU, put a command similar to the following in the AUTOEXEC.NCF file:

```
LOAD SCHDELAY DBNLM = 2
```

To view the list of processes and determine which processes are hoarding CPU time, follow this procedure.

Procedure



1. At the server console prompt, type

LOAD MONITOR <Enter>

2. Select “Scheduling Information” from the “Available Options” menu.

A list of server processes appears, with four columns of information about each process.

- ◆ The “Sch Delay” column shows the current SCHDELAY value assigned to the process. The default is 0.
- ◆ The “Time” column shows the CPU time spent executing the process code.
- ◆ The “Count” column shows how many times the process ran during the sample period. (The “Total Sample Time” is displayed at the bottom of the screen.)
- ◆ The “Load” column shows the percent of total CPU time spent on the process.

3. Select a process that has a high number in the “Load” column.

Look for processes that have consistently high values, rather than high peak values.

4. Use the plus (+) or minus (–) key to increase or decrease the SCHDELAY value for the selected process.

The value in the “Load” column changes accordingly.

5. Experiment with higher or lower SCHDELAY values until the CPU load, as displayed in the “Load” column, is within acceptable limits.

What is acceptable depends upon your system and your needs. The goal is to balance the load so that processes run to your satisfaction.



To make the SCHDELAY value effective each time the server boots, put the SCHDELAY command in the AUTOEXEC.NCF file.


6. **To exit the scheduling information window, press <Esc>.**

You do not need to reboot the server.

Checking a Server's Processor Utilization

One possible cause of network problems is the monopolizing of the CPU by a single process. This procedure helps you determine which processes are using the CPU and how much load they are putting on the CPU's time.

Procedure



1. **At the server console prompt, type**

LOAD MONITOR <Enter>

2. **Choose "Processor Utilization" from the "Available Options" menu.**
3. **Press <F3> to view a list of the currently loaded processes and interrupts.**

For each item on the list, notice the percentage given in the "Load" column. This is the percentage of time the CPU spent servicing this process during the one-second sample period.

Compare processes to determine which processes are taking the most CPU time.

4. **(Optional) To redistribute processor utilization, you can do either or both of the following:**
 - ◆ Unload noncritical NLM programs that are taking a large percentage of CPU time.
 - ◆ Assign processes a lower priority, using SCHDELAY. For more information, see "Prioritizing Server Processes" on page 486.

Viewing a Server's Packet Traffic

You can view the packet traffic coming in and going out of the server by using `TRACK ON` at the server console prompt. Use `TRACK OFF` to stop viewing the packet traffic.

See “`TRACK ON`” and “`TRACK OFF`” in *Utilities Reference* for more information about these commands.

Viewing LAN Driver Statistics

Reviewing the LAN driver statistics on your server can help you regulate network traffic and troubleshoot system problems.

For example, you can use the LAN driver statistics to determine


- ◆ The total traffic since the LAN driver was loaded
- ◆ Which drivers handle the most traffic
- ◆ The frequency of various errors

Many statistics count the instances of a particular error. For example, the Send Packet Retry Count statistic indicates how many times the server tried to send a packet but failed because of a hardware problem.

The number and kinds of statistics available depend on the LAN drivers installed on your system. For descriptions of the most common statistics, refer to Appendix B, “LAN Driver Statistics,” of *Utilities Reference*.

Follow this procedure to view the statistics.

Procedure



1. At the server console prompt, type

LOAD MONITOR <Enter>

2. Select “LAN/WAN Information” from the “Available Options” menu.

3. Select a LAN driver from the “LAN Driver Information” menu.

The statistics for the selected LAN driver are displayed. Press <PageUp> and <PageDown> to scroll through the information in the window. For an explanation of the statistics, see Appendix B, “LAN Driver Statistics,” of *Utilities Reference*.

Improving Server Performance

You can increase server performance two ways:

- ◆ By using magneto-optical media to extend server storage capacity. Such storage enables the server to migrate seldom-used files from the hard disk, freeing space needed by other files and processes.

For procedures related to magneto-optical media, see Chapter 6, “Migrating Data Using the High Capacity Storage System.”

- ◆ By increasing the size of the packets received by the server.

The method of increasing packet size is explained in the following procedure.




Use this procedure only if the network boards in your server can transmit more than 512 bytes of data per packet. Refer to the documentation that came with the network boards to determine the packet size.

Prerequisites



- Large-packet network boards installed in the server. Refer to the documentation that comes with the boards to determine the available packet size.

Procedure



1. **At the server console prompt, type**

LOAD SERVMAN <Enter>

2. **Choose “Server Parameters” from the “Available Options” menu.**
3. **Choose “Communications” from the “Categories” menu.**
4. **Choose “Maximum Physical Receive Packet Size” from the “Communications” menu.**
5. **Increase the value and press <Enter>.**

Refer to the documentation that comes with the boards to determine the maximum value.

6. **Press <Esc> twice to reach the “Update Options” menu.**
7. **Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

A window appears, indicating the path to the STARTUP.NCF file.

8. **Press <Enter> to update the file.**
9. **When you want the changes to take effect, reboot the server.**

Improving Disk Writes

The following sections describe methods for improving the speed of disk writes.

Increasing the Number of Concurrent Writes

You can increase the speed and efficiency of disk cache writes by increasing the number of write requests that can be executed at one time.

To determine if you need to increase the number of concurrent writes, first compare the number of dirty cache buffers to the total number of cache buffers. Dirty cache buffers contain data that has not yet been written to disk. These statistics are found on the “General Information” screen in MONITOR.

If the number of dirty cache buffers is greater than 70 percent of total cache buffers, use the following procedure to increase the number of concurrent write requests.


For an explanation of file caching and dirty cache buffers, refer to “Cache memory” in *Concepts*.



Note

Increasing the number of concurrent disk cache writes slows disk cache reads. You will want to balance the speed of disk writes and reads to meet your own needs.

Procedure



Procedure

1. At the server console prompt, type

```
LOAD SERVMAN <Enter>
```

2. Choose “Server Parameters” from the “Available Options” menu.
3. Choose “File Caching” from the “Categories” menu.
4. Choose “Maximum Concurrent Disk Cache Writes” from the “File Caching” menu.

5. Increase the parameter value and press <Enter>.

If the parameter is currently at the default value of 50, try increasing it to 100.

6. Press <Esc> twice to reach the “Update Options” menu.

7. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

8. If desired, press <Enter> to update the file.

If you press <Enter>, the system writes the parameter to the AUTOEXEC.NCF file or updates the parameter if it is already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter change lasts only until you reboot the server.

Changing Disk and Directory Caching for Faster Writes

Use this procedure if network users frequently make many small write requests and the server is slow to respond to the requests.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose “Server Parameters” from the “Available Options” menu.

3. Choose “File Caching” from the “Categories” menu.

4. Choose “Dirty Disk Cache Delay Time” from the “File Caching” menu.

5. Increase the value and press <Enter>.

This parameter specifies how long the system waits before writing a not-completely-dirty cache buffer to disk.

If the value is low, the system writes to disk more frequently, but writes fewer requests each time. If the value is high, the system waits longer before writing to disk, but executes more write requests with each operation. A higher value provides greater efficiency in writing to disk.

If the parameter is currently at the default value of 3.3 seconds, try increasing the value to 7 seconds.

6. Press <Esc> to reach the “Categories” menu.

7. Choose “Directory Caching” from the “Categories” menu and press <Enter>.

8. Choose “Dirty Directory Cache Delay Time.”

This parameter specifies how long the system keeps a directory table write request in memory before writing it to disk.

Increasing the parameter provides slightly faster performance, but may increase the chance of directory tables becoming corrupted.

9. Increase the value and press <Enter>.

If the parameter is currently at the default value of 0.5 second, try increasing the value to 2 seconds.

10. Choose “Maximum Concurrent Directory Cache Writes.”

This parameter determines how many write requests from directory cache buffers are executed at one time. Increasing this value increases the efficiency of directory cache write requests.



Increasing the number of concurrent directory cache writes decreases the speed of directory cache reads. You will want to balance the speed of writes and reads to meet your own needs.

11. Increase the value and press <Enter>.

If the parameter is currently at the default value of 10, try increasing the value to 25.

12. Press <Esc> twice to reach the “Update Options” menu.

13. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

14. If desired, press <Enter> to update the file.


If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.

Turning Off Read-After-Write Verification

Read-after-write verification is an important means of protecting the data on your system. Normally, you should not disable it.

However, if your disks are mirrored and reliable, you may choose to disable read-after-write verification because disabling almost doubles the speed of disk writes.




Turning off read-after-write verification may increase the risk of data corruption on the server's hard disk. You should use the following procedure only if your disks are mirrored and reliable, and you understand the risk.

Prerequisites



Mirrored hard disks

Procedure



1. At the server console prompt, type

LOAD MONITOR <Enter>

2. Choose “Disk Information” from the “Available Options” menu and press <Enter>.

3. **Select a disk drive from the “System Disk Drives” menu.**
4. **Select “Read After Write Verify” from the “Driver Status” menu.**
5. **Select “Disable Verify” from the “Read After Write Status” menu.**



You can also disable read-after-write verification by setting the Enable Disk Read After Write Verify SET parameter to OFF. However, this setting only affects those disks loaded after the parameter value is changed. It does not change the setting for currently loaded disks.

Improving Disk Reads

The following sections describe methods for improving the speed of disk reads.

Changing Concurrent Disk and Directory Writes for Faster Reads

Use this procedure if your server is slow to respond to read requests.



This procedure requires you to decrease the values of the Maximum Concurrent Disk Cache Writes parameter and the Maximum Concurrent Directory Cache Writes parameter. Although decreasing these values increases the speed of read requests, it may decrease the speed and efficiency of write requests.

Procedure



1. **At the server console prompt, type**
LOAD SERVMAN <Enter>
2. **Choose “Server Parameters” from the “Available Options” menu.**
3. **Choose “File Caching” from the “Categories” menu.**
4. **Choose “Maximum Concurrent Disk Cache Writes.”**
5. **Decrease this value and press <Enter>.**

If the parameter is currently set to the default value of 50, try setting the value to 10.

- 6. Press <Esc> to reach the “Categories” menu.**
- 7. Choose “Directory Caching” from the “Categories” menu.**
- 8. Choose “Maximum Concurrent Directory Cache Writes.”**
- 9. Decrease this value and press <Enter>.**

If the parameter is currently set to the default value of 10, try setting the value to 5.

- 10. Choose “Directory Cache Buffer NonReferenced Delay.”**

This parameter specifies how long a directory entry must be cached before it can be overwritten. Increasing this value causes the system to allocate more directory cache buffers and thus speeds directory access.

- 11. Increase this value and press <Enter>.**

If the parameter is currently set to the default value of 5.5 seconds, try setting the value to 60 seconds.

- 12. Press <Esc> twice to reach the “Update Options” menu.**

- 13. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

A window appears, indicating the path to the AUTOEXEC.NCF file.

- 14. If desired, press <Enter> to update the file.**

If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.


Changing the Turbo FAT Wait Time for Faster Reads

When a program randomly accesses a file that contains more than 64 File Allocation Table (FAT) entries, the system builds a turbo FAT index for the file so that the information in the file can be accessed quickly.

The Turbo FAT Re-Use Wait Time parameter specifies how long a turbo FAT index remains in memory after the indexed file is closed. When the turbo FAT index is in memory, files can be opened and information accessed faster.

If network users frequently access files larger than 64 blocks, use this procedure to increase the time the index is kept in memory.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose “Server Parameters” from the “Available Options” menu.

3. Choose “File System” from the “Categories” menu.

4. Select “Turbo FAT Re-Use Wait Time.”

5. Increase this value and press <Enter>.

You must specify the value in seconds.

If the parameter is currently set to the default value of 329.5 seconds (5 minutes 29.6 seconds), try setting the value to 600 seconds (10 minutes).

6. Press <Esc> twice to reach the “Update Options” menu.

7. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

8. If desired, press <Enter> to update the file.

If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.

Controlling Resource Allocation with Locks

A lock prevents a file or record from being updated by more than one user at a time. By controlling the number of file and record locks available to a workstation or a server, you control access to files and records.

You may choose to limit the number of locks to prevent overuse of file resources. Or you may choose to increase the number of locks if workstations can't open files.

Keeping Workstations from Overusing File Resources

If workstations are overusing file resources by opening and locking too many files or records at one time, use the following procedure to limit the number of locks for the server and for each workstation.

Before limiting file and record locks, make sure the workstations have an adequate number of locks to access the files and records they need. Limiting needed locks can cause applications to generate errors.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose "Server Parameters" from the "Available Options" menu.

3. Choose "Locks" from the "Categories" menu.

4. **Decrease the following parameter values to limit the total number of file and record locks allowed for the server:**

Maximum Record Locks

Maximum File Locks

5. **Decrease the following parameter values to limit the total number of file and records locks allowed for each workstation:**

Maximum Record Locks Per Connection

Maximum File Locks Per Connection

6. **Press <Esc> twice to reach the “Update Options” menu.**
7. **Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

A window appears, indicating the path to the AUTOEXEC.NCF file.

8. **If desired, press <Enter> to update the file.**


If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.

Increasing File Record Locks

Use this procedure if applications fail because they cannot open enough files or because not enough record locks are available.

Procedure



1. **At the server console prompt, type**

LOAD SERVMAN <Enter>

2. **Choose “Server Parameters” from the “Available Options” menu.**

3. **Choose “Locks” from the “Categories” menu.**

4. **Increase the following parameter values to increase the number of file and record locks allowed for the server:**

Maximum Record Locks

Maximum File Locks

5. **Increase the following parameter values to increase the number of file and record locks allowed for each workstation:**

Maximum Record Locks Per Connection

Maximum File Locks Per Connection

6. **Press <Esc> twice to reach the “Update Options” menu.**

7. **Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

A window appears, indicating the path to the AUTOEXEC.NCF file.

8. Press <Enter> to update the file.

➤ If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.

Displaying Alert Messages

You can use SET parameters to control the display of alert messages on the server console.

Turning on Hardware and NLM Alerts

To display alerts about hardware devices and NLM programs, add the following SET commands to the beginning of the STARTUP.NCF file:

```
SET Display Spurious Interrupt Alerts=ON  
SET Display Lost Interrupt Alerts=ON  
SET Display Relinquish Control Alerts=ON  
SET Display Old API Names=ON
```


To temporarily turn off an alert message display (when it is filling up the console screen, for example), use the SET command to change the relevant parameter setting to OFF. When you reboot the server, the setting is changed to ON again.

Turning on Early Warning Alerts for Cache Buffer Memory Allocation

All memory not allocated for other processes is given to file caching. As memory is requested for other processes, the server gives up file cache buffers. However, there is a minimum number of file cache buffers that cannot be given up. This minimum is specified by the Minimum File Cache Buffers SET parameter.

Use the following procedure to specify when the system will warn you that the number of buffers is approaching the minimum level.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose “Server Parameters” from the “Available Options” menu.

3. Choose “File Caching” from the “Categories” menu.

4. Select Minimum File Cache Buffers.

View this parameter to learn the minimum number of file cache buffers currently set.

5. Choose “Minimum File Cache Report Threshold” from the “File Caching” menu.

This parameter specifies how many unused cache buffers must be left above the minimum value for the system to warn you that cache buffers are low.

6. Change the parameter value as needed and press <Enter>.

For example, if the Minimum File Cache Buffers parameter is set to 20 and the Minimum File Cache Report Threshold is set to 25, you are warned when all but 45 cache buffers are allocated for other processes.

7. Press <Esc> twice to reach the “Update Options” menu.

8. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

9. If desired, press <Enter> to update the file.

If you press <Enter>, the system writes the parameter to the AUTOEXEC.NCF file or updates the parameter if it is already in the file. The file is updated immediately. You do not need to reboot the server.


If you do not press <Enter>, the parameter change lasts only until you reboot the server.

Turning on Early Warning Alerts for a Full Volume

Use this procedure to enable early warning that a volume is nearly full.

If a volume is becoming full, you may need to remove unnecessary files or add a new hard disk.

Procedure



1. At the server console prompt, type

LOAD SERVMAN <Enter>

2. Choose “Server Parameters” from the “Available Options” menu.

3. Choose “File System” from the “Categories” menu.

4. Select “Volume Low Warn All Users” and change the value to ON.

This causes the system to send a warning message to users when a volume is becoming full. The next two steps specify when the warning is sent and when it is removed.

5. Select “Volume Low Warning Threshold” and set it to the desired value.

This parameter determines when the system will warn users that the volume is becoming full.

The threshold specifies how much free space, in blocks, must be left on the volume for the warning to be issued.

6. Select “Volume Low Warning Reset Threshold” and set it to the desired value.

This parameter determines when the warning will be removed.

It specifies the amount of space that must be freed on the volume for the warning message to be cleared and the warning mechanism to be reset.

For example, both the Volume Low Warning Threshold and the Volume Low Warning Reset Threshold might be set to 256 blocks. In this case, the system sends a warning when 256 blocks of free space is left on the volume. The volume must then gain an additional 256 blocks (for a total of 512 blocks of free space) before the warning message disappears and the warning mechanism is reset.

Once the mechanism is reset, the free space on the volume must dip below 256 blocks again before another warning is sent.

7. Press <Esc> twice to reach the “Update Options” menu.

8. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

9. Press <Enter> to update the file.

If you press <Enter>, the system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file. The file is updated immediately. You do not need to reboot the server.

If you do not press <Enter>, the parameter changes last only until you reboot the server.

Viewing the Server Error Log

The server error log is a text file called SYS\$LOG.ERR in the server's SYS_servername:SYSTEM directory. All system messages and alerts that appear on the server console are recorded in the SYS\$LOG.ERR file. You should view this file periodically to see what kinds of errors are occurring on your server.

NetWare security violations are also recorded in the SYS\$LOG.ERR file. Check this file daily if you are concerned about security at your site.



Note

You should regularly clear the SYS\$LOG.ERR file to keep it from using too much server disk space.

You can view the server error log from a workstation using a text editor or from the console using EDIT.NLM. For more information about using EDIT, see "Editing Text Files from the Server Console" on page 460.

You can also view or clear the server error log with the NETADMIN and FILER utilities, as described here.


Prerequisites



Checklist

- A workstation running DOS 3.30 or later.
- A minimum of 512 KB of memory available on the workstation.
- Supervisor or equivalent directory right to SYS:SYSTEM.

Procedure for NETADMIN



Procedure

1. **At the workstation prompt, type**
NETADMIN <Enter>
2. **From the "NetAdmin Options" screen, choose "Manage Objects."**
3. **From the browser, select the Server object whose error file you want to view and press <F10>.**
4. **Choose "View or Edit Properties of This Object."**

5. Choose “View Server Error Log File.”

The server error log appears.

6. To delete or exit the server error log, press <Esc>.

A prompt to clear the error log file appears.

7. To clear the error log file, choose “Yes.”

8. To exit NETADMIN, press <Alt>+<F10>.

Procedure for FILER



1. At the workstation prompt, type

FILER <Enter>

2. From the “Available Options” screen, choose “Manage Files and Directories.”

3. From the “Directory Contents” screen, browse the tree to select the SYS: volume on the Server object whose error file you want to view, and press <F10>.

4. Choose the error log file you want to view.

5. From the “File Options” screen, choose “View File.”

The server error log appears.

6. To exit the server error log, press <Esc> twice.

7. To clear the error log file, be sure the correct filename is highlighted and press <Delete>

A prompt to clear the error log file appears.

8. To clear the error log file, choose “Yes.”

9. To exit FILER, press <Alt>+<F10>.

Managing Error Log Files

NetWare creates and maintains three error log files:

- ◆ SYS\$LOG.ERR, for server errors
- ◆ VOL\$LOG.ERR, for volume errors
- ◆ TT\$LOG.ERR, for data backed out by the Transaction Tracking System™

To keep error log files from using too much disk space, NetWare 4 includes the following SET parameters:


```
SET Server Log File State  
SET Server Log File Overflow Size  
SET Volume Log File State  
SET Volume Log File Overflow Size  
SET Volume TTS Log File State  
SET Volume TTS Log File Overflow Size
```

See "SET" in *Utilities Reference* for information on using these parameters to limit error log file size.

Checking for Disk Errors


You should check for disk errors regularly. Any new errors should be cause for concern, because such errors can lead to a failed file system and loss of data.

If you check for disk errors regularly, you can replace a disk before it causes data loss.



If your server's disk drive performs read-after-write verification and automatically redirects bad blocks, the following procedure does not work. It only works if NetWare does the redirection, using Hot Fix™.

Procedure



1. **At the server console prompt, type**

LOAD MONITOR <Enter>

2. **Choose “Disk Information” from the “Available Options” menu.**
3. **Choose the disk drive you want to check from the “System Disk Drives” menu.**

The upper part of the screen displays information about the disk device that you selected, including Hot Fix and Redirected Block statistics.

4. **Write down the number of Redirected Blocks and the number of Redirection Blocks shown in the upper window of the screen.**

When NetWare cannot reliably write a block of data to the disk, it marks the disk block as bad and writes the data to the Hot Fix Redirection Area.

The Redirected Blocks count shown in MONITOR is the number of blocks redirected since the server was installed.

The Redirection Blocks count is the number of blocks designated as the Hot Fix Redirection Area.

5. **Evaluate these statistics and proceed according to the following steps.**

- 5a. **Compare the current number of redirected blocks to the number you recorded the last time you checked this number.**

If the number of redirected blocks is the same as the previously noted number, you do not need to take any corrective action.

If the number of redirected blocks has increased significantly, or if the number of redirected blocks is over half the number of redirection blocks, continue with Step 5b.

If the number of redirected blocks has slightly increased over the previous number, monitor the server closely.

- 5b. Send a message to all users to close files and log out.
- 5c. Back up the server data and bring down the server.
- 5d. Troubleshoot the disk drive, driver, controller, and host bus adapter, using the hardware documentation.

Additional Information

For more information about	Refer to
Data protection	"Data protection" in <i>Concepts</i>
MONITOR utility	"MONITOR" in <i>Utilities Reference</i> Online help in MONITOR
Redirection, Hot Fix, and read-after-write verification	"Hot Fix" in <i>Concepts</i>

Testing for Bad Disk Blocks

Use the surface test described in this section to check a NetWare partition's hard disk for bad blocks. The surface test runs in the background, so you can access other hard disks during the test.



Note

Use the surface test only if your hardware is several years old. Newer hardware is of such reliability that the surface test is not helpful. Problems with newer hardware should be addressed according to the manufacturer's recommendations.

Some drivers may not support this test.

This test can take several hours.



Important

Test for bad blocks when there is no activity on the NetWare server. Otherwise, client workstations connected to the server may hang.

You can run two types of tests on a NetWare partition:

- ◆ **Destructive.** This test destroys all data on the hard disk as it reads and writes test patterns. After the test is over, the disk still retains the test patterns. This method is much faster than the nondestructive surface test.

Use this procedure if you are planning to install a new NetWare partition and you have backed up the data that is currently on the partition.

- ◆ **Nondestructive.** This surface test not only fixes the bad blocks. It is much slower than the destructive method, but it protects the data on the disk.

This test pre-reads the data on the disk, writes and reads its test patterns, and then restores the original data. The nondestructive test takes 20 percent longer to run than the destructive test.

Prerequisites

Checklist



- A NetWare partition to test
- A backup of all data on the hard disk to be tested
- All volumes dismounted on the disk to be tested
- Unmirrored disk partitions
- All network users logged out

Procedure

Procedure



Complete the following steps to test your hard disk for bad blocks.

- 1. At the server console prompt, dismount each volume on the hard disk by typing**

```
DISMOUNT volume_name <Enter>
```

Replace *volume_name* with a name of a volume on the disk. Repeat this step for each volume on the disk.

- 2. Load INSTALL by typing**

```
LOAD INSTALL <Enter>
```

- 3. Select “Disk Options” from the “Installation Options” menu.**
- 4. Select “Perform Surface Test” from the “Available Disk Options” menu.**

5. If you have more than one hard disk on your server, select the disk you want to test from the “Available Disk Drives” list.

The system displays a “Disk Surface Test Status” window and a “Surface Test Options” menu, as shown in the following illustration.

Disk Surface Test Status	
Test Type: None	Blocks Tested:
Bad Blocks: 0	Total Blocks: 16358

Surface Test Options
Begin surface test
Stop surface test
Return to previous menu

As the test progresses, the “Test Type” field on the “Disk Surface Test Status” screen can show the following status information:

Status	Explanation
None	The partition is not being tested.
Destructive	A destructive test was selected. Data on the hard disk will be lost.
Nondestructive	A nondestructive test was selected. Data on the hard disk will be preserved.
Failed	A disk error occurred that Hot Fix could not correct.
Completed	The test was completed successfully.
Terminated	The test was stopped before completion.

6. Select “Begin Surface Test” from the “Surface Test Options” menu.

7. Select the type of test to be performed (destructive or nondestructive).



Do not exit INSTALL during the test. All testing stops if you exit, and your data may not be preserved.

To stop the test before it finishes, select “Stop Surface Test” from the “Surface Test Options” menu. If you select this option, you must start the test from the beginning.



To perform other NetWare server operations while the surface test is running (such as testing other hard disks), press <Esc> to move to the other options in INSTALL, or press <Alt>+<Esc> to go to the console prompt.

At the end of the test, the “Disk Surface Test Status” window displays the number of blocks tested, the number of bad blocks found, and the total blocks tested.

The “Test Type” field displays “Terminated” if the test was stopped before completion, “Failed” if a disk error occurred that could not be fixed, or “Completed” if the test was successful.

8. Press <Alt>+<F10> to exit INSTALL.

Maintaining Volumes

This section describes how to create, delete, rename, mount, dismount, and repair NetWare volumes. It also includes procedures for adding volume segments and changing volume size.

In NetWare Directory Services™, each volume is also an object in the Directory. When you create a volume with INSTALL.NLM, it puts a Volume object in the same context as the NetWare Server object within the Directory tree. By default, INSTALL names the Volume object *servername_volumename*.

You can change the context of Volume objects and rename Volume objects with either the NETADMIN text utility or the NetWare Administrator graphical utility.



To rename or delete a volume, you must use INSTALL to rename or delete the volume on the server, and then use NETADMIN or NetWare Administrator to change the Volume object in the Directory.

Creating Volumes

With NetWare 4.1, you can create a new volume on any hard disk that has a NetWare 4.1 partition.

You use INSTALL to create volumes.


Your NetWare server can have from 1 to 64 NetWare volumes, depending on the needs of your users. NetWare requires only one volume, SYS:

Prerequisites



- A NetWare partition on the disk where you want to create the volume
- An existing volume SYS:

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. From the “Installation Options” menu, choose “Volume Options.”

If any volumes exist, they are now listed.

3. Press <Insert>.

The “Volume Disk Segment List” is displayed.

4. Select any existing free space and press <Enter>.

Free space is any NetWare disk partition space to which no volume assignment has been made.



By default, a new volume takes up all remaining free space of a NetWare disk partition when it is created. However, you can decrease the volume size to leave some free space on the partition for adding NetWare volumes in the future.

If no free space is available, you must add a new hard disk, or add magneto-optical disk storage, before you can create new volumes.

5. Type a new volume name in the box provided and press <Enter>.

The volume name can be up to 15 characters long (A through Z, 0 through 9, and underscore characters are allowed).

The newly created volume now appears in the “Volume Disk Segment List.”

If you want the new volume to use all of the free disk space, skip to Step 8. If not, continue with Step 7.

6. If you don’t want the new volume to use all of the free disk space, do the following:

6a. Select the volume you just created and press <Enter>.

6b. Enter a new volume size, expressed in megabytes (MB), and press <Enter>.

- 6c. **Press <Esc> and save the settings.**
- 6d. **To assign the free space you just created to another volume, repeat Step 4 and Step 5.**
- 6e. **Skip to Step 9.**
- 7. **If you want the new volume to use all of the free disk space, either press <Esc> to continue with other volume tasks or press <F10> to save volume information to disk.**
- 8. **From the menu, choose “Mount/Dismount an Existing Volume.”**
- 9. **Choose “Mount.”**

Additional Information

For more information about	Refer to
Creating NetWare disk partitions	“Creating NetWare Disk Partitions” on page 561
Changing the size of a volume	“Modifying the Size of a Volume” on page 529
Adding a hard disk	“Adding a Hard Disk to the NetWare Server” on page 551
Adding magneto-optical disk storage	Chapter 6, “Migrating Data Using the High Capacity Storage System”

Mounting and Dismounting Volumes

By mounting a volume, you make its information available to network users. You can mount or dismount one or all volumes while the NetWare server is running.

You may want to leave rarely used volumes dismounted because each mounted volume takes up memory and reduces the memory available for file caching.


You can mount and dismount volumes in two ways:

- ◆ Use the MOUNT and DISMOUNT console commands. (See “MOUNT” and “DISMOUNT” in *Utilities Reference*.)
- ◆ Use INSTALL as described in the following procedure.



To mount CD-ROM volumes, use the CD command. For more information, see “Using a CD-ROM as a NetWare Volume” on page 530.

Procedure



1. **At the server console prompt, type**

LOAD INSTALL

2. **From the “Installation Options” menu, choose “Volume Options” and press <Enter>.**

A list of existing volumes is displayed.

3. **Select the volume you want to mount or dismount and press <Enter>.**

The “Volume Information” screen is displayed.

4. **Using the arrow keys, highlight the “Status” field.**

Depending on the status of the volume, this field displays “Mounted,” “Not Mounted,” or “New, not mounted.”

5. **If the status is “New, not mounted,” press <Esc> and then <F10> to save volume changes to disk before the volume can be mounted.**

6. Press <Enter> to display a menu of available actions.

7. Depending on your situation, select either “Mount” or “Dismount” and press <Enter>.

NetWare mounts or dismounts your volume.



If volumes fail to mount, you may not have enough RAM installed to accommodate the volume. For more information, see “Assessing Server RAM” on page 481.


Deleting Volumes



Deleting one volume segment deletes all existing data on the volume! You cannot delete only part of a volume.

This procedure tells you how to delete an entire volume.

Procedure



1. If the volume you want to delete contains HCSS directories, unload the HCSS media. For information on unloading the media, see “Reformatting Media” on page 420.

2. Back up the data stored on the volume you want to delete.

3. Dismount the volume you want to delete.

4. At the server console prompt, type

```
LOAD INSTALL <Enter>
```

5. From the “Installation Options” menu, choose “Volume Options” and press <Enter>.

A list of existing volumes is displayed.

6. From the list of existing volumes, select the volume you want to delete and press <Delete>.

A warning similar to the following appears:

```
Volume volume_name may contain valuable data that
will be lost if you confirm "Yes" in the box that
follows this message and if you save volume changes
on exit from the volumes list.
```

```
<Press ENTER to continue>
```

7. Press <Enter> if you want to continue.
8. When the "Delete existing volume?" prompt appears, select "Yes" and press <Enter>.




Deleting volumes with INSTALL does not affect Volume *objects*. You must use NETADMIN or NetWare Administrator to delete Volume objects.

Additional Information

For more information about	Refer to
HCSS volumes	Chapter 6, "Migrating Data Using the High Capacity Storage System"
Backing up a volume	Chapter 9, "Backing Up and Restoring Data."
Deleting NDS Volume objects	"Deleting Objects from the Directory Tree" on page 86
Using NetWare Administrator to delete a volume	"Deleting Objects Using NetWare Administrator" on page 88
Using NETADMIN to delete a volume	"Deleting Objects Using NETADMIN" on page 89

Renaming Volumes

Procedure



1. **Dismount the volume you want to rename.**

2. **At the server console prompt, type**

LOAD INSTALL <Enter>

3. **From the “Installation Options” menu, choose “Volume Options” and press <Enter>.**

All existing volumes are listed.

4. **Select the volume whose name you want to change and press <Enter>.**

The “Volume Information” screen is displayed.

5. **Use the arrow keys to highlight the “Name” field, and press <Enter>.**

6. **Backspace to erase the old name; then type in a new name and press <Enter>.**



Do not change the name of volume SYS: to another name. A volume called SYS: is mandatory.

7. **Press <Esc> twice; then press <F10> to save the volume information.**

If you have not dismounted the volume it is automatically dismounted now.

8. **When prompted, answer “Yes” to mount the volume with the new name.**

9. **To set the Volume object name and context in NetWare Directory Services, you are prompted to log in to the Directory.**

10. **After you log in, verify that the displayed context and Volume object name are correct.**



Note

Renaming a volume with INSTALL creates a Volume object with the new name. It does not delete the Volume object with the old name.

11. Use NETADMIN or NetWare Administrator to delete the old Volume object from the Directory tree.

Additional Information

For more information about	Refer to
Dismounting volumes	"Mounting and Dismounting Volumes" on page 518
Using NetWare Administrator to delete the old Volume object	"Deleting Objects Using NetWare Administrator" on page 88
Using NWADMIN to delete the old Volume object	"Deleting Objects Using NETADMIN" on page 89

Setting Up a Volume to Store Non-DOS Files

By default, NetWare volumes support DOS naming conventions. To store non-DOS files (such as for Macintosh, OS/2, or UNIX files) on a NetWare volume, you must load the appropriate name space NLM and add the add name space support to that volume.



Important

Each name space added to a volume requires additional server memory. If you add name space support to a volume and do not have enough memory, that volume cannot be mounted.

Calculating Memory Required for Name Space Support

Use the following formula to calculate the additional memory required for each added name space on a non-DOS volume:

$$0.032 \times \text{volume_size}(\text{in MB}) / \text{block_size}(\text{in MB})$$

Round the result to the next highest megabyte.

For example, adding Macintosh name space to a 100MB volume with a block size of 4 KB would require 1 MB of additional memory, as shown:

$$0.032 \times 100 \text{ MB} / 4 = 0.8 \text{ MB}$$

Adding a Name Space

Prerequisites

Checklist



- A mounted volume
- Sufficient memory

Procedure

Procedure



1. Load the appropriate name space NLM by typing

```
LOAD [path]name_space <Enter>
```

For example, to load the name space module for Macintosh support, type

```
LOAD MAC.NAM <Enter>
```

2. Add name space support to the volume by typing

```
ADD NAME SPACE name to volume_name <Enter>
```

Replace *name* with the name space NLM. Replace *volume_name* with the name of the volume that will store the non-DOS files.

For example, to add the Macintosh name space to the MACVOL volume, type

```
ADD NAME SPACE MAC TO MACVOL <Enter>
```



You only need to add a name space to a volume once. You don't need to add it each time the server comes up. The name space module autoloads each time the server comes up.

To verify that the name space has been added, view a list of all added name spaces by typing

ADD NAME SPACE <Enter>

3. To see a list of current volumes and name spaces, type

VOLUMES <Enter>

Additional Information

For more information about	Refer to
ADD NAME SPACE command	"ADD NAME SPACE" in <i>Utilities Reference</i>
Name space	"Name space support" in <i>Concepts</i>

Removing Name Space Support

Once a name space is added to a volume, the name space can be removed either by deleting the volume and re-creating it or by using VREPAIR.

To use the VREPAIR method, see the next section, "Repairing Volumes." To delete a volume, see "Deleting Volumes" on page 519.

Repairing Volumes

Typically, you can't mount a volume if it has even minor damage. Occasionally, however, a damaged volume mounts, but causes errors in the process.

Use VREPAIR to correct volume problems or to remove name space entries from File Allocation Tables (FATs) and Directory Entry Tables (DETs).

You can run VREPAIR on a bad volume while other volumes are mounted. Following are typical instances when VREPAIR can help:

- ◆ A hardware failure either prevented a volume from mounting or caused a disk read error.



Although VREPAIR can't fix hardware problems, VREPAIR can sometimes fix related volume damage.

- ◆ A power failure caused a corrupted volume.
- ◆ The server console displays a mirroring error when the server boots. This mirroring refers to the two copies of FATs and DETs that the operating system keeps (if disks are mirrored, NetWare keeps four copies).

If a volume fails to mount as the server is booting, VREPAIR loads automatically and attempts to repair the volume.

When VREPAIR autoloads, it uses the default options. If you want to use an alternate option, load VREPAIR manually and set the alternate option before running VREPAIR.



If you don't want VREPAIR to automatically repair a volume that won't mount, use the "Automatically Repair Bad Volumes" SET parameter to change the default. (See "SET" in Utilities Reference.)


Prerequisites



- The volume you want to repair must be dismounted.
- If the volume to be repaired has name space support, the corresponding VREPAIR name space module (*V_namespace.NLM*) must be located in either the SYS:SYSTEM directory or in a search path directory.

Example modules include V_MAC.NLM and V_OS2.NLM.

Procedure



1. At the server console prompt, type

```
LOAD VREPAIR [volume name] [logfile name] <Enter>
```

(Optional) Replace *volume name* with the name of the volume to repair. If there is only one volume that is dismounted, you don't need to specify this parameter, since VREPAIR will attempt to repair that volume.

(Optional) If you want to save the error log, replace the *logfile name* with the name of the file you want VREPAIR to create. VREPAIR creates a log of errors it finds. VREPAIR displays the errors on screen and will write them to a file if you specify a filename.

When you load VREPAIR, an "Options" menu is displayed.

2. Accept the default options, or select alternate options, as appropriate.

The first time you try to repair a volume, accept the default options. If the default options fail to repair the volume, select alternate options.

2a. To accept the default options, continue with Step 3.

2b. To set alternate options at the "Options" menu, choose "Set VRepair Options" by typing

```
2 <Enter>
```

For more information on the options that are displayed, see "VREPAIR" in *Utilities Reference*.

3. To begin the repair process, choose “Repair A Volume” from the “Options” menu.

- ◆ If more than one volume is dismounted, select the volume to repair from those listed.
- ◆ If only one volume is dismounted, VREPAIR assumes it is the volume that needs repairing and begins the repair.

As the volume is being repaired, the server console screen displays a message indicating VREPAIR activity.

4. (Optional) Modify error log settings after the repair has started.

If VREPAIR finds many errors during the repair process, you might want to change some of the run-time error settings. To modify these settings after the repair has started, press <F1> to display the “Current Error Settings” menu.

- ◆ Select Option 1 if you do not want VREPAIR to pause after each error.
- ◆ Select Option 2 if you want VREPAIR to log errors in a text file.
- ◆ Select Option 3 to stop the repair of the volume.
- ◆ Select Option 4 to continue with a volume repair after you have stopped it.

5. When the repair is complete, answer “Y” when prompted to write repairs to the disk.

6. If VREPAIR has found errors, run VREPAIR again by repeating Steps 2 through 6. Repeat until VREPAIR finds no errors.

If you are unable to mount the volume after running VREPAIR a number of times, you must delete the volume, re-create the volume using INSTALL, and then restore the data from backups.

Additional Information

For more information about	Refer to
Using VREPAIR	"VREPAIR" in <i>Utilities Reference</i>
Dismounting a volume	"Mounting and Dismounting Volumes" on page 518
Creating a volume	"Creating Volumes" on page 515
Setting the Automatically Repair Bad Volume parameter	"SET" in <i>Utilities Reference</i>
Setting and viewing the console search path	"SEARCH" in <i>Utilities Reference</i>


Adding a Segment to an Existing Volume

Prerequisites



- An existing NetWare partition with free space. (To create a NetWare partition, see "Creating NetWare Disk Partitions" on page 561.)

Procedure



1. **At the server console prompt, type**
LOAD INSTALL <Enter>
2. **From the "Installation Options" menu choose "Volume Options."**
3. **Press <Insert> or <F3> to view existing volume segments.**
4. **Select a segment that has free space (no volume assignment) and press <Enter>.**



If no free space exists, you can't add a segment to a volume.


5. To add this segment to an existing volume, select “Make This Segment Part of Another Volume” and press <Enter>.

A list of existing volumes appears.

6. Select the volume you want to add this segment to, and press <Enter>.
7. Verify that the segment has been added in the “Volume Disk Segment List.” Then press <Esc>.
8. To save volume assignments to disk, press <F10>.

Modifying the Size of a Volume

A volume can have multiple segments spanning multiple hard disks. This arrangement allows you to add a hard disk when you need to expand a given volume.




Warning

You can add segments to a volume without destroying data, but removing any segment from a volume destroys *all* of the data on the volume.

Keep the following requirements in mind if you plan to expand volumes:

- ◆ The maximum number of segments allowed per volume is 32.
- ◆ The maximum number of segments you can create on one hard disk is 8.
- ◆ The maximum number of volumes allowed is 64.

Procedure



Procedure

1. At the server console prompt, type

LOAD INSTALL <Enter>

2. From the “Installation Options” menu choose “Volume Options.”

All existing volumes are listed.

3. Press <Insert>.

The “Volume Disk Segments List” is displayed.

4. Select the volume segment to modify and press <Enter>.

The status of that segment is displayed.

5. From the status display, determine what kind of modification you can do:

- ◆ If the status of the volume segment is “N” (new), you can change the size of the segment. For instructions, see Step 6.
- ◆ If the status of the volume segment is “E” (existing), you can’t change the size of the segment.

You can increase the volume size by adding new segments. (See “Adding a Segment to an Existing Volume” on page 528.)

- ◆ If you want to reduce the size of the segment or volume, you must delete the entire volume with all segments, and then re-create the volume. See “Deleting Volumes” on page 519.

6. Type the new volume size in megabytes (or a fraction thereof) on the “Disk Segment Parameters” screen.

7. Press <Esc> twice; then press <F10> to save the settings.

Using a CD-ROM as a NetWare Volume

CDROM allows the NetWare server to use a CD-ROM disc as a NetWare volume.



Treat the CD-ROM as a read-only volume. Do not enable block suballocation or use file compression on the volume. These actions will corrupt the CD-ROM volume index data.

If you do so by mistake, use the CD command to rebuild the volume’s index file. For command syntax, see “CD” in *Utilities Reference* or type CD HELP at the server console prompt.

NetWare 4.1 supports CD-ROMs that are mounted with the MAC and NFS name spaces.

The CDROM.NLM program supports High Sierra and ISO 9660 formats. It does not support HFS (Apple) file systems.

Mounting a CD-ROM as a NetWare Volume

To mount a CD-ROM disc as a NetWare volume, follow these steps.

Prerequisites



- Volume SYS: mounted
- An installed host bus adapter (HBA) that is NetWare compatible and supports CD-ROM devices



There are no serial, parallel, or IDE device drivers for CD-ROM support. Use a SCSI adapter.

- The NetWare Peripheral Architecture™ (NPA) module, NWPA.NLM.
- The disk driver file and necessary support modules




Some disk drivers consist of more than one file and some HBA devices require additional support modules for CD-ROM functionality. These files should accompany the HBA. For specific file requirements, consult your adapter documentation.

NetWare 4.1 includes the third-party HBA device drivers that are Novell-certified for NetWare 4.1. If the HBA is certified to support CD-ROM devices, the necessary support modules are also included with NetWare.

You can copy these files during installation, or you can copy them from the NetWare distribution media using NWXTRACT. (For help, see “NWXTRACT” in *Utilities Reference*.)

Procedure



1. At the server console prompt, load the disk driver by typing

```
LOAD [path]disk_driver <Enter>
```

Replace *disk_driver* with the name of the disk driver specified in the HBA documentation.

You may be prompted to supply command line values such as a port or slot number for the HBA. Consult your HBA documentation for this information.

For example, to load the disk driver for the Adaptec AHA-1522 SCSI HBA, type

```
LOAD [path]AHA1520.DSK <Enter>
```

```
LOAD [path]ASPICD.DSK <Enter>
```

Some files may be automatically loaded. In this example, ASPITRAN.DSK is automatically loaded when AHA1520.DSK is loaded.

If your disk driver autoloads the ASPITRAN.DSK driver, you need to load either one of the following: ASPICD.DSK or CDNASPI.DSK. In this example, we used ASPICD.DSK.

2. Load the NWPANLM (NetWare Peripheral Architecture) driver by typing

```
LOAD NWPANLM <Enter>
```

This is the device-independent software that interfaces with the Media Manager.



CDROM.NLM will not load unless the NWPANLM driver has been loaded.

3. Load CDROM.NLM by typing

```
LOAD CDROM <Enter>
```



When a CD-ROM volume is being mounted or a CD-ROM disc is being changed, CD-ROM devices might be deactivated. Do not be alarmed. This deactivation occurs because device configuration information is being updated.

4. Mount the CD-ROM as a volume by typing

```
CD MOUNT [object number] | [volume name]  
          [name space] [/option] <Enter>
```



The *object number* is the Media Manager object number (hereafter called the *object number*).

Replace *object number* with the object number or replace *volume name* with the volume name of the CD-ROM disc. (You can use the CD DEVICE LIST or CD VOLUME LIST command to see the object numbers and volume names.)

If appropriate, add name space support for the volume by replacing *name space* with /MAC or /NFS name space, or use both name spaces.

If appropriate, replace *option* with any of the following

- ◆ /R: Rebuilds the index file for a CD-ROM volume.
- ◆ /G= *group number*: Sets the default volume group access rights while mounting the volume.
- ◆ /X: Excludes a directory, at the root level, on the CD-ROM.
- ◆ /I: Mounts the volume even if importing errors are found and not all files are available.

For more information on the CD MOUNT command and its options, see “CD” in *Utilities Reference*.

For example, to mount the NetWare_41 CD-ROM, type

```
CD MOUNT NETWARE_41 <Enter>
```



It may take several minutes to mount the volume, depending on the size of the CD-ROM and on the speed of your computer.

5. (Optional) To mount the CD-ROM as a NetWare volume each time the server comes up, do the following:

5a. Edit your STARTUP.NCF file to add the disk driver.

For example, to load the driver described in Step 1, add the following lines to the STARTUP.NCF file:

```
LOAD [path]disk_driver  
LOAD [path]ASPICD.DSK  
LOAD [path] NWPA.NLM
```

If your disk driver autoloads the ASPITRAN.DSK driver, load either one of the following: ASPICD.DSK or CDNASPI.DSK. In this STARTUP.NCF file example, we used ASPICD.DSK.

5b. Edit your AUTOEXEC.NCF by adding the following:

```
LOAD CDROM  
CD MOUNT [object number] | [volume name]  
[name space] [option]
```

Additional Information

For more information about	Refer to
Using the CD command	"CD" in <i>Utilities Reference</i> , or type CD HELP at the console
Using CDROM.NLM	"CDROM" in <i>Utilities Reference</i>
Editing .NCF files	"Creating or Editing a Server Batch (.NCF) File" on page 459
Novell-certified device drivers for NetWare 4	Call 1-800-NETWARE (1-800-637-9273)

Administering a CD-ROM as a NetWare Volume

Using the CD console commands associated with CDROM, you can do the following:

- ◆ Monitor the status of CD-ROM devices and NetWare volumes
- ◆ Change the media in a CD-ROM device
- ◆ Display the root directory on a NetWare CD-ROM volume
- ◆ Mount and dismount CD-ROM volumes
- ◆ Add and delete group names

For information on command syntax and for examples of these procedures, see “CD” in *Utilities Reference*.

You can also type

```
CD HELP <Enter>
```

at the server console prompt to get information.

Managing Server Hard Disks

This section describes how to determine available disk space, how to increase available disk space by compressing files and purging files, and how to add a hard disk to a NetWare server.

Checking Available Disk Space with NDIR

You should monitor available disk space regularly and keep a log so you can track disk usage over time. This information helps you make the best use of your disk space management options, such as adding a new hard disk, compressing files, and migrating data to an offline system such as an optical disc library.

Prerequisites

Checklist



- A workstation running DOS 3.30 or later, or OS/2 v2.x.

Procedure

Procedure



1. At the workstation, type

```
NDIR volume: /VOL <Enter>
```

Statistics for the volume you selected appear on the screen.

For more information about the NDIR utility, see “NDIR” in *Utilities Reference*.

2. Write down the percentage shown for Space remaining on volume.

3. Repeat this procedure weekly for each volume.

Tracking this percentage over time helps you predict when a volume may run out of disk space.

Suggestion



You may want to enter the disk space information into a spreadsheet to create a graph of disk usage over time.

Using File Compression

File compression allows you to store more information on a server hard disk.

File compression is enabled or disabled at the volume level. Once enabled, it can be turned on and off for individual files and directories by setting compression attributes with FLAG, FILER, or the NetWare Administrator graphical utility.

See “Setting Compression for a File or Directory” on page 539.

When you install NetWare 4.1, the default is to enable file compression. However, you can override the default and disable file compression during installation.

If you disable file compression, you can enable it later. Once enabled, however, file compression cannot be turned off unless you first re-create the volume.

Managing File Compression on Files and Directories

It isn't necessary to separate applications from data for file compression, because the Days Untouched Before Compression SET parameter can eliminate the compression of applications that are used regularly.

This parameter specifies the number of days that must pass without access to a file before the file can be compressed. The parameter uses the last accessed date to gauge whether a file should be compressed or not.

To avoid the overhead of uncompressing files that don't compress well, the system calculates the compressed size of a file before actually compressing it.

If no disk space will be saved by compression, or if the size difference does not meet the value specified by the Minimum Percentage Compression Gain SET parameter, the file is not compressed.

For a file to be uncompressed, there must be enough free space on the volume to accommodate the uncompressed file size.

Enabling File Compression on a Volume

When you install NetWare or create a new volume, file compression is enabled by default. You may override the default and disable file compression when the volume is created.

The following procedure explains how to enable file compression on a volume for which compression was disabled.



Once file compression is enabled for a volume, you can't disable it unless you first re-create the volume.



Warning

If you have a CD-ROM volume mounted as a NetWare volume, you should treat it as a read-only volume.

Do not use file compression on a CD-ROM volume. This will corrupt the CD-ROM volume index data. If you do so by mistake, use the CD command to rebuild the volume's index file. For command syntax, see "CD" in *Utilities Reference*.

Prerequisites



Checklist

- 1 MB of available RAM on the server during the time file compression is running.

Procedure



Procedure

1. **At the server console prompt, type**

LOAD INSTALL <Enter>

2. **Select "Volume Options" from the "Installation Options" menu.**

A list of volumes appears.

3. **Select the volume for which you want to enable file compression.**

The Volume Information screen appears, showing the current settings for the volume.

4. **Use the Down-arrow key to move the cursor to the File Compression field and press <Enter>. Change the setting from Off to On.**

5. **Press <Esc> twice to save the new value and redisplay the Installation Options menu.**

Setting Compression for a File or Directory

Once file compression is enabled for a volume, turn it on or off for individual files and directories by using FLAG, FILER, or NetWare Administrator.

You may want to group the files you want to compress into different directories from those you don't want to compress. For example, if you have data you don't access often, such as quarterly reports, you can create a directory called QUARTER and set it for immediate compression. Then when files are moved into the QUARTER directory, they stay compressed until the next report time, taking up much less disk space.

Using FLAG to Set Compression Attributes

You can use the FLAG utility from a workstation to set attributes for a directory or file. The attributes are Immediate Compress (IC) and Don't Compress (DC).

Prerequisites



- A workstation running DOS 3.30 or later, or OS/2 v2.x
- File compression enabled for the volume
- Modify File or Directory right to the file or directory you are setting

Procedure



1. At the workstation, type

FLAG <Enter>

The attributes for the current drive appear.

If you are not in the directory where you want to set the file compression attributes, enter the drive letter or the complete path after the FLAG command.

2. Enter the file compression option you want to apply to the file or directory.

For example, if you want the files in the subdirectory LETTERS to be compressed each time a file is created or moved into the directory, type the Immediate Compress option:

FLAG SYS:DOC\LETTERS IC <Enter>

If you have an APPS directory for files that you do not want compressed at any time, type the Don't Compress option:

FLAG SYS:APPS DC <Enter>

For more information, see "FLAG" in *Utilities Reference*.

Using FILER to Set Compression Attributes

You can set file compression attributes with the FILER menu utility. The attributes are Immediate Compress (IC) and Don't Compress (DC).


Online help appears at the bottom of each screen to help you use FILER, and you can press <F1> for more detailed help.

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 512 KB of memory available on the workstation.
- File compression enabled for the volume.
- Modify File or Directory right to the file or directory you are setting.

Procedure



1. At the workstation, type

FILER <Enter>

A list of available options appears.

2. Select “Manage Files and Directories” from the “Available Options” menu and press <Enter>.

A list of files appears in the “Directory Contents” window.

3. Select the file you want to set file compression attributes for, and press <F10>.

4. Select “View/Set File Information” from the “File Options” menu and press <Enter>.

5. Select the “Attributes” list in the “File Information” display and press <Enter>.

The “Current File Attributes” window appears.

6. Press <Insert> to see other file attributes.

7. Select “Immediate Compress” or “Don’t Compress” from the “Other File Attributes” window and press <Esc>.

The compression attribute you selected (Ic or Dc) appears in the attributes list.

8. To exit FILER, press <Alt>+<F10>.

Using NetWare Administrator to Set Compression Attributes

You can set file compression attributes with the NetWare Administrator. The attributes are Immediate Compress (IC) and Don't Compress (DC).

Prerequisites

Checklist



- A 386 or 486 workstation running NetWare Administrator
- 6 MB of available RAM on the Windows workstation, or 12 MB of available RAM on the OS/2 workstation
- File compression enabled for the volume
- Modify File or Directory right to the file or directory you are setting

Procedure

Procedure



- 1. Choose the “NetWare Administrator” icon from the Windows Program Manager or the OS/2 desktop.**
- 2. Select the file you want to set file compression attributes for.**

To see a file in the NetWare Administrator, choose a volume object and then choose the directory the file resides in.
- 3. From the “Object” menu, choose “Details.”**
- 4. Choose the “Attributes” button.**
- 5. From the “File Attributes” list, choose “Immediate Compress” or “Don't Compress.”**
- 6. Choose “OK.”**


Additional Information

For more information about	Refer to
Changing SET parameters	“SET” and “SERVMAN” in <i>Utilities Reference</i>
File compression	“File compression” in <i>Concepts</i>
Viewing file compression attributes for a directory or a file	“Viewing Other Information about a Directory or File” on page 156 “FLAG” and “NDIR” in <i>Utilities Reference</i>
Viewing file size	“Viewing Other Information about a Directory or File” on page 156 “FILER” and “NDIR” in <i>Utilities Reference</i>

Suspending File Compression

You can suspend file compression temporarily by using the Enable File Compression SET parameter. Follow this procedure.

Procedure



1. **At the server console prompt, load SERVMAN by typing**

```
LOAD SERVMAN <Enter>
```

2. **Select “File System” from the “Categories” menu.**
3. **Select “Enable File Compression” and change the value to OFF.**

If this parameter is set to OFF, any files that would normally be compressed are queued and will be compressed as soon as the parameter is reset to ON.

If you have many files to compress, server performance may be affected when you turn file compression back on. If so, make sure you set file compression to be activated during nonpeak hours. See the Compression Daily Check Stop Hour and the Compression Daily Check Starting Hour parameters in Table 7-3.

4. Press <Esc> twice to display the “Update Options” menu.
5. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

You are prompted to update the STARTUP.NCF file.

6. If the SET parameters are in the STARTUP.NCF file, press <Enter> to update the values in the file. Otherwise, press <Esc>.

You are prompted to update the AUTOEXEC.NCF file.

7. If the SET parameters are in the AUTOEXEC.NCF file, press <Enter> to update the values. Otherwise, press <Esc>.

Changes to the parameters take effect immediately.

If you do not update either of the .NCF files, changes to the parameters last only until you reboot the server.

Using SET to Control File Compression

You can control various aspects of file compression by using the file compression SET parameters. For example, the Compression Daily Check Starting Hour parameter determines when the file compressor begins scanning volumes for files that need to be compressed. File compression SET parameters do not affect volumes on which compression is disabled.

To change SET parameters, execute the SET command at the server console prompt, or use the SERVMAN utility.

For more information about executing the SET command, see “Using the SET Command to Change File Compression SET Parameters” on page 547.

For more information about using SERVMAN, see “Using SERVMAN to Change File Compression SET Parameters” on page 548.



In most cases, the default values for the SET parameters are the most efficient. We recommend you read this section before changing the default values.

Add the file compression SET parameters to your AUTOEXEC.NCF or STARTUP.NCF file so that they are executed whenever the server is brought up.

You can use EDIT or INSTALL to edit the AUTOEXEC.NCF or the STARTUP.NCF file. You can use SERVMAN to change the value of parameters already in either file.

Table 7-3 lists the file compression SET parameters and their default values. The settings apply to all files and directories in compression-enabled volumes on the server. For more detailed information about SET parameters, see "SET" in *Utilities Reference*.

Table 7-3
SET Parameters for File Compression

SET Parameter and Default	Explanation
Compression Daily Check Stop Hour=6 (6:00 a.m.)	Sets the hour when the file compressor stops scanning volumes for files that need to be compressed. Hours are specified by a 24-hour clock. Supported values: 0 to 23.
Compression Daily Check Starting Hour=0 (midnight)	Sets the hour when the file compressor starts scanning volumes for files that need to be compressed. Hours are specified by a 24-hour clock. Supported values: 0 to 23.
Minimum Compression Percentage Gain=2 (%)	Specifies the minimum percentage a file must compress in order to remain compressed.
Enable File Compression=ON	Allows file compression to occur on compression-enabled volumes. If you set this parameter to OFF, compression is suspended temporarily. Files flagged for compression are queued until the parameter is reset to ON.
Maximum Concurrent Compressions=2	Specifies the number of volumes that can compress files at the same time. Increasing this setting may slow server performance.
Convert Compressed To Uncompressed Option=1	Determines how the server stores a file after the file has been accessed. <ul style="list-style-type: none"> ◆ Option 0 always leaves the file compressed. ◆ Option 1 leaves the file compressed after the first access since a file compression. The second access leaves the file uncompressed. ◆ Option 2 always leaves the file uncompressed.
Decompress Percent Disk Space Free To Allow Commit=10(%)	Specifies the percentage of free disk space required on a volume before committing an uncompressed file to disk. This prevents newly uncompressed files from filling up the volume.
Decompress Free Space Warning Interval=31 minutes 18.5 seconds	Specifies the interval for displaying warning alerts when the volume has insufficient free disk space for uncompressed files. To turn off the alerts, set the parameter to 0.

Table 7-3 *continued*

SET Parameters for File Compression

SET Parameter and Default	Explanation
Deleted Files Compression Option=1	<p>Determines how the server handles unpurged deleted files.</p> <ul style="list-style-type: none"> ◆ Option 0 doesn't compress deleted files. ◆ Option 1 compresses deleted files during the next search. ◆ Option 2 compresses deleted files immediately.
Days Untouched Before Compression=7	<p>Specifies the number of days that must pass with no access to the file before the file can be compressed.</p>


Using the SET Command to Change File Compression SET Parameters

Prerequisites



- File compression enabled for the volume

Procedure



1. **At the server console prompt, use the following format to set a file compression parameter:**

```
SET parameter = value <Enter>
```

For example, the following command specifies 6 am as the time when the system stops scanning for files to compress:

```
SET Compression Daily Check Stop Hour = 6 <Enter>
```

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

Using SERVMAN to Change File Compression SET Parameters

Prerequisites

Checklist



- File compression enabled for the volume

Procedure

Procedure



1. **At the server console prompt, type**

LOAD SERVMAN <Enter>

2. **Choose “Server Parameters” from the “Available Options” menu.**
3. **Choose “File System” from the “Categories” menu.**
4. **Change the values of the compression parameters in the “File System” menu as desired.**

Be sure to press <Enter> after each change you make.

5. **Press <Esc> twice to reach the “Update Options” menu.**
6. **Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.**

You are prompted to update the STARTUP.NCF file.

7. **If the SET parameters are in the STARTUP.NCF file, press <Enter> to update the values in the file. Otherwise press <Esc>.**


You are prompted to update the AUTOEXEC.NCF file.

8. **If the SET parameters are in the AUTOEXEC.NCF file, press <Enter> to update the values. Otherwise, press <Esc>.**

Changes to the parameters take effect immediately. You do not need to reboot the server.

Purging Files from a Disk

Purging deleted files frees disk space on the NetWare server's hard disk. You can purge files manually, or you can set up automatic purging.



Purged files cannot be salvaged.

Manually Purging Files


To purge deleted files manually, use the NetWare Administrator graphical utility, or the FILER or PURGE text utility. See "Purging Files" on page 165.

Setting Up Automatic Purging with SERVMAN

To set up automatic purging of all files on a NetWare server's hard disk, use the SERVMAN utility as described in this section.

You can also use the SET utility to set each of the relevant SET parameters manually. For more information, see "SET" in *Utilities Reference*.

Procedure



1. **At the server console prompt, type**

```
LOAD SERVMAN <Enter>
```

2. **Choose "Server Parameters" from the "Available Options" menu.**
3. **Choose "File System" from the "Categories" menu.**



4. Change the following parameters in the “File System” menu to the indicated values:

These settings cause the system to purge deleted files immediately. Therefore, the files cannot be salvaged. Be sure you have backups of important files.

Immediate Purge Of Deleted Files=ON

File Delete Wait Time=0

Minimum File Delete Wait Time=0

5. Press <Esc> twice to reach the “Update Options” menu.

6. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

The system displays the path to the AUTOEXEC.NCF file.

7. If desired, press <Enter> to update the file.

The new parameter values take effect immediately. You do not need to reboot the server.

Additional Information

For more information about	Refer to
Purging files	“Salvageable files” in <i>Concepts</i>
Using the FILER utility	“FILER” in <i>Utilities Reference</i>
Using the PURGE utility	“PURGE” in <i>Utilities Reference</i>

Adding Optical Storage for File Migration

The High Capacity Storage System (HCSS) extends the storage capacity of a NetWare server by integrating one or more optical library units, or jukeboxes, into the NetWare file system.

HCSS moves files between the server's hard disk and optical disks in a jukebox. HCSS uses hard disk free space to cache the files stored on optical disk. When the hard disk free space reaches a configurable threshold, HCSS migrates files to optical disk.

If you have installed an HCSS system on your network, you can set it to use as much or as little of the server's hard disk space as you want.

HCSS procedures are described in Chapter 6, "Migrating Data Using the High Capacity Storage System."

For more information about	Refer to
HCSS functionality	"High Capacity Storage System" and "Data migration" in <i>Concepts</i>
HCSS installation and operating procedures	Chapter 6, "Migrating Data Using the High Capacity Storage System"

Adding a Hard Disk to the NetWare Server

The procedures in this section explain how to add the following:

- ◆ An additional IDE hard disk ("Adding an Additional IDE Hard Disk" on page 552)
- ◆ An internal SCSI hard disk and adapter ("Adding an Internal SCSI Hard Disk and Adapter" on page 553)
- ◆ A SCSI hard disk to an existing subsystem ("Adding a SCSI Hard Disk to a Subsystem" on page 555)

Select the procedure that fits your situation.

Adding an Additional IDE Hard Disk

To add an additional internal IDE hard disk, complete the following steps.

Prerequisites



- All users logged out of the server
- Access to the documentation that came with the hard disk
- Access to the documentation that came with the computer

Procedure



1. Bring down the server by typing

DOWN <Enter>

EXIT <Enter>

2. Turn off power to the server.

3. Install and cable the hard disk.

Make sure that the jumpers are configured so that one disk is the master and one the slave.

Refer to the documentation that came with the hard disk for more information.

4. Configure the computer to recognize the new disk.

For example, you may need to run the CMOS or EISA configuration utility. Refer to the documentation that came with the computer for information about configuration methods and requirements.

5. Reboot the server.

6. At the server console prompt, type

LOAD INSTALL <Enter>

7. Use **INSTALL** to create a NetWare partition.

If you need help, see “Creating NetWare Disk Partitions” on page 561.

8. Create a new volume or designate the partition as a new segment of an existing volume.

If you need help, see “Creating Volumes” on page 515 or “Adding a Segment to an Existing Volume” on page 528.

Adding an Internal SCSI Hard Disk and Adapter


To install an internal SCSI hard disk and new adapter, complete the following steps.

Prerequisites



- All users logged out of the server
- Access to the documentation that came with the hard disk and adapter board
- Access to the documentation that came with the computer

Procedure



1. Bring down the server by typing

DOWN <Enter>

EXIT <Enter>

2. Turn off power to the server.

3. Configure and install the new adapter.

- ◆ Some adapters require you to set jumpers as part of the configuration process. Set jumpers on the adapter if necessary. Refer to the documentation that came with the board to learn what jumper settings may be required.
- ◆ Make sure both ends of the SCSI bus are terminated. Refer to the documentation that came with the adapter for more information.

4. **Install the hard disk, connect the power cable to the disk, and cable the disk to the new adapter.**
5. **If necessary, configure the computer to recognize the new disk and adapter.**

Refer to the documentation that came with the computer for more information about configuration methods and requirements.

6. **Reboot the server.**
7. **At the server console prompt, type**

LOAD INSTALL <Enter>

8. **Use INSTALL to load the disk driver.**

For more information about loading drivers, see "Loading Disk Drivers" on page 559.

9. **Use INSTALL to create a NetWare partition.**

If you need help, see "Creating NetWare Disk Partitions" on page 561.

10. **Do one of the following steps:**

- 10a. **Mirror the disk.**

For information about disk mirroring, see "Mirroring and Duplexing a Hard Disk" on page 565.

- 10b. **Create a new volume.**

If you need help, see "Creating Volumes" on page 515.

- 10c. **Designate the partition as a new segment of an existing volume.**

See "Adding a Segment to an Existing Volume" on page 528.

Adding a SCSI Hard Disk to a Subsystem

A subsystem contains additional hard disks. To add a SCSI hard disk to an existing adapter in a subsystem, complete the following steps.

Prerequisites



- All users logged out of the system
- Access to the documentation that came with the hard disk
- Access to the documentation that came with the computer

Procedure



1. Bring down the server by typing

DOWN <Enter>

EXIT <Enter>

2. Turn off power to the subsystem.

3. Install the hard disk and cable it to the adapter.

4. Set the SCSI ID for the drive and terminate the SCSI bus.

Refer to the documentation that came with the hard disk for more information about SCSI IDs and termination requirements.

5. Turn on the power to the disk subsystem.

6. Reboot the server.

7. At the server console prompt, type

LOAD INSTALL <Enter>

8. Use **INSTALL** to create a NetWare partition.

If you need help, see “Creating NetWare Disk Partitions” on page 561.

9. Do one of the following steps:

9a. Mirror the disk.

For information about disk mirroring, see “Mirroring and Duplexing a Hard Disk” on page 565.

9b. Create a new volume.

If you need help, see “Creating Volumes” on page 515.

9c. Designate the partition as a new segment of an existing volume.

See “Adding a Segment to an Existing Volume” on page 528.

Replacing a Hard Disk

If a hard disk becomes unreliable or unusable, follow this procedure to remove the disk from the network.

Once you start this procedure, you must complete it; otherwise, the system hangs as it looks for the missing hard disk.

If you have been backing up your data consistently and verifying its integrity, you will be able to reload data for the volumes affected by the disk failure.




Although some hard disks (called “hot plug” devices) can be removed while the system is running, we do not recommend this practice.

If you do try to replace a mirrored hot-plug disk while the system is running, you must unmirror the disk first or run SCAN FOR NEW DEVICES as soon as the disk is removed. Failure to do one of these steps could cause the loss of all data on the disk’s mirrored partner when you install the replacement disk.

See “Unmirroring Hard Disks” on page 567 and “SCAN FOR NEW DEVICES” in *Utilities Reference*.

Procedure



1. Make a backup copy of the data on the hard disk and verify its validity.

Make sure the backup copy contains uncorrupted versions of all files and directories on the hard disk.

2. If the disk is mirrored, dismount the volumes that reside on the disk by using either the DISMOUNT command or the INSTALL NLM.

- ◆ To use DISMOUNT, repeat the following command for each volume on the disk and then go to Step 3.

```
DISMOUNT [volume_name] <Enter>
```

- ◆ To dismount the volume using INSTALL, follow the procedure in “Mounting and Dismounting Volumes” on page 518. Repeat the process for all remaining volumes on the disk and then continue with Step 3.

3. Unmirror the hard disk if it is mirrored.

See “Unmirroring Hard Disks” on page 567.

4. At the server console, bring down the server by typing

```
DOWN <Enter>
```

```
EXIT <Enter>
```


5. Turn off the system containing the hard disk to be replaced.

6. Remove the hard disk and install the replacement.

Use the instructions that came with the hard disk and see “Adding a Hard Disk to the NetWare Server” on page 551.

7. Turn on power and reboot the server.

8. If the replacement disk has been used before, delete any existing partitions on the replacement disk.



Deleting disk partitions destroys *all* data on the deleted partitions. Be sure you delete the partitions on the new disk, not an existing disk.

9. Create a partition on the replacement disk.

See "Creating NetWare Disk Partitions" on page 561.

10. Complete one of the following steps:

10a. If the disk was not mirrored, create the volumes that previously resided on the hard disk and restore the data from a backup.

See "Creating Volumes" on page 515.

10b. If the hard disk was mirrored, remirror the disk.

See "Mirroring and Duplexing a Hard Disk" on page 565.



If you replaced the "Not Mirrored" disk rather than the "Out of Sync" disk, recover the volume data from the "Out of Sync" disk before remirroring the disks.

Loading Disk Drivers

After you add or replace a hard disk on your NetWare server, you must load the corresponding disk driver.

Loading a disk driver enables communication between the disk controller and the server's CPU.

Load the disk driver once for each adapter you want to support.

Most NetWare 4 disk drivers have a description file that appears on the screen as you select the driver. Refer to these descriptions or to Table 7-4 to determine which driver to load.


Table 7-4
NetWare Disk Drivers

Computer architecture	Device Interface	Disk driver
Industry Standard Architecture (ISA)	AT, ESDI, MFM	ISADISK
	IDE (ATA)	IDE
	SCSI	SCSI
	ISA Vendor Proprietary	See Vendor
Microchannel (MCA)	ESDI (PS/2)	PS2ESDI
	IBM SCSI	PS2SCSI
Extended Industry Standard Architecture (EISA)	AT class	ISADISK
	IDE (ATA)	IDE
	SCSI	SCSI
	EISA Vendor Proprietary	See Vendor
Peripheral Component Interconnect(PCI)		See Vendor
Personal Computer Memory Card International Association (PCMCIA)		See Vendor



For maximum performance, use the ISADISK driver only if your interface is the old-style AT, EDSI or MFM.

Procedure



1. **At the server console prompt, load INSTALL by typing**
LOAD INSTALL <Enter>
2. **Select “Driver Options” from the “Installation Options” menu.**
3. **Choose “Configure Disk and Storage Device Drivers” from the “Driver Options” menu.**
4. **Choose “Load an Additional Driver” from the “Additional Drivers Actions” menu.**
5. **Select the disk driver you want to load and press <Enter>.**

For example, if you are using an AT controller, the disk driver filename is ISADISK.DSK and its description filename is ISADISK.DDI.



Some drivers do not have a description file (a configuration file that's appended to the driver). These drivers have to be loaded manually, at the system console. To load these drivers, follow the screen prompts or press <F1> for help.

6. **Choose “Select/Modify Driver Parameters.”**

The cursor becomes active in the parameter window for the selected driver. Default parameters may be displayed in the fields.

7. **Use the arrow keys to move from field to field. Press <Enter> to change a value in a field.**

When you move the cursor to a field, the system displays a help window that describes the field.

When you press <Enter> to change a value, the system usually displays a pop-up list of values from which you select the desired one.

8. **When finished setting parameter values, press <F10>.**

9. **Select “Save Parameters and Load Driver.”**

The system displays a message asking if you want to load another driver.

10. If you want to load an additional driver, select “Yes” and repeat Steps 5 through 9.

You can load the same driver more than once if you have additional disk controller boards of the same type. (Remember to check controller board settings for conflicts).

11. To exit INSTALL, press <Alt>+<F10>.

Additional Information

For more information about	Refer to
Expansion bus computer architecture	“Expansion bus” in <i>Concepts</i>
SCSI or IDE buses	“SCSI bus” and “IDE” in <i>Concepts</i>
Novell-certified device drivers for NetWare 4	Call 1-800-NETWARE (1-800-637-9273)

Creating NetWare Disk Partitions

NetWare partitions can be created on any hard drive and can coexist with other partitions such as DOS, OS/2, and UNIX.


If you have partitions from previous versions of NetWare that you are no longer using, you can delete them and create a new NetWare partition. (See “Deleting NetWare Disk Partitions” on page 563 for more information.)



Note

Some machine vendors such as COMPAQ* create a small partition that setup and configuration utilities can be run from. Don't delete this partition.

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Select “Disk Options” from the “Installation Options” menu.

3. Select “Modify Disk Partitions and Hot Fix” from the “Available Disk Options” menu.

If you have more than one disk, the “Available Disk Drives” list appears. If this list appears, select the disk you want to partition.

A list of existing partitions appears, along with the “Disk Partition Options” menu.

4. Select “Create NetWare Disk Partition” from the “Disk Partition Options” menu.

NetWare allows only one NetWare partition per disk (unless the machine vendor presents multiple hard drives as one logical drive). If the disk has more than one free partition area, select an area for the NetWare partition.

A screen containing disk partition and Hot Fix information with their default values appears. The cursor is on the “Partition Size” field.

5. If you don’t want to use the whole partition for NetWare, type a smaller value in the “Partition Size” field.

The difference between the original size and the smaller size is added to the free space on the disk.

6. To change the size of the data area, use the arrow keys to move the cursor to one of the “Data Area” fields, and then type in a new value.

One field contains the size of the data area in blocks, the other in megabytes. When you change the value in one field, the system changes the corresponding value in the other field.

7. **To change the size of the Hot Fix Redirection Area, use the arrow keys to move the cursor to one of the “Hot Fix Redirection Area” fields, and then type in a new value.**

One field contains the size of the area in blocks, the other as a percentage of the partition. When you change the value in one field, the system changes the corresponding value in the other field.

8. **Press <Esc> to display the “Create NetWare Partition?” window.**
9. **Select “Yes.”**
10. **Press <Esc> twice to redisplay the “Available Disk Options” menu.**
11. **To create NetWare partitions on additional disks, repeat Steps 3 through 10 for each disk. When finished, go to Step 12.**
12. **If you want to mirror a disk, see “Mirroring and Duplexing a Hard Disk” on page 565. Otherwise, press <Esc> to exit.**

Deleting NetWare Disk Partitions



Note

The partition table displays partitions such as OS/2, UNIX, and XENIX* as “Unknown Partition Type #.” Don’t delete these unknown partition types unless you know what is on them.

Some machine vendors such as COMPAQ create a small partition that setup and configuration utilities can be run from. Don’t delete this partition.

Prerequisites



Checklist

- A backup copy of data on the partition, if needed
- All volumes on the partition dismounted and deleted

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Select “Disk Options” from the “Installation Options” menu.

3. Select “Modify Disk Partitions and Hot Fix” from the “Available Disk Options” menu.

If you have more than one disk, the “Available Disk Drives” list appears. If this list appears, select the disk you want from the “Available Disk Drives” list.

A list of disk partitions on the selected disk appears at the top of the screen. The “Disk Partition Options” menu appears at the bottom of the screen.

4. Select “Delete Any Disk Partition” from the menu.

A list of the available disk partitions is displayed in a pop-up menu.

5. Select the disk partition you want to delete.

If you did not delete existing volumes from the disk partition, a message similar to the following appears:

WARNING: The selected disk partition may contain valuable data; all data on the partition will be lost.

<Press ENTER to continue>

6. If you are ready to delete the partition, press <Enter>; otherwise, press <Esc>.

If you press <Enter>, a confirmation box is displayed containing the message “Delete Disk Partition?”

7. Select “Yes.”

The space occupied by the deleted disk partition now appears as Free Space in the list of disk partitions at the top of the screen. You can create a new disk partition in this free space.

Mirroring and Duplexing a Hard Disk

NetWare 4 enables you to protect data from hard disk failure by duplicating, or “mirroring,” one hard disk’s data on one or more other hard disks.

When you mirror hard disks over different disk channels or host bus adapters, this is called *duplexing*.

The INSTALL processes for mirroring and duplexing are the same. The term “mirroring” is used in all menus to refer to both mirroring and duplexing.



If you want mirroring with IDE/ATA disks, you must duplex them by mirroring them on separate host adapters or on different ports of a multiport IDE/ATA host adapter, with each port or adapter configured to support a single drive. This requirement is necessary because of a design limitation of the IDE drives that prevents them from continuing normally after one of the drives fails.

Following are some considerations to keep in mind when using mirrored disks:

- ◆ Mirroring requires the NetWare partitions to be exactly the same size. If they are not the same size, NetWare adjusts them to the same size during the mirroring process.
- ◆ Although you can mirror eight partitions together, mirroring two partitions is usually sufficient fault tolerance for most systems.
- ◆ If a hard disk fails and cannot be accessed by the server, you can unmirror the hard disks and salvage the volume from the functional disk. See “Recovering Data from an Unmirrored (Out of Sync) Hard Disk” on page 591.
- ◆ If you want to remove a hot-plug mirrored disk without bringing down the server, you must unmirror the disk first. See “Unmirroring Hard Disks” on page 567.



When mirroring two disks, one of the disks should be blank. Delete all partitions from the blank disk and then create a new NetWare partition on it. Do not create any volumes. When you mirror the disks, the disk containing data will be mirrored to the blank disk.


To mirror two disks to each other, follow these steps.

Prerequisites



- One blank disk containing one new NetWare partition with no volumes
- A disk containing data you want to mirror

Procedure



1. At the server console prompt, type

```
LOAD INSTALL <Enter>
```

2. Choose “Disk Options” from the “Installation Options” menu.

3. Choose “Mirror/Unmirror Disk Partitions” from the “Available Disk Options” menu.

A “Disk Partition Mirroring Status” window appears. This window lists each device and its mirrored status. For example:

```
Mirrored: Device 1  
Not Mirrored: Device 3  
Out of Sync: Device 4
```

The meaning of each status is as follows:

Status	Explanation
Mirrored	The partition is mirrored to another partition.
Not Mirrored	The partition is not mirrored to another partition.
Out of Sync	The partition was mirrored to another partition, but is currently unmirrored.

4. From the “Disk Partition Mirroring Status” list, choose one of the two devices (disk partitions) you want to mirror to each other.

See “Device numbering” in *Concepts* for an explanation of how logical partitions relate to the installed hard disks.

A “Mirrored Disk Partitions” list appears. This window displays any hard disks that are currently mirrored to the partition you selected.

You must now specify another partition to be mirrored to the selected partition.

5. Press <Insert>.

The system displays the “Available Disk Partitions” list—a list of all partitions that can be mirrored to the first partition.

6. Select the other partition from the “Available Disk Partitions” list.

The system mirrors the two partitions to each other so that the blank partition contains the data from the other partition.



If you are mirroring hard disks, the status on the “Mirrored Disk Partitions” list is now “Out of Sync” for one of the disks. As soon as mirroring is complete, the status of both disks changes to “In Sync.”


Unmirroring Hard Disks

You must unmirror mirrored hard disks before you can delete a partition or conduct surface tests on a disk.



If you plan to replace a hot-plug mirrored disk while the system is running, you must first unmirror the disk. Failure to unmirror the disk could cause the loss of all data on its mirrored partner when you install a replacement disk.

Procedure



1. At the server console prompt, type

LOAD INSTALL <Enter>

2. Choose “Disk Options” from the “Installation Options” menu.

3. Choose “Mirror/Unmirror Disk Partitions” from the “Available Disk Options” menu.

The system displays the “Disk Partition Mirroring Status” list.

4. Select a mirrored partition from the list.

After you select a partition, the system displays the list of disk partitions in the mirrored set.

5. Select the hard disk partition you want to unmirror from the first disk and press <Delete>.

The system displays a message warning that the partition contains volume information.

6. Press <Enter> to clear the message.

7. Choose “Yes” to salvage the information on the partition or “No” if you do not want to salvage the data.

- ◆ If you choose “Yes,” the system warns that the selected partition contains a volume with a name and segment number matching an existing volume. Continue with Steps 7a and 7b.

7a. Press <Enter> to clear the message.

The system prompts you to enter a new name for the volume.

7b. Enter the new volume name.

The “Mirrored Disk Partitions” list is redisplayed.

- ◆ If you choose “No,” the data is destroyed and the “Mirrored Disk Partitions” list is redisplayed.

8. Press <Esc> from the “Mirrored Disk Partitions” list.

The “Partition Mirroring Status” list is displayed. The “Out of Sync” status appears next to the disk you unmirrored.

You can now

- ◆ Finish salvaging the information on the volume (if you chose to salvage information in Step 7) by pressing <F3> and following the prompts.
- ◆ Repair or replace the hard disk.

- ◆ Format the hard disk.
- ◆ Delete the partition on the hard disk.
- ◆ Allocate the unmirrored partition as a new volume segment.

You must first remove any existing volume information from the partition, if necessary. See “Recovering Data from an Unmirrored (Out of Sync) Hard Disk” on page 591 for information about removing existing volume information.

Then see “Creating Volumes” on page 515 for instructions on creating the new volume segment.

Formatting a Hard Disk

You don’t need to do a low-level format on the hard disk for a NetWare partition. A low-level format is done by the manufacturer.



Perform a format on the hard disk only if the disk vendor recommends it! Use tools supplied by the vendor. If you perform a low-level format of a modern SCSI or IDE hard drive you may void the drive’s warranty.

If the manufacturer recommends it, consider formatting a hard disk in the following instances:

- ◆ You have a hard disk that is unreliable and you hope to salvage the disk. Formatting destroys the data on the disk and may not solve the problem.
- ◆ You have changed the controller type on a hard disk, and the new controller uses a different format for data storage.

Disaster Prevention and Recovery

The NetWare 4™ operating system includes security and fault tolerance features to protect network data from intruders, power problems, and system failures. This section describes ways to prevent network data loss, reduce system vulnerability, and recover from failures.

Securing the Server Console

The server console is most secure when locked in a place where no one can reboot it. You can gain an additional level of security by using the SECURE CONSOLE utility.


SECURE CONSOLE provides the following security features, while still allowing authorized operators to use the console:

- ◆ Prevents NetWare Loadable Module programs from being loaded from any directory other than SYS:SYSTEM. This means no one can load an invasive NLM from a server's diskette drive or DOS partition.
- ◆ Prevents keyboard entry into the operating system debugger. This restricts the ability to access secure data directly.
- ◆ Prevents anyone from changing the date and time. Some security and accounting features depend on date and time for their enforcement.

SECURE CONSOLE also removes DOS from the server.

SECURE CONSOLE does not lock the server console; use the "Lock File Server Console" menu option in MONITOR to prevent keyboard entry at the console.


Procedure



1. At the server console prompt, type

SECURE CONSOLE <Enter>

2. (Optional) To secure the console whenever the server is booted, add the **SECURE CONSOLE** command to the server's **AUTOEXEC.NCF** file.



To remove **SECURE CONSOLE**, you must down the NetWare server and reboot it. If the **SECURE CONSOLE** command is in the **AUTOEXEC.NCF** file, you must first remove it.

Preventing Virus Infection

Keep viruses off the network by educating users about virus dangers and by enforcing procedures that reduce virus risks, such as the following:

- ◆ Back up data frequently.
- ◆ Maintain layers of archived backups so you can retrieve a backup from a preinfected file.
- ◆ Keep a write-protected, bootable diskette with the latest virus scan and removal software for all servers and workstations.
- ◆ Keep a backup of executable files and flag them Execute Only.
- ◆ Educate yourself about the infection techniques of the latest viruses.
- ◆ Educate network users about how to detect viruses.
- ◆ Warn users of the dangers of viruses. Discourage them from using diskettes and files that have been in computers away from work.
- ◆ Teach users to power down their workstations immediately upon encountering a virus.
- ◆ Restrict access to a server's diskette drives by locking the server in a secure room. Put tape over the drive openings to remind you not to use them unnecessarily.
- ◆ Avoid using the Supervisor account when possible. The fewer privileges your login account has, the less power a virus has to destroy data and to spread.

Preventing Packet Forgery

NetWare 4 includes a security feature called NCP Packet Signature that protects servers and clients using the NetWare Core Protocol™ (NCP) services.

NCP™ Packet Signature prevents packet forgery by requiring the server and the client to “sign” each NCP packet. The packet signature changes with every packet.

NCP packets with incorrect signatures are discarded without breaking the client’s connection with the server. However, an alert message about the invalid packet is sent to the error log, the affected client, and the server console. The alert message contains the login name and the station address of the affected client.

NCP Packet Signature Options

Because the packet signature process consumes CPU resources and slows performance both for the client and the NetWare server, NCP Packet Signature is optional.

Several signature options are available, ranging from never signing NCP packets to always signing NCP packets. NetWare servers and NetWare clients each have four settable signature levels.

The signature options for servers and clients combine to determine the level of NCP packet signature on the network.

You can choose the packet signature level that best meets both your system performance needs and network security requirements.



Some combinations of server and client packet signature levels may slow performance. However, low-CPU-demand systems may not show any performance degradation.

When to Use NCP Packet Signature

NCP Packet Signature is not necessary for every installation. You may choose not to use NCP packet signature if you can tolerate security risks in situations such as

- ◆ When only executable programs reside on the server

- ◆ You know and trust all network users
- ◆ Data on the NetWare server is not sensitive and loss or corruption of this data would not affect operations

NCP Packet Signature is recommended for security risks such as

- ◆ An untrustworthy user at a workstation on the network
- ◆ Easy physical access to the network cabling system
- ◆ An unattended, publicly accessible workstation

Server Signature Levels

Before changing the NCP Packet Signature level, read this entire section to understand all of the options.

To determine the server's current signature level, type

SET NCP Packet Signature Option <Enter>

To set a server's packet signature level, type

SET NCP Packet Signature Option = *number*

Replace *number* with 0, 1, 2, or 3. The default is 1.

Number	Explanation
0	Server does not sign packets (regardless of the client level).
1	Server signs packets <i>only</i> if the client requests it (client level is 2 or higher).
2	Server signs packets if the client is capable of signing (client level is 1 or higher).
3	Server signs packets and requires all clients to sign packets or logging in will fail.

You can use the SET console command to change the signature level from a lower to a higher level. You cannot change from a higher to a lower level unless you first reboot the server.

For example, if the current signature level is 2, you can't set the signature level to 1 by using the SET command at the console.

To change the signature level from 2 to 1, you must first reboot the server and then set the level to 1 by typing:

```
SET NCP Packet Signature Option = 1
```

You can add this SET command to your AUTOEXEC.NCF file to set the signature level each time the server is brought up.

Client Signature Levels

To set DOS or MS Windows client signature levels, add the parameter to the workstation NET.CFG file. The format is as follows:

```
signature level = number
```

To set OS/2 client signature levels, add the parameter to the NET.CFG file under the NetWare Requester option. The format is as follows:

```
signature level number
```

Replace *number* with 0, 1, 2, or 3. The default is 1.

Number	Explanation
0	Client does not sign packets.
1	Client signs packets <i>only</i> if the server requests it (server level is 2 or higher).
2	Client signs packets if the server is capable of signing (server level is 1 or higher).
3	Client signs packets and requires the server to sign packets or logging in will fail.

Effective Packet Signature

The NCP Packet Signature levels for the server and the client interact to create the effective packet signature for the network. Some combinations of server and client levels do not allow logging in.

Figure 7-1 shows the interactive relationship between the server packet signature levels and the client signature levels.

Figure 7-1
Effective Packet
Signature of Server
and Client

lf	Server = 0	Server = 1	Server = 2	Server = 3
Client = 0	○	○	○	∅
Client = 1	○	○	●	●
Client = 2	○	●	●	●
Client = 3	∅	●	●	●

- Packet signature
- No packet signature
- ∅ No logging in

Examples of Signature Levels in Different Situations

The default NCP Packet Signature level is 1 for clients and 1 for servers. In general, this setting provides the most flexibility while still offering protection from forged packets. Following are some examples of situations requiring different signature levels.

All Information on the Server Is Sensitive

If an intruder gains access to *any* information on the NetWare server, it could damage the company.

Recommendation: Set the server to level 3 and all clients to level 3 for maximum protection.

Sensitive and Nonsensitive Information Reside on the Same Server

The NetWare server has a directory for executable programs and a separate directory for corporate finances (such as Accounts Receivable).

Recommendation: Set the server to level 2 and the clients that need access to Accounts Receivable to level 3. All other clients remain at the default, level 1.

Users Often Change Locations and Workstations

You are uncertain which employees will be using which workstations, and the NetWare server contains some sensitive data.

Recommendation: Set the server to level 3. Clients remain at the default, level 1.

A Workstation Is Publicly Accessible

An unattended workstation is set up for public access to nonsensitive information, but another server on the network contains sensitive information.

Recommendation: Set the sensitive server to level 3 and the unattended client to level 0.

Changing the Signature Level for an NLM using CLIB

NLM programs that use CLIB are assigned a default NCP Packet Signature level that corresponds to the current signature level of the server.

To change the packet signature level for *all* NLM programs using CLIB, use the following command format when you load CLIB:

```
LOAD CLIB /Lnumber
```

Replace *number* with 0, 1, 2, or 3.



To make sure CLIB uses the correct signature level when it is automatically loaded by other NLM programs, put the command in the AUTOEXEC.NCF file.

To change the packet signature level for a single NLM, use the following command format when you load the NLM:

```
LOAD NLM [CLIB_OPT] /Lnumber
```

Replace *number* with 0, 1, 2, or 3.

For information about Packet Signature Levels, see “Server Signature Levels” on page 573.

Packet Signature Considerations for Job Servers

A job server is a server that performs a task and then returns the completed task. Most job servers are third-party products.

You should be aware that some job servers do not support NCP Packet Signature. A job server may produce unsigned sessions if

- ◆ It does not operate on top of DOS
- ◆ It does not use standard NetWare shells
- ◆ It is not an NLM
- ◆ It uses its own implementation of the NCP engine (such as embedded print servers in printers)

Minimizing Risks

To minimize security risks associated with job servers:

- ◆ Install queues only on servers with signature level 3.
- ◆ Do not allow privileged users to put jobs in queues on servers with signature levels below 3.
- ◆ Make sure the job server's account is unprivileged.
- ◆ Disable the job server's ability to change to client rights.

Disabling Change to Client Rights

To prevent a job server from assuming the rights of a client, add the following SET command to the server's STARTUP.NCF file:

SET Allow Change to Client Rights = OFF

The default is ON, because certain job servers and third-party applications cannot function without changing to client rights. Refer to the documentation that comes with the job server to determine whether the job server can function without client rights.

Additional Information

For more information about	Refer to
NCP Packet Signature	"NCP Packet Signature" in <i>Concepts</i>
SET parameters	"SET" and "SERVMAN" in <i>Utilities Reference</i>

Activating UPS Monitoring

An uninterruptible power supply (UPS) is an indispensable part of your network. Not only does it prevent damage to your computers by eliminating power surges and brownouts, but it also prevents data loss during power outages.

Use the following procedure to load UPS.NLM, which allows you to configure the UPS monitoring feature in NetWare 4. You can specify the time the server functions on battery power and how long the battery needs to charge.



Note

UPS.NLM supports the following ports: mouse port, Novell® keycard, disk coprocessor board (DCB), and standalone UPS cards. If you are using a UPS machine connected to the server by serial port, you need NetWare-compatible third-party UPS software.

Prerequisites



Checklist

- A UPS hardware driver must be loaded on the server before UPS is activated.
- The DCB.DSK device driver must be loaded if you are using a DCB to connect the UPS machine to the server.
- The UPS hardware must be installed according to the vendor's instruction manual.

Procedure



Procedure

1. **At the server console prompt, load the UPS NLM by typing**

```
LOAD UPS <Enter>
```

2. **Set the type value by entering**

```
TYPE = typename <Enter>
```

Replace *typename* with the appropriate adapter card that you have connected to the UPS machine. Valid type names include DCB, mouse (IBM PS/2*-style), standalone, and keycard.

For example, to load UPS.NLM with the DCB device, type

```
TYPE = DCB <Enter>
```

3. When prompted, enter a valid I/O port number.



Valid I/O port numbers vary depending on your hardware. A list of valid port numbers is displayed. Refer to “UPS” in *Utilities Reference* for valid port numbers.

4. Select a discharge time, a recharge time, and a wait time.

For example, type

```
DISCHARGE=20 RECHARGE=60 WAIT=15 <Enter>
```

See “UPS” in *Utilities Reference* for valid values for these fields.

5. (Optional) To automatically load UPS.NLM when the server boots up, add the LOAD UPS command to your AUTOEXEC.NCF file.

You can edit the AUTOEXEC.NCF file by using INSTALL or EDIT.



You can set all required parameters when you load UPS.NLM. For example, add the following statement to your AUTOEXEC.NCF file:

```
LOAD UPS TYPE=DCB PORT=346 DISCHARGE=20 RECHARGE=120  
WAIT=10
```

6. (Optional) To view the status of your UPS after loading UPS.NLM, type

```
UPS STATUS <Enter>
```

7. (Optional) To change the discharge or recharge settings for UPS.NLM, use the UPS TIME command at the server console.

For example, you might use a command similar to the following:

```
UPS TIME DISCHARGE=10 RECHARGE=90 <Enter>
```

Remember to change the setting in the AUTOEXEC.NCF file if you want the server to retain the new setting after rebooting.

Additional Information

For more information about	Refer to
Using UPS.NLM	"UPS" in <i>Utilities Reference</i>
Using UPS STATUS	"UPS STATUS" in <i>Utilities Reference</i>
Using UPS TIME	"UPS TIME" in <i>Utilities Reference</i>

Protecting the Operating System's Memory

NetWare 4 has memory protection features that guard the operating system memory from corruption by NetWare Loadable Module (NLM) programs. These features allow you to run such programs in a separate memory domain called the OS_PROTECTED domain.

Once you have loaded an NLM in the OS_PROTECTED domain and found it is safe, you can load it into the OS domain, where it can run more efficiently.

If you have third-party NLM programs such as a database or backup NLM, check the documentation that came with the NLM to determine the recommended memory domain for the NLM.

Using Domains

Use the following procedure to enable the operating system memory protection and load an NLM in the OS_PROTECTED domain.


This procedure requires you to reboot the NetWare server. You may want to do it after business hours or when no one is accessing the server.



DOMAIN.NLM must be located in your boot directory and must be loaded from your STARTUP.NCF file before you load any other NLM programs or LAN drivers. DOMAIN will not load if any NLM files or LAN drivers have previously been loaded.

To load DOMAIN from the STARTUP.NCF file, you must have the following files in your bootup directory: DOMAIN.NLM, NWTIL.NLM, NWTILR.NLM, and DOMAIN.MSG. (The server installation process places these files in the bootup directory.)

Procedure



1. Edit the **STARTUP.NCF** file to put the following command in the first line of the file:

```
LOAD DOMAIN
```

2. Reboot the server.



In NetWare 4.1, when you load **DOMAIN.NLM**, the following files are automatically loaded: **NWTIL.NLM** and **NWTILR.NLM**. Once all your NLM programs have been loaded, you can unload **NWTIL.NLM** to reclaim memory. **NWTIL.NLM** is only needed by **DOMAIN.NLM** when NLM programs are loaded into the **OS_PROTECTED** domain.

3. Change to the **OS_PROTECTED** memory domain by typing

```
DOMAIN=OS_PROTECTED <Enter>
```



You can place this command in the **STARTUP.NCF** or **AUTOEXEC.NCF** file, or you can type it at the server console.

This command, and the **DOMAIN** commands that follow, are available only after **DOMAIN.NLM** has been loaded.

4. (Optional) Load an NLM in the **OS_PROTECTED** domain.

Following is a sample load command:

```
LOAD EXAMPLE.NLM <Enter>
```

The NLM is now loaded in the **OS_PROTECTED** domain.

5. (Optional) To change back to the OS domain, type

DOMAIN=OS <Enter>

6. (Optional) To view memory domains and the NLM programs loaded into them, type

DOMAIN <Enter>

Additional Information

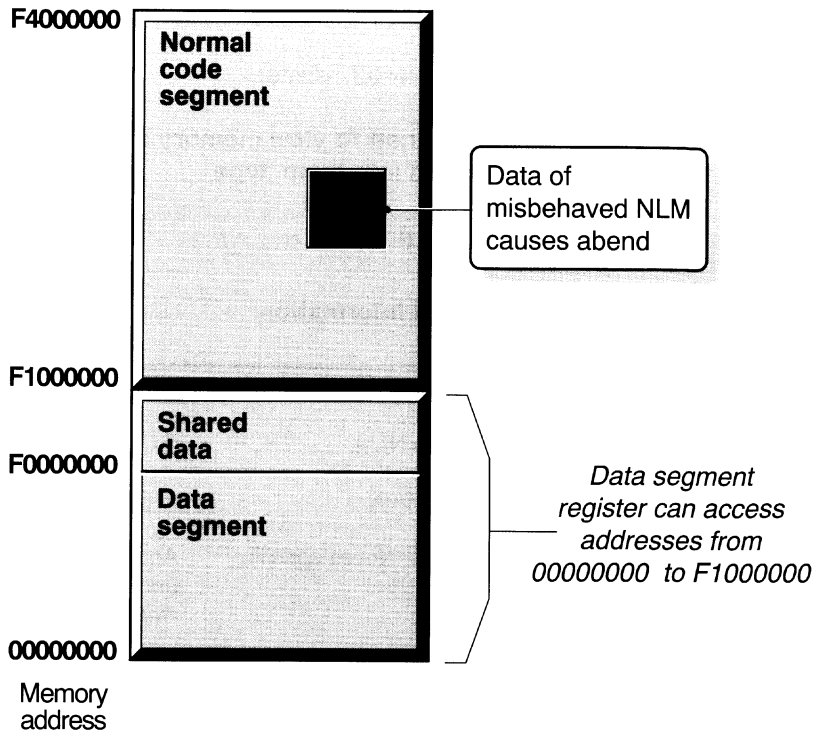
For more information about	Refer to
Loading an NLM	"LOAD" in <i>Utilities Reference</i>
Using DOMAIN	"DOMAIN" in <i>Utilities Reference</i>
Domains for Novell-specific NLMs	Appendix C, "Memory Domains for NLMs and Autoloaded NLMs," of <i>Utilities Reference</i>
Domains	"Domains" in <i>Concepts</i>
Memory protection	"Memory protection" in <i>Concepts</i> "Paging" in <i>Concepts</i>

Loading a Misbehaved NLM

This section is for the advanced network supervisor. Consider using one of the following procedures only if you have a general understanding of segmented memory.

If you load an NLM, and you receive an abend that indicates a segmentation violation, the NLM may be accessing data that exists in the code segment in memory. Figure 7-2 illustrates such an NLM.

Figure 7-2
Logical Memory
Map



Note

Please inform the NLM vendor if you have an NLM that causes a domain to be quarantined or causes an abend.

If you want to load the NLM despite this problem, you can use either of the two procedures that follow.

Debugging the NLM is easier if you use Option 1 to load the NLM; however, this procedure affects memory management for all NLM programs that you load in the OS_PROTECTED domain.

If you do not debug the NLM, or if you have already debugged it, load the NLM using Option 2. This procedure affects memory management only for the errant NLM.

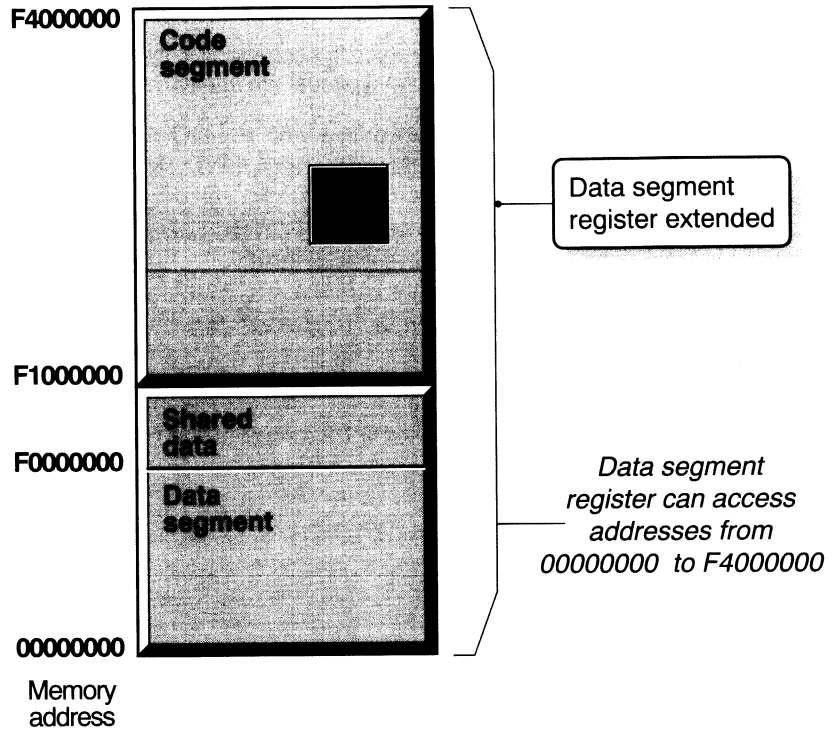
Option 1: Loading a Suspect NLM

You can load the NLM in the OS_PROTECTED domain by using the data segment register to access the code segment. Figure 7-3 illustrates this process.




If you use this option, all NLM programs that are subsequently loaded in the OS_PROTECTED domain will use this method of memory management by allowing access to the code segment using the data segment register.

Figure 7-3
Logical Memory
Map with the
DOMAIN -E
Parameter



Procedure



1. Create the OS_PROTECTED domain by placing the following command in the STARTUP.NCF file:

```
LOAD DOMAIN -E
```

2. Reboot the server.
3. Change to the OS_PROTECTED memory domain by typing

```
DOMAIN=OS_PROTECTED <Enter>
```



You can place this command in the STARTUP.NCF or AUTOEXEC.NCF file, or you can type it at the server console.

This command, and the DOMAIN commands that follow, are available only after DOMAIN.NLM has been loaded.

4. From the OS_PROTECTED domain, load the NLM, using the following command:

```
LOAD nlm_name <Enter>
```

where

nlm_name is the name of the NLM you want to load.

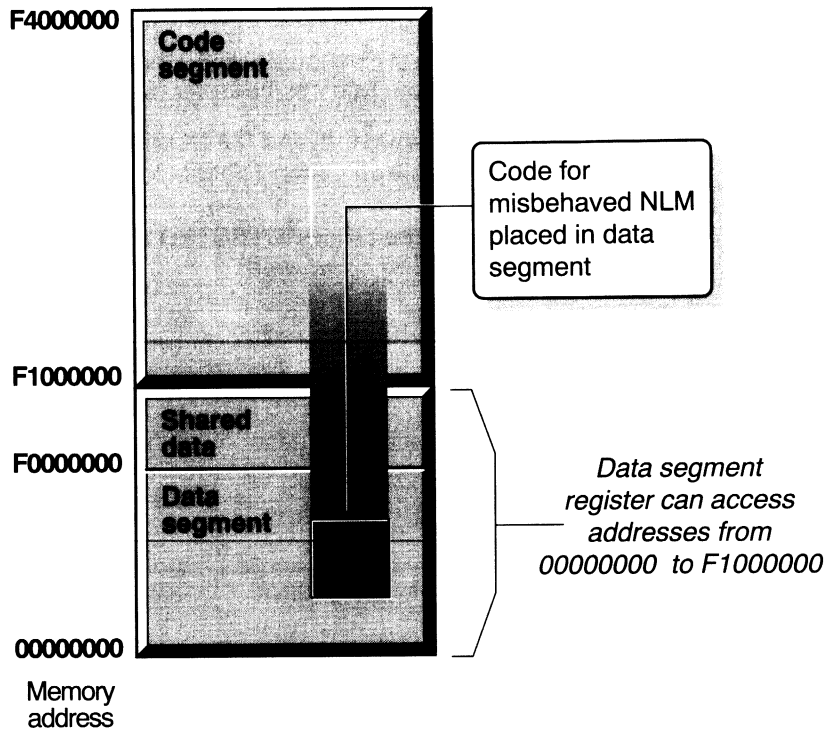
Option 2: Loading a Suspect NLM

In this procedure, the NetWare operating system loads the NLM code in the data segment. Figure 7-4 illustrates this process.




If you use this option, you will affect memory management only for the NLM that you load with the `-L` parameter.

Figure 7-4
Logical Memory
Map with the LOAD
`-L` Parameter



Procedure



1. Create the OS_PROTECTED domain by placing the following command in the STARTUP.NCF file:

```
LOAD DOMAIN
```

2. Change to the OS_PROTECTED memory domain by typing

```
DOMAIN=OS_PROTECTED <Enter>
```



You can place this command in the STARTUP.NCF or AUTOEXEC.NCF file, or you can type it at the server console.

This command, and the DOMAIN commands that follow, are available only after DOMAIN.NLM has been loaded.

3. From the OS_PROTECTED domain, load the NLM, using the following command:

```
LOAD -L nlm_name
```

where

nlm_name is the name of the NLM you want to load.

Additional Information

For more information about	Refer to
Using DOMAIN	"DOMAIN" in <i>Utilities Reference</i>
Domains for Novell-specific NLMs	Appendix C, "Memory Domains for NLMs and Autoloaded NLMs," of <i>Utilities Reference</i>
Domains	"Domains" in <i>Concepts</i>
Memory protection	"Memory protection" in <i>Concepts</i> "Paging" in <i>Concepts</i>

Protecting Database Integrity with TTS

The Transaction Tracking System (TTS) can prevent data corruption by backing out of incomplete transactions and keeping a record of backed-out data.

By default, TTS™ is enabled.

The NetWare server automatically disables TTS if one of the following happens:

- ◆ The SYS: volume becomes full. (The SYS: volume is the TTS backout volume.)
- ◆ The NetWare server has insufficient memory to operate TTS.

If TTS has been disabled and you have solved the problems that led to its disabling, use the `ENABLE TTS` command to enable TTS again. Type the following:

```
ENABLE TTS <Enter>
```

Automatically Backing Out of Incomplete Transactions


Use this procedure to enable the server to automatically back out of any incomplete transactions, without being prompted.

Prerequisites



- TTS enabled

Procedure



1. **At the server console prompt, type**

```
LOAD SERVMAN <Enter>
```

2. **Choose “Server Parameters” from the “Available Options” menu.**

3. **Choose “Transaction Tracking” from the “Categories” menu.**

4. Choose “Auto TTS Backout Flag” from the “Transaction Tracking” menu.
5. Press <Enter> to change the value to ON.
6. Press <Esc> twice to reach the “Update Options” menu.
7. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the STARTUP.NCF file.

8. Press <Enter> to update the file.
9. When you want the changes to take effect, reboot the server.

Keeping a Log of Backed-Out Data


Use this procedure to keep a log (TTS\$LOG.ERR on volume SYS:) of all data that is backed out by TTS.

Prerequisites



- TTS enabled

Procedure



1. At the server console prompt, type
`LOAD SERVMAN` <Enter>
2. Choose “Server Parameters” from the “Available Options” menu.
3. Choose “Transaction Tracking” from the “Categories” menu.
4. Choose “TTS Abort Dump Flag” from the “Transaction Tracking” menu.
5. Press <Enter> to change the value to ON.
6. Press <Esc> twice to reach the “Update Options” menu.

7. Choose “Update AUTOEXEC.NCF and STARTUP.NCF Now” from the “Update Options” menu.

A window appears, indicating the path to the AUTOEXEC.NCF file.

8. If desired, press <Enter> to update the file.

The system writes the parameters to the AUTOEXEC.NCF file or updates the parameters if they are already in the file.

If you do not update the AUTOEXEC.NCF file, the parameter changes last only until the server is rebooted.

Mirroring and Duplexing to Protect Data

NetWare 4 enables you to protect data from hard disk failure by duplicating, or “mirroring,” one hard disk’s data on one or more other hard disks.

If a mirrored hard disk fails and cannot be accessed by the server, you can unmirror the hard disks and salvage the volume from the functional disk.


How to mirror or duplex hard disks is explained under “Mirroring and Duplexing a Hard Disk” on page 565. How to recover data from a mirrored disk is explained in the next section, “Recovering Data from an Unmirrored (Out of Sync) Hard Disk.”

Recovering Data from an Unmirrored (Out of Sync) Hard Disk

Once a hard disk is unmirrored, its status is listed as either “Not Mirrored” or “Out of Sync” on the “Disk Partition Mirroring Status” list. To display this list, load INSTALL, select “Disk Options,” and then choose “Mirror/Unmirror Disk Partitions.”

When a hard disk is listed as “Out of Sync,” the operating system does not recognize any volume information on it. Use this procedure to recover data from an “Out of Sync” partition.

Procedure



1. At the server console prompt, type

```
LOAD INSTALL <Enter>
```

2. Select “Disk Options” in the “Installation Options” menu.
3. Select “Mirror/Unmirror Disk Partitions” in the “Available Disk Options” menu.
4. Select the “Out of Sync” partition and press <F3>.

The following warning may appear if another volume segment with the same name exists on another disk:

```
Warning! The selected partition contains "volume  
<name> segment <number>" and that volume is  
already defined.
```

```
<Press ESCAPE To Continue>
```

- 4a. Press <Esc> and respond “Yes” to the “Rename the Volume Segment?” prompt.
- 4b. Enter a unique name for the new volume segment.
5. Respond “Yes” to the prompt to salvage the volume segment.

Now the segment can be mounted as an independent volume.

Managing Network Time Synchronization

Time synchronization ensures that all servers in a Directory tree report the same time and order Directory Services events correctly.

NetWare 4.1 uses the TIMESYNC NLM, which automatically loads when the server is booted, to control time synchronization. The TIMESYNC.CFG file in SYS:SYSTEM can contain the parameters for the TIMESYNC NLM, if the default settings are not adequate.

To change the time synchronization configuration on a NetWare server, use the EDIT server utility to modify the parameters in the TIMESYNC.CFG file and then reboot the server.



Do not put SET time synchronization parameters in the STARTUP.NCF file, because the TIMESYNC.CFG file is read *after* the STARTUP.NCF file.

SET Parameters for TIMESYNC

Table 7-5 shows the SET parameters used to configure time synchronization. The defaults are appropriate for most small networks. For large network TIMESYNC considerations, see “Configuring Time Synchronization for Large Networks” on page 601.

Table 7-5
TIMESYNC SET parameters

SET parameter and default	Recommended usage
TIMESYNC ADD Time Source	Use EDIT rather than this parameter to add a time server to the time source list in the TIMESYNC.CFG file.
TIMESYNC Configuration File Default = SYS:SYSTEM\	Use this parameter to specify the file path if the TIMESYNC.CFG file is <i>not</i> located in SYS:SYSTEM.
TIMESYNC Configured Sources Default = OFF	Turn this parameter ON if you are using a custom-configured list of time sources.
TIMESYNC Directory Tree Mode Default = ON	Turning this parameter OFF allows SAP packets from other Directory trees to influence time synchronization on this server's Directory tree.
TIMESYNC Hardware Clock Default = ON	Turn this parameter OFF <i>only</i> if this server uses an external time source (such as a radio clock). All servers in the same tree should have the same setting for this parameter.
TIMESYNC Polling Count Default = 3 (time packets)	Do not increase this setting. Doing so adds unnecessary traffic to the network.
TIMESYNC Polling Interval Default = 600 (seconds)	Decrease this setting if you need to poll time servers more often than every 10 minutes. All servers in the same tree should have the same setting for this parameter.
TIMESYNC REMOVE Time Source	Use EDIT rather than this parameter to remove a time server from the time source list in the TIMESYNC.CFG file.
TIMESYNC RESET Default = OFF	Use EDIT rather than this parameter to reset values in the TIMESYNC.CFG file and to remove the list of time sources.
TIMESYNC Restart Flag Default = OFF	Turn this parameter ON <i>only</i> if you want to reload the TIMESYNC NLM (because you edited TIMESYNC.CFG) without rebooting the server.

Table 7-5 *continued*

TIMESYNC SET parameters *continued*

SET parameter and default	Recommended usage
TIMESYNC Service Advertising Default = ON	Turn this parameter OFF if you are using a custom-configured list of time sources.
TIMESYNC Synchronization Radius Default = 2000 (milliseconds)	Increase this parameter to allow a wider margin of error for time synchronization between servers. If you decrease the radius too much (below 1,000), some time servers may never achieve synchronization.
TIMESYNC Time Adjustment	Use this parameter sparingly to correct network-wide time errors, or you may corrupt time synchronization and the order of events on your network.
TIMESYNC Time Source	Use EDIT rather than this parameter to add a time server or to display the time source list in the TIMESYNC.CFG file.
TIMESYNC Type Default = Secondary	Use EDIT rather than this parameter to change the default time server type in the TIMESYNC.CFG file.
TIMESYNC Write Parameters Default = OFF	Use EDIT rather than this parameter to change the settings in the TIMESYNC.CFG file.
TIMESYNC Write Value	Use EDIT rather than this parameter to modify the TIMESYNC.CFG file.

Editing the TIMESYNC.CFG File

Time synchronization parameters are stored in their own configuration file, by default named TIMESYNC.CFG in the SYS:SYSTEM directory.

To change the time synchronization configuration on a NetWare server, modify the parameters in the TIMESYNC.CFG file, and then reboot the server or use SET to turn ON the “TIMESYNC Restart Flag” parameter.



Note You can use SET or SERVMAN to change the TIMESYNC parameters; however, unless you also edit the TIMESYNC.CFG file, changes made in SET or SERVMAN will be lost the next time the server boots.

Prerequisites

Checklist



- Access to the server console, or a remote console session with the server.

Procedure

Procedure



1. At the server console prompt, type

LOAD EDIT <Enter>

A prompt to “Enter file to edit or press <Esc> to exit” appears.

2. At the prompt, type

SYS:SYSTEM\TIMESYNC.CFG <Enter>

The current TIMESYNC.CFG file appears.

3. Use the arrow keys to move the cursor to the lines you need to edit.

4. Modify the settings for the TIMESYNC parameters you want to change.

5. To save your edits, press <Esc>.

A “Do you want to save SYS:SYSTEM\TIMESYNC.CFG” prompt appears.

6. To save the TIMESYNC.CFG file, choose “Yes.”

7. To exit EDIT, press <Esc>.

Suggestion



If you have several parameters in your TIMESYNC.CFG file, you may want to copy the file from server to server. Be sure that each copied TIMESYNC.CFG file contains the correct default time server type and configured-sources list for the server it configures.

Creating a Custom Time Source Configuration

Time servers use one of two methods to find each other on the internetwork: SAP (Service Advertising Protocol) and custom time source configuration.

SAP Method

By default, Primary, Reference, and Single Reference time servers use SAP to announce their presence on the network. Secondary time servers do not advertise, but they do “listen to” SAP information to find the nearest time source.

The SAP method allows for quick installation without regard to the network layout. It also allows automatic reconfiguration if new servers are added to the network.

Custom Time Source Method

The custom time source method lists specific, authorized time servers that each server contacts to determine network time. In a custom time configuration, servers do not listen for SAP information from other time servers, and they do not advertise with SAP.

A configured-sources list resides in the TIMESYNC.CFG file of each time server on the internetwork. The order of the server names in the list is the polling order the server follows to get the correct time.

The custom configuration method requires careful planning before NetWare installation, but it gives the network supervisor complete control of the time synchronization environment.

In most circumstances, you can use the default SAP method and do not need to create a custom time source configuration. However, you may need a custom configuration if

- ◆ You have more than one Reference or more than one Primary time server on your network.
- ◆ You frequently have temporary or “test” servers added to or removed from your network.
- ◆ You need to prevent unauthorized time sources from determining network time.

Use this procedure to create a custom time source configuration. These steps must be repeated for each time server on your network.

Prerequisites



- Access to the server console, or a remote console session with the server.
- Prioritized list of time sources.

Procedure



1. At the server console prompt, type

LOAD EDIT <Enter>

A prompt to “Enter file to edit or press <Esc> to exit” appears.

2. At the prompt, type

SYS:SYSTEM\TIMESYNC.CFG <Enter>

The current TIMESYNC.CFG file appears.

3. Add the following parameters to the TIMESYNC.CFG file:

Type = x
Service Advertising = OFF
Configured Sources = ON

Replace *x* with this server's time server type (Reference, Single Reference, Primary, or Secondary).

4. Add the list of time sources for this server as follows:

Time Source = x
Time Source = y
Time Source = z

Replace *x* with the first time server this server should contact (probably a Reference server or the nearest Primary time server). Replace *y* and *z* with the other time servers this server should contact, in order of importance.



The order of the time servers in the time source list determines their priority when advertising network time.

5. To save your edits, press <Esc>.

A "Do you want to save SYS:SYSTEM\TIMESYNC.CFG" prompt appears.

6. To save the TIMESYNC.CFG file, choose "Yes."

7. To exit EDIT, press <Esc>.



If you have several parameters in your TIMESYNC.CFG file, you may want to copy the file from server to server. Be sure that each copied TIMESYNC.CFG file contains the correct default time server type and time source list for the server it configures.

For more information about	Refer to
SAP	"Service Advertising Protocol" in <i>Concepts</i>
Time server types	"Understanding Time Synchronization in NDS" in Chapter 4 of <i>Introduction to NetWare Directory Services</i>
Time synchronization	"Time synchronization" in <i>Concepts</i>

Adjusting a Fast or Slow Reference Time Server

If your Reference time server is several minutes off, all other time servers on your network will be inaccurate also.

To correct the time on the Reference Server and, therefore, network-wide, type the following SET command at the Reference Server's console:

```
SET TIMESYNC Time Adjustment=[+|-] hour:minute:second  
at month/day/year hour:minute:second [AM|PM] <Enter>
```

For example, if your Reference Server is 10 minutes fast and you want to adjust the time at 2 a.m. on May 1, 1995, you would type

```
SET TIMESYNC Time Adjustment=-00:10:00 at 5/1/95  
2:00:00 AM <Enter>
```

Important



Adjusting network time affects the order of Directory Services events. It is best to schedule time adjustments during low-usage periods, such as late evening or early morning.

To cancel a time adjustment, type the following SET command at the Reference Server's console:

```
SET TIMESYNC Time Adjustment = cancel <Enter>
```

Important



The cancel command must be issued *before* the scheduled time adjustment and at the *same* server where the original time adjustment was set, or it will not cancel the time adjustment.

Configuring Time Synchronization for Large Networks

Some of the default settings for TIMESYNC are not suitable for large networks. This section explains time synchronization considerations for configuring large installations.

Time Providers and Time Consumers

Time synchronization servers are either time providers or time consumers. Secondary servers are time consumers. Single Reference, Reference, and Primary servers are time providers.

Reference and Primary servers must be able to exchange information with another time provider in order to synchronize. A Reference or Primary server cannot claim to be synchronized if it is the only time provider on the network, because it requires at least one other time provider to synchronize with.

The Single Reference server, as its name implies, is a special type of Reference server that does not need another time provider in order to synchronize.

Secondary servers (time consumers) require at least one time provider.

Time Synchronization and the Directory Tree

Although time synchronization and NetWare Directory Services are independent features of NetWare 4.1, time synchronization can be configured to follow the Directory tree structure.

The default NetWare 4.1 installation makes the root server in a Directory tree a Single Reference time server and all other servers in the tree Secondary time servers.

The default time synchronization parameter, "SET TIMESYNC Directory Tree Mode = ON," recognizes the Directory tree structure and only synchronizes time within that tree.

In a large installation with more than one Directory tree, the network supervisor can change the defaults to configure time synchronization across Directory trees.

The default does not require custom time synchronization parameters during installation. However, the default relies on a Single Reference server, thus permitting only a single point of failure and increasing network traffic at that server.

In addition, the default installation relies on Service Advertising Protocol (SAP) packets to identify the time providers.

On a large network, relying on SAP is not only inefficient, it can also lead to errors—especially if it is common for test servers (which may be misconfigured) to come and go on the network.

The custom time source method described under “Custom Time Source Method” on page 597 is recommended for large networks because it does not rely on SAP.

The time synchronization algorithm always attempts to contact explicitly named time sources before listening to SAP. This makes it possible to synchronize across Directory trees, because the Directory Tree Mode only applies to SAP packets.

Network supervisors with custom-configured time sources may choose to eliminate the use of SAP entirely.

Using an .NCF File to Configure Multiple Time Servers

One easy way to configure several servers is to create a custom .NCF file containing the SET parameters required to configure time sources.

Using the custom .NCF file avoids repetitive typing at each server console. Include the “SET TIMESYNC Write Parameters = ON” parameter in the custom .NCF file to save the time source parameters to each server’s TIMESYNC.CFG file.

Most servers will have identical TIMESYNC configuration parameters with the exception of the time source list, where the first entry in the list likely varies from server to server.

To change the first server in the time source list, edit the TIMESYNC.CFG file or issue the “SET TIMESYNC Time Source” parameter from the server console before saving the configuration parameters.

Avoiding SET TIMESYNC Parameter Inconsistencies

Several of the SET TIMESYNC parameters interact with each other. Because these parameters are set on each server, it is important to configure the entire network consistently.

Understanding the relationship of the parameters described here can help avoid time synchronization problems on large networks.

SET TIMESYNC Service Advertising = ON/OFF

When ON, time providers advertise using SAP. This parameter has no effect on time consumers (Secondary time servers). When OFF, time providers do not advertise.

We recommend all servers be configured the same (ON or OFF) for this parameter to be effective.

SET TIMESYNC Configured Sources = ON/OFF

When ON, time servers do not listen to SAP information. They listen only to the time providers in the configured sources list (initially read from the TIMESYNC.CFG file). When OFF, time servers listen for SAP information.

SET TIMESYNC Directory Tree Mode = ON/OFF

When ON, time providers do not advertise under the name of the host server. Instead, they advertise under the name of their Directory Services tree. When ON, servers listening for SAP packets will reject packets that are not from within their own tree.

When OFF, time providers advertise under their host server name, and listening servers will accept packets from any time provider.

This parameter also interacts with the "SET TIMESYNC Service Advertising" and the "SET TIMESYNC Configured Sources" parameters. Directory Tree Mode has no effect if Service Advertising is OFF and Configured Sources is ON.

SET TIMESYNC Time Source and SET TIMESYNC Add Time Source

These parameters identify a time provider by server name. When configured sources are listed in a server's TIMESYNC.CFG file, their time information is always used before any SAP information is considered.

SET Default Time Server Type

This parameter provides a method to override the default TIMESYNC type, which is Secondary, without using a TIMESYNC.CFG file.

However, when present, the "SET TIMESYNC Type" parameter always overrides the "SET Default Time Server Type" parameter.

If you use INSTALL to change the default type, it may have no effect if you have already created a TIMESYNC.CFG file with the overriding "TIMESYNC Type" parameter in it.

Example of Synchronizing Two Directory Trees

Suppose you have two Directory trees on your network and you want all the servers in both trees to synchronize to the same time.

You used the default installation, so the root servers of each tree (ROOT1 and ROOT2) are Single Reference servers, Tree Mode is ON, and Service Advertising is ON.

Each tree is synchronized within itself, but time between the trees differs by a few minutes. To synchronize both trees together, create a custom .NCF file containing the following commands and execute it on the respective servers:

Server ROOT1

```
SET TIMESYNC Type = REFERENCE  
SET TIMESYNC Time Source = ROOT2  
SET TIMESYNC Write Value = 3  
SET TIMESYNC Write Parameters = ON
```

Server ROOT2


```
SET TIMESYNC Type = PRIMARY  
SET TIMESYNC Time Source = ROOT1  
SET TIMESYNC Write Value = 3  
SET TIMESYNC Write Parameters = ON
```

The root servers of both trees will synchronize, which in turn will synchronize all of the other servers in both trees.

Since the roots are now both time providers, it is necessary to change the "TIMESYNC Types" from "Single" to "Reference" on Server ROOT1 and "Primary" on Server ROOT2.

Server ROOT1 is a Reference server so that there will be a central point for controlling time on the network.

The Write Parameters command causes the TIMESYNC.CFG file to be rewritten so the changes will be in place if the servers are rebooted.



When you synchronize two Directory trees, you basically double the size of the time synchronization network. You should examine the physical layout of the network and add other Primary servers, or consider hand-configuring the time sources and eliminating the use of SAP.

For more information about	Refer to
Time server types	"Understanding Time Synchronization in NDS" in Chapter 4 of <i>Introduction to NetWare Directory Services</i>
Time synchronization	"Time synchronization" in <i>Concepts</i>
SAP	"Service Advertising Protocol" in <i>Concepts</i>

Using Remote Console to Manage a Server

The NetWare 4 operating system allows you to use a workstation as a server console when needed. A workstation functioning as a console is called a *remote console*.

A remote console gives you greater server security because you can lock servers in a safe place and remove the keyboards and monitors. You can then start a remote console session from a workstation whenever you need to accomplish console tasks.

When finished, you can exit the remote console session so the computer functions as a workstation again.

You can perform the following functions from a remote console:

- ◆ Use console commands as you would at the server console.
- ◆ Scan directories and edit text files in both NetWare and DOS partitions on a server.
- ◆ Transfer files to, but not from, a server.
- ◆ Bring down or reboot a server.
- ◆ Install or upgrade NetWare.

This section describes how to establish a remote console session and how to accomplish some essential management tasks from a remote console.

Understanding Connection Types

Remote console sessions can be run on the network or through a modem.

- ◆ A network connection is called a *direct connection*.
- ◆ A connection through a modem or null modem cable is called an *asynchronous connection*. A null modem cable is a direct serial connection between two computers, as opposed to a network connection.

Running a Remote Console Session Through a Direct Connection

To run a remote console session over a direct connection you must do the following:

1. Load the REMOTE NLM and RSPX NLM on the server.
2. Run RCONSOLE.EXE from the workstation to access the server.

Each of these tasks is described in the following procedures

Loading the REMOTE NLM and RSPX NLM on the server

Prerequisites



- The REMOTE NLM, the remote password, and the RSPX NLM. You don't need the Supervisor object right to the remote server because RCONSOLE does not use NetWare Directory Services.
- A server running NetWare 3.11 or a later version
- 75 KB of available memory

Procedure



1. At the server console prompt, type

LOAD REMOTE [*password*] <Enter>

If you do not specify a password as a parameter, you are prompted to enter one. A password is required. The password you establish here must be entered again when the remote console session is established.

2. At the server console prompt, type

LOAD RSPX <Enter>

You must load the RSPX NLM before you can run a remote console session through a direct connection.

When you load RSPX, you have the option to require packet signatures to ensure security. The default is ON, which means that packet signatures are required.

However, packets with signatures are not compatible with NetWare 3.11. If you run NetWare 3.11, set packet signatures OFF when you load RSPX. To set packet signatures OFF, type

LOAD RSPX SIGNATURES OFF <Enter>

For more information about packet signatures, see "Preventing Packet Forgery" on page 572.

Running RCONSOLE.EXE from a Network Workstation

Once the REMOTE NLM and RSPX NLM are loaded on the server, you can start a remote session from a workstation by executing the RCONSOLE.EXE file and specifying the remote password. You don't need the Supervisor object right to the remote server, since RCONSOLE does not use NetWare Directory Services.

The following procedure explains how to run RCONSOLE from a network drive. You can also run RCONSOLE from your local hard drive.

Prerequisites




- A workstation running DOS 3.30 or later version
- A search drive mapped to the directory where the RCONSOLE.EXE file is located (usually SYS:SYSTEM)
- 300 KB of available memory on the workstation
- The remote console password



In NetWare 4.1, when you run RCONSOLE, you must enter the password that you establish when you load REMOTE.NLM. The supervisor password does not allow remote access, unless it is set as the password when you load REMOTE.NLM.

Procedure



1. Log in to the network.

1a. If you are logging in to a NetWare 4 network, type

```
LOGIN username <Enter>
```



If you do not have a context specified in your NET.CFG file, see "LOGIN" in *Utilities Reference*.

1b. If you are logging in to a NetWare 3 server, type

```
LOGIN servername/username <Enter>
```

Enter your password when prompted.

2. Make sure you have a drive mapped to the directory where the RCONSOLE.EXE file is located (usually SYS:SYSTEM).

If you don't have a drive mapped, type

```
MAP S16:=sys:system <Enter>
```

3. Start RCONSOLE by completing one of the following steps:

- 3a. If you know the name of the server, type the following command, and then skip to Step 6:**

```
RCONSOLE servername <Enter>
```

- 3b. If you know part of the server name, type the partial name followed by an asterisk:**

```
RCONSOLE partial_servername*
```

Select a server from the displayed list and then continue with Step 6.

- 3c. If you do not know the name of the server, type the following command, and then continue with Steps 4 and 5:**

```
RCONSOLE <Enter>
```

4. Select SPX from the “Connection Type” menu.

A list of servers available for a remote console session appears.

5. Connect to the server, using one of the following methods:

- 5a. Select a server from the list, and then go to Step 6.**

- 5b. Type the name of the server, and then go to Step 6.**

- 5c. If you know the IPX internal network number of the server, do the following:**

◆ Press <Insert>.

◆ Enter the IPX internal network number in hexadecimal notation, then continue with Step 6.

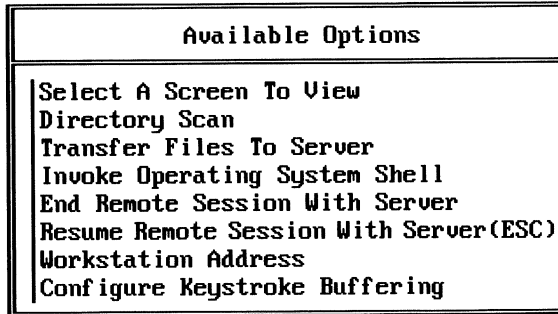
6. Type the server’s remote console password at the prompt.

This is the password you established when you loaded REMOTE.NLM.

The remote console session is now open. The workstation screen displays the same information you would see on the server console screen.

7. To view a menu of options for the remote session, press <Alt>+<F1>.

The remote console "Available Options" menu appears.



During an RCONSOLE session you can use the keys described in Table 7-6. All the other keys function as if you were at the server console.


**Table 7-6
RCONSOLE Keys**

Key	Description
<Alt>+<F1>	Access the "Available Options" menu.
<Alt>+<F2>	Exit Remote Console.
<Alt>+<F3>	Cycle backward through the current console screens.
<Alt>+<F4>	Cycle forward through the current console screens.
<Alt>+<F5>	Show the address of the workstation you are using for this session.
<F1>	Display Remote Console help when the Remote Console "Available Options" menu is on the screen. (When another screen is showing, this key displays the relevant server help information.)

Ending a Remote Console Session

You can exit the remote session from either the “Available Options” menu or the console prompt.

Procedure



- 1. From the “Available Options” menu, select “End Remote Session With Server,” or from the console prompt, press <Alt>+<F2>.**
- 2. Choose “Yes” from the “Quit Remote Console Session?” box.**

The next screen to appear depends upon how you invoked RCONSOLE.

If you specified the name of a specific server as a parameter when you executed RCONSOLE, the remote session with that server is ended and the DOS prompt is redisplayed.

If you selected the server from a list of available servers, the available servers list is redisplayed. To exit from the list of servers, press <Esc> and then select “Yes” from the “Exit Remote Console?” menu.

Additional Information

For more information about	Refer to
Passwords	“Password” in <i>Concepts</i>
RCONSOLE.EXE	“RCONSOLE” in <i>Utilities Reference</i>
Remote console	“Remote console” in <i>Concepts</i>

Running a Remote Console Session Over a Modem

This section explains how to run a remote console session over an asynchronous connection—a connection using a modem or null modem cable. To run the session, you must complete the following steps:

1. On the server, load REMOTE.NLM, AIO.NLM, and RS232.NLM. Also load the appropriate communications port driver, such as AIOCOMX.NLM.
2. Create a callback file, if desired.
3. Load necessary workstation files.
4. Execute RCONSOLE.EXE from the workstation and, on first execution, configure the modem.

Each of these tasks is explained in the following procedures.


Loading Modules on the Server

Prerequisites



- The following NLM files:
 - REMOTE.NLM
 - RS232.NLM
 - AIO.NLM
- A communications port driver such as AIOCOMX

Procedure



1. At the server console prompt, type

```
LOAD REMOTE [password] <Enter>
```

Enter the password you want to use when accessing the server from the remote workstation. A password is required.

2. Complete the following steps to load the communications interface module and driver:

2a. Load the communications port interface module by typing

```
LOAD AIO <Enter>
```

2b. Load the communications port driver by typing

```
LOAD AIOCOMX <Enter>
```

This driver is provided with NetWare. You can replace it with an equivalent driver if you choose.

2c. Load the asynchronous connection NLM by typing

```
LOAD RS232 [comm_port] [modem_speed] [N] [C]  
<Enter>
```

Replace *comm_port* with the communications port number (1 or 2).

Replace *modem_speed* with the baud rate (2400, 4800, or 9600).

Use the "N" parameter if you are using a null modem cable.

Use the "C" parameter if you want to use the callback option. The callback option is explained in the next procedure.

Creating a Callback List

A callback list enables you to create a list of authorized modem numbers that can be used to access the server.

When a connection attempt is made, the server notes the number of the modem that is calling and then terminates the connection. The server then compares the number to the numbers in the callback list.

If the number is in the list, the server calls the modem at that number and reestablishes the connection. If it is not in the list, the server ignores the call.

If you included the callback parameter when you loaded RS232.NLM in the preceding section ("Loading Modules on the Server"), complete the following steps:

Procedure



- 1. Create a CALLBACK.LST file in the SYS:SYSTEM directory on the server.**

Use a text editor or the EDIT utility to create the file. If you need help creating the file, see "Editing Text Files from the Server Console" on page 460.

- 2. In the file, enter a list of modem numbers that are authorized to start a remote console session.**

For example, you might list three modem numbers:


```
18015552257
5554321
4269
```

The first number authorizes a modem for a long distance number (with an area code), the second number is a local telephone number, and the third number is for a modem where only the extension number is required.

- 3. Save the text file.**

Preparing the Workstation

Procedure



1. **Prepare the workstation as a remote console by creating a directory and copying the following RCONSOLE run files to it:**

RCONSOLE.EXE
RCONSOLE.HEP
RCONSOLE.MSG
IBM_RUN.OVL
_RUN.OVL
IBM_AIO.OVL
_AIO.OVL
TEXTUTIL.HEP
TEXTUTIL.IDX
TEXTUTIL.MSG

It doesn't matter where you place the directory or what you name it. The run files come with your NetWare system.


Running RCONSOLE.EXE and Configuring the Modem

Prerequisites



- A workstation with a Hayes*-compatible modem (2400, 4800, or 9600 baud) or a null modem cable connected to the server
- 300 KB of available memory on the workstation
- REMOTE.NLM, AIO.NLM, RS232.NLM, and the communications port driver (such as AIOCOMX) loaded on the server
- All necessary files loaded on the workstation
- The remote console password

Procedure



1. **At the workstation, change to the directory you created for the remote files (described in the preceding section, “Preparing the Workstation”) and type**

RCONSOLE <Enter>

2. **Select “Asynchronous” from the “Connection Type” menu.**
3. **If this is the first session from this workstation, set up your modem by selecting “Configuration” from the “Asynchronous Options” menu, and then complete Steps 4 through 6. Otherwise, skip to Step 7.**
4. **In the “Current Modem Configuration” window, set the options according to your modem and workstation specifications.**
 - 4a. **Use the arrow keys to highlight a field, and press <Enter>.**

Either the cursor becomes active in the field or a selection list is displayed.
 - 4b. **If the cursor is active in the field, type in the appropriate information. If a selection list is displayed, select the appropriate item. Press <Enter>.**



The “UserID” and the “Call Back” fields must be filled in.

Choose any string for the UserID. For example, you may use your name or your phone number. This ID is displayed on the server console during the remote session.

In the “Call Back” field, enter the phone number you are calling from. If the “callback” option is used, this number must be in the callback list.

5. **Press <Esc> to exit the window.**
6. **Choose “Yes” from the “Save Changes?” box to save your configuration.**
7. **Choose “Connect to Remote Location” from the “Asynchronous Options” menu.**

A list of available servers is displayed.

8. Select a server.

If you are using the callback option, the server terminates the connection and compares the number in the modem configuration file with the numbers in the callback list.

If the number in the modem configuration file is in the list, the server dials the number to establish the connection and then displays the console prompt. If the number is not in the list, the server displays an error message.

9. Once the console prompt is displayed, press <Alt>+<F1> to view the “Available Options” menu for the remote console session.

During the RCONSOLE session, you may use the keystrokes shown in Table 7-6 on page 611. All other keys function as if you were at the server console.




RCONSOLE does not provide guaranteed packet delivery. When sending or receiving packets from a server with extensive disk or LAN activity, RCONSOLE may “time out.”

If this happens, you receive a “No response from server” or an “Unable to send request to server” message. If you receive one of these messages, retry the operation.

Ending a Remote Console Session

You can exit the remote session from either the “Available Options” menu or the console prompt.

Procedure



1. From the “Available Options” menu, select “End Remote Session,” or from the console prompt, press <Alt>+<F2>.
2. Select “Yes” from the “Quit Remote Console Session?” box.

Loading Remote Modules on Bootup

You can put commands in the AUTOEXEC.NCF file to load remote console modules each time the server is booted.

Since these modules require the remote password, you may choose to place an encrypted version of the password in the file, too. The encrypted password secures the server console by hiding the remote console password from users who access the AUTOEXEC.NCF file.

The following procedure explains how to set an encrypted password and what commands to place in the AUTOEXEC.NCF file.

Prerequisites



- The remote console password

Procedure



1. **To encrypt a password, type the following:**

```
LOAD REMOTE <Enter>
```

2. **Execute the following command:**

```
REMOTE ENCRYPT <Enter>
```

The server prompts you to enter the password to be encrypted.

3. **Enter the password you want to use for remote console sessions.**

The system displays the encrypted value and a message asking if the LOAD REMOTE command should be written to the SYS:SYSTEM\LDREMOTE.NCF file.

4. **Respond “Yes.”**

When you respond “Yes,” the system places a LOAD REMOTE command into this file with the encrypted password as a parameter.

You can execute this file whenever you want to load the REMOTE NLM, or you can place the command to execute the file into your AUTOEXEC.NCF file. This feature saves you from having to type in the encrypted password, which is often quite long.

5. Use either INSTALL or EDIT to open the AUTOEXEC.NCF file.

For information about editing .NCF files, see "Creating or Editing a Server Batch (.NCF) File" on page 459.

6. Move the cursor to the end of the file and type the following:

LDREMOTE

6a. For direct connections, type the following after LDREMOTE:

LOAD RSPX

6b. For asynchronous connections, type the following after LDREMOTE:

LOAD AIO

LOAD AIOCOMX

The AIOCOMX driver is provided with NetWare. You can replace it with an equivalent driver if you choose.

LOAD RS232 [comm_port] [modem_speed] [N] [C]

Replace *comm_port* with the communications port number (1 or 2).

Replace *modem_speed* with the baud rate (2400, 4800, or 9600).

Use the "N" parameter if you are using a null modem cable.

Use the "C" parameter if you want to use the callback option. For information and instructions, see "Creating a Callback List" on page 615.

7. Exit and save the AUTOEXEC.NCF file.

The server will now automatically load the necessary remote modules whenever it is booted.

Additional Information

For more information about	Refer to
EDIT utility	“EDIT” in <i>Utilities Reference</i>
LOAD command	“LOAD” in <i>Utilities Reference</i>
RCONSOLE utility	“RCONSOLE” in <i>Utilities Reference</i>
REMOTE utility	“REMOTE” in <i>Utilities Reference</i>
RS232 utility	“RS232” in <i>Utilities Reference</i>
RSPX utility (including the Packet Signatures option)	“RSPX” in <i>Utilities Reference</i>

Upgrading Files on a Remote Server


You can use a remote console to copy new NetWare files to a server. Files can be copied from diskettes, a CD-ROM drive, a local drive, or from a network directory. To copy new NetWare files to volume SYS: on a remote server, follow these steps.

Prerequisites



- NetWare diskettes or CD-ROM drive, if needed.
- Workstation and server prepared for remote console sessions. See “Running a Remote Console Session Through a Direct Connection” on page 607 and “Running a Remote Console Session Over a Modem” on page 613.

Procedure



1. Start a remote console session with the server.

To start a session over a LAN, see “Running RCONSOLE.EXE from a Network Workstation” on page 608.

To start a session over a modem connection, see “Running RCONSOLE.EXE and Configuring the Modem” on page 616.

2. At the remote console, type

LOAD INSTALL <Enter>

3. Choose “Copy Files Option” from the “Installation Options” menu.

4. Choose a remote path.

This is the path from which the files will be installed. You may load files from NetWare diskettes, a local drive, a CD-ROM, or a network drive.

5. Follow the subsequent prompts.

Installing or Reconfiguring Applications on a Remote Server


You can use a remote console session to install or reconfigure Novell or third-party products on a remote server. Follow these steps.

Prerequisites



- Application diskettes, if necessary.
- Workstation and server prepared for remote console sessions. See “Running a Remote Console Session Through a Direct Connection” on page 607 and “Running a Remote Console Session Over a Modem” on page 613.

Procedure



1. Start a remote console session with the server.

To start a session over a LAN, see “Running RCONSOLE.EXE from a Network Workstation” on page 608.

To start a session over a modem connection, see “Running RCONSOLE.EXE and Configuring the Modem” on page 616.

2. At the remote console, type

LOAD INSTALL <Enter>

- 3. Choose “Product Options” from the “Installation Options” menu.**
- 4. Select “View/Configure/Remove Installed Products” from the “Other Installation Actions” menu.**

A list of currently installed products (if any) appears.
- 5. To reconfigure an installed product, select that product and press <Enter>. Follow screen instructions.**
- 6. To install a new product, press <Insert>.**

A message states that the product will be installed from the A: drive, unless otherwise indicated.
- 7. If necessary, insert the product installation diskette into the drive.**
- 8. Press <Enter> to accept the default, drive A:, or press <F3> to enter a different installation source path.**
- 9. Follow the prompts to install the product.**

Rebooting a Remote Server

You can reboot a NetWare server from a remote console by first modifying the server's AUTOEXEC.BAT file and then executing another .NCF file that you create for this purpose.


Preparing the Server

Prerequisites



- Commands in the server AUTOEXEC.NCF file that load the remote modules. See "Loading Remote Modules on Bootup" on page 619.

Procedure



- 1. To automatically reboot the server, put commands to change to the NetWare 4.1 directory and execute SERVER.EXE in the server's AUTOEXEC.BAT file.**

For example, if the SERVER.EXE file is in the NET_4.1 directory on drive C:, put the following lines in the server's AUTOEXEC.BAT file:

```
CD C:\NET_4.1  
SERVER.EXE
```

- 2. Start a remote console session with the server.**

To start a session over a LAN, see "Running RCONSOLE.EXE from a Network Workstation" on page 608.

To start a session over a modem connection, see "Running RCONSOLE.EXE and Configuring the Modem" on page 616.

- 3. Load the EDIT NLM on the server by typing**

```
LOAD EDIT <Enter>
```

4. **Create an .NCF file and include the following commands in the order given:**


```
REMOVE DOS  
DOWN  
EXIT
```

See "Creating or Editing a Server Batch (.NCF) File" on page 459.

5. **Save the file in the SYS:SYSTEM directory.**

Rebooting the Server

Procedure



1. **Start a remote console session with the server.**

To start a session over a LAN, see "Running RCONSOLE.EXE from a Network Workstation" on page 608.

To start a session over a modem connection, see "Running RCONSOLE.EXE and Configuring the Modem" on page 616.

2. **Send a message to all users to inform them that the server is going down.**

See "Sending Console Messages to Workstations" on page 444 if you need help sending a message.

3. **At the system prompt on the remote console, enter the name of the .NCF file you created.**

This reboots the server. The keyboard is disabled for a few minutes while the file is processed.

When the file has finished processing, a "connection is lost" message appears.

4. **Press <Esc> to exit the remote session.**
5. **Open a new remote console connection with the server.**

Administering Accounting

This section describes how to set up NetWare accounting, calculate charge rates for using server resources, and view various system statistics.

Setting Up Accounting

Accounting allows you to monitor NetWare server usage and charge users for server resources.


Setting Up Accounting Using the NetWare Administrator

Prerequisites



- A 386 or 486 workstation running Windows 3.1 or OS/2 v2.x and NetWare Administrator.
- 6 MB of available memory on the Windows workstation, or 12 MB of available memory on the OS/2 workstation.
- Supervisor object right to the Server object where you are setting up accounting.

Procedure



- 1. From the Windows Program Manager or the OS/2 desktop, choose the “NetWare Administrator” icon.**
- 2. Select the Server object you want to set up accounting for.**
- 3. From the “Object” menu, Choose “Details.”**
- 4. From the bottom of the dialog box, choose “Accounting.”**
- 5. To confirm that you want to install accounting on this server, choose “Yes.”**

Accounting page buttons appear on the right side of the dialog box with the other page buttons.

6. Use the accounting buttons to edit the accounting pages.

7. Choose "OK."


Setting Up Accounting Using NETADMIN

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 512 KB of memory available on the workstation.
- Supervisor object right to the Server object you want to set up accounting for.

Procedure



1. At the workstation prompt, type
NETADMIN <Enter>
2. From the "NetAdmin Options" menu, choose "Manage Objects."
3. Select the Server object you want to set up accounting for by completing the following steps:
 - 3a. If the server appears on the list, select it and press <F10>.
 - 3b. If the server is not on the list, browse the directory by selecting objects and pressing <Enter> until you see the server you want. Select it and press <F10>.
4. Choose "View or Edit Properties of this Object."
5. Choose "Accounting."
6. From the "Install Accounting" prompt, choose "Yes."
7. Choose "Accounting Servers."
8. Press <Insert>.

9. From the list of possible servers, choose "NetWare Server."
10. Enter the name of the server, or press <Insert> and select a server from the list.
11. To return to the "Accounting Options" menu, press <Esc>.
12. Choose the options you want to set up from the menu.
Press <F1> if you need help with any of the options.
13. To exit NETADMIN, press <Alt>+<F10>.

Calculating Charge Rates for Accounting

Charge rates for using server resources are specified as multipliers and divisors. To set the charge amount to "No Charge," enter 0 as the multiplier. Use the following formula to calculate a charge rate:

$$\frac{\text{CHARGE (charge rate multiplier)}}{\text{ESTIMATED USAGE (charge rate divisor)}} = \text{CHARGE RATE}$$

For example, if you want to charge \$200 per month for server disk storage, and there are usually 100,000 blocks used each month on your hard drive:

- ◆ 20,000 would be the multiplier (\$200 in cents)
- ◆ 100,000 would be the divisor (disk block usage per month)
- ◆ Therefore, users would be charged two-tenths of one cent per disk block used.

$$\frac{20,000 \text{ cents}}{100,000 \text{ blocks read}} = 1 \text{ cent} / 5 \text{ blocks read}$$

Viewing Accounting Totals

The ATOTAL utility provides a total for the accounting services used on your network. ATOTAL compiles information from the system accounting records and lists the following:

- ◆ Total connect time in minutes
- ◆ Total service requests
- ◆ Total blocks read
- ◆ Total blocks written
- ◆ Total disk storage in blocks per day

Prerequisites



- A workstation running DOS 3.30 or later, or OS/2 v2.x.
- Accounting set up on the NetWare server. (See “Setting Up Accounting” on page 626.)
- Supervisor object right to the Server object you want to see accounting information for.

Procedure



1. **Change to the SYS:SYSTEM directory.**
2. **To view daily and weekly totals for accounting services, type**
ATOTAL <Enter>
3. **To redirect ATOTAL data to a file, type**
ATOTAL > filename <Enter>




What Is NetWare Auditing?

In general, auditing means examining records to make sure that transactions are accurate and that confidential information is secure.

NetWare® auditing allows individuals, acting independently of network supervisors and other users, to audit network transactions. A *transaction* is defined here as any action that changes the NetWare Directory Services™ (NDS) database or a volume's content.

Auditors can audit NDS™ events as well as those specific to a file system volume or a server. Some of the events that can be audited are shown in the following figure.

Figure 8-1
Events You Can
Audit



Auditor

File or directory events

- Create, modify, delete directories or files
- Salvage, move, rename directories or files
- Create, delete, service print queues

Server events

- Bring down a server
- Create or delete bindery objects
- Mount or dismount volumes
- Modify security rights

Directory Services events

- Add, delete objects
- Move, rename objects
- Add, remove security equivalence
- Track User object logins and logouts

Auditors can track events and activities on the network, but they do not have rights to open or modify network files, other than Audit files, unless they are granted rights by the network supervisor.

Auditing is enabled at the volume level for file system auditing. It is enabled at the container level for auditing of NDS events.

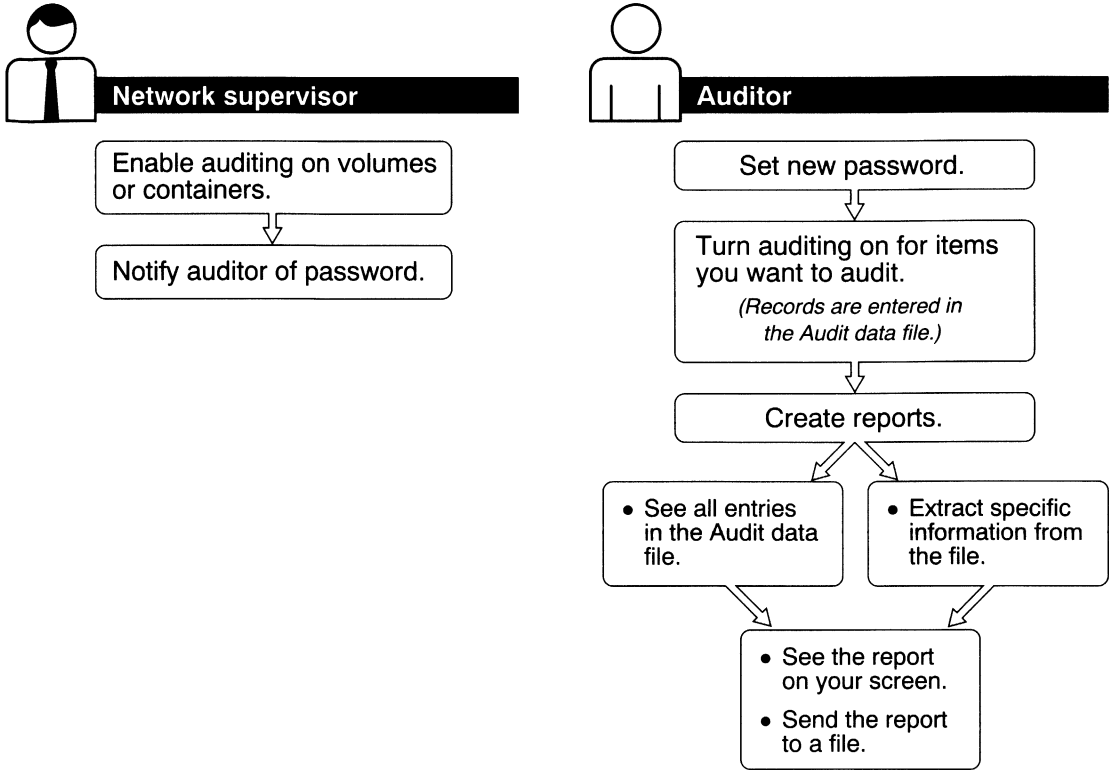


The AUDITCON program files are installed automatically on your system when you install or upgrade to NetWare 4™.

The procedures in this chapter address two audiences. “Enabling Auditing” on page 634 is for the network supervisor. “Auditor Procedures” on page 639 and “Creating Audit Reports” on page 655 are for those assigned to audit NDS events or changes to NetWare volumes.

Figure 8-2 illustrates the responsibilities of the network supervisor and the auditor.

Figure 8-2
Network Supervisor and Auditor Tasks



Audit files are automatically created for volumes and containers that have auditing enabled. The files keep records much like a system or error log file. All activity tracked by the auditing utility is entered into the Audit files.

The files continue to accept records until auditing is disabled or the file gets full.

For more information about	Refer to
Auditing	“Auditing” in <i>Concepts</i> “AUDITCON” in <i>Utilities Reference</i>
Audit files and how they work	“Maintaining Audit Files” on page 662

Enabling Auditing


Creating a Directory and an Auditor User Object

Prerequisites



- A workstation running DOS 3.30 or later
- The Supervisor or equivalent rights in the file system where you want to create the auditor’s directory
- The Create object right in the container where you will create the auditor as a User object

Procedure



1. Create a directory in the file system for the auditor to use.

The auditor needs space to store audit reports. Consider creating a directory in SYS:HOME or wherever you have created other user account directories.

2. Create the auditor as a User object.

If the auditor will audit NetWare Directory Services events, assign him or her the Browse object right in the container objects to be audited.

3. Give the auditor trustee rights to the directory you created for the auditor.

You can grant the auditor the Supervisor trustee right to the directory, or you can assign the auditor all rights but the Supervisor right. Either method produces the same result.

4. Map a drive to the directory containing the audit program files.

If you have not included a search drive mapping to SYS:PUBLIC in the system login script, create a user login script for the auditor and map a drive to this directory.

Also give the auditor the Browse object right and the File Scan directory trustee right to SYS:PUBLIC. (AUDITCON and Unicode files are in SYS:PUBLIC, unless you move them).

Additional Information

For more information about	Refer to
Assigning rights to directories and files	"Understanding File and Directory Rights" on page 131 and "Adding a Trustee to a Directory or File" on page 134
Creating a file system directory	"Creating Directories and Copying Files" on page 109
Creating a User object	"Creating Leaf Objects" on page 35
Assigning trustee rights to a user	"Adding a Trustee to a Directory or File" on page 134
Creating user login scripts	"Creating, Modifying, Copying, and Printing Login Scripts" on page 175
Modifying the system login script	"Creating, Modifying, Copying, and Printing Login Scripts" on page 175
Assigning rights to objects	"Rights Needed to Create and Manage Objects" on page 17.

Enabling Auditing for a Volume

Prerequisites

Checklist



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- Supervisor object right to the Volume object where you are enabling auditing

Procedure

Procedure



1. Access the auditing utility by typing

AUDITCON <Enter>

Your current server and volume are displayed at the top of the screen. If you want a different server or volume than the one shown, complete Step 2.

2. (Optional) Change the server or volume.

2a. Select “Change Current Server” or “Change Current Volume” from the “Available Audit Options” menu.

2b. Press <Insert> to see a list of available servers.

2c. Select the server, and then select the volume.

3. Select “Enable Volume Auditing” from the “Available Audit Options” menu.

4. Enter a password for the volume.

5. Enter the password again when prompted.

6. Notify the auditor of the password.

After auditing is enabled, the independent auditor assumes control of the program. If the auditor changes the password, the supervisor no longer has access to auditing information.

See “Auditor Procedures” on page 639.

Additional Information

For more information about	Refer to
Auditing	“Auditing” in <i>Concepts</i> “AUDITCON” in <i>Utilities Reference</i>
Passwords	“Password” in <i>Concepts</i>

Enabling Auditing for a NetWare Directory Services Container

When you enable auditing for a container object (Organization or Organizational Unit), it does not enable auditing for subordinate container objects.

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- Supervisor object rights to the container holding the Server object

Procedure



- 1. Log in to the network from the workstation.**
- 2. Access the auditing utility by typing**

AUDITCON <Enter>

Your current server and volume are displayed at the top of the screen.

- 3. Select “Audit Directory Services” from the “Available Audit Options” menu.**

Your current context is displayed at the top of the screen. If you want a different context than the one shown, complete Step 4.


4. (Optional) Change the context.

4a. Select “Change Session Context” from the “Audit Directory Services” menu.

4b. Enter the context for the container where auditing is being enabled, and press <Enter>.

Type the Organizational Unit name *before* the displayed organization name, and separate the two by a period. For example, a graphical representation of a directory tree with the “Sales” Organizational Unit is shown below.

Figure 8-3
Sample Directory
Tree



The context for the container is:

SALES . NOVELL

4c. Press <Esc> to get back to the “Available Audit Options” menu.

5. Select “Audit Directory Tree” from the “Audit Directory Services” menu.

6. Browse the Directory tree to locate the container where you are enabling auditing, and press <F10>.

7. Select “Enable Container Auditing” from the “Available Audit Options” menu.

8. Enter a password for the container.

9. Enter the password again when prompted.

10. Notify the auditor of the password.

11. Select “Display Audit Status.”

12. Verify that “Auditing Status” is On.

After auditing is enabled, the independent auditor assumes control of the program. If the auditor changes the password, the supervisor no longer has access to auditing information.

Auditor procedures begin in the following section.

Additional Information

For more information about	Refer to
Auditing	“Auditing” in <i>Concepts</i> “AUDITCON” in <i>Utilities Reference</i>
Passwords	“Password” in <i>Concepts</i>

Auditor Procedures

To keep your audit information secure, we suggest that you perform the following tasks before you begin auditing:

- ◆ Change the auditing password given to you by the network supervisor. (See “Setting Passwords” on page 640.)
- ◆ Log in to the volume or container using the new password. (See “Logging In” on page 643.)
- ◆ Set your preferences for the auditing environment. (See “Setting Up Your Work Environment” on page 646.)

From this point, you can select what you want to audit and create reports based on the audit records you collect. (See “Selecting Items to Audit” on page 650.)

Setting Passwords

Changing the Password for a Volume

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- The password assigned to the volume by the network supervisor or the password currently being used
- The Browse right to the volume you are auditing

Procedure



- 1. Log in to the network from the workstation.**
- 2. Access the auditing utility by typing**

AUDITCON <Enter>

Your current server and volume are displayed at the top of the screen. If you want a different server or volume than the one shown, complete Step 3.

- 3. (Optional) Change the server or volume.**
 - ◆ To change to a different server, select "Change Current Server" from the "Available Audit Options" menu, and then select the server you want.
 - ◆ To change to a different volume, select "Change Current Volume" from the "Available Audit Options" menu, and then select the volume you want.
- 4. Select "Auditor Volume Login" from the "Available Audit Options" menu.**

A password prompt for the volume appears.
- 5. Type in the current auditing password for the volume.**

6. Select “Auditing Configuration” from the “Available Audit Options” menu.
7. Select “Change Audit Password” from the “Auditing Configuration” menu.
8. Enter the current password.
9. Enter the new password.
10. Retype the new password.

The password is changed.

Changing the Auditor Password for a NetWare Directory Services Container

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- The password assigned to the container
- The Browse right to the container you are auditing.

Procedure



1. Log in to the network from the workstation.
2. Access the auditing utility by typing

AUDITCON <Enter>

3. Select “Audit Directory Services” from the “Available Audit Options” menu.

Your current context is displayed at the top of the screen. If you want a different context than the one shown, complete Step 4.

4. (Optional) Change the current context.
 - 4a. Select “Change Session Context” from the “Audit Directory Services” menu.
 - 4b. Enter the context to the container where you are enabling auditing, and press <Enter>.
5. Select “Audit Directory Tree” from the “Audit Directory Services” menu.
6. Locate the container where you are enabling auditing by pressing <Enter> to move around in the Directory tree.
7. Select the container you want and press <F10>.



You will see only containers to which you have been given the Browse right. If you do not see the container you want to audit, contact the network supervisor.

8. Select “Auditor Container Login” from the “Available Audit Options” menu.
9. Enter the current container password.
10. Select “Auditing Configuration” from the “Available Audit Options” menu.
11. Select “Change Audit Password” from the “Auditing Configuration” menu.
12. Type in the current password for the container.
13. Enter the new password.
14. Retype the new password.

The password is changed.

Setting and Changing Second-Level Passwords

As an auditor, you can set an additional level of security for your audit records by enabling a second-level password.

If you set this option, a second password is required before any changes to the audit configurations or report filters can be made.

Instructions for enabling and changing a second-level password are in "Setting Up Your Work Environment" on page 646.

Logging In

Logging In to a Volume

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- The password assigned to the volume

Procedure



1. **Log in to the network.**
2. **Access the auditing utility by typing**

AUDITCON <Enter>

Your current server and volume are displayed at the top of the screen. If you want a different server or volume than the one shown, complete Step 3.

3. (Optional) Change the server or volume.

3a. To change to a different server, select “Change Current Server” from the “Available Audit Options” menu, and then select the server you want.

3b. To change to a different volume, select “Change Current Volume from the “Available Audit Options” menu, and then select the volume you want.

4. Select “Auditor Volume Login” from the “Available Audit Options” menu.

A password prompt for the volume appears.

5. Type the auditing password for the volume.


Logging in to a NetWare Directory Services Container

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- The password assigned to the container by the network supervisor or the password currently being used
- Browse right to the container you want to audit

Procedure



- 1. Log in to the network.**
- 2. Access the auditing utility by typing**

AUDITCON <Enter>

The current server and volume are displayed at the top of the screen.

3. Select “Audit Directory Services” from the “Available Audit Options” menu.

Your current context is displayed at the top of the screen. If you want a different context than the one shown, complete Step 4.

4. (Optional) Change the current context.


- 4a. Select “Change Session Context” from the “Audit Directory Services” menu.

- 4b. Enter the context to the container where you are enabling auditing, and press <Enter>.

Type the Organizational Unit name *before* the displayed organization name, and separate the two by a period.

For example, a graphical representation of a Directory tree with the “Sales” Organizational Unit is shown below.

Figure 8-4
Sample Directory
Tree



The context for the container is:

SALES . NOVELL

- 4c. Press <Esc> to get back to the “Available Audit Options” menu.
5. Select “Audit Directory Tree” from the “Audit Directory Services” menu.
6. Locate the container where you are enabling auditing by using the Up- and Down-arrow keys and pressing <Enter> to move around in the Directory tree.



7. Select the container you want and press <F10>.

You will see only containers to which you have been given the Browse right. If you do not see the container you want to audit, contact the network supervisor.

8. Select “Auditor Container Login” from the “Available Audit Options” menu.

A password prompt for the container appears.

9. Type the auditing password for the container.

You are now logged in to the container.

Setting Up Your Work Environment

You can set several options that determine how auditing will work in the volume or container you are auditing. You can


- ◆ Specify the size of the Audit Data file
- ◆ Determine what action is taken if the Audit Data file reaches the maximum size
- ◆ Set a second-level password for additional security

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- Browse right to the container, if you want to audit NetWare Directory Services events

Procedure



1. Log in to the volume or container.

If you need help, see “Logging In” on page 643.

2. Select “Auditing Configuration” from the “Available Audit Options” menu.

3. Select “Audit Options Configuration” from the “Auditing Configuration” menu.

The “Audit Configuration” screen appears.

4. Select the applicable options and enter the settings you want for your work environment.

The fields are explained in Table 8-1.

5. To save the settings, press <Esc> and choose “Yes” at the prompt.

Table 8-1
Setup Options for the Auditing Environment

Option	Explanation
Audit File Maximum Size	Enter the file size limit (in bytes) you want for the Audit Data file. Discuss this issue with the network supervisor before changing the default size because this affects available disk space for other network services.
Audit File Threshold Size	Enter the file size threshold (in bytes) beyond which AUDITCON sends warning messages to the server console and the system log file. The threshold setting should be 90% of the maximum size. For example, a maximum setting of 1,000,000 bytes should have a threshold setting of 900,000 bytes.
Automatic Audit File Archiving	<p>Auto-archiving provides automated backup of your audit files. Automating the process provides protection of your audited information; it also requires large (duplicate) amounts of disk space, depending on the number of old audit files you decide to keep on the server. Use in conjunction with the next three options.</p> <p>Discuss this option with the network supervisor, since it affects available disk space on the server. Make sure the network supervisor has not set your storage limits too low.</p>

Table 8-1 *continued*

Setup Options for the Auditing Environment

Option	Explanation
Days Between Audit Archives (1–255)	Set the number of days to collect data before the file is automatically archived. You can select between 1–255 days; the default is 7 days. This option is valid only when the “Automatic Audit File Archiving” option is set to Yes.
Hour of Day to Archive (0–23)	<p>Select the hour of the day for the auto-archive to take place. You can select any hour of the day, using a 24 hour clock (0–23). The default is 0 (midnight).</p> <p>This option is valid only after you have turned on the auto-archive option and selected the number of days between archives. Discuss this option with the network supervisor because it affects network traffic.</p>
Number of Old Audit Files to Keep (1–15)	<p>This option works in conjunction with the “Automatic Audit File Archiving” option and also with the error recovery feature of the “Archive Audit File” option. Either option closes the current audit files, archives copies of the current files, then opens new copies of the audit files. Those files that are archived are automatically stored on the server. When the maximum number of old audit files is reached, the oldest of the old files is deleted.</p> <p>Discuss this option with the network supervisor to determine possible disk storage limitations.</p>
Allow Concurrent Auditor Logins	Choose “Yes” if you want to allow more than one auditor to log in to a volume or container at the same time. All auditors use the same password.
Broadcast Error Messages to All Users	Choose “No” if you want error messages sent only to the server console. We recommend that you do not broadcast these messages, because they increase network traffic and lock users’ workstation screens until they press <Ctrl>+<Enter>.
Force Dual-Level Audit Passwords	Choose “Yes” if you want to require two passwords. You are prompted to enter a password when you set this field to “Yes” for the first time. Whenever you change an auditing setting, you are prompted for a second password before the change is made.

Table 8-1 *continued*

Setup Options for the Auditing Environment

Option	Explanation
Archive Audit File	<p>This is one of three options available when the Audit Data file is full. The three options are: "Archive Audit File," "Dismount Volume," and "Disable Event Recording." Discuss these options with the network supervisor.</p> <p>Consider the "Number of Old Audit Files to Keep" option when selecting this option. When the Audit Data file reaches its maximum size, it is automatically closed, archived, and replaced with a new Audit Data file. No audit data is lost, and users continue to access the container or volume.</p> <p>Compare the "Archive Audit File" option with the "Dismount Volume" and "Disable Event Recording" options before deciding which to use.</p>
Dismount Volume (volume auditing only)	<p>Dismounts the audited volume if the Audit Data file reaches the maximum size limit or an unrecoverable write error occurs. <i>Do not</i> set this option to "Yes" without discussing it with the network supervisor, because dismounting a volume locks out all users. If the volume is dismounted, the network supervisor must remount it.</p> <p>NOTE: If the volume has been dismounted and the auditor has changed the assigned password, the network supervisor must have the auditor's password to remount the volume.</p> <p>See "Remounting a Volume Dismounted by Auditing" on page 665 for more information. The default option is "Disable Event Recording."</p>
Disable Audited Events (NetWare Directory Services auditing only)	<p>When this option is set to "Yes," all audited events applicable to the container object are disabled when the Audit Data file reaches the maximum size limit or an unrecoverable write error occurs. <i>Do not</i> set this option to "Yes" without discussing it with the network supervisor, because disabling NDS events prevents users from performing any task that is being audited. For example, if logins are being audited, all attempts to log in to the container fail. The default option is "Disable Event Recording."</p>

Table 8-1 *continued*

Setup Options for the Auditing Environment

Option	Explanation
Disable Event Recording	<p>This is the default of the three possible options. The program stops entering new records into the Audit Data file when the maximum size limit is reached or an unrecoverable write error occurs. You must reset the file to enable event recording again.</p> <p>Until you reenable event recording, no further auditing information is recorded. However, this option allows users to continue to access the volume and perform their work.</p> <p>When you reset the file, the information is saved in the old Audit Data file, and the auditing function resumes.</p>
Minutes Between Warning Messages	<p>Warnings are sent to the server console at this frequency if the Audit Data file is full and the "Disable Event Recording" option is set to "Yes."</p>

Selecting Items to Audit

This section explains how to set auditing to gather the information you want from a volume or container.


After you have logged in and accessed the "Available Auditing Options" menu, you can move between volumes and containers to customize auditing procedures for all areas to which you have a valid password and the appropriate rights.

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 470 KB of memory available on the workstation.
- Have AUDITCON started and be logged in to the volume or container with the appropriate password. (See "Logging In" on page 643 if you need help.)

Procedure



1. Select “Auditing Configuration” from the “Available Auditing Options” menu.
2. Select the appropriate options from the “Auditing Configuration” menu.


You might need to set auditing in more than one area to get the results you want. If you need help understanding the options and how they interact, read the following explanations, and then see Figure 8-5.

- ◆ **User and File** means you have limited the auditing of an event to a specific user and file. An auditing record will be entered only if a file in the audited directory is opened or written to by the audited user.
- ◆ **User or File** means you want to audit all occurrences of an event in a directory regardless of the user who performed the event, or you want to audit all occurrences of an event by a user regardless of the directory in which the event took place.
- ◆ **Global** means you are auditing for all instances of an event throughout the volume.




See “Auditing Examples” on page 666 to see how “User and File” and “User or File” options are used.

Figure 8-5
User, File, and Global Options



Records a file write only if User1 writes to a file in DirectoryB.


-
1. Turn on File Write User and File.
 2. Turn on auditing for DirectoryB.
 3. Turn on auditing for User1.



Records all file writes to DirectoryB, by any user.

Records file writes by User1 to any directory.

-
1. Turn on File Write User or File.
 2. Turn on auditing for DirectoryB or User1.



Audits all writes that occur in the volume or container. Includes all files and all users.

-
1. Turn on File Write Global.

Options you can select for auditing volumes are described in Table 8-2. Options for auditing NetWare Directory Services events are described in Table 8-3.

Table 8-2
Auditing Options for Volumes

Option	Explanation
Audit by Event	Allows you to select the types of events to audit. If you select this option, AUDITCON prompts you to select from the event types explained below. Events are dependent on the files/directories and users audited. Unless you are auditing an event globally, you must also turn on auditing for the files/directories and users you want to audit for the event. For more information, see "Auditing Examples" on page 666.
File Events	Includes file or directory activities such as directory create or delete, and file create, open, close, read, write, and salvage requests. Events can be audited globally (all files/directories are audited for the event regardless of user), by user or file (the event is recorded only when it applies to an audited user or file), or user and file (the event is recorded when it applies to an audited user in an audited file or directory). For more information about how file events interact with users and files, see "Auditing Examples" on page 666.
QMS Events	Includes activities affecting queues, such as requests to create or destroy a print queue.
Server Events	Includes requests that affect a specific server, such as mounting a volume or bringing down a server. Server events are audited globally (all instances of the audited event are recorded).
User Events	Includes activities such as logins and logouts, trustee assignment changes, and creating or deleting users.
Audit by File/Directory	Gives you a list of all files and directories from which you can select those you want to audit. If you audit a file or directory, a record is added to the Audit Data file each time an audited event occurs for that file. If you set a parent directory to "Audited," it has the same effect as browsing up one level in the Directory tree and setting the actual directory name as audited. For more information, see "Auditing Examples" on page 666.
Audit by User	Gives you a list of users from which you can select those you want to audit. If you turn on auditing for a user, a record is entered in the Audit Data file each time the user performs a user event that is audited.

Table 8-2 *continued*

Auditing Options for Volumes

Option	Explanation
Audit Options Configuration	Allows you to configure the audit file parameters and specify what action will be taken if the Audit Data file reaches the specified size limit. For instructions on using this option to set up your working environment, see "Setting Up Your Work Environment" on page 646.
Change Audit Password	Allows you to change the auditor password for the volume.
Disable Volume Auditing	Disables auditing for this volume. If you select "Yes," the network supervisor will have to enable auditing again if it is needed in the future.
Display Audit Status	Displays information about a volume's Audit Data and Audit History files.

If you are auditing NetWare Directory Services events, use the options explained in Table 8-3 to set up auditing.

Table 8-3

Auditing Options for NetWare Directory Services

Option	Explanation
Audit by DS Event	Lists NetWare Directory Services events from which you can select those you want to audit. When you turn on auditing for an event, all occurrences of the event within the container are recorded.
Audit by User	Lists Directory tree objects. You can select the users to audit from the list by selecting the user's name and pressing <F10>. If you turn on auditing for a user, a record is entered in the Audit Data file each time the user performs a NetWare Directory Services event that is audited.
Audit Options Configuration	Allows you to configure the audit file parameters and specify what action is taken if the Audit Data file reaches the specified size limit. For instructions on using this option to set up your working environment, see "Setting Up Your Work Environment" on page 646.
Change Audit Password	Allows you to change the auditor password for the container.
Change Audit Password Two	Only shows if "Force dual-level password" has been set in the Auditing Configuration screen. Allows you to change the second-level password for the container.

Table 8-3 *continued*

Auditing Options for NetWare Directory Services

Option	Explanation
Disable Container Auditing	Disables auditing for this container. If you select "Yes", the network supervisor must reenabling auditing if it is needed in the future. If "Force Dual Level Password" was set to "Yes," the second password is required to reenabling auditing.
Display Audit Status	Displays general information about the Audit Data file for the container.

Creating Audit Reports

You have several options available for creating and viewing reports from the records in the audit files. When creating a report, you can

- ◆ See all records in the Audit Data file
- ◆ Create filters to extract specific information from the Audit Data file

When viewing a report, you can

- ◆ See a report on your screen
- ◆ Send the report to a file in your default directory or to a local drive

Creating Report Filters

AUDITCON lets you create filters so you can select information from the Audit Data file to create reports. If you view a report without applying a filter, the entire contents of the file are included.


You can create as many filters as you want to screen information in the Audit Data file. Then, you can select and apply the filter you want when you generate a report.

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 470 KB of memory available on the workstation.
- Have AUDITCON running and be logged in to the volume or container with the appropriate password. (See “Logging In” on page 643 if you need help.)

Procedure



1. **Select “Auditing Reports” from “Available Auditing Options.”**
2. **Select “Edit Report Filters” from “Available Auditing Options.”**

The “Edit Filter” list appears, which includes all filters previously created.

3. **Select a filter from the list or press <Insert> to create a new filter.**

The “Edit Report Filter” menu appears.

4. **Select options you want to include in the filter from the “Edit Report Filter” menu.**

The options are explained in Table 8-4. For help with any of the options, press <F1>.



If a filter contains conflicts between “include” and “exclude” options, the “exclude” always takes priority.

Table 8-4
Options for Creating Report Filters

Option	Explanation
Report by Date/Time	Allows you to enter a time period for which you want to create a report. All recorded events that fall within the dates and times entered appear in the report.
Report By Event	<p>Allows you to select the types of audited events to include in the report. If you select this option, AUDITCON prompts you to select from the events that have been audited. Events are dependent on the files/directories and users included in or excluded from the report. Event types you can select are:</p> <ul style="list-style-type: none"> ◆ File Events ◆ QMS Events ◆ Server Events ◆ User Events
Report Exclude Paths/Files	Allows you to specify files or directories you do not want to include in the report. Press <Insert> and enter the name of the file or directory you want to exclude. All other audited files and directories are included.
Report Exclude Users	Allows you to specify the name of an audited user that you do not want to include in the report. Press <Insert> and enter the name of the user you want to exclude, or press <Insert> again to select from a list. All other users are included.
Report Include Paths/Files	Allows you to specify an audited directory path or file that you want to include in the report. Press <Insert> and enter the directory path (or the filename, if the file is in the default directory) that you want to include, or press <Insert> again to select from a list. The default is that all audited directories and files are included. This is indicated by the asterisk (*). All other paths and files are excluded.
Report Include Users	Allows you to specify the name of an audited user that you want to include in the report. The default is that all users are included. This is indicated by the asterisk (*). If you don't want to include all users, press <Insert> and enter the name of the user you want to include; all other users are excluded.

5. Press <Esc> until you get back to the “Edit Report Filter” menu.
6. At the “Save Filter Changes” prompt, complete one of the following steps:
 - 6a. To save the change, select “Yes” and press <Enter>. If you have created a new filter, backspace to delete “_no_filter_” from the list and enter a name (maximum eight characters) for the filter.
 - 6b. To exit the filter without saving the changes, select “No” and press <Enter>.
 - 6c. To continue editing the filter, press <Esc>.

Sending Reports to a File

AUDITCON lets you send auditing information to a directory on the network.


You can send the information to your default network directory, a directory where you have [RWCEMF] rights, or a local or floppy drive at your workstation.

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 470 KB of memory available on the workstation.
- Have AUDITCON running and be logged in to the volume or container with the appropriate password. (See “Logging In” on page 643 if you need help.)

Procedure



1. **Select "Auditing Reports" from the "Available Auditing Options" menu.**
2. **Select "Report Audit File" from "Available Auditing Options."**
3. **Specify where you want the file to be sent by completing one of the following:**



If you send reports to a network directory, a user with Supervisor rights in the file system can see the file. If you want to secure your reports, you should send them to a local drive or a floppy drive on your workstation.

- ◆ Press <Enter> to copy the report to a file in your default directory.

The default filename is AUDITDAT.TXT.

- ◆ Enter the directory path and filename where you want the file to be sent, and then press <Enter>.

Specify an existing directory path and filename. (The program will not create a default file for you.)

4. **Select a filter from the "Select Filter List" and press <Enter>. To create a report without a filter, select "_no_filter."**
5. **(Optional) Create or edit a filter. To edit a filter in the list, press <F10>. To create a filter, press <Insert>.**

The list item "_no_filter" cannot be edited.

For help creating and editing filters, see "Creating Report Filters" on page 655 or press <F1> for help.

6. Specify how you want the filter to be used by completing one of the following steps:

- ◆ To save the filter changes, select "Save Filter Changes."

If it is a new filter, enter a name (maximum eight characters).

- ◆ To use the filter once without permanently changing the filter, select "Use Filter without Saving Changes."

The report is generated with the new filter settings, and then the filter reverts to its original settings. New filters are not saved.

- ◆ To discard filter changes, select "Discard Filter Changes".

Edited filters revert to the original settings. New filters are deleted.

The filtered report is sent to the directory and file you specified. Reports are in DOS text file format. The report file can be edited or printed using the same methods you use for other DOS text files.


Viewing a Report on the Screen

Prerequisites



- A workstation running DOS 3.30 or later.
- A minimum of 470 KB of memory available on the workstation.
- Have AUDITCON running and be logged in to the volume or container with the appropriate password. (See "Logging In" on page 643 if you need help.)

Procedure



1. **Select “View Audit File” from the “Auditing Reports Menu.”**
2. **Select a filter from the “Select Filter List” and press <Enter>. To see a report without a filter, select “no_filter” from the list.**
3. **(Optional) Create or edit a filter. To edit a filter in the list, press <F10>. To create a filter, press <Insert>.**

The list item “_no_filter” cannot be edited.

For help creating and editing filters, see “Creating Report Filters” on page 655, or press <F1> for online help.

4. **Specify how you want the filter to be used by completing one of the following:**
 - ◆ To save new changes to the filter, select “Save Filter Changes.”
If you created a new filter, enter a name (maximum eight characters).
 - ◆ To use the filter for this report only, select “Use Filter Without Saving Changes.”
The report is generated with the new filter settings, and then the filter reverts to its original settings. New filters are not saved.
 - ◆ To discard filter changes, select “Discard Filter Changes.”
Edited filters revert to the original settings. New filters are deleted.
5. **Press the arrow keys to scroll through the report. To exit, press <Esc> and answer “Yes” at the “Exit View” prompt.**

Maintaining Audit Files

Understanding the Audit Files

AUDITCON automatically creates audit files when it is enabled. Each volume or container using AUDITCON will have its own audit files.

The Audit Data File

This file keeps records of all audited transactions. The auditing configuration you set determines the type of records entered into the data file.

This file operates like a system log or error file. Records are automatically entered into the file whenever an audited event occurs.

When you audit container objects for NetWare Directory Services events, the audit history information is combined with audit data into this file.

The Audit History File

This file keeps a record of the auditor's activities in a volume, such as auditing configuration changes and auditor logins and logouts.

In NetWare Directory Services, the auditor's activities are recorded in the Audit data file. There is no separate file for this information.

You can, however, use the menu options such as "View Audit History" to see auditor records just as you would if you were auditing a volume.

When you reset an audit file, existing records are moved to an Old Audit Data or Old Audit History file. The original audit file continues to function as it did before.


The maintenance options for the audit files are explained in the following table.

Table 8-5
Audit Files Maintenance Options

Option	Explanation
Close Old Audit File	Allows you to close an Old Audit Data file. All Audit Data files are kept open by the system for security reasons. To copy or delete an old audit file, first close it, then copy or delete it.
Copy Old Audit File	Allows you to copy an Old Audit Data file to a file or directory as a compressed, nonreadable file. You must first close the old audit file. You are prompted for the name of the destination file. The default name is AUDITOLD.DAT.
Delete Old Audit File	Allows you to delete the Old Audit Data file. You must first close the old audit file.
Display Audit Status	Displays basic information about the audit files, such as file size and file size limits
Reset Audit Data File	Resets the Audit Data file and moves the data in the file to the Old Audit Data file. Old audit files are limited to 15 at a time. When the maximum number of old audit files is reached, the oldest old file is deleted and number 14 becomes number 15, and so forth. The Audit Data file continues to collect records based on the configuration settings, as it did before it was reset.
Reset Audit History File	For auditing volumes only. Resets the Audit History file and moves the data in the file to the Old Audit History file. Old audit files are limited to 15 at a time. When the maximum number of old audit files is reached, the oldest old file is deleted and number 14 becomes number 15, and so forth. The Audit History file continues to collect records, as it did before it was reset.


Resetting the Audit Data File

Prerequisites



- A workstation running DOS 3.30 or later
- A minimum of 470 KB of memory available on the workstation
- Have AUDITCON running and be logged in to the volume or container with the appropriate password. (See “Logging In” on page 643 if you need help.)

Procedure



- 1. Select “Audit Files Maintenance” from the “Available Auditing Options” menu.**
- 2. Select “Reset Audit Data File” from the “Audit Files Maintenance” menu.**

A warning message appears on the screen to notify you that the current contents of the file will be moved to the Old Audit Data file. If you already have data in an old file, it will be moved to the next old audit file.

- 3. Select “Yes” to reset the file.**

The Audit Data file clears existing entries and then continues to record new entries based on your auditing configuration.

The following fields on the “Audit Status” screen are reset to zero:

- ◆ Audit file size
- ◆ Audit record count

Remounting a Volume Dismounted by Auditing

If a volume is dismounted because the Audit Data file reached its maximum size limit and the “Dismount Volume” option is set, the network supervisor must remount the volume and enable auditing again.

Prerequisites



- Access to the server console, or a remote console session with the server whose volume was dismounted.
- The auditing password to the volume.

Procedure



1. At the server console, mount the volume by typing

```
MOUNT volume_name <Enter>
```

To continue, you must enter the auditing password for that volume.

2. At the console, enter the current auditing password for the volume.

This is the password the auditor assigned to the volume.

3. Reenable auditing.

Follow the instructions in “Enabling Auditing for a Volume” on page 636.

The auditor then needs to log in and change the password again by following the instructions in “Setting Passwords” on page 640.


Auditing Examples

Examples in this section are taken from the following scenario. You may need to refer to the illustrations and explanations as you work through the examples.

The examples assume that you have logged in to a volume as the auditor and are at the “Available Audit Options” menu.

Suppose you are auditing a company called XYZ Sales. The records you will audit are on volume SYS:, which is set up as shown in Figure 8-6:

Figure 8-6
File System for XYZ
Sales (Example)



You have four users: Thomas, who is the network supervisor, Ann, Fred, and Sue.

SYS:RECORDS is used by the network supervisor to keep confidential employee records.


SYS:ACCOUNTS has subdirectories containing information about each company that does business with XYZ Sales. Ann, Fred, and Sue maintain these files.

Examples of Selecting Items to Audit

Auditing Events by User and File

Suppose you have a directory with sensitive and critical information in it. Even though you have given Write and Erase rights to Ann, Fred, and Sue, you want to audit all their file open and file write requests in the SYS:ACCOUNTS directory.

Procedure



1. **Select “Auditing Configuration” from the “Available Audit Options” menu.**
2. **Select “Audit by Event.”**
3. **Select the file events you want to audit.**

Since you want to audit file events in a particular directory by certain users, select “File open—user and file” and “File write—user and file” options.

For this example, you would select the events shown in Figure 8-7:

Figure 8-7
Auditing File Events (Example)

Auditing Configuration	
Audit By Event	
Audit By File/Directory	
Audit By U	
Audit Opti	
Change Aud	
Disable Uo	
Display Au	

Audit By Event	
Audit By File Events	
Audit By QMS Events	
Audit By Server Events	
Audit By User Events	

Audit by file events	
File delete - user and file	off
File delete - user or file	off
File open - global	off
File open - user and file	on
File open - user or file	off
File read - user and file	off
File read - user or file	off
File rename/move - global	off
File rename/move - user and file	off
File rename/move - user or file	off
File salvage	off
File write - user and file	on
File write - user or file	off
Modify directory entry - global	off
Modify directory entry - user and file	off
Modify directory entry - user or file	off


User and File means you have limited the auditing of an event to a specific user and file. An auditing record will be entered only if a file in the audited directory is opened or written to by the audited user.

User or File means you want to audit all occurrences of an event in a directory regardless of the user who performed the event, or you want to audit all occurrences of an event by a user regardless of the directory in which the event took place.


Global means that you are auditing for all instances of an event throughout the volume.

Examples of these options and how they affect the type of information you gather are shown in more detail in Figure 8-8.


Figure 8-8
Auditing File/Directory Options (Example)



Records a file write only if ANN writes to a file in Home.




Records all file writes to Home, by any user.



Audits all writes that occur in the volume or container. Includes all files and all users.

-
1. Turn on File Write User and File.
 2. Turn on auditing for Home.
 3. Turn on auditing for ANN.



Records file writes by ANN to any directory.

-
1. Turn on File Write Global.

-
1. Turn on File Write User or File.
 2. Turn on auditing for Home or ANN.

4. Select the subdirectories you want to audit.


For this example, you want to audit the SYS:ACCOUNTS directory. Figure 8-9 shows how to select this directory.



Note

When you audit a directory, existing subdirectories do not inherit the Audited attribute. Subdirectories created after the directory is audited do inherit the attribute.

Figure 8-9
Auditing Directories (Example)



Since these subdirectories were already created, select each of them for auditing.


To audit subdirectories that might be created in the future under SYS:ACCOUNTS, also select the ACCOUNTS directory (parent) as audited. Any subdirectories created after that time would automatically be audited.

5. Select the users you want to audit.

For this example, you want to audit the three users who maintain records about the companies who do business with XYZ Sales.

Figure 8-10 shows how to select these users.

Figure 8-10
Auditing Users (Example)



Auditing All Activities in a Directory

Suppose you want to audit the RECORDS directory and all activities that occur in that directory.

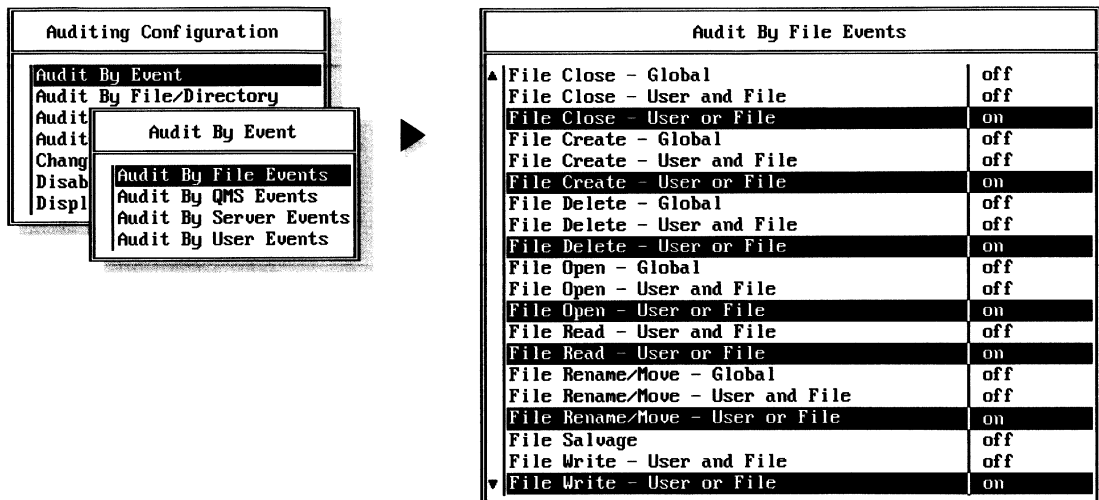
Procedure



1. Select “Auditing Configuration” from the “Available Audit Options” window.
2. Turn on auditing for file/directory events.

Audit all of the events that have a User or File option as shown in Figure 8-11.

Figure 8-11
Auditing All Events in a Directory (Example)



Since the network supervisor (Thomas) is the only user who has rights to the RECORDS directory, you could select all events by user and file, and then audit Thomas.


However, to make sure you get a complete record of all activities in the directory, you can select all events to be audited by user or file. Then records would be entered for the audited events regardless of the user.

3. Turn on auditing for the directory.

On the “Audit by File/Directory” screen, go to the SYS:RECORDS directory (SYS:RECORDS appears at the top of the screen).

In this directory, SYS:RECORDS is the parent, meaning all of the files in the list belong to it. Turn on auditing for SYS:RECORDS by auditing (parent) as shown in Figure 8-12.

Figure 8-12
Auditing a Parent Directory (Example)




Examples of Creating a Report Filter

Sorting Information about Activity in a Subdirectory

Suppose you want to see the activities of the three users who use the Acme Products directory. You could create a report filter to extract this information from the Audit Data file.

Procedure



1. Select “Auditing Reports” from the “Available Audit Options” menu.
2. Select “Edit Report Filters” from the “Auditing Reports” menu.

A list of filters appears. When you create and name new filters, they will be added to this list.

The “_no_filter” item in the list is the one you select when you want to see the entire contents of the Audit Data file.

3. Press <Insert> to create a new filter.

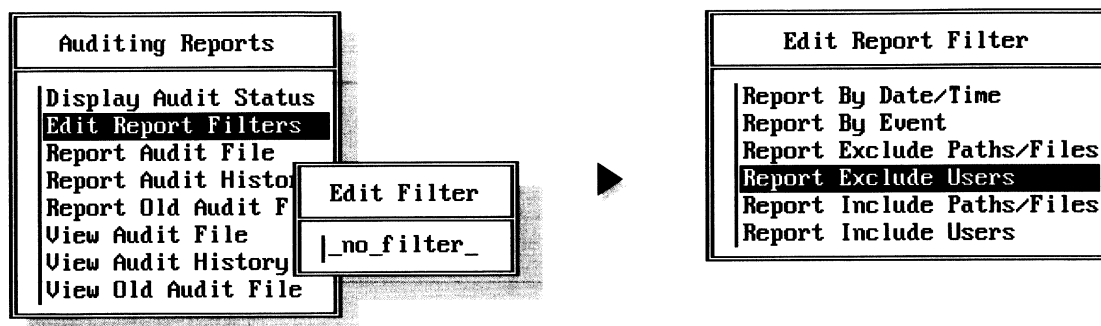
Since the default for filters is to include all audited users, select “Edit Report Filters” and then press <Insert> from the filter list to create a new filter. (The “_no_filter_” option means you want your report to include the entire contents of the Audit Data file.)

4. Select the appropriate users and files.

Since you want to see the activities of the three users, you only need to exclude the network supervisor (Thomas) from the report. Selecting “Report Exclude Users” and pressing <Insert> twice brings up a list of users, where you can select Thomas to be excluded. See Figure 8-13.

Figure 8-13

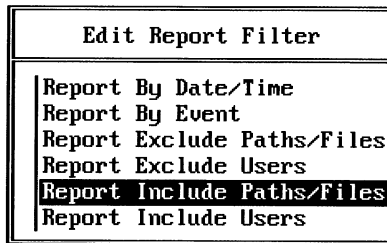
Creating a Report Filter (Example)



Since you want to see the users’ activities in only the Acme Products files, include the ACME subdirectory, which excludes all other directories by default.

Select “Report Include Paths/Files,” and then select ACME from the list that appears. See Figure 8-14.

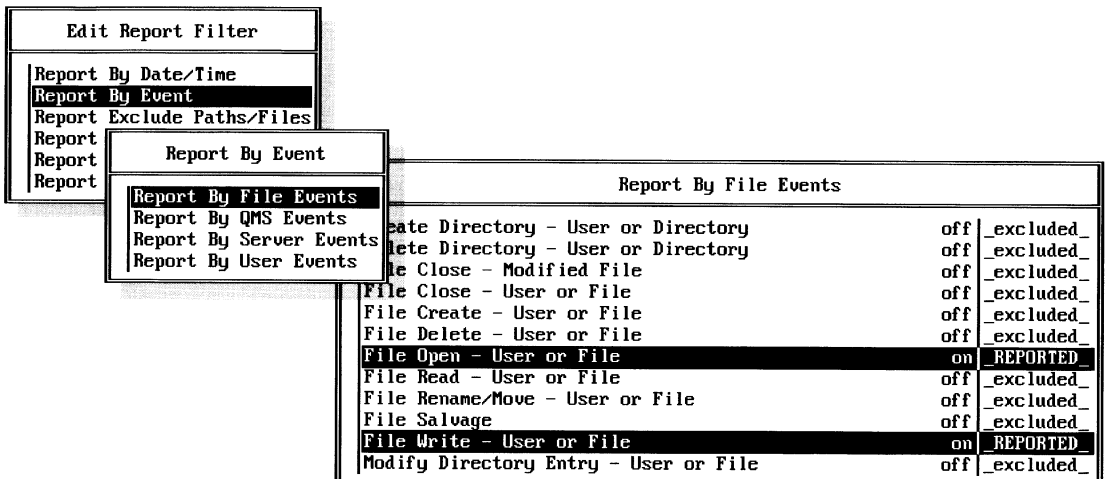
Figure 8-14
Including a
Directory in a
Report (Example)



5. Select the file events

Include the file events you audited in the report filter. See Figure 8-15.

Figure 8-15
Including File Events in a Report (Example)



6. Save the Filter.


Now that you have selected the items you want for the report, press <Esc> until you get to the "Save Filter Changes?" prompt. Save your changes and name the filter by replacing "_no_filter_" with the filter name.

The filter is added to the list for you to use whenever you want to create a report with this information. For this example, we've called the filter "ACMEUSER", as shown in Figure 8-16.

Sorting Information about a User's Activity

If you want to extract information about the network supervisor's audited activities, you can create another filter for this purpose.

Procedure




1. Select "Edit Report Filters" from the "Auditing Reports" menu.
2. Press <Insert> to create a filter.
3. Include the appropriate items.

When you select "Report Include Users" and press <Insert>, a list of users appears. Turn on auditing for user Thomas, the network supervisor, by selecting his name from the list. (Make sure the asterisk is deleted from the beginning of his name.) All of the other users are automatically excluded.

The following figure illustrates this procedure.

Figure 8-16
Including Users in a Report (Example)



4. Press <Esc>.

Continue pressing <Esc> until you get to the "Save Filter Changes?" prompt. Save your changes and name the filter by replacing "_no_filter_" with the filter name.


The new filter is added to your list. For this example, we've called the filter "SUPERVIS" as shown in Figure 8-17.

Example of Using a Filter without Saving the Changes

Using an Existing Filter to Sort Information by Date and Time

Suppose you want to see Thomas's activities for the past two days, but you don't want to permanently alter the filter you created to sort the network supervisor's activity.


Procedure



1. Select "View Audit File" from the "Auditing Reports" menu.
2. Select the filter you want to use and press <F10>.

In this case, select the filter you created to sort network supervisor Thomas's activities in the previous example. (See the example in the following figure.)

Figure 8-17
Using an Existing
Report Filter
(Example)




3. Select **“Report by Date/Time”** and press **<Enter>**; then press **<Insert>**.

Enter the beginning and ending days and times for the period you want to include in the report, as shown in Figure 8-18.

Figure 8-18

Sorting Data by Date and Time (Example)



4. Press **<Enter>** to leave the time settings, and then back out to the **“Save Filter Menu”** by pressing **<Esc>** twice.

The **“Save Filter Menu”** appears. Since you want to use the modified filter for this report only, select the option as shown in Figure 8-19.

Figure 8-19

Using a Filter without Saving Changes (Example)



5. Press **<Enter>**.

Your report appears on the screen, showing only the supervisor’s activity for the days you specified.



chapter

9

Backing Up and Restoring Data

What SBACKUP Does

SBACKUP is a NetWare® Loadable Module™ (NLM) that allows you to back up and restore data for NetWare® 2, 3, and 4 (including NetWare Directory Services™, bindery, and volume restrictions), and OS/2 network servers, as well as workstations or other services you select.

SBACKUP copies data that you select, stores the data on the storage media of your choice (see “Supported Storage Devices and Drivers” on page 685), and retrieves the data later. SBACKUP also allows you to

- ◆ Choose from four different backup types: full, incremental, differential, or custom. (See “Backup Types” on page 699 for definitions.)
- ◆ Perform backups immediately or schedule them for a later time.
- ◆ Back up a server or workstation even when other users are working on it. (However, files in use are not backed up.)
- ◆ Back up and restore multiple name space formats defined on a volume (including DOS, FTAM, Macintosh, NFS, and OS/2 files), retaining file attributes associated with each format.
- ◆ Restore all or part of your data to either its original location or to an alternate location in the directory structure.

Backup and Restore Concepts

Before using SBACKUP, you should be familiar with the following concepts (see *Concepts*):

- “Backup”
- “Data set”
- “Host”
- “Major resource”
- “Media Manager”
- “Minor resource”
- “Modify bit”
- “Restore”
- “Storage Management Services”
- “Target”
- “Target Service Agent”
- “Target Service Agent resources”

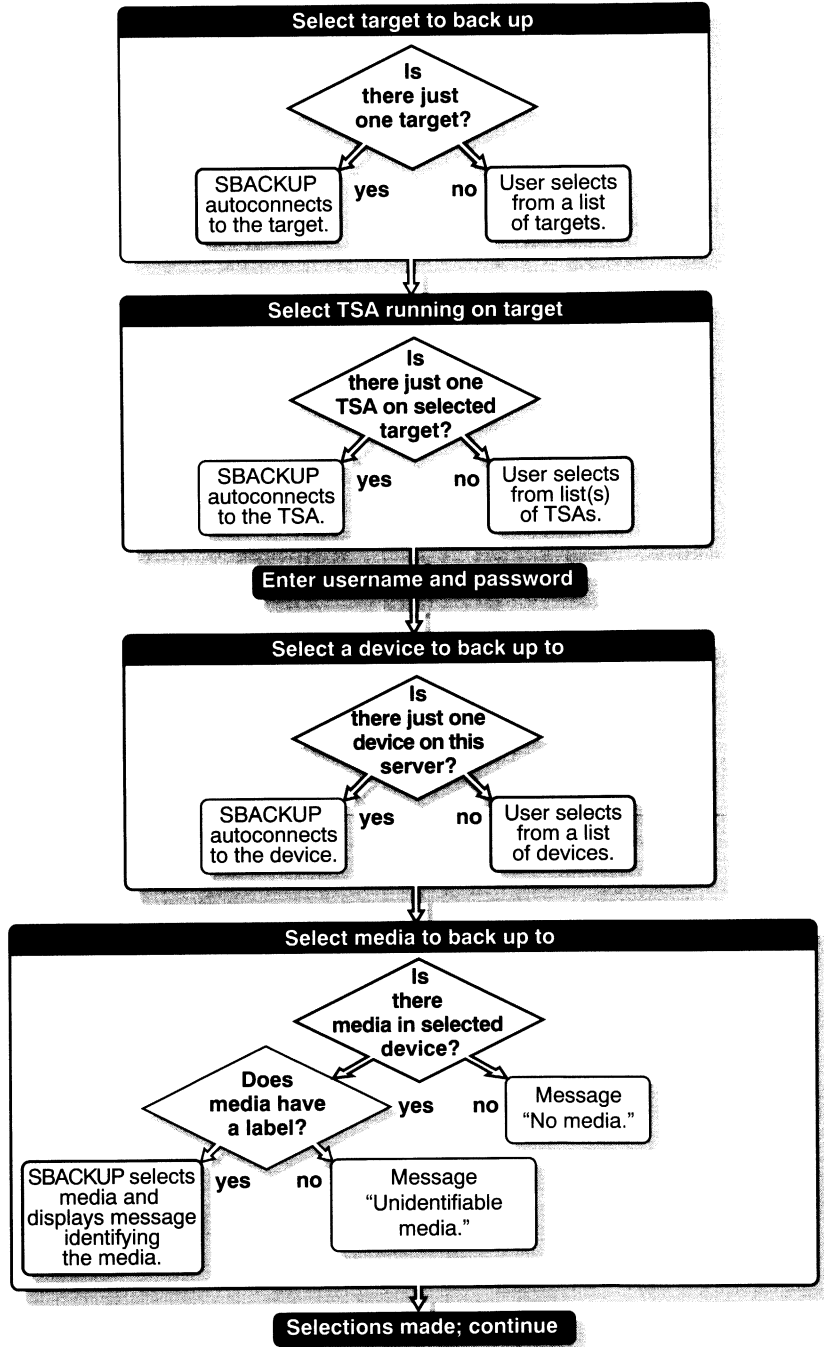
How SBACKUP Presents Choices

Because every network environment is different, SBACKUP offers many different ways of performing a backup or restore. In some cases, SBACKUP streamlines the process for you.

You can have one or more targets, Target Service Agents, storage devices, or storage media. If there is only one of any item, or if a choice you previously made precludes other choices that exist under other circumstances, SBACKUP selects the default and then continues without asking you to make a specific selection.

Figure 9-1 shows an example of this kind of operation.

Figure 9-1
SBACKUP Decision
Process



Rules for Using SBACKUP

The next two sections, “Using SBACKUP” and “Using Target Service Agents,” list rules that will help you avoid problems and confusion when using SBACKUP and its Target Service Agents.

Using SBACKUP

- ◆ **Use the latest version of NLM programs:** Use the NetWare 4.1 version to ensure that SBACKUP loads and works properly.

For example, to back up a NetWare 3.11 target server, use the TSA311.NLM and SMDR.NLM files that shipped with NetWare 4.1.

- ◆ **Log in as ADMIN:** When connecting to a target (server, workstation, or database), log in as ADMIN, or as a network-supervisor equivalent for that target.

For security reasons, many SBACKUP options are limited to the network supervisor. You will have limited success performing backups and restores if you log in without supervisor rights.

- ◆ **Set common context:** Set targets to the same context as the host server to avoid having to specify the context during a backup or restore session.
- ◆ **Reserve disk space for temporary files:** Make sure you have disk space available (1 to 2 MB) on the target server’s volume SYS: to accommodate temporary files.

SBACKUP creates temporary files on the target server during backup. If you have linked UNIX files or files with extended attributes, the temporary files might be larger than a megabyte.

- ◆ **Avoid restoring Netware Directory Services from media when possible:** We recommend using replicas instead. (See Chapter 5, “Managing the NetWare Directory Tree.”)
- ◆ **Don’t mount or dismount volumes during a session:** The data might be corrupted or an abend might occur at the host server.

- ◆ **Use the correct name space and name space formats:** If you don't use the correct name space and name space formats when entering paths and filenames, files can't be restored.

You can only restore to a new location in the name space that is supported by the drive you are restoring to.



On an HPFS drive you can only rename in the OS/2 name space, and on a File Allocation Table (FAT) file system you can only rename in the DOS name space. However, SBACKUP displays both in the log file even though only one is valid.

Name spaces: DOS, FTAM, Macintosh, NFS, and OS/2.

Name space formats:

- ◆ For Macintosh, use *Volume::directory:directory:filename*
- ◆ For all others, use *Volume:/directory/directory/filename*
- ◆ **Exit SBACKUP before unloading drivers:** If you unload a manually loaded driver (such as AHA1740.DSK or ASPITRAN.DSK) before exiting SBACKUP, you might cause the host server toabend.
- ◆ **Use original case for non-DOS names:** Non-DOS pathnames and filenames are case-sensitive. NetWare recognizes DOS pathnames and filenames in uppercase, lowercase, or mixed case. If you're not sure of the original case, refer to your log file.

Using Target Service Agents

This section highlights a few things you should make note of before using Target Service Agents to back up or restore workstations or databases.



When you are backing up or restoring a workstation or database, some of the screens you see will vary from those displayed in this document, which are examples of NetWare 4.1 target server screens; however, the same general principles apply.

- ◆ **Load Target Service Agents automatically:** You can load these files automatically each time a workstation is booted by placing the load commands in the workstation's AUTOEXEC.BAT file (for DOS workstations) or the Startup file (for OS/2 workstations).
- ◆ **Restore NetWare Directory Services (NDS) from replicas:** The only absolute way to ensure that NDS™ can be fully restored is through partition replication, with one master partition and associated replicas on multiple servers.

For more information on partitions and replicas, see Chapter 5, "Managing the NetWare Directory Tree."

- ◆ **Avoid using SBACKUP to restore NDS:** Restoration from SBACKUP should be used *only* in the event of a catastrophic data loss, where all replicas have become corrupted.


If you *must* restore NDS from SBACKUP, load the NDS Target Service Agent (TSANDS) on the same server as SBACKUP (the host server).

Installing and Setting Up SBACKUP

This section explains storage device and driver support as well as memory requirements you should be aware of before loading SBACKUP files on a server or workstation.

Supported Storage Devices and Drivers

The NetWare 4.1 version of SBACKUP supports 0.25-inch, 4mm, and 8mm storage devices. However, we recommend pretesting any tape drive device.



To ensure reliable operations, pretest all media storage devices that are not Novell® certified with the appropriate NetWare device driver and SBACKUP backup/restore utility.

If you are using 4mm tape, use only DDS certified, computer grade tapes.

The following driver files are included with NetWare 4.1. Use the files recommended by your hardware manufacturer.

- ◆ **TAPEDAI.DSK:** This is the Novell generic tape device driver file. It works with most ASPI-compatible SCSI controllers, but contains specific driver information for the tape drives listed in the following table. These drives may or may not be Novell certified, but they are more likely to offer correct functionality than those not listed.

Manufacturer	Product
Archive	Python
Cipher	T826 (0.5 inch, 2.6 GB), T860 (0.5 inch, 6 GB)
Exabyte	EXB-8500 (This is the only Exabyte tape drive that works with TAPEDAI.DSK.)
Hewlett-Packard	HP35450A, HP35470A, HP35480A

Manufacturer	Product
Tandberg	TDC 3800, TDC 4100, TDC 4200
WANGDAT	1300, 3200
WANGTEK	5525ES, 6130, WANGTEK_6130_HS



The generic driver included in TAPEDAI.DSK may not provide the optimal performance from the tape device. Follow the tape driver manufacturer's instructions for specific driver information.

- ◆ **MNS16S.DSK, MNS8MM.DSK, and MNSDAT.DSK:** These files work with Mountain Network Solutions, Inc. SCSI controllers and tape systems, when they are used with SBACKUP.NLM.
- ◆ **PS2SCSI.DSK:** This file works with IBM 2.2GB, 8mm tape devices that use PS2 SCSI controller boards, when they are used with SBACKUP.NLM.
- ◆ **AHA1540.DSK, AHA1640.DSK, AHA1740.DSK, and ASPITRAN.DSK:** These files work with Adaptec 1540, 1640, and 1740 controllers, when they are used with SBACKUP.NLM.

For more information on devices and drivers

- ◆ Contact your Novell Authorized Reseller^{CLM} representative or consultant, or call 1-800-NETWARE for the list of currently supported devices.
- ◆ Download, from NetWireSM forum, a list of currently supported devices. (For more information about NetWire, see "NetWire" in *Concepts* or contact your Novell Authorized Reseller.)
- ◆ Call the Novell LabsTM facility at 1-801-429-5544 for a list of tape drives that are certified as "Yes, Tested and Approved" with NetWare 4.1.

Memory

To run SBACKUP, the host server requires a minimum of 3 MB of memory in addition to the memory required to run NetWare.

If an additional 3 MB of memory is not available, you might be able to set the storage buffers lower than the default and still run SBACKUP.

For memory and buffer settings, see “Enhancing Host Server Performance” on page 783.

Loading Files Needed to Run SBACKUP


Use the appropriate procedure in the following table to load device and controller drivers and to load files on your target. When file loading is completed, the SBACKUP “Main Menu” is displayed on your host server’s console.

To back up	See these procedures
Servers	“Loading Procedures for Drivers” on page 688 and “Loading Procedures for Servers” on page 689
Workstations	“Loading Procedures for Drivers” on page 688 and “Loading Procedures for Workstations” on page 690
NetWare Directory Services	“Loading Procedures for Drivers” on page 688 and “Loading Procedures for NetWare Directory Services” on page 694
SFT III	“Loading Procedures for SFT III Users” on page 695

A summary of all load commands can be found in Table 9-2, but the table should be used only after you are familiar with the procedure for each target.


Loading Procedures for Drivers

Prerequisites



- Ensure that you have installed your controller and storage device according to the hardware manufacturer's instructions before proceeding.

Procedure



1. Load the drivers for controller and storage devices.

To load device driver(s), at the server console prompt, type:

LOAD *controller_device_driver_name(s)* <Enter>

LOAD *storage_device_driver_name* <Enter>

For example, to load AHA1540.DSK and TAPEDAI.DSK, type

LOAD C:AHA1540 <Enter>

LOAD C:TAPEDAI <Enter>

To have these device drivers load automatically, place these commands in the server's STARTUP.NCF file.

2. To register the storage device with the system, type the following at the server console prompt:

SCAN FOR NEW DEVICES <Enter>



If you load TAPEDAI.DSK from within the STARTUP.NCF file, you do not need the command "SCAN FOR NEW DEVICES."


Loading Procedures for Servers

Prerequisites



- Ensure that you have installed your controller and storage device according to the hardware manufacturer's instructions before proceeding.
- Ensure that your host server's controller and storage device drivers have been loaded. (See "Loading Procedures for Drivers" on page 688.)
- Ensure that you have registered the new device with the system by typing "SCAN FOR NEW DEVICES" at the server console prompt.

Procedure



1. Using the following table, load the necessary files for your target server (NetWare 4.1, 4.0, 3.12, or 3.11).

To back up this	At this console	Enter this command
4.1 host server	4.1 host server	LOAD TSA410
	4.1 host server	LOAD SBACKUP <i>[options]</i>
4.1 target server	Target server	LOAD TSA410
	4.1 host server	LOAD SBACKUP <i>[options]</i>
4.0 target server	Target server	LOAD TSA400
	4.1 host server	LOAD SBACKUP <i>[options]</i>
3.12 target server	Target server	LOAD TSA312
	4.1 host server	LOAD SBACKUP <i>[options]</i>
3.11 target server	Target server	LOAD TSA311
	4.1 host server	LOAD SBACKUP <i>[options]</i>

The SBACKUP options are buffer size and number of buffers. To use these options, at the host server console, type

LOAD [path]SBACKUP SIZE=xxx BUFFER=x <Enter>

The x's represent the buffer size and number of buffers.

Option	Supported Values
Size	16, 32, 64, 128, or 256 KB (Default=64 KB)
Buffers	2 to 10 (Default=4)

If you type "LOAD SBACKUP" without specifying options, SBACKUP loads the default buffer size (64 KB) and the default number of buffers (4).

For more on changing buffer sizes and numbers, see "Enhancing SBACKUP Performance" on page 782.

To have the files load automatically, place the commands in the server's STARTUP.NCF file, and in the workstation's NET.CFG file or AUTOEXEC.BAT file (for DOS) or Startup folder (for OS/2).


Loading Procedures for Workstations

Prerequisites



- Ensure that you have installed your controller and storage device according to the hardware manufacturer's instructions before proceeding.
- Ensure your host server's controller and storage device drivers have been loaded. (See "Loading Procedures for Drivers" on page 688.)
- Ensure that you have registered the new device with the system by typing "SCAN FOR NEW DEVICES" at the server console prompt.

Procedure



1. Using the following table and procedures, load the necessary files for your target workstation.

To back up this	At this console	Enter this command
DOS workstation	Host server	LOAD TSADOS
	Target workstation	TSASMS [<i>parameters</i>]
	Host server	LOAD SBACKUP [<i>options</i>]
OS/2 workstation	Host server	LOAD TSAPROXY
	Target workstation	TSAOS2
	Host server	LOAD SBACKUP [<i>options</i>]

1a. At the host server, load either TSADOS or TSAPROXY, depending on your target workstation.

1b. At the target workstation, reject all incoming messages.

Receiving messages break the SPX™ connection and all activity on a workstation stops until you press <Ctrl>+<Enter>.

To reject all messages and avoid breaking the SPX connection, at the target workstation prompt, type

SEND /A=N <Enter>



After completing a backup, a counter-command is necessary to allow the workstation to receive messages. See “Unloading SBACKUP Files” on page 781 for unloading procedures.

1c. At the same workstation, change to the client software directory where TSASMS.COM or TSAOS2 resides.

An example of a NetWare client software directory is “C:\NWCLIENT”.

1d. At the same workstation, enter the load command for your Target Service Agent (either TSASMS [*parameters*] or TSAOS2).

To load TSASMS, use the following syntax and Table 9-1.

```
TSASMS /SE=host_server_name
/P=workstation_password (or /T)
/D=drive_letter /N=workstation_name
```

Table 9-1 explains how to use the TSASMS parameters at DOS workstations.

Table 9-1
TSASMS Parameters

Parameter	Rules for Use	Explanation
/SE[<i>rver</i>]= <i>server_name</i>	Must specify the host server.	Specifies the name of the NetWare server you want the target workstation to connect to for backup and restore services.
/P[<i>assword</i>]= <i>password</i>	Required if /T[<i>rust</i>] parameter is not used. Limited to 10 characters.	Identifies the target workstation's unique password, or is used initially to specify a password.
/T[<i>rust</i>]	Required if /P= <i>password</i> parameter is not used. Use only if bindery services are used on host server. User must be ADMIN or a supervisor equivalent.	Allows a supervisor equivalent to back up and restore without the workstation password if bindery services are used on the server (if "Set bindery services" is in the AUTOEXEC.NCF or on the server console).
/D[<i>rive</i>]= <i>x/xx</i>	Must specify at least one (<i>x</i> =C, D or E). If two or more, do not space between letters (example: /D=CD).	Specifies the hard drive(s) being backed up or restored. TSASMS will not back up floppy drives or network drives.
/B[<i>uffers</i>]= <i>n</i>	Not required, but can be changed (<i>n</i> =1 through 30).	Increases the number of buffers, which increases throughput speed but requires more RAM on the workstation. The default is one buffer.
/N[<i>ame</i>]= <i>workstation_name</i>	Required the first time TSASMS is loaded for a particular server. Limited to 10 characters.	Identifies the target workstation's unique name, or is used initially to name a workstation.

Table 9-1 *continued*
TSASMS Parameters

Parameter	Rules for Use	Explanation
/ST[ack]= <i>nnn</i>	Not required. Use only if RAM is extremely limited or you receive "Stack Overflow" messages. Replace <i>nnn</i> with a number (512 through 4096).	Specifies stack size. The default is 2048 bytes.
/U[nload]	Can be used only at the workstation command line.	Unloads TSASMS from the workstation's memory.
/H[elp]	Can be used only at the workstation command line.	Displays all parameters.

For example, to load TSASMS where the following is true for the workstation being backed up,

- (1) CHICAGO is the host server name
 - (2) The /T[rust] parameter is used
 - (3) Drive C is specified as the target drive
 - (4) BASEBALL is the workstation name
- at the workstation command line, type

TSASMS /SE=CHICAGO /T /D=C /N=BASEBALL <Enter>

If loading multiple terminate-and-stay resident (TSR) programs on a DOS workstation, load TSASMS last.

To avoid loading the workstation Target Service Agent with each session, place it in the workstation's AUTOEXEC.BAT file or NET.CFG file (for DOS), or Startup folder (for OS/2).

If you decide to use NET.CFG, be aware that

- (1) You must be using IPXODI and LSL.
- (2) Command line parameters override the NET.CFG.
- (3) /H and /U are only available at the command line.

For more information on the NET.CFG file, see *NetWare Client for DOS and MS Windows User Guide*.

- 1e. **At the host server, load SBACKUP, using the following syntax:**

LOAD SBACKUP [*options*]

If you type "LOAD SBACKUP" without specifying options, SBACKUP will assume you want the default buffer size (64 KB) and the default number of buffers (4).

To avoid loading the commands with each session, place them in either the server's AUTOEXEC.NCF or STARTUP.NCF file.

For information on SBACKUP options, see "Loading Procedures for Servers" on page 689.

Loading Procedures for NetWare Directory Services

Before using these procedures to back up your NetWare Directory Services (NDS) database, you should understand the importance of having both replicas and media backup sessions available.

Restore NDS from Replicas

The only absolute way to ensure that the NDS database can be fully restored is through partition replication, with one master partition and associated replicas on multiple servers.

For more information on partitions and replicas, see Chapter 5, "Managing the NetWare Directory Tree."

Restore NDS from Media Backup Session

Restore an NDS database by using SBACKUP *only* in the event of a catastrophic data loss, where all replicas have become corrupted.


If you must restore NDS from SBACKUP, load the NDS Target Service Agent (TSANDS) on the same server as SBACKUP (the host server).

Prerequisites



- Ensure that you have installed your controller and storage device according to the hardware manufacturer's instructions before proceeding.
- Ensure that your host server's controller and storage device drivers have been loaded. (See "Loading Procedures for Drivers" on page 688.)
- Ensure that you have registered the new device with the system by typing "SCAN FOR NEW DEVICES" at the server console prompt.

Procedure



1. **At the host server console prompt, enter the load commands for an NDS Target Service Agent:**


LOAD TSANDS <Enter>

LOAD SBACKUP [*options*] <Enter>

Loading Procedures for SFT III Users

The NetWare 4.1 SFT III™ operating system handles device input/output and file system operations in separate engines, the IOEngine and the MEngine. Enter load commands in the engine indicated.

Procedure



1. **Change to the IOEngine prompt of the SFT III server where the storage device is installed.**
2. **Load the device drivers for the controller and storage device. (See "Loading Procedures for Drivers" on page 688.)**

For example, type

LOAD AHA1740 <Enter>

LOAD TAPEDAI <Enter>

To have the drivers load automatically, put the LOAD commands in the IOSTART.NCF file of the server with the storage device.

3. Change to the MEngine prompt.

4. Enter the load commands for your target. (See “LOAD Commands Summary” on page 696.)

For example, type

SCAN FOR NEW DEVICES <Enter>

LOAD TSA410 <Enter>

LOAD SBACKUP <Enter>

To have the modules load automatically, put the commands in the MSAUTO.NCF file.

LOAD Commands Summary

Table 9-2 is a summary of the load commands for backup or restore targets. Use this table for reference only after you are familiar with the target’s loading procedures.

Type the commands at the specified console in the order given.

Table 9-2

LOAD Commands for Servers, Workstations, or NDS Databases

To back up this	At this console	Enter this command
4.1 host server	4.1 host server	LOAD TSA410
	4.1 host server	LOAD SBACKUP <i>[options]</i>
4.1 target server	Target server	LOAD TSA410
	4.1 host server	LOAD SBACKUP <i>[options]</i>
4.0 target server	Target server	LOAD TSA400
	4.1 host server	LOAD SBACKUP <i>[options]</i>
3.12 target server	Target server	LOAD TSA312
	4.1 host server	LOAD SBACKUP <i>[options]</i>
3.11 target server	Target server	LOAD TSA311
	4.1 host server	LOAD SBACKUP <i>[options]</i>

Table 9-2 *continued*

LOAD Commands for Servers, Workstations, or NDS Databases

To back up this	At this console	Enter this command
DOS workstation	Host server	LOAD TSADOS
	Target workstation	TSASMS [<i>parameters</i>]
	Host server	LOAD SBACKUP [<i>options</i>]
OS/2 workstation	Host server	LOAD TSAPROXY
	Target workstation	TSAOS2
	Host server	LOAD SBACKUP [<i>options</i>]
NDS database	Host server	LOAD TSANDS
	Host server	LOAD SBACKUP [<i>options</i>]

For information on SBACKUP options (changing buffer size and number of buffers), see “Changing the Size and Number of Buffers” on page 783.

For information on TSASMS parameters, see Table 9-1 on page 692.



If loading multiple TSRs on a DOS workstation, load the DOS workstation Target Service Agent (TSASMS) last. For example, load IPXODI before TSASMS.



For SFT III 3.11, TSA311 loads at the MEngine prompt, and SBACKUP loads in the IOEngine associated with the tape device.

Backing Up Data

Introduction to Backing Up Data

This section provides an introduction to

- ◆ Backup files
- ◆ Backup types
- ◆ Backup options
- ◆ Backup logbook
- ◆ Backup schedules
- ◆ Effect of backup type and schedule on the backup and restore process

Backup Files

Each backup session produces three types of files:

- ◆ **Data files:** Copied to the selected storage media.
- ◆ **Log file:** Placed in a directory on the host server and accessed through the SBACKUP “Main Menu” screen.
- ◆ **Error file:** Placed in a directory on the host server and accessed through the SBACKUP “Main Menu” screen.

The log and error files are labeled with the same description you give the session and are automatically updated with every session.

Both log and error files contain pertinent information, such as the session’s date, time, and media identification, but the error file also contains a list of any errors that occurred during the backup session, such as files that were not backed up.

Backup Types

SBACKUP has four types of backup sessions, which affect each file's modify bit in different ways.

When the modify bit (or Archive Needed attribute) is set, it indicates that the file or directory needs archiving. When the modify bit is cleared, the file will be passed over during certain types of backups.

- ◆ **Full backup:** Backs up all data and clears the modify bit.
- ◆ **Differential backup:** Backs up all data modified since the last full backup and does not clear the modify bit.
- ◆ **Incremental backup:** Backs up all data modified since the last full or incremental backup (whichever was last) and clears the modify bit.
- ◆ **Custom backup:** Backs up only the data you specify and allows you to specify whether to clear the modify bit. Use this option when you want to select portions of the directory structure as the target, using include and exclude options.

Differential and incremental backup are not interchangeable. See "Backup Schedules" on page 707 for examples of use.

Backup Options

For a custom backup, you can specify exactly what to back up by selecting options from the "Custom Backup Options" screen.

A custom backup session allows you to

- ◆ Choose subsets of data to back up
- ◆ Scan what you are backing up

Choosing Subsets of Data to Back Up

You can choose specific subsets of a data set to exclude from or include in the backup session by selecting major resources (such as volumes) or minor resources (such as directories, paths, or files) to exclude from or include in the backup.

Understanding the Exclude and Include Options

Whenever you perform a custom backup, you can use the exclude and include options to select subsets of what you want to back up.

Using Exclude

To back up most of the directory structure while omitting only a small part, use the exclude option to omit the part you don't want to back up. Everything that you don't specifically exclude is included.

After you exclude part of the directory structure (like a volume or directory), you cannot include any subdirectories or files beneath that excluded volume or directory.

Using Include

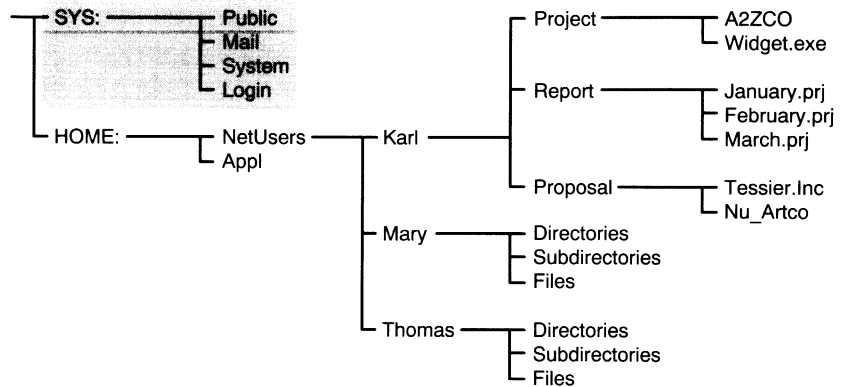
To back up a relatively small part of the directory structure, use the include option to specify the data you want. Everything you don't specifically include is excluded.

When you select only part of the directory structure to include (such as a volume), all directories, subdirectories, and files under that selection are included in the backup by default.

Figure 9-2 illustrates this use of the include option. In the figure, volume SYS: is selected as an include option. All other areas of the directory structure will be excluded from the backup.

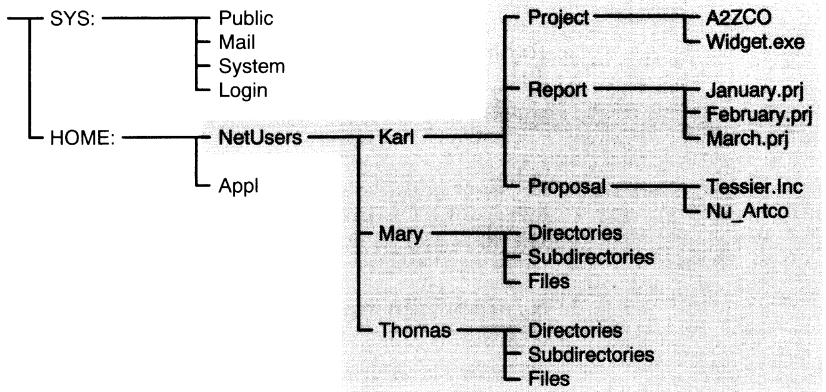
However, you can specifically exclude some subdirectories or files beneath your selection if you need to.

Figure 9-2
SBACKUP Include
Option: Specific
Volume Included, All
Others Excluded



The same principle applies when you specify a directory with the include option. Figure 9-3 shows that all directories, subdirectories, and files under the NetUsers directory will be included in the backup. All other areas of the directory structure are excluded from the backup.

Figure 9-3
SBACKUP Include
Option: Specific
Directory Included,
All Others Excluded



The reverse is true when you select a major Target Service Agent resource, a directory, or a file as an exclude option. All other areas of the directory structure will be included in the backup.

Using Include and Exclude Options Together

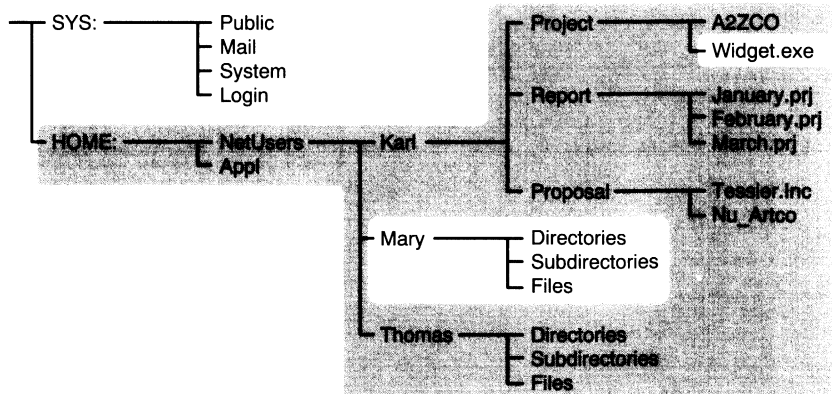
By combining the include and exclude options, you can more precisely control what is backed up.

For example, the following command sequence results in volume HOME: being included in the backup, with the exception of the MARY directory and the WIDGET.EXE file.

Include major TSA resources HOME:
Exclude directories (full path): HOME:NETUSERS/MARY
Exclude path/files
HOME:NETUSERS/KARL/PROJECT/WIDGET.EXE

Figure 9-4 illustrates this example.

Figure 9-4
Combining
SBACKUP Include
and Exclude
Options



Scanning Data Sets

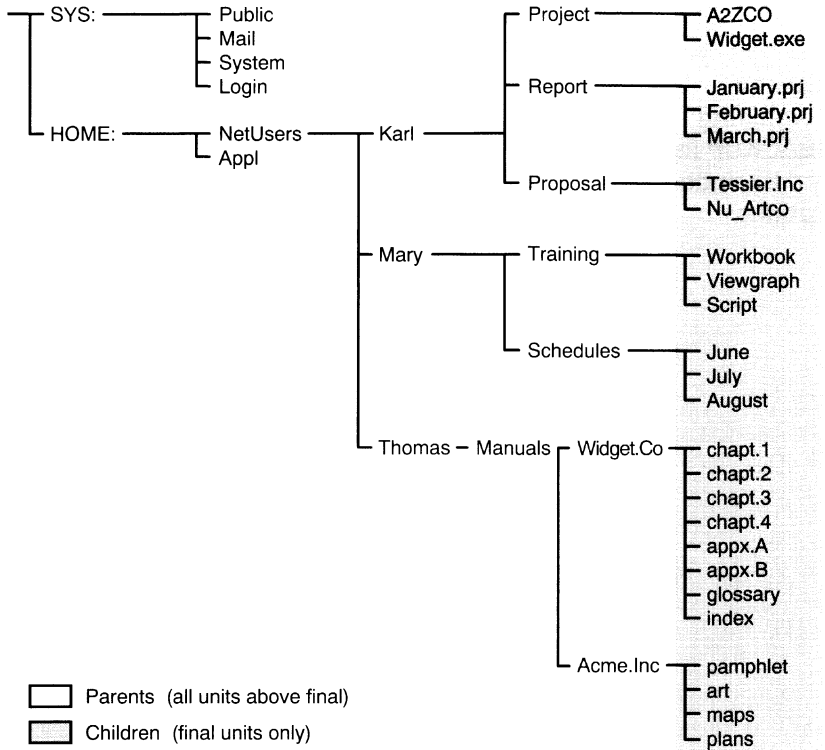
You can specify different types of data sets to be scanned.

A data set is a group of data that can be manipulated by SBACKUP. Each data set in the directory structure can be classified as a parent or a child, and each class includes different types of data items.

Within SBACKUP, a *parent* might be a server, a volume, or a directory. A *child* is a file, which is the lowest level of the directory structure.

As shown in Figure 9-5, the unit below a parent is not necessarily a child; it might be another parent, or the line might end with the parent. The unit above a child must always be a parent.

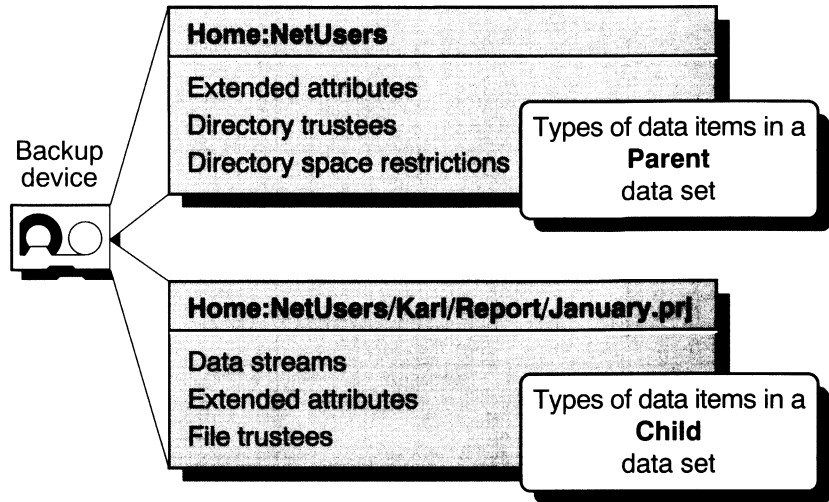
Figure 9-5
Parent and Child
Levels in a Directory
Structure



Items in a data set for either a parent or child should be items that do not frequently change. You may choose to exclude from the backup session one or more of the items in your target's data set.

Figure 9-6 shows typical NetWare data sets.

Figure 9-6
Data Sets
Containing Different
Types of Data Items



Backup Logbook

You should keep a hard-copy log of your backups in case your online log and error files become corrupted. The log should contain the following information:

- ◆ Source backed up (such as the server or workstation name)
- ◆ Full path for session log and error files directory
- ◆ Full path for backed-up data
- ◆ Label of the media on which the backup is stored
- ◆ Description of the session
- ◆ Name or initials of the person performing the backup
- ◆ Date and time of the backup session

Figure 9-7 shows a sample backup log page.

Backup Schedules

Before you begin actual backup procedures, plan a backup schedule based on your needs, considering such factors as the number of users and frequency of changes to files.

You may want to perform different types of backups on different schedules. For example,

- ◆ **Daily:** Perform an incremental or differential backup after the close of business. If revisions are heavy and rapid, consider several backup sessions each day.
- ◆ **Weekly:** Perform an incremental or differential backup after the close of business on the last day of the week for three of the four weeks in the month.
- ◆ **Monthly:** Perform a full backup on the last business day of the month (for example, the first or last Friday).
- ◆ **Major changes:** Perform a full backup before and after you change your configuration, and before and after you upgrade your server to a new version of NetWare.
- ◆ **Application changes:** Perform a custom backup before and after you modify applications.

Effect of Backup Type and Schedule on the Backup and Restore Process

Each type of backup has a different effect on the backup and restore process. In planning your backup schedule, consider all of the variables in Table 9-3 before determining the schedule that is right for you.



Don't interchange differential backups and incremental backups. If you do, the differential backup won't contain all changes since the last full backup. Use full backups interspersed with differential backups, or full backups interspersed with incremental backups.

Table 9-3

Comparative Speed and Media Use of Different Backup Types

Variable	Incremental and Full Backups Interspersed	Differential and Full Backups Interspersed
Media usage for storage	Less media used	More media used
Backup speed	Faster	Slower
Restore speed	Slower	Faster

Media usage and backup speed: Each differential backup uses more media and is slower than an incremental backup because it backs up more files. However, this helps increase the speed of the restore.

Restoring after incremental backups: If you have done full and incremental backups and need to restore data, you must restore the last full backup as well as all subsequent incremental backups.

Restoring after differential backups: If you have done full and differential backups and need to restore data after an unexpected loss, you need to restore only the last full and the last differential backup.

Table 9-4 illustrates what happens in a series of full and incremental backups.

Table 9-4

A Typical Backup Schedule of Full and Incremental Backups

Date (day)	Type	Total Number of Files	New or Revised Files Since Last Full or Incremental Backup	Total Number of Files Backed Up
Sept 4 (Fri)	Full	50	N/A	50
Sept 7 (Mon)	Incremental	60	25	25
Sept 8 (Tue)	Incremental	80	25	25
Sept 9 (Wed)	Incremental	100	30	30
Sept 10 (Thu)	Incremental	125	40	40
Sept 11 (Fri)	Full	150	N/A	150

Incremental backup sessions back up only files that have the modify bit set, which are files that have changed since the last full or incremental backup session (when the modify bit was cleared).

For example, if your backup schedule resembled Table 9-4 and you experienced a loss of data before completing the full backup on Sept 11, you would have to restore the sessions listed in the following table.

Date of Backup Session	Type of Backup	Number of Files Backed Up	Cumulative Number of Files Restored
Sept 4 (Fri)	Full	50	50
Sept 7 (Mon)	Incremental	25	75
Sept 8 (Tue)	Incremental	25	100
Sept 9 (Wed)	Incremental	30	130
Sept 10 (Thu)	Incremental	40	170

Table 9-5 illustrates what happens in a series of full and differential backups.

Table 9-5
A Backup Schedule of Full and Differential Backups

Date (Day)	Type	Total Number of Files	New or Revised Files Since Last Full Backup	Total Number of Files Backed Up
Sept 4 (Fri)	Full	50	N/A	50
Sept 7 (Mon)	Differential	60	25	25
Sept 8 (Tue)	Differential	80	25	50
Sept 9 (Wed)	Differential	100	30	80
Sept 10 (Thu)	Differential	125	40	120
Sept 11 (Fri)	Full	150	N/A	150

When you perform a differential backup, the modify bit isn't cleared after the backup, as it is after a full or incremental backup. So, all files modified since the last full backup are included in the backup (unless they have been deleted).

For example, if your backup schedule resembled Table 9-5 and you experienced a loss of data before completing the full backup on Sept 11, you would have to restore the sessions listed in the following table.

Date of Backup Session	Type of Backup	Number of Files Backed Up	Cumulative Number of Files Restored
Sept 4 (Fri)	Full	50	50
Sept 10 (Thu)	Differential	120	170

Compare this example with the previous example. In both cases the total number of files restored is 170. However, to restore all of the lost data, five backup sessions have to be restored in the first example compared with only two backup sessions in the second.

For more on backup scheduling, see "Backup" in *Concepts*.

Prerequisites to Backing Up



- Understand the terms listed in "Backup and Restore Concepts" on page 680.
- Understand the rules and guidelines for using SBACKUP. (See "Rules for Using SBACKUP" on page 682)
- If you are backing up an OS/2 or DOS workstation, know the workstation's password if the Target Service Agent was originally loaded with the "/Password" parameter instead of the "/Trust" parameter, as explained in Table 9-1 on page 692.
- Make sure the required files for your target are loaded. (See Table 9-2 on page 696.)
- Know what type of backup you want to perform: Full, Differential, Incremental, or Custom. (See "Backup Types" on page 699.)

- Have Read and File Scan rights to the directories/files you are backing up.
- For a custom backup, know the directory structure of the target you are backing up. You will be prompted for the paths to the volumes and directories you want to include or exclude from your backup.

Procedures for Full, Differential, or Incremental Backups

The following steps detail a full, differential, and /or incremental backup from beginning to end.

For a custom backup, see “Procedures for Custom Backups” on page 721.

The figures and choices in this section depict choices available to you when using the NetWare 4.1 File System Target Service Agent (TSA410). If you are backing up a workstation or database, rather than a NetWare 4.1 server, your screen displays will vary slightly.

SBACKUP displays prompts and messages at the bottom of the screen. For instance, you can press <F1> for help at any time.




If you are performing a series of backups and need to change your target, follow the procedures in “Changing Targets” on page 779 before proceeding.

Prerequisites




- Ensure that you have met the prerequisites detailed in “Prerequisites to Backing Up” on page 710.
- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



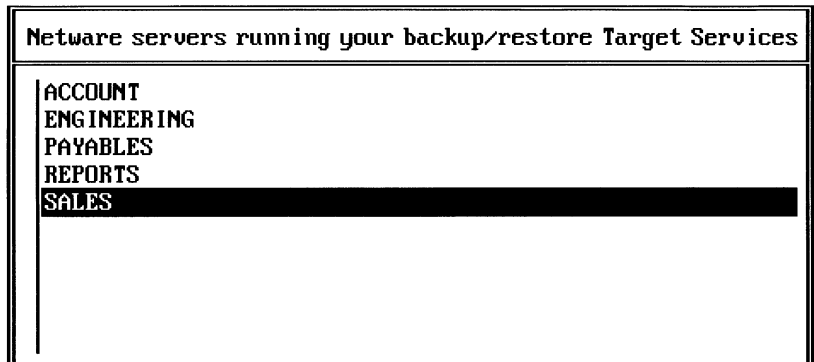
1. From the SBACKUP “Main Menu,” choose “Backup” (Figure 9-8).

Figure 9-8
SBACKUP “Main Menu”



2. Choose a NetWare server running your backup/restore Target Service Agent (similar to Figure 9-9).

Figure 9-9
List of Servers Available for Backup (Typical)



Note

This screen will not appear if you have only one server with a TSA loaded.

If you don't see the target you want on the list, check for these possible causes:

- ◆ The proper Target Service Agent isn't loaded on either the host or the target. (See Table 9-2 on page 696 for LOAD commands.)
- ◆ The host and target are physically far apart, or network traffic is heavy. Press <Esc> to return to the “Main Menu,” and then press <Enter> again. Your target should now be listed.

If	Go To
Your chosen target has more than one Target Service Agent loaded	Step 3
Your chosen target has only one Target Service Agent loaded	Step 4

3. From the Target Services list (similar to Figure 9-10), choose a target.

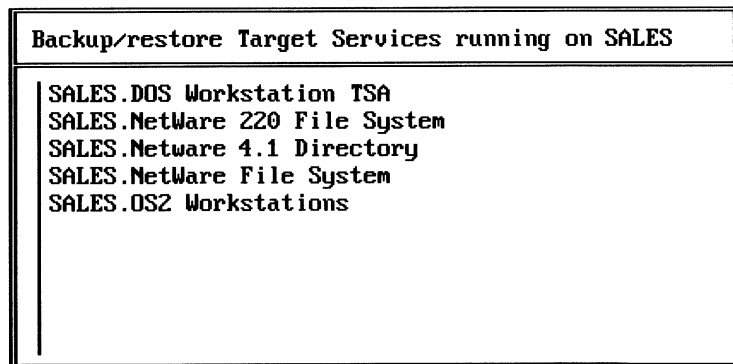
If the target you choose has more than one Target Service Agent loaded, SBACKUP will show you a list of their full names.

For example, Figure 9-10 shows that the NetWare server named SALES is running the following Target Service Agents:

- DOS Workstation (TSADOS)
- NetWare 2.2 File System (TSA220)
- NetWare 4.1 Directory (TSANDS)
- NetWare [4.1] File System (TSA410)
- OS/2 Workstation (TSAPROXY).

You can back up data residing in any of these environments.

Figure 9-10
Display of Multiple
Target Service
Agents on a Target



4. When prompted for the Target username, enter your username (and context if required) as network supervisor for the target.

If SBACKUP rejects the supervisor username you entered, you probably need to include the context of where the user object is located. For example, instead of entering "ADMIN" as the username, enter ".CN=ADMIN.O=*company_name*".

You must include the context in the username at this point if the username you use exists in a directory container that is different from the context set on the server running SBACKUP.

5. If a password is requested, enter the password for the target.


SBACKUP will take a few moments to attach to the target. Wait for the confirmation box, then press any key to continue.

6. Select the device and media for the backup.

- ◆ If your host has only one storage device attached, SBACKUP selects the device for you and notifies you of the device and media that will be used.
- ◆ If your device has more than one storage device attached, choose an available device from those listed.
- ◆ If the device contains media, SBACKUP automatically selects it for the backup.
- ◆ If the storage media does not have a label, SBACKUP displays a message saying that the media cannot be identified. SBACKUP continues, but prompts you later to enter the media label.

7. Specify a location for your session log and error files (see Figure 9-11).

Figure 9-11
The Default Directory for SBACKUP Log Files




You can create different directories to log backup or restore sessions done for different types of targets, such as one directory for workstation backups and another for server backups. Or you might want to keep backup logs for different organizational units in different directories.

To specify a location, use one or more of the following methods.


- ◆ Press <Enter> to accept the default location; then go to Step 8.
- ◆ Press <Insert> to choose from a list of network directories (see Figure 9-12), choose a directory, and press <Esc> to return to the previous window; then go to Step 8.
- ◆ Backspace over the path shown (or a portion of the path), then either type in a new directory or path, or press <Insert> to choose from a list; then go to Step 8.

Figure 9-12
Specifying a Different Location for Log and Error Files



8. From the “Type of Backup” menu (Figure 9-13), choose “Full,” “Differential,” or “Incremental” backup.

Figure 9-13
“Type of Backup”
Menu



If you want to customize your backup with include and exclude features, see “Procedures for Custom Backups” on page 721.

- ◆ A full backup backs up the entire file system, regardless of whether the data has been changed since the last backup, and clears the modify bit after the backup.

If you have installed a migration system built on RTDM, such as Novell’s High Capacity Storage System (HCSS), you will have two full backup options: (1) “Full: All of the Target Including Migrated Data” and (2) “Full: All of the Target Except Migrated Data.” If you choose to include migrated data, ensure your backup media capacity matches the jukebox media capacity.

- ◆ A differential backup backs up only data that has been changed since the last full backup and leaves the modify bit set after the backup.
- ◆ An incremental backup backs up only data that has been changed since the last full or incremental backup (whichever was last) and clears the modify bit after the backup.

For more information on the types of backup, see “Backup Types” on page 699.

9. At the “Backup Options” screen (Figure 9-14), complete the form.

Figure 9-14

The “Backup Options” Screen

Backup Options
Description of what you are backing up: Append to previous sessions on media: No

9a. Type a description for the backup session.

The description should help you identify the specific backup session so that, if a restore is necessary, you can easily identify the session you need.

SBACKUP automatically adds the date, time, number of the media in the media set series, and source (such as the NetWare server or workstation name, but not the volumes, directories, etc.) of the data backed up.



Although no requirements exist for what to enter as the session description, you may want to include the full path of the data, such as SYS:HOME\REPORTS\APRIL.95, which you need to know if a restore session is necessary.

9b. (Optional) If your form has an “Append to Previous Sessions on Media” option, specify whether or not to append.

The option will appear on your “Backup Options” form only if your device supports appending to previous sessions on the media.

Accept the default “No” if you want to label new media, or change the media label and overwrite all sessions on the media. The media rewinds and existing data is overwritten by the next backup.

If you want to append this session, press <Y> and then <Enter>. The backup session is appended to the media at the end of the previous session. Each appended session has separate backup and error logs.



Warning

If you are appending a backup session to a media set (two or more tapes), insert the last tape in the set. If you try to append to a tape other than the last one, you may not be able to restore the data on the later tapes in the set.



Note


Labeled media is not considered empty media. There need not be any sessions on the media for you to be able to set the "Append" option to "Yes."

Delete the log and error files for any session you overwrite to avoid trying to restore a session that no longer exists.

10. Press <F10> to save your options and continue with the backup.

The "Proceed With The Backup?" box is displayed (Figure 9-15).

**Figure 9-15
Specifying When to
Back Up**



11. Choose when you want the backup to start, and then go to the step indicated below.

If you chose	Go to
"No, Return to the Options Form."	Step 9, or continue to exit by pressing <Esc>.
"Start the Backup Now."	Step 14
"Start the Backup Later."	Step 12

12. When the dialog box shown in Figure 9-16 is displayed, set the date and time for the backup to begin.

Figure 9-16
Specifying the Time
for a Delayed
Backup

Start the Backup at:
Start date: 07-15-93
Start time: 11:59:59pm
System time: 07-15-93 10:24:47am

You can choose the default date and time shown for the backup to begin, or you can type a new date and time. The default is midnight of the current day.

Press <F1> for help with the format for entering a date and time for the delayed backup.


13. Press <F10>, and then press <Enter> to save your changes.



To help safeguard your network, SBACKUP exits when a delayed backup is complete. If the delayed backup session cannot fit on the media, SBACKUP prompts you to insert additional media.

If additional media is not inserted, the backup does not finish and SBACKUP does not exit, thus compromising security.

14. (Optional) If you set the “Append to Previous Sessions on Media” option to “No” in Step 9b, after the media is mounted enter the media label.



If you set this option to “No,” the data in the current session and the new label you are prompted to enter will overwrite the existing data and label. All previous sessions contained on the media will be lost.



Note the media label on the media cartridge so that you can find it easily when SBACKUP asks for it during a restore session.

If the media does not have enough empty space to hold the entire backup session, SBACKUP prompts you to insert additional media when the current media is full.

SBACKUP breaks the session after completing a file and “spans” the next file to the next media. SBACKUP then copies the media label from the previous one and increments the media number.

15. (Optional) To abort a backup any time during the session, press <Esc> and choose “Yes” from the “Abort Menu.”

SBACKUP finishes backing up the current data set and then displays a “Dismounting Media” message. Dismounting will take a few minutes.

When the media has been dismounted, the “Backup Terminated” message is displayed. Press <Enter> to return to the “Type of Backup” menu.



When a backup session is aborted, SBACKUP finishes backing up any files in progress at the time of the abort.

16. Verify that the backup was successful.

When the backup has finished, your screen display should look similar to Figure 9-17.

If an error message is displayed, check the error file for more specific information about the backup session.



If you have performed a delayed backup, you will not see this screen. For security reasons, SBACKUP exits to the “Main Menu” after a delayed backup.

Figure 9-17

The Backup Process Completed (Typical Display)

```
SYS:ETC/I01/
SYS:ETC/I02/
SYS:ETC/SAMPLES/
  ATYPES.CFG
  BUILTINS.CFG
  GATEWAYS
  HOSTS
  NETPROTO.CFG
  NETWORKS
  PROTOCOL
  SERVICES
  SNMP.CFG
Parent name: SAMPLES
Child name: TRAPTARG.CFG
Parents: 30
Children: 1306
Elapsed time: 00:04:18
Total written: 40594 KB
Media ID: 1
```

The backup process was completed normally.
Press <Enter> to continue.

17. To return to the “Main Menu,” press <Enter> and then <Esc>.

The <Esc> key allows you to exit SBACKUP one screen at a time until you reach the “Main Menu.” If you want to exit SBACKUP, press <Esc> once again and answer the confirmation prompt.

Procedures for Custom Backups

The following steps detail a single, custom backup from beginning to end.

A custom backup allows you to specify exactly what you want to back up, such as a server, volume, directory database, workstation, directory, or file. You can even specify subsets of groups.

The figures and choices in this section depict choices available to you when using the NetWare 4.1 File System Target Service Agent (TSA410). If you are backing up a workstation or database, rather than a NetWare 4.1 server, your screen displays will vary slightly.

SBACKUP displays prompts and messages at the bottom of the screen. For instance, you can press <F1> for help at any time.



Note

If you are performing a series of backups or restores and need to change your target, follow the procedures in “Changing Targets” on page 779 before proceeding.


Prerequisites



Checklist

- Ensure that you have met the prerequisites detailed in “Prerequisites to Backing Up” on page 710.
- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.


Procedure



Procedure

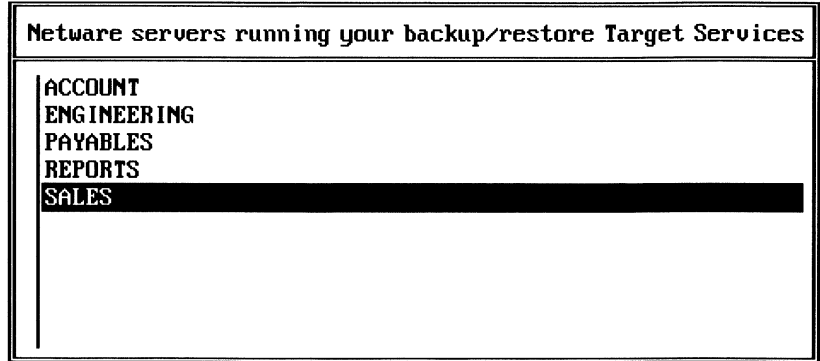
1. From the SBACKUP “Main Menu,” choose “Backup” (Figure 9-18).

Figure 9-18
SBACKUP “Main
Menu”



2. Choose a NetWare server running your backup/restore Target Service Agent (similar to Figure 9-19).

**Figure 9-19
List of Servers
Available for
Backup**



If you don't see the target you want on the list, check for these possible causes:

- ◆ The proper Target Service Agent isn't loaded on either the host or the target. (See Table 9-2 on page 696 for LOAD commands.)
- ◆ The host and target are physically far apart, or network traffic is heavy. In these circumstances, it might take SBACKUP a few moments to detect all of the targets. Press <Esc> to return to the "Main Menu," and then press <Enter> again. Your target should now be listed.

If	Go to
Your chosen target has more than one Target Service Agent loaded.	Step 3
Your chosen target has only one Target Service Agent loaded.	Step 4

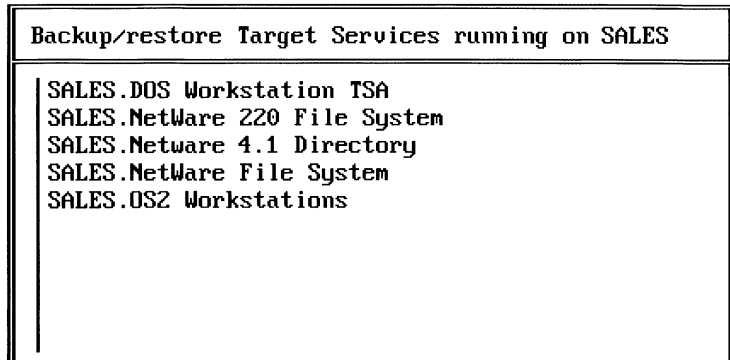
3. From the Target Services list (similar to Figure 9-20) choose a target.

If the target you chose has more than one Target Service Agent loaded, SBACKUP will show you a list of their full names.

Figure 9-20 shows that the NetWare server named SALES is running the following Target Service Agents:

DOS Workstation (TSADOS)
NetWare 2.2 File System (TSA220)
NetWare 4.1 Directory (TSANDS)
NetWare [4.1] File System (TSA410)
OS/2 Workstation (TSAPROXY).

Figure 9-20
Display of Multiple
Target Service
Agents on a Target



4. When prompted for the Target username, enter your username as network supervisor (and context if required) for the target.

If SBACKUP rejects the supervisor username you entered, you probably need to include the context of where the user object is located. For example, instead of entering "ADMIN" as the username, enter ".CN=ADMIN.O=*company_name*".

You must include the context in the username at this point if the username exists in a directory container that is different from the context set on the server running SBACKUP.

5. If a password is requested, enter the password for the target.

SBACKUP will take a few moments to attach to the target. Wait for the confirmation box; then press any key to continue.

6. Select the device and media for the backup.

- ◆ If your host has only one storage device attached, SBACKUP selects the device for you and notifies you of the device and media that will be used. Press any key to continue.
- ◆ If your device has more than one storage device attached, choose an available device from those listed.
- ◆ If the device contains media, SBACKUP automatically selects it for the backup.
- ◆ If the storage media does not have a label, SBACKUP displays a message saying that the media cannot be identified. SBACKUP continues but prompts you later to enter the media label.

7. Specify a location for your session log and error files (see Figure 9-21).

Figure 9-21

The Default Directory for SBACKUP Log Files

Set Location of Log and Error Files
SYS:SYSTEM/TSA/LOG




You can create different directories to log backup or restore sessions done for different types of targets (such as one directory for workstation backups and another for server backups), for different media series, or for different organizational units.

To specify a location, use one or more of the following methods.

- ◆ Press <Enter> to accept the default location; then go to Step 8.
- ◆ Press <Insert> to choose from a list of network directories (see Figure 9-22), choose a directory, and press <Esc> to return to the previous window; then go to Step 8.
- ◆ Backspace over the path shown (or a portion of the path), then type in a new directory or path; then go to Step 8.

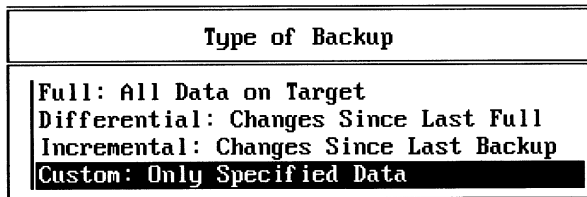
You can also press <Insert> during this process to choose existing parts of the path from a list.

Figure 9-22
Choosing a Different Location for Log and Error Files



8. From the "Type of Backup" menu, choose "Custom Backup" (Figure 9-23).

Figure 9-23
"Type of Backup"
Menu



9. At the “Custom Backup Options” screen (similar to Figure 9-24), complete the form.

Figure 9-24

“Custom Backup Options” Screen

Custom Backup Options
What do you want to back up? Press <Enter> to view the selection. Subsets of what you are backing up: Press <Enter> to view the list. How to scan what you are backing up: Press <Enter> to view the list. Description of what you are backing up: Clear the Modify bit on source files? No Append to previous sessions on media: No



At least two levels of help screens are available to you in the “Custom Backup Options” screen. To see help for a specific line item, select the item and press <Enter> to get in editing mode, and then press <F1>.



9a. Specify what you want to back up

Do not use the following characters in the “What To Back Up” field: `?`, `*`, `/`, or `C:/` `*`. SBACKUP accepts them but does not back up anything.

You can type the path to your target and press <Enter>, or press <Enter> to see the backup default. If the default is what you want to back up, press <Esc> and go to Step 9b.

If you want to see a selection list, press <Insert>. A screen similar to Figure 9-25 is displayed.

Figure 9-25
“Select” List Screen from “What to Back Up”
Screen

What to Back Up	
Select	
NetWare server SYS:	

When you press <Insert>, the selection list will vary depending on your target. The selections represent the top of each directory structure for either the server, NDS database, or workstation.

The “NetWare Server” option backs up the entire server, except the NDS database. If you want the NDS database backed up, you must change your target to “NetWare 4.1 Directory.”

The “DOS Workstation” or “OS/2 Workstation” will be the default option if either of these is your target. Each one backs up the entire workstation.

You can (1) type a path in the box and press <Enter>; (2) select an item, press <Esc>, and press <Enter>; or (3) continue down the directory structure by pressing <Enter> until you find a particular branch of the directory to select, then press <Esc> and then press <Enter>.

9b. (Optional) Choose subsets of what you are backing up.

At the “Custom Backup Options” screen, choose “Subsets of What You Are Backing Up” to see a list of available subsets of the directory structure (similar to Figure 9-26).

The subset options are related to the Target Service Agent on the target to be backed up.

Figure 9-26
Subsets of Data

Choose subsets of what you are backing up	
Exclude major TSA resources	Default
Include major TSA resources	Default
Exclude directories (full path)	Default
Include directories (full path)	Default
Exclude files	Default
Include files	Default
Exclude path/files	Default
Include path/files	Default

The “Default” setting means that everything in the “What to Back Up” screen (Figure 9-25) is backed up—nothing specific has been included or excluded.

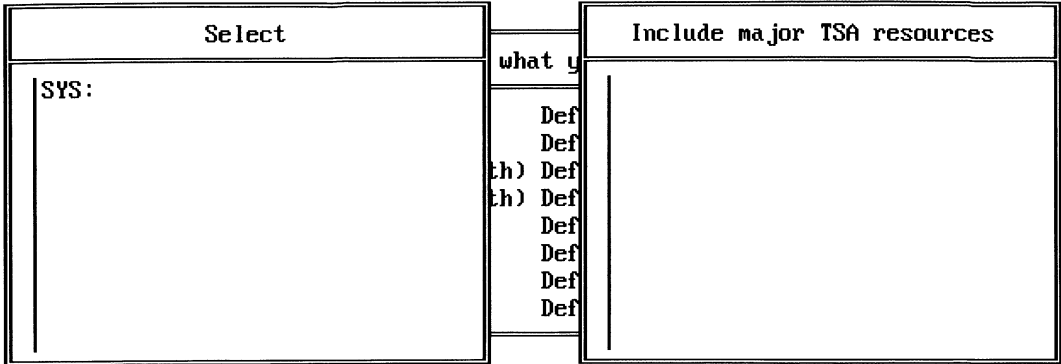
To specify subsets of what you are backing up, choose any subset to get an “Exclude...” or “Include...” screen; then use the options at the bottom of the screen to add the subset to the “Exclude...” or “Include...” screen.

You can add, delete or modify an entry. For example, to exclude a directory in volume SYS: from the backup

- (1) From the “Choose Subsets of What You Are Backing Up” screen, choose “Include Major TSA Resources.”
- (2) Press <Insert>, choose “Volume SYS:,” and press <Esc>. (When you press <Insert>, a screen similar to Figure 9-27 is displayed.)
- (3) Choose “Exclude Directories (Full Path)” and press <Insert>.
- (4) If requested, enter name space type (DOS, FTAM, Macintosh, NFS, or OS/2).
- (5) Type the full path to the directory, including the volume name. For example, a Macintosh path is “Volume::directory:subdirectory:filename.”

Figure 9-27

Select a Major Target Service Agent Using Options Listed at the Bottom of the Screen



The target server's directories and files are not affected by what is specified in the Include or Exclude field.

When you are finished, press <Esc> or <F10> to return to the "Custom Backup Options" screen.

9c. (Optional) Specify how SBACKUP should scan the data to be backed up.

Press <Enter>. The list of data set items available to scan is displayed, similar to Figure 9-28.

Your screen might appear different, depending on which subsets you chose and which Target Service Agent you are using.

The data set contains various types of data, as listed on the screen. You can exclude user account information or types of files and directories (such as hidden files).

The default "No" means that nothing will be excluded. If you want to accept the default, press <Esc>.

If you have installed a migration system built on RTDM, such as Novell's High Capacity Storage System (HCSS), you can choose to exclude the migrated file data in the "How To Scan What You Are Backing Up" field.

Figure 9-28
List of Data Set
Items for Scanning

How to Scan What You Are Backing Up	
Exclude subdirectories:	No ▲
Exclude files that have not changed:	No ■
Exclude hidden files:	No
Exclude hidden directories:	No
Exclude system files:	No
Exclude system directories:	No
Exclude file trustees:	No
Exclude directory trustees:	No
Exclude volume restrictions:	No
Exclude directory space restrictions:	No ▼

If you want to exclude a particular type of data from the backup, move the cursor to the appropriate line, press <Enter> to access an edit mode (blinking cursor), and type "Y." Press <Enter> to save the settings. (The "No" is replaced with "Yes.")

For information about how SBACKUP scans data, see "Data set" in *Concepts*.



If you have compressed files or a compressed volume that you might want to restore to a noncompressed volume, choose the "Backup Compressed Files as Expanded Files" option.

When all data items are set as you want them to be, press <Esc> or <F10>.

9d. Type a description for the backup session.

A session description helps you locate a session when you restore data. This description becomes the header for log and error files.




SBACKUP automatically records the date and time of the session and the source of the data (such as the server or workstation name), as well as the number of this media in the media set series.



Although no mandatory rules exist for what you enter as the session description, you should include the full path of the data (such as SYS:HOME\REPORTS\APRIL.95).

9e. Specify whether to clear the modify bit on the selected data after the backup.



If you clear the modify bit, these files will not be recognized as having been changed at the time of the next backup, unless they are changed again between this backup and the next one.

(For more information, see “Backup Types” on page 699, or see “Modify bit” and “Backup” in *Concepts*.)

The default is “No.” If the entry in this field is “No,” the modify bit is not changed after files are backed up. Files are backed up in the next backup session, even if they were not modified.

If you want to reset the modify bit, type “Y” and press <Enter>. If the entry in this field is “Yes,” the modify bit is cleared after files are backed up.

9f. (Optional) If your form has an “Append to Previous Sessions on Media” option, specify whether or not to append.

The option will appear on your “Custom Backup Options” form only if your device supports appending to previous sessions on the media.

Press <Enter> to accept the default “No” if you want to label new media, or change the media label and overwrite all sessions on the media. The media rewinds and existing data is overwritten by the next backup.

If you want to append this session, type “Y” and press <Enter>. The backup session is appended to the media at the end of the previous session. Each appended session has separate backup and error logs.



If you are appending a backup session to a media set (two or more tapes), insert the last tape in the set. If you try to append to a tape other than the last one, you may not be able to restore the data on the later tapes in the set.



Labeled media is not considered empty media. There need not be any sessions on the media for you to be able to set the “Append” option to “Yes.”

Delete the log and error files for any session you overwrite to avoid trying to restore a session that no longer exists.

10. When you have finished editing the screen, press <F10> to save your options and continue.

The "Proceed With The Backup?" prompt is displayed (Figure 9-29).

Figure 9-29
Specifying When to Back Up

Proceed With the Backup?
No, Return to the Options Form.
Start the Backup Now.
Start the Backup Later.

11. Choose when you want the backup to start, and then go to the step indicated below.

If you chose	Go to
"No, Return to the Options Form"	Step 9, or continue to exit by pressing <Esc>.
"Start the Backup Now"	Step 14
"Start the Backup Later"	Step 12

12. When the "Start the Backup at" dialog box is displayed (Figure 9-30), set the date and time for the backup to begin.


Figure 9-30
Specifying When to Start a Delayed Backup

Start the Backup at:
Start date: 07-15-93
Start time: 11:59:59pm
System time: 07-15-93 10:24:47am

You can choose the default date and time shown for the backup to begin, or you can type a new date and time. The default is midnight of the current day.

Press <F1> for help with the format for entering a date and time for the delayed backup.

13. Press <F10>, and then press <Enter> to save your changes.



To help safeguard your network, SBACKUP exits when a delayed backup is complete. If the delayed backup session cannot fit on the media, SBACKUP prompts you to insert additional media. If additional media is not inserted, the backup does not finish and SBACKUP does not exit, thus compromising security.

14. (Optional) If you set the “Append to Previous Sessions on Media” option to “No” (in Step 9f), after the media is mounted enter the media label.



If you set “Append to Previous Sessions on Media” to “No,” the data in the current session and the new label you are prompted to enter will overwrite the existing data and label. All previous sessions contained on the media will be lost.

If the media does not have enough empty space to hold the entire backup session, SBACKUP will prompt you to insert additional media when the current one is full. SBACKUP will break the session after completing a file and will “span” the next file to the next media. Then SBACKUP will copy the media label from the previous one and increment the media number.

15. (Optional) To abort a backup session any time during the session, press <Esc> and choose “Yes” from the “Abort Menu.”

SBACKUP finishes backing up the current data set and then displays a “Dismounting Media” message. Dismounting will take a few minutes.

When the media is dismounted, the “Backup Terminated” message is displayed. Press <Enter> to return to the “Type of Backup” menu.



When a backup session is aborted, SBACKUP finishes backing up any files in progress at the time of the abort.

16. Verify that the backup was successful.

When the backup has finished, your screen display should look similar to Figure 9-31.

Figure 9-31
The Backup Process Completed (Typical Display)

```
SYS:ETC/I01/  
SYS:ETC/I02/  
SYS:ETC/SAMPLES/  
  ATYPES.CFG  
  BUILTINS.CFG  
  GATEWAYS  
  HOSTS  
  NETPROTO.CFG  
  NETWORKS  
  PROTOCOL  
  SERVICES  
  SNMP.CFG  
  
Parent name: SAMPLES  
Child name: TRAPTARG.CFG  
Parents: 30  
Children: 1306  
Elapsed time: 00:04:18  
Total written: 40594 KB  
Media ID: 1  
  
The backup process was completed normally.  
Press <Enter> to continue.
```

If an error message is displayed, view the error file for more specific information about the backup session.



Note

If you have performed a delayed backup, you will not see this screen. For security reasons, SBACKUP exits to the “Main Menu” after a delayed backup.

17. To return to the “Main Menu,” press <Enter> and then <Esc>.

The <Esc> key allows you to exit SBACKUP one screen at a time until you reach the “Main Menu.” If you want to exit SBACKUP, press <Esc> once again and answer the confirmation prompt.



Note

To unload SBACKUP and applicable Target Service Agents, see “Unloading SBACKUP Files” on page 781.

Restoring Data

Introduction to Restoring Data

This section provides an introduction to

Restore sessions

Restore options

Choosing subsets of data to restore

Scanning data sets

Restoring to a different location

Overwriting a parent or child

Restoring NetWare Directory Services™

How Subsets, Scan, and Overwrite Work Together

Restore Sessions

Use SBACKUP to retrieve and reinstate data you have backed up to a storage media. Typically, you would perform a restore if data has been lost or corrupted since the backup was made.

A restore session restores data from a backup. The restore session produces the requested data, which is retrieved from the storage media and restored to the location you specify. If an error occurs during the restore session, an error message is appended to the error file on the host server, if an error file exists.

The error file is labeled with the same description that you give the backup session (such as "Friday's Full Backup") and is accessed through the "Log/Error File Administration" option of the SBACKUP "Main Menu."

The error file might be located in the default directory or in another directory that you specified during the backup or restore session.

Restore Options

For a custom restore, you can specify exactly which data to restore. Several options work together to allow you maximum flexibility in your restore session. These options allow you to do the following:

- Choose subsets of data to restore.

- Scan what you are restoring.

- Restore to a different location.

- Restore migrated files.

- Overwrite an existing parent (such as a directory) or child (such as a file).

Choosing Subsets of Data to Restore

You can choose specific subsets of a backup session to include in or exclude from the restore session by selecting major resources (such as volumes) or minor resources (such as directories, paths, or files).

Understanding the Exclude and Include Options

Whenever you perform a custom restore, you can use the exclude and include options to select subsets of what you want to restore.

Using Exclude

To restore most of the directory structure while omitting only a small part, use the exclude option to omit the part you don't want to restore. Everything that you don't specifically exclude is included.

After you exclude part of the directory structure (such as a volume or directory), you cannot include any subdirectories or files beneath that excluded volume or directory.

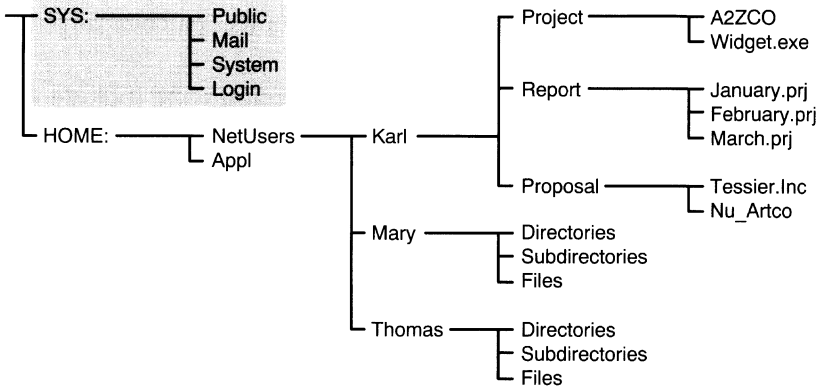
Using Include

To restore a relatively small part of the directory structure, use the include option to specify the data you want. Everything you don't specifically include is excluded.

When you select only part of the directory structure to include (such as a volume), all directories, subdirectories, and files under that selection are included in the restore by default. However, you can specifically exclude some of the subdirectories or files beneath your selection if you need to.

Figure 9-32 illustrates this use of the include option. The figure shows a directory structure with volume SYS: selected as an include option. All other areas of the directory structure will be excluded from the restore.

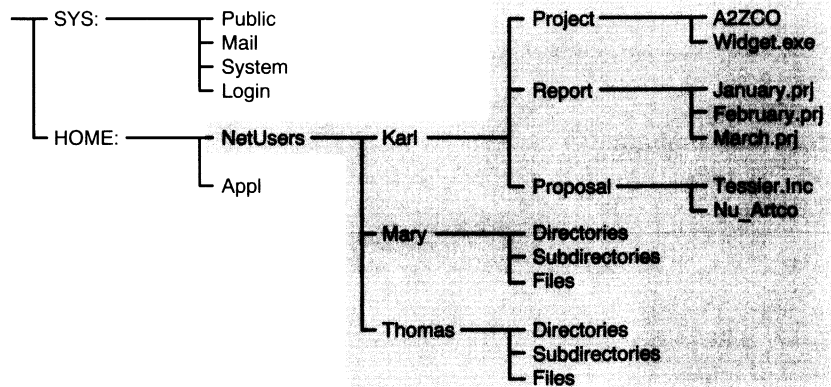
Figure 9-32
SBACKUP Include
Option: Specific
Volume Included, All
Others Excluded



The same principle applies when you specify a directory with the include option. Figure 9-33 shows that all directories, subdirectories, and files under the NetUsers directory will be included in the restore. All other areas of the directory structure are excluded from the restore.

Figure 9-33

**SBACKUP Include
Option: Specific
Directory Included,
All Others Excluded**



The reverse is true when you select a major resource, directory, or file as an exclude option. All other areas of the directory structure will be included in the restore.

Using Include and Exclude Options Together

By combining the include and exclude options, you can more precisely control what is restored.

For example, the following command sequence results in volume HOME: being included in the restore, with the exception of the MARY directory and the WIDGET.EXE file.

Include major TSA resources HOME:

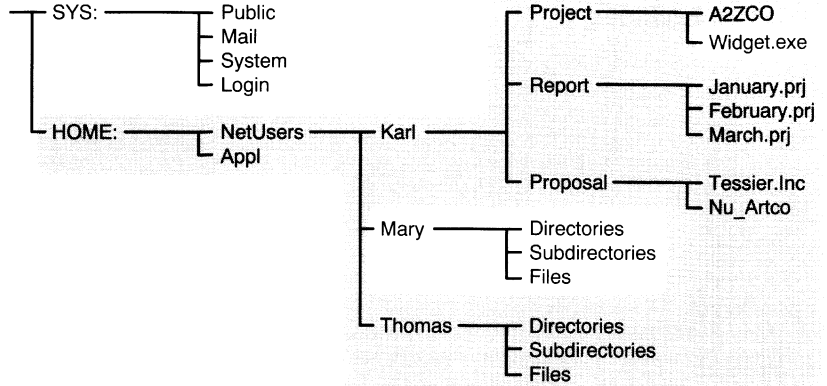
Exclude directories (full path): HOME:NETUSERS/MARY

Exclude path/files

HOME:NETUSERS/KARL/PROJECT/WIDGET . EXE

Figure 9-34 illustrates this example.

Figure 9-34
Combining
SBACKUP Include
and Exclude
Options



Scanning Data Sets

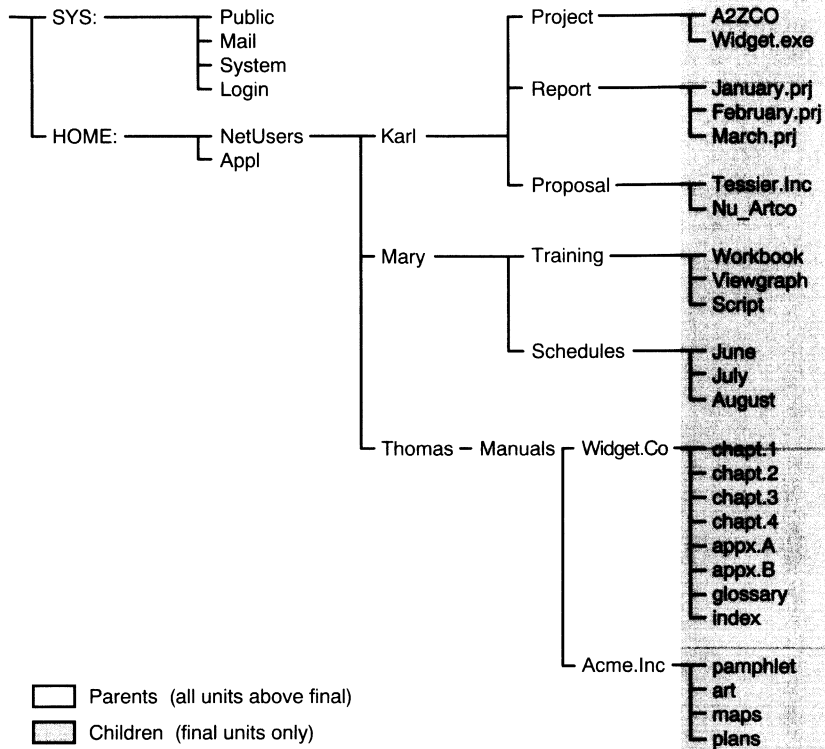
You can specify different types of data sets to be scanned.

A data set is a group of data that can be manipulated by SBACKUP. Each data set in the directory structure can be classified as a parent or a child, and each class includes different types of data items.

Within SBACKUP, a parent might be a server, a volume, or a directory. A child is a file, which is the lowest level of the directory structure.

As shown in Figure 9-35, the unit below a parent is not necessarily a child; it might be another parent, or the line might end with the parent. The unit above a child must always be a parent.

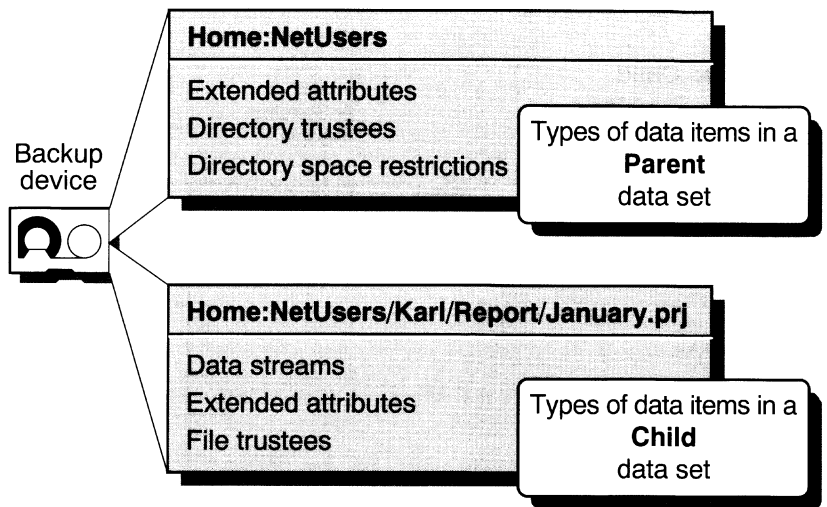
Figure 9-35
Parent and Child
Levels in a Directory
Structure



Items in a data set for either a parent or child should be items that do not frequently change. You may choose to exclude from the restore session one or more of the items in your target's data set.

Data sets can contain different types of data items. Figure 9-36 shows typical NetWare data sets.

Figure 9-36
Data Sets
Containing Different
Types of Data Items



Restoring to a Different Location

You have the option of restoring data to the same location in the directory structure as the backup source, or of restoring data to a new location.

To specify a new location, you must specify the full directory path. If the location specified does not exist, SBACKUP will create the new directory structure.

Restoring Migrated Files

You can restore migrated files if you:

1. Ensure HCSS (or your migration system) is not loaded on the server.
2. Load SBACKUP and restore the migratable directories using SBACKUP while *excluding* the files from the restore (see Figure 9-46 on page 757).
3. Load HCSS (or your migration system) at the server console prompt.
4. Return to SBACKUP and restore the migratable files using SBACKUP while *including* the files in the restore.

You can choose to exclude files from the restore by using an option in the “How to Scan the Session to be Restored” screen (see Figure 9-46 on page 757).

Overwriting a Parent or Child

SBACKUP allows you to overwrite all existing parents or children.

It also allows you to specify that children be overwritten only if the date of the data set on the hard disk is more recent than the date of the data set that was backed up.

Restoring NetWare Directory Services

The only absolute way to ensure that your NetWare Directory Services (NDS) database can be fully restored is through partition replication, with replicas of the entire database on multiple servers. For more information on replicas and partitions, and on creating partitions, see Chapter 5, “Managing the NetWare Directory Tree.”

Use SBACKUP to restore the NDS™ database *only* in the event of a catastrophic data loss where all replicas have become corrupted.

If data has become corrupted, you should

1. Delete the corrupted NDS data
2. Allow time for the deletion to propagate throughout the network
3. Restore the NDS data

If the objects are not deleted before the restoration, the NDS database will be corrupted by the old data.

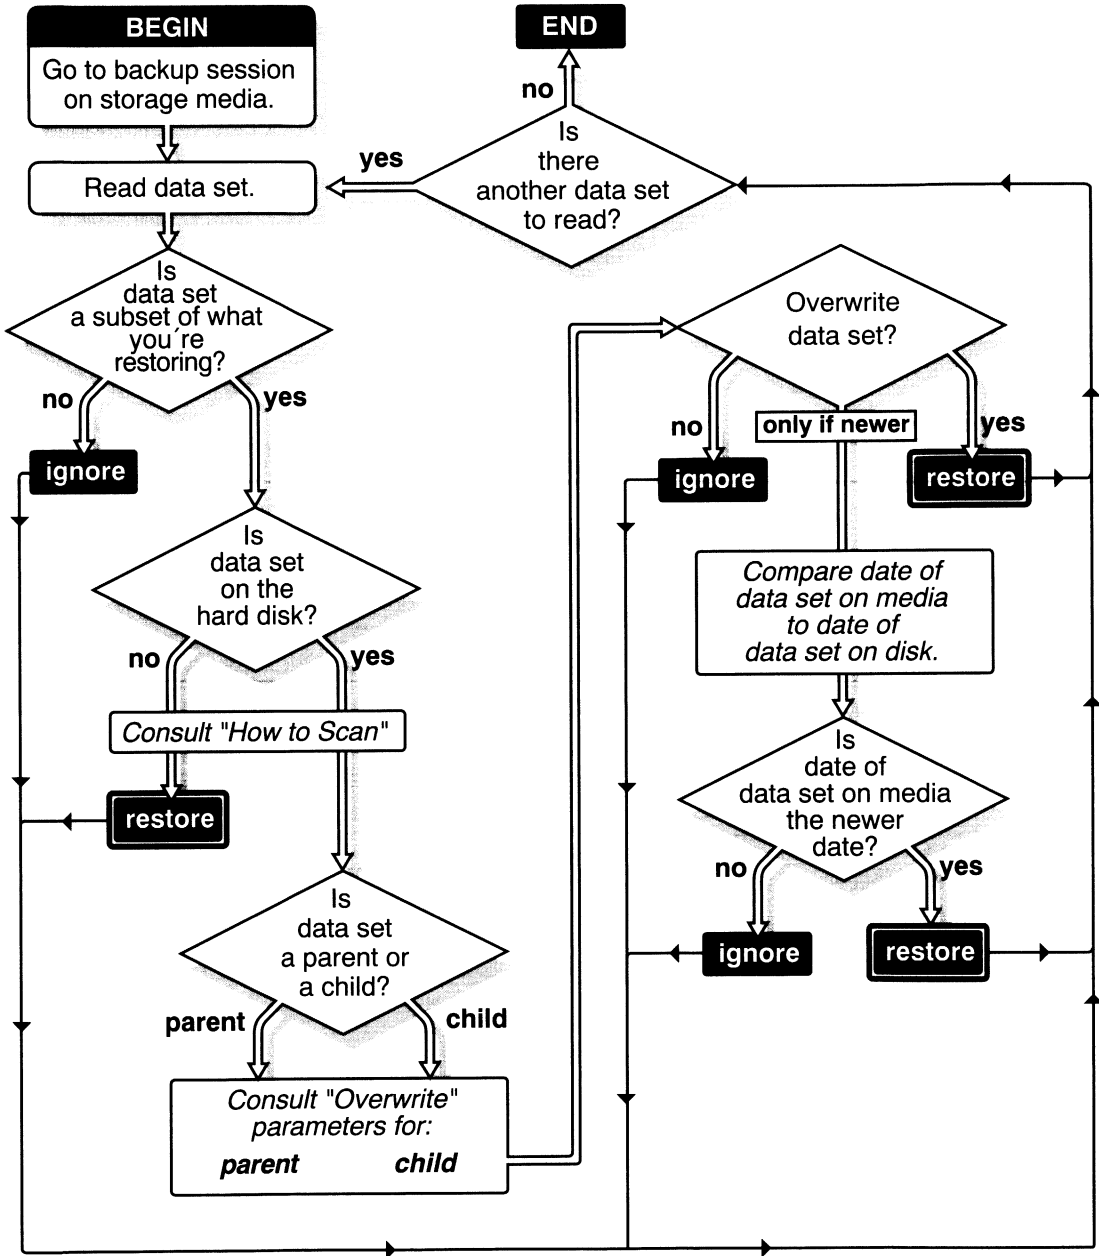
How Subsets, Scan, and Overwrite Work Together

During a restore session, SBACKUP reads the backup storage media and the Target Service Agent compares the media data set to the existing hard disk data set. The Target Service Agent evaluates each data set according to the following criteria:

- ◆ Is this data set a subset of what is being restored?
- ◆ Is this data set found on the hard disk?
- ◆ Which parts of the data set are subject to restoring?
- ◆ Is this data set a parent or a child, and is the Overwrite parameter set to "Yes," "No," or "Overwrite Only if Newer?"
- ◆ If the parameters for a child are set to "Overwrite Only if Newer," does the backup copy have a more recent date than the existing copy?

Figure 9-37 shows how SBACKUP evaluates data sets before performing a restore.

Figure 9-37
**How SBACKUP Evaluates Data
 Being Restored**



Prerequisites to Restoring

Checklist



- Understand the entries listed in “Backup and Restore Concepts” on page 680.
- Understand the SBACKUP rules. (See “Rules for Using SBACKUP” on page 682.)
- Have network supervisor rights (such as user ADMIN) or equivalent rights to the directories/files you are restoring.
- If you are restoring to an OS/2 or DOS workstation, know the workstation’s password if the Target Service Agent was originally loaded with the “/Password” parameter instead of the “/Trust” parameter, as explained in Table 9-1 on page 692.
- Have all required files loaded. (See Table 9-2 on page 696.)
- Know the username and password for the target server or workstation to which you want to restore data.
- Know the session description that you want to restore. You can get this information from your session log and error files. (See “Log and Error Files” on page 762.)
- For a custom restore, know the directory structure of the data you are restoring. You will be prompted for specific paths and filenames while setting the restore options. (For more information, see “Enhancing Host Server Performance” on page 783.)
- If you want to restore data to a new location (different from where the original data was located), you must specify the full path to both the original data and the new location. If the new location does not exist, SBACKUP will create a new directory structure.

Note



You can get the original path from your backup logbook or from the session log files if you noted the path at the time the backup was performed.

Procedures for a Restore

The following steps detail a single restore from beginning to end.

The figures and choices in this section depict choices available to you when using the NetWare 4.1 File System Target Service Agent (TSA410). If you are restoring a workstation or database, rather than a NetWare 4.1 server, your screen displays will vary slightly.

SBACKUP displays prompts and messages at the bottom of the screen. For instance, you can press <F1> for help at any time.

For field-specific help, use the arrow keys to move the cursor to the field you want, press <Enter>, then press <F1>.




If you are performing a series of restores and need to change your target, follow the procedures in “Changing Targets” on page 779 before proceeding.

Prerequisites




- Ensure that you have met the prerequisites detailed in “Prerequisites to Restoring” on page 746.
- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



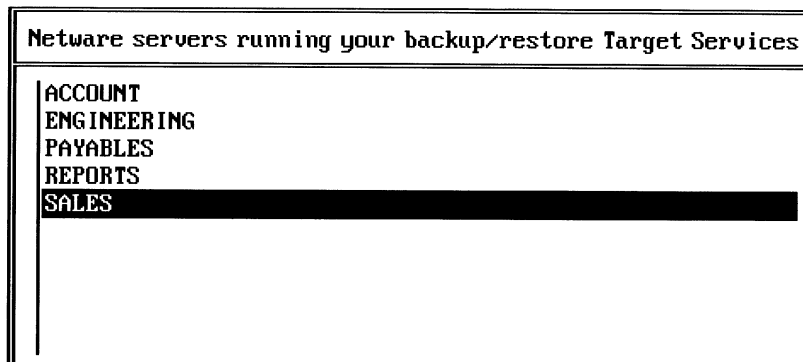
1. From the SBACKUP “Main Menu,” choose “Restore” (Figure 9-38).

Figure 9-38
SBACKUP “Main
Menu”



2. Choose a NetWare server running your backup/restore Target Services (similar to Figure 9-39).

Figure 9-39
List of Servers
Available as Targets
(Running Target
Services)



If you don't see the target you want on the list, check for these possible causes:

- ◆ The proper Target Service Agent isn't loaded on either the host or the target. (See Table 9-2 on page 696 for loading instructions.)
- ◆ The host and target are physically far apart, or network traffic is heavy. In these circumstances, it might take SBACKUP a few moments to detect all of the targets. Press <Esc> to return to the "Main Menu," and then press <Enter> again. Your target should now be listed.

If	Go to
Your chosen target has more than one Target Service Agent loaded	Step 3
Your chosen target has only one Target Service Agent loaded	Step 4

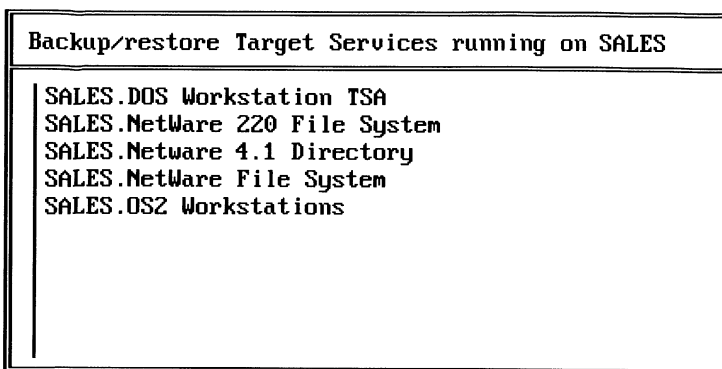
3. From the Target Services list (similar to Figure 9-40), choose a target.

If the target you choose has more than one Target Service Agent loaded, SBACKUP will show you a list of their full names.

Figure 9-40 shows that the NetWare server named SALES is running the following Target Service Agents:

DOS Workstation (TSADOS)
NetWare 2.2 File System (TSA220)
NetWare 4.1 Directory (TSANDS)
NetWare [4.1] File System (TSA410)
OS/2 Workstation (TSAPROXY).

Figure 9-40
Display Multiple
Target Service
Agents on a Target



You can restore data to any of these environments.

4. When prompted for the Target username, enter your username (and context if required) as network supervisor for the target.

If SBACKUP rejects the supervisor username you entered, you probably need to include the context of where the user object is located. For example, instead of entering "ADMIN" as the username, enter ".CN=ADMIN.O=*company_name*".

You must include the context in the username at this point if the username exists in a directory container that is different from the context set on the server running SBACKUP.

5. If a password is requested, enter the password for the target.

SBACKUP will take a few moments to attach to the target. Wait for the confirmation box; then press any key to continue.

6. From the “Restore Menu,” select a restore option, using the information in the following table to make your choice.

Choose one of the following	Then
Select “Choose a Session to Restore” if you want to select a session from a list of session log headers.	Go to Step 6a.
Select “Restore Without Session Files” if you need to select a session directly from the backup media. This is necessary when log and error files are corrupted or accidentally deleted.	Go to Step 7. (Later in the restore session, SBACKUP will display session headers that you can accept or reject one at a time.)

6a. When the prompt shown in Figure 9-41 is displayed, specify the path to the session log file of the session you want to restore.

Figure 9-41
Specify the Path to the Backup Session


Where are the Session Log Files?
SYS:SYSTEM/TSA/LOG

To specify a path, use one or more of the following methods.

- ◆ Press <Enter> to accept the default; then go to Step 6b.
- ◆ Press <Insert> to choose from a list of directories, press <Esc> to return to the previous window, and then go to Step 6b.
- ◆ Backspace over the path shown (or a portion of the path), type in a new directory or path, and then go to Step 6b.

Press <Insert> anytime during this process to choose parts of the existing path from a list.

- 6b. When a list of sessions is displayed (similar to Figure 9-42), choose the one you want to restore, and then skip to Step 8.**



Don't restore an NDS session from a tape backup session unless all replicas have been corrupted.

Figure 9-42
Specifying Sessions to Restore

Select the Session You Want to Restore			
Description	Date	Time	Source
SALES-SYS:SYSTEM/PETE/OLDNLMS	12-15-92	10:44:30am	SALES
Full Merry's media	12-15-92	10:25:28am	SALES

- 7. If you are restoring without session files, set the location of log and error files.**

You can use one or more of the following methods.

- ◆ Press <Enter> to accept the location; then go to Step 8.
- ◆ Press <Insert> to choose from a list of network directories, choose a directory, press <Esc> to return to the previous window, and then go to Step 8.
- ◆ Backspace over the path shown (or a portion of the path), then type in a new directory or path, and then go to Step 8.


Press <Insert> anytime during this process to choose parts of the existing path from a list.

- 8. Press any key to see your selection of device and media, and decide where you want to restore from.**



If the media on which the data resides is not loaded on the device you choose, you will be prompted to insert the correct media.

- 8a. If a list of devices and media is displayed, choose the one containing the session you want to restore.
- 8b. If only one device is attached, a message is displayed indicating the default device and media. Press <Enter>.
- 9. From the “Restoring From *Backup_Session_Name*” screen, choose the type of restore you want, and then go to the step indicated on the following table.



Warning

TSA410 backs up files in compressed or uncompressed format, as specified by the user. However, if you try to restore a compressed file to a volume without compression, the file is corrupted and no error message is given.



Note

If you have installed a migration system built on RTDM (such as HCSS) and need to restore data after a catastrophic data loss, refer to the steps in “Restoring Migrated Files” on page 743, and then choose “Perform a Custom Restore” and specify when to exclude and include files in the “Subsets of the Session to be Restored” field.

If you select	Then go to
“Restore One File or Directory”	Step 10
“Restore an Entire Session”	Step 11
“Perform a Custom Restore”	Step 12

- 10. To restore one file or directory, complete the form in the “Restore Options” screen (Figure 9-43).

Figure 9-43
Restore Options when Restoring One File or Directory

Restore Options
<p>Name of file:</p> <p>Name of directory (full path):</p> <p>Include subdirectories: No</p> <p>Different location to restore to (full path):</p> <p>Name space of file and directory names: DOS</p>

10a. To restore one file

- (1) Enter the exact name at "Name of file."
- (2) Enter the full path to the file, including directory and volume names.
- (3) Skip the "Include Subdirectories" line.

10b. To restore one directory

- (1) Skip the "Name of file" line.
- (2) Enter the full directory path, including volume name.
- (3) Specify whether or not to include subdirectories.

10c. (Optional) To restore to a new location, enter the full path to the target directory, including the volume name.

- (1) Enter the full directory path, including volume name, to the new location (a location different from where the data was backed up). For example, type "*Volume:/directory/directory.*"

If the new location you are restoring to is a different NetWare server:

- (1) Back up the data from the server where it resides
- (2) Exit to the "Main Menu"
- (3) Choose "Change Target to Back Up From or Restore To" to change your target to the server you want to restore the data to.


10d. Enter the name space for the name space format you used when entering the file or directory name.

The name space allows SBACKUP to correctly interpret the path information entered earlier in this menu in either the directory field or the location field.



If files are restored to a new location, the name spaces must be the same or an error message will occur.

You can only restore to a new location in the name space that is supported by the drive you are restoring to.



On an HPFS drive you can only rename in the OS/2 name space, and on a "fat file" system you can only rename in the DOS name space. However, SBACKUP displays both in the log file when only one is valid.

The supported name spaces are DOS, FTAM, Macintosh, NFS, and OS/2.

SBACKUP provides a list of the name spaces currently loaded on the target you are restoring to.

For example, if restoring to a new Macintosh location, enter a path formatted like "SYS::TEST:MONDAY," then enter "Macintosh" as your name space.

10e. Press <F10> to begin the restore.

10f. At the "Proceed With the Restore?" prompt, choose "Yes." Skip to Step 13.

11. To restore an entire session, at the "Proceed With the Restore?" prompt, choose "Yes." Skip to Step 13.

12. To perform a custom restore, at the "Restore Options" screen (Figure 9-44), complete the form.

Figure 9-44
"Restore Options"
Screen

Restore Options
Subsets of the session to be restored: Press <Enter> to view the list.
How to scan the session to be restored: Press <Enter> to view the list.
Restore data to different location: No
Overwrite existing parent: Yes
Overwrite existing child: Yes



At least two levels of help screens are available to you while you are in the "Restore Options" screen. To see help for a specific line item, select the item, press <Enter> to get in editing mode, then press <F1>.

For more information on a specific option, see "Restore Options" on page 737.

12a. (Conditional) If you want to restore a subset of a backup session only, choose “Subsets of the Session to be Restored” and then edit the form (similar to Figure 9-45).

The screen lists major and minor resources. The word “DEFAULT” indicates that none of that type of subset has been specified. To specify a subset

- (1) Choose any item on the screen.
- (2) Edit the “Include” or “Exclude” screen using <Insert>, <Delete> and/or <Enter>.
- (3) Press <F10> to save your changes.

**Figure 9-45
Subsets of a
Session**

Choose subsets of the session to be restored	
Exclude major TSA resources	Default
Include major TSA resources	Default
Exclude directories (full path)	Default
Include directories (full path)	Default
Exclude files	Default
Include files	Default
Exclude path/files	Default
Include path/files	Default

For more information about major and minor resources, see “Major resource” and “Minor resource” in *Concepts*.

You can include or exclude any or all items listed. For more information about including and excluding items, see “Choosing Subsets of Data to Restore” on page 737.

12b. (Conditional) If you want to exclude certain types of data from the restore, choose “How to Scan the Session to be Restored” and then edit the form (similar to Figure 9-46).

The backed-up data set contains the various types of data items listed on the screen. By default, all types of data items are included in the restore.

Figure 9-46
Scanning a Session

How to Scan the Session to be Restored	
Exclude data streams:	No
Exclude extended attributes:	No
Exclude directory trustees:	No
Exclude file trustees:	No
Exclude volume restrictions:	No
Exclude directory space restrictions:	No

If you want to exclude any item, highlight it, type “Y,” and press <Enter> to save your changes. Press <F10> when you are finished editing the form.

If you have installed a migration system built on RTDM, such as Novell’s High Capacity Storage System (HCSS) and you have already restored your HCSS directory and choose to include migrated files in the restore, choose “Include File Data From Secondary Storage” and change the “No” to “Yes.”

12c. (Optional) To restore data to a different workstation or location on a different NetWare server,

- (1) Back up the data from the server where it resides.
- (2) Exit to the “Main Menu.”
- (3) Choose “Change Target to Back Up From or Restore To” to change your target to the server you want to restore the data to. (See “Changing Targets” on page 779.)

12d. (Optional) To restore data to a different workstation or location within the backed-up server’s directory structure, use the following procedures.

Important 

Some important rules about restoring to a new location:

(1) When you specify a particular portion of the directory structure as the source location, it does not necessarily mean that this will be the only data restored. You can influence what is restored by using the "Include" and "Exclude" options on the "Choose Subsets Of What You're Restoring" screen.

(2) You can only restore to a new location in the name space that is supported by the drive you are restoring to.

(3) On an HPFS drive you can only rename in the OS/2 name space, and on a "fat file" system you can only rename in the DOS name space. However, SBACKUP displays both in the log file when only one is valid.

(4) If you do not want to overwrite any subdirectories that might exist in the area you have included, you must specifically exclude them by their full path.

To restore to a new location,

(1) Select the option, type "Y" and press <Insert>.

(2) Enter the full directory path of the source, including the volume name.

(3) Enter the full directory path to the new destination.

(4) Press <Esc> to save your changes.

Note 

The NetWare Directory Services Target Service Agent (TSANDS) does not support restoration of data to a different area of the directory structure.

12e. Specify whether or not to overwrite the existing parent or child.

A parent might be a server, a volume, or a directory. A child is a file, which is the lowest level of the directory structure.

Use the table below to help you decide.

If	Then
You want to overwrite the parent or child that exists on the hard disk, regardless of whether the existing version or the backed-up version has the latest date	Accept the “Yes” default for parent or child or both.
You don’t want to overwrite the parent or child that exists on the hard disk, regardless of whether the existing version or the backed-up version has the latest date	Select the option (parent or child or both), type “N” and press <Enter>.
You want to overwrite the child that exists on the hard disk only if the date of the backed-up copy on the media is later than the date of the copy on the hard disk	Highlight “Overwrite Existing Child,” press <Enter>, and choose the option.

12f. Press <F10> to save your changes and begin the restore.

12g. At the “Proceed With The Restore?” prompt, press <Enter>.

Three windows are displayed.

13. (Conditional) If you selected “Restore Without Session Files” earlier (Step 6), select a session now.

Although this method of selecting a session (selecting it directly from the media) is more time consuming, it is useful when log and error files are corrupted or deleted for some reason.



Don’t restore an NDS session from a tape backup session unless all replicas have been corrupted.

SBACKUP goes through the media and displays the session headers one at a time, prompting you to specify whether to restore the displayed session or skip to the next session (assuming other sessions have been appended).

Use the table below to assist you.

If	Then
The session you want is displayed	Choose "Yes, Restore This Session" to begin the restore. After a short delay (perhaps as much as a few minutes) the ongoing status is displayed in the windows. Go to Step 14.
The session you want is not displayed	Choose "No, Go On to the Next Session on the Media," and repeat this step until the session you want is displayed.
You want to quit the restore session	Choose "No, Quit Restore." The "Restoring From <i>Backup_Session_Name</i> " screen displays. Go back to Step 9 or continue to exit SBACKUP by pressing <Esc>.

14. (Optional) To abort a restore session after choosing "Yes, Restore This Session," press <Esc> and choose "Yes" from the "Abort Menu."

SBACKUP finishes restoring the current data set, then displays a "Dismounting Media" message. Dismounting might take a few minutes.

When the media is dismounted, the "Restore Terminated" message is displayed. Press <Enter> and the "Restore Menu" is displayed.

15. Verify that the restore was successful.

When the restore has finished, your screen display should look similar to Figure 9-47.

Figure 9-47

The Restore Process Completed (Typical Display)

```
SYS:ETC/I01/
SYS:ETC/I02/
SYS:ETC/SAMPLES/
  ATYPES.CFG
  BUILTINS.CFG
  GATEWAYS
  HOSTS
  NETPROTO.CFG
  NETWORKS
  PROTOCOL
  SERVICES
  SNMP.CFG
Parent name: SAMPLES
Child name: TRAPTARG.CFG
Parents: 30
Children: 1306
Elapsed time: 00:04:18
Total written: 40594 KB
Media ID: 1
The restore process was completed normally.
Press <Enter> to continue.
```

16. To return to the “Main Menu,” press <Enter> and then <Esc>.

The <Esc> key allows you to exit SBACKUP one screen at a time until you reach the “Main Menu.” If you want to exit SBACKUP, press <Esc> once again and answer the confirmation prompt.



To unload SBACKUP and applicable Target Service Agents, see “Unloading SBACKUP Files” on page 781.

Log and Error Files

Introduction to Log and Error Files

SBACKUP generates a log file and an error file each time a backup session is performed and records information in these files specifically about the session.

SBACKUP also creates a default directory (such as SYS:SYSTEM/TSA/LOG) for the log and error files. If you prefer, you can create any directory for the log and error files before the backup session or restore session, as long as it resides on the host server.



You might want to create individual log directories for the different types of backup or restore session targets or different organizational units. For example, you could create one directory for workstation backups and another for server backups.

SBACKUP keeps a list of all the log and error files. These lists show

- ◆ The description you enter for the session
- ◆ The date and time the session was begun or, in the case of a delayed backup session, the time the session was scheduled
- ◆ The name of the target the data was backed up from

The Log File

The log file is created on the host the first time a particular set of data is backed up. This log contains

- ◆ The session date and time, and the description you entered
- ◆ The target from which the data was backed up
- ◆ Media set identification information
- ◆ The area of the directory structure that was backed up (such as the volume name, directory name, etc.)

- ◆ The names of files that were backed up
- ◆ The numerical location of the data on the storage media

The Error File

The error file is created on the host the first time a particular set of data is backed up. It contains a list of any errors that occurred during a backup or restore session.

This file contains

- ◆ The session date and time, and the description you entered
- ◆ The target from which the data was backed up
- ◆ Media set identification information
- ◆ The area of the directory structure that was backed up
- ◆ The total number of parents and children that were backed up
- ◆ The names of files that were not backed up, along with any error messages or information
- ◆ Skipped data sets (any file that is opened when a backup session begins is not backed up and is listed as a skipped data set, unless opened in a read-only mode)

Viewing a Log File

Prerequisites



Checklist

- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.


Procedure



Procedure

1. From the SBACKUP “Main Menu,” choose “Log/Error File Administration.” The menu shown in Figure 9-48 is displayed.

Figure 9-48
“Log/Error File Administration”
Menu



2. Choose “View a Log File.”
3. If necessary, set the location of the log and error files.

If you have already set the location, go to Step 4.

To set the location, use one or more of the following methods.

- ◆ Press <Enter> to accept the default location, and then go to Step 4.
- ◆ Press <Insert> to choose from a list of network directories, choose a directory, and press <Esc> to return to the previous window; then go to Step 4.

- ◆ Backspace over the path shown (or a portion of the path), then type in a new directory or path, and go to Step 4.

You can also press <Insert> during this process to choose existing parts of the path from a list.

4. From the list presented, choose the log file you want to view.

A backup log is displayed, similar to the log shown in Figure 9-49.

Figure 9-49

Sample Backup Log for Weekly Full Backup

```

Backup Log for Session: Monthly full backup
Session date and time: 5-18-94 11:09:54am Log
Description: Monthly full backup
Target: SALES (NetWare 4.10)
Media Set ID: MediaSet.label #1 5-18-94 10:53:50am
Backed up: NetWare server
:1e98 07ca 05 12 0b 09 36 1e98 07ca 05 12 0a 35 32

0001 00000001 000000C5
                DOS: >SYS:
0001 00000001 000001BC
                DOS: >VOL$LOG.ERR
0001 00000001 00000403
                DOS: >TTS$LOG.ERR

```

Use the Up- and Down-arrow keys to view the file, and the <Ctrl>+<PageUp> or <Ctrl>+<PageDown> keys to get to the top or bottom of the file.

The backup log displays all supported name space types. The DOS name space is the first in the group. A right angle bracket (>) appears next to the file's creator name-space type.

If you have linked UNIX files, SBACKUP retains both symbolic and hard links created under the NFS name-space type, but links are not listed in the backup log. Restoring the original file restores associated links.

5. To return to the “Main Menu,” press <Enter> and then <Esc>.

The <Esc> key allows you to exit SBACKUP one screen at a time until you reach the “Main Menu.” If you want to exit SBACKUP, press <Esc> once again and answer the confirmation prompt.



To unload SBACKUP and applicable Target Service Agents, see “Unloading SBACKUP Files” on page 781.

Viewing an Error File


An error file records any errors that occurred during the backup or restore session.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.


Procedure



- 1. From the SBACKUP “Main Menu,” choose “Log/Error File Administration.”**

The “Log/Error File Administration” menu is displayed (Figure 9-50).

Figure 9-50
“Log/Error File Administration” Menu



2. Choose “View an Error File.”

3. If necessary, set the location of log and error files.

If you have already set the location, go to Step 4.

To set the location, use one or more of the following methods.

- ◆ Press <Enter> to accept the default location, and then go to Step 4.
- ◆ Press <Insert> to choose from a list of network directories, choose a directory, and press <Esc> to exit; then go to Step 4.
- ◆ Backspace over the path shown (or a portion of the path), and type a new directory or path; then go to Step 4.

You can also press <Insert> during this process to choose existing parts of the path from a list.

4. From the list presented, choose the error file you want to view.

An error file is displayed, similar to the file shown in Figure 9-51.

Figure 9-51

Sample Error File for Weekly Full Backup

```
Backup and Restore Errors for Session: Monthly full backup      ▼ 73
Session date and time: 5-18-94 11:09:54am Errors
Description: Monthly full backup
Target: SALES (NetWare 4.10)
Media Set ID: MediaSet.label #1 5-18-94 10:53:50am
Backed up: NetWare server
:1e98 07ca 05 12 0b 09 36  1e98 07ca 05 12 0a 35 32

(SBACKUP-4.02-171) SBACKUP successfully backed up 30 parents and 1316
children. Total written: 40684 KB. Elapsed time: 00:04:17.

***** The following files or directories were not backed up. *****
SYS:SYSTEM/TSA/LOG/1CB2593B.ERR (In use)
SYS:SYSTEM/TSA/LOG/1CB2593B.LOG (In use)
```

Use the Up- and Down-arrow keys to view the file, and the <Ctrl>+<PageUp> or <Ctrl>+<PageDown> keys to get to the top or bottom of the file.

Restore errors are appended to a session's backup error file. A "Restore Session Begins" message marks the beginning of the restore errors, if any exist.

If you have linked UNIX files, symbolic and hard links created under the NFS name-space type can fail to be restored if the path name is not recognized. When this happens, the error file contains messages similar to the following:

```
Error restoring name space specific information of  
"XXXX:tmp/hosts" in NFS name space, error 0x7!
```

```
Unable to allocate directory handle for  
"XXXX:tmp/test"!
```

5. When you are finished viewing the file, press <Esc>.

6. To return to the "Main Menu," press <Esc> three times.

The <Esc> key allows you to exit SBACKUP one screen at a time until you reach the "Main Menu." If you want to exit SBACKUP, press <Esc> once again and answer the confirmation prompt.



To unload SBACKUP and applicable Target Service Agents, see "Unloading SBACKUP Files" on page 781.

Setting Location of Log and Error Files

Every time you do a backup or restore, you are required to set a location (or directory path) for the log and error files.


You can either set the location when requested in the backup or restore procedures, or you can use the following method prior to beginning the backup or restore procedures.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



- 1. From the SBACKUP “Main Menu,” choose “Log/Error File Administration.”**

The “Log/Error File Administration” menu is displayed.

- 2. Choose “Set Location of Log and Error Files,” and set the location.**

To set the location, use one or more of the following methods.

- ◆ Press <Enter> to accept the default location, and then press <Esc> to return to the “Main Menu.”
- ◆ Press <Insert> to choose from a list of network directories, choose a directory, and press <Esc> twice to return to the “Main Menu.”
- ◆ Backspace over the path shown (or a portion of the path), type a new directory or path, press <Enter>, and then press <Esc> to return to the “Main Menu.”


You can also press <Insert> during this process to choose existing parts of the path from a list.

- 3. Press <Esc> when finished to return to the “Main Menu.”**

Deleting the Log Files for Overwritten Sessions

SBACKUP allows you to overwrite sessions and reuse storage media.

When you delete the log file for a session, using the following procedures, you automatically delete the corresponding error file as well.




To prevent outdated session information from taking up disk space and to help prevent confusion when you select sessions to view or restore, delete the session log files of any sessions that have been overwritten.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



- 1. From the SBACKUP “Main Menu,” choose “Log/Error File Administration.”**

The “Log/Error File Administration” menu is displayed.

- 2. Choose “View a Log File.”**

3. If you have not set the location of log and error files, do so now.

You can use one or more of the following methods.

- ◆ Press <Enter> to accept the default location; and then go to Step 4.
- ◆ Press <Insert> to choose from a list of directories, choose a directory, and press <Esc> to return to the previous window; then go to Step 4.
- ◆ Backspace over the path shown (or a portion of the path), and type a new directory or path; then go to Step 4.

You can also press <Insert> during this process to choose existing parts of the path from a list.

4. Select the description of the session that you have overwritten (or plan to overwrite) and press <Delete>.

To delete several session files at once, select each one, using <F5>; then, when finished marking, press <Delete>.

A confirmation prompt is displayed.

5. To delete the log and error file for that session, choose “Yes.”

6. Press <Esc> when finished to return to the “Main Menu.”

Storage Device Administration

You can see a list of available storage devices by choosing “Storage Device Administration” from the SBACKUP “Main Menu.” For each listed device, you can access the following information, using the options listed at the bottom of the screen:

- ◆ Viewing media lists
- ◆ Checking status of a device
- ◆ Checking status of media in a device
- ◆ Media utilities
- ◆ Renaming a storage device
- ◆ Managing storage media

Viewing Media Lists


This option lists the available media for the highlighted device, and provides the location and media name.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



1. **From the SBACKUP “Main Menu,” choose “Storage Device Administration.”**

The “List of Devices” screen is displayed.

2. **Select one device name, and press <Insert>.**

If a message is displayed saying that no media is in the device, but you know that there is, return to the device list and press <Enter>. This will force the media to be read.

3. **Press <Esc> when finished to return to the “Main Menu.”**

Checking Status of a Device

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



1. **From the SBACKUP “Main Menu,” choose “Storage Device Administration.”**

The “List of Devices” screen is displayed.

2. Select one device, and press <Tab>.

The “Status of: *Device_Name*” screen is displayed, listing the following:

- ◆ **Current operation:** This option indicates whether the device is currently reading, writing, formatting the media, or none of these.
- ◆ **Device mode:** This option indicates whether the device is selected for reading, writing, or both.
- ◆ **Maximum capacity:** This option indicates total storage capacity of the media currently in the device.

3. Press <Esc> when you are finished to return to the “Main Menu.”


Checking Status of Media in a Device

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



1. From the SBACKUP “Main Menu,” choose “Storage Device Administration.”

The “List of Devices” screen is displayed.

2. Select one device name, and press <Insert>.

The “Media for: *Device_Name*” screen is displayed, listing the names of any media in the device you selected.

If, instead, a message is displayed saying that no media is in the device, but you know that there is, return to the device list and press <Enter>. This will force the media to be read.

3. Select one media name, and press <Tab>.

The “Status of: *Media_Name*” screen is displayed, listing the following:

- ◆ **Media owner:** If the media in the device is blank or if it was written to in a data format other than SIDF, the owner is listed as “Unidentified.”



If the media is written by an SIDF application, the media owner is not displayed on the status screen.

- ◆ **Number in media set:** This number indicates the sequential number of this media within a particular media set.

SBACKUP automatically labels and appends an incrementing number to any backup sessions that spans multiple media.

- ◆ **Creation time:** This shows the time that the current media label was first used on this media.
- ◆ **Mount status:** This indicates whether the media is “Mounted” (ready for reading or writing), “Mount Pending” (request to mount is waiting to be processed), or “Not Mounted.”
- ◆ **Media mode:** This indicates whether the media is selected for reading, writing, or both, or whether the mount request is pending.
- ◆ **Media type:** This indicates the type of media in the device, such as 4mm DDS (Digital Data Storage) tape, 8mm tape, etc.
- ◆ **Total capacity:** This indicates the total capacity of this media.

4. **Press <Esc> when you are finished to return to the “Main Menu.”**

Media Utilities


Use the following procedures to erase media headers, to erase all data on the media, or to retension media in tape devices.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



1. **From the SBACKUP “Main Menu,” choose “Storage Device Administration.”**

The “List of Devices” screen is displayed.

2. **Select one device name, and press <Insert>.**

The “Media for: *Device_Name*” screen is displayed, listing the names of any media in the device you selected.

If, instead, a message is displayed saying that no media is in the device, but you know that there is, return to the device list and press <Enter>. This will force the media to be read.

3. **Select one media name, and press <Insert>.**

4. From the “Utilities For: *Media_Name*” menu, choose one option.

The following table explains each option.

Option	Explanation
“Erase the Media Header”	This is a quick method of making the media appear blank. Although this is less secure than erasing all of the data, this method is quick and effective.
“Erase all Data on the Media”	This is a security measure that might take up to 2 hours, depending on the size of the media.
“Retension this Media”	This option is for tape devices that allow retensioning (0.25-inch cartridges, for example). This will fast-forward and rewind the tape, which might resolve a problem reading the tape.

Renaming a Storage Device


The default names for storage devices are controlled by the devices. These names often are not descriptive. If you have several devices, it might be difficult to remember which name goes with which device. For this reason, renaming storage devices might be helpful.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



1. From the SBACKUP “Main Menu,” choose “Storage Device Administration.”

The “List of Devices” screen is displayed.


2. Select the device name that you want to rename, and press <F3>.
3. Type a new name or identifier for the device, and then press <Enter>.
4. Press <Esc> when finished to return to the “Main Menu.”



Any new device names that you assign are lost if the NetWare server goes down. When you load the SBACKUP module after the server has been down, the device names revert to the default.

Managing Storage Media

If the storage media you are using does not have enough empty space for the entire backup, SBACKUP will span media, prompting you to insert additional media when the first is full.



Always have extra media on hand in case the backup session spans more than one media.

SBACKUP designates the first media that has a particular label as #1 and increments by 1 any subsequent media with the same label.



To help safeguard your network, SBACKUP exits when a delayed backup is complete. If the delayed backup session cannot fit on the media, SBACKUP prompts you to insert additional media. If additional media is not inserted, the backup does not finish and SBACKUP does not exit, thus compromising security.

Changing Targets

Anytime you are performing a series of backup or restore sessions, you may need to change your target. This is easily done from the SBACKUP “Main Menu.”


If you don’t specify a new target, SBACKUP automatically selects the target used during the last session and continues without asking you to make a selection.

Prerequisites



- Ensure that you have loaded the drivers for your specific device and controller board.
- Ensure that you have loaded the SBACKUP files for your specific target. (See “Loading Files Needed to Run SBACKUP” on page 687.)
- Ensure that media is inserted into your storage device.
- Ensure that the SBACKUP “Main Menu” is displayed on your server console.

Procedure



- 1. From the SBACKUP “Main Menu,” choose “Change Target to Back Up From or Restore To.”**

The “NetWare Servers Running Your Backup/Restore Target Services” screen is displayed.

- 2. Choose one server.**
- 3. From the list of Target Service Agents, similar to Figure 9-52, choose a target.**

Figure 9-52
Display of Multiple
Target Service
Agents on a Target

```
Backup/restore Target Services running on SALES
SALES.DOS Workstation TSA
SALES.NetWare 220 File System
SALES.Netware 4.1 Directory
SALES.NetWare File System
SALES.OS2 Workstations
```

If the target you choose has more than one Target Service Agent loaded, SBACKUP will show you a list of their full names.

Figure 9-52 shows that the NetWare server named SALES is running the following Target Service Agents:

- DOS Workstation (TSADOS)
- NetWare 2.2 File System (TSA220)
- NetWare 4.1 Directory (TSANDS)
- NetWare [4.1] File System (TSA410)
- OS/2 Workstation (TSAPROXY).


You can back up from or restore data to any of these environments.

- 4. Type in the username and password for the target you just selected.**

Unloading SBACKUP Files

You can unload SBACKUP if you want to have more memory available on your host or target.

Procedure



1. Exit SBACKUP.

Press <Esc> until you reach the SBACKUP “Main Menu”; then press <Esc> again and choose “Yes” at the confirmation prompt.

2. Unload the server files.

To unload the server files, type the following commands at the server console prompt in the order shown:



The SMSDI, WSMAN, and SMDR files are automatically loaded by SBACKUP and other Storage Management Services modules, but they need to be manually unloaded.

```
UNLOAD SMSDI <Enter>
UNLOAD TSA410 <Enter>
UNLOAD TSADOS <Enter> (if applicable)
UNLOAD TSAPROXY <Enter> (if applicable)
UNLOAD TSANDS <Enter> (if applicable)
UNLOAD SMDR <Enter>
```

3. Unload the workstation files, if applicable.

To unload the DOS workstation file, type the following command at the workstation prompt.

```
TSASMS /U <Enter>
```

If the workstation had the message reception turned off (using the command “SEND /A=N”), turn the message reception back on using the following command:

```
SEND /A=A <Enter>
```

Enhancing SBACKUP Performance

Compressed Files

When a NetWare®4.1 volume is mounted, file compression is set to ON by default. When you perform a backup using TSA410, you can specify whether to keep already compressed files in a compressed state for the backup, or to back them up in an uncompressed state.



Note

File compression is specified on the form “How To Scan What You Are Backing Up” found in Figure 9-28 on page 731.

To help you make your decision, remember these points:

- ◆ The backup is faster if files already compressed are left compressed.
- ◆ Compression is not supported in some environments (for example, on a NetWare 3.11 server, or a DOS workstation). Therefore, if you intend to restore a file that is currently compressed to an environment that does not support compression, you should back it up in an uncompressed state.
- ◆ SBACKUP itself has no compression feature, so it cannot compress a file that is currently uncompressed.
- ◆ If volume compression is turned on and you try to back up compressed files in an uncompressed state, restore speed is degraded if you overwrite existing files. To improve restore speed, delete the files you no longer want from the hard disk before restoring them from the backup media. If you leave compressed files in a compressed state to restore them, speed is not affected.
- ◆ You might run out of disk space if you restore uncompressed files to a volume that compresses them.



Warning

SBACKUP and file compression should not be run simultaneously. However, the default time for both delayed SBACKUP sessions and compression is midnight.

If you want to perform a delayed backup that includes files flagged for compression, schedule the delayed backup after the compression time to allow time for the completion of compression.

Enhancing Host Server Performance

The speed of SBACKUP depends upon the host server's configuration and whether or not the server is backing up its own data or that of another server or workstation.

For information on server configuration and setting parameters, see "Communication Parameters" and "File Caching Parameters" under "SET" in *Utilities Reference*.

A server backing up its own data runs about twice as fast as a server backing up data from another server.

This section includes four options using LOAD and SET commands that can enhance your host server's performance:

- ◆ Changing the size or number of buffers
- ◆ Setting minimum packet receive buffers
- ◆ Setting subdirectory levels and minimum cache buffers
- ◆ Setting reserved buffers below 16 MB

Changing the Size and Number of Buffers

If your server has insufficient free memory, or if it has memory not being used, you can speed up a backup session by increasing the number and size of the host server's buffers when you load SBACKUP.

You can change the size and number of buffers when you load SBACKUP. To do this, type

```
LOAD SBACKUP SIZE=XXX BUFFER=X <Enter>
```

Select the size and number of buffers from the following chart:



Warning

If you use an 8-bit or 16-bit host adapter and a driver that does not force the use of memory below 16 MB, do not specify a buffer-size and memory-size combination that exceeds 16 MB. If you do, corrupted data is copied to the backup media.

Option	Range
Size	16, 32, 64, 128, or 256 KB (Default=64 KB)
Buffers	2 to 10 (Default=4)




Note

For more about changing the size or number of buffers, see your device driver's documentation.

Setting Minimum Packet Receive Buffers

The packet receive buffers represent the space in the server memory dedicated to handling network traffic. If the buffers are set too low, the server performance may degrade.

Procedure



Procedure

1. Include the following SET command in your STARTUP.NCF file:

```
SET MINIMUM PACKET RECEIVE BUFFERS = XXX
```


The *x*'s represent the number of buffers. Set two buffers for each user on the system. The minimum is 10.

2. To activate the changes, bring down the server (type DOWN at the server console prompt) and reboot it.

Setting Maximum Subdirectory Level and Minimum Cache Buffers

If you change the maximum subdirectory level in the server's STARTUP.NCF file, you must also change the minimum cache buffers. The maximum subdirectory level default is 25.

Procedure



1. **Include the following commands in the AUTOEXEC.NCF file:**

```
SET MINIMUM FILE CACHE BUFFERS=X  
SET MAXIMUM DIRECTORY CACHE BUFFERS=X
```

Replace the *x* with a number equal to or greater than the subdirectory level you set in the STARTUP.NCF file. The default is 20 for the file buffers and 500 for the directory buffers.

2. **To activate the changes, down the server (type DOWN at the server console prompt) and bring it back up.**

Setting Reserved Buffers Below 16 MB

If your storage device driver requires memory below 16 MB and the server has more than 16 MB available to it, you must reserve memory below 16 MB for the driver.

Procedure



1. **Include the following command in the STARTUP.NCF file:**

```
SET RESERVED BUFFERS BELOW 16 MEG=X
```

Replace the *x* with a number between 8 and 300. The default is 200.

2. **To activate the changes, bring down the server (type DOWN at the server console prompt) and reboot it.**

Troubleshooting

If you have trouble while using SBACKUP, check Table 9-6 to see if you can solve the problem.

Table 9-6
Troubleshooting Guide for SBACKUP Problems

Problem	Possible Cause	For more information, see
The media owner is "Unidentified."	Media is blank or was written in a non-SIDF format.	"Checking Status of Media in a Device" on page 774
Backup speed is slow.	Compressed files are being backed up in an uncompressed format.	"Enhancing SBACKUP Performance" on page 782
Restore speed is slow.	File compression and SBACKUP are running at the same time. Compressed files are being overwritten with uncompressed files.	"Enhancing Host Server Performance" on page 783
Data is corrupted.	Volumes were mounted or dismounted or drivers were unloaded during a session.	"Rules for Using SBACKUP" on page 682
Data is corrupted, but no error message was given.	Compressed files were restored to a volume without compression.	"Prerequisites to Restoring" on page 746
The server is in an abend condition.	Volumes were mounted or dismounted or drivers were unloaded during a session.	"Rules for Using SBACKUP" on page 682
The specified file or directory cannot be found.	Wrong spelling or wrong case was used for case-sensitive name spaces.	"Rules for Using SBACKUP" on page 682
The target you want is not listed.	The Target Service Agent isn't loaded on the desired target. SBACKUP has not had time to "find" the target. The host and target are physically far apart and network traffic is heavy.	"Backing Up Data" on page 698 (Not applicable) (Not applicable)

Table 9-6 *continued*

Troubleshooting Guide for SBACKUP Problems

Problem	Possible Cause	For more information, see
You (network supervisor) can't back up a workstation.	The "/Password" option is set and is denying access to data on the workstation.	Table 9-1 on page 692)
	Other CPU-intensive applications were running, and they did not allow the workstation TSR to obtain resources.	(Not applicable)
	Backup was attempted from two different servers at the same time. One succeeded and the other one failed.	(Not applicable)
A backup does not contain all changes.	Differential and incremental backups were combined. Use one or the other of these types in conjunction with full backups.	"Backup Types" on page 699
	The modify bit was cleared after the last custom backup, so changed files are not recognized.	"Backup Types" on page 699
You cannot find the session log file you want.	It might be in a different directory, or you might have deleted it accidentally.	"Deleting the Log Files for Overwritten Sessions" on page 770
Pressing <Insert> on device list does not produce a media list.	No medium has been selected.	"Select" option at the bottom of the screen.
"The time form is invalid" message is displayed when you try to schedule a delayed backup.	The date has been incorrectly entered in the time form. Press <F1> for help with the time or date format.	(Not applicable)
Files were restored but the error file contains a message specifying which name space formats were not restored.	The file attributes and name-space formats are not configured on the volume you restored to.	(Not applicable)



appendix

A

Troubleshooting the Network

This appendix provides general hardware and network troubleshooting information for resolving equipment-related errors on a server running NetWare 4™ networking software.

Troubleshooting Hardware and Network Problems

Making Computer Memory Available for Network Drivers


To increase the amount of workstation memory available for network drivers:

- ◆ Add more memory.
- ◆ Unload unneeded TSR (terminate-and-stay-resident) programs (DOS clients only).
- ◆ Modify the NET.CFG file to load only necessary drivers.
- ◆ Modify the CONFIG.SYS file to load only necessary drivers.

For more information about freeing up memory on a workstation, see “Resolving Workstation Memory Problems” on page 802.

For information about freeing up memory on a server, see “Resolving Server Memory Problems” on page 798.

Resolving Server Hardware Problems



Before servicing the NetWare® server, notify users that the server is going down and they may lose data if they don't log out. If possible, do a complete system backup.

To diagnose and resolve problems with NetWare server hardware:

- ◆ Make sure the server's configuration settings and network address are correct.
- ◆ Make sure the servers and workstations are using the same packet frame type (such as Ethernet 802.2 or Ethernet 802.3).
- ◆ If hardware was recently installed on the server, remove the hardware and reboot the server.

If the error does not reoccur, run diagnostics on the removed hardware. Check for conflicts (interrupts, memory addresses, I/O ports, DMA, etc.) between the new and existing computer hardware.

- ◆ Check the system board, memory board, memory chips/hard disk.

The following tools provide server diagnostic capabilities:

- ◆ LANalyzer®
- ◆ MONITOR.NLM
- ◆ SERVMAN.NLM
- ◆ Manufacturer's diagnostic programs

Resolving Network Problems

To diagnose and resolve network communications problems:

- ◆ Make sure configuration settings and network numbering are configured correctly.

List server and workstation settings in a binder. Include a printout of the current configuration files.

- ◆ Make sure network numbering is set up correctly on the network.
 - ◆ Each network segment must have a *unique* IPX™ external network number.
 - ◆ On that network segment, each node must have the *same* IPX external network number and a *unique* node number.

(For more information, see “Network numbering” in *Concepts*.)

- ◆ Make sure the servers and workstations are using the same packet frame type (for example, Ethernet 802.2 or Ethernet 802.3).
- ◆ Maintain an error log recording previous errors and problems experienced by the servers and workstations on the network. Whenever possible, get a screen dump when an error occurs and include it in the binder.

Resolving Cabling Problems

- ◆ Use the proper cabling for your network topology as specified by IEEE. Make sure cable segments do not exceed the recommended lengths.
- ◆ Make sure cable segments are properly terminated for the type of cabling being used.
- ◆ Make sure terminators and in-line cable connectors are working properly.

If you are not sure whether a terminator or connector is working properly, replace it. If the new components work properly, discard the old ones.

- ◆ Make sure there are no breaks in the cable or shield. Use a time delay reflectometer (TDR), a LANalyzer, or a volt ohm meter (VOM) to test cabling for breaks in the cable conductor or shield.
- ◆ Make sure cabling is routed away from devices that produce high electric or magnetic fields, such as fluorescent lights, microwaves, radar, X-rays, copy machines, etc.

Power Supply Errors

An inconsistent power source is the most common cause of hardware problems. It also produces the most devastating results.

Power outages cause workstations, network servers, print servers, and backup devices to reboot. When this happens, all information stored in RAM is lost, and sometimes hardware is damaged.

Power spikes and “brownouts” can also cause a variety of hardware errors.

You can have reliable network performance only if you plan for power outages and fluctuations and protect against them. The following tips can help you do this:

- ◆ Add a dedicated power feed and ground line from your breaker box to critical equipment. Make sure the ground line connects to earth ground.
- ◆ Install an uninterruptible power supply (UPS) or a standby power system (SPS) to provide power to critical equipment for 15 minutes after a power outage.

The capacity of such power supplies is limited, so you may not want to plug nonvital hardware (such as monitors or printers) into the UPS or SPS line.

- ◆ Train users to save data and log out of the network when the lights go out. Then you can bring down the NetWare server in an orderly fashion without forcing users off the system while the server is running on UPS or SPS power.
- ◆ Install a surge suppressor or power conditioner on all power lines that are used by computers. Many UPS and SPS devices already have this feature.
- ◆ Do not allow anyone to plug fans, printers, copy machines, vacuum cleaners, or other motor-driven appliances into the dedicated line or into any power line with computers.

Insert dummy plugs into open outlets to prevent people from plugging such appliances into computer power sources.

Troubleshooting the NetWare Server

This section provides troubleshooting suggestions for typical NetWare® server problems such as disk I/O errors, insufficient disk space, and insufficient memory.

Resolving Abends

Overview of Server Abend Messages

The NetWare 4™ operating system is very resilient, but errors can and will arise. Serious problems are usually accompanied by abend (abnormal end) messages.

Abend messages are usually caused by consistency check errors, but they also can be caused by insufficient memory, DMA (Direct Memory Access) conflicts, or hardware and software interrupts.

Abend messages are displayed in a format similar to the following:

Abend: `SERVER-4.10-message_number message_string`

Abend messages are documented in *System Messages*.

Consistency Check Errors

Consistency check errors are internal tests placed in the NetWare operating system by Novell software engineers.

The primary function of consistency checks is to ensure the stability and integrity of internal operating system data.

Numerous consistency checks are interlaced throughout NetWare to validate critical disk, memory, and communications processes.

Consistency check errors might be caused by a corrupted operating system file, by corrupted or outdated drivers and NetWare Loadable Module™ (NLM) programs, or by hardware failure.

These errors can also be associated with defective memory chips, static discharges, faulty power supplies, or power surges or spikes.

What to Do When You Get an Abend

If an abend occurs, note the abend message for reference. Then try one or more of the following

- ◆ Reboot the server.
- ◆ Make sure the server has sufficient memory. See “Resolving Server Memory Problems” on page 798.
- ◆ Check the system board, memory chips, power supplies, and power conditioning equipment. Replace any faulty hardware.
- ◆ Make sure the server has the latest drivers and NLM™ programs. Update everything that is not current.
- ◆ If new hardware has been recently installed in the server, remove the hardware. If the error does not recur, run diagnostics on the removed hardware.

Make sure there are no conflicts (interrupts, memory addresses, I/O ports, DMA) between the hardware in the computer and the new hardware.

- ◆ If the hardware is good, unload individual NLM programs, reboot, and repeat the procedures that created the original abend. You may be able to isolate an NLM that is causing the problem.
- ◆ Reinstall the operating system, drivers, or NLM programs from the master diskettes. If the server gets a serialization error (abend messages 1418, 1419, 1420, and 1421), reinstall SERVER.EXE from the master diskettes.

Make sure the copy of NetWare being installed is a correct, legal (nonpirated) original.

If none of these remedies resolve the problem, record all the hardware installed on the server and all actions you have taken to correct the problem. Then contact your Novell Authorized Reseller^{CLM} representative for technical support.

Preparing a Core Dump

Upon contacting your Novell Authorized Reseller representative, you may be asked for a “core dump” that may be analyzed in Novell’s diagnostic laboratory. A core dump is a “snapshot” of the server’s memory at the time of the abend.

Often, though not always, Novell’s technicians can diagnose the problem by analyzing the core dump, searching in memory for clues as to what caused the abend.

If you are asked to prepare a core dump, make sure you have sufficient diskettes on hand to copy all of the computer’s memory.

For example, to copy 12 MB of memory, you need nine 3.5-inch high-density (1.44MB) diskettes.

Also, if you can recreate the problem and describe exactly what steps led up to the abend, you should record this information and send it in along with the diskettes.

If Novell’s technicians are able to correct the problem, and if the problem has been caused by a NetWare software bug, they debug the program and send a patch for the problem.

Resolving Server I/O Errors

To resolve a general disk I/O error on the server, try one or more of the following remedies:

- ◆ Check the disk subsystems to make sure the power is on and the cables are correctly connected between the controller and the subsystems.
- ◆ Make sure the subsystem cables are terminated correctly.
- ◆ Make sure the disks are installed correctly.
- ◆ Make sure you have current NetWare 4-certified disk driver loaded. Many drivers can be found on the NetWare CD-ROM or on the master diskettes. Updated drivers may also be available from third-party disk driver manufacturers.

- ◆ Make sure the interrupt parameters, I/O port settings, slot settings, etc. for the driver match those for the hardware. Also, make sure custom parameters have been set correctly for your hardware.

- ◆ At the server console, type the following command:

SCAN FOR NEW DEVICES <Enter>

This causes the operating system to request controller information about all devices.

- ◆ Load INSTALL, selecting "Maintenance/Selective Install," "Disk Options," and then "Modify Disk Partitions and Hot Fix." Make sure the device is visible and has a valid partition.
- ◆ Increase the Hot Fix™ Redirection Area, using INSTALL.NLM.

To change the Hot Fix Redirection Area on an existing drive, back up all the data on the partition, delete the volumes on the partition, and delete the partition; then re-create it.

Assign the partition a different percentage for the Hot Fix Redirection Area; then recreate the volumes and restore the data.

If you have tried all the preceding suggestions without success, contact your Novell Authorized Reseller representative or drive manufacturer.

Resolving Server Disk Space Problems

To resolve an insufficient disk space error, you should do one or more of the following:

- ◆ Delete unnecessary files and directories from the volume.
- ◆ Change the Minimum File Delete Wait Time SET parameter in your server's AUTOEXEC.NCF file equal to 0 so that files can be purged immediately rather than being retained in a salvageable state on the volume.
- ◆ Use the FILER or NetWare Administrator utility to purge deleted files if they cannot be purged automatically. (The purged files are using up directory table space.)
- ◆ Increase the volume size and/or add more disks to the volume.

- ◆ Increase the percentage of disk space that can be used by a directory.
- ◆ If the disk or volume has space available, check the disk drives and disk channel to see if a failure has occurred.
- ◆ Use SET to increase the percentage of disk space that can be used by a directory.
- ◆ Delete NLM programs that you no longer use.
- ◆ Check and adjust or remove any volume restrictions placed upon users.
- ◆ If the disk or volume has space available, check the disk drives and disk channel to see if a failure has occurred.

For information on SET parameters, see “Managing Server Hard Disks” on page 535 and “SET” in *Utilities Reference*.

Resolving Server Memory Problems

To free up server memory temporarily (until you can add more memory to the server), do one or more of the following:

- ◆ Use the FILER or NetWare Administrator utility to purge deleted files on the specified directory that cannot be purged automatically. (Deleted files are using up directory table space.)
- ◆ Use REMOVE DOS or SECURE CONSOLE to free the memory in the server that has been reserved for DOS.
- ◆ Unload NLM programs, such as INSTALL or MONITOR, that are not currently needed.
- ◆ Dismount volumes that are not being used.
- ◆ Reduce the size or number of volumes that the server supports.
- ◆ Delete unused files and directories on the specified volume.

- ◆ Change the Minimum File Delete Wait Time SET parameter in the AUTOEXEC.NCF file so that files can be purged immediately rather than being retained in a salvageable state on the volume.
- ◆ As a last resort, back up all files in your volume, bring down your server, and use INSTALL to reinitialize the volume. Specify block size of 64 KB and turn the “Block Suballocation” option to “Off.”

(This setting uses a lot of disk space but increases the amount of memory available.)

For more information, see Chapter 7, “Maintaining the NetWare Server,” and “SET” in *Utilities Reference*.

Resolving Locked Device Errors

To resolve a locked device error, try one or more of the following:

- ◆ Wait for a while; the task in process may complete and free the device.
- ◆ Retry the action that resulted in the error.
- ◆ Load MONITOR.NLM, delete all user connections, and disable logins.
- ◆ For a disk device error, unload NLM programs (other than the device driver) that may be using the disk. Dismount all volumes on the disk.

If you have tried all of the above without success, contact your Novell Authorized Reseller representative or the drive manufacturer.

Resolving File I/O Errors

To resolve a file I/O error, try one or more of the following:

- ◆ Check to see that the volume (especially volume SYS:) is mounted. To check volumes, type VOLUMES at the server console prompt.
- ◆ If the volume is out of disk space, error messages will appear on the console screen indicating that the volume is almost out of disk space. Check the console screen.

- ◆ Type DIR at a DOS workstation to see how much space remains.

To increase the amount of free space, do one or more of the following:

- ◆ Delete extraneous files (if you can log in from a workstation).
- ◆ Type SET IMMEDIATE PURGE OF DELETED FILES = ON at the console prompt, and retry the action.
- ◆ If you have an additional disk, increase the size of the volume by creating an additional segment of the volume on the disk.

See also “Resolving Volume I/O Errors” on page 800.

Resolving Volume I/O Errors

To resolve a volume I/O error, try one or more of the following:

- ◆ Make sure *all* devices that contain the volume are online. (Volumes may span multiple devices.) See “Resolving Server I/O Errors” on page 796.
- ◆ Load and execute VREPAIR.
- ◆ Load INSTALL, selecting “Selective Install/Maintain” and “Volume Options.” Make sure the volume is visible.

If you have tried all of the above without success, contact your Novell Authorized Reseller representative or disk drive manufacturer.

Resolving Event Control Block Allocation Errors

Event control block allocation system messages can occur when you first bring up the server or after the server has been running for some time.

These messages indicate that the server was unable to acquire sufficient packet receive buffers, usually called *event control blocks* (ECBs). Running out of ECBs is not a fatal condition.

Servers that run for several days where high loads occur in peaks might exceed the set maximum number of ECBs, causing the system to generate ECB system messages.

If these situations are caused by *occasional* peaks in the memory demand, you should probably maintain your current maximum ECB allocation and allow the message to be generated at those times.

On the other hand, if your server memory load is very high and you receive frequent ECB allocation errors, you should probably set your maximum ECB allocation higher. Use the following SET command in the STARTUP.NCF file:

```
SET MAXIMUM PACKET RECEIVE BUFFERS=number
```



Memory allocated for ECBs cannot be used for other purposes.

The minimum number of buffers available for the server can also be set in the STARTUP.NCF file with the following command:

```
SET MINIMUM PACKET RECEIVE BUFFERS=number
```

Troubleshooting Workstations

This section provides troubleshooting suggestions for typical NetWare® client (workstation) problems, including ways to resolve workstation memory problems. This section also lists common NetWare workstation error messages (the 900 series of messages).

Resolving Workstation Memory Problems

DOS Workstations

To free up DOS workstation memory temporarily (until you can add more memory to the workstation), do one or more of the following.

- ◆ Unload any unneeded TSR (terminate-and-stay-resident) programs.
- ◆ Optimize memory usage by loading DOS and other programs into high memory.
- ◆ Modify the CONFIG.SYS file to reduce the number of files that can be open at the same time, the number of buffers allocated for disk drives, and the memory size allocated by the shell for the DOS environment (the /E option).

The following settings are sufficient for normal workstation operation, but the values can be reduced further until problems occur:

FILES=20

BUFFERS=20

SHELL=C:\COMMAND.COM /E:640 /P

Be sure to reboot the machine after modifying the CONFIG.SYS file.

OS/2 Workstations

Check the CONFIG.SYS file to see if there are any devices that can be unloaded; then try to load the driver again.

(For more information about the CONFIG.SYS file, see *NetWare Client for OS/2 User Guide*.)

Resolving Workstation Disk Space Problems

To resolve an insufficient disk space error, do one or both of the following:

- ◆ Delete unnecessary files and directories from the volume.
- ◆ Unload TSR (terminate-and-stay-resident) programs that have swap files on the hard disk.



appendix

B

Managing NetWare Server for OS/2

NetWare Server for OS/2 and NetWare 4

Use NetWare® Server for OS/2* like any other NetWare 4™ server, with a few exceptions:

- ◆ The server is loaded using NWOS2.EXE instead of SERVER.EXE.
- ◆ Different keystrokes are used for navigating through server screens.
- ◆ DOMAIN.NLM is not supported and will not load.
- ◆ The PMMON utility is available in addition to MONITOR for monitoring the server.


This appendix provides an explanation of these differences.

Loading NetWare Server for OS/2

Load NetWare Server for OS/2 using NWOS2.EXE instead of SERVER.EXE.

Using NWOS2.EXE to Load the Server from the Command Line

Procedure



1. **Open an OS/2 full screen or OS/2 window session.**
2. **Change to the directory where the server files are installed.**

If you used the defaults during the installation, the directory is \NWSERVER and is located on your OS/2 boot drive. If you did not use the defaults, the directory has the name and location you assigned to it.


3. **At the command prompt, type**

NWOS2 <Enter>

NetWare Server for OS/2 loads in the full screen or window session you opened.

Using the OS/2 Graphical Interface to Load the Server

Procedure



1. **Choose the “Novell” group icon on your desktop.**

The “Novell” group icon was created when you installed NetWare Server for OS/2 or NetWare Client for OS/2.

2. **Choose the “NetWare Server for OS/2” icon.**

By default, NetWare Server for OS/2 loads in a full screen session. NetWare Server for OS/2 can be loaded in an OS/2 window by using the OS/2 “Settings” feature on the “NetWare Server for OS/2” icon.

For more information about the “Settings” feature, see your OS/2 documentation.

Navigating in NetWare Server for OS/2

Use the following key sequences for accessing screens and menus in NetWare Server for OS/2:

Task	NetWare Server for OS/2 Key Sequence
Switching between screens	<Alt>+<N>
Accessing selection menu	<Ctrl>+<N>

If you release the <Alt> key too quickly (without pressing another key), you will activate the OS/2 system menu for the window session rather than the next NetWare screen.

Using Loadable Modules in NetWare Server for OS/2

All NetWare Loadable Module™ (NLM) programs and console utilities work with NetWare Server for OS/2 the same as they do on any NetWare 4 server except for the following:

- ◆ **DOMAIN.NLM.** This NLM™ does not work with NetWare Server for OS/2 and cannot be loaded. Memory management is handled by OS/2.
- ◆ **NPRINTER.EXE.** If your computer is running NetWare Server for OS/2 and NetWare Client for OS/2, we recommend that you use NPRINTER.NLM rather than NPRINTER.EXE for printing. NPRINTER.EXE will work, but NPRINTER.NLM will work better for you. Don't load NPRINTER.NLM and NPRINTER.EXE on a computer that is running both the client and server.
- ◆ **VGADISP.NLM.** This NLM is used to display Japanese characters on a monitor screen. VGADISP.NLM does not work on NetWare Server for OS/2 and should not be loaded. OS/2 will take care of Japanese characters.

Monitoring NetWare Server for OS/2

To monitor NetWare Server for OS/2, you can use MONITOR or PMMON. (PMMON is a monitoring utility that runs under the OS/2 Presentation Manager.)



Although MONITOR works with NetWare Server for OS/2, the server utilization information is not correct. For a more accurate report of server utilization, use PMMON.


An icon for PMMON is placed in the “Novell” group icon on the desktop when you install NetWare Server for OS/2.

Loading PMMON

PMMON can be loaded either by choosing the “PMMON” icon or by entering the command at the command line.

Using the OS/2 Desktop

Procedure




1. Choose the “Novell” group icon.

This opens the Novell Icon View.

2. Choose the “PMMON” icon.

The PMMON screen appears (see Figure B-1).

Figure B-1
**PMMON Graphical
 Server Monitor**



Using the Command Line

Procedure



1. Open an OS/2 full screen or OS/2 window session.
2. Change to the directory where the server files from the NetWare Server for OS/2 installation are installed.

If you used the defaults during the installation, the directory is \NWSERVER and is located on your OS/2 boot drive. If you did not use the defaults, the directory has the name and location you assigned to it.

3. Type

PMMON <Enter>

The PMMON screen appears.

Using PMMON

PMMON includes pull-down menus that can be used to set memory sharing options, set performance tuning, free base server memory, access online help, and close the utility.

The following statistics are gathered and displayed in PMMON:

- NetWare Server for OS/2 utilization
- Original cache buffers
- Total cache buffers
- Packet receive buffers
- Directory cache buffers
- Service processes
- The percent of OS/2 CPU usage
- The percent of NetWare CPU usage
- Dirty cache buffers
- Connections in use
- Open files

There are six buttons on the right side of the PMMON screen. These buttons are used to set values and “watermarks” for the bar graphs at the bottom of the utility. For more information about each button, see the online help.

PMMON tracks many of the same statistics as MONITOR, but there is approximately a 5-percent performance loss when PMMON is running instead of MONITOR.



MONITOR does not correctly account for sharing the server with OS/2. For example, in MONITOR the server use may constantly read 100 percent even though the server is not busy.

Setting Performance Tuning


You can use PMMON to set the performance tuning level for NetWare Server for OS/2.

Setting the performance tuning level allows you to specify how much processing time is allocated to NetWare. The higher the performance tuning level, the more processing time is allocated to the server.

Whatever processing time is not allocated to NetWare Server for OS/2 is used by OS/2.

The default setting should work fine for your needs. However, you may want to lower the performance tuning level if OS/2 operation is too slow because of the amount of traffic on the server.

Procedure



1. Choose “Performance Tuning” from the “Configuration” menu.

This activates a submenu listing 10 levels of performance tuning.

2. Choose the level you want for your server.

The performance tuning default level is 8. Changing the level in PMMON overrides the default while the server is running.

If the server is brought down, the default performance tuning level will return to 8 when the server is brought back up (unless you have another level set in the NET.CFG file).

Performance tuning can also be set in the NET.CFG file. (See “Performance Tuning” on page 823.)

Setting Memory Sharing

You can use PMMON to set manual and automatic memory sharing options for NetWare Server for OS/2 and OS/2.

- ◆ Manual memory sharing allows you to allocate more or less memory to NetWare Server for OS/2 or OS/2 while the server is running. You determine the amount of memory each operating system needs.

For example, if the NetWare® operating system is used significantly more on your computer than OS/2, you would allocate more memory to NetWare Server for OS/2.


- ◆ Automatic memory sharing determines the proper amounts of memory NetWare Server for OS/2 or OS/2 will need as the server is running.

Manual Memory Sharing

Use the “Memory Sharing” dialog box to allocate memory manually to NetWare Server for OS/2 or to OS/2.

The allocation changes you make in the “Memory Sharing” dialog box will override any memory settings in the NET.CFG file while the server is running. When the server is brought down and back up, the memory settings in the NET.CFG file will again take effect.

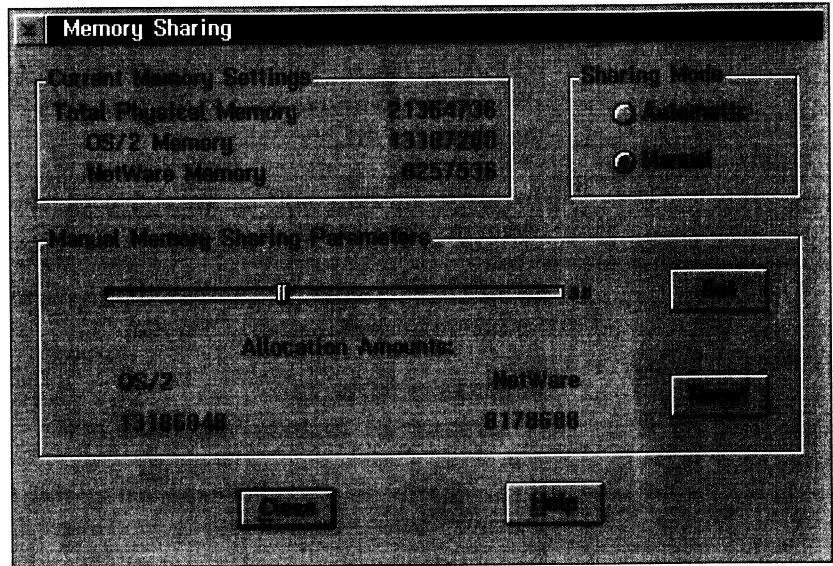
Procedure



1. **Load PPMON.**
2. **Choose “Memory Sharing” from the “Configuration” menu.”**

The “Memory Sharing” dialog box appears (see Figure B-2).

Figure B-2
"Memory Sharing"
Dialog Box



This box displays the current memory settings and the amount of memory used by NetWare Server for OS/2 and OS/2.

3. **Click on the slider and move the mouse to the left or right until you get the allocation ratio you want.**

The amount of memory allocated to NetWare Server for OS/2 and to OS/2 changes as indicated in the "Allocation Amounts" field of the "Memory Sharing" dialog box.



You can also click on the left and right arrow buttons next to the slider to change memory settings.

4. **Choose "Set."**

The changes you make to memory allocation do not take effect until the "Set" button is selected.

Choosing "Reset" switches the memory allocation to the amount of memory that the server had when you opened the "Memory Sharing" dialog box.

5. **Choose "Close" to close the "Memory Sharing" dialog box.**

Automatic Memory Sharing

Automatic memory sharing is active as a default. Automatic memory sharing can be deactivated in the "Memory Sharing" dialog box or in the NET.CFG file.

To deactivate automatic memory sharing, choose "Manual" in the "Memory Sharing" dialog box.

To deactivate automatic memory sharing in the NET.CFG file, enter the following lines in the NET.CFG file:

```
NETWARE FOR OS/2  
    AUTOMATIC MEMORY SHARING OFF
```

If you deactivate automatic memory sharing from the "Memory Sharing" dialog box, that setting overrides the default ON setting in the NET.CFG file while the server is running. If the server is brought down and back up, the setting in the NET.CFG file again takes effect.


For more information about setting automatic memory sharing in the NET.CFG file, see "Automatic Memory Sharing" on page 826.

Freeing Base Server Memory

You can use PMMON to free the base server memory that NetWare Server for OS/2 allocates from OS/2. The memory is released back to OS/2 when the server is brought down.

You can also free the server memory by using the REMOVE BASE SERVER MEMORY option in the NET.CFG file (see "Remove Base Server Memory" on page 825).

Procedure



1. **Load PMMON.**
2. **Choose “Free Base Server Memory” from the “Configuration” menu.**

When “Free Base Server Memory” is active, a check mark appears next to it on the menu. To deactivate it, choose “Free Base Server Memory” again from the menu and the check mark will disappear.


Releasing Server Memory

You can use PMMON to release the base server memory immediately back to OS/2 if the server is not running.

For example, if BASE SERVER MEMORY in your NET.CFG file is set at 8, OS/2 will allocate an 8 MB block of memory to NetWare Server for OS/2 when the computer boots. If you do not run the server, OS/2 will not be able to use that 8 MB block of memory because it is reserved for NetWare Server for OS/2.

If you want OS/2 to be able to use this memory when the server is not running, use PMMON to release it.

Procedure



1. **Load PMMON when the server is not running.**

A message will appear indicating that the server is down.

2. **Choose “Cancel” to cancel the message from the screen.**

Do not choose “OK” to close the message because that will close PMMON.

3. **Choose “Release Server Memory” from the “Options” menu.**

A message appears that reads, “If you select OK, memory for the server will be given to OS/2.”

4. **Choose “OK” in the message box to release the memory back to OS/2.**

Supporting Workstations

NetWare Server for OS/2 supports all NetWare® clients.

If you are using the OS/2 High Performance File System (HPFS), you need to load the OS/2 and Macintosh name spaces in order for the server to support OS/2 and Macintosh clients. For information about loading name spaces, see “Setting Up a Volume to Store Non-DOS Files” on page 522.

If you run Extended Attributes on your OS/2 workstation, enter the following SET parameters at the server console:

SET maximum Extended Attributes per file or path = 32

**SET maximum percent of volume space allowed for
Extended Attributes = 10**

For information about SET parameters, see “SET” in *Utilities Reference*.

Configuring NetWare Server for OS/2

This section contains instructions for changing the configuration of NetWare Server for OS/2 using the NET.CFG file. It also includes an alphabetical listing of the options used for NetWare Server for OS/2 in the NET.CFG file.

The information in this section is also found in the configuration help screens in the NetWare Server for OS/2 installation program.

You may want to change the configuration of NetWare Server for OS/2 for any of the following reasons:

- ◆ To allow NetWare Server for OS/2 and NetWare Client for OS/2 to use the same network board
- ◆ To allocate a specific amount of memory to NetWare Server for OS/2
- ◆ To remove the memory that NetWare Server for OS/2 has allocated from OS/2 when the server is brought down

- ◆ To specify the proportion of processing time allocated to NetWare Server for OS/2
- ◆ To activate or deactivate automatic memory sharing between NetWare Server for OS/2 and OS/2

NET.CFG File Format

Modify the NetWare® Server for OS/2 configuration by entering options and settings in the NET.CFG file.

The NET.CFG file contains options and settings that define how NetWare Server for OS/2 and NetWare Client for OS/2 will function on the network.

Each NET.CFG option has settings. If you use the NetWare Server for OS/2 installation program to modify the NET.CFG file, these settings are alphabetized under the option to which they apply in the “NET.CFG Options” window.

If you use an OS/2 text editor to modify NET.CFG, see “NET.CFG Options Reference” in *NetWare Client for OS/2 User Guide* for more about each option and setting.

Conventions for Entering NET.CFG Options

You can create or edit the NET.CFG file through the NetWare Server for OS/2 installation program or by using an OS/2 text editor.

Use the following rules to create or modify the NET.CFG file. These rules are illustrated in Figure B-3.

- ◆ Type options at the left margin with no spaces before or after them.
- ◆ Type options in either uppercase or lowercase (they are not case-sensitive).
- ◆ Type one option per line.
- ◆ Type settings, one per line, below the option to which they apply.

- ◆ Use the Spacebar to indent these settings. Settings *must* be indented at least one space.

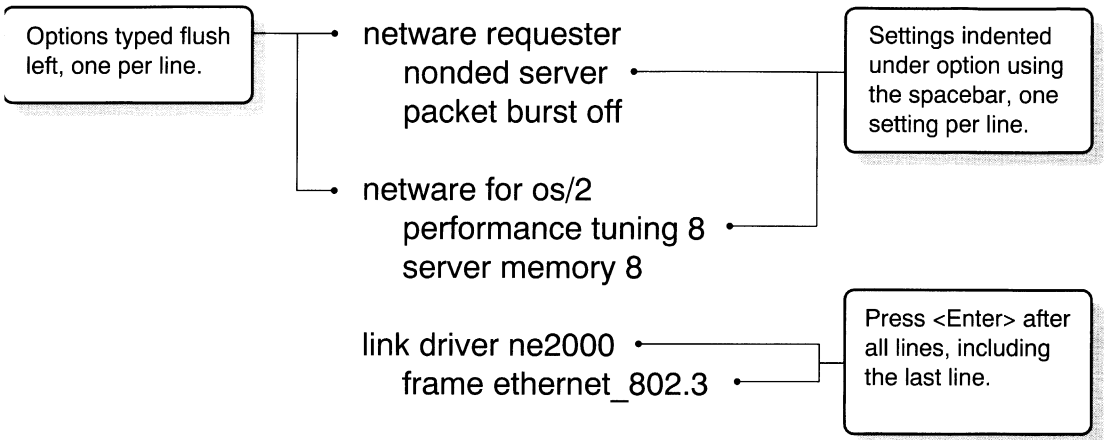


Note

Do not use <Tab> to indent the setting if you are using the graphical NetWare Server for OS/2 utility to edit the NET.CFG file. Pressing <Tab> activates different windows and buttons in the NET.CFG Options box.


- ◆ Press <Enter> at the end of every line in the file, *including the last line*. If you don't press <Enter> at the end of the last line, the line is ignored.
- ◆ Precede comments with a semicolon.

Figure B-3
Format of NET.CFG
File



Editing the NET.CFG File Using the Installation Program

Procedure



1. **Start the installation program by choosing the “Install” icon in the “Novell” group on the OS/2 desktop.**
2. **In the installation program, choose “This Workstation” from the “Configuration” menu.**

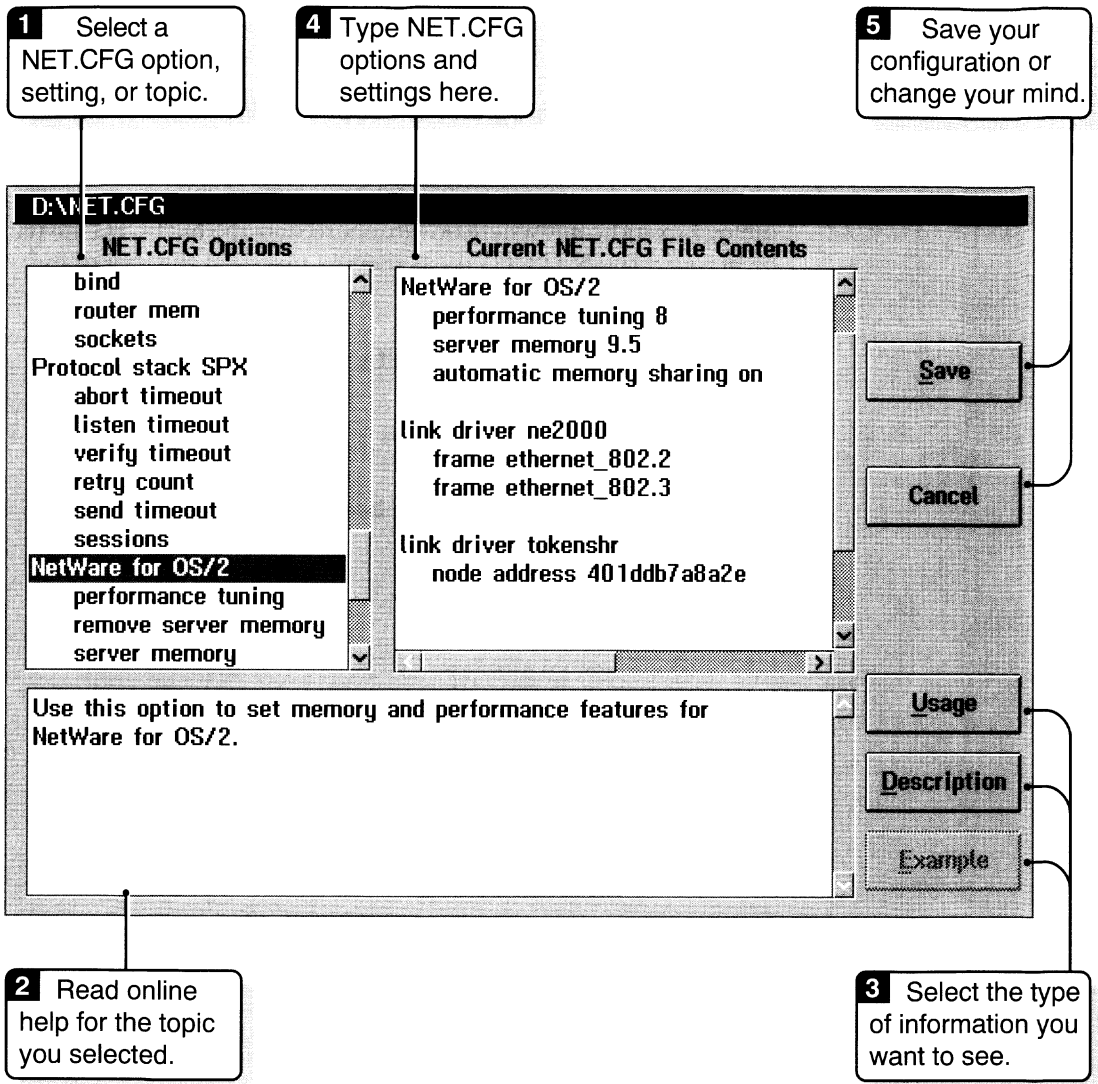
The “Default Location for NET.CFG File” dialog box appears. This dialog box indicates where on your hard disk the NET.CFG file will be placed. (The NET.CFG file is usually placed at the root of the OS/2 boot drive.)

3. **Choose “Edit” from the “Default Location for NET.CFG File” dialog box.**

A dialog box allowing you to edit the NET.CFG file appears.

4. **Use the “Edit NET.CFG” dialog box to add configuration options (see Figure B-4).**

Figure B-4
 “Edit NET.CFG”
 Dialog Box



Editing the NET.CFG File Using an OS/2 Text Editor

Procedure

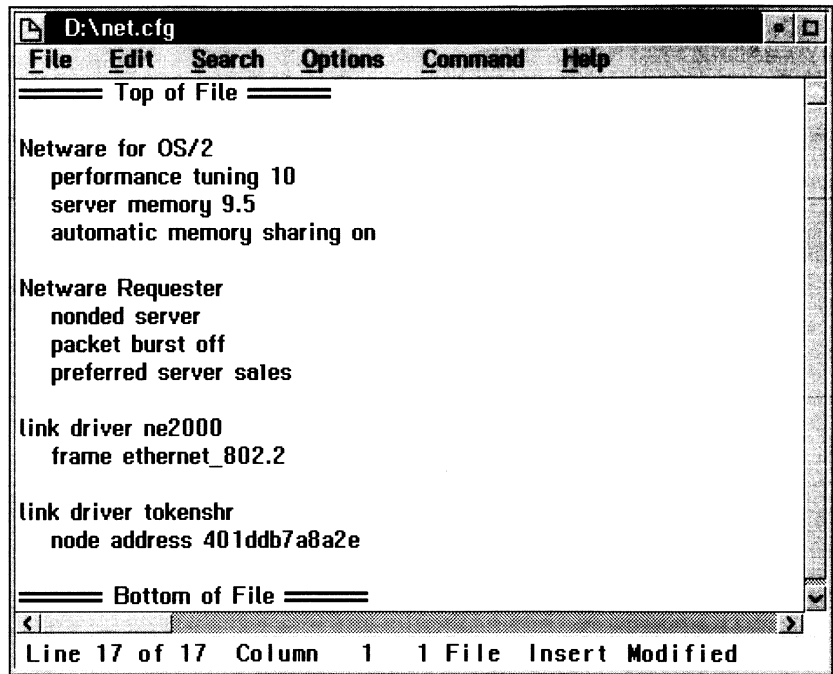


1. Open an OS/2 full screen or window session.
2. At the command line prompt, type

EPM NET.CFG <Enter>

An OS/2 text editor appears (see Figure B-5).

Figure B-5
OS/2 Text Editor



3. Use the editor to add the configuration options you need.

NET.CFG Options and Settings

The following sections explain the NET.CFG options and settings you can use to enhance NetWare Server for OS/2 performance. Each example includes a description of the option and setting, the default setting, and an example of the proper syntax for entering the option and setting in the NET.CFG file.

The NetWare for OS/2 Option

Use this option to set memory and performance features for NetWare Server for OS/2.

Syntax

```
NETWARE FOR OS/2  
PERFORMANCE TUNING number  
INITIALIZATION SCREEN DELAY number  
BASE SERVER MEMORY number  
REMOVE BASE SERVER MEMORY  
ALLOCATE MEMORY value  
AUTOMATIC MEMORY SHARING value  
MAXIMUM SERVER MEMORY number  
BLOCK ALLOCATION SIZE number
```

Settings

Performance Tuning

Use this setting to specify how much processing time is allocated to NetWare Server for OS/2. Whatever is not allocated to NetWare Server for OS/2 is left for OS/2.

Syntax **NETWARE FOR OS/2**
 PERFORMANCE TUNING *number*

Replace *number* with a number from 1 to 10.

The higher the number, the more processing time is allocated to NetWare Server for OS/2 and the faster your NetWare server will run.

Default 8

Example To set the performance tuning level to 7:

```
NETWARE FOR OS/2
PERFORMANCE TUNING 7
```

The default setting should be acceptable for your needs. However, you may want to lower the performance tuning level if OS/2 operation is too slow because of the amount of traffic on the server.

Initialization Screen Delay

Use this setting to specify the number of seconds you want the NetWare Server for OS/2 configuration options screen to remain displayed when OS/2 boots. (The screen displays information about device driver configuration. This screen cannot be used to change the configuration; it is for information only.)

Syntax **NETWARE FOR OS/2**
 INITIALIZATION SCREEN DELAY *number*

Replace *number* with a number expressed in seconds between 0 and 180.

Default 8

Example To set a 3-second delay:

```
NETWARE FOR OS/2
  INITIALIZATION SCREEN DELAY 3
```

Base Server Memory

Use this setting to allocate a base amount of memory for the server at boot time. If volumes on your server do not mount, increase the amount of base server memory using this setting.

Syntax **NETWARE FOR OS/2**
 BASE SERVER MEMORY *number*

Replace *number* with a number expressed in megabytes. Be sure the number you specify is large enough to run the server, but small enough to leave adequate memory for OS/2.

If NetWare Server for OS/2 cannot allocate the amount of memory you specify, it displays a warning message and uses the highest amount it can obtain.

Default 4

Example To allocate 8.5 MB of base server memory for NetWare Server for OS/2:

```
NETWARE FOR OS/2
  BASE SERVER MEMORY 8.5
```

Remove Base Server Memory

Use this setting to release the base server memory allocated to NetWare Server for OS/2 so the memory will be available for OS/2 to use. The memory is released when the server is brought down.

Syntax **NETWARE FOR OS/2**
REMOVE BASE SERVER MEMORY.

Default Base server memory is not released.

Example To release base server memory upon bringing down the server:

```
NETWARE FOR OS/2
  REMOVE BASE SERVER MEMORY
```

Allocate Memory

Use this setting to request that memory allocated from OS/2 to the server be allocated above or below the first 16 MB.

Syntax **NETWARE FOR OS/2**
ALLOCATE MEMORY *value*

Replace *value* with HIGH or LOW.

Default None.

Example To allocate the memory above 16 MB:

```
NETWARE FOR OS/2
  ALLOCATE MEMORY HIGH
```

To allocate the memory below 16 MB:

```
NETWARE FOR OS/2
  ALLOCATE MEMORY LOW
```

Automatic Memory Sharing

Use this setting to deactivate automatic memory sharing (it is active as a default). When automatic memory sharing is active, the system determines the proper amounts of memory for NetWare Server for OS/2 or OS/2 as the server is running.

Syntax **NETWARE FOR OS/2**
AUTOMATIC MEMORY SHARING *value*

Replace *value* with OFF.

Default ON

Example To deactivate automatic memory sharing:

```
NETWARE FOR OS/2
  AUTOMATIC MEMORY SHARING OFF
```


Maximum Server Memory

Use this setting to set the maximum amount of memory that NetWare Server for OS/2 can allocate from OS/2 when manual or automatic memory sharing is active. This prevents NetWare Server for OS/2 from taking too much memory from OS/2.

Syntax **NETWARE FOR OS/2**
 MAXIMUM SERVER MEMORY *number*

Replace *number* with a number expressed in megabytes.

Default None.

Example To set the maximum amount of memory that can be used by NetWare Server for OS/2 to 12 MB:

```
NETWARE FOR OS/2
MAXIMUM SERVER MEMORY 12
```

Block Allocation Size

Use this setting to specify the number of memory pages you want to include in a total block of memory that NetWare Server for OS/2 allocates from OS/2. One memory page is equal to 4 KB of memory.

For example, if you use the default of 8, NetWare Server for OS/2 allocates 8 memory pages in each block of memory. This means that the total size of the memory block is 32 KB (8 memory pages at 4 KB each = 32 KB).

When memory sharing is active, NetWare Server for OS/2 allocates blocks of memory from OS/2 in this size.

Syntax **NETWARE FOR OS/2**
 BLOCK ALLOCATION SIZE *number*

Replace *number* with a number between 4 and 1024.

Default 8

Example To set the Block Allocation Size so that NetWare Server for OS/2 allocates 200 KB in each memory block:

```
NETWARE FOR OS/2
  BLOCK ALLOCATION SIZE 50
```

(50 x 4 KB = 200 KB)

The NetWare Requester Option

Use this option to control network requests from your workstation to a NetWare server.



There are several settings that can be used with the NetWare Requester™ option. Only the ones that you may want to use with NetWare Server for OS/2 are described in this section.

For information about all the available NET.CFG options and settings, see “NET.CFG Options Reference” in *NetWare Client for OS/2 User Guide*.

Syntax

```
NETWARE REQUESTER  
NONDED SERVER  
PACKET BURST OFF
```

Settings

Nonded Server

Use this setting to allow NetWare Server for OS/2 and NetWare Client for OS/2 to use the same network board.

If you specify this setting, NetWare Server for OS/2 must be running on your computer before you can attach to the network using NetWare Client for OS/2 (otherwise, the client displays a connection error).

Syntax **NETWARE REQUESTER**
 NONDED SERVER

Default NetWare Server for OS/2 and NetWare Client for OS/2 use separate network boards.

Example To allow NetWare Client for OS/2 and NetWare Server for OS/2 to use the same board:

```
NETWARE REQUESTER
NONDED SERVER
```

Note: If your client software will not attach to a NetWare network, it may be because option has not been set.

Packet Burst Off

Use this setting to disable the Packet Burst feature. For more information about Packet Burst, see "Improving Speed by Using Packet Burst" in *NetWare Client for OS/2 User Guide*.

Syntax **NETWARE REQUESTER**
 PACKET BURST OFF

Default Packet Burst is enabled.

Example To disable Packet Burst:

```
NETWARE REQUESTER
PACKET BURST OFF
```



appendix

C

SFT III Management Tips

This appendix describes tips for managing and troubleshooting a NetWare® 4.1 SFT III™ network. The information is divided into these categories:

- Server Synchronization
- Mirrored Server Link (MSL)
- Server Configuration (.NCF) Files
- Server Memory
- Server Consoles
- Server Hard Disks
- Network Clients
- Network Performance
- Troubleshooting

The SET command is mentioned frequently in these sections. For more information on the SET command and its parameters, see “SET” in *Utilities Reference*.

Server Synchronization

This section discusses possible solutions to software failures related to SFT III server synchronization.

Detecting Server Synchronization Errors

For a comprehensive check of MEngine outputs on each server, turn on the Comprehensive MEngine Synchronization Check SET parameter. By default, a less intrusive check is performed. Using the comprehensive check may affect server performance.

Because this parameter must be set before the servers are activated (only settable at startup), put this command in the IOSTART.NCF file for each server:

```
SET Comprehensive MEngine Synchronization Check=ON
```

Handling Secondary Restarts Immediately After Synchronization

If the secondary server restarts immediately after synchronization, increase the IPX Internet Down Wait Time and the MSL Deadlock Detect Wait Time SET parameters, using the following syntax:

```
SET IPX Internet Down Wait Time=variable  
SET MSL Deadlock Detect Wait Time=variable
```

If this does not solve the problem, halt the server by changing all mirrored servers' error recovery options to 0. See "Servers Restart for No Apparent Reason" on page 857 for specific options. After halting the server, contact your support representative.

Notifying Users of Server Synchronization

To send a broadcast message to all logged-in network users when the mirrored servers begin synchronizing, type the following.

```
SET Notify All Users Of Mirrored Server  
Synchronization=ON
```

The default setting for this parameter is OFF.

Reducing the Time for Resynchronizing and Remirroring

To shorten server synchronization time, reduce the amount of memory that must be synchronized, using the following syntax:

```
SET New End Address For Unclaimed Memory  
Block=variable
```



Using this SET parameter to reduce the amount of unclaimed memory may prevent some loadable modules from loading in the MSEngine. If the MSEngine memory is too small, some volumes may not mount.

If you have disk arrays, you may be able to speed up disk mirroring with the Concurrent Remirror Requests SET parameter.

Responding to “Invalid Mirrored Server Initialization” Messages

If this message appears when the servers are in test mode, no action is necessary as long as the servers continue to synchronize.

To prevent server initialization problems, increase the Restart Minimum Delay Amount, using the following syntax:

```
SET Restart Minimum Delay Amount=variable
```

Also, verify that the MSL boards and cables are functional. Then unload and reload the MSL driver on each server.

If you are using NE2000™ boards as MSL boards, make sure that the interrupt settings on the MSL boards have a higher priority than the network boards, and that the MSL boards are cabled correctly.



Note

Interrupt priorities from high to low are 0, 1, 2 or 9, A, B, C, D, E, F, 3, 4, 5, 6, 7, and 8. Interrupt 2 is the highest priority you can assign to an MSL board because interrupts 0 and 1 are reserved.

Responding to “Mirrored Server Engine Already Loaded” Messages

This message means you have already executed the `ACTIVATE SERVER` console command, either at the IOEngine console or in the `IOSTART.NCF` file of the other server. Delete “`ACTIVATE SERVER`” from the `IOSTART.NCF` file.



Important

Putting `ACTIVATE SERVER` in the `IOSTART.NCF` file may cause synchronization problems. If both servers execute `ACTIVATE SERVER` simultaneously, the utility may load two separate, unsynchronized MEngines, and both servers would assume the primary server role.

Mirrored Server Link (MSL)

This section discusses possible solutions to hardware and software failures related to the high-speed cable links between the mirrored SFT III™ servers.

Determining Appropriate MSL Cable Lengths

The maximum cable length between MSL boards is determined by the cable manufacturer. Ranges are from 30 to 100 meters for coaxial cable, and from 1 to 40 kilometers for fiber optic cable.

Handling MSL Cable or MSL Board Failure

When an MSL cable or board fails, the secondary server restarts but cannot synchronize with the primary server until the MSL problem is corrected.

To prevent future MSL problems and loss of server mirroring, install redundant, alternate MSL boards and cabling in each server. The alternate MSL driver must be loaded before it can take over for an active MSL that fails.

Put the LOAD command for the alternate MSL driver in the IOSTART.NCF files of both servers, after the LOAD command for the first MSL driver. Use the following syntax:

```
LOAD alternate_driver_name
```

The order the drivers are loaded determines which MSL is the default and which is the first alternate, second alternate, and so forth.

Responding to MSL Communications Error Messages

Check the MSL cable and board connections to make sure they are properly connected and that the boards are firmly seated. Also check for kinks or damage in the cabling.

Responding to “MSL Deadlock Delivering Data” Messages

This message means that neither MSL can transmit data because of a “holdoff.” A holdoff occurs when the IOEngine receives a packet but cannot process it, or when one MSL sends a packet but doesn’t receive an acknowledgment.

Increase the MSL Deadlock Detect Wait Time, using the following syntax:

```
SET MSL Deadlock Detect Wait Time=variable
```

If the servers are very busy, increase this parameter in 1-second increments until the error disappears.



Make sure the value of the MSL Deadlock Detect Wait Time parameter is at least 1 second longer than the IPX Internet Down Wait Time parameter.

Server Configuration (.NCF) Files

SFT III reads from these server configuration files, in the following order, when you power on or restart the servers:

- ◆ IOSTART.NCF (two files—one for each server)
- ◆ MSSTART.NCF
- ◆ MSAUTO.NCF
- ◆ IOAUTO.NCF (two files—one for each server)

Most of these files are created during the installation process. Use **INSTALL** to create or edit server configuration files. You can also use **EDIT** to edit .NCF files.

Example: IOSTART.NCF file (on DOS partition)

```
ioengine name SFT3_I01
ioengine ipx internal net 7654321
load isadisk port=1f0 int=e
load nmsl
```

Example: MSSTART.NCF file (on DOS partition)

```
set Concurrent Remirror Requests=11
```

Example: MSAUTO.NCF file (on volume SYS:)

```
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=0=Novell
msengine name SFT3
msengine ipx internal net 1234567
mount all
```

Example: IOAUTO.NCF file (on DOS partition or volume SYS:)

```
sys:etc\iol\initsys.ncf
load tcpip
```



Note To edit INITSYS.NCF, see "INETCFG" in *Utilities Reference*.

Server Memory


This section discusses how to install additional server memory without disrupting the network and how to respond to the message "Secondary server is missing RAM."

Adding Memory Without Bringing the MEngine Down

A NetWare 4.1 SFT III system allows you to upgrade server hardware without loss of service to clients.

Use the procedure below to add memory to your servers without bringing down the MEngine.

Procedure



Procedure

1. **Halt the secondary server and turn it off.**

The primary server is still running.

2. **Add memory to the secondary server and turn it on.**

3. **Reconfigure the hardware, using the hardware-specific configuration procedure specified by the manufacturer.**

4. **From the DOS prompt for drive C: on the secondary server, type**

```
MSERVER <Enter>
```

5. Wait for resynchronization to complete.

During resynchronization, the primary server does not recognize the additional memory in the secondary server.

6. After the disks are remirrored, halt the primary server and turn it off.

The secondary server becomes the new primary server.

7. Add memory to the new secondary server and turn it on.

8. Reconfigure the hardware, using the hardware-specific configuration procedure specified by the manufacturer.



To ensure synchronization, install the same amount of memory in both SFT III servers.

9. From the DOS prompt for drive C: on the secondary server, type

MSERVER <Enter>

10. Wait for resynchronization to complete.

During synchronization, the primary server now recognizes the additional memory in the secondary server.

Responding to “Secondary Server Is Missing RAM” Messages


This message means that the primary server cannot synchronize memory with the secondary server for one of the following reasons:

- ◆ The primary server has more RAM than the secondary server.
- ◆ The primary and secondary have the same amount of RAM installed, but the memory is noncontiguous and the memory “holes” don’t match.

In this situation, do one of the following:

- ◆ Add memory to the secondary server so that its RAM is equal to the RAM in the primary server.
- ◆ Run the EISA configuration utility (if applicable) so that both servers are in the same mode (linear or compatible mode).
- ◆ Align the server memory holes using the procedure below.

Procedure



1. **Change to the IOEngine prompt on the secondary server.**
2. **Check the memory addresses by typing**

MEMORY MAP <Enter>

A display similar to the following appears.

System memory map:

0-12288 (DOS)

12288-159744 (DOS)

159744-654336 (IOEngine)

1048576-5767168 (IOEngine)

5767168-17170432 (Unclaimed)

3. **From the “Unclaimed” range on the last line of the display, write down the “start” and “end” memory addresses.**



Unclaimed memory is used by the MEngine, so it must be identical on both servers.

4. **From the IOEngine console on the *primary* server, type**

HALT <Enter>

5. Add the following commands to the primary server's IOSTART.NCF file:

```
SET New Start Address for Unclaimed Memory Block=x  
SET New End Address for Unclaimed Memory Block=x
```

The values following the = (equals) sign are the start and end of the unclaimed memory range on the secondary server (the numbers you wrote down in Step 3).

6. Restart the primary server.

Now that the unclaimed memory is aligned on both servers, they can synchronize. SFT III allocates any extra memory to the IOEngine on that server.

Server Consoles

This section discusses solutions for keeping track of the three NetWare 4.1 SFT III server consoles—the two IOEngines and the MEngine—and for managing errors displayed at these consoles.

Logging and Viewing All Console Messages

Some SFT III console messages are never seen because you can't view all three consoles (two IOEngines and one MEngine) at once.

To keep track of these messages, use CONLOG to capture all console messages and to write them to SYS:\ETC\CONSOLE.LOG.

To log messages from all three consoles, add the line LOAD CONLOG to the IOSTART.NCF files of both servers and to the MSAUTO.NCF file.



The LOAD CONLOG command must be the first line in each .NCF file from which you want to log messages.

To view the CONSOLE.LOG file, use EDIT or INETCFG.

Handling Lost Interrupt Alerts on the Console

If lost interrupt alerts are filling up the console screen, at the appropriate IOEngine console, type

```
SET Display Lost Interrupt Alerts=OFF <Enter>
```

A lost interrupt message indicates a problem driver or faulty hardware. To find the driver with the interrupt problem, turn on the Display Lost Interrupt Alerts parameter as shown above. Then unload all the drivers from the appropriate IOEngine, and reload them one at a time. Contact the vendor of the problem driver.

Responding to “Loader Cannot Find...” Messages

This message means you tried to load an NLM in the wrong engine (IOEngine or MEngine), or that the NLM you tried to load depends on other NLMs that aren't loaded.

Use the <Alt>+<Esc> keys to toggle to the other engine's console, and attempt to load the NLM again.

If the message still appears, you may have loaded an outdated or unsupported NLM.

Shortening the Console Prompt

Long SET parameter names may scroll behind the console prompt if you have a long IOEngine or MEngine name in the prompt. However, scrolling does not affect the execution of the SET utility.

If you want a two-character console prompt (IO: or MS:), type the following at the appropriate console:

```
SET Replace Console Prompt With Server Name=OFF <Enter>
```

Server Hard Disks

This section discusses solutions for hard disk errors and storage shortages on a NetWare 4.1 SFT III server.

Controlling the Size of Log Files on Volume SYS:

To limit the size of SFT III log files, IO\$LOG.ERR and MSSTATUS.DMP, use the following commands:

```
SET IOEngine Error Log File Overflow Size= number  
SET Status Dump File Overflow Size= number
```

Replace *number* with the size limit in bytes.


The system will delete or rename the log files when they meet or exceed the size limit, depending on the SET IOEngine Error Log File State and SET Status Dump File State parameter values.

Handling Primary Hard Disk Failure

The primary server remains active if a primary hard disk fails.


In this situation, clients continue to communicate with the primary server, but disk requests use only the secondary server's disk (instead of splitting the seeks between the servers).

Use the procedure below to restore fault tolerance to your server storage.



Make sure the failed disk is on the primary server and that the secondary server's disk is functional before attempting this procedure, or you may lose data.

Procedure



1. Change to the primary server's IOEngine console.
2. Force a server switchover by typing


HALT <Enter>

The secondary server takes over the primary server role.

3. Correct the hard disk problems and resynchronize the servers.
4. Use the Server Failure Notification Name SET parameter to notify you immediately in case of another server disk failure. Add this parameter to the MSAUTO.NCF file:

```
SET Server Failure Notification Name=group_name |  
user_name
```

Mounting a CD-ROM as a NetWare 4.1 SFT III Volume



Access to data on a CD-ROM volume will be lost if the SFT III servers switch from primary to secondary unless you mount the same CD on both servers.

To mount a CD-ROM disc as a NetWare 4.1 SFT III volume, follow these steps.

Prerequisites




- A mounted volume SYS:
- An installed Host Bus Adapter (HBA) that is NetWare compatible and supports CD-ROM devices
- The NPAIO.DSK and NPAMS.NLM NetWare Peripheral Architecture (NPA) modules
- The disk driver files and necessary support modules for the CD-ROM device



Some disk drivers consist of more than one file and some HBA devices require additional support modules for CD-ROM functionality. These files should accompany the HBA. For specific file requirements, consult your adapter documentation.

Procedure



1. Change to the IOEngine prompt of the SFT III server where the HBA is installed.

2. Load the disk driver by typing

```
LOAD [path]disk_driver <Enter>
```

Replace *disk_driver* with the name of the disk driver specified in the HBA documentation.

For example, to load the disk driver for the Adaptec AHA-1522 SCSI HBA, type

```
LOAD [path]AHA1520.DSK <Enter>
```

```
LOAD [path]ASPICD.DSK <Enter>
```

3. Load CDROM.NLM by typing

```
LOAD CDROM <Enter>
```

This auto-loads the NPAMS module



When a CD-ROM is mounted or a CD-ROM disc is changed, some CD-ROM devices may be deactivated. This deactivation occurs because device configuration information is being updated.

4. View the device number and volume name by typing

```
CD DEVICE LIST <Enter>
```

5. Mount the CD-ROM as a volume by typing

```
CD MOUNT [device number] | [volume name] <Enter>
```

Replace *device number* with the device number or replace *volume name* with the volume name of the CD-ROM disc.

For example, to mount the NetWare_41 CD-ROM, type

```
CD MOUNT NETWARE_41 <Enter>
```



It may take several minutes to mount the volume the first time, depending on the size of the CD-ROM and the speed of your computer.

The standard volume mount messages appear.

6. To mount the CD-ROM as a NetWare volume each time the server comes up, edit the IOSTART.NCF file of the SFT III server where the HBA is installed. Add these commands to the IOSTART.NCF file:

```
LOAD [path]disk_driver
```

7. Edit the MSAUTO.NCF file, adding these commands:

```
LOAD CDROM  
CD MOUNT [device number] | [volume name]
```

Notifying Users of Disk or Server Failure

To notify a user or group of users about a disk or server failure, use the Server Failure Notification Name SET parameter in the MEngine. Also put this parameter in the MSAUTO.NCF file. For example:


```
SET SERVER FAILURE NOTIFICATION NAME = ADMIN
```

When a failure happens, the specified user or group receives a broadcast message. To send this broadcast message to all logged-in users, use group EVERYONE.

Recovering an Orphaned Partition

Use the following procedure to recover an orphaned partition:

Procedure



1. Load INSTALL.
2. From the "Installation Options" menu, choose "Disk Options."
3. From the "Available Disk Options" menu, choose "Mirror/Unmirror."

The "Partition Mirroring Status" appears, showing one mirrored and one out of sync. The out-of-sync partition is the "orphan."

4. Choose the mirrored partition.

The "Mirrored Disk Partitions" menu displays the partition number and device number of each mirrored set. The device number for the out-of-sync partition is an "unavailable device."

5. Remove the partition from the set by highlighting the out-of-sync partition and pressing <Delete>.

6. Return to "Disk Partition Mirroring Status" by pressing <Esc>.

"Not mirrored" appears by the good partition, and "Out of sync" appears by the other partition. Note the number of the "not mirrored" partition.

7. Highlight the out-of-sync partition and press <F3>.

A warning message similar to the following appears:

```
Warning!! The selected partition contains "Volume
SYS Segment 0" and that volume is already defined.
```

8. Select the "No Salvage" option.

9. Select the good partition (as noted in Step 6).

10. Press <Insert> for a list of the available partitions.

11. Highlight the previously orphaned partition and press <Enter>.

12. Continue by pressing <Esc>.

The unavailable partition is deleted. After a brief delay, the remirroring process begins. Because the entire partition is remirrored, the process takes several minutes or hours, depending on the partition size. You can check remirroring status with the MIRROR STATUS command or by watching the install screen.

Network Clients

This section discusses solutions for workstation problems that may be encountered on NetWare 4.1 SFT III networks.

Finding a Server from an Ethernet Workstation

Some Ethernet clients may have problems finding or attaching to a server.

If you use the 802.3 frame type, change the value of the Enable IPX Checksums SET parameter to 0 in the MSEngine by typing the following:

```
SET Enable IPX Checksums = 0 <Enter>
```

The default setting is 0. Also, make sure volume SYS: is mounted.

Finding the SFT III Server First

If another NetWare server is on the same network segment as an SFT III server, clients will not find the SFT III server first.

Because the IOEngine routes the get-nearest-server request to the MSEngine, the SFT III server appears to be two hops away while the other NetWare server appears to be closer.

To find the SFT III server first, include the following statement in the client's NET.CFG file:

```
Preferred Server=<MSEngine_name>
```

Handling DOS IPX Session Timeouts During Server Switchover

If DOS IPX™ clients are timing out or losing connections when the servers switch over, increase the IPX retry count parameter in the NET.CFG file on clients using IPX. For example:

```
ipx retry count 40
```

Logging in ARCnet or Token Ring Clients Immediately After a Halt

Some clients on token-passing networks may have problems logging in after a server has been halted. Wait for the client to time out; then try to log in again.

Responding to “LAN Driver Loopback Error Detected” Messages

This message indicates that two or more ARCnet boards have the same node address. Reconfigure each board to a unique address.

Turning Off the IPX Checksum Option for Token Ring

Token ring doesn't support the enabling of IPX checksums. If you are on a token ring network, change the value of the Enable IPX Checksums SET parameter in the MSEngine to 0 by typing the following:

```
SET Enable IPX Checksums = 0 <Enter>
```

Network Performance

This section discusses methods of optimizing network performance.

Handling Frequent Ups and Downs of the IPX Internet

This situation indicates a faulty network connection or a problem with communication between the servers.

Use SET to increase the IPX Internet Down Wait Time and the MSL Deadlock Detect Wait Time parameters by 0.5 second or more.

```
SET IPX Internet Down Wait Time = variable
```

```
SET MSL deadlock Wait Time = variable
```

The value of *variable* should be 0.5 second greater than the current value. If this solves the problem, put the SET commands in the IOSTART.NCF files. If it does not solve the problem, increase *variable* again.

Also, check the network cabling attached to each server and the network connections (including routers and bridges) between the servers.

Make sure the interrupt priority of the MSL board is higher than the interrupt priority of the network boards in each NetWare server.

In most cases, this means setting the MSL interrupt to a lower number. (See the MSL board manufacturer's documentation for more information on setting interrupts.)



Interrupt priorities from high to low are 0, 1, 2 or 9, A, B, C, D, E, F, 3, 4, 5, 6, 7, and 8. Interrupt 2 is the highest priority you can assign to an MSL board because interrupts 0 and 1 are reserved.

Responding to “IPX Network No Longer Returning Status” Messages

This message indicates a problem with a network connection: either a faulty network board in one of the servers, a cabling problem in the network, or an incorrectly bound network protocol.

If this server is the primary server, and the system determines that the other server has more functional network boards, that server will become the primary.

To determine if a server's network board has failed, type the CONFIG command at each server's IOEngine prompt.

A “not sending” or “not receiving” message indicates a bad board or a faulty network connection.

If this message frequently recurs with the message, “IPX Network is now returning status check packets,” increase the IPX Internet Down Wait Time SET parameter to give the system a little more time before determining that a network board has failed. Use the following syntax:

```
SET IPX Internet Down Wait Time = variable
```

Handling Primary Network Board Failure

If a network board in the primary server stops transmitting or receiving packets, the secondary server assumes the primary role *only* if the following conditions are true:

- ◆ The secondary server's network boards are more functional than the primary's network boards.
- ◆ The servers are fully synchronized and their disks are completely mirrored.

Three SET parameters in the IOEngine can help detect and prevent downtime caused by network board failure: Check LAN Option, Check LAN Extra Wait Time, and Use Diagnostic Responder to Validate LAN Functionality.

SET Check LAN Option=2

This setting (2 is the default) forces a server switchover by restarting the primary server if a network board fails in the primary.

SET Check LAN Extra Wait Time=10

This setting adds 10 extra seconds to the time the system waits before forcing a switchover because of a bad network board in the primary server.

By default this setting is 0, but if you have a large network or heavy traffic on the network, you may want to increase the wait time to prevent a premature server switchover.

SET Use Diagnostic Responder to Validate LAN Functionality=ON

This setting broadcasts an IPX diagnostic request to verify that a network board is functional.

By default this setting is OFF because the diagnostic request adds traffic overhead and can hurt performance of large networks.

However, if you want to know whether a network board is bad or just slow, set this parameter to ON.

Troubleshooting

This section suggests solutions to various problems that can occur with NetWare 4.1 SFT III networks.

Both Servers Are Primary; No Secondary Console Display

ACTIVATE SERVER was executed from both IOEngines, and both servers have assumed the primary server role. Delete "ACTIVATE SERVER" from the IOSTART.NCF file. Restart one server.

Make sure the node addresses on the network boards are unique to each server.

Disks Won't Mirror After Servers Synchronize

Verify that the correct disk driver is loaded in each server's IOEngine. Also, put the load command for the disk driver in each server's IOSTART.NCF file.

Check the mirror status in the INSTALL.NLM for orphaned drives.

"Error IPX internet may be too slow" Message Appears

A busy IPX network or busy routers between SFT III servers can cause this message:

```
IPX internet may be too slow to notify secondary
server if MSL fails... increase secondary take over
delay amount.
```

If "I'm alive" packets take too long to travel between the servers over IPX, the secondary server may prematurely take over for the primary server.

To prevent this, use the SET command in the IOEngine to increase the Secondary Take Over Wait Time. Use the following syntax.

```
SET Secondary Take Over Wait Time = variable
```

Also check for problems with network connection that may be slowing IPX.

“Error transferring IOEngine error log to MEngine” Message Appears

Volume SYS: may not have mounted. If the volume is mounted, check the file attributes of the IO\$LOG.ERR file to verify the file isn't flagged as Read Only.

If the problem persists, there may not be room for the file on volume SYS:. Delete or rename the IO\$LOG.ERR file on the DOS partition.

“Inactive device associated with mirror partition” Message Appears

This message is caused by one of the following situations:

- ◆ The MEngine was brought down, but only one server was brought back up.
- ◆ The disk driver on one server didn't load correctly.
- ◆ A hard disk or controller failed.

Correct any hardware problems associated with the failure and bring both servers back up.

IOEngine Network Number Prompt Appears After Executing IOSTART.NCF

Check both IOSTART.NCF files to make sure they assign unique internal network numbers to each IOEngine.

Check the spelling and syntax of the IPX internal network number command. There should *not* be an equal sign (=) in the command.

Make sure the IOSTART.NCF file is in the same directory as the MSERVER.EXE file on each server.

Keyboard Is Slow or Frozen After Loading ARCnet on IRQ 2

Make sure there are no I/O port or memory address conflicts.

LOGIN Fails, Even If Correct Password Is Given

If checksumming is disabled at the workstation, it must also be disabled at the server. Change the value of the Enable IPX Checksums SET parameter in the MEngine to 0, using the following syntax:

```
SET Enable IPX Checksums = 0
```

This may also indicate a problem with the user's Directory context.

LOGIN Fails for an NLM in the IOEngine

This happens only in the case of a backup NLM that logs in to the server using a different user name and would happen only on the secondary IOEngine. Type the following.

```
SET Reply To Get Nearest Server = ON <Enter>
```

MEngine Name Prompt Appears Even Though MSAUTO.NCF Exists

Volume SYS: did not mount, or the command was typed incorrectly in the MSAUTO.NCF file.

There should *not* be an equal sign (=) in the MEngine name command.

Type in the MEngine name and IPX internal network number. Edit the MSAUTO.NCF file if necessary.

Make sure the disk drivers are loaded in both IOEngines. Mount volume SYS: from the command line or with INSTALL. If volume SYS: is unable to mount, run VREPAIR and try mounting it again.

MSL Drivers Are Loaded But Servers Do Not Synchronize

You might have installed the MSL boards incorrectly. Check the following on each server:

- ◆ Interrupt settings on the MSL boards (for conflicts with the network boards)
- ◆ MSL cables and board connections
- ◆ MSL drivers (both servers must have the same version of the driver)

MSL Isn't Activated

Verify that the same MSL driver is loaded on each server and that the MSL boards and cables are installed correctly.

MSL Times Out When a Device Driver is Loaded

Load the device driver before loading the MSL driver, or use the SET utility in the IOEngine to increase the MSL Error Wait Time value. Use the following syntax:

```
SET MSL Error Wait Time = variable
```

If this solves the problem, put the SET command shown above in the IOSTART.NCF file, above the LOAD command for the MSL driver.

“Other Server Requested This Server to Halt” Message Appears

The following sequence of events causes this message to appear:

1. The secondary server stops receiving “I’m alive” packets from the primary server over the IPX network.
2. The MSL between the two servers fails.
3. The secondary server begins to take over as primary (because of no “I’m alive” packets).

4. The secondary server then receives an "I'm alive" packet from the primary server.
5. Because two primary servers can't coexist, the secondary server halts.

Check for hardware problems on the IPX network and the MSL connection. If there are no hardware problems, use SET in the IOEngine to increase the Secondary Take Over Wait Time, using the following syntax:

```
SET Secondary Take Over Wait Time = variable
```

Primary Server Doesn't Know IPX Route to Secondary Server

Verify that

- ◆ Both IOEngines are active
- ◆ The LAN drivers are loaded in both IOEngines
- ◆ The protocols are bound to the network boards in both servers

Use SET to increase the IPX Internet Down Wait Time and the MSL Deadlock Detect Wait Time parameters, using the following syntax:

```
SET IPX Internet Down Wait Time = variable  
SET MSL Deadlock Detect Wait Time = variable
```

Check the network cabling and the network boards and routers between the SFT III servers to verify they are functioning properly.

Queue Overrun Abend Occurs

Decrease the number of seconds in the IPX Internet Down Wait Time and the MSL Deadlock Wait Time SET parameters in the IOEngine, using the following syntax:

```
SET IPX Internet Down Wait Time = variable  
SET MSL Deadlock Detect Wait Time = variable
```


Secondary Doesn't Switch Over When Primary LAN Driver Is Unloaded

The secondary server is designed to take over for the primary server if the secondary server detects a network board failure in the primary.

However, an unloaded LAN driver is not considered a hardware failure, since it was explicitly requested by a user at the server console.

Therefore, if a LAN driver in the primary server is unloaded, both the primary and the secondary server will report they aren't receiving "I'm alive" packets, but the secondary server will not switch over.

When the LAN driver is reloaded, the primary server will continue to function as primary.



Any activity in progress when the LAN driver is unloaded from the primary server will be suspended, and client connections will time out unless the driver is promptly reloaded.

Server Failure Occurs

Because the MSEngines on both SFT III servers are mirrored and the IOEngines are not, you can assume the following about most server failures:

- ◆ If both servers fail or have an abend condition, the failure is probably related to software running in the MEngine.
- ◆ If one server fails and the remaining server assumes the role of primary server, the failure is probably related to hardware, or to a driver or NLM loaded in the failed server's IOEngine.

Servers Restart for No Apparent Reason

You may have the wrong values for the recovery option SET parameters, or server test mode is causing constant switchovers. Check the IO\$LOG.ERR and MSSTATUS.DMP files for the cause of the problem.

Use these SET parameter values in *both* IOEngines to halt the server (without restarting it) so you can see the error:

```
Test Mode=0  
Secondary Server MSL Send Blocked Recovery Option=0  
Primary Server MSL Send Blocked Recovery Option=0  
MSEngine Abend And Processor Exception Recovery  
Option=0  
IOEngine Abend And Processor Exception Recovery  
Option=0  
Machine Check Recovery Option=0  
Memory Parity Error Recovery Option=0  
Secondary Server MSL Hardware Failure Recovery  
Option=0  
Primary Server MSL Hardware Failure Recovery Option=0  
MSEngine Outputs Different Recovery Option=1  
Secondary Server MSL Consistency Error Recovery  
Option=0  
Primary Server MSL Consistency Error Recovery  
Option=0  
Secondary Server MSL Deadlock Recovery Option=0  
Primary Server MSL Deadlock Recovery Option=0
```

SFT III Error Log Files

When a failure occurs, SFT III updates three error log files in the SYS:SYSTEM directory:

- ◆ IO\$LOG.ERR records the activity of both IOEngines.
- ◆ SYS\$LOG.ERR records the MSEngine activity.
- ◆ MSSTATUS.DMP records status dumps of engine states, synchronization and communications states, IOEngine to MSEngine requests, and other information following a failure or server switchover.

Use these error log files to track the events that occurred prior to a failure or following a switchover.



Note

The IO\$LOG.ERR file on the failed server is written to its DOS partition until the servers come back up. Then, the IO\$LOG.ERR file from the DOS partition is appended to the IO\$LOG.ERR file on volume SYS:.

“Should This Machine Become the Primary Server?” Message Appears

The message above appears on the secondary server’s console preceded by

```
All communication channels with the primary server
have failed. Since the IPX network communication
channel failed before the mirrored server link
failed, the secondary is unable to determine if the
primary server is still active.
```

Verify that the primary server has failed, and then type “Y.” If the primary server is still active, type “N.”

Test Mode Is No Longer Working

You can set the Test Mode parameter from the command line. For example

```
SET Test Mode = variable
```

However, Test Mode resets to the default (no test) the first time that server is automatically downed and rebooted.

To keep the server in test mode, put the SET command line shown above, using the appropriate parameter, in the IOSTART.NCF file of both SFT III servers.

If your drives are not mirrored, the primary server will not initiate test mode. In this case use INSTALL to remirror the drives.

“Unknown Command” Message Appears

Use the <Alt>+<Esc> keys to toggle to the other engine’s console; then execute the command again.

Use the LOAD command to execute modules (such as INSTALL and MONITOR).

“Unknown SET Parameter Name” Message Appears

Use the <Alt>+<Esc> keys to toggle to the other engine’s console; then retype the SET command.


Check the spelling and syntax of the SET parameter and retype it correctly.

Volume SYS: Does Not Mount

Load VREPAIR. If VREPAIR does not load, halt the primary server by typing

HALT <Enter>

Execute MSERVER with the -ns parameter. Then reload the disk driver and load VREPAIR.



When VREPAIR is loaded, Option 2 must be set to "Write all directory and FAT entries out to disk" and Option 3 must be set to "Write changes immediately to disk."

Make sure you are using the latest disk driver.

Workstations Cannot Find or Attach to a Server

Make sure that

- ◆ Volume SYS: is mounted
- ◆ The appropriate network drivers are loaded on each SFT III server
- ◆ NetWare Directory Services™ (DS.NLM) is loaded
- ◆ The LAN protocols are bound to the network boards
- ◆ The workstation and server are using the same frame type

Workstations Time Out During Server Resynchronization

Change the IPX retry count in the NET.CFG file on the workstations to a higher value. For example:

ipx retry count 40

Turn on the SET parameter Notify All Users Of Mirrored Server Synchronization in the MEngine to broadcast a message to all users when synchronization occurs. Use the following syntax.

**SET Notify All Users Of Mirrored Server
Synchronization = <group name>**



Trademarks

Novell, Inc. has attempted 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 Trademarks

DR DOS is a registered trademark of Novell, Inc.

Hot Fix is a trademark of Novell, Inc.

Internetwork Packet Exchange and IPX are trademarks of Novell, Inc.

LANalyzer is a registered trademark of Novell, Inc.

Media Support Module and MSM are trademarks of Novell, Inc.

NE/2 is a trademark of Novell, Inc.

NE/2-32 is a trademark of Novell, Inc.

NE1000 is a trademark of Novell, Inc.

NE2000 is a trademark of Novell, Inc.

NE2100 is a trademark of Novell, Inc.

NE3200 is a trademark of Novell, Inc.

NE32HUB is a trademark of Novell, Inc.

NetWare 3 is a trademark of Novell, Inc.

NetWare 4 is a trademark of Novell, Inc.

NetWare 4.1 SFT III is a trademark of Novell, Inc.

NetWare Application Notes is a trademark of Novell, Inc.

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

NetWare Directory Services and NDS are trademarks of Novell, Inc.

NetWare DOS Requester and NDR are trademarks of Novell, Inc.

NetWare Link Services Protocol and NLSP are trademarks of Novell, Inc.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

NetWare Message Handling Service and NetWare MHS are trademarks of Novell, Inc.

NetWare Peripheral Architecture and NPA are trademarks of Novell, Inc.

NetWare Requester is a trademark of Novell, Inc.

NetWire is a service mark of Novell, Inc.

Novell Authorized Reseller is a collective mark of Novell, Inc.

Novell Labs is a trademark of Novell, Inc.

Personal NetWare is a trademark of Novell, Inc.

RX-Net is a trademark of Novell, Inc.

RX-Net/2 is a trademark of Novell, Inc.

SFT (System Fault Tolerant) and SFT III are trademarks of Novell, Inc.

SPX is a trademark of Novell, Inc.

Storage Management Services and SMS are trademarks of Novell, Inc.

Transaction Tracking System and TTS are trademarks of Novell, Inc.

UNIX is a registered trademark of Novell, Inc. in the United States and other countries, licensed excessively through X/Open Company, Ltd.

WordPerfect is a registered trademark of Novell, Inc.

Third-party trademarks

Apple and AppleTalk are registered trademarks of Apple Computer, Inc.

ARCnet is a registered trademark of Datapoint Corporation.

AT&T is a registered trademark of American Telephone & Telegraph.

COMPAQ is a registered trademark of Compaq Computer Corporation.

Hayes is a registered trademark of Hayes Microcomputer Products, Inc.

Hewlett Packard is a registered trademark of Hewlett-Packard Company

IBM is a registered trademark of International Business Machines Corporation.

Lotus is a registered trademark of Lotus Development Corporation.

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.

NFS is a registered trademark of Sun Microsystems, Inc.

OS/2 is a registered trademark of International Business Machines Corporation.

PS/2 is a registered trademark of International Business Machines Corporation.

Unicode is a registered trademark of Transoft Ltd.

Word is a registered trademark of Microsoft Corporation.

XT is a trademark of International Business Machines Corporation.



Index

login script command, using 189
%n identifier variable, using to shift LOGIN parameters 229, 241
* character, using in login scripts 224
; character, using in login scripts 224
<variable> identifier variable, explained 238

A

Abends, resolving 794
Aborting partition operation
 explained 299
 with NetWare Administrator 300
 with PARTMGR 301
Access server, checking for, in login scripts 239
ACCESS_SERVER identifier variable, explained 239
Account expiration date, changing 96, 100
Account Restrictions property, setting with NETADMIN 100
Account, ADMIN, unlocking 473
Accounting
 calculating charge rates for 628
 setting up 626, 627
 viewing totals 629
ACTIVATE SERVER utility, deleting from IOSTART.NCF file 851
ADD NAME SPACE utility, using 523
Add or Delete Self property right, explained 21. *See also* Rights, property

Address
 node, specifying in login scripts 238
 postal, changing 96, 101
 workstation network, specifying in login scripts 238
ADMIN User object
 default location in Directory tree 3
 default login script, executing 243
 default rights, listed 3
 logging in as 5
 rights available to 17
 unlocking account 473
Administration utilities, setting up after installation 5
AFP Server object, explained 36
AIO port interface module, loading 614
AIOCOMX communications port driver, loading 614
Alerts
 displaying on server console 502
 full volume 504
 hardware and NLM, turning on 502
 insufficient disk space for compression 546
 insufficient memory, handling 482
 turning off 502
 viewing in server error log 506
Alias object
 creating after moving container 76
 creating when renaming Container object 90
 deleting 88
 explained 36

- Allocating memory in NetWare Server for OS/2 825
- Alternate MSL boards, installing 834
- AM_PM identifier variable, explained 237
- Appending backup sessions 778
- Application
 - assigning trustee rights to 124
 - failure due to insufficient locks 501
 - guidelines for installing on network 123
 - installing on network 117
 - mapping fake root in login scripts 213
 - mapping fake root to, examples 219
- Archive Needed attribute, using with SBACKUP 732
- Arcnet LAN drivers, listed 450
- ASCII file. *See* Text file
- Asterisk, using in login scripts 224
- Asynchronous connection
 - explained 607
 - NLM for remote console 614
- ATOTAL utility, using 629
- ATTACH login script command, using 191
- Attributes, directory. *See also* Attributes, file
 - modifying with FILER 148
 - modifying with NetWare Administrator 147
- Attributes, file. *See also* Attributes, directory
 - displaying, with FILER 541
 - explained 132
 - modify bit (Archive Needed), using with SBACKUP 732
 - modifying 147, 148
- Audit Data file
 - closing old 663
 - copying old 663
 - creating reports from 655
 - explained 658, 662
 - resetting 663
 - setting size limits 647
- Audit History file, explained 662
- Audit reports
 - creating 655
 - creating filter for 655
 - creating filter, example 672
 - sending to file 658
 - using filter temporarily, example 676
 - viewing 660
- AUDITCON utility. *See also* Auditing
 - auditing networks with 631
 - files installed automatically 632
 - remounting volumes dismounted by 665
- Auditing
 - available events, listed 631
 - configuring environment for 646
 - creating auditor account 634 (*see also* Auditor)
 - creating reports 655 (*see also* Audit reports)
 - dismount volume option 649
 - enabling for container object 637
 - enabling for volume 636
 - events, selecting 650
 - examples 666
 - explained 631
 - files (*see* Audit Data file; Audit History file)
 - remount volumes dismounted by AUDITCON 665
 - using 634
- Auditor
 - account, creating 634
 - changing password for container object 641
 - changing password for volume 640
 - creating directory for 634
 - creating user object 634
 - logging in to container object 644
 - logging in to volume 643
 - logins, setting options 648
 - procedures before auditing 639
 - setting level-two password 643
 - setting up work environment 646
 - tasks, explained 639
- AUTOEXEC.BAT file, executing SERVER.EXE in 624
- AUTOEXEC.NCF file
 - creating or editing 459
 - file compression SET parameters in 545
 - internal network number in 471
 - SCHDELAY command in 487
 - server name in 471

B

Backup

- custom options 727
- files, explained 698
- target, missing from SBACKUP list 712, 723
- types of, explained 699

Backup devices

- checking status of 773
- renaming 777
- supported, listed 685

Backup log

- defined 762
- logbook, sample 706

Backup media

- amount required 778
- checking status of 774
- labeling 778
- numbers, incrementing 778
- out of space, handling 778
- spanning 778
- utilities, using 776

Backup MSL. *See* Alternate MSL boards, installing

Backup power. *See* UPS monitoring

Backup procedures. *See also* Restoring

- custom (selected portions of file system) 721
- differential (changes since last full backup) 711
- full (all data) 711
- incremental (changes since last backup) 711
- prerequisites 710
- troubleshooting 786

Backup session

- choosing subsets 699
- log and error files, deleting 771
- logbook, maintaining hard-copy 705
- overwriting parent and child data 743
- overwriting sessions on media 770
- scheduling 707

Backup types

- effect on restore sessions 707
- explained 699
- sample schedule, full and differential 709
- sample schedule, full and incremental 708

Bad blocks, testing hard disks for 511

- Batch file. *See also* AUTOEXEC.BAT file; AUTOEXEC.NCF file; STARTUP.NCF file
- adding search path for 454
- server, creating or editing 459

BATCH option for ITEM menu command, explained 262

Battery power. *See* UPS monitoring

Bindery services, rules for naming objects viewed through 42

Bindery. *See* Directory tree

Binding LAN drivers 451

Block suballocation on CD-ROM volume 530

Blocks

- surface test for bad 511
- viewing redirected 509

Booting

- remote server 624
- with EXIT command 624

Bottlenecks, avoiding server 480

BREAK login script command

- using 192
- using with GOTO 203

Browse object right, explained 19. *See also* Rights, object

Buffers

- cache, reporting low 502
- controlling with SET 784, 785

C

Cabling, resolving problems 792

Cache buffer

- alerts, turning on 502
- hits, tracking 481

Caching

- directory, changing 494
- file, changing 503

Callback list, creating 615

CD-ROM

- disk, using as read-only volume 530, 538
- extracting NetWare files from 457
- making client diskettes from 6

- CDROM NLM, loading 532
- CD-ROM volume
 - enabling block suballocation 530
 - enabling file compression 530, 538
 - mounting 530
- Charge rates, accounting, formula 628
- CHDIR option for ITEM menu command, explained 263
- Children of parent directory, for backup or restore
 - explained 703, 740, 759
 - overwriting 743
- CLEAR STATION command, using 474
- CLIB NLM, changing packet signature level for 577
- Client
 - install diskettes, creating from CD-ROM 6
 - rights, disabling 578
- CLS login script command, using 192
- Code page, setting keyboard type with 468
- COMMAND.COM file, specifying location of, in login scripts 193
- Commands
 - executing after stopping login script 199
 - executing from within login scripts 189
 - login script (*see* Login script commands)
- Comments
 - displaying during login with WRITE 233
 - using in login scripts with REMARK 224
- Communications port driver (AIOCOMX), loading 614
- Communications problems, resolving 791
- Compaq computer
 - long machine name, specifying in login scripts 211, 238
 - short machine name, specifying in login scripts 238
- Compare property right, explained 21. *See also* Rights, property
- COMPATIBLE login script command. *See* PCCOMPATIBLE login script command
- Compiling menus 257
- Compression, file. *See* File compression
- Computer object, explained 37
- Computer. *See also* Workstation
 - long machine name, specifying in login scripts 211, 238
 - short machine name, specifying in login scripts 238
- COMSPEC login script command
 - using 193
 - using with OS/2 workstations 194
- Conditional statements, using in login scripts 204
- Configuration files for SFT III, explained 836
- Configuring
 - audit options 646 (*see also* Auditing)
 - NetWare Server for OS/2 816
 - server products 456
 - server, sample files for 105
 - time synchronization for large networks 601
- Connection number
 - determining workstation 474
 - workstation, specifying in login scripts 238
- Connections to server
 - bottlenecks, resolving 478
 - clearing 473
 - logins, enabling and disabling 473
 - monitoring with watchdog packets 475
 - packet receive buffers, increasing 477, 478
- Console
 - display, changing SFT III 859
 - prompt, changing 841
- Container login script, explained 170
- Container objects
 - changing auditor password to 641
 - enabling auditing for 637
 - explained 25
 - listed 27
 - logging in to as auditor 644
 - naming rules for 28
 - property values, changing 94, 99
 - renaming 90
- Container objects, creating
 - with INSTALL 30
 - with NETADMIN 32
 - with NetWare Administrator 30

- Container objects, moving
 - explained 76
 - with NETADMIN 83
 - with NetWare Administrator 78, 81
- CONTEXT login script command, using 194
- Context, name
 - changing, in login scripts 194
 - displaying context user exists in 237
 - displaying context when login started 238
 - specifying for UIMPORT 343
- Control file
 - creating for UIMPORT 338
 - import control parameters, listed 340
- Control parameters, using in UIMPORT 340
- COPY command (DOS), verifying copies, in login scripts 197
- Copy, of Directory partition. *See* Replica
- Copying files
 - from installation diskettes 457
 - NetWare 456
 - verifying copies made with DOS COPY 197
 - with FILER 114
 - with NetWare Administrator 113
- Copying login scripts
 - with NETADMIN 181, 185
 - with NetWare Administrator 183
- Core dump, preparing 796
- Country object, explained 27
- Create Home Directory, control parameter for UIMPORT 340
- Create object right, explained 19. *See also* Rights, object
- Creating
 - Directory Map objects 127, 129
 - NetWare disk partitions 561
 - volumes 515
- Creating directories
 - HCSS domain 406
 - HCSS media-label 408
 - with FILER 111
 - with NetWare Administrator 109
- Creating Directory partitions
 - with NetWare Administrator 287
 - with PARTMGR 289

- Creating login scripts
 - with NETADMIN 180
 - with NetWare Administrator 177
- Creating replicas
 - explained 302
 - with NetWare Administrator 303
 - with PARTMGR 304
- Custom backup. *See* Backup procedures
- CX login script command. *See* CONTEXT login script command

D

- DAT file for menus 257
- Data file
 - auditing (*see* Audit Data file)
 - creating for UIMPORT 334
 - embedded punctuation 335
 - examples of, for UIMPORT 356, 357
 - updating 357
- Data migration. *See also* Optical media
 - creating HCSS domain directories for 406
 - deleting HCSS domain directory 418
 - explained 394
 - formatting for HCSS 411
 - HCSS-compatible hardware 400
 - reformatting optical media for 420
 - reloading optical media in jukebox 418
 - threshold example 394
- Data protection. *See* Security
- Database
 - Directory, repairing 361
 - importing user information from, with UIMPORT 334
 - integrity, protecting (*see* Transaction Tracking System)
 - sample, for UIMPORT 356
- Date
 - server, setting and viewing 469
 - specifying for delayed backup 719, 733
 - specifying in login scripts 237
- DAY identifier variable, explained 237
- DAY_OF_WEEK identifier variable, explained 237

- Deadlock, troubleshooting MSL 835
- Decompressing files. *See* File compression
- Default
 - drive, specifying in login scripts 198
 - objects and rights, NetWare 4.1 table 2
 - user rights, listed 3
- Default directory, backup log
 - changing 715, 725, 751
 - explained 762
- Default login script. *See also* Login scripts
 - commands in 243
 - explained 170
 - preventing from executing 220
- Default user values. *See* User template.
- Definitions, field, using in UIMPORT 345
- Delayed backup
 - security considerations 719
 - setting date and time 719, 733
 - with compressed files 782
- Delete object right, explained 19. *See also* Rights, object
- Delete property, control parameter for UIMPORT 341
- Delete Self property. *See* Add or Delete Self property
- Deleted files, purging
 - automatically with SERVMAN 549
 - with FILER 167
 - with NetWare Administrator 165
- Deleted files, salvaging
 - with FILER 163
 - with NetWare Administrator 161
- DELETED.SAV directory, explained 105
- Deleting objects
 - cautions 86
 - explained 86
 - with NETADMIN 89
 - with NetWare Administrator 88
- Description file for disk drivers, using 559
- Device, checking for new, on NetWare server 401, 402
- Diagnostic tools, listed 790
- Differential backup. *See* Backup procedure
- Directory
 - auditing activities in, example 671 (*see also* Auditing)
 - default, specifying in login scripts 198
 - home, specifying for UIMPORT 340
 - mapping drives in login scripts 212
 - mapping fake root in login scripts 213
 - name space information, viewing 156
 - rights, explained 18
 - searches, improving speed of 484
 - viewing information about 157, 158, 159
- Directory attributes. *See* Attributes, directory
- Directory caching
 - changing for faster writes 493
 - controlling 482
 - decreasing 496
 - increasing 484
- Directory Entry Table, updating with VREPAIR 525
- Directory Map object
 - creating 127, 129
 - explained 37, 125
 - mapping drive in login scripts 217
- Directory partition. *See* Partition, Directory
- Directory security
 - adding trustees 134, 135
 - changing owner 150, 151
 - deleting trustees 137, 139
 - modifying trustee rights 140, 142
 - rights, explained 131 (*see also* Rights, file system)
- Directory Services. *See* NetWare Directory Services
- Directory tree
 - changing name context in login scripts 194
 - configuring time synchronization for 601
 - creating partitions 286 (*see also* Partition, Directory)
 - database, repairing 361
 - example 4
 - example of new 4
 - merging 373
 - moving objects in 76
 - synchronizing two trees 604

- Directory, creating
 - automatically, during installation 105
 - planning guidelines for 105
 - structure, example 108
 - with FILER 111
 - with NetWare Administrator 109
 - workstation operating systems 117
- Directory, default
 - changing for backup log 715, 751
 - for backup log and error files 715
- Directory, deleted
 - purging 165, 167
 - salvaging 161, 163
- Disaster prevention and recovery 570, 793
- Disk drivers
 - listed 559
 - loading 559
- Disk duplexing 565
- Disk errors
 - checking for, with MONITOR 508
 - surface test for 511
- Disk mirroring
 - explained 565
 - recovering orphaned partitions 845
- Disk partition. *See* Partition, disk
- Disk reads, improving
 - by changing concurrent disk and directory writes 496
 - by changing turbo FAT Re-Use Wait Time 498
- Disk space
 - assigning to other volumes 517
 - checking, with NDIR 535
 - errors in server, resolving 797
 - errors in workstation, resolving 803
 - increasing, with file compression 536
 - managing server 535
 - required to add volume segment 528
 - required to create volumes 516
 - required to uncompress files 537
 - warning for insufficient 546
- Disk writes, improving
 - by changing disk and directory caching 493
 - by increasing concurrent writes 492
 - by turning off read-after-write verification 495
- Disk, hard. *See* Hard disk
- Diskette, client, making from CD-ROM 6
- Dismounting volume
 - with INSTALL 518
 - with AUDITCON 649
- DISPLAY login script command, using 195
- Distribution List object, explained 37
- Domain
 - loading NLMs in 581
 - misbehaved NLMs, using for 583
 - third-party NLMs, using for 581
- Domain directory, creating HCSS for data migration 406, 408
- DOMAIN NLM, loading 582
- DOS
 - code page, setting keyboard types with 468
 - installing on network 117
 - name space supported by SBACKUP 683
 - returning server console to 447
 - text files, editing 460
- DOS in login scripts
 - allowing <Ctrl><Break> 196
 - COMMAND.COM, specifying location of 193
 - directory, mapping search drive to 218
 - DOS BREAK command, using 196
 - DOS COPY command, verifying copies 197
 - DOS VERIFY command, using 197
 - environment variables, setting 226
 - prompt, setting 227
 - type, specifying 238
 - version, specifying 238
- DOS SET login script command. *See* SET login script command
- DOS workstation. *See also* Workstation
 - environment variables for, setting in login scripts 226
 - environment variables, setting in login scripts 238
 - installing software 15
 - installing, starting NETADMIN 15
 - making client installation diskettes from CD-ROM 6
 - non-IBM_PC, using PCCOMPATIBLE in login scripts 199, 222

- DOS workstation (*continued*)
 - password option for SBACKUP 692
 - resolving memory problems 802
 - running DOS from network 117
- DOWN utility, using 446
- DRIVE login script command, using 198
- Drive mapping
 - creating, in login scripts 212
 - deleting, in login scripts 213
 - with Directory Map object 125
 - Directory Map object, explained 37
 - to fake root, in login scripts 213
 - in login scripts, examples 218
 - to physical volume, in login scripts 213
 - recommended order in login scripts 215
 - using with Directory Map object, example 217
- Drive, default, specifying in login scripts 198
- Drivers. *See* Disk drivers; LAN drivers
- DSMERGE utility
 - checking servers 379
 - checking time synchronization 381
 - merging Directory trees with 373
 - merging trees 382
 - options explained 378
 - overview 373
 - preparing trees 376
 - renaming trees 385
- DSREPAIR utility
 - checking replica synchronization 368
 - checking time synchronization 365
 - error log, viewing and editing 371
 - options explained 363
 - overview 361
 - repairing Directory database with 361
 - running unattended 364
- DSTRACE utility
 - explained 389
 - using 389
 - using SET parameters 389
- Duplexing disks 565

E

- ECB available count errors, handling 477
- EDIT NLM
 - editing text files from console 460
 - function keys listed 461
- Editing NET.CFG file for OS/2 829
- Effective rights
 - trustee 153
 - viewing 154
- End address for unclaimed memory, setting 840
- Environment
 - auditing options, listed 647
 - setting for auditor 646
 - variables (DOS and OS/2), setting in login scripts 226, 238
- Error log
 - backup, changing 725
 - backup, explained 736
 - defined 762
 - explained 763
 - managing 508
 - MSEngine, explained 858
 - server, viewing 506
 - session errors appended to file 768
 - TTS, creating 590
 - viewing 766
- Error messages. *See* Alerts; Disk errors; *System Messages* manual
- ERROR_LEVEL identifier variable, explained 239
- ETC directory, explained 105
- Ethernet
 - 802.2 frame type 452
 - LAN drivers, listed 450
 - workstations, finding SFT III server 847
- Event control block errors, resolving 801
- Events, auditing. *See also* Auditing
 - creating reports of 655
 - explained 631
 - listed 653
 - selecting 650

- Examples
 - data files 356
 - Import Control Files 356
- Excluding data
 - in backups 700
 - in restores 737, 739, 756
 - while including portions 702
- EXEC menu command, explained 266
- Executable file
 - executing after stopping login script 199
 - executing from within login script 189
- EXIT login script command
 - using 199
 - using on non-IBM computers 199
- EXIT utility, using 447
- Exporting media, HCSS 415
- External Entity object, explained 38
- External hard disk, adding to server 555. *See also*
 - Hard disk
- Extracting NetWare files from installation
 - CD-ROM 457

F

- Fake root
 - mapping, in login scripts 213
 - mapping, in login scripts, examples 219
- FAT. *See* File Allocation Table
- FDISPLAY login script command, using 201
- Field definitions, using in UIMPORT 345
- File Allocation Table (FAT)
 - changing turbo FAT re-use wait time 498
 - updating with VREPAIR 525
- File attributes. *See* Attributes, file
- File caching, controlling 482
- File compression
 - backing up compressed files 782
 - backing up uncompressed files 782
 - daily checks for 543, 546
 - enabling, on volume 537
 - explained 536
 - grouping files for 539
 - on CD-ROM volume 530, 538
 - restoring files to unsupported environment 782
 - SET parameters for 544
 - setting, for files and directories 539
 - suspending 543
 - uncompressing files 537
 - using SERVMAN to manage 548
 - using SET command to manage 547
- File events
 - auditing 653
 - reporting 655
- File information
 - name space, viewing 156
 - size, viewing 156
 - viewing 157, 158, 159
- File locks, limiting 499
- File security. *See also* Rights, file system
 - adding trustee 135
 - adding trustees 134
 - changing owner 150
 - changing owner of 151
 - deleting trustees 137, 139
 - rights, explained 18
- File server. *See* NetWare server
- File system rights. *See* Rights, file system
- File, batch. *See* Batch file
- FILE_SERVER identifier variable, explained 238
- FILER utility. *See* specific task
- Files
 - database, protecting with TTS 589
 - extracting from CD-ROM 457
 - import control and data, creating for UIMPORT 338
 - NetWare, location of 105
 - purging from disk 549
 - server batch (.NCF), creating or editing 459
 - text, editing 460
 - word-processed, displaying during login 201
- Files, copying
 - NetWare 456
 - with FILER 114
 - with NetWare Administrator 112
- Filter. *See* Inherited Rights Filter

- Filters, audit report. *See also* Auditing
 - creating report 655
 - creating temporary, to sort audit information, example 676
 - creating, example 672
 - options for audit report 657
- FIRE PHASERS login script command, using 202
- FLAG utility, using to change file compression attributes 539
- Formatting
 - hard disks 569
 - optical media for HCSS 412
- Formula
 - accounting charge rates 628
 - name space memory requirements 523
- Frame type, Ethernet 802.2 452
- FTAM name space supported by SBACKUP 683
- Full backup. *See* Backup procedures
- Full volume alerts, setting 504
- FULL_NAME identifier variable, explained 237

G

- GETP menu command, variables 271
- GETx menu command, using 270
- Global auditing options. *See also* Auditing
 - examples of 652, 666
 - explained 647
- GOTO login script command, using 203
- GREETING_TIME identifier variable, explained 237
- Group object
 - adding members 51, 52
 - deleting members 57, 58
 - explained 38
 - giving rights to file system 53, 55
 - managing User objects as Group 48
 - specifying in login scripts 238

H

- Hard disk space. *See* Disk space
- Hard disk. *See also* Disk drivers
 - adding SCSI to subsystem 555
 - bad blocks, testing for 511
 - drivers, loading 559
 - formatting 569
 - internal IDE, adding 552
 - internal SCSI disk and adapter, adding 553
 - mirroring, explained 565
 - out of sync, recovering 591
 - replacing 556
 - unmirroring, explained 567
- Hardware
 - checking for new, on NetWare server 401, 402, 439
 - troubleshooting 790
- HCSS media
 - assigning to directories 413
 - checking status 420
 - exporting 416
 - formatting 412
 - importing 409
 - reformatting 420
 - reloading 418
- HCSS parameters
 - explained 421
 - methods of modifying 427
 - setting, at server console 434
 - setting, using NetWare Administrator 427
- HCSS. *See* High Capacity Storage System (HCSS)
- High Capacity Storage System (HCSS)
 - creating volume 402
 - deleting domain directories 418
 - explained 106
 - file system structure 392
 - hardware, installing compatible 400
 - limitations 397

- loading on server 403, 440
- managing file system, procedures 415
- overview 398
- setting up file system 405
- setting up workstation 404

High-speed link. *See* Mirrored server link (MSL)

History file, auditing. *See* Audit History file

Home directory, specifying for UIMPORT 340

Hot Fix Redirection Area, viewing blocks with MONITOR 509

HOUR identifier variable, explained 237

HOUR24 identifier variable, explained 237

HPFS files, removing name space for 524

I

I/O errors, resolving

- event control block allocation 801
- file 799
- server 796
- volume 800

I'm alive packets, troubleshooting SFT III 854

IBM computer. *See also* Non-IBM computer; Workstation

- compatible with, specifying in login scripts 222
- long name, specifying in login scripts 211, 238
- short name, specifying in login scripts 238

IBM PC Network LAN driver, listed 450

IDE disk. *See* Hard disk

Identifier variables

- %n, explained 241
- %n, shifting LOGIN parameters with SHIFT 229
- conventions in login scripts 177
- displaying value during login, with WRITE 233
- examples 240
- listed 237
- using 240
- using with IF...THEN 204

IF...THEN login script command

- using 204
- using with identifier variables 204, 240

Import Control File 338

Import control file

- control parameters, listed 340
- creating for UIMPORT 338
- examples 355
- field definitions, explained 338
- parameters, explained 338

Import mode, control parameter for UIMPORT 342

Importing

- media, HCSS 409
- user information with UIMPORT 333

INCLUDE login script command, using 208

Including data

- in backups 701
- in restores 738, 739, 756
- while excluding portions 702

Incomplete transactions, backing out of 589

Incremental backup. *See* Backup procedures

Inherited Rights Filter (IRF), viewing and modifying 144, 145

INSTALL NLM

- adding volume segment 528
- changing internal network number 471
- changing server name 471
- changing volume size 529
- copying files 456
- creating disk partitions 562
- creating volume 515
- deleting disk partitions 564
- deleting volume 519
- dismounting volume 518
- editing NCF files 459
- installing, uninstalling, and configuring server products 455
- loading disk drivers 560
- loading LAN drivers 451
- mirroring disks 565
- mounting volume 518
- recovering out-of-sync hard disks 592
- renaming volume 521
- testing hard disk surfaces 512
- unmirroring disks 567

Installation CD-ROM, extracting files from 457

Intruder Detection property, changing 96, 100

- IO\$LOG.ERR file, limiting size of 842
- IOAUTO.NCF file, explained 836
- IOEngine
 - capturing console messages 840
 - log file, limiting size of 842
 - network number, troubleshooting 852
- IOSTART.NCF file
 - explained 836
 - setting test options in 859
- IPX internal network number
 - changing 471
 - troubleshooting 852
- IPX protocol, binding to 451
- IRF. *See* Inherited Rights Filter (IRF)
- ITEM command options
 - BATCH, explained 262
 - CHDIR, explained 263
 - listed 262
 - PAUSE, explained 263
 - SHOW, explained 263

J

- Job servers, packet signature considerations for 577
- Joining Directory trees with DSMERGE 373

K

- KEYB NLM, using 467
- Keyboard, setting language for 467

L

- LAN drivers
 - increasing available memory for 789
 - listed 450
 - loading and binding 449
 - statistics for 489

- Language
 - changing server language 464
 - copying language files to server 462
 - setting, for NLMs 466
 - setting, for server keyboard 467
- LANGUAGE utility, using 467
- LAST_NAME identifier variable, explained 237
- LASTLOGINTIME login script command, using 210
- Leaf objects
 - changing property values 95, 100
 - deleting 88, 89
 - explained 35
 - listed 36
 - mandatory fields when creating 45
 - naming rules for 41
 - renaming 90
- Leaf objects, creating
 - explained 35
 - with NETADMIN 45
 - with NetWare Administrator 44
- Leaf objects, moving
 - with NETADMIN 79
 - with NetWare Administrator 77, 78
 - with NetWare Administrator (drag and drop) 78
- Links, UNIX
 - error restoring 768
 - retained by SBACKUP 765
- LOAD menu command, explained 267
- LOAD utility. *See* Loading drivers; Loading NLMs
- Loader errors, SFT III (troubleshooting) 841
- Loading
 - NetWare Server for OS/2 806
 - network applications, guidelines for 117
- Loading drivers
 - device, backup 688
 - device, controller 401, 438, 688
 - device, optical media storage 401, 438
 - disk 559
 - LAN 450

- Loading HCSS
 - drivers 401, 438
 - on server 403, 440
 - on workstation 404
- Loading NLMs
 - determining protection domain for 448
 - procedure 449
 - setting language for 466
- Loading SBACKUP files
 - for NDS 694
 - for servers 689
 - for SFT III 695
 - for workstations 690
 - general requirements 687
 - summary table 696
- Locked ADMIN user account, enabling 473
- Locked device errors, resolving 799
- Locks, file and record, increasing 501
- Log file
 - restore, changing directory or path 751, 752
 - server, viewing 506
 - TTS, creating 590
- Log file, backup
 - changing directory or path 715, 725
 - creating directories for 715
 - deleting overwritten sessions from 770
 - explained 762
- Logging in
 - ADMIN User object 5
 - attaching to server from login script 191
 - auditor, to container 644
 - auditor, to volume 643
 - changing name context in login scripts 194
 - enabling and disabling 473
 - last time of, displaying during login 210
 - to backup target 714, 724, 750
- Login script commands
 - # (executing commands) 189
 - ATTACH (attaching to servers) 191
 - CLS (clearing screen) 192
 - CONTEXT (changing context) 194
 - DISPLAY (displaying text file) 195
 - DRIVE (specifying default drive) 198
 - FDISPLAY (displaying word-processed file) 201
 - IF...THEN (executing conditional action) 204
 - INCLUDE (executing subscript) 208
 - LASTLOGINTIME (displaying time of last login) 210
 - MACHINE (sets DOS machine name) 211
 - MAP (mapping drives and search drives) 212
 - NO_DEFAULT (preventing default script) 220
 - NOSWAP (preventing memory swapping) 221
 - PAUSE (pausing login script execution) 222
 - PCCOMPATIBLE (specifying non-IBM computers) 222
 - REMARK (inserting comments) 224
 - SET_TIME (synchronizing with server) 228
 - SHIFT (shifting LOGIN parameters) 229
 - SWAP (swapping memory) 231
 - WRITE (displaying messages) 233
- Login Script property
 - changing with NETADMIN 100
 - changing with NetWare Administrator 96
- Login scripts
 - commands and variables, listed 187
 - container, example 245
 - conventions, listed 176
 - default, explained 170
 - default, preventing execution of 220
 - examples 243
 - executing one script as part of another 208
 - exiting to menu from 199, 276
 - explained 169
 - hints for using 176
 - identifier variables, conventions 177
 - identifier variables, listed 237
 - identifier variables, using 236
 - location of, explained 172
 - objects that contain, listed 171
 - previous NetWare versions, using 174
 - profile, explained 170
 - profile, sample 248
 - types, explained 170
 - user, explained 170
 - user, sample 250

- Login scripts, creating
 - command syntax 187
 - copying 181, 183, 185
 - examples 243
 - planning 171
 - profile, assigning user to 178, 183
 - rules 176
 - upgrading from previous versions 174
 - with NETADMIN 180
 - with NetWare Administrator 177
 - Login scripts, mapping drives
 - command format 212
 - to Windows directories 122
 - with Directory Map object 126
 - LOGIN utility parameters, shifting with %n in
 - login scripts 229
 - LOGIN_CONTEXT identifier variable, explained 237
 - LOGIN_NAME identifier variable, explained 237
 - Long machine name. *See also* Machine name
 - non-IBM_PC, using PCCOMPATIBLE with 199
 - specifying in login scripts 211, 238
 - Lost interrupt alerts, troubleshooting SFT III 841
- ## M
- MACHINE identifier variable, explained 238. *See also* Machine name
 - MACHINE login script command, using 211
 - Machine name
 - long, specifying in login scripts 211, 238
 - non-IBM_PC, using PCCOMPATIBLE with 199
 - short, specifying in login scripts 238
 - Macintosh files
 - adding name space for 522
 - name space supported by SBACKUP 683
 - removing name space for 524
 - viewing information about 157, 158, 159
 - Magneto-optical media. *See* Optical media
 - MAIL directory, explained 105
 - Managing NetWare Server for OS/2 805
 - MAP login script command
 - examples 217
 - using 212
 - Map object. *See* Directory Map object
 - Mapping drives. *See* Drive mapping
 - Master replica. *See also* Replica
 - changing to another type 321
 - explained 282, 284
 - Maximum directory retries, control parameter for UMIPORT 343
 - Media
 - magneto-optical (HCSS), labeling 393
 - status, HCSS 420
 - Media utilities
 - erasing all data on media 777
 - erasing media header 777
 - retensioning tape 777
 - MEMBER OF "group" identifier variable, explained 238
 - Memory
 - domain, explained 581
 - requirements for SBACKUP 687
 - MEMORY MAP utility
 - displaying unclaimed memory with 839
 - sample display 839
 - Memory sharing. *See* NetWare Server for OS/2
 - Memory, NetWare Server for OS/2
 - allocating 825
 - freeing 814
 - sharing 812
 - Memory, server
 - adding to SFT III servers 837
 - allocation, controlling 482
 - cache buffer alerts 502
 - freeing up RAM 482
 - increasing available 482
 - protecting OS domain 581
 - resolving problems 798
 - setting new end address for unclaimed 840
 - SFT III, aligning 839
 - unclaimed, displaying 839

- Memory, workstation
 - increasing available amount 789
 - not swapping, with login scripts 221
- Menu commands. *See also* NMENU utility
 - EXEC (specifying actions) 266
 - GETx (specifying user input) 269
 - ITEM (specifying menu options) 262
 - LOAD (executing one menu from another) 267
 - MENU (naming menus) 261
 - SHOW (displaying submenu) 268
- Menu program. *See* NMENU utility
- MENU utility, converting to NMENU 277
- MENUCNVT command, using 277
- MENUMAKE utility, compiling converted menu program 277
- Menus. *See* NMENU utility
- Merging partitions
 - explained 290
 - with NetWare Administrator 291
 - with PARTMGR 292
- Merging two trees
 - after running the DSMERGE utility 387
 - using the DSMERGE utility 382
- Message Routing Group object, explained 38
- Messages
 - displaying during login 233
 - sending from console with SEND 444
- Messages, error. *See* Alerts; Disk errors; *System Messages* manual
- Messaging Server object, explained 38
- Microsoft Windows. *See* Windows
- MINUTE identifier variable, explained 237
- Mirrored server link (MSL)
 - cable lengths, listed 834
 - failure message, example 858
 - failure, troubleshooting 834
 - troubleshooting 834
- Mirroring disks 565. *See also* Unmirroring hard disks
- Missing RAM error, troubleshooting 838
- MNU filename extension, explained 277. *See also* NMENU utility
- Modem
 - null modem cable, using 614, 620
 - running remote console session over 613
- Modify bit (Archive Needed attribute), using with SBACKUP 732
- MONITOR.nlm
 - checking cache utilization 481
 - checking for disk errors 509
 - checking processor utilization 488
 - clearing workstation connection 474
 - determining connection number or username 444
 - Packet Receive Buffer value 478
 - prioritizing server processes 487
 - turning off read-after-write verification 495
 - viewing LAN driver statistics 490
- MONTH identifier variable, explained 237
- MONTH_NAME identifier variable, explained 237
- Mounting volume
 - with INSTALL 530
 - dismounted by AUDITCON 665
 - failure due to insufficient RAM 519
- Moving objects
 - with NETADMIN 79, 83, 297
 - with NetWare Administrator 81, 294
 - with NetWare Administrator (drag and drop) 78
- Moving partitions
 - explained 293
 - with NETADMIN 297
 - with NetWare Administrator 294
- MSAUTO.NCF file, explained 836
- MSEngine
 - already loaded message, explained 834
 - error log, explained 858
 - memory range, setting 840
 - name, troubleshooting 853
 - outputs, checking 832
 - unclaimed memory used by 839
- MSL boards
 - alternate, installing 834
 - troubleshooting 854

MSL deadlock, troubleshooting 835
MSL. *See* Mirrored server link (MSL)
MSSTART.NCF file, explained 836
MSSTATUS.DMP file, limiting size of 842

N

Name

- container object, rules for 28
- leaf object, rules for 41
- server, changing 471
- user's full, specifying in login scripts 237
- user's last, specifying in login scripts 237
- user's login, specifying in login scripts 237
- volume, assigning 516

Name context

- changing, in login scripts 194
- control parameter for UIMPORT 343

Name space

- adding to volume 522
- calculating memory required for 525
- NFS, retains symbolic and hard links 765, 768
- removing from volume 524
- SBACKUP supported formats 683
- viewing directory or file 156

Name, machine

- long, specifying in login scripts 211, 238
- short, specifying in login scripts 238

NCF files for SFT III, explained 836

NCP packet signature

- CLIB NLMs and 577
- client levels 574
- effective 575
- examples of 576
- explained 572
- job servers and 577
- options for 572
- server levels 573
- turning off for RSPX 608
- when to use 572

NCUPDATE utility

- updating NET.CFG files after moving
Container object 76

- updating NET.CFG files after renaming
Container object 91

NDAY_OF_WEEK identifier variable, explained 237

NDIR utility

- checking available disk space with 535
- viewing information about directories and files
with 159

NDS synchronizationsstatus

- viewing and managing 389

NDS Trace, explained 389

NDS. *See* NetWare Directory Services

NET.CFG file, setting OS/2 nonded server option 829

NETADMIN utility, installing on DOS workstation 15. *See also* specific task

NetWare Administrator utility. *See also* specific task

- creating Windows icon for 10, 14

- file compression, setting 542

- installing on OS/2 workstation 12

- installing on Windows workstation 9

- starting 9

NetWare Directory Services

- auditing NDS events 654 (*see also* Auditing)

- managing Directory database 281

- partitions, explained 282

- repairing with DSREPAIR 361

- trees, merging with DSMERGE 373

NetWare Directory Services, restoring

- from backup session 694

- from replicas 694

- with partitions 743

- without partitions 743

NetWare files

- extracting from installation CD-ROM 457

- location of 105

- upgrading from remote console 621

NetWare Loadable Modules

- domains for 448, 581

- language, specifying 466

- loading 448

- loading misbehaved 583

- NetWare Requester for DOS, version, specifying in login scripts 238
- NetWare Requester for OS/2
 - installing, on OS/2 workstation 12
 - SYS:PUBLIC directory, mapping drive to in login scripts 215
 - version, specifying in login scripts 238
- NetWare server
 - abends, resolving 794
 - accounting, setting up 626
 - auditing events 647 (*see also* Auditing)
 - bindery-based, attaching to from login script 191
 - booting 447
 - changing server language 464
 - copying language files to 462
 - diagnostic tools, listed 790
 - error logs, viewing 506
 - hardware, checking for new 401, 439
 - insufficient disk space, resolving 797
 - managing connections to 473
 - monitoring and optimizing 481
 - name and IPX network number, changing 471
 - password for remote console 608
 - processes, prioritizing 486
 - processor utilization, checking 488
 - remote 606
 - resolving I/O errors 796
 - resolving memory problems 798
 - specifying in login scripts 238
 - troubleshooting 794
- NetWare Server for OS/2
 - allocating memory 825
 - automatic memory sharing, setting 814
 - configuring 816
 - freeing server memory 814
 - loading 806
 - loading PMMON 808
 - managing 805
 - manual memory sharing, setting 812
 - memory sharing, setting 812
 - monitoring 808
 - navigating in 807
 - NLMs, using 807
 - NWOS2.EXE, loading with 806
 - performance tuning 811
 - PMMON utility, using 810
 - printing 807
 - removing base server memory 825
 - using with NetWare 4 805
- NetWare Server for OS/2, NET.CFG file
 - automatic memory sharing, setting 826
 - base server memory, setting 824
 - block allocation size, setting 827
 - entering/editing options 817
 - format 817
 - initialization screen delay, setting 824
 - maximum server memory, setting 827
 - nonded server option in 829
 - setting performance tuning in 823
 - syntax for entering options 822
- NetWare Server object
 - deleting 330, 331, 333
 - explained 39
 - naming restrictions for 43
- NetWare server, configuring
 - keyboard type, changing 467
 - time and time zone, viewing and setting 469
- NetWare shell, specifying version in login scripts 238
- NETWARE_REQUESTER identifier variable, explained 238
- Network
 - address, specifying in login scripts 238
 - applications, guidelines for installing 123
 - applications, installing 117
 - auditing, explained 631 (*see also* Auditing)
 - board, LAN drivers for 450
 - directory structure, sample 108
 - installing DOS on 117
 - installing operating systems on 117
 - installing Windows on 119
 - supervisor, auditing tasks explained 633
 - traffic, monitoring 489
 - troubleshooting communications problems 791
- Network number, internal, changing 459
- Network time synchronization. *See* Time synchronization

- NETWORK_ADDRESS identifier variable, explained 238
 - NFS name space supported by SBACKUP 683
 - NLIST utility, explained 75. *See also* SEARCH utility
 - NLM. *See* NetWare Loadable Modules
 - NMENU utility
 - compiling menus 257
 - complex menu example 258
 - converting MENU program to NMENU 277
 - elements, explained 254
 - examples 257, 258, 272
 - executing menu program from login scripts 199
 - explained 251
 - illustrated 253
 - making menus work 275
 - maximum item lines 259
 - maximum number of menus 259
 - maximum number of windows 259
 - memory requirements 251
 - name length 259
 - numbering 261
 - old menu files, converting 277
 - planning 256
 - planning menus 259
 - setting up user environments 275
 - simple menu example 257
 - starting from login script 276
 - steps in creating menus 256
 - submenus, numbering 261
 - syntax 252
 - NMENU utility, commands
 - control, listed 265
 - GETx, using 270
 - ITEM options, explained 262
 - MENUCNVT, using 277
 - organizational, listed 260
 - types, explained 260
 - wrapping 259
 - NO_DEFAULT login script command, using 170, 220
 - Node address, specifying in login scripts 238
 - Non-IBM computers
 - long name, specifying in login scripts 211, 238
 - short name, specifying in login scripts 238
 - NOSWAP login script command
 - using 221
 - using with # command 189
 - NOT MEMBER OF "group" identifier variable, explained 238
 - NWPA NLM, using with CDROM NLM 531
 - NWPAR.DLL file, restricting access to 286
 - NWXTRACT utility, using 457
- ## O
- Object rights. *See* Rights, object
 - Objects
 - renaming 90, 91
 - Objects, moving
 - explained 76
 - with NETADMIN 79, 83
 - with NetWare Administrator 77, 78, 81
 - Objects. *See also* specific objects
 - cautions when deleting 86
 - deleting from Directory tree, explained 86
 - Directory Map object, mapping drive in login scripts 217
 - renaming 91, 92
 - searching for 71
 - trustee assignments to, managing 22
 - with login scripts 171
 - Operating systems, installing on network 117
 - Optical media. *See also* Data migration
 - formatting for HCSS 411
 - labeling HCSS 413
 - reformatting for HCSS 420
 - storing files on, explained 392
 - Organization object
 - changing property values 94, 99
 - enabling auditing for 637
 - explained 27
 - logging in to, as auditor 644
 - login script (*see* Login scripts)

- Organizational Role object
 - creating 64, 66
 - explained 39
 - managing Users with 63
- Organizational Unit object
 - changing property values 94, 99
 - explained 27
 - login script (*see* Login scripts)
- OS domain. *See* Domain
- OS identifier variable, explained 238
- OS/2 HPFS files
 - adding name space for 522
 - name space supported by SBACKUP 683
 - removing name space for 524
 - viewing information about 156, 158, 159
- OS/2 workstation. *See also* NetWare Requester for OS/2; Workstation
 - environment variables for, setting in login scripts 226, 238
 - installing NetWare Requester software 12
 - installing, starting NetWare Administrator 12
 - making client installation diskettes from CD-ROM 6
 - resolving memory problems 803
 - using COMSPEC with 194
- OS/2. *See* NetWare Server for OS/2
- OS_PROTECTED domain. *See* Domain
- OS_VERSION identifier variable, explained 238
- Out of sync disk, recovering data from 591
- Overwriting
 - deleting overwritten sessions from backup log 770
 - existing children during backup 743
 - existing parent directory during backup 743
 - session on backup media 770
- Owner, directory or file
 - changing 150, 151
 - viewing information about 156

P

- P_STATION identifier variable, explained 238
- Packet receive buffers
 - during peak usage 480
 - explained 477
 - increasing maximum 477
 - increasing minimum 479
 - New Packet Receive Buffer Wait Time 480
- Packet. *See also* NCP packet signature; Packet receive buffers
 - increasing size of 490
 - preventing forgery 572
 - viewing traffic 489
- Parameters
 - HCSS, explained 421
 - HCSS, setting 427
 - TSASMS parameters for SBACKUP, explained 692
- Parent directory, for backup or restore
 - explained 703, 740, 759
 - overwriting 743
- Partition Manager, restricting access to 286
- Partition, Directory. *See also* Replica
 - aborting operation, explained 299
 - cautions when moving 294
 - creating, explained 286
 - deleting server containing 330, 331
 - explained 282
 - guidelines for managing 283
 - listing 315, 317
 - merging, explained 290
 - moving, explained 293
 - planning 282
 - rules for using 283
- Partition, disk
 - creating 561
 - deleting 563
 - mirroring 565
 - orphaned, recovering 845

- PARTMGR utility. *See* specific task
- Password
 - changing auditor, for container object 641
 - changing auditor, for volume 640
 - displaying expiration time during login 238
 - setting level-two, for auditing 643
 - using with SBACKUP 692
- Password for remote console
 - encrypting 619
 - entering 608
- Password Restrictions property, changing 95, 100
- PASSWORD_EXPIRES identifier variable, explained 238
- Path, home directory, specifying for UIMPORT 341
- PAUSE login script command, using 222
- PAUSE option for ITEM menu command, explained 263
- PCCOMPATIBLE login script command
 - using 222
 - using with EXIT 199
- Performance tuning
 - setting for NetWare Server for OS/2 811
 - setting in NET.CFG 823
- Phaser sound (FIRE PHASERS login script command) 202
- Planning
 - directory and file system 105
 - file system, illustrated 108
 - login scripts 171
 - menus 256
 - partitions and replicas 282
- PMMON utility
 - loading 808
 - using 810
 - using to monitor NetWare for OS/2 808
- Postal Address property, changing 96, 101
- Power outage. *See* UPS monitoring
- Power supply errors, troubleshooting 792
- Preparing
 - checking time synchronization before merge 377
 - merging trees 376
- Primary server, SFT III
 - determining 858
 - failure, explained 856
 - hard disk failure, handling 842
 - IPX route to secondary, troubleshooting 855
 - network board failure, handling 850
 - troubleshooting two 834, 851
 - unloading LAN driver from 856
- Print Queue object, explained 39
- Print Server object, explained 39
- Printer object, explained 40
- Problems, system. *See* Disaster prevention and recovery; Troubleshooting
- Processes, server
 - prioritizing 486
 - viewing 488
- Processor utilization, server
 - checking 488
 - prioritizing processes 486
 - redistributing 488
- Profile login script. *See also* Login scripts
 - creating 180, 183
 - example 248
 - explained 170
- Profile object
 - assigning user to 178, 181
 - changing Login Script property 97, 101
 - creating 60, 61
 - explained 40
 - login script, example of 248 (*see also* Login scripts)
 - managing User objects with 59
- Prompt, setting in login scripts 227
- Properties, object. *See* Account Restrictions property; Intruder Detection/Lockout property; Login Script property; Password Restrictions property; Postal Address property; Security Equivalences property
- Property rights. *See* Rights, property
- Protecting data. *See* Security
- Protecting OS memory 581
- Protocol, binding communication 451

PUBLIC directory
explained 106
mapping drive for OS/2, in login scripts 215
mapping drive in login scripts 216
Purging deleted files. *See also* Salvaging deleted files
automatically, with SERVMAN 549
with FILER 167
with NetWare Administrator 165

Q

QMS events. *See also* Auditing
auditing 653
reporting 657
Queue Management Services events. *See* QMS events
Queue object. *See* Print Queue object
Quote, control parameter for UIMPORT 343

R

RAM, server. *See* Memory, server
RAM, workstation
required for NMENU 251
required for SBACKUP 692
RCONSOLE utility
connection types 607
navigating with <Alt> and function keys 611
time out error, handling 618
using with asynchronous (modem) connection 613
using with direct connection 607
Read property right, explained 21. *See also* Rights, property
Read/write replica. *See also* Replica
changing to another type 321, 323
explained 284
Read-after-write verification, turning off 495
Read-only replica. *See also* Replica
changing to another type 321, 323
explained 283, 284

Reads (disk), improving. *See* Disk reads, improving
Rebooting
NetWare server 447
remote servers 624
Record locks, limiting 499
Recovering deleted files with FILER 163
Redirected blocks, recording 509
REM (REMARK) login script command, using 224
Remote console
callback list, creating 615
ending session 612
function keys 611
loading modules on bootup 619
modem configuration for 617
modules required for asynchronous (modem) connection 613, 616
modules required for direct connection 607
password encryption 619
password, entering 608
starting session 607, 613
upgrading NetWare files from 621
REMOTE NLM, using for remote console 608
Remote server
installing or reconfiguring applications on 622
managing from remote console 606
rebooting 624
upgrading files on 621
REMOVE DOS utility, freeing up RAM with 482
Rename object right, explained 19. *See also* Rights, object
Renaming objects
explained 90
with NETADMIN 92
with NetWare Administrator 91
Renaming trees
explained 385
using the DSMERGE utility 385
Repairing Directory database with DSREPAIR 362
Replace value, control parameter for UIMPORT 343
Replica
changing type 321, 323
creating, explained 302

- Replica (*continued*)
 - defaults created at installation 283
 - deleting server containing 329
 - deleting, explained 306
 - explained 282
 - listing 318, 319
 - planning 283
 - rules and guidelines for using 283
 - types, explained 284
 - using to manage NDS 282
- Replica updates
 - receiving 312, 314
 - sending 309, 310
- Report, audit. *See* Audit reports
- REQUESTER_CONTEXT identifier variable, explained 238
- RESTART SERVER utility, using 447
- Restore session
 - choosing subsets 737
 - errors appended to error file 768
- Restoring
 - deleted directories and files 161, 163
 - differential backups, example 708
 - Directory database with DSREPAIR 361
 - full and differential backups, example 710
 - full and incremental backups, example 709
 - incremental backups, example 708
 - NDS from backup session 694
 - NDS from replicas 694
 - options, explained 737
 - prerequisites 746
 - procedures 747
 - to new location 742, 746, 754, 757
- Resynchronizing servers, reducing time for 833
- Rights, file system
 - assigning to operating system directories 124
 - assigning to trustee 131
 - default 2
 - effective, viewing trustee 153, 154
 - explained 131
 - modifying trustee 140, 142
 - viewing and modifying Inherited Rights Filter 144, 145

- Rights, object
 - default for User 3
 - default for User ADMIN 3
 - explained 18
 - listed 19
 - trustee, changing 23
- Rights, property
 - explained 20
 - listed 21
- Root, fake
 - mapping in login scripts 213
 - mapping in login scripts, examples 219
- RS232, loading 614
- RSPX NLM, loading 607, 608
- RX-Net LAN driver, listed 450

S

- Salvaging deleted directories and files
 - with FILER 163
 - location 105
 - with NetWare Administrator 161
- SAP. *See* Service Advertising Protocol (SAP)
- SBACKUP utility, procedures
 - backing up data, irregular (custom) 721
 - backing up data, regular (full, differential or incremental) 711
 - changing targets 779
 - checking device status 773
 - checking media status 774
 - deleting log files 770
 - enhancing host server performance 783
 - erasing or retensioning media 776
 - loading drivers 688
 - renaming storage device 777
 - restoring data 747
 - setting log and error file location 768
 - troubleshooting 786
 - unloading SBACKUP files 781
 - viewing error file 766
 - viewing log file 764
 - viewing media lists 772

- SBACKUP utility. *See also* Backup
 - children and parent directories, illustrated 704
 - compressed files, handled by 782
 - error files, explained 763
 - include and exclude restore options, explained 700, 737
 - installing 685
 - log files, explained 762
 - overwriting, explained 743
 - restore, explained 736
 - scanning, explained 703, 740
- Scanning data to back up 730
- SCHDELAY utility, using 486
- Scripts. *See* Login scripts
- SCSI disk. *See* Hard disk
- Search drive
 - deleting, in login scripts 213
 - mapping in login scripts 212
 - mapping to fake root, in login scripts 213
 - mapping to physical volume, in login scripts 213
 - recommended order, in login scripts 216
- Search path, adding server 454
- SEARCH utility
 - creating searchable container objects 29
 - creating searchable leaf objects 43
 - searching with 72, 74, 75
 - using 454
- SECOND identifier variable, explained 237
- Secondary server, SFT III
 - missing RAM error explained 838
 - restarts, troubleshooting 832
- SECURE CONSOLE utility, using 570
- Security
 - intruder prevention. *See* NCP packet signature
 - mirroring and duplexing hard disks 565
 - preventing packet forgery 572
 - preventing virus infection 571
 - remote console 606
 - risks associated with job servers 577
 - securing server console 570
 - Transaction Tracking System 589
 - uninterruptible power supply 579
- Security Equivalences property, setting for User 96, 101
- Segment, adding volume 528
- Semicolon, using in login scripts 224
- SEND command 444
- Separator, control parameter for UIMPORT 344
- Server console. *See also* NetWare server; Remote console
 - editing text files from 460
 - securing 570
- Server memory. *See* Memory, server
- Server object
 - caution when deleting 86
 - explained 2
- Server products, installing, uninstalling, and configuring 455
- Server utilities. *See* specific utility name
- Server. *See* NetWare server
- Service Advertising Protocol (SAP)
 - Directory tree mode for 594
 - method of time synchronization, explained 597
- SERVMAN NLM
 - purging deleted files 549
 - using to set file compression parameters 548
- SET login script command, using 226
- SET TIME utility, using 469
- SET TIMESYNC parameters, listed 594
- SET utility
 - controlling file compression with 544
 - displaying alert messages with 502
 - minimum cache buffers 785
 - minimum packet receive buffers 784
- SET_TIME login script command, using 228
- SFT III log files
 - explained 858
 - limiting size of 842
- SFT III server
 - clients, finding 847
 - configuration files, explained 836
 - console messages, logging 840
 - failure, troubleshooting 856
 - loading SBACKUP 695
 - memory, troubleshooting 837
 - testing 859

- Shell, NetWare, specifying version in login scripts 238
- SHELL_TYPE identifier variable, explained 238
- SHIFT login script command, using 229
- Short machine name, specifying in login scripts 238. *See also* Machine name
- SHORT_YEAR identifier variable, explained 237
- SHOW menu command, explained 268
- SHOW option for ITEM menu command, explained 263
- Signature level. *See* NCP packet signature
- SMACHINE identifier variable, explained 238. *See also* Machine name
- Sorting audit events
 - by date and time 677
 - by directory, example 671
 - by user activity, example 675
 - by user and file, example 667
 - for reports 655
 - instructions for 650 (*see also* Auditing)
- SRC file extension for menus 256
- Stack size, changing 693
- Standby MSL. *See* Alternate MSL boards, installing
- Standby power. *See* UPS monitoring
- STARTUP.NCF file, creating or editing 459
- STATION identifier variable, explained 238
- Storage, file. *See* Data migration; Optical media
- Subdirectory. *See* Directory
- Submenus, defining and numbering 261
- Subordinate replica, explained 285
- Subscribe, executing as part of login script 208
- Subsets of data to restore, specifying 728
- Subsystem. *See* Hard disk
- Supervisor object right, explained 19. *See also* Rights, object
- Supervisor property right, explained 21. *See also* Rights, property
- Surface test
 - for bad blocks 511
 - status explained 513
- SWAP login script command
 - compared to NOSWAP 221
 - using 231
 - using with # command 189
- Synchronization of network time. *See* Time synchronization
- SYS: volume
 - directories in, created automatically 105
 - required for NetWare 515
 - troubleshooting SFT III problems 852, 860
- SYS:DELETED.SAV directory, explained 105
- SYS:DOC directory, explained 106
- SYS:ETC directory, explained 105
- SYS:LOGIN directory, explained 105
- SYS:MAIL directory, explained 105
- SYS:PUBLIC directory
 - explained 106
 - mapping drive in login scripts for OS/2 215
 - mapping search drive in login scripts 216
- SYS:SYSTEM directory, explained 106
- SYS:SYSTEM/TSA/LOG directory, explained 762
- SYS_LOG.ERR file, viewing 506
- SYSICON utility. *See* NetWare Administrator; NETADMIN; specific task
- SYSTEM directory, explained 106
- System files, extracting from installation CD-ROM 457
- System login script. *See also* Login scripts
 - creating 180, 183
 - example 245
- System messages. *See* Alerts; *System Messages* manual

T

- Target Service Agent (TSA)
 - DOS workstation, command 691
 - DOS workstation, parameters 692
 - loading automatically 684
 - NDS, special instructions 684
 - NDS, when not to use 684
 - OS/2 workstation, command 691
- Target, backup or restore
 - changing targets 779
 - not listed in SBACKUP 712, 723, 749
- TEMP SET login script command. *See* SET login script command

- Template. *See* User template
- Test Mode, SFT III
 - recovery options, listed 857
 - troubleshooting 859
- Testing for bad blocks 511
- Text file
 - displaying during login 195
 - editing from server console 460
 - executing as subscript of login script 208
 - using to create users, with UIMPORT 334
- Text, displaying during login 233
- Time
 - specifying for delayed backup 719, 734
 - specifying in login scripts 237
 - workstation, synchronizing with server in login scripts 228
- Time and time zone
 - changing server time zone 470
 - setting server time 469
 - viewing server time 469
- Time out error, RCONSOLE, handling 618
- Time server type. *See* Time synchronization
- Time synchronization
 - adjusting time servers 600
 - configuration file, editing 595
 - creating custom configuration 598
 - custom configuration method, explained 597
 - example with two Directory trees 604
 - managing network 593
 - parameters listed 594
 - problems, avoiding 603
 - recognizing Directory tree structure 601
 - SAP method, explained 597
- TIME utility, using 469
- TIMESYNC NLM
 - automatically loaded 593
 - parameters listed 594
- TIMESYNC.CFG file, editing 595
- Token-Ring LAN driver, listed 450
- Traffic, viewing server 489
- Transaction Tracking System (TTS)
 - backing out of transactions automatically 589
 - enabling 589
 - logging backed-out data 590
- Tree. *See* Directory tree
- Troubleshooting
 - backup and restore sessions 786
 - hardware and network 789
 - memory 798
 - NetWare server 794
 - workstations 802
- Troubleshooting SFT III
 - I'm alive packets 854
 - IOEngine number 852
 - IPX internal network number 852
 - IPX route to secondary server 855
 - loader errors 841
 - lost interrupt alerts 841
 - missing RAM 838
 - MSEngine name 853
 - MSL boards 854
 - MSL problems 834
 - secondary server restarts 832
 - server failure 856
 - server memory 837
 - server Test Mode 859
 - SYS: volume problems 852, 860
 - two primary servers 834, 851
 - unknown commands 859
 - workstation problems 848
- Trust option, using with SBACKUP 692
- Trustee rights. *See also* Rights; Trustee
 - effective, viewing 153, 154
 - explained 131
- Trustee. *See also* Rights; Trustee rights
 - adding to directory or file 134, 135
 - assignment, explained 17
 - deleting from directory or file 137, 139
 - explained 17
 - list, explained 17
 - managing assignments to objects 22
 - modifying rights 140, 142
 - object, managing assignments 22
- TSA. *See* Target Service Agent (TSA)
- TTS. *See* Transaction Tracking System (TTS)

U

UIMPORT

- Control Parameters 340
- Control parameters 338
- Creating Import Control File 338
- Field definitions 338
- file examples 355

UIMPORT utility

- control file, creating 333
- control parameters, listed 340
- creating user objects 359
- examples of files 356
- explained 333
- field definitions, listed 345
- generating data file 334
- importing data to NDS with 333, 334
- importing process overview 333
- process, illustrated 334
- running after setup 359
- using 333

Unclaimed memory

- displaying 839
- setting new end address for 840
- used by MSEngine 839

Uninstalling server products 455

UNIX files

- adding name space for 522
- backing up linked 682
- removing name space for 524
- viewing information about 157, 158, 159

Unknown command error, troubleshooting SFT III 859

Unloading

- optical media 415
- SBACKUP files 781

Unmirroring hard disks 567

Updating replicas

- receiving 312
- sending 309, 310

Upgrading

- files on remote server 621
- login scripts 174

UPS monitoring

- changing settings for 580
- installing 579
- viewing status of 580

User ADMIN. *See* ADMIN User object

User login script. *See also* Login scripts

- creating 177, 180
- example 250
- explained 170
- NO_DEFAULT command, using 170

User object

- auditing network activities of 651 (*see also* Auditing)
- cautions when deleting 87
- changing property values 95, 100
- creating with UIMPORT 333, 359
- default rights, listed 3
- explained 40

User object ADMIN. *See* ADMIN User object

User object in login scripts

- full name, specifying 237
- ID number, specifying 238
- last name, specifying 237
- login name, specifying 237

User Space Limits property, changing in volume 98

User template

- control parameter for UIMPORT 344
- creating 69, 70
- explained 67

USER_ID identifier variable, explained 238

Utilities in login scripts

- DOS, mapping drive to 216
- executing 187
- OS/2, mapping drive to 215

Utilization

- cache, viewing 481
- processor, checking 488

V

Variables

- <variable> identifier variable, explained 238
- environment (DOS and OS/2), setting in login scripts 226
- GETP command, example 271
- login script, explained 187

VERIFY. *See* DOS VERIFY login script command

Virus infection, preventing 571

Volume

- adding name space support to 522
- changing size of 529
- creating 515
- creating for HCSS 402
- deleting 519
- dismount option, in auditing 649
- dismounting 518
- enabling auditing for 636 (*see also* Auditing)
- full alerts, setting 504
- mapping drive in login scripts 213
- mounting 518
- remounting, when dismounted by AUDITCON 665
- removing name space support from 524
- renaming 521
- repairing 525
- segment, adding 528
- using CD-ROM as 530

Volume object

- created automatically 515
- explained 41
- INSTALL cannot delete 520
- INSTALL cannot rename 522

Volume Restrictions property, changing 98, 102

Volume SYS:, troubleshooting SFT III problems 852, 860

VREPAIR NLM

- removing name space support 524
- repairing volumes 525

W

Warnings. *See* Alerts; *System Messages* manual

Watchdog parameters, changing 475

Windows

- files required to run with NetWare 120
- icon for NetWare Administrator, creating 10, 14
- installing on network 119
- mapping search drive in login scripts 216
- network installation methods, explained 119
- running on network 120

Windows workstation. *See also* Workstation

- installing Windows workstation software 9
- making client installation diskettes from CD-ROM 6

Workstation problems, troubleshooting SFT III 848

Workstation, specifying in login scripts

- clearing screen 192
- connection number 238
- machine name 211
- network address 238
- node address 238
- PCCOMPATIBLE command 199
- time, synchronizing with server 228

Workstation. *See also* DOS workstation; OS/2

- workstation; Windows workstation
- clearing connection 473
- increasing available memory 789
- network address, specifying in login scripts 238
- operating system directories, explained 107
- sending messages to 444
- troubleshooting 802

WRITE login script command, using 233

Write property right, explained 21. *See also* Rights, property

Y

YEAR identifier variable, explained 237

User Comments

We want to hear your comments and suggestions about this manual. Please send them to the following address:

Novell, Inc.
Technical Publications
MS C-23-1
122 East 1700 South
Provo, UT 84606
U.S.A

NetWare 4.1
*NetWare 4.1 Supervising the
Network*
Part #100-002075-001
December 1994

Fax: (801) 429-3002

For technical support issues, contact your local dealer.

Your name and title: _____

Company: _____

Address: _____

Phone number: _____ Fax: _____

I use this manual as an overview a tutorial a reference a guide _____

	Excellent	Good	Fair	Poor
Completeness	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Readability (style)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization/Format	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Accuracy	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Illustrations	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Usefulness	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please explain any of your ratings: _____

In what ways can this manual be improved? _____

You may photocopy this comment page as needed so that others can also send in comments.