

Advanced Novell® Network Management

COURSE 3005

Novell Education

www.novell.com

INSTRUCTOR GUIDE

Proprietary Statement

Copyright © 2003 Novell, Inc. All rights reserved.

No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express prior consent of the publisher. This manual, and any portion thereof, may not be copied without the express written permission of Novell, Inc.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606-2399

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of this manual, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose.

Further, Novell, Inc. reserves the right to revise this publication and to make changes in its content at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any NetWare software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose.

Further, Novell, Inc. reserves the right to make changes to any and all parts of NetWare software at any time, without obligation to notify any person or entity of such changes.

This Novell Training Manual is published solely to instruct students in the use of Novell networking software. Although third-party application software packages are used in Novell training courses, this is for demonstration purposes only and shall not constitute an endorsement of any of these software applications.

Further, Novell, Inc. does not represent itself as having any particular expertise in these application software packages and any use by students of the same shall be done at the students' own risk.

Software Piracy

Throughout the world, unauthorized duplication of software is subject to both criminal and civil penalties.

If you know of illegal copying of software, contact your local Software Antipiracy Hotline.

For the Hotline number for your area, access Novell's World Wide Web page at <http://www.novell.com> and look for the piracy page under "Programs."

Or, contact Novell's anti-piracy headquarters in the U.S. at 800-PIRATES (747-2837) or 801-861-7101.

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, Inc. Trademarks

NetWare, the N-Design, and Novell are registered trademarks of Novell, Inc. in the United States and other countries.

CNA, CDE, CNA are service marks and CNE is a registered service mark of Novell, Inc. in the United States and other countries.

Certified Novell Administrator, Certified Novell Engineer, eDirectory, IPX, NCP, NetWare Core Protocol, NLM, Novell Certificate Server, Novell Client, Novell Cluster Services, Novell Distributed Print Services, Novell iFolder, and Novell Storage Services are trademarks of Novell, Inc.

ConsoleOne, GroupWise, ManageWise, NDPS, NDS, Novell Directory Services, and ZENworks are registered trademarks of Novell, Inc.

Other Trademarks

Adaptec is a registered trademark of Adaptec, Inc. AMD is a trademark of Advanced Micro Devices. Pentium is a registered trademark of Intel Corporation. Windows is a registered trademark of Microsoft Corporation. NetScanTools is a trademark of Northwest Performance Software, Inc. Oracle is a registered trademark of Oracle Corporation. Btrieve is a registered trademark of Pervasive Software, Inc. Norton AntiVirus is a trademark of Symantec Corporation. VMware is a trademark of VMware, Inc.

Contents

Course Setup

Set Up the Classroom	Setup-1
Setup Requirements	Setup-2
Setup Diagram	Setup-5
Setup Time-Saving Procedure	Setup-6
Setup Instructions	Setup-7
Setup for Section 12	Setup-50
Customize the Exercise to Meet Student Needs	Setup-50
Set Up the Network for the Exercise	Setup-51

Introduction

Audience	Intro-1
Prerequisites	Intro-1
Certification	Intro-2
Relationship to Other Courses in the Curriculum	Intro-2
Agenda	Intro-2
Course Feedback	Intro-4
Scenario	Intro-4
Digital Airlines Company Information	Intro-4
Digital Airlines eDirectory Tree	Intro-7
Digital Airlines NetWare 6 Implementation Plan	Intro-8

MODULE 1	Migrate NetWare 4 and NetWare 5 Servers to NetWare 6	
SECTION 1	Migrate NetWare 4 and NetWare 5 Servers to NetWare 6	
	Objectives	1-1
	Introduction	1-1
	In-Place Upgrade	1-2
	NetWare Accelerated Upgrade	1-3
Objective 1	Review How to Prepare for a Server	
	Migration to NetWare 6	1-4
	Prepare the Workstation	1-5
	Prepare the Source (Original) Server	1-6
	Prepare the Destination (New) Server	1-7
	Prepare Server Application Files	1-9
Objective 2	Review How to Implement Novell Licensing	1-10
	Server and User Licensing Models	1-11
	How UAL Coexists with SCL	1-13
	How the Licensing Models Differ	1-15
	License Types	1-16
Objective 3	Identify How to Perform a Migration.	1-18
Objective 4	Perform Post-Migration Tasks	1-39
	Run the External Reference Check Process	1-39
	Upgrade Existing NSS Volumes	1-40
	Perform Other Post-Installation Tasks	1-41
	Exercise 1-1 Upgrade a NetWare 4.11 Server to NetWare 6	1-42
Objective 5	Troubleshoot Post-Installation Issues.	1-65
	Summary	1-68

MODULE 2 Troubleshoot and Resolve Novell Network Problems

SECTION 2 Identify Tools for Troubleshooting Novell Network Performance Issues

	Objectives	2-1
	Introduction	2-1
Objective 1	Upgrade Novell Network Management Tools	2-3
	Exercise 2-1 Upgrade Your Novell Network Management Tools	2-5
Objective 2	Identify the Troubleshooting Features of Novell NetWork Management Tools	2-8
	ConsoleOne Reports	2-9
	Novell iMonitor	2-11
	NetWare Remote Manager	2-12
	Novell iManager	2-14
Objective 3	Identify the Purpose and Function of IP/IPX Troubleshooting Tools	2-17
	NetWare IP/IPX Troubleshooting Tools	2-17
	Client IP Troubleshooting Tools	2-25
	TCP/IP Troubleshooting Example	2-42
	Exercise 2-2 Test Your Network	2-44
	Protocol Analyzers	2-47
	TCP/IP Toolkits	2-51
	IP Addressing Calculators	2-54
Objective 4	Identify Additional Network Troubleshooting Resources . . .	2-57
	Novell Web Site Resources	2-57
	Shareware and Freeware	2-59
	Summary	2-60
	Exercise Answers	2-61

SECTION 3	Troubleshoot and Resolve NetWare Server Issues	
	Objectives	3-1
	Introduction	3-1
Objective 1	Identify Server Hardware and Operating System Components	3-2
	Identify Server Hardware	3-2
	Identify Operating System Components	3-8
	Exercise 3-1 Determine Hardware and Operating System Components	3-18
Objective 2	Troubleshoot and Resolve NetWare Server Issues	3-21
	Identify the Top Novell Technical Support Server Issues and How to Resolve Them	3-21
	Identify Problems after Installation	3-22
	Resolve Console Lock Ups	3-23
	Resolve Hard Disk Errors and Access Problems	3-24
	Resolve Application Monopolizing Server CPU	3-25
	Resolve Server Memory Problems	3-26
	Resolve Slow Server Response	3-30
	Identify Multiprocessing Problems	3-33
	Find Tools for Managing Servers	3-35
	Exercise 3-2 Resolve Server Problems	3-39
Objective 3	Troubleshoot and Resolve Critical Server Abends	3-47
	What an Abend Is	3-48
	What Types of Abends Occur	3-49
	What an ABEND.LOG File Is	3-50
	What a Core Dump Is	3-51
	How to Respond to an Abend	3-52
	How to Create and Submit a Core Dump for Analysis	3-57
	Exercise 3-3 Submit an ABEND.LOG File for Analysis	3-64
	Exercise 3-4 Create a Core Dump	3-66

Objective 4	Troubleshoot and Resolve Server Communication Issues . . .	3-70
	Resolve Server-to-Server Communication Problems	3-71
	Resolve Workstation-to-Server Communication Problems . . .	3-72
	Identify Preventative Maintenance Tasks	3-73
	Exercise 3-5 Resolve Communication Problems	3-76
	Summary	3-79
	Exercise Answers	3-82
SECTION 4	Monitor and Troubleshoot eDirectory	
	Objectives	4-1
	Introduction	4-1
Objective 1	Identify eDirectory Databases and Processes	4-2
	eDirectory 8.7 Databases	4-2
	eDirectory Processes	4-4
	Post-Migration or Upgrade Issues That Affect eDirectory Databases and Processes	4-6
Objective 2	Identify eDirectory Troubleshooting Steps	4-8
Objective 3	Identify Partition and Replication Placement Design	4-15
	Exercise 4-1 Adding Replicas with iManager	4-17
Objective 4	Use iMonitor Reports to Obtain Server and eDirectory Information	4-20
	How to Review Report Options	4-21
	How to Run a Report	4-23
	How to View Saved Reports	4-25
	Exercise 4-2 Verify eDirectory Status Using Reports	4-26

Objective 5	Perform Health Checks	4-28
	Health Check Items	4-28
	iMonitor Health Check Features	4-30
	iMonitor Health Check Procedure	4-30
	How to Run the Agent Health Report	4-34
	How to Perform a Trace with iMonitor	4-39
	How to Perform Directory Service Repair	4-43
	Exercise 4-3 Verify Network Health	4-44
	Exercise 4-4 Evaluate an eDirectory Problem	4-49
	Summary	4-61
	Exercise Answers	4-64

MODULE 3 Demonstrate Advanced Novell Network Storage Management Skills

SECTION 5 Perform Advanced Novell Storage Services Tasks

	Objectives	5-1
	Introduction	5-1
	The Scenario	5-2
Objective 1	Expand an NSS Storage Space	5-2
	Increase the Number of Logical Volumes in a Storage Pool	5-2
	Overbook the Storage Pool	5-3
	Increase the Size of a Storage Pool	5-4
Objective 2	Configure NSS Volume Attributes	5-9
Objective 3	Mount a DOS Partition as an NSS Volume	5-11
	How DOSFAT.NSS Works	5-12
	How to Mount a DOS Partition Using DOSFAT.NSS	5-13
	How to Verify That the DOS Partition Is Mounted	5-14

Objective 4	Use VCU to Create an NSS Volume from a Traditional Netware Volume	5-15
	How VCU Works	5-15
	How to Copy a Traditional Volume to an NSS Volume	5-16
	The Correct VCU Syntax	5-17
	Exercise 5-1 Perform Advanced NSS Storage Management Tasks	5-18
Objective 5	Resolve Common NSS Errors	5-20
	Determine the Cause of the Problem and List Possible Solutions	5-20
	Use VERIFY to Determine the Integrity of an NSS Pool	5-23
	Assess Possible Solutions	5-23
	Use REBUILD as a Last Resort Solution	5-24
	Use Third-Party Software or Services to Recover Data	5-25
	Exercise 5-2 Resolve NSS Error Codes	5-25
	Part I: Research an Error Code and Find a Possible Solution	5-26
	Part II: Implement the Solution	5-26
Objective 6	Restore a Deleted Logical Volume.	5-27
Objective 7	Describe Storage Area Networks and Network Attached Storage	5-29
	Identify How a SAN Works	5-29
	Identify How NAS Works	5-31
	List SAN and NAS Design and Implementation Considerations	5-33
	Summary	5-34

SECTION 6	Configure and Troubleshoot a RAID Solution Using NSS	
	Objectives	6-1
	Introduction	6-1
Objective 1	Implement RAID 0 with NSS	6-2
	What RAID Is	6-2
	Hardware RAID versus Software RAID	6-2
	Hardware and Software RAID Levels	6-3
	How to Configure Software RAID 0 in NSS	6-6
	How to Use NSSMU to Create RAID Arrays	6-10
	Exercise 6-1 Configure a Software RAID Solution	6-10
Objective 2	Configure Partition Mirroring and Duplexing in NSS	6-15
	Configure Partition Mirroring and Duplexing in NSS	6-16
	Troubleshoot Software RAID and Mirroring in NSS	6-19
	Exercise 6-2 Mirror an NSS partition	6-20
	Summary	6-22
SECTION 7	Perform Advanced iFolder Tasks and Troubleshooting	
	Objectives	7-1
Objective 1	Describe iFolder Configuration Files	7-1
	How to Edit iFolder Configuration Files	7-2
	When to Edit iFolder Configuration Files	7-9
Objective 2	Perform iFolder Management Tasks	7-14
	Stop and Start the iFolder Server	7-14
	Set iFolder Client and Server Policies	7-16
	Change the Location of iFolder User Data	7-17
	Add Contexts	7-18
	Add Additional Administrators	7-19
	Change the iFolder Server IP Address	7-20
	Exercise 7-1 Perform Advanced iFolder Management Tasks	7-22

Objective 3	Maintain and Troubleshoot the iFolder Client	7-40
	The Benefit of the iFolder Client	7-40
	How the iFolder Client Works	7-40
	Common Issues Involving the Client	7-43
Objective 4	Maintain and Troubleshoot the iFolder Server	7-46
	Adjust the Number of Threads per Child	7-47
	The Admin Cannot Access the Server Management Console	7-48
	Port Conflict with iPrint Secure Port	7-49
	Restoring User Accounts When Pass Phrases Are Forgotten	7-50
	LDAP Incorrectly Configured for Non-secure Port	7-53
	Summary	7-54
MODULE 4	Deliver High Availability Services with Novell Cluster Services	
SECTION 8	Design and Set Up an NCS Cluster Configuration	
	Objectives	8-1
	Introduction	8-1
Objective 1	Identify the Purpose and Advantages of Implementing an NCS Solution	8-2
	High Availability Terms	8-3
	High Availability Definition	8-4
	Computer System Outage Factors	8-7
	Benefits and Features of an NCS High Availability Solution	8-8

Objective 2	Design and Set Up an NCS Cluster Configuration	8-10
	Basic Clustering System Terms	8-10
	NCS Cluster Components	8-12
	Typical NCS Shared Disk System Cluster Configurations . .	8-13
	NCS System Terms	8-15
	Rules for Managing an NCS SCSI SAN	8-27
	Troubleshooting a 2-Node NCS SCSI SAN	8-28
	Exercise 8-1 Design and Set Up a 2-Node SCSI Clustering Configuration	8-32
	Summary	8-40
	Exercise Answers	8-42
SECTION 9	Install and Test NCS on a 2-Node Cluster	
	Objectives	9-1
	Introduction	9-1
Objective 1	Verify NCS System Requirements	9-2
	Hardware Requirements	9-2
	Software Requirements	9-3
	License Requirements	9-3
	Shared Disk System Requirements	9-4
Objective 2	Create a Cluster by Installing NCS	9-5
Objective 3	Check Cluster Configuration Settings	9-13
	Cluster ADMIN Object	9-14
	Cluster Object	9-14
	Master IP Address Resource Object	9-23
	Cluster Server Node Objects	9-24
	Exercise 9-1 Install and Check NCS on a 2-Node Cluster . . .	9-25
Objective 4	Test and Monitor the Cluster	9-34
	Cluster State and Cluster Status Views	9-34
	Console Prompt Commands	9-38
	Exercise 9-2 Test the SBD Partition and Heartbeats	9-39

	Summary	9-46
	Exercise Answers	9-49
SECTION 10	Configure and Test High Availability File Access	
	Objectives	10-1
	Introduction	10-1
Objective 1	Configure NCS for High Availability File Access	10-2
	Create a Shared Disk Partition	10-2
	Create and Cluster-Enable an NSS Volume and Pool on a Shared Storage Device	10-3
	Cluster-Enable an Existing Pool or Volume on the Shared Disk System	10-6
	Exercise 10-1 Create a Cluster-Enabled Volume for High Availability File Access	10-10
Objective 2	Manage Resources in an NCS Cluster	10-18
	How to Migrate Resources	10-18
	How to Troubleshoot Resource States	10-20
	Exercise 10-2 Test High Availability File Access on the 2-Node Cluster	10-23
	Summary	10-29
SECTION 11	Configure and Test High Availability Services	
	Objectives	11-1
	Introduction	11-1

Objective 1	Identify Cluster-Aware and Cluster-Naive Applications . . .	11-2
Objective 2	Identify How to Cluster-Enable an Application	11-3
Objective 3	Identify How to Assign Nodes to a Resource	11-5
Objective 4	Identify How to Set Start, Failover, and Failback Modes . . .	11-6
Objective 5	Identify How to View and Edit Load and Unload Scripts . . .	11-8
Objective 6	Identify How to Find NCS Configuration and Troubleshooting Information	11-10
	Exercise 11-1 (Optional) Cluster-Enable and Test DHCP Server on Your 2-Node Cluster.	11-13
	Exercise 11-2 Cluster-Enable and Test iFolder on Your 2-Node Cluster.	11-20
	Summary	11-30
	Exercise Answers	11-32

MODULE 5 Troubleshoot a NetWare 6 Network

SECTION 12 Troubleshoot a NetWare 6 Network

	Objectives	12-1
	Introduction	12-1
Objective 1	Create a Disaster Recovery Plan	12-1
	Planning for Hardware Failures	12-2
	Planning for Calamities	12-3
Objective 2	Troubleshoot Network Problems	12-4
	Exercise 12-1 Troubleshoot Network Problems	12-4
	Summary	12-5

APPENDIX A	Network Components	
APPENDIX B	The Network Communication Process	
	ISO Layers and the Communication Process	B-1
	IP Routing	B-4
APPENDIX C	Protocol Analyzers	
	Protocol Analyzer Elements	C-1
	Protocol Analyzer Types	C-6
	Standalone Analyzers	C-6
	Distributed Protocol Analyzers	C-7
	Hardware and Software Analyzers	C-8
	Protocol Analyzer Placement	C-9
	Hubbed Network	C-9
	Bridged Network	C-10
	Switched Network	C-11
	Routed Network	C-13
	WAN Links	C-14

Index

Course Setup

Read this section before you set up instructor and student workstations and servers.



This course was tested for NetWare® 6 with Support Pack 2. If you use a later support pack, you might need to adjust exercise steps.

Set Up the Classroom

Classroom setup can take several hours. Allow adequate time to set up the hardware and software and to test the setup.

As part of your testing, complete the exercises to ensure that students can complete the exercises without unexpected problems.

To set up your classroom, use the following:

- [Setup Requirements](#)
- [The following is a diagram of the classroom setup for the course:](#)
- [Setup Time-Saving Procedure](#)
- [Setup Instructions](#)

Setup Requirements

The following are the classroom setup requirements:

Table Setup-1

Setup	Minimum Requirements
Setup Time	<ul style="list-style-type: none"> ■ 8 hours for setup (one task at a time sequentially) ■ 4 hours for setup (following the procedure under "Setup Time-Saving Procedure" on Setup-6) ■ .5 hours for testing
Servers	8 servers: 6 for students; 2 for instructor
Hardware	<ul style="list-style-type: none"> ■ Pentium® II 266 Mhz ■ 512 MB RAM ■ CD drive ■ 3.5-inch diskette drive ■ Ethernet board ■ DA1 and DA4 - DA9 servers (see Figure Setup-1) require 2 hard disks with 8 GB space each ■ DA1, DA2, and DA4 - DA9 servers (see Figure Setup-1) need Adaptec® 2940 or 29160 SCSI adaptors. ■ DA2 server requires 2 network boards. ■ External SCSI Drive (2 GB) housed in an external SCSI enclosure for every 2 servers (including DA1 and DA2) ■ 2 SCSI cables to connect the Adaptec adapters from every 2 servers to the 2 connectors on the back of the SCSI enclosure

Table Setup-1 (continued)

Setup	Minimum Requirements
Software	<ul style="list-style-type: none"> ■ NetWare 6 OS CD ■ NetWare 6 Support Pack 2 CD ■ eDirectory™ 8.7 CD ■ eDirectory Webapps CD ■ 3005.LDIF (on the Enhanced Learning CD or CNI Net) ■ 3005LicenseFolders.EXE (NetWare 6 server and user licenses in DA1\DATA\SETUP on the DA1 VMware™ server or CNI Net) or the DAx license folders on the Enhanced Learning CD (Setup)
Workstations	7 workstations: 6 for students; 1 for instructor
Hardware (students)	A Pentium II 266 Mhz computer, with <ul style="list-style-type: none"> ■ 128 MB RAM ■ 2 GB hard disk space ■ Ethernet board ■ CD drive ■ 3.5-inch disk drive
Hardware (instructor)	Pentium III 750 Mhz computer, with <ul style="list-style-type: none"> ■ 1 GB RAM ■ 10 GB hard disk space ■ Ethernet board ■ CD drive ■ 3.5-inch disk drive

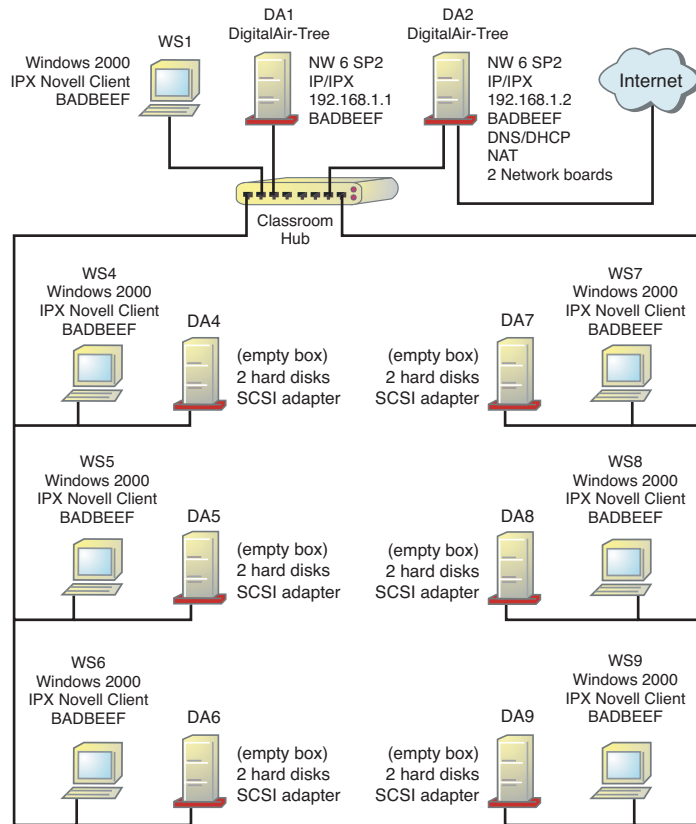
Table Setup-1 (continued)

Setup	Minimum Requirements
Software (students and instructor)	<ul style="list-style-type: none"> ■ Windows® 2000 Professional with Service Pack 3 ■ Internet Explorer 6 ■ Novell Client™ for Windows 2000 (start the course with version 4.81, from the NetWare 6 Client CD) ■ Novell Client for Windows 2000 4.83 or later (used for Section 1) ■ ConsoleOne® CD ■ ConsoleOne snap-ins for NetWare 6 ■ NCI Client 2.4 CD ■ WinZip
Software (instructor only)	<ul style="list-style-type: none"> ■ VMware Workstation installed on the instructor's workstation (evaluation copy and license available at www.vmware.com) ■ VMware servers DA1 and DA3-DA9 (NetWare 4.11) installed on the instructor's workstation (available for download from CNI Net, on DVD from distributor or student kit)

Setup Diagram

The following is a diagram of the classroom setup for the course:

Figure Setup-1 (slide)



Arrange placement of each pair of servers (DA1 and DA2, DA4 and DA5, DA6 and DA7, DA8 and DA9) so they are close enough for the SCSI cables to reach from each server to the SCSI hard disk for clustering.

Setup Time-Saving Procedure

The tasks listed under “[Setup Instructions](#)” on [Setup-7](#) provide a complete set of steps for setting up the classroom.

Although these tasks are listed consecutively, you can save hours by performing these tasks as follows:

1. Install Windows 2000 on the instructor workstation (WS1) and a student workstation (WS4).

You can use a Microsoft scripted installation file (SIF) to run the installation unattended while performing other setup tasks.



For details on SIF files, see <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windows2000pro/deploy/depop/easydep.asp>.

2. While the Windows 2000 installations are running, complete a NetWare 6 pre-migration installation on server TEMP1 (for VMware server DA1).

At this point, do not migrate the server; only perform the pre-migration installation.
3. While you are waiting for the NetWare 6 files to copy to the servers, return to WS1, make sure Windows 2000 is completely installed, and begin copying the VMware server files to the workstation.

Keep the copying process running until all VMware server files are copied in the correct folders.
4. While the VMware servers are copying on WS1, return to WS4 and complete any remaining steps under “[Set Up Student Workstations](#)” on [Setup-47](#) to prepare the workstation for the course.
5. With WS4 prepared, use imaging software to image workstations WS5 through WS9 from WS4.

After imaging the workstations, make sure each workstation has a different SID (Windows 2000/NT Security ID).

6. Make sure all VMware servers are copied to WS1; then complete remaining steps under “[Set Up the Instructor Workstation](#)” on [Setup-8](#).

At this point, all workstations are set up and ready, and the VMware servers are on and running.

7. Finish the steps under “[Set Up the Instructor Servers](#)” on [Setup-14](#) to migrate DA1; install DA2; and update services and management tools.

At this point, the instructor servers are ready.

8. Follow the steps under “[Set Up Student Servers](#)” on [Setup-46](#) and “[Test the Setup](#)” on [Setup-49](#).

You have completed the course setup.

Setup Instructions

Complete the following:

- [Gather Necessary CDs](#)
- [Set Up the Instructor Workstation](#)
- [Set Up the Instructor Servers](#)
- [Set Up Student Servers](#)
- [Set Up Student Workstations](#)
- [Test the Setup](#)

Gather Necessary CDs

You need the following SEL CDs to perform this setup:

- January 2002 CD6A and CD7B (NetWare 6 OS and Novell Client 4.81)
- April 2002 CD3A (ConsoleOne 1.3.3)
- August 2002 CD4A (NICI 2.4)
- November 2002 CD2A and CD2B (eDirectory 8.7 and Web Apps)

In addition, you need to have the following:

- NetWare 6 Support Pack 2 CD

You can create this CD by downloading the files from <http://support.novell.com/tools/csp/csplist.html>.

- NetWare 4.11 VMware server files on DVD or CD

These are available on DVD from a distributor (such as Kinko's) or in the student kit.

You can also download these files from CNI Net and create your own CDs (or DVD).

Set Up the Instructor Workstation

Perform the following on the instructor workstation:

1. Install Windows 2000 Professional using the following details:
 - Partition type: **NTFS**
 - Computer name: **WS1**
 - Default log in: **Administrator**
 - Administrator password: **novell**
 - OS update: **Service Pack 3 for Windows 2000** (or later)
 - Do not backup files.

- ❑ IP address: **192.168.1.21**
- ❑ Subnet mask: **255.255.255.0**



You configure your workstation to use DHCP after DHCP is installed and configured on server DA2.

2. Install **Internet Explorer 6** or later on the workstation.
3. Turn off the Friendly HTTP error messages option in Internet Explorer:
 - a. From the Internet Explorer menu bar, select **Tools > Internet Options**.
 - b. Select the **Advanced** tab.
 - c. Scroll down and deselect **Show friendly HTTP error messages**.
 - d. Select **Apply**.
4. Set up the home page in Internet Explorer to show Web Manager:
 - a. In the Internet Options dialog, select the **General** tab.
 - b. In the Address field under Home Page, enter **https://192.168.1.1:2200**.
 - c. Select **Apply**.
 - d. Select **OK**.
5. Install Novell Client version 4.81:



Install Novell Client version 4.81 so students can upgrade to 4.83 in a later exercise. Do not install the latest version at this time.

- a. Insert the **Novell Client 4.81** CD (January 2002 SEL CD7B).
- b. Select *your language*.

- c. After the CD autoruns, select **Novell Client 4.81 for Windows NT/2000/XP**.
 - d. In the License Agreement screen, select **Yes**.
 - e. Select **Custom Installation**; then select **Next**.
 - f. Deselect all components, if any are selected; then select **Next**.
 - g. Select **IP and IPX**; then select **Next**.
 - h. Select **NDS (NetWare 4.x or later)**; then select **Next**.
 - i. Select **Finish**.
 - j. When the installation is complete, select **Reboot**.
6. Install NICI 2.4:
- a. Insert the **NICI** CD (August 2002 SEL CD4A).
 - b. Browse to and double-click **NICI24\Win_all\WCNICIU0**.
 - c. In the Welcome window, select **Next**.
 - d. In the License Agreement window, select **Yes**.
 - e. In the Setup Complete window, select **Finish**.
7. Install WinZip according to the manufacturer's instructions.
8. Install ConsoleOne 1.3.3:
- a. Insert the **ConsoleOne** CD (April 2002 SEL CD3A).
 - b. Browse to and double-click **ConsoleOne133\C1_NW_Win**.
 - c. In the WinZip self-extractor window, select **Setup**.
 - d. In the Welcome window, select **Next**.
 - e. In the License Agreement window, select **Accept**.
 - f. In the Location window, accept the default location by selecting **Next**.
 - g. In the ConsoleOne Components window, make sure **Reporting Snapins** is selected; then select **Next**.
 - h. In the Next window, select *your language* (English is always installed); then select **Next**.

- i. In the Summary window, select **Finish**.
 - j. When the installation is complete, select **Close**.
9. Install ConsoleOne snap-ins:
- a. Insert the **Enhanced Learning** CD in WS1.
 - b. Locate and run
D:\SETUP\CONSOLEONE 1.3.3 SNAPINS FOR 3005.EXE
(This file is also available from CNI Net or in DATA:\SETUP on your DA1 VMware server.)
 - c. Copy the snap-in files to the default location (C:\NOVELL\CONSOLEONE\1.2) on your hard drive by selecting **UNZIP**.
10. On your administrator workstation (WS1) install VMware Workstation according to the product instructions (follow the wizard and accept the defaults).
11. (Conditional) If you are installing VMware Workstation 3.2, fix the conflict with Novell Client that causes the workstation to lock up after logging in.
- To fix this problem, you configure the VMware Virtual adaptors on the workstation to use static IP addresses.
- Do the following:
- a. When you reboot the workstation after installing VMware Workstation, press **F8** to start Windows 2000 in **Safe Mode with Networking**.
 - b. in Notepad, open
C:\WINNT\SYSTEM32\VMNETDHCP.CONF.
 - c. Search for the fixed IP addresses for host **VMnet1** and host **VMnet8**; then record the addresses:

Table Setup-2

VMnet1 IP Address	VMnet8 IP Address

- d. Exit Notepad without saving changes to **VMNETDHCP.CONF**.
- e. Select **Start > Settings > Network and Dial-up Connections**.
- f. Right-click **VMware Virtual Ethernet Adapter (basic host-only support for VMnet1)**; then select **Properties**.
- g. Select **Internet Protocol (TCP/IP)**; then select **Properties**.
- h. Enter the *VMnet1 IP address* from the **VMNETDNCP.CONF** file; then for the subnet mask enter **255.255.255.0**.
- i. Select **OK** twice.
- j. Repeat Steps **f - i** for **VMware Virtual Ethernet Adapter (Network Address Translation (NAT) for VMnet8)** using *VMnet8 IP address*.
- k. Restart the computer and log in.
You should now be able to run both the Novell Client and VMware workstation without conflicts.



For additional VMware support information, check the Knowledge Base and news groups available at www.vmware.com.

12. Copy VMware server files:

- a. On WS1 create the following directory:
C:\PROGRAM FILES\VMWARE\DIGITALAIR
- b. Do one of the following:
 - If you are using CDs (created from files available on CNI Net), extract the contents of **DA1VM.EXE** (the **DA1** folder) on your **DA1VM** CD to the **DIGITALAIR** directory you created in Step **a**. Repeat this process for CDs **DA3VM**, **DA4VM - DA9VM**, extracting and copying folders **DA3** and **DA4-DA9**.

- If you are using the Course 3005 VMware Servers DVD, copy the **DAVM1.EXE** and **DAVM2.EXE** files to a temporary directory on your hard drive; then open and unzip the file folders to the **DIGITALAIR** directory you created in Step **a**.
- c. Remove the **Read-only** attribute from all VMware server folders and files copied from the CDs or DVD.



The DA1VM — DA9VM CDs are those you create by downloading the VMware server files from CNI Net, and then burning the CDs. The DVD is available through distribution channels (such as Kinko's) or in the student kit.

You can also download the VMware server files to a computer or portable storage device with enough space (650 MB per VMware server); then copy the files from that computer or storage device to WS1.

13. Launch and turn on VMware servers (1, 3-9):

- a. From your administrator workstation desktop, start VMware Workstation.
- b. Select **File > Open**; then browse to and double-click **C:\PROGRAM FILES\VMWARE\DIGITALAIR\DA1\DA1.VMX**.
Server DA1 is listed in the left column.
- c. Select **DA1**; then select the **Power On** button.
- d. Wait for the DA1 NetWare 4.11 server to load; then repeat steps **a - c** for **C:\PROGRAM FILES\VMWARE\DIGITALAIR\DA3\DA3.VMX** and all other VMware servers.



Any VMware servers you do not migrate during Setup and Section 1 must stay on and running throughout the course. The servers are included in the eDirectory tree and are part of the replica ring.

Set Up the Instructor Servers

Do the following to set up servers for the instructor:

- [Create License Diskettes](#)
- [Prepare the NetWork for NetWare 6:](#)
- [Set Up Server TEMP1](#)
- [Migrate DA1 to TEMP1](#)
- [Upgrade eDirectory and Management Tools for DA1](#)
- [Set Up Server DA2](#)
- [Upgrade eDirectory and Management Tools for DA2](#)
- [Install User Licenses for DA1 and DA2](#)
- [Configure DNS/DHCP on DA2](#)
- [Import 3005LDIF](#)
- [Configure NAT on DA2](#)
- [Prepare for Clustering](#)

Create License Diskettes

Do the following:

1. Locate the **3005LicenseFolders.EXE** file from one of the following:
 - CNI Net
 - DA1 VMware server (DATA:\SETUP)
2. Double-click the file and extract the **DAx** folders to a temporary directory.

(You can also copy the DAx folders from the SETUP\LICENSE FILES directory on the Enhanced Learning CD).

3. Copy folders **DA1** and **DA2** (for the instructor servers) and **DA4** to **DA9** (for the student servers) to individual floppy disks labeled with the same name.

Prepare the NetWork for NetWare 6:

Do the following:

1. At the instructor workstation, insert the **NetWare 6 OS CD** (January 2002 SEL CD6A).
2. If it does not autorun, browse to the CD and launch **NWDEPLOY.EXE**.
3. Double-click **Network Preparation**.
4. Double-click **Step 2: View and Update NDS Versions**.
5. Enter **netware://Tree/DIGITALAIR-TREE**.
6. Select **Include Subordinate Containers**; then select **Next**.
You should find that no servers need to be updated.
7. Select **Exit**.
8. Double-click **Step 3: Prepare for NDS eDirectory 8.6**.
9. Enter **netware://Tree/DIGITALAIR-TREE**.
10. Select **Next**.
11. Select **DA1**; then select **Next**.
12. When the update is complete, select **Exit**.
13. Close Deployment Manager by selecting **Cancel** and then **Yes**.

Set Up Server TEMP1

Do the following:

1. Begin a pre-migration installation of NetWare 6:
 - a. Start the computer and insert the **NetWare 6 OS CD**.

- b. When prompted, reboot by selecting one of the following:
 - To install from your IDE CD, press **I**.
 - To install from your SCSI CD, press **S**.
 - To install both, press **B**.
 - c. Install your language by selecting *your language*.
 - d. In the Welcome to NetWare Server Installation window, use your arrow keys to select **Accept License Agreement**.
 - e. Select **Create a New Bootable Partition**.
 - f. In the First Hard Disk window, select enough space for your DOS partition using the following formula:
200 + amount of RAM in machine
 - g. Select **Continue**.
 - h. Verify you want to create a boot partition by using the arrow keys to select **Continue**.
 - i. Reboot your computer by pressing any key on the keyboard.
 - j. Allow your computer to reboot.
 - k. When prompted, reboot by selecting one of the following:
 - To install from your IDE CD, press **I**.
 - To install from your SCSI CD, press **S**.
 - To install both, press **B**.
2. Configure basic installation parameters:
- a. In the License Agreement for Jreport Runtime screen, press **F10** to accept the license.
 - b. In the Welcome screen, select **Express**; then press **Enter** to switch to **CUSTOM**.
 - c. In the Welcome screen, select **New Server**; then press **Enter** to switch to **PRE-MIGRATION**.
 - d. Press **Tab**.
 - e. Select **Continue**.
 - f. In the Pre-Migration Installation screen, select **Continue**.

3. Continue the pre-migration installation using the following information:
 - ❑ Volume SYS: **2009 MB** (NetWare partition size)
 - ❑ Server name: **TEMP1**
 - ❑ IP address: **192.168.1.1**
 - ❑ Subnet mask: **255.255.255.0**
 - ❑ Protocol: **IPX**
 - ❑ Host and domain name: **DA1.digitalairlines.com**
 - ❑ Domain name server: **192.168.1.2**
 - ❑ License diskette: **DA1 (SERVER LICENSE directory)**
 - ❑ Volume: **DATA (traditional, 2000 MB)**
 - ❑ Tree: **TEMPTREE1**
 - ❑ Server and admin context: **O=TEMP**
 - ❑ Password: **novell**
4. In the NDS Summary screen, select **Next**.
5. Configure licensing and LDAP:
 - a. In the Licenses screen, select **Next**.
 - b. In the LDAP Configuration screen, select **Next**.
6. In the Summary screen, select **Finish** and allow the files to copy.
7. When prompted that the installation is complete, remove the Installation CD from the server; then select **Yes** and allow the server to reboot.
8. Make sure the GUI resolution is set to **800 x 600**.

Migrate DA1 to TEMP1

After all VMware servers are on and running NetWare 4.11, prepare DA1 by doing the following:

1. Using DSREPAIR, run the **Unattended Full Repair and Report Synchronization Status** options:

- a. At the server prompt, enter **LOAD DSREPAIR**.
 - b. In Available Options, select **Unattended Full Repair**.
 - c. In Available Options, select **Report Synchronization Status**.
 - d. Press **Esc**.
 - e. Select **Time Synchronization**.
If time is not synchronized for all servers, troubleshoot it now.
 - f. Close DSREPAIR.
2. Install the migration wizard on your workstation:
- a. On WS1, select **Start > Run**.
 - b. Select **Browse**.
 - c. On the NetWare 6 OS CD, browse to **PRODUCTS\MIGRTWZD\MIGRTWZD.EXE**.
 - d. Select **Open**.
 - e. Select **OK**.
 - f. Allow the files to extract to your workstation.
 - g. In the Choose Setup Language dialog, select *your language*; then select **OK**.
 - h. In the Welcome window, select **Next**.
 - i. In the Software License Agreement window, select **Yes**.
 - j. In the Choose Destination Location window, accept the default location by selecting **Next**.
 - k. Allow the migration wizard to install.
 - l. In the Setup Complete window, select **Finish**.
3. Attach to both servers using IPX:
- a. At the TEMP1 server console, enter **CONFIG** and note the *board name* that the IP address is bound to.
 - b. Enter **UNBIND IP *board_name***.

- c. At the server console, enter **CONFIG**.
You should see that IP is no longer bound to your network board.
 - d. From workstation WS1, right-click the red **N** icon in the system tray; then select **NetWare Login**.
 - e. In the Novell Login window, select **Advanced**.
 - f. Enter the following information; then select **OK**:
 - Username: **admin**
 - Password: **novell**
 - Tree: **DIGITALAIR-TREE**
 - Context: **DIGITALAIR**
 - Server: **DA1**
 - g. Right-click the red **N** icon in the system tray; then select **NetWare Login**.
 - h. In the Novell Login window, select **Advanced**.
 - i. Enter the following information; then select **OK**.
 - Username: **admin**
 - Password: **novell**
 - Tree: **TEMPTREE1**
 - Context: **TEMP**
 - Server: **TEMP1**
4. Verify that you are connected to both servers with IPX:
 - a. Right-click the red **N** icon in the system tray; then select **NetWare Connections**.
 - b. The Trans type for your connections to DA1 and TEMP1 should be IPX. If this is not the case, repeat Step 3.

5. Add long name space to DA1 at the server console prompt by entering the following:
LOAD LONG
ADD NAME SPACE LONG TO SYS
ADD NAME SPACE LONG TO DATA
6. Run the Migration Wizard by selecting **Start > Programs > Novell > NetWare Migration Wizard > NetWare Migration Wizard**.
7. Create a migration project:
 - a. In the About Novell NetWare Migration Wizard Startup window, select **OK**.
 - b. In the Novell NetWare Migration Wizard Startup window, make sure **Create a New Project** is selected; then select **OK**.
 - c. In the Migration Type window, make sure **NetWare 4, 5, or 6** is selected; then select **OK**.
 - d. In the Create Project: Setup Tasks window, select **Next**.
 - e. In the Project Filename field, enter **NetWare 4**; then select **Next**.
 - f. In the Select the Source NDS Tree window, make sure **DigitalAir-Tree** appears in the drop-down field; then select **Next**.
 - g. In the Select the Source Server window, select **DA1.IS.SLC.DIGITALAIR**; then select **Next**.
 - h. In the Select the Destination NDS Tree window, make sure **TEMPTREE1** appears in the drop-down field; then select **Next**.
 - i. In the Select the Destination Server window, select **TEMP1.TEMP**; then select **Next**.
 - j. Save your project and access the Project window by selecting **Create**.
The Project window now appears.

8. Copy volumes from DA1 to TEMP1:
 - a. In the Project window, select **Copy Volumes**.
 - b. In the Select Volumes to Copy window, select **SYS**.
 - c. With volume **SYS** selected, select **No** in the Copy This Volume field.
 - d. In the Copy column for volume **DATA**, make sure **Yes** appears; then select **Next**.
 - e. In the Duplicate Files window, make sure **Copy the Source File If It Is Newer** is selected; then select **Next**.
 - f. In the Disable Login window, select **Disable Login**; then select **Next**.
 - g. In the Source Tree Password field, enter **novell**.
 - h. In the Destination Tree Password field, enter **novell**; then select **Next**.

Allow the verification process to run.

- i. In the Error Resolution window, make sure you receive no critical errors; then do one of the following:
 - If there are no critical errors, select **Next**.
 - If there are critical errors, select **Cancel**, correct the errors, and start the process over by selecting **Copy Volumes**.



Most critical errors involve SMDR issues. Make sure Novell client connections to both DA1 and TEMP1 are IPX™.

- j. In the Ready to Copy Files window, copy the file system to the destination volume tree by selecting **Migrate**.

Allow the file trustees to back up and the volume files to migrate.
- k. On VMware server DA1, notice the message that login was disabled.

The migration wizard is in the process of closing down server DA1.

- l. In the Copy Volumes Status window, verify that the file copy was completed with no critical errors.
 - m. In the Copy Volumes Status window, review the error log.
 - n. Select **View Success Log**.
 - o. Scroll to the end of the success log and verify that volume DATA was migrated.
 - p. In the Copy Volumes Status window, select **Done**.
9. Edit configuration files:
- a. Select **Edit Configuration Files**.
 - b. In the list of configuration files, select **AUTOEXEC.NCF**; then select **Edit File**.
 - c. In the Compare Configurations window for the destination server (TEMP1), change the **FILE SERVER NAME** command to use **DA1** instead of TEMP1.
 - d. Copy the **IPX INTERNAL NET** command line from the source file (DA1) and insert it below the **FILE SERVER NAME** command in the destination file (TEMPx).
For example, if your destination server is TEMP4, your destination file should look similar to the following:

```
FILE SERVER NAME DA4
IPX INTERNAL NET 3E760BF1
```
 - e. Select **Save & Close**.
 - f. In the Compare Configurations windows, select **Close**.
10. Begin the eDirectory migration:
- a. At the workstation, in the Project window, select **Begin NDS Migration**.
 - b. In the Migrate NDS window, select **Next**.
 - c. In the Install License window, select **An MLA Is Already Installed**; then select **Next**.

Before you begin the eDirectory migration, VMware server DA3 must be running because it holds a replica of ROOT.

- d. In the Update Schema window, update the source server's schema by making sure **Yes** is selected; then select **Next**.
- e. In the Verify Novell Directory Services Tree window, verify that eDirectory is in good health by selecting **Yes**; then select **Next**.
- f. In the Delete Connections window, delete all user connections (except your own) to the source and destination servers by selecting **Next**.
- g. In the Password Verification window, enter *your password* in the Source Tree Password field.
- h. In the Destination Tree Password field, enter *your password*; then select **Next** and wait for the verification process to complete.
- i. Make sure no critical errors exist.
- j. In the Migrate NDS Verification Results window, select **Next**.
- k. In the Ready To Migrate NDS window, begin the eDirectory migration by selecting **Migrate**.
Notice that your source server (DA1) shuts down during migration, because it has been moved to your destination server.
- l. In the Migrate NDS Results window, view the Error and Success logs; then select **Done**.
- m. Close the current project by selecting **Close**.
- n. Check the former destination server and verify that it has restarted and has taken on the name of the source server.
- o. Close the migration wizard.
A message indicates that you need to reboot your workstation to clear the Novell Client cache.
You can continue without rebooting the workstation (and powering back on all the VMware servers).

11. Finish eDirectory Migration:

- a. Open the migration wizard by selecting **Start > Programs > Novell > NetWare Migration Wizard > NetWare Migration Wizard**.
- b. In the About Novell NetWare Migration Wizard window, select **OK**.
- c. Make sure **Open Last Project** is selected; then select **OK**.
- d. In the Getting Started Migrating window, select **Close**.
- e. In the Project Window, select **Finish NDS Migration**.
- f. In the Continue NDS Migration window, select **Yes**; then select **Next**.
- g. In the Password field, enter **novell**; then select **Next**.
- h. In the Ready To Continue Migrate NDS window, finish the eDirectory migration by selecting **Continue**.
- i. In the Continue Migrate NDS Results window, view the error log.
- j. Select **View Success Log**.
- k. Scroll to the bottom of your success log and verify that the migration completed; then close the **log**.
- l. In the Continue Migrate NDS Results window, select **Done**.
- m. Close the current project by selecting **Close**.
- n. Close the migration wizard.

12. Check the HOSTS and HOSTNAME files:

- a. From the workstation, use Windows Explorer to navigate to **DA1\SYS\ETC**.
- b. Double-click **HOSTS** and open with Notepad.
- c. Verify that 192.168.1.1 is mapped to **DA1.DIGITALAIRLINES.COM**.
- d. (Conditional) If you made changes, save the file and close Notepad.
- e. Double-click **HOSTNAME** and open with Notepad.

- f. Verify that 192.168.1.1 is mapped to DA1.DIGITALAIRLINES.COM.
- g. (Conditional) If you made changes, save the file and close Notepad.

13. Restart server DA1.

14. Install products from the NetWare 6 OS CD:

- a. Insert the **NetWare 6 OS CD**.
- b. Mount the CD at the server console by entering **CDROM**.
- c. At the server's graphical interface, select **Novell > Install**.
- d. Select **Add**.
- e. In the source path field, enter **NETWARE6:\;** then select **OK**.
- f. In the components screen, select **Clear All**.
- g. Select the following products to install:
 - Novell Certificate Server**
 - NDS iMonitor Services**
 - NetWare Remote Manager**
 - ConsoleOne 1.3.2**
 - NetWare Web Manager**
 - Novell iFolder Storage Services**
 - eDirectory iManage Service**
- h. Select **Next**.
- i. Authenticate as
 - User Name: **admin**
 - User Password: **novell**
 - User Context: **digitalair**
- j. In the Certificate Server screen, select **Next**.
- k. In the Organizational CA warning, select **OK**.
- l. In the LDAP Configuration Screen, select **Next**.

- m. In the iFolder Server screen, select **Next**.
 - n. In the iManager screen, select **Next**.
 - o. In the installation summary screen, select **Finish**.
 - p. After the file copy is complete, restart **DA1**.
15. Install NetWare 6 Support Pack 2:
- a. Mount the **NetWare 6 SP2** CD as a NetWare volume on DA1.



You can create a NetWare 6 SP2 CD by downloading the files from <http://support.novell.com/tools/csp/csplist.html>.

- b. At the server console, enter **NWCONFIG**.
- c. In Configuration Options, select **Product Options**.
- d. In Other Installation Actions, select **Install a Product Not Listed**.
- e. To specify the directory path, press **F3**.
- f. In Specify a Directory Path, change A:\ to **NW6SP2**: (include the colon).
- g. Press **Enter**.
- h. In the Novell Terms and Conditions screen, press **Esc** to continue.
- i. Accept the license agreement by selecting **Yes**.
- j. In the License Agreement for JReport Runtime JInfonet software, press **Esc** to continue.
- k. Accept the license agreement for JReport Runtime by selecting **Yes**.
- l. Install NetWare Support Pack version 6.0.2 by pressing **Enter**.
- m. On the Backup Files Replaced by NetWare Support Pack screen, select **No**.

- n. On the Do You Want to Update the Storage/LAN/PSM/WAN Drivers Currently in Use screen, do one of the following:
 - If you are using newer devices (such as CD drives), select **Yes**.
 - If you are using older equipment in the classroom, and want to avoid problems with using updated drivers, select **No**.
 - o. To reboot your server after the file copy, select **Yes**.
 - p. In the Warning screen, press **Enter** to continue.
 - q. When prompted, authenticate using your full context and password.
Allow files to copy and your server to reboot.
 - r. When prompted, do not press any key to exit.
16. Complete eDirectory clean-up tasks by using iMonitor to perform a health check and repair problems.
- These procedures are described in [“Perform Health Checks” on 4-28](#).

Upgrade eDirectory and Management Tools for DA1

Do the following:

1. Upgrade eDirectory:
 - a. Mount the **eDirectory 8.7** CD (November 2002 SEL CD2A) as a NetWare volume.
 - b. At the server console, load **NWCONFIG**.
 - c. In the Available Options menu, select **Product Options**.
 - d. Select **Install a Product Not Listed**.
 - e. To specify the path to the CD, press **F3**.
 - f. Specify the path to the directory where the installation program can find the NDS8.IPS file by entering *volume name:NW*.

For example: EDIR_8_7:NW.

- g. Allow the files to copy.
 - h. On the Software License Agreement screen, press **Esc** to continue.
 - i. Accept the license agreement.
 - j. On the License Agreement for JReport Runtime JInfonet Software screen, press **Esc**.
 - k. Accept the Reporting license agreement.
 - l. Continue by pressing **Esc**.
 - m. Read the warning; then press **Esc** and allow the files to copy.
 - n. In the Administrator Name field, enter *your full distinguished name*.
 - o. In the Password field, enter the *your password* and allow the files to copy and your server to reboot.
 - p. Authenticate to the Directory and allow the files to copy.
 - q. In the Are You Installing Remotely Through RConsole screen, select **No-Local**.
 - r. In the Novell Certificate Server 2.40 Objects screen, select **Next**.
 - s. In the LDAP Configuration screen, select **Next**.
 - t. In the Novell Modular Authentication Service screen, select **Next**.
 - u. In the Next screen, select **Next**.
 - v. In the Components screen, select **Next**.
 - w. In the Summary screen, select **Finish** and allow the files to copy.
 - x. When the Installation Complete screen appears, remove your CD; then select **Yes** and allow your server to reboot.
2. Install iManager 1.5 on DA1:
 - a. Mount the **Web Apps** CD (November 2002 SEL CD2B) as a NetWare volume.

Before you can start the iManager installation, you must change the volume name in WEBAPP.NCF from eDirWebapps to CD2B.

- b. Edit the **WEBAPP.NCF** file:
 - a. At the server graphical interface, select **Novell > Utilities > File Browser**.
 - b. Double-click **CD2B**.
 - c. Right-click **WEBAPP.NCF** and select **Edit File**.
 - d. Select **File > Save As**; then save the file in **SYS:SYSTEM**.
 - e. Right-click **SYS:SYSTEM\WEBAPP.NCF** and select **Edit File**.
 - f. Change `java -cp eDirWebapps:\` to `java -cp CD2B:\`.
 - g. Select **File > Save**.
- c. Start the iManager 1.5 installation by double-clicking **SYS:SYSTEM\WEBAPP.NCF**.
- d. Select *your language* from the drop-down menu; then select **OK**.
- e. In the Introduction screen, select **Next**.
- f. Accept the license agreement; then select **Next**.
- g. Deselect Novell eGuide; then select **Next**.
- h. In the Selected Applications screen, select **Install**.
- i. In the eDirectory iManager screen, select **OK**.
- j. In the Novell iManager Introduction screen, select **Next**.
- k. In the Detection Summary screen, select **Next**.
- l. In the Pre-Installation Summary screen, select **Install**.
- m. When the installation is complete, select **Done**.

Set Up Server DA2

Server DA2 must have 2 network boards installed and a SCSI adaptor installed before proceeding with the NetWare installation.

Do the following:

1. Prepare the network for NetWare 6:
 - a. At the workstation, insert the **NetWare 6 OS CD**.
 - b. (Conditional) If it does not autorun, browse to the CD and launch **NWDEPLOY.EXE**.
 - c. Double-click **Network Preparation**.
 - d. Double-click **Step 2: View and Update NDS Versions**.
 - e. Enter **netware://Tree/DIGITALAIR-TREE**.
 - f. Select **Include Subordinate Containers**; then select **Next**.
You should find that no servers need to be updated.
 - g. Select **Exit**.
 - h. Double-click **Step 3: Prepare for NDS eDirectory 8.6**.
 - i. Enter **netware://Tree/DIGITALAIR-TREE**.
 - j. Select **Next**.
 - k. Select **DA1**; then select **Next**.
 - l. When the update is complete, select **Exit**.
 - m. To close Deployment Manager select **Cancel**; then select **Yes**.
2. Install NetWare 6 on server DA2:
 - a. Start the computer and insert the **NetWare 6 OS CD**.
 - b. Select *your language*.
 - c. Accept the license agreement.
 - d. Create a new boot partition and modify the size to equal **200 MB** + the *amount of RAM* on the computer.
 - e. After the computer reboots and the boot partition is created, accept the license agreement.

- f. Select **Custom** and **New Server**.
- g. Continue through the installation using the following information:
 - NetWare partition size: **2100 MB**
 - Server name: **DA2**
 - DNS name: **DA2.DIGITALAIRLINES.COM**
 - Domain name server: **192.168.1.2**
 - License diskette: **DA2**
 - Volume: **DATA (traditional with 2000 MB)**
 - IP address: **192.168.1.2**
 - Subnet mask: **255.255.255.0**
 - Additional protocol: **IPX**
 - Tree: **DIGITALAIR-TREE** (existing tree)
 - Context: **CORPORATE.SLC.DIGITALAIR**
 - When prompted, authenticate as **admin.digitalair** with the password **novell**.
- h. After NDS® installs, make sure the tree name and context are correct; then select **Next**.
- i. Select the browse button to the right of the License field, browse to and select the server license file; then select **Next**.
- j. In the Components Summary screen, select **Clear All**.
- k. Select **Novell DNS/DHCP Services** and **Novell iFolder Services**; then select **Next**.
- l. In the Certificate Server screen, select **Next**.
- m. In the LDAP Configuration screen, select **Next**.
- n. In the iFolder Server Options screen, select **Next**.
- o. In the iManage screen, select **Next**.
- p. In the Installation Summary screen, select **Finish**.
- q. After the installation is complete, remove the CDs; then select **Reboot**.

3. Configure TIMESYNC on DA2:
 - a. At the server console, enter **MONITOR**.
 - b. Select **Server Parameters > Time**.
 - c. Scroll down to **TIMESYNC Configured Sources**; then press **Enter** to set this option to **On**.
 - d. Scroll down to **TIMESYNC Time Sources**; then press **Enter**.
 - e. In the TIMESYNC Time Sources window enter **DA1**;
Make sure you include the semicolon (;).
 - f. Scroll up to **TIMESYNC Restart Flag**; then press **Enter**.
 - g. Select **Yes**; then press **Enter**.
 - h. Press **Enter** again.
 - i. Press **Esc** 3 times; then select **Yes** to exit MONITOR.
4. Apply Support Pack 2 to DA2:
 - a. Mount the **NetWare 6 SP2 CD** on DA2.
 - b. At the server console, enter **NWCONFIG**.
 - c. Select **Product Options**; then select **Install a Product Not listed**.
 - d. For the path enter **NW6SP2:**.
Make sure you include the colon (:).
 - e. Press **Esc** to continue.
 - f. Select **Yes** to accept the license agreement.
 - g. Press **Esc** to continue.
 - h. Select **Yes** to accept the ConsoleOne Reporting Tool License agreement.
 - i. Press **Enter** to continue.
 - j. Select **No** to not replace backup files.
 - k. (Conditional) If you see an update drivers message, select **Yes**.

- l. Select **Yes** to reboot the server after the files are copied.
- m. Press **Enter** to continue.
- n. When prompted, authenticate as **admin.digitalair** with the password **novell**.
- o. When the file copy is complete, wait for the server to restart.

Upgrade eDirectory and Management Tools for DA2

Do the following

1. Upgrade eDirectory:
 - a. Mount the **eDirectory 8.7** CD as a NetWare volume.
 - b. At the server console, load **NWCONFIG**.
 - c. In the Available Options menu, select **Product Options**.
 - d. Select **Install a Product Not Listed**.
 - e. To specify the path to the CD, press **F3**.
 - f. Specify the path to the directory where the installation program can find the NDS8.IPS file by entering *volume name:NW*.
For example: **EDIR_8_7:NW**.
 - g. Allow the files to copy.
 - h. On the Software License Agreement screen, press **Esc** to continue.
 - i. Accept the license agreement.
 - j. On the License Agreement for JReport Runtime JInfonet Software screen, press **Esc**.
 - k. Accept the Reporting license agreement.
 - l. Continue by pressing **Esc**.
 - m. Read the warning; then press **Esc** and allow the files to copy.
 - n. In the Administrator Name field, enter *your full distinguished name*.

- o. In the Password field, enter *your password* and allow the files to copy and your server to reboot.
 - p. Authenticate to the Directory and allow the files to copy.
 - q. In the Are You Installing Remotely Through RConsole screen, select **No-Local**.
 - r. In the Novell Certificate Server 2.40 Objects screen, select **Next**.
 - s. In the LDAP Configuration screen, select **Next**.
 - t. In the Novell Modular Authentication Service screen, select **Next**.
 - u. In the Next screen, select **Next**.
 - v. In the Components screen, select **Next**.
 - w. In the Summary screen, select **Finish** and allow the files to copy.
 - x. When the Installation Complete screen appears, remove your CD, select **Yes**, and allow your server to reboot.
2. Install iManager 1.5 on DA2:
- a. Mount the **Web Apps** CD as a NetWare volume.
Before you can start the iManager installation, you must change the volume name in WEBAPP.NCF from eDirWebapps to CD2B.
 - b. Edit the **WEBAPP.NCF** file:
 - a. At the server graphical interface, select **Novell > Utilities > File Browser**.
 - b. Double-click **CD2B**.
 - c. Right-click **WEBAPP.NCF** and select **Edit File**.
 - d. **Select File > Save As**; then save the file in **SYS:SYSTEM**.
 - e. Right-click **SYS:SYSTEM\WEBAPP.NCF** and select **Edit File**.
 - f. Change `java -cp eDirWebapps:\` to `java -cp CD2B:\`.

- g. Select **File > Save**.
 - c. Start the iManager 1.5 installation by double-clicking **SYS:SYSTEM\WEBAPP.NCF**.
 - d. Select *your language* from the drop-down menu; then select **OK**.
 - e. In the Introduction screen, select **Next**.
 - f. Accept the license agreement; then select **Next**.
 - g. Deselect Novell eGuide; then select **Next**.
 - h. In the Selected Applications screen, select **Install**.
 - i. In the eDirectory iManager screen, select **OK**.
 - j. In the Novell iManager Introduction screen, select **Next**.
 - k. In the Detection Summary screen, select **Next**.
 - l. In the Pre-Installation Summary screen, select **Install**.
 - m. When the installation is complete, select **Done**.
 - n. When the installation is complete again, select **Done** again.
 - o. When the installation is complete again, select **Done** again.
3. Update ConsoleOne to version 1.3.4:
- a. From the instructor workstation, right-click **Start**.
 - b. Select **Explore**.
 - c. Browse to **DA1\SYS:\PUBLIC\MGMT\CONSOLEONE**.
 - d. Copy the **1.2 folder**.
 - e. Paste the 1.2 folder on your workstation's hard drive at **C:\NOVELL\CONSOLEONE**.

This should replace the older version of the 1.2 folder on your workstation and allow you to use ConsoleOne 1.3.4.
 - f. Test ConsoleOne by starting the program from your workstation desktop.

Install User Licenses for DA1 and DA2

Do the following:

1. At the workstation use Internet Explorer to access iManager at **HTTPS://192.168.1.2:2200**.
2. Under eDirectory iManager select **DA2**.
3. Use the following information to authenticate:
 - User name: **admin**
 - Password: **novell**
 - Context: **Digitalair**
 - Tree: **Digitalair-tree**
4. From the navigation frame on the left, expand **License Management**; then select **Install a License**.
5. At the right of the **Load license file** field, select **Browse**.
6. From folder **DA1** on your license diskette, select the *user license* file; then select **Open**.
7. Continue by selecting **Next**.
8. Select an available *user license*; then select **Next**.
9. In the **Location** field, enter **SLC.DIGITALAIR**; then select **Install**.
10. When the license is installed, select **Done**.

Configure DNS/DHCP on DA2

Do the following:

1. Configure DA2 as a DHCP server:
 - a. At the workstation use Internet Explorer to access iManager at **HTTPS://192.168.1.2:2200**.
 - b. Select **DA2** under eDirectory iManager.

- c. Use the following information to authenticate:
 - User name: **admin**
 - Password: **novell**
 - Context: **Digitalair**
 - Tree: **Digitalair-tree**
 - d. Expand **DHCP Management**.
 - e. Select **DNS/DHCP Scope Settings**.
 - f. For Context of DNS/DHCP Locator object, enter **CORPORATE.SLC.DIGITALAIR**.
 - g. For Administrator Scope, enter **DIGITALAIR**; then select **OK**.
 - h. When prompted that the operation is successful, select **OK**.
 - i. Select **DHCP Server Management**.
 - j. Select **Create Server**; then select **OK**.
 - k. Enter **DA2.CORPORATE.SLC.DIGITALAIR**; then select **Create**.
 - l. When prompted that the operation is successful, select **OK**.
2. Configure a subnet:
 - a. Select **Subnet Management**.
 - b. Select **Create Subnet**; then select **OK**.
 - c. Configure the subnet with the following information:
 - Subnet name: **3005Subnet**
 - eDirectory context:
CORPORATE.SLC.DIGITALAIR
 - Subnet IP address: **192.168.1.0**
 - Subnet mask: **255.255.255.0**
 - Default DHCP server:
DHCP_DA2.CORPORATE.SLC.DIGITALAIR
 - d. Select **Create**.

- e. When prompted that the operation is successful, select **OK**.
3. Configure an address range:
 - a. Select **Address Range Management**.
 - b. Select **Create Address Range**; then select **OK**.
 - c. Configure the address range using the following:
 - Subnet: **3005Subnet**
 - Address range name: **3005AddressRange**
 - Start address: **192.168.1.41**
 - End address: **192.168.1.100**
 - d. Select **Create**.
 - e. When prompted that the operation is successful, select **OK**.
4. Configure DHCP_DA2 to ping ahead before assigning IP addresses:
 - a. Select **DHCP Server Management**.
 - b. Select **View/Modify Server**; then select **OK**.
 - c. Select **DHCP_DA2.CORPORATE.SLC.DIGITALAIR**; then select **OK**.
 - d. Select **Next**.
 - e. Select **Ping Enable**; then select **Done**.
5. Configure DHCP_DA2 to dynamically provide DNS configuration information:
 - a. Select **Global DHCP Configuration**.
 - b. Select **View/Set Global Preferences**; then select **OK**.
 - c. Select **Modify**.
 - d. From the Available DHCP Options window select **00006 Domain Name Server**; then select **Add**.
 - e. When the Domain Name Server window appears, select **Add**.
 - f. In the IP address window enter **192.168.1.2**; then select **OK**.

- g. Select **Done**.
 - h. Select **Next** 3 times to move through the remaining configuration option windows.
 - i. From the Global DHCP Preferences screen scroll down to the bottom and select **Done**.
 - j. When prompted that the operation is successful, select **OK**.
6. Start the DHCP service on DA2:
- a. At the DA2 server console, enter **DHCPSRVR**.
 - b. Edit AUTOEXEC.NCF and place **DHCPSRVR** anywhere after the MOUNT ALL command so the DHCP service launches when the server is restarted.
7. Configure IP on the instructor workstation to obtain an IP address automatically.
8. Verify that DHCP is functioning:
- a. At the command prompt enter **IPCONFIG /RELEASE**; then enter **IPCONFIG /RENEW**.
 - b. Verify the workstation is receiving an IP address assignment from the DHCP server.
9. Configure DA2 as a DNS server:
- a. In the left frame, expand **DNS Management**.
 - b. Select **DNS Server Management**.
 - c. In the drop-down list, select **Create Server > OK**.
 - d. In the Enter NCP Server Name field, enter **DA2.CORPORATE.SLC.DigitalAir**.
 - e. In the Enter Host Name field, enter **DA2**.
 - f. In the Enter Domain Name field, enter **DigitalAirlines.com**; then select **Create**.
 - g. When prompted that the operation is successful, select **OK**.
10. Create a DNS zone for the classroom network:
- a. In the left frame, select **Zone Management**.

- b. In the drop-down list, select **Create Zone > OK**.
 - c. Select **Create New Zone**.
 - d. In the Specify eDirectory Context field, enter **CORPORATE.SLC.DigitalAir**.
 - e. In the Enter Zone Domain Name field, enter **DigitalAirlines.com**.
 - f. Select **Primary**.
 - g. In the Select Assigned Authoritative Zone Server field, select **DNS_DA2.CORPORATE.SLC.DigitalAir**.
 - h. Select **Create**.
- 11.** Using Step 10, do the following:
- a. Create an **IN-ADDR.ARPA (reverse) zone** for the DigitalAirlines.com zone you just created.
 - b. For the Zone Domain Name, enter **192.168.1.0**.
 - c. When prompted that the operation is successful, select **OK**.
- 12.** Create A resource records for each server on the classroom network and the classroom printer:
- a. In the left frame, select **Resource Record Management**.
 - b. In the drop-down list, select **Create Resource Record**; then select **OK**.
 - c. In the Select Domain Name drop-down list, select **DigitalAirlines.com**; then select **Create**.
 - d. In the Specified Host Name field, enter **DA1**.
 - e. In Select RR Type, make sure **A** is selected.
 - f. In the Enter IP Address field, enter **192.168.1.1**; then select **Create**.
 - g. When you are notified that the request was successful, select **OK**.

- h. Repeat steps **c - g** for each server in the classroom using the following server names and IP addresses:

Table Setup-3

IP Address	Server Host Name
192.168.1.1	DA1.DigitalAirlines.com
192.168.1.2	DA2.DigitalAirlines.com
192.168.1.4	DA4.DigitalAirlines.com
192.168.1.5	DA5.DigitalAirlines.com
192.168.1.6	DA6.DigitalAirlines.com
192.168.1.7	DA7.DigitalAirlines.com
192.168.1.8	DA8.DigitalAirlines.com
192.168.1.9	DA9.DigitalAirlines.com
192.168.1.31	iFolder1.DigitalAirlines.com
192.168.1.32	iFolder2.DigitalAirlines.com
192.168.1.34	iFolder4.DigitalAirlines.com
192.168.1.35	iFolder5.DigitalAirlines.com
192.168.1.36	iFolder6.DigitalAirlines.com
192.168.1.37	iFolder7.DigitalAirlines.com
192.168.1.38	iFolder8.DigitalAirlines.com
192.168.1.39	iFolder9.DigitalAirlines.com

13. Load the DNS service at the server console by loading **NAMED.NLM**.
14. Add **NAMED** to your server's **AUTOEXEC.NCF** file by placing it after the **STARTX** command.

Import 3005LDIF

You build the classroom tree using the 3005.LDIF file included on the Enhanced Learning CD.

1. If necessary, authenticate from your workstation as **Admin**.
2. Start **ConsoleOne**.
3. Import the 3005.LDIF file from the SETUP directory on the Enhanced Learning CD:
 - a. In ConsoleOne, browse to the **DigitalAir** container.
 - b. Select **Wizards > NDS Import/Export**.
 - c. Select **Import LDIF File > Next**.
 - d. From the Select Source LDIF File screen, browse to and select the **3005.LDIF** file in the SETUP directory on the Enhanced Learning CD.

(This file is also available from CNI Net or in DATA:\SETUP on the DA1 server.)
 - e. Select **Advanced**.
 - f. Deselect **Exit on Error**; then select **OK**.
 - g. Select **Next**.
 - h. From the Select Destination LDAP Server screen, select **New**.
 - i. In the Description field, enter **DA Import**.
 - j. In the Server DNS Name/IP Address field, enter **192.168.1.1**.
 - k. In the Port field, enter **636**.
 - l. In the Der File Containing Server Key Used for SSL Communications field, browse to and select the **RootCert.der** file from SYS:\PUBLIC on DA1.
 - m. In the User DN field, enter **cn=admin,o=DigitalAir**.
 - n. Select **OK**.
 - o. From the Select Destination LDAP Server screen, select **DA Import**.

- p. In the Password field, enter **novell**.
- q. Select **Advanced**.
- r. Select **Allow Forward References**.
- s. Deselect **Use LBURP**.
- t. Select **OK**.
- u. Select **Next**.
- v. In the Summary screen, select **Finish**.



If you receive an error stating that the client couldn't connect to the LDAP server, reboot the server and workstation; then run the import again.

Text similar to the following appears:

```
Source Handler: ICE LDIF handler for Novell
eDirectory
8.6.0 version: 10110.05
Destination Handler: ICE LDAP handler for
Novell
eDirectory 8.6.0 version: 10110.05
ICE log file: ice.log
Start time: Tuesday, Oct 9, 2001 9:53:54 am
Operation in progress ...

Total entries processed: 257
Total number of errors: 1
End time: Tuesday, October 9, 2001 9:53:55 am
Total Time: 0:00:01.107
Time per entry: 00:00.046
```



You will see a few errors. These are instances where a container or user object exists in the tree. You can disregard the errors. However, if more than 18 errors occur, there was a problem with the import.

- w. Select **Close**.

4. Refresh your tree view by pressing **F5** and then verify that the new containers and user objects were created.

Configure NAT on DA2

In this course, you use DA2 to route between the classroom network and the external network to provide Internet access. To configure NAT, do the following:

1. Transfer control of the server LAN driver configuration to INETCFG:
 - a. At the DA2 server console, enter **INETCFG**.
 - b. When prompted to transfer LAN configuration, select **Yes**.
 - c. When prompted to leave INETCFG and restart the server, select **Yes** and wait while the server restarts.
2. At the server console prompt, enter **CONFIG** and make sure of the following:
 - Drivers for both network boards in DA2 are loaded
 - A private IP address is assigned to the private board
 - A public IP address is assigned to the public board
3. Configure dynamic NAT to enable public network access from your private network:
 - a. At the server console, enter **INETCFG**.
 - b. Select **Yes, Use the Fast Setup Method**.
 - c. Press Esc; then select **Go to INETCFG Main Menu**.
 - d. Select **Bindings**.
 - e. Select the **TCP/IP binding** for the network board connected to your organizational network segment.
 - f. Select **Configure TCP/IP Bind Options**.
 - g. Select **Expert TCP/IP Bind Options**.
 - h. Select **Network Address Translation**.
 - i. Change the status to **Dynamic Only**.

- j. Update the configuration by pressing **Esc** 4 times.
 - k. When prompted to update the TCP/IP configuration, select **Yes**.
 - l. Return to the main menu by pressing **Esc**.
4. Using INETCFG, configure a static route to the next router on your organizational network:
- a. From the Internetworking Configuration menu, select **Protocols**.
 - b. Select **TCP/IP**.
 - c. Change the status of LAN Static Routing to **Enabled**.
 - d. Select **LAN Static Routing Table**.
 - e. Press **Insert**.
 - f. Change the Route Type to **Default Route**.
 - g. On the Next Hop Router on Route line, enter the *IP address* of the default gateway for your organizational network.
 - h. Press **Esc** twice.
 - i. When prompted to update the database, select **Yes**.
 - j. Press **Esc**.
 - k. When prompted to update the TCP/IP configuration, select **Yes**.
 - l. Press **Esc**.
5. Reinitialize the system to apply changes:
- a. From the INETCFG Main Menu, select **Reinitialize System**.
 - b. When prompted, select **Yes**.
 - c. When you are notified that the new configuration will take effect, press **Enter**.
 - d. Verify that the correct bindings are created; then switch to the INETCFG screen and exit INETCFG.



If reinitializing fails, restart the server. You might also have to reset the router after the server loads.

Prepare for Clustering

You must prepare the hardware for performing the clustering exercises. You might also want to perform some of the clustering exercises from Sections 8, 9, and 10 during the classroom setup so that you are prepared to demonstrate clustering.

Do the following:

1. Install SCSI adapters in DA1 and DA2 and make sure there is a SCSI cable for each.
2. Arrange placement of DA1 and DA2 so they are close enough for the SCSI cables to reach from each server to the SCSI hard disk.
3. (Optional) Perform clustering setup steps from Sections 8, 9, and 10.

Set Up Student Servers

Student servers must have 2 hard disk drives with nothing installed on them and must be prepared for the clustering exercises.

Do the following:

1. Run FDISK on each hard disk so there are no partitions configured on them.
2. Install SCSI adapters in each server and make sure there is a SCSI cable for each.
3. Make sure there is an external SCSI hard disk for every 2 servers.
4. Arrange the placement of servers so DA4 and DA5 are close enough for the SCSI cables to reach from each server to the SCSI hard disk; do the same for DA6 and DA7 and for DA8 and DA9.

Set Up Student Workstations

Do the following on all workstations:

1. Install Windows 2000 Professional using the following information:
 - Partition type: **NTFS**
 - Computer name: **WSx**
 - Administrator password: **novell**
 - OS update: **Service Pack 3 for Windows 2000** (or later)
 - Obtain an IP address automatically.
 - Obtain DNS server address automatically.
2. Install **Internet Explorer 6** or later on the workstation.
3. Turn off the Friendly HTTP Error Messages option in Internet Explorer:
 - a. From the Internet Explorer menu bar, select **Tools > Internet Options**.
 - b. Select the **Advanced** tab.
 - c. Scroll down and deselect **Show Friendly HTTP Error Messages**.
 - d. Select **Apply**.
4. Set up the home page in Internet Explorer to show Web Manager:
 - a. In the Internet Options dialog, select the **General** tab.
 - b. In the Address field under Home Page, enter **https://192.168.1.1:2200**.
 - c. Select **Apply**.
 - d. Select **OK**.
5. Install **Novell Client 4.81**:
 - a. Insert the **Novell Client** CD (January 2002 SEL CD7B).
 - b. Select *your language*.

- c. After the CD autoruns, select **Novell Client 4.81 for Windows NT/2000/XP**.
 - d. In the License Agreement window, select **Yes**.
 - e. Select **Custom Installation**; then select **Next**.
 - f. Do not select any components to install; then select **Next**.
 - g. Select **IP and IPX**; then select **Next**.
 - h. Select **NDS (NetWare 4.x or later)**; then select **Next**.
 - i. Select **Finish**.
 - j. When the installation is complete, select **Reboot**.
6. Install NICI 2.4:
- a. Insert the **NICI** CD (August 2002 SEL CD4A).
 - b. Browse to and double-click **NICI24\Win_all\WCNICIU0**.
 - c. In the Welcome window, select **Next**.
 - d. In the License Agreement window, select **Yes**.
 - e. In the Setup Complete window, select **Finish**.
7. Install WinZip according to the manufacturer's instructions.
8. Install ConsoleOne 1.3.3:
- a. Insert the **ConsoleOne** CD (April 2002 SEL CD3A).
 - b. Browse to and double-click **ConsoleOne133\C1_NW_Win**.
 - c. In the WinZip self-extractor window, select **Setup**.
 - d. In the Welcome window, select **Next**.
 - e. In the License Agreement window, select **Accept**.
 - f. In the Location window, accept the default location by selecting **Next**.
 - g. In the ConsoleOne Components window, deselect **Reporting Snapins**; then select **Next**.
 - h. In the Next window, select *your language* (English is always installed); then select **Next**.
 - i. In the Summary window, select **Finish**.

- j. When the installation is complete, select **Close**.
9. Install ConsoleOne snap-ins:
- a. Insert the **Enhanced Learning** CD in WS1.
 - b. Locate and run
D:\SETUP\CONSOLEONE 1.3.3 SNAPINS FOR 3005.EXE

(This file is also available from CNI Net or in DATA:\SETUP on your DA1 VMware server.)
 - c. Copy the snap-in files to the default location (C:\NOVELL\CONSOLEONE\1.2) on your hard drive by selecting **UNZIP**.
10. Copy the “Power to Change” Novell marketing video files to C:\MARKETING VIDEO on each student workstation.
- These files are in a Marketing Video folder in D:\EXERCISES\SECTION 10 on the Enhanced Learning CD or in a MarketingVideo.EXE file on CNI Net.

Test the Setup

1. Boot each workstation and make sure it loads properly.
2. Boot the instructor server and make sure NetWare 6 loads properly.
3. Run IPCONFIG.EXE on a workstation and verify that DHCP is working properly.

This completes the steps you must do before teaching this class. The following topic contains steps you perform after teaching **SECTION 11** and before teaching **SECTION 12**.

Setup for Section 12



Do not perform the following steps during initial classroom setup.

The exercise in Section 12 tests the skills of the students in troubleshooting the classroom network. Because there is only one network (or LAN), students should work as a group to plan and troubleshoot the problems in the exercise.

As the instructor for the course, you must do the following:

- [Customize the Exercise to Meet Student Needs](#)
- [Set Up the Network for the Exercise](#)

Customize the Exercise to Meet Student Needs

The following are suggestions for making the troubleshooting experience relevant and successful for your students:

- You are not limited to the setup steps listed under [Set Up the Network for the Exercise](#). Evaluate the technical expertise of your students and make adjustments by eliminating some setup tasks and adding others.
- As you make modifications to the standard setup for the troubleshooting exercise, make sure you note to students which network problems in Exercise 12-1 are no longer valid, which network problems have been modified, and any additional network problems they need to resolve.
- During the first 4 days of the course, students might encounter problems that they cannot resolve or do not have the time to resolve.

By including these as part of the troubleshooting exercise, you give students the chance to successfully resolve the problems and feel good about the classroom experience.

Set Up the Network for the Exercise

The following is a standard setup for Exercise 12-1.

As instructor for the class, you are welcome to modify this setup to include your own troubleshooting tasks or to customize the setup to meet the needs of the students.



If you do not migrate all NetWare 4.11 servers (DA4 - DA9), you must modify this setup and the student network problems in Exercise 12-1 to match the current state of your network.

For example, if you have not migrated DA8, you must implement step 1e on another server (such as DA4) or eliminate the step.

Do the following:

1. Introduce communication issues into the LAN:
 - a. On DA4, edit AUTOEXEC.NCF and change the subnet mask of the server to **255.255.255.252**.
 - b. On DA5, edit AUTOEXEC.NCF and change the IP address of the server to **10.0.0.4**.
 - c. On DA6, unplug the LAN cable and edit AUTOEXEC.NCF to change the subnet mask of the server to **255.255.255.252**.



You might want to wait until you have completed all setup tasks before unplugging any LAN cables.

- d. On DA7, edit AUTOEXEC.NCF to comment-out the LAN driver **LOAD** and protocol **BIND** statements for IP and IPX.
- e. On DA8, edit AUTOEXEC.NCF and change the subnet mask of the server to **255.255.255.252**.
- f. On 3 student workstations, reinstall the Novell Client using the **IP with IPX Compatibility** protocol option.

2. Introduce time synchronization issues into the LAN:
 - a. At the DA1 server console prompt, change server DA1 to secondary time type by entering the following 2 commands:
SET TIMESYNC TYPE = SECONDARY
SET DEFAULT TIME SERVER TYPE = SECONDARY
 - b. Set *time* on DA4, DA5, and DA6 to 1 week in the future.
3. Introduce eDirectory issues into the LAN:
 - a. On 3 servers holding master replicas in the tree, create a file in SYS:\SYSTEM called **APPSTART.NCF**.
 - b. Edit APPSTART.NCF and add the following command:
UNLOAD DS.NLM
 - c. Add **APPSTART.NCF** towards the end (but not at the end) of AUTOEXEC.NCF.
4. Introduce server issues into the LAN.
 - a. On DA 4 use NWCONFIG to deselect (remove) the **SCSIHD.CDM** driver from the list of drivers.
 - b. On DA5 use the SCSI controller card BIOS utility (for AHA2940 cards it is Ctrl+A on boot) to assign the same SCSI ID number to the hard drive and the controller card.
 - c. On DA7, edit HTTPD.CONF and change the Listen IP address parameter to an *incorrect IP address*.
 - d. On DA8 and DA9, edit HTTPD_ADDITIONS_NW.CONF and change the iFolder root directory to **SYS:\PUBLIC**.
 - e. On DA1, edit the DHCP clustering load script to use an *incorrect tree name*.
5. Reboot all servers and workstations.

Introduction

Duration: 30 minutes

In this course you learn advanced NetWare® network management and troubleshooting skills for NetWare operating system environments, directory services, and data storage and services.

You also learn how to set up an NCS SAN to test high availability of resources.

Audience

This course is for students who have entry-level experience in managing small LAN or WAN networks and meet the necessary prerequisite knowledge.

You should have Certified Novell AdministratorSM (CNASM) certification in NetWare 4, NetWare 5, or NetWare 6 (or equivalent experience).

Prerequisites

You must have an understanding of the following:

- Foundations of Novell Networking, Novell Course 3001
- Novell Network Management, Novell Course 3004
- Terminology, hardware, and practices commonly used in medium to large enterprise networks, such as routers, hubs, switches, backbones, and subnets

Certification

This course helps you prepare for the following tests:

Table Intro-1

Certification	CNE Test Number	CNI Test Number
Certified Novell Engineer SM (CNE [®])	050-682	050-882

Arrange to take a test within 6 weeks of completing or acquiring the course. Thereafter, the test might be replaced by a test based on an updated version of the course.

Tests apply toward Novell professional certifications, such as the CNE, CDESM, and Specialist certifications. For more about Novell certification programs, see www.novell.com/education/certinfo.

Relationship to Other Courses in the Curriculum

This course satisfies one of the course requirements for the Certified NetWare Engineer (CNE) certification requirement.

Agenda

This is a 5-day course.

Table Intro-2

	Module	Duration
Day 1	Introduction	00:30
	MODULE 1: Migrate NetWare 4 and NetWare 5 Servers to NetWare 6:	
	<ul style="list-style-type: none"> ■ Migrate NetWare 4 and NetWare 5 Servers to NetWare 6 	05:00

Table Intro-2 (continued)

	Module	Duration
	MODULE 2: Troubleshoot and Resolve Novell Network Problems:	
Days 1 and 2	<ul style="list-style-type: none"> ■ Identify Tools for Troubleshooting Novell Network Performance Issues 	02:30
Day 2	<ul style="list-style-type: none"> ■ Troubleshoot and Resolve NetWare Server Issues 	03:30
	<ul style="list-style-type: none"> ■ Monitor and Troubleshoot eDirectory 	03:00
Day 3	MODULE 3: Demonstrate Advanced Novell Network Storage Management Skills:	
	<ul style="list-style-type: none"> ■ Perform Advanced Novell Storage Services Tasks 	01:30
	<ul style="list-style-type: none"> ■ Configure and Troubleshoot a RAID Solution Using NSS 	01:00
	<ul style="list-style-type: none"> ■ Perform Advanced iFolder Tasks and Troubleshooting 	04:00
Day 4	MODULE 4: Deliver High Availability Services with Novell Cluster Services:	
	<ul style="list-style-type: none"> ■ Design and Set Up an NCS Cluster Configuration 	02:00
	<ul style="list-style-type: none"> ■ Install and Test NCS on A 2-Node Cluster 	02:30
	<ul style="list-style-type: none"> ■ Configure and Test High Availability File Access 	01:30
Day 5	<ul style="list-style-type: none"> ■ Configure And Test High Availability Services 	02:00
	MODULE 5: Troubleshoot a NetWare 6 Network	
	<ul style="list-style-type: none"> ■ Troubleshoot a NetWare 6 Network 	04:00

Course Feedback

Your feedback is valuable to Novell Education. To provide feedback on the course materials, use the web services tool at <http://www.novell.com/education/courses/feedback/index.html>.

After you submit your feedback, it will be entered into a database and assigned to a Novell Education course developer for resolution.

Scenario

This scenario is based on a WAN configuration. However, for classroom exercise purposes, all servers are connected over a LAN.

This results in students performing some tasks in class (such as creating a 2-node cluster) that would not normally be done over a WAN.

As you teach the course, make sure students understand (when appropriate) the limitations of performing these tasks over a WAN.

Digital Airlines, Inc. is a flight business that provides luxury charter and scheduled flight services for executives, government officials, athletic teams, and others needing private, flexible, secure, and catered air travel.

The following provides information about the Digital Airlines network, eDirectory tree structure, and tasks branch office network administrators must perform to begin upgrading servers to NetWare 6:

- [Digital Airlines Company Information](#)
- [Digital Airlines eDirectory Tree](#)
- [Digital Airlines NetWare 6 Implementation Plan](#)

Digital Airlines Company Information

Digital Airlines, Inc. has been in business for 5 years. They have

- 20 aircraft
- 435 employees:
 - 50 pilots
 - 300 flight attendants

- 10 ground crew per terminal
- 2 network administrators per terminal
- 73 other employees, including executive officers and administrative assistants

The following is additional information about the company:

- [Digital Airlines Offices](#)
- [Digital Airlines Office Departments](#)
- [Digital Airlines Executive Staff](#)

Digital Airlines Offices

Digital Airlines is headquartered in Salt Lake City, Utah and has terminals and offices in the following airports:

- Delhi Indira Gandhi International Airport (DEL)
- New York La Guardia International Airport (LGA)
- London Heathrow International Airport (LON)
- Salt Lake City International Airport (SLC)
- Sydney Australia International Airport (SYD)
- Berlin Tegel International Airport (TXL)
- Tokyo Narita International Airport (TYO)

Digital Airlines Office Departments

- Corporate (Salt Lake City only)
- Customer Service
- Flight Operations
- IS

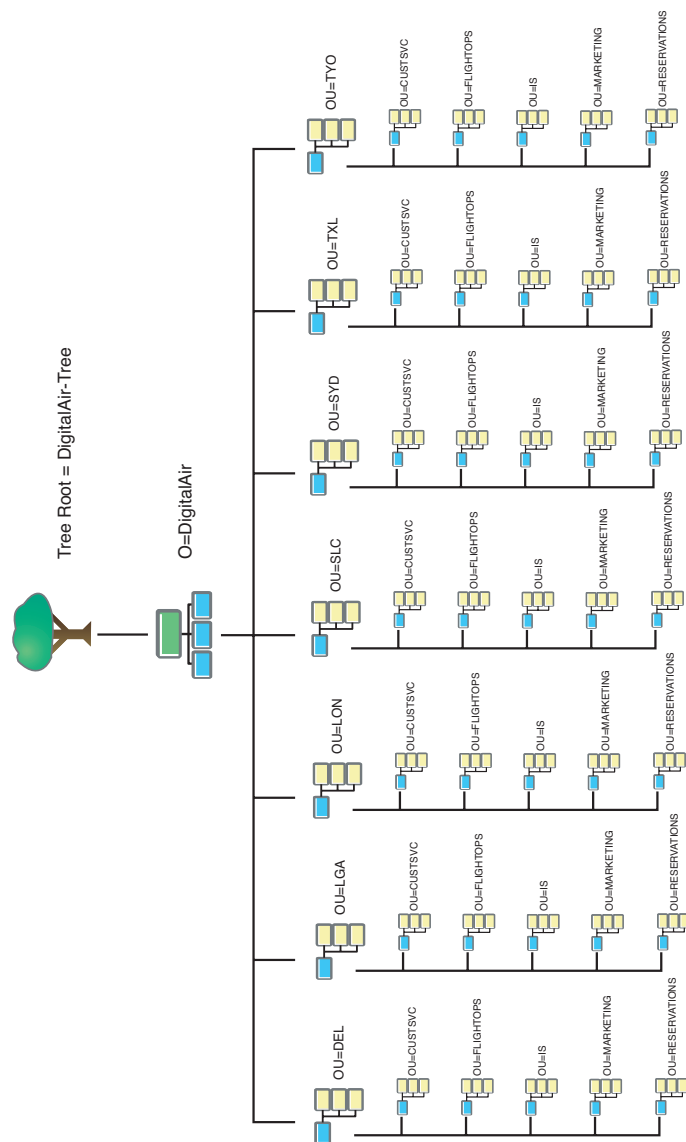
- Marketing
- Reservations

Digital Airlines Executive Staff

- Brian Edward Johnson, CEO
- George Sanders, CFO
- Christie Zervos, COO
- Craig Jenkins, CIO
- Catherine Burt, Director of Human Resources
- Cindy Valdez, Administrative Assistant
- LaVerl Tracy, Director of New York Terminal
- Leah Morgan, Director of London Terminal
- Wolfgang Mozart, Director of Berlin Terminal
- Pradeep Rathi, Director of Delhi Terminal
- Kaoru Tsunoda, Director of Tokyo Terminal
- Nathan Wadsworth, Director of Sydney Terminal

Digital Airlines eDirectory Tree

Figure Intro-1 (slide)



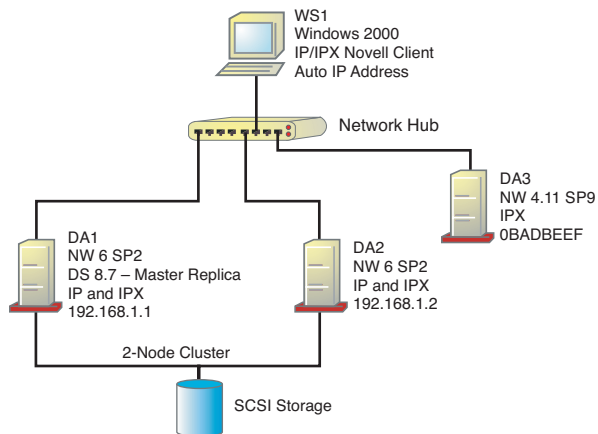
Digital Airlines NetWare 6 Implementation Plan

The executive board of Digital Airlines, under advisement from the Digital Airlines' CIO, has authorized a phased upgrade of corporate NetWare servers to NetWare 6.

This will let Digital Airlines take advantage of technologies and services such as Novell Storage Services™ (NSS), Novell iFolder™, and Novell Cluster Services™ (NCS).

The network administrator for the Salt Lake City office administers the following NetWare servers on a local LAN from a Windows 2000 workstation:

Figure Intro-2



Two of the servers (DA1 and DA2) are cluster-enabled to provide high availability to services such as DNS/DHCP and iFolder.

The servers provide various resources and are located in the following eDirectory containers:

Table Intro-3

DA1 and DA3 currently hold replicas for servers across the WAN.

If this design flaw is noted by the students, congratulate them and note that this is a purposeful design flaw that will be corrected by them in Section 4.

Server	Container	Resources
DA1	IS.SLC.DIGITALAIR	<ul style="list-style-type: none"> ■ Stores the ROOT master replica for DigitalAir-Tree and currently holds master replica for all servers in the tree ■ Provides network backup for the Salt Lake City LAN ■ Provides cluster-enabled iFolder services for the Salt Lake City office
DA2	CORPORATE.SLC .DIGITALAIR	<ul style="list-style-type: none"> ■ Stores a ROOT R/W replica for DigitalAir-Tree ■ Provides cluster-enabled DHCP/DNS services for all Digital Airlines offices ■ Provides iFolder services for Salt Lake City corporate employees
DA3	RESERVATIONS.SLC. DIGITALAIR	<ul style="list-style-type: none"> ■ Stores a ROOT R/W replica for DigitalAir-Tree and a R/W replica for all servers in the tree ■ Hosts the Salt Lake City office reservations system


Although the network administrator has recently migrated DA1 from a NetWare 4.11 server to a NetWare 6 server, the Reservations manager is apprehensive about upgrading DA3 to NetWare 6.

The reservations system on DA3 currently provide excellent performance and service (with very little downtime).

Because it will take some time to determine the impact of upgrading DA3 (and because there is no urgency to upgrade), the network administrator decides to leave DA3 as a NetWare 4.11 server and help the Digital Airlines branch offices upgrade their servers to NetWare 6.

Each branch office currently stores important files in a DATA volume on a NetWare 4.11 server. The server is administered from a Windows 2000 workstation:

Figure Intro-3



The network administrator for Digital Airlines decides to have all the branch offices begin the transition to NetWare 6 by migrating this server and the volume DATA files to NetWare 6 on a new server box.

The servers are located in the following eDirectory containers:

Table Intro-4

Server	eDirectory Container
DA4	IS.DEL.DIGITALAIR
DA5	IS.LGA.DIGITALAIR
DA6	IS.LON.DIGITALAIR
DA7	IS.SYD.DIGITALAIR

Table Intro-4 *(continued)*

Server	eDirectory Container
DA8	IS.TXL.DIGITALAIR
DA9	IS.TYO.DIGITALAIR

After the migration process is complete, all branch office network administrators must perform advanced networking tasks such as checking the health and performance of the LAN, the migrated NetWare 6 server, and eDirectory.

In addition, the network administrators must implement iFolder and familiarize themselves with cluster-enabling by testing a 2-node NCS cluster in a lab setting.

MODULE 1

Migrate NetWare 4 and NetWare 5 Servers to NetWare 6

Section 1 Migrate NetWare 4 and NetWare 5 Servers to NetWare 6

SECTION 1 Migrate NetWare 4 and NetWare 5 Servers to NetWare 6

Duration: 5 hours

In this section you learn how to migrate NetWare 4.x and NetWare 5.x servers to NetWare 6.

Objectives

1. Review How to Prepare for a Server Migration to NetWare 6
2. Review How to Implement Novell Licensing
3. Identify How to Perform a Migration
4. Perform Post-Migration Tasks
5. Troubleshoot Post-Installation Issues

Introduction

In this section you perform a NetWare 6 migration.

You perform a migration rather than an upgrade when you need to upgrade your hardware.

Before performing a NetWare 6 migration, it's helpful to review the following alternatives for upgrading to NetWare 6:

- In-Place Upgrade
- NetWare Accelerated Upgrade

(Although you learned about migrations in a prerequisite course, you did not perform a migration. You performed an in-place upgrade.)

In-Place Upgrade

Briefly review the in-place upgrade material. Students performed an in-place upgrade in Course 3004.

You can use the NetWare 6 installation program to perform an in-place upgrade of an existing NetWare 4 or NetWare 5 server to NetWare 6.

An in-place upgrade is like a typical server installation, except the installation uses existing tree and volume information to complete the upgrade to NetWare 6.

When you perform the upgrade, you complete the following:

- Meet system and software requirements
- Prepare the network and the computer
- Specify hardware and software settings
- Create additional disk volumes (if required)
- Select and install networking protocols
- Set up Novell eDirectory
- Install other networking products

The upgrade program automates the following tasks:

- Device drivers and LAN drivers for NetWare 6 are loaded. Outdated drivers are matched with and replaced by new drivers included with NetWare 6.
- eDirectory is upgraded.
- NetWare 6 information is added to the AUTOEXEC.NCF and STARTUP.NCF files.
- The NetWare 6 files are copied to the server.

After you access the NetWare 6 installation files, you follow the instructions for installing a server found at www.novell.com/documentation.

NetWare Accelerated Upgrade

The NetWare Accelerated Upgrade utility is an advanced utility used to upgrade a NetWare 4.11, NetWare 4.2, or NetWare 5 server to NetWare 6.

NetWare Accelerated Upgrade is intended for use by network administrators who are skilled at troubleshooting and installing NetWare networks.

You can run NetWare Accelerated Upgrade from a Windows workstation or at the server console.

After you upgrade or install the first NetWare 6 server using the NetWare 6 installation program, you can then use NetWare Accelerated Upgrade to upgrade other servers in the tree.

Although NetWare Accelerated Upgrade is quicker than the standard installation process, it does not install additional network products, licensing services, or license certificates.

For example, when using NetWare Accelerated Upgrade to upgrade from NetWare 5 to NetWare 6, the NetWare 5 version of ConsoleOne is not upgraded. To get the NetWare 6 version, you must install it after completing the upgrade.

The following summarizes the advantages and disadvantages of using NetWare Accelerated Upgrade:

Table 1-1

Advantages	Disadvantages
<ul style="list-style-type: none"> ■ Runs from a Windows workstation or from the server console ■ Is quicker than the standard NetWare 6 installation program 	<ul style="list-style-type: none"> ■ Cannot be used to upgrade the first NetWare 6 server on the network, so you must use the NetWare 6 installation program ■ Does not install or upgrade products or licensing certificates, so you must install them after completing the upgrade ■ Does not provide backout procedures that restore the server or trustee assignments if the upgrade fails



After the accelerated upgrade you must install Apache Web Server for the following to operate: iFolder, NetWare Web Search, iManager, NetWare Web Access, and NetStorage.

Objective 1 Review How to Prepare for a Server Migration to NetWare 6

This information was discussed in Course 3004. Consider briefly reviewing it.

Properly preparing your server is critical for a successful migration. Although this information was covered in Course 3004, it is important to review the process when preparing for a migration.

The process not only provides guidelines for how much hardware you'll need, but it also outlines the correct steps you must take to prepare your file system and Directory for migration.

By improperly preparing your server or network for migration, you risk losing important files in your file system and Directory. You also risk losing all data on your server.

To prepare for a migration, you must do the following:

- [Prepare the Workstation](#)
- [Prepare the Source \(Original\) Server](#)
- [Prepare the Destination \(New\) Server](#)
- [Prepare Server Application Files](#)

Prepare the Workstation

Your workstation must be prepared so the migration wizard can run.

Make sure your workstation meets the following requirements:

- The workstation must run Windows 98 or Windows NT 4/2000 with 50 MB of available disk space:
 - The Windows 98 workstation must be running Novell Client for Windows 98 version 3.3 or later.
NetWare Migration Wizard 6 does not run on Windows 95 workstations.
 - The Windows NT 4/2000/XP workstation must be running Novell Client for Windows NT version 4.8 or later.
- Install the latest service pack for Windows 98/NT/2000/XP on your workstation.

- If you are migrating from NetWare 4, IPX should be configured on your Novell Client workstation.
- For better performance, run the source (original) server, destination (new) server, and workstation on the same LAN segment.

Prepare the Source (Original) Server

Define *source server*. Help students understand that the source server is the original (old) server.

The *source server* is the original NetWare server that contains the files, volumes, and eDirectory objects to be copied to the destination server. Valid source servers can NetWare 4.11, 4.2, 5.0, 5.1, or 6 servers.

Do the following:

1. Verify that you have the Supervisor right to the source server's file system and Directory tree.
2. Make sure the destination (new) server can communicate with the source server:
 - a. On the destination server, if it is using IPX, enter **Display Servers** and make sure the source server is listed.
 - b. On the destination server, if it is using IP, enter **Display SLP Services** and make sure the source server is listed.

Point out the note. It is very important.



All servers and workstations involved in the migration must have common protocols bound. Also, if using IPX, make sure a common frame type is bound.

3. Apply the **NW6NSS1A** patch to update the version of NSS.



See TID 2961749 - Post SP1 NSS modules for NetWare 6. Despite what the abstract implies, this patch can be applied before applying Support Pack 1.

4. Update the source server with the latest NetWare support pack, available at support.novell.com.
5. (Conditional) If you are migrating data from NetWare 4, make sure the source server's volumes have long name space support added to all volumes to be copied.

To add long name space support to a NetWare 4.11 or NetWare 4.2 volume, enter the following at the server console:

LOAD LONG
ADD NAME SPACE LONG TO *volume name*

6. Load **DSREPAIR** and run the following options:

Unattended Full Repair
Time Synchronization
Report Synchronization Status

Make sure these finish with no errors. However, in a mixed NetWare 4.x and 5.x environment, it is possible to finish a full unattended repair with errors due to schema mismatches.

You perform these operations to ensure that the Directory is healthy and stable before you migrate the server.

7. Make 2 full, verified backups of the eDirectory tree and the file system.
8. Check reference material to verify that you are using the supported hardware.

Point out that it is possible to finish a full unattended repair with errors due to schema mismatches.

Prepare the Destination (New) Server

The *destination server* is the new computer that receives the data from the source server. This server must be installed into a temporary tree.

Define *destination server*. Help students understand that the destination server is the new server.

After data is migrated from source to destination server, the destination server reboots and the migration wizard modifies the destination server's AUTOEXEC.NCF file to include the source server's name and internal IPX number or server ID.



The migration wizard does not support migration to NetWare 4 destination servers.

Review the following before installing the NetWare OS on the destination server:

1. Verify that you have the Supervisor right to the destination server's file system and Directory tree.
2. Make sure you create a temporary eDirectory tree with a temporary eDirectory tree name.

Point out the note. This is very important.



Do not use the same name as the source server's name or eDirectory tree name; otherwise, the destination server cannot assume the identity of the source server after the migration.

Point out that if students don't create the same volumes on the destination server that exist on the source server, the volume's files won't be migrated to the new server.



3. During installation of the destination server, create volumes on the destination server that are the same size as, or larger than, volumes on the source server.

Volume names on the destination server must be the same as the volume names on the source server.

The migration wizard migrates compressed volumes. If you are migrating compressed volumes to uncompressed volumes, the migration wizard decompresses the volumes during migration.

The decompression process is CPU and time intensive, so allow enough time to complete the operation.

Make sure you have room on the uncompressed volume to accommodate the source volumes after they are decompressed.

Point out that the destination server must be installed using the Pre-Migration installation option.

This option is available only in NetWare 6, not NetWare 5.1. If you were installing NetWare 5.1 as the destination server, you would not have a pre-migration option.



4. If you are migrating from NetWare 4, install and configure IPX. IPX must be bound to the destination server for the migration to work. (You can remove IPX after the migration.)
IP addresses for the source server are not migrated. Instead, you manually change IP addresses during migration.
5. Install the NetWare OS on the destination server using the **Pre-Migration** installation option.
6. Prevent time synchronization issues:
 - a. Configure the destination server as a secondary time source by entering the following commands at the console prompt:
SET TIMESYNC TYPE = SECONDARY
SET CONFIGURED SOURCES = ON
SET TIMESYNC TIME SOURCES = SOURCE
SERVER NAME; or IP ADDRESS;

If your source server is a NetWare 4 server, use the server name.

- b. Turn on the Timesync Debugger screen by entering **SET TIMESYNC DEBUG = 7**.
- c. Set the Timesync Restart Flag to restart TIMESYNC by entering **SET TIMESYNC RESTART FLAG = ON**.
- d. Make sure the destination server is the same time or later than the source server by entering **TIME**.

This prevents critical time synchronization errors during migration.

Prepare Server Application Files

Make sure your server application files are ready to migrate by verifying that both your source server and destination servers are running the same applications.

For example, if you are running GroupWise®, Oracle®, or backup software databases and virus software, make sure both the destination and source servers are running these applications.

Preparing application files is the biggest issue you face when preparing for a migration because it is the most time-consuming.



Always test your migration in a lab environment before performing a migration in your production environment.

Objective 2 **Review How to Implement Novell Licensing**

Students learned about licensing in Course 3004, but because the licensing model changed in NetWare 6, it is important to review it again to ensure a successful migration.

The 2 key points for students to understand about licensing are (1) 1 user license will connect a user to more than 1 NetWare 6 server, and (2) place user licenses above the user objects in the tree.



Novell Licensing Service (NLS) lets you manage license units that comply with the licensing requirements of Novell.

In NetWare 6, you use iManager to install license certificates when you add NetWare servers and users to the eDirectory tree.

Other features of iManager let you delete and move license units for NetWare and other NLS-enabled products. (BorderManager is an example of an NLS-enabled product.)

You are bound by the licensing terms and conditions of your agreement with Novell to manually determine if usage is exceeding the licensing agreement.

To identify how server and user licensing works in NetWare 6, you must understand the following:

- [Server and User Licensing Models](#)
- [How UAL Coexists with SCL](#)
- [How the Licensing Models Differ](#)
- [License Types](#)

Server and User Licensing Models

In addition to a server license, NetWare 6 requires a user license for each user who accesses the network and uses services provided by NetWare 6 servers.

Server and user licenses are separate files. They can be installed anywhere in the eDirectory tree.

However, Novell recommends that licenses be installed in a container higher in the tree than the objects that use the license units.

(You can install licenses in the same container as the objects that use them if all your users exist in a single container.)

To manage the licenses required in a NetWare environment, you must understand the following:

- **The server connection license (SCL) model.** Prior to the release of NetWare 6, Novell used the SCL model to regulate licensed usage of NetWare and its services.

In the SCL model, users were granted access to network services on a server basis. This meant that a single user might use several connection licenses if the user concurrently connected to multiple servers.

With the SCL model, you had to estimate the number of connections a user might need to perform his or her job to determine the number of connection licenses needed.

The number of nonuser objects requiring connections also had to be accounted for in the license unit count.

Point out that the SCL model is not available in NetWare 6.

Novell required the purchase of SCLs to be installed on each server. These licenses could be purchased in bundles, such as a 100-license bundle that allowed 100 server connections on a single server.

Only 100 server connections at a time could be made for the organization to be in compliance with its licensing agreement.

Even printers and other nonuser requestors of services, such as an NDPS® printer or a ZENworks® workstation, used a connection license.

This model limited an organization's ability to provide services on a server to its users because a single user could monopolize several connection licenses through multiple drive mappings and NDPS objects.

- **The UAL model.** With the release of NetWare 6, Novell has implemented a user access licensing (UAL) model. In the UAL model, user objects are assigned a license unit that allows a user unlimited access to NetWare 6 servers and their services.

In the UAL model, each user is assigned a user license as they initially log in to a NetWare 6 server. The user can then connect to any other NetWare 6 server in a single eDirectory tree without requiring another license.

After the first assignment of a user license to a user object, that license is reserved for that user as long as the user continues to authenticate to the network.



If you install user licenses in the same container as your users, all your users must be in that container for the licenses to work.

If the user doesn't log in again for 90 days or more, the license is released and made available to the next user who needs a license assignment.

Point out that a user license allows a user unlimited access to NetWare 6 servers and their services.



User licenses can be released from the originally assigned user through iManager. This is helpful when license units have been assigned to users who log in infrequently and you want to release the license for use by others who require more regular network access.

In the UAL model, server licenses are still required for NetWare 6 servers.

During the installation of a NetWare 6 server, you install a server license. Each server in the tree must have a unique server license.

You can download additional server licenses from www1.novell.com/eld/LRequest.jsp?ENCRYPTION=NW6.

Licenses downloaded from this site are demo server licenses and are provided at no cost. This allows you to get a server installed and running for your organization.

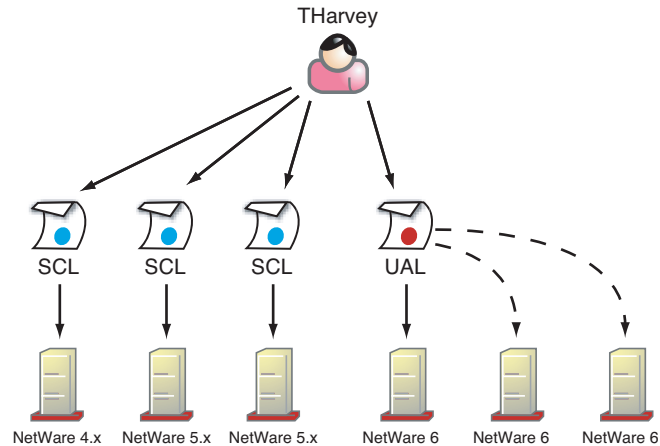
To purchase and download licenses that are not demo licenses, you must establish a license agreement with Novell or an authorized reseller.

How UAL Coexists with SCL

UAL can coexist on a network that is using SCL. Depending on the resources being accessed, a user might use a UAL unit and an SCL unit simultaneously.

For example, suppose a company has 6 NetWare servers, as shown in the following figure. Of the 6 servers, 1 is a NetWare 4.x server, 2 are NetWare 5.x servers, and 3 are NetWare 6 servers:

Figure 1-1 (slide)



When THarvey logs in to a NetWare 6 server, he uses one UAL unit, which allows connections to any number of NetWare 6 servers THarvey needs to access.

For each NetWare 4.x or 5.x server THarvey logs in to, a NetWare SCL unit is used.

To assess the license needs for a network consisting of NetWare 4.x, NetWare 5.x and NetWare 6 servers, you must understand each user's need for services provided by the network servers.

For example, if THarvey has a drive mapping to each of the non-NetWare 6 servers and also logs in to each of those servers every day, THarvey uses 2 SCL units every day on the 3 non-NetWare 6 servers, and he uses 1 UAL unit every day to access the 3 NetWare 6 servers.

As the network administrator, you benefit by upgrading NetWare 4.x and 5.x servers to NetWare 6 servers because you need the same number of license units as there are users in the eDirectory tree.

You don't have to take into account the multiple types of connections that might be using an SCL unit.

How the Licensing Models Differ

The following shows the differences between the SCL model and the UAL model:

Table 1-2

Feature	UAL Model	SCL Model
License packaging	Server and user licenses are available together or separately.	Server and user licenses are available in the same license envelope or separately.
Search for license	A search starts at the user's context and goes up the tree. If your users are installed in different containers, install the licenses above the users' containers.	A search starts at the server's context and goes up the tree.

Point out that a search for a license starts looking up the tree for the license.

Table 1-2 (continued)

Feature	UAL Model	SCL Model
Context of licenses	Install license certificates high in the tree to accommodate all users who need a NetWare 6 user license. By default, licenses are installed in the same context as the server.	Install license certificates high in the tree relative to servers' contexts.
License released when user logs out	No.	Yes.
Connection-oriented objects (like NDPS Printers and ZENworks objects) use a user or connection license	No.	Yes, until the current Support Pack is installed.

Point out that a UAL unit is not released when a user logs out.

License Types

The UAL model allows for the following 2 license types:

- **License agreement licenses.** Large companies that require many user licenses sign a license agreement with Novell. The agreement stipulates the number of license units that can be used before more licenses must be purchased.

Novell's licensing agreements provide pricing breaks according to the size of an organization.

Note: server licenses are free and user licenses are paid for.

The following are types of licensing agreements:

- **Master License Agreement (MLA).** Designed for large, worldwide organizations, the MLA offers a direct partnership with Novell that lets customers take advantage of Novell support services.
License purchases are set up between the organization and Novell to establish pricing, support services, and auditing responsibility terms.
- **Corporate License Agreement (CLA).** The CLA is designed for medium to large organizations and is available only through CLA resellers.
License purchases are set up between the organization and the CLA reseller to establish pricing, support services, and auditing responsibility terms.
- **Volume License Agreement (VLA).** The VLA lets small to medium organizations purchase licenses through any Novell reseller without a signed contract.
- **Clustering User License Agreement (CUAL).** The CUAL is installed when you install NetWare Cluster Services; by default it is placed in the same context as the cluster object.
For users to connect to servers in the cluster, CUALs must be accessible to user objects. This means individual CUALs must be placed at or above the user's context in the eDirectory tree so they are accessible.

- **Retail licenses.** Companies that purchase a copy of NetWare through the Novell distribution channel, ShopNovell, receive a licensing disk in the box with the product. NetWare purchased through this channel is called a Red Box product.

If you need more licenses for your Red Box NetWare product, you can purchase additional licenses from Novell.

Objective 3 Identify How to Perform a Migration

Because students were introduced to this information in Course 3004, briefly review this material.

Point out that a migration is performed when you are upgrading hardware.

Define *temporary tree*.

As an alternative to upgrading a server, you can also migrate data from an existing NetWare server to a new NetWare 6 server.

You perform a migration when you need to upgrade your hardware.

When you migrate data, the migration wizard copies the file system and eDirectory database from an existing NetWare 4, 5, or 6 server to a newly installed NetWare 6 server.

After the original server's file system and eDirectory database are migrated, the original server is brought down and the NetWare 6 server reboots and assumes the name and identity of the original server on your network.

Before you can migrate your data, you must first install a NetWare 6 server in a temporary eDirectory tree.

A *temporary tree* is a tree that contains one server with a basic installation of NetWare and no additional products (other than SMS or any other default products).

To migrate data from NetWare 4, NetWare 5, or NetWare 6 after you prepare the network and server, you do the following:

- [Run the Migration Wizard](#)
- [Copy Volumes](#)
- [Edit Configuration Files](#)
- [Begin the eDirectory Migration](#)
- [Finish eDirectory Migration](#)

Run the Migration Wizard

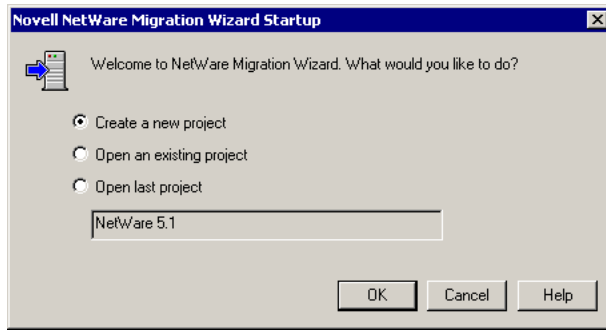
Do the following:

1. Install the migration wizard from the NetWare 6 OS CD.

2. Run the migration wizard on the workstation by selecting **Start > Programs > Novell > Netware Migration Wizard > NetWare Migration Wizard**.
3. Read the Welcome screen; then select **OK**.

The following appears:


Figure 1-2



4. In the Novell NetWare Migration Wizard Startup window, select **Create a New Project**.

The following appears:

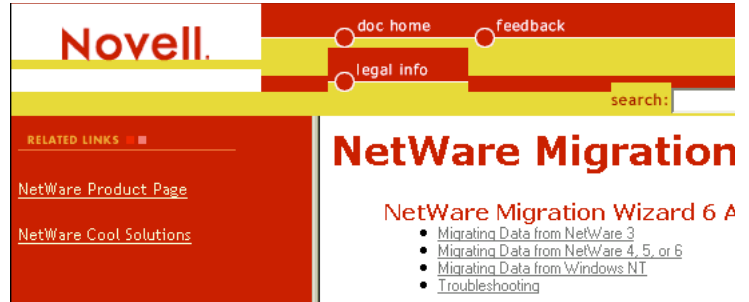
Figure 1-3



5. Select **NetWare 4, 5, or 6**; then select **OK**.
6. Select **View Setup Tasks**.

Selecting View Setup Tasks launches your default web browser and takes you to the Novell Migration Wizard 6 online documentation, as shown in the following:

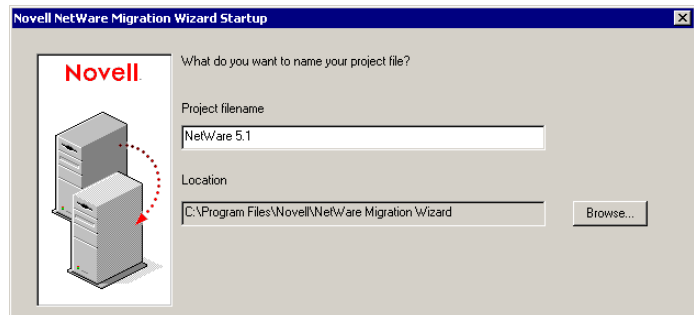
Figure 1-4



7. Select **Migrating Data from NetWare 4, 5, or 6** and make sure you have completed the system and software requirements.
8. Close your browser; then select **Next**.

The following appears:

Figure 1-5

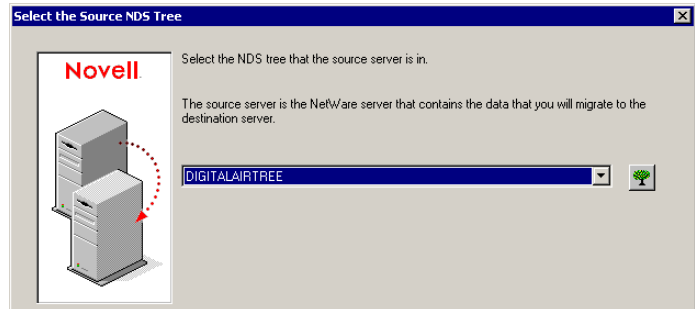


9. Name the project and choose a place to save it; then select **Next**.

By default, the migration wizard saves all projects to **C:\PROGRAM FILES\NOVELL\NETWARE MIGRATION WIZARD**.

The following appears:

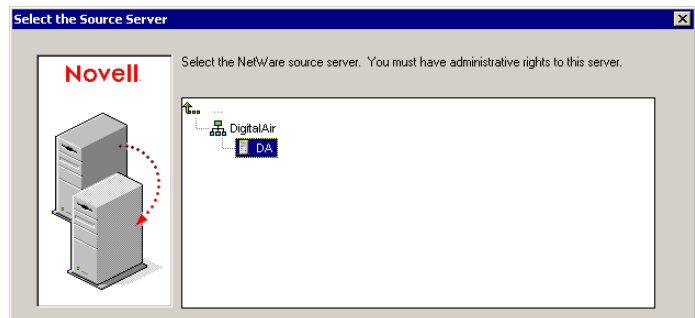
Figure 1-6



10. In the Select the Source NDS Tree window, select the *Directory tree* that contains your source server; then select **Next**.

The following appears:

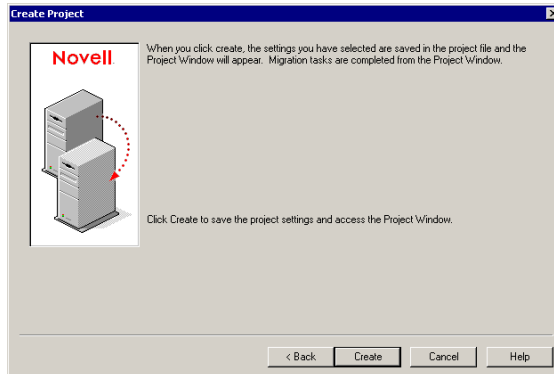
Figure 1-7



11. In the Select the Source Server window, select your *source server* from the Directory tree.
12. In the Select the Destination NDS Tree window, select the *eDirectory tree* that contains your destination server.
13. In the Select the Destination Server window, select your *destination server* from the destination eDirectory tree.

The following appears:

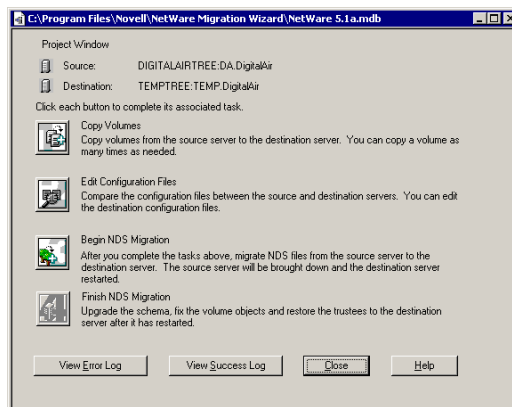
Figure 1-8



14. Save your project and access the Project Window by selecting **Create**.

The Project Window now appears:

Figure 1-9



Copy Volumes

Selecting Copy Volumes from the Project Window lets you copy volumes from the source server to the destination server. You can copy a volume as many times as you need to complete this step.

Before you can copy a volume from the source server to the destination server, you must have a volume with the same volume name created on the destination server.

For example, if you want to migrate volume DATA from your source server to your destination server, you must create a volume DATA on the destination server and give it enough space to hold all the files from the source server.

In the migration wizard, copy your volumes by doing the following:

1. In the Project Window, select **Copy Volumes**.

Before the migration wizard starts copying files, it backs up your directory and file trustees and saves them in files on the source and destination servers.

After the migration is complete, the migration wizard restores the trustees from the files it stored on the destination server.

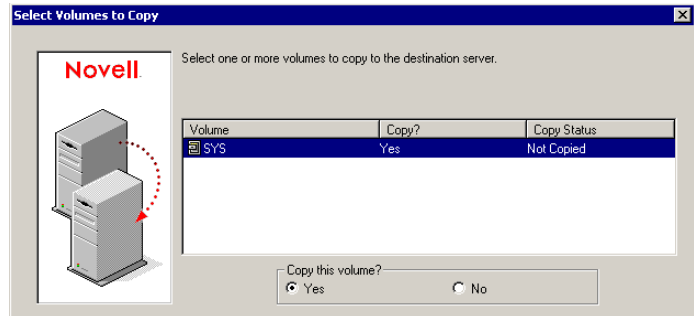
You do not need to copy all volumes at the same time. You can select volumes to copy now and then copy other volumes later by reopening the project file. Remember that open files are not migrated.



If you copy volumes in phases, at the final volume copy make sure you select all volumes that you previously copied; then select Cancel. Otherwise, the migration wizard restores trustee assignments only to the last volumes that were copied.

The following appears:

Figure 1-10



2. Select each *volume* you want to copy and select **Yes**; then select **Next**.



If you decide not to copy any volumes, select **No** for all volumes; then select **Next** and continue with [“Edit Configuration Files” on 1-29](#).

If you selected **Yes** for one or more volumes, continue with [Step 3](#).

Keep the following alternatives for copying volumes in mind:

- If you have big volumes or slow LAN connections, or if you want to reconfigure your data by putting existing directories into different folders on the destination server, consider using a backup tape to copy your volumes.
- If you use a tape backup, do not restore the source server’s standard SYS directories to the destination server if you are upgrading NetWare to a later version.

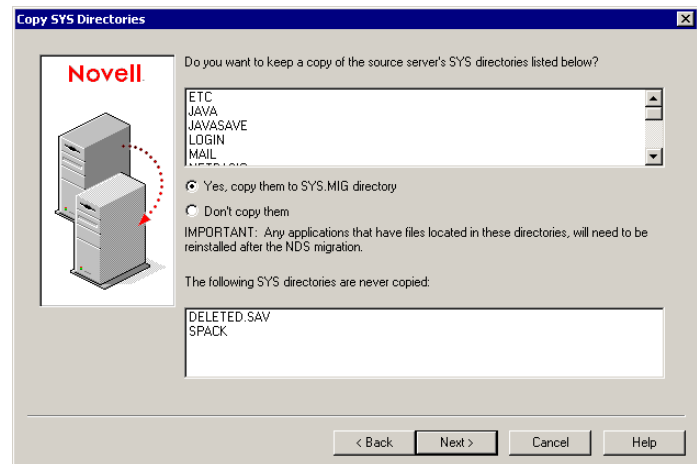
The standard SYS directories, like SYS:SYSTEM and SYS:PUBLIC, were created during NetWare installation.

If you decide to use tape backup, select **No** for all volumes in the Select Volumes to Copy window and select **Next**; then continue with “[Edit Configuration Files](#)” on 1-29.

- If you are migrating data to new hardware and keeping the same version of NetWare, restore the source server’s standard SYS directories to the destination server.

The following appears:

Figure 1-11



3. In the Copy SYS Directories window, decide if you want to copy the source server’s SYS directories to the destination server’s SYS:SYS.MIG directory; then select **Next**.

The migration wizard never overwrites the SYS directories on the destination server.

If you migrate the source server’s SYS directories, the migration wizard migrates them to SYS:SYS.MIG on the destination server.

If there are files in the source server's SYS directories that you want to use on the destination server, after the migration is completed, copy the files from SYS:SYS.MIG into the appropriate SYS directory on the destination server.



Remember, any applications that have NLM™ programs in this directory must be reinstalled after the migration.

4. In the Duplicate Files window, determine how you want to handle duplicate filenames between the source server and the destination tree by selecting one of the following; then select **Next**:
 - Don't copy over existing files
 - Copy the source file if it is newer
 - Always copy the source file
5. In the Disable Login window, determine how you want to copy your volumes; then select **Next**.

You have the following choices:

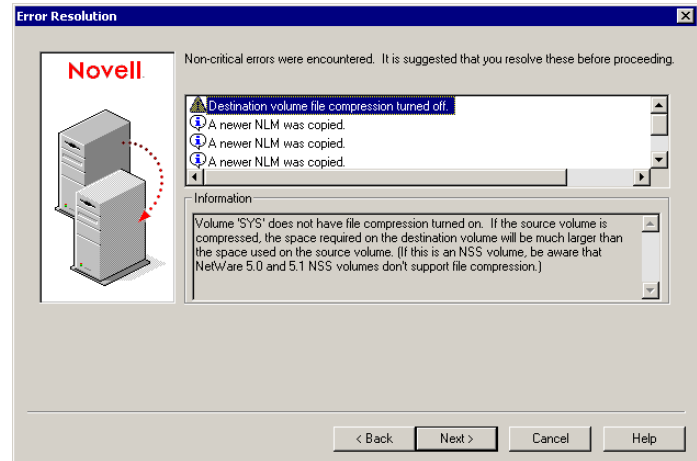
- Copy volumes with users logged in
- Disable login

The migration wizard does not copy open files. If you disable user login, no other users can log in and open files during the file copy.

6. In the Password Verification window, enter the *passwords* for the source and destination trees.

The following appears:

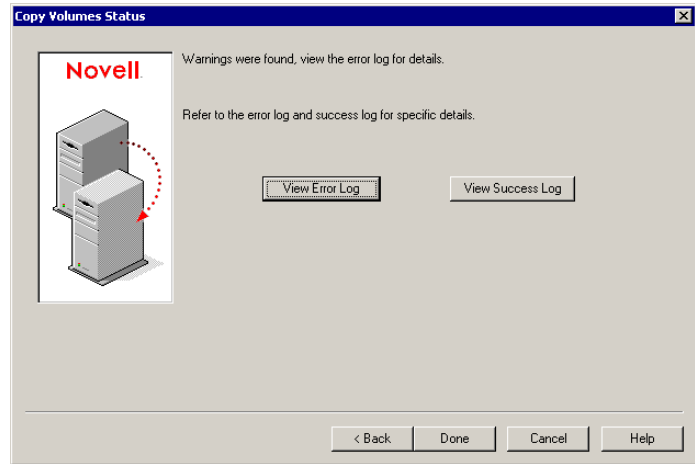
Figure 1-12



7. If prompted, resolve any errors or warnings; then select **Next**.
Noncritical errors are identified by yellow triangles; critical errors are identified by red circles with a white X in the center.
If you receive critical errors, you cannot proceed with the migration until you resolve those errors.
8. In the Ready to Copy Files window, select **Migrate** to copy the file system to the destination tree.
After you select Migrate, the following happens:
 - ❑ File trustees are backed up
 - ❑ Volume files are migrated to the destination server

When the volume migration completes, the following appears:

Figure 1-13

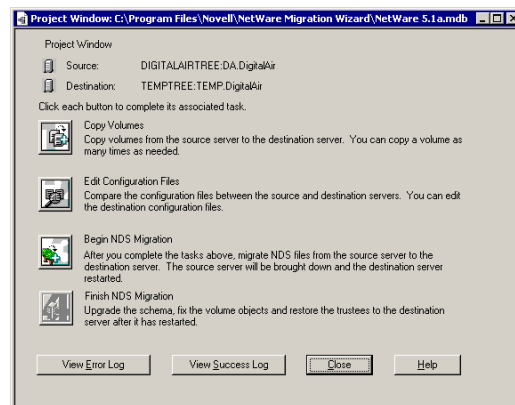


9. View the error and success logs; then select **Done**.

The error log is very helpful for identifying and fixing errors that occur during migration. The success log is useful for verifying how much of the migration completed.

The following appears:

Figure 1-14



Edit Configuration Files

Next, you compare the source and destination server configuration files. You also edit the configuration files on the destination server.

To edit the configuration files, do the following:

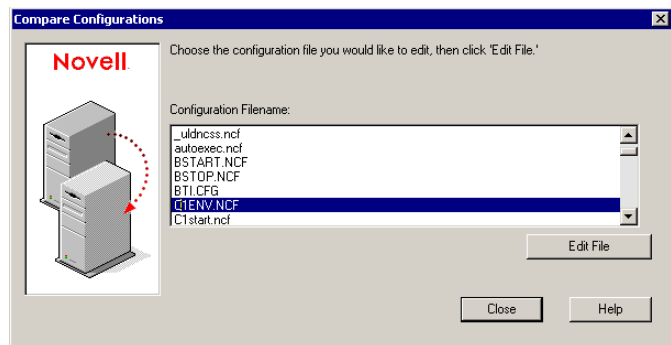
1. In the Project Window, select **Edit Configuration Files**.

The migration wizard lets you modify any NCF or CFG files on the destination server. These files contain default LOAD statements and parameters.

If you are editing the AUTOEXEC.NCF file, make sure the file is closed before you migrate your eDirectory database.

The following appears:

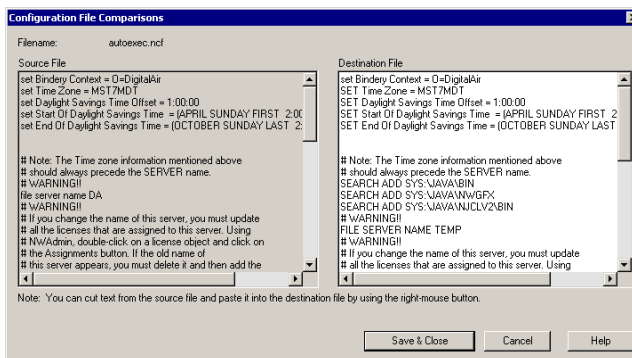
Figure 1-15



2. Select the configuration file you want to edit; then select **Edit File**.

The following appears:

Figure 1-16



3. Copy and paste the commands from the configuration file on the source server to the corresponding configuration file on the destination server.

To change the IP address on your destination server to be the same as the IP address of your source server, you must change the IP address in 3 places: **AUTOEXEC.NCF**, **SYS:\ETC\HOSTNAME**, and **SYS:\ETC\HOSTS**:

- a. Copy the *source server's IP address* and paste it into the destination server's AUTOEXEC.NCF file.



This works only if your IP LOAD and BIND statements are in AUTOEXEC.NCF. If you use INETCFG to assign IP addresses, change the IP addresses after the migration is complete.

- b. Using EDIT.NLM at the server console of the destination server, change the *IP address* and all instances of the *server name* in SYS:\ETC\HOSTNAME and SYS:\ETC\HOSTS.
4. When you finish modifying your configuration file, select **Save & Close**.
 5. Close the Compare Configurations dialog.

Begin the eDirectory Migration

After you complete the volume copy and edit configuration files, you can migrate eDirectory from the source to the destination server.

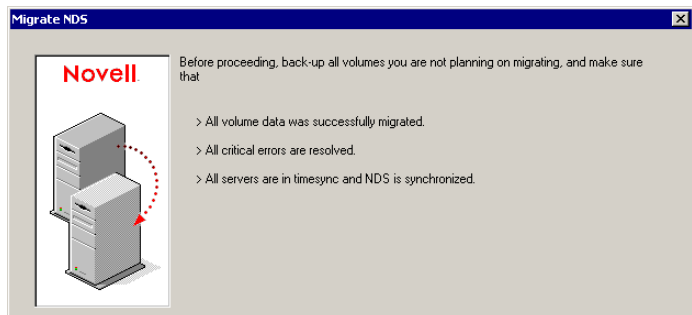
During this phase, the source server is brought down and the destination server is restarted. The destination server takes the source server's name and place in your network.

To begin the eDirectory migration, do the following:

1. In the Project Window, select **Begin NDS Migration**.

The following appears:

Figure 1-17



2. Back up all volumes that you are not planning to migrate and complete the following tasks before continuing; then select **Next**:
 - Make sure all volume data migrated successfully.
 - Make sure all critical errors from the file copy are resolved.
 - Make sure the time is synchronized on the servers in your source tree and that eDirectory is synchronized.

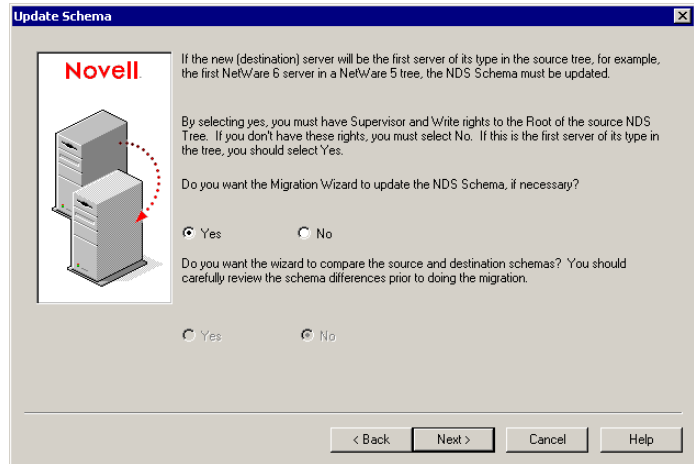
For more information, see Step 2 in [“Finish eDirectory Migration”](#) on 1-37.

3. When the Install License window appears, insert the license disk and browse to or enter the path to the license file; then select **Next**.

If you have an MLA, you can select the MLA instead of inserting the license disk.

The following appears:

Figure 1-18



4. Update the source server's schema by selecting **Yes**; then select **Next**.

The migration wizard updates the source server's schema to include the eDirectory classes of the default applications that are installed on the destination server.



If you select No to update the schema, by default the second Yes option is enabled, which causes the migration wizard to compare the source and destination schemas.

This comparison helps you determine how to extend your schema in preparation for the migration.

5. In the Verify Novell Directory Services Tree window, verify that you have run DSREPAIR to verify that the eDirectory tree containing the source server is functioning correctly; then select **Yes** or **No** to acknowledge that your tree is healthy; then select **Next**.

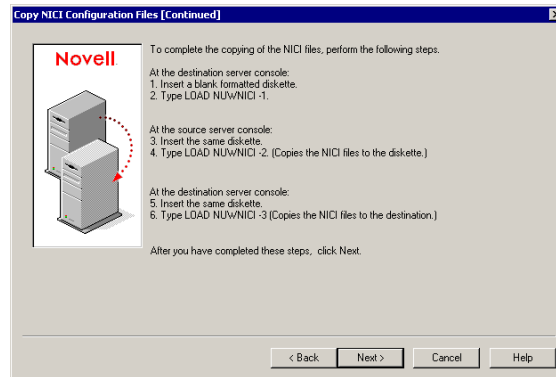


The migration wizard does not check the health of the tree and will not prevent you from continuing if your tree is unhealthy. If your tree is unhealthy the migration might not complete successfully.

6. In the Copy NCI Configuration Files window, select **Copy NLM** to begin the process of copying your NCI files.
7. When you see the message that NUWNCI.NLM was copied to the destination server from the source server, select **OK**.

The following appears:

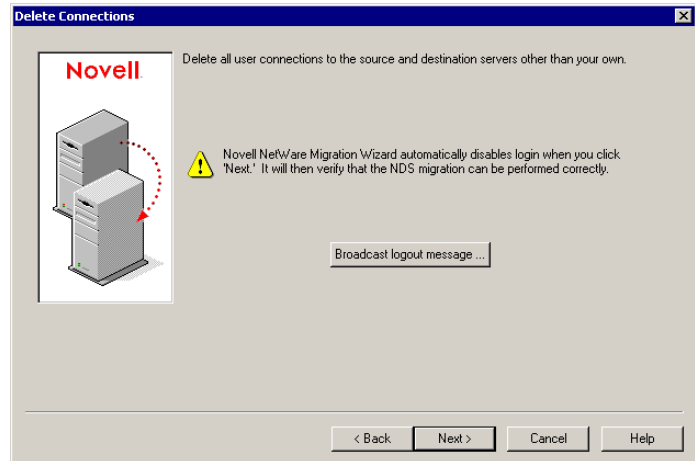
Figure 1-19



- Copy the NCI files from the source server to the destination server by following the on-screen instructions.

When the NCI configuration is complete, the following appears:

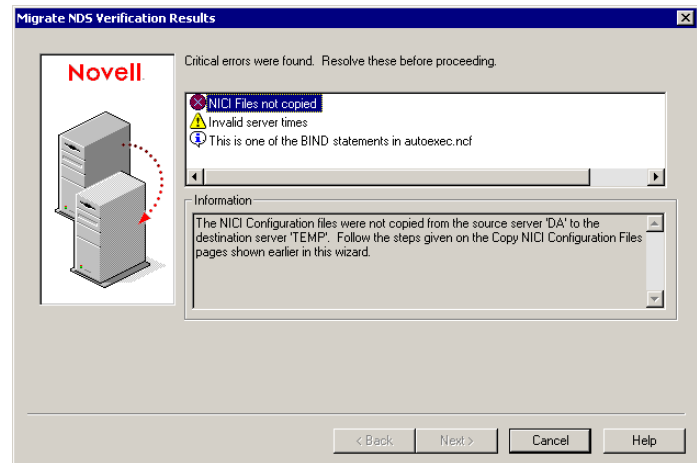
Figure 1-20



- Delete all user connections (except your own) to the source and destination servers; then select **Next**.
- In the Password Verification window, enter the *passwords* for the source and destination trees; then select **Next**.

The following appears:

Figure 1-21



11. Resolve any critical errors or warnings shown in the Migrate NDS Verification Results screen; then select **Next**.
12. In the Ready to Migrate NDS window, select **Migrate** to begin the migration.

At the end of the migration, the source server is brought down and the destination server reboots and takes over the name and identity of the source server.

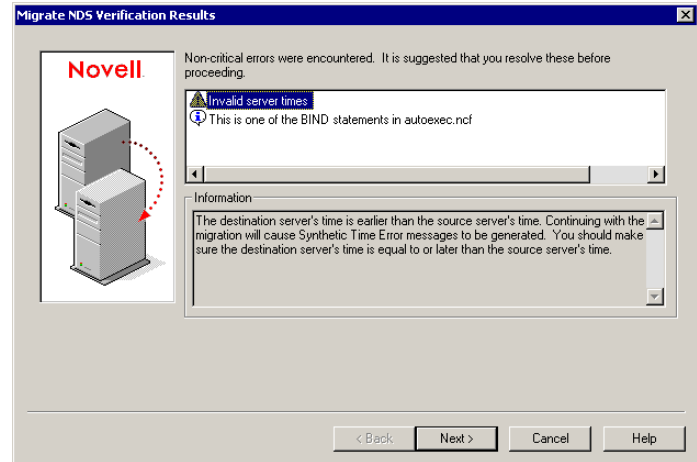
The migration wizard modifies the following items in the destination server's AUTOEXEC.NCF file:

- ❑ The server name changes to the name of the source server.
- ❑ The time zone is changed to the time zone in the source server's AUTOEXEC.NCF file.
- ❑ The server ID changes to the server ID in the source server's AUTOEXEC.NCF file.
- ❑ The default time server type is changed to the value stored in the source server's AUTOEXEC.NCF file.

- The bindery context is changed to the bindery context stored in source server's AUTOEXEC.NCF file.

When the migration is complete, the following appears:

Figure 1-22



13. View the Error and Success logs; then select **Done**.

Use the Error log to see errors that occurred during migration. If there were errors, use the Success log to determine how far the migration progressed.

If migration failed, restore your servers to their original configuration.

14. Check the destination server and verify that it has restarted and taken on the name of the source server.
15. Reboot your workstation and log in to the former destination server.

Finish eDirectory Migration

In this phase you upgrade the schema, fix volume objects, and restore trustees to the destination server after the server is restarted.

To finish the migration, do the following:

1. In the Project Window, select **Finish NDS Migration**.

In the Project Window at the Finish NDS Migration phase, notice that all buttons but the Finish NDS Migration button are gray.

These options are gray because the source server has been migrated to the destination server. You cannot go back and redo these options.

2. Read the Continue NDS Migration window and make sure the following has happened; then select **Next**:

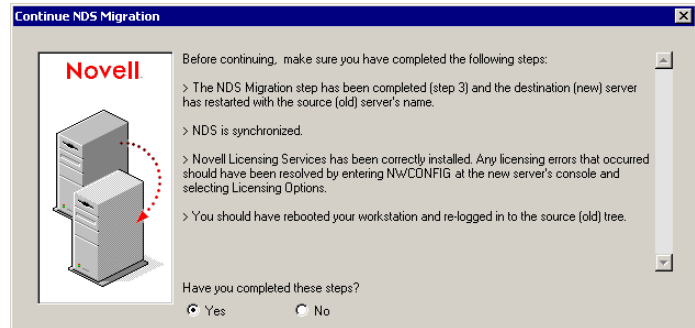
- The former destination server has restarted and has the name and identity of the source server.
- Licensing is installed.
- eDirectory is synchronized on all servers in the tree.

To check eDirectory synchronization status, enter **DSREPAIR** at the server console and run the Report Synchronization Status and Time Synchronization options.

If the destination server does not contain a Read/Write or Master replica, check eDirectory synchronization by running DSREPAIR on another server in the eDirectory tree that has one of these replicas.

The following appears:

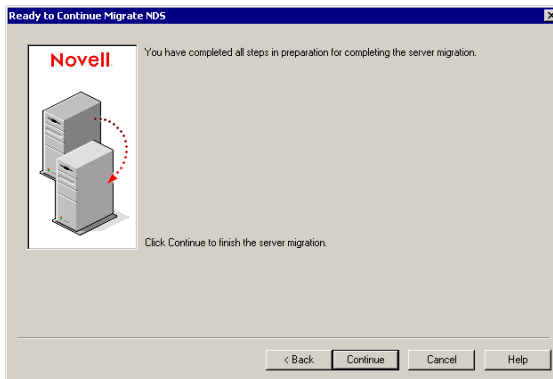
Figure 1-23



3. Verify that you have completed the tasks by selecting **Yes** or **No**; then select **Next**.
4. In the Password Verification window, enter the *password* for the source tree; then select **Next**.
5. Resolve warnings or errors; then select **Next**.

The following appears.

Figure 1-24



6. Finish the eDirectory migration by selecting **Continue**.
During this stage, the migration wizard restores the file trustees and updates the Volume objects in the tree.
7. View the error and success logs; then select **Done** to close the current project.

Objective 4 Perform Post-Migration Tasks

After the migration is complete, you perform the following to ensure that your new NetWare 6 network is running properly:

- [Run the External Reference Check Process](#)
- [Upgrade Existing NSS Volumes](#)
- [Perform Other Post-Installation Tasks](#)

Run the External Reference Check Process

Define *external reference check process* and *external reference*.

The *external reference check process* is an eDirectory process that verifies external references.

An *external reference* is a pointer to an eDirectory object on another server. An external reference indicates that an object in a replica has an ID on a server where the replica doesn't exist.

When migrating, trustee assignments might not be fully restored to user and group objects. NSS volumes depend on the eDirectory external reference check process to create ID information on objects before trustees can be associated with those objects.

To manually run the external reference check process, enter the following at the server console:

```
SET DSTRACE=ON
SET DSTRACE=+BLINK
SET DSTRACE=*B
```

If you switch to the Directory Services screen on the server, you can watch the objects as they are processed.



You can also run the external reference check process using iMonitor.

These commands must be repeated until no more user or group objects appear in the external reference list on the Directory Services screen. When complete, reboot the server.

Upgrade Existing NSS Volumes

If the original server had NSS volumes, you need to upgrade them to be compatible with the version of NSS shipped with NetWare 6.



If you migrate volume SYS, SYS remains a traditional volume until you upgrade it to NSS.

To upgrade existing NSS volumes, do the following for each NSS volume:

1. (For an in-place upgrade only) At the command prompt enter **NSS /ZLSSVOLUMEUPGRADE=All**
2. At the server console, enter **SET NLS SEARCH TYPE**
3. Make sure the value for NLS SEARCH TYPE is set to **0**; if not, at the console prompt enter **SET NLS SEARCH TYPE = 0**

4. At the server console, enter
**SET STORE NETWARE 5 CONN SCL MLA USAGE IN
NDS**
5. Make sure the value is set to **OFF**; if not, enter
**SET STORE NETWARE 5 CONN SCL MLA USAGE IN
NDS = OFF**
6. At the server console prompt, enter
FLUSH CDBE

Perform Other Post-Installation Tasks

After a NetWare 6 migration or upgrade, you must also perform the following:

1. Run DSREPAIR on the destination server and select **Unattended Full Repair**.
2. Make sure user information was migrated or upgraded successfully.
3. Reinstall applications that have files or NLMs associated with the following SYS directories: SYSTEM, PUBLIC, MAIL, ETC, and NETBASIC.

Because the source server's identity replaces the destination server during a migration, eDirectory objects on the destination server, including those representing applications, are removed and replaced by objects that exist on the source server.

However, the NLMs associated with the destination server's applications still exist.

Because you can't manage applications without their eDirectory objects, you must reinstall the applications to restore their objects to the tree.

**3 hours 30 minutes**

Assign a VMware server to each pair of students. For example, DA4 should be assigned the students with WS4, and so on.

Because DIGITALAIR-TREE recognizes DA4 - DA9, any VMware servers not migrated need to be running to avoid “can’t communicate with DAX” error messages and beeps.

Provide students with their Admin context found in the Introduction.

Alternative Setup: Consider loading Remote Console on the 4.11 servers; then have the students use RCONSOLE to access their server from their workstation to complete Part II of the exercise.

At the end of this exercise, perform a full unattended repair to resolve any network problems caused by the migrations.

Exercise 1-1 Upgrade a NetWare 4.11 Server to NetWare 6

In this exercise you migrate a NetWare 4.11 server into NetWare 6.

Recall that the executive board of Digital Airlines has authorized an upgrade of selected corporate NetWare servers to NetWare 6 to take advantage of technologies and services such as Novell Storage Services (NSS), iFolder, and Novell Cluster Services (NCS).

The network administrator for Digital Airlines has decided to have all branch offices begin the transition to NetWare 6 by migrating your server to NetWare 6 on a new machine.

You must migrate your server to a new machine because the hardware on your existing server will not support NetWare 6. Your hardware has already been upgraded, so your job is to migrate your server and Directory data to the new box.

Your DA_x (NetWare 4.11) server is your source (original) server. DA_x contains the files, volumes, and eDirectory objects to be copied to the NetWare destination (new) server, TEMP_x.

To perform the migration and update your network services, do the following:

- [Part I: Install and Configure Your Destination Server](#)
- [Part II: Prepare Your Source Server](#)
- [Part III: Prepare Your Server for eDirectory 8.6 Using Deployment Manager](#)
- [Part IV: Install the Migration Wizard](#)
- [Part V: Run the Migration Wizard](#)
- [Part VI: Copy Volumes](#)
- [Part VII: Edit Configuration Files](#)

- [Part VIII: Begin the eDirectory Migration](#)
- [Part IX: Finish eDirectory Migration](#)
- [Part X: Complete eDirectory Cleanup Tasks](#)
- [Part XI: Install NetWare Products and Services](#)
- [Part XII: Install the Latest Support Pack](#)
- [Part XIII: Upgrade Directory Services](#)
- [Part XIV: Update the Novell Client](#)
- [Part XV: Install User Licenses](#)

Part I: Install and Configure Your Destination Server

To migrate your NetWare 4.11 server to NetWare 6 across the wire, you must install a pre-migration NetWare 6 server.

Make sure your server is on; then do the following:

1. On your TEMP_x server, insert your **NetWare 6 OS** CD into the CD drive and allow your machine to boot from the CD.
2. When prompted, select one of the following:
 - To install from your IDE CD, press **I**.
 - To install from your SCSI CD, press **S**.
 - To install both, press **B**.
3. Install your language by selecting your *language*.
4. In the Welcome to NetWare Server Installation window, use your arrow keys to select **Accept License Agreement**.
5. Select **Create a New Boot Partition**.
6. In the First Hard Disk window, select **Modify**.

7. Enter enough space for your DOS partition using the following formula:
 $200 + \textit{amount of RAM in machine}$
Ask your instructor for the amount of RAM in your machine.
8. Select **Continue**.
9. Verify that you want to create a boot partition by using the arrow keys to select **Continue**.
10. Reboot your computer by pressing any key on the keyboard; then allow your computer to reboot.
11. When prompted, reboot by selecting one of the following:
 - To install from your IDE CD, press **I**.
 - To install from your SCSI CD, press **S**.
 - To install both, press **B**.
12. Configure basic installation parameters:
 - a. In the License Agreement for Jreport Runtime screen, press **F10** to accept the license.
 - b. In the Welcome screen, select **Express**; then press **Enter** to switch to **CUSTOM**.
 - c. In the Welcome screen, select **New Server**; then press **Enter** to switch to **PRE-MIGRATION**.
 - d. Press **Tab**.
 - e. Select **Continue**.
 - f. In the Pre-Migration Installation screen, select **Continue**.
 - g. In the Server Settings screen, select **Continue**.
 - h. In the Regional Settings screen, select **Continue**.
 - i. In the mouse and video selection screen, select **Continue** and allow the files to copy.
13. Configure NetWare device drivers:
 - a. In the disk driver screen, select **Continue**.

- b. In the device driver screen, select **Continue** and allow the driver to copy and load.



If your drivers are not auto-detected, manually install them.

14. Modify partitions:

- a. In the Volume SYS and Partition Properties screen, select **Modify**.
- b. On the NetWare Partition Size line, press **Enter**; then enter **2009**.
- c. Press **Enter** again.

Make sure file compression remains off.

- d. Leave the remaining fields as their default value.
- e. Save your settings by pressing **F10**.
- f. Select **Continue**.

Allow the files to copy and the graphical installation utility to load.

15. Configure advanced server installation parameters:

- a. In the Server Properties screen, enter **TEMP x** (where x = your server number) in the Server Name field; then select **Next**.

For example, if your server is DA4, enter TEMP4.

- b. Insert your license disk.
- c. In the Encryption screen, browse to and select the *server license* file in the **SERVER LICENSE** directory.
- d. Select **Open**; then select **Next**.
- e. Select **Free Space**.
- f. Select **Create**.
- g. In the New Volume screen, enter **DATA** in the Volume Name field.

- h. In Volume Type, mark **Traditional**.
 - i. In the Space to Use field, enter **2000**.
 - j. Select **Apply to Volume**; then select **OK**.
 - k. In the Configure File System screen, select **Next**.
 - l. In the Mount Volumes screen, make sure **Yes** is marked; then select **Next**.
 - m. In the Protocols screen, make sure your *network board* is selected.
 - n. Mark **IP**.
 - o. In the IP Address field, enter **192.168.1.x** (where *x* = your server number).
For example, if your server is DA4, your IP address is 192.168.1.4.
 - p. In the Subnet Mask field, make sure **255.255.255.0** appears.
 - q. Make sure the Router field empty.
 - r. Mark **IPX**; then select **Next**.
 - s. In the Domain Name Service screen, enter **DAx** in the Hostname field (where *x* = your server number).
 - t. In the Domain field, enter **DigitalAirlines.com**.
 - u. In the Name Server 1 field, make sure **192.168.1.2** appears in the field.
 - v. Select **Next**.
- 16.** Configure time synchronization:
- a. In the Time Zone screen, select your *time zone*.
 - b. Select **Advanced**.
 - c. In the Time Server Type field, make sure **Secondary** is selected.
 - d. Mark **Use TIMESYNC Configured Sources**.
 - e. In the Time Source 1 line, enter **DA1**; (include the semicolon).

- f. Select **OK**.
- g. Select **Next**.

17. Configure eDirectory using the following information:

Table 1-3

Server	eDirectory Context
DA4	IS.DEL.DIGITALAIR
DA5	IS.LGA.DIGITALAIR
DA6	IS.LON.DIGITALAIR
DA7	IS.SYD.DIGITALAIR
DA8	IS.TXL.DIGITALAIR
DA9	IS.TYO.DIGITALAIR

- a. In the NDS Install screen, mark **New NDS Tree**; then select **Next**.
- b. In the Tree Name field, enter **TEMPTREE x** (where x = your server number).
- c. In the Context field, enter your *context* found in Table 1-3; then press **Tab**.
- d. Make sure your *context* appears in the Admin Context field.
- e. In the Password field, enter **novell**.
- f. In the Retype Password field, enter **novell**; then select **Next** and allow eDirectory to install.
- g. In the NDS Summary screen, select **Next**.



If you receive an invalid license file dialog, select **OK**; then remove the user license and continue with the exercise.

18. Configure licensing:

- a. In the Licenses screen, select **Next**.

- b. In the LDAP Configuration screen, select **Next**.
19. In the Summary screen, select **Finish** and allow the files to copy.
20. When prompted that the installation is complete, remove the installation CD and license disk from the server; then select **Yes** and allow the server to reboot.



Ignore time sync messages saying that single servers are incompatible with Reference and Primary servers.

You can have students perform these steps using RCONSOLE by loading RSPX and REMOTE on each VMware 4.11 server.

Part II: Prepare Your Source Server

From the instructor's workstation WS1, prepare your NetWare 4.11 server by doing the following:

1. Using DSREPAIR, do the following:
 - a. At the server prompt, enter **LOAD DSREPAIR**.
 - b. In Available Options, select **Advanced options menu**.
 - c. Select **Repair local DS database**.
 - d. Begin the repair by selecting **F10**; then select **Yes**.
 - e. When the repair is finished, view the current log file by pressing **Enter**.
If you receive errors, have your instructor assist you.
 - f. Continue by pressing **Esc** twice; then press **Enter**.
 - g. Unlock the database file by pressing **Esc**; then return to the Available options menu by pressing **Esc**.
 - h. In Available Options, select **Report synchronization status**.
If you receive errors, have your instructor assist you.
 - i. Exit DSREPAIR.

2. Verify that the time is synchronized on your NetWare 4.11 server:
 - a. At your NetWare 4.11 server console prompt, enter **Time**.
 - b. Verify that time synchronization is active and time is synchronized to the network.
3. At the server console prompt, enter the following commands:

LOAD LONG
ADD NAME SPACE LONG TO DATA
4. At the server console, verify that long name space has been added to volume DATA by entering **Volume**.

Part III: Prepare Your Server for eDirectory 8.6 Using Deployment Manager

Use Deployment Manager to prepare your NDS 6.33 files for eDirectory 8.6 by doing the following:

1. At the TEMP_x server console, enter **CONFIG**.
2. Verify that IPX is bound.
3. In the following space, specify the network board name bound to IP:
 4. At the server console, enter **UNBIND IP *network_board_name*** using the network board name from Step 3.
 5. On WS_x, log in to DIGITALAIR-TREE using the following parameters:

Table 1-4

Field	Value
Username	Admin

Table 1-4 (continued)

Field	Value
Password	novell
Tree	DigitalAir-Tree
Context	<i>Your server context</i>
Server	DAx

6. On WS_x, log in to TEMPTREE_x using the following parameters:

Table 1-5

Field	Value
Username	Admin
Password	novell
Tree	TempTree _x
Context	<i>Your server context</i>
Server	TEMP _x

7. On WS_x, insert the **NetWare 6 OS** CD into the CD drive.
8. (Conditional) If the CD does not auto-start, do the following:
- Select **Start > Run**.
 - Select **Browse**.
 - Browse to the **NetWare 6 OS** CD.
 - At the root of the CD, select **NWDEPLOY**.
 - Select **Open**.
 - Select **OK**.
9. In Deployment Manager, double-click the **Network Preparation** folder.
10. Under Network Preparation, select **Step 2: View and Update NDS Versions**.

Remind students that no updates are necessary because NWDEPLOY was run when DA2 was installed.

11. In the Update NDS window, select the **browse** button at the right of the text field.
12. In the NDS Tree Browser window, expand **Novell Network**.
13. Expand **Novell Directory Services**.
14. Expand **DigitalAir-Tree**.
15. Expand **DigitalAir**.
16. Expand the **containers** until you get to your server's container.
17. Select **OK**.
18. In the Update NDS window, select **Next**.
19. When you receive a message saying that no servers were found that require an update, select **OK**.
Normally you must perform the NDS preparation steps when preparing NDS. If you are updating NDS in your work environment, do the following; then continue with Step 20:
 - a. In the Select Servers to Update field, make sure **Update NDS** is marked on your *server's* row.
 - b. Select **Next** and allow files to copy.
 - c. Make sure **Restart NDS** is marked; then select **Next**.
 - d. After server DA x is updated, select **Exit**.
20. Select **Exit**.
21. Under Network Preparation, select **Step 3: Prepare for NDS eDirectory 8.6**.
22. In the NDS Tree window, select the **browse** button at the right of the text field.
23. In the NDS Tree Browser window, expand **Novell Network**.
24. Expand **Novell Directory Services**.
25. Select **DigitalAir-Tree**.
26. Select **OK**.

If students are using Remote Console and receive errors, cancel the operation and abort the installation.

27. In the Update NDS window, select **Next**.
28. Make sure **DA1** appears in both the Available Servers and Selected Server fields.

Normally you should see your server name appear in these fields.
29. Select **Next**.
30. When you receive the message that the NDS tree is prepared for eDirectory 8.6, select **Exit**.

Because your server does not contain a replica, you cannot confirm this process.

However, if you were performing this in your work environment, you would make sure you received the message **NDSEM Process: Complete** on your server.

This message verifies that eDirectory on your server is prepared for eDirectory 8.6.
31. Close Deployment Manager.

Part IV: Install the Migration Wizard

The migration wizard must be installed on your workstation before you can run it.

Do the following:

1. On **WS_x**, select **Start > Run**.
2. Select **Browse**.
3. On the **NetWare 6 OS CD**, browse to **PRODUCTS\MIGRTWZD\MIGRTWZD.EXE**.
4. Select **Open**.
5. Select **OK** and allow the files to extract to your workstation.

6. In the Choose Setup Language dialog, select your *language*; then select **OK**.
7. In the Welcome window, select **Next**.
8. In the Software License Agreement window, select **Yes**.
9. In the Choose Destination Location window, accept the default location by selecting **Next**; then allow the migration wizard to install.
10. In the Setup Complete window, select **Finish**.

Part V: Run the Migration Wizard

Run the migration wizard by doing the following:

1. Run the migration wizard by selecting **Start > Programs > Novell > NetWare Migration Wizard > NetWare Migration Wizard**.
2. In the About Novell NetWare Migration Wizard Startup window, select **OK**.
3. In the Novell NetWare Migration Wizard Startup window, make sure **Create a new project** is selected; then select **OK**.
4. In the Migration Type window, make sure **NetWare 4, 5, or 6** is selected; then select **OK**.
5. In the Create Project: Setup Tasks window, select **Next**.
6. In the Project Filename field, enter **NetWare 4.11**; then select **Next**.
7. In the Select the Source NDS Tree window, make sure **DigitalAir-Tree** appears in the drop-down field; then select **Next**.
8. In the Select the Source Server window, select **DAx**; then select **Next**.

9. In the Select the Destination NDS Tree window, make sure **TempTree** appears in the drop-down field; then select **Next**.
10. In the Select the Destination Server window, select **TEMP**; then select **Next**.
11. Save your project and access the Project Window by selecting **Create**.

The Project Window now appears.

Part VI: Copy Volumes

Copy volume SYS by doing the following:

1. In the Project Window, select **Copy Volumes**.
2. In the Select Volumes to Copy window, select volume **SYS**.
3. Under Copy this Volume, select **No**.
4. In the Select Volumes to Copy window, select volume **DATA**.
5. Under Copy this Volume, make sure **Yes** is marked; then select **Next**.
6. In the Duplicate Files window, make sure **Copy the source file if it is newer** is selected; then select **Next**.
7. In the Disable Login window, select **Disable login**; then select **Next**.
8. In the Source Tree Password field, enter **novell**.
9. In the Destination Tree Password field, enter **novell**; then select **Next** and allow the verification process to run.
10. In the Error Resolution window, make sure you receive no critical errors; then select **Next**.

Remember, if you receive critical errors, you cannot proceed with the migration until you resolve those errors.

If you receive a critical error, resolve the error; then begin with Step 1 of this part of the exercise. You can only perform Step 11 when you receive no critical errors.

If you receive an SMDR critical error, make sure you are logged in with IPX on both of your connections; then begin with Step 1 of this part of the exercise. If you still receive critical errors, ask your instructor for assistance.

11. In the Ready to Copy Files window, copy the file system to the destination volume tree by selecting **Migrate**; then allow the file trustees to back up and the volume files to migrate.
12. Notice on your server the message that login was disabled.
13. In the Copy Volumes Status window, verify that the file copy was completed with no critical errors.
14. In the Copy Volumes Status window, review the Error log by selecting **View Error Log**.
15. In the Copy Volumes Status window, review the Success log by selecting **View Success Log**.
16. Scroll to the end of the Success log and verify that volume DATA was migrated.

Remember that open files cannot be migrated. This is normal.

17. In the Copy Volumes Status window, select **Done**.

Part VII: Edit Configuration Files

Edit AUTOEXEC.NCF by doing the following:

1. In the Project Window, select **Edit Configuration Files**.
2. In the Configuration Filename field, select **AUTOEXEC.NCF**.
3. Select **Edit File**.

4. In the Configuration File Comparisons window, replace the file server name line in your destination AUTOEXEC.NCF file with the following 2 lines from your source AUTOEXEC.NCF file by copying and pasting:

File Server Name line
IPX internal net number line



Do not replace the Server ID number with the IPX internal net number.

5. On your workstation in the migration wizard, select **Save & Close**.
6. In the Compare Configurations window, select **Close**.
7. Using EDIT.NLM at the server console of the destination server, verify your *server name* was changed from TEMP x to **DA x** in the SYS:\ETC\HOSTS file.
8. Using EDIT.NLM at the server console of the destination server, verify your *server name* was changed from TEMP x to **DA x** in the SYS:\ETC\HOSTNAME file.

Make sure both instances of TEMP x were changed in the HOSTNAME file.

9. Save and exit EDIT.NLM by pressing **Esc**.

Part VIII: Begin the eDirectory Migration

Begin the eDirectory migration by doing the following:

1. In the Project Window, select **Begin NDS Migration**.
2. In the Migrate NDS window, select **Next**.
3. In the Install License window, mark **An MLA is already installed**; then select **Next**.

4. In the Update Schema window, update the source server's schema by making sure **Yes** is marked; then select **Next**.
5. In the Verify Novell Directory Services Tree window, verify that eDirectory is in good health by selecting **Yes**; then select **Next**.
6. In the Delete Connections window, delete all user connections (except your own) to the source and destination servers by selecting **Next**.
7. In the Password Verification window, enter your *password* in the Source Tree Password field.
8. In the Destination Tree Password field, enter your *password*; then select **Next**.
9. In the Migrate NDS Verification Results window, make sure no critical errors exist; then select **Next**.
10. In the Ready to Migrate NDS window, begin the eDirectory migration by selecting **Migrate**.

Notice that your source server (DA x) shuts down during migration, because it has been moved to your destination server.
11. In the Migrate NDS Results window, view the Error and Success logs; then select **Done**.
12. Close the current project by selecting **Close**.
13. Check the destination server and verify that it has restarted and taken on the name of the source server.
14. Close the migration wizard.
15. Remove any disks and CDs from your workstation.
16. Reboot your workstation and log in as **admin.IS.xxx.DIGITALAIR** (where xxx = your location container) to your NetWare 6 server.

To clear up any schema sync issues at the DA1 server console, enter the following:

```
Set DSTRACE=+SCHEMA
Set DSTRACE=+SYNC
Set DSTRACE=*SCHEMA
Set DSTRACE=*SS
Set DSTRACE=*H
```

Keep the VMWare servers that are not migrated running for the rest of the course.

Turn off only the migrated VMWare servers when the migration is complete.

Part IX: Finish eDirectory Migration

Finish the eDirectory migration by doing the following:

1. On your workstation, open the migration wizard by selecting **Start > Programs > Novell > NetWare Migration Wizard > NetWare Migration Wizard**.
2. In the About Novell NetWare Migration Wizard window, select **OK**.
3. Make sure **Open Last Project** is selected; then select **OK**.
4. In the Getting Started Migrating window, select **Close**.
5. In the Project Window, select **Finish NDS Migration**.
6. In the Continue NDS Migration window, mark **Yes**; then select **Next**.
7. In the Password field, enter **novell**; then select **Next**.
8. In the Ready to Continue Migrate NDS window, finish the eDirectory migration by selecting **Continue**.
9. In the Continue Migrate NDS Results window, select **View Error Log**.
10. Select **View Success Log**.
11. Scroll to the bottom of the success log and verify that the migration completed; then close the **log**.
12. In the Continue Migrate NDS Results window, select **Done**.
13. Close the current project by selecting **Close**.
14. Close the migration wizard.
15. Restart the DA_x server.

Part X: Complete eDirectory Cleanup Tasks

Complete eDirectory cleanup tasks by doing the following:

1. At the server console, enter
SET DSTRACE=ON
SET DSTRACE=+BLINK
SET DSTRACE=*B

This ensures that all your trustees are restored.
2. Switch to the Directory Services screen and verify that the external reference check process succeeded.
3. From your server console, enter **RESET SERVER**.

Part XI: Install NetWare Products and Services

To install necessary NetWare products and services, do the following:

1. Insert and mount the **NetWare 6 OS CD**.
2. From the GUI, select **Novell > Install**.
3. Select **Add**.
4. Select volume **NetWare6**; then select **OK** and allow the files to copy.
5. Select **OK**.
6. Deselect the following:
Storage Management Services
Novell Native File Access Pack
Novell Advanced Audit Service
7. Select **Next**.
8. Log in as Admin to your server.
9. In the Novell Certificate Server screen, select **Next**.

10. In the LDAP Configuration screen, select **Next**.
11. In the eDirectory iManage Install Options screen, select **Next**.
12. In the Summary screen, select **Finish**.
13. When the installation is complete, select **Close**.

Part XII: Install the Latest Support Pack

After upgrading your server's OS, install the latest support pack for NetWare 6.

Do the following:

1. Mount the **NetWare 6 SP2** CD as a NetWare volume on **DAx**.
2. At the server console, enter **NWCONFIG**.
3. In Configuration Options, select **Product Options**.
4. In Other Installation Actions, select **Install a product not listed**.
5. To specify the directory path, press **F3**.
6. In Specify a directory path, change A:\ to **NW6SP2:** (include the colon).
7. Press **Enter**.
8. In the Novell Terms and Conditions screen, press **Esc** to continue.
9. Accept the license agreement by selecting **Yes**.
10. In the License Agreement for JReport Runtime JInfonet software, press **Esc** to continue.
11. Accept the license agreement for JReport Runtime by selecting **Yes**.
12. Install the NetWare Support Pack version 6.0.2 by pressing **Enter**.

Alternative Installation: During classroom setup, copy the Support Pack CD to DA1\DATA\NW6SP2.

Have students install the support pack by entering DA1\DATA\NW6SP2 instead of completing Steps 1 and 6.

IDE CD Drives

If you are using older IDE CD drives in the classroom, instruct students to select **No** for step 14.

This prevents newer drivers from causing problems with older equipment.

13. In the Backup Files Replaced by NetWare Support Pack screen, select **No**.
14. In the Do You Want to Update the Storage/LAN/PSM/WAN Drivers Currently in Use screen, select **Yes**.
15. Reboot your server after the file copy by selecting **Yes**.
16. In the Warning screen, press **Enter** to continue.
File copy begins.
17. Authenticate using your full context and password; then allow files to copy and your server to reboot.
18. (Conditional) If prompted, do not press a key to exit.

Part XIII: Upgrade Directory Services

During the upgrade, you installed eDirectory 8.6. Now you need to upgrade eDirectory 8.6 to eDirectory 8.7.

Do the following:

1. Mount the **eDirectory 8.7** CD as a NetWare volume.
2. At the server console, load **NWCONFIG**.
3. From the Available Options menu, select **Product Options**.
4. Select **Install a Product Not Listed**.
5. (Conditional) If you receive the Close the Previously Specified Paths screen, press **Esc**.
6. Specify the path to the CD by pressing **F3**.
7. Specify the path to the NW directory where the installation program can find the NDS8.IPS file by entering **volume name:NW**.
For example, **EDIR_8_7:NW**.
8. Allow the files to copy.

Alternative Installation: During classroom setup, copy the eDirectory 8.7 CD to DA1\DATA\EDIR_8_7.

Have students install eDirectory by entering DA1\DATA:\EDIR_8_7\NW instead of completing Steps 1 and 7.

9. In the Software License Agreement screen, press **Esc** to continue.
10. Accept the license agreement.
11. In the License Agreement for JReport Runtime JInfonet Software screen, press **Esc**.
12. Accept the Reporting license agreement.
13. Continue by pressing **Esc**.
14. Read the warning; then press **Esc** and allow the files to copy.
15. In the Administrator Name field, enter your *full distinguished name*.



If you cannot log in as your admin, log in as ADMIN.DIGITALAIR.

16. In the Password field, enter your *password* and allow the files to copy and your server to reboot.
17. Authenticate to the Directory and allow the files to copy.
18. In the Are You Installing Remotely through RConsole screen, select **No-Local**.
19. In the Novell Certificate Server 2.40 Objects screen, select **Next**.
20. In the LDAP Configuration screen, select **Next**.
21. In the Novell Modular Authentication Service screen, select **Next**.
22. In the Next screen, select **Next**.
23. In the Components screen, select **Next**.
24. In the Summary screen, select **Finish** and allow the files to copy.
25. (Conditional) If the SNMP Object Creation Error dialog appears, select **OK**.

26. When the Installation Complete screen appears, remove the CD; then select **Yes** and allow the server to reboot.
27. At the server console, enter the following:
SET DSTRACE=ON
SET DSTRACE=+BLINK
SET DSTRACE=*B

This ensures that all your trustees are restored.
28. Switch to the Directory Services screen and verify that the external reference check process has succeeded.
29. From your server console, enter **RESET SERVER**.

Part XIV: Update the Novell Client

Point out that students update the client to show them the tasks involved in performing a thorough migration.

Now you upgrade the Novell Client to enable IP and IPX.

Do the following:

1. On your workstation, insert the **Novell Client** CD into your CD drive.
2. In the Client Installation window, select your *language*.
3. Select **Novell Client 4.83 for Windows NT/2000/XP**.
4. In the Novell Client Installation window, mark **Custom Installation**; then select **Next**.
5. In the Components to Install window, select **Next**.
6. In the Protocol Preference window, make sure **IP and IPX** is marked; then select **Next**.
7. In the Login Authenticator window, make sure **NDS** is selected; then select **Next**.
8. Complete the installation by selecting **Finish**; then allow the files to copy.
9. On the Installation Complete window, select **Reboot**.

Part XV: Install User Licenses

Install a user license certificate by doing the following:

1. On your workstation, launch iManager.
2. Log in as *admin* with a password of *novell* to your *xxx container* (where *xxx* = your location container).

For example, if you are the admin of IS.LON.DIGITALAIR, your location container is LON.DIGITALAIR.
3. In the left frame, select **License Management**.
4. Select **Install a License**.
5. On your workstation, insert your license disk.
6. Next to the Load License File field, select **Browse**.
7. On your license disk, browse to and select the *NLF file* in **DAx**.
(Don't select the NLF file in the SERVER LICENSE folder.)
8. Select **Open**.
9. Select **Next**.
10. Mark **Select Certificates**.

This selects your user license.
11. Select **Next**.
12. In the Location field, enter *xxx.DIGITALAIR* (where *xxx* = your location container).
13. Browse to and select your **DAx** server.
14. Select **Install**.
15. Verify that your license was successfully installed; then select **Done**.
16. Close the iManager window.

At the end of this exercise, perform a full unattended repair on DA1 to resolve network problems caused by the migrations.

(End of Exercise)

Objective 5 Troubleshoot Post-Installation Issues

In this objective you learn about the following post-installation issues:

Table 1-6

Issue	Solution
Missing device drivers	<p>The server installation program copies to a startup directory (C:\NWSERVER) only drivers (such as HAMs, CDMs and PSMs) for devices that were autodetected during the installation.</p> <p>If you attempt to load a HAM, CDM, or PSM that was not autodetected during installation and it fails to load, copy the appropriate driver from the C:\NWSERVER\DRIVERS directory to the C:\NWSERVER directory and then load the driver again.</p>
Status of old LAN and WAN files	<p>After an upgrade to NetWare 6, old LAN and WAN files are not deleted. These old files might not be supported in a NetWare 6 environment.</p>

Table 1-6 (continued)

Issue	Solution
Speeding up the post-installation utility	<p>If performance of the post-installation utility and other Java applications is slow, change the VM Cache Pool Percentage SET parameter by entering the following at the server console:</p> <p>SET VM CACHE POOL PERCENTAGE = 30</p> <p>The performance of the post-installation program and some Java applications improves significantly with this change.</p> <p>The changes are saved by the OS and remain even if the server is rebooted.</p>
Update SMS components	<p>A new release of Storage Management Services (SMS) is available on the Novell Support web site. This release contains important fixes to the SMS components delivered with NetWare 6.</p> <p>The fixes ensure compatibility between earlier versions of NetWare and the NetWare 6 SMS modules.</p> <p>In addition, the patch includes updates to SMS components (including SMDR and TSA), that make the product more stable and robust.</p> <p>To maintain backup and restore services on NetWare 6 and your network, install the patch. The patch supersedes the SMS modules installed by default during the NetWare 6 installation.</p> <p>For more information, see the Readme included with the patch.</p>

Table 1-6 (continued)

Issue	Solution
Agent installation for GroupWise 6 replaces LDAP files needed by iFolder	<p>The agent installation program for GroupWise 6 lets you overwrite LDAP modules that disable iFolder running on NetWare 6.</p> <p>To avoid this problem, select No when prompted to overwrite the LDAP modules during GroupWise agent installation. The LDAP module includes</p> <p>LDAPSDK.NLM LDAPSSL.NLM LDAPX.NLM</p> <p>If these files are overwritten (by selecting Yes), you must manually copy the files from the NetWare 6 OS CD before iFolder will run on NetWare 6.</p>
Missing user licenses	<p>NetWare 6 requires you to install user licenses separately from the server license.</p> <p>After migrating to NetWare 6, use iManager to install license certificates when you add NetWare servers and users to the eDirectory tree.</p>



For issues that can happen before and during installation and migration, see Known Issues at the NetWare 6 documentation web site at www.novell.com/documentation/lg/nw6p/index.html.

Summary

The following is a summary of the objectives in this section:

Objective	What You Learned
1. Review How to Prepare for a Server Migration to NetWare 6	<p>To prepare for a migration, you must do the following:</p> <ul style="list-style-type: none">■ Prepare the workstation■ Prepare the source (original) server■ Prepare the destination (new) server■ Prepare server application files
2. Review How to Implement Novell Licensing	<p>Novell's licensing technology lets you manage license units that are required to comply with the licensing requirements of Novell.</p> <p>To identify how server and user licensing works in NetWare 6, you must understand the following:</p> <ul style="list-style-type: none">■ Server and user licensing models■ How UAL coexists with SCL■ How the licensing models differ■ License types
3. Identify How to Perform a Migration	<p>To migrate data from NetWare 4, NetWare 5, or NetWare 6 after you prepare the network and server, do the following:</p> <ul style="list-style-type: none">■ Run the migration wizard■ Copy volumes■ Edit configuration files■ Begin the eDirectory migration■ Finish eDirectory migration

Objective	What You Learned
4. Perform Post-Migration Tasks	After the migration is complete, you do the following to ensure that your new NetWare 6 network is running properly: <ul style="list-style-type: none"><li data-bbox="954 401 1409 434">■ Run the external reference check process<li data-bbox="954 443 1305 476">■ Upgrade existing NSS volumes<li data-bbox="954 485 1354 518">■ Perform other post-installation tasks
5. Troubleshoot Post-Installation Issues	You learned how to resolve the following issues: <ul style="list-style-type: none"><li data-bbox="954 569 1214 602">■ Missing device drivers<li data-bbox="954 611 1321 644">■ Status of old LAN and WAN files<li data-bbox="954 653 1380 686">■ Speeding up the post-installation utility<li data-bbox="954 695 1250 728">■ Update SMS components<li data-bbox="954 737 1429 791">■ Agent Installation for GroupWise 6 replaces LDAP files needed by iFolder<li data-bbox="954 800 1208 833">■ Missing user licenses

MODULE 2

Troubleshoot and Resolve Novell Network Problems

- Section 2** Identify Tools for Troubleshooting Novell Network Performance Issues
- Section 3** Troubleshoot and Resolve NetWare Server Issues
- Section 4** Monitor and Troubleshoot eDirectory

SECTION 2 Identify Tools for Troubleshooting Novell Network Performance Issues

Duration: 3 hours 30 minutes

In this section you learn about the tools available for troubleshooting a Novell network and how to troubleshoot problems in a mixed IP/IPX LAN environment.

Objectives

1. Upgrade Novell Network Management Tools
2. Identify the Troubleshooting Features of Novell NetWork Management Tools
3. Identify the Purpose and Function of IP/IPX Troubleshooting Tools
4. Identify Additional Network Troubleshooting Resources

Introduction

As you consider troubleshooting problems on the network, you realize how complex your network is. The following identifies common network communication problems:

- Workstations can't communicate with the server
- Connections are dropped periodically
- The web browser cannot access a web site
- Slow network response time

To troubleshoot a problem, you must consider if the problem is the result of one or more of following factors:

- Hardware
- Software
- Configuration
- Human error

When troubleshooting networking problems you should focus on the 3 main areas (in order) that Novell Customer Services uses for troubleshooting: the LAN, the server, and eDirectory.

To troubleshoot LAN, server, and eDirectory issues, you must have a solid understanding of what a network is and the components that comprise a network.

With the knowledge you gained from previous courses, you should have a firm understanding of the components that make up a network, and be able to determine where in the network communication process a specific problem might occur.



For a review of the fundamental network components, see *Appendix A: Network Components*.

For a review of the network communication process, see *Appendix B: The Network Communication Process*.

A critical factor in identifying and troubleshooting your network is the network administration tools you use.

The following reviews the Novell network administration management tools you should include as part of any NetWare 6 upgrade or migration.

Objective 1 Upgrade Novell Network Management Tools

Whenever you upgrade NetWare or any of its components by installing a new version or support pack, you should also upgrade the Novell network administration management tools.

By upgrading these tools, you can take advantage of improved performance, remote management capabilities, and additional features that can enhance your network management capabilities and reduce the time required to manage the network.

The following are key Novell network management tools you should upgrade:

Table 2-1

Tool	Prerequisites	Area of Focus	Limitations
ConsoleOne	Platform specific	Available on the following platforms: <ul style="list-style-type: none"> ■ Windows ■ NetWare ■ Linux ■ Solaris ■ Tru64 UNIX Provides <ul style="list-style-type: none"> ■ eDirectory administration ■ File and volume management ■ eDirectory partition and replication management 	<ul style="list-style-type: none"> ■ Does not provide DS process tracking or repair tools.

Table 2-1 *(continued)*

Tool	Prerequisites	Area of Focus	Limitations
ConsoleOne Reports (Windows only)	<ul style="list-style-type: none"> ■ Available only on the Windows version of ConsoleOne ■ One NetWare volume ■ 128 MB RAM on the Windows client 	<ul style="list-style-type: none"> ■ eDirectory ■ File and volume management 	<ul style="list-style-type: none"> ■ You must install ConsoleOne locally or use a drive map. ■ You must extend the schema and install report capabilities separately.
iMonitor	<ul style="list-style-type: none"> ■ Web browser 	<ul style="list-style-type: none"> ■ eDirectory 	<ul style="list-style-type: none"> ■ Server specific, with proxy capabilities for most functionality. For example, when in proxy mode, you cannot perform a Repair on that server.
Novell Remote Manager	<ul style="list-style-type: none"> ■ Web browser 	<ul style="list-style-type: none"> ■ Server maintenance 	<ul style="list-style-type: none"> ■ Server specific. ■ Requires Port configuration when accessing through a firewall.

Table 2-1 *(continued)*

Tool	Prerequisites	Area of Focus	Limitations
iManager	<ul style="list-style-type: none"> ■ Web browser 	<ul style="list-style-type: none"> ■ eDirectory administration ■ Server maintenance ■ Licensing ■ DNS and DHCP ■ Dynamic group management ■ Partition and replication management ■ Rights management 	<ul style="list-style-type: none"> ■ Server specific. ■ Requires Port configuration when accessing through a firewall. ■ Interface limitations.



You must use iManager navigational buttons. The browser's back and forward buttons do not work.

Exercise 2-1 Upgrade Your Novell Network Management Tools

As network administrator for your Digital Airlines office, you have just completed migrating a NetWare 4.11 server to NetWare 6. You have also installed SP2 and eDirectory 8.7.

To make sure you have the latest NetWare management tools available, you need to do the following:

- [Part I: Update ConsoleOne](#)
- [Part II: Update iManager](#)

Part I: Update ConsoleOne

Prior to the eDirectory upgrade, you had ConsoleOne 1.3.3 installed on your DAX server. With the upgrade, ConsoleOne 1.3.4 was installed. Now you need to update ConsoleOne on your workstation by doing the following:

1. Right-click **Start**.
2. Select **Explore**.
3. Browse to **DAX\SYS:\PUBLIC\MGMT\CONSOLEONE**.
4. Copy the **1.2** folder.
5. Paste the 1.2 folder on your workstation's hard drive at **C:\NOVELL\CONSOLEONE**.
6. Overwrite all existing files by selecting **Yes to All**.

This should replace the older version of the 1.2 folder on your workstation and allow you to use ConsoleOne 1.3.4.

7. From your workstation desktop, start **ConsoleOne** and check the splash screen.

Make sure you see 1.3.4 for the version number.

8. Exit **ConsoleOne**.

Alternative Installation: For Step 2, you can also enter **JAVA -CP {volume_name}: installs \nwMonitorInstall.jar install.**

If students have problems accessing iManager, do the following:

1. Open the **HOSTNAME** file and verify that all instances of TEMPx have been changed to DAx.
2. Open **SYS:\APACHE\CONF**.
3. Open **ADMINSERV.CONF**; then search and replace any instance of TEMPx with DAx. Two or 3 instances should be changed in the virtual host information.
4. Reset the server and access iManager again.



Part II: Update iManager

Update iManager 1.2 to iManager 1.5 by doing the following:

1. Mount the **Web Apps** CD as a NetWare volume.
2. At the server prompt, enter **EDIRWEBAPPS:WEBAPP.NCF**.
3. Select your *language*; then select **OK**.
4. In the Novell eDirectory Web Applications screen, select **Next**.
5. Accept the license agreement by marking **I accept the terms of the License Agreement**; then select **Next**.
6. Deselect **Novell eGuide - eDirectory White Pages**; then select **Next**.

For a demonstration of Novell eGuide, see www.novell.com.

7. Make sure Novell iManager is listed to be installed; then select **Install** and allow the eDirectory iManager wizard to launch.
8. Select your *language*; then select **OK**.
9. In the iManager screen, select **Next**.
10. In the Detection Summary screen, accept the default settings by selecting **Next**.
11. In the Pre-Installation Summary screen, select **Install** and allow files to install.
12. After iManager is installed, select **Done**.
13. After the web applications are installed, launch the web browser by selecting **Done**.
14. In your browser, select the **Getting Started** link.
15. From your workstation, launch your *browser*.
16. In the Location field, enter **HTTPS:\\192.168.1.x:2200** (where *x* = *your server number*).

17. In the Security Alert window, select **Yes**.
18. Under eDirectory iManager, select *your server*.
19. On the Login window, enter the *appropriate information*.
20. Select **Login**.
21. Set up role-based services:
 - a. In the Container field, enter **IS.xxx.DIGITALAIR** (where *xxx = your location container*); then select **Next**.
For example, if you are the admin in IS.LON.DIGITALAIR, you enter IS.LON.DIGITALAIR in the Container field.
 - b. In the Scope field, enter **xxx.DIGITALAIR**.
For example, if you are the admin in IS.LON.DIGITALAIR, your scope is LON.DIGITALAIR.
 - c. Select **Start**.
 - d. When you receive the message to close the Role Based Service window, select **Close**.

(End of Exercise)

Objective 2 Identify the Troubleshooting Features of Novell NetWork Management Tools

Students are familiar with some of the functionality of the tools. This section helps them identify which tool can be used for troubleshooting networking components.

You can use the following Novell network management tools to administer your network, monitor processes, and troubleshoot problems:

- [ConsoleOne Reports](#)
- [Novell iMonitor](#)
- [NetWare Remote Manager](#)
- [Novell iManager](#)



You can access all Novell web management tools (except iMonitor) from NetWare Web Manager (<https://your server IP address:2200>). This includes management tools for services such as iFolder.

However, if you change the port numbers for services (such as iFolder) you must change the port numbers in SERVERS.ORG and SERVERS.XML (in SYS:\WEBAPPS\WEBADMIN) for the management tool link to work properly in the NetWare Web Manager interface.

Introduce ConsoleOne reports as a troubleshooting tool.

ConsoleOne Reports

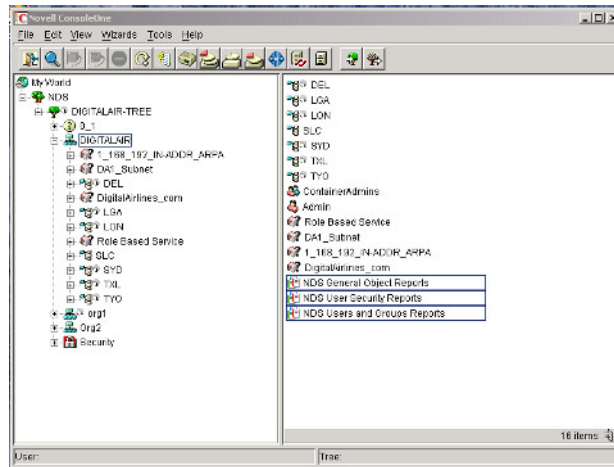
With the release of ConsoleOne 1.3 for Windows, you can generate predefined reports to help you troubleshoot network problems. The other platforms do not support the Reports functionality.

To run reports, you must meet the following prerequisites:

- ConsoleOne must be installed on a Windows workstation with 128 MB of RAM.
- Your eDirectory tree must contain a NetWare volume to install the report catalog files.

After you extend the schema to support reporting you must install the desired reports. After the reports are installed, they appear in the tree, as shown in the following:

Figure 2-1



You must extend the schema and install reporting before you can demonstrate the reports option in ConsoleOne.

The following predefined report categories are available:

- **eDirectory General Object Reports.** The reports for general objects include NetWare file servers, print servers, and printers. These reports provide information and status for each object.
- **eDirectory User Security Reports.** This report catalog contains report forms that let you generate reports on eDirectory login and rights security for users in your tree. The following reports are available:
 - Disabled User Accounts
 - Users Locked by Intruder Detection
 - Security Equivalence
 - Template Security Settings
 - Trustee Security Settings
 - Trustee Assignments

- ❑ User Password Requirements
- ❑ Users Not Logged In
- ❑ Users with Expired Password
- ❑ Users with Multiple Workstation Logins
- **eDirectory User and Group Reports.** This report catalog contains report forms that let you generate reports on the users, groups, and organizational roles in your eDirectory tree.



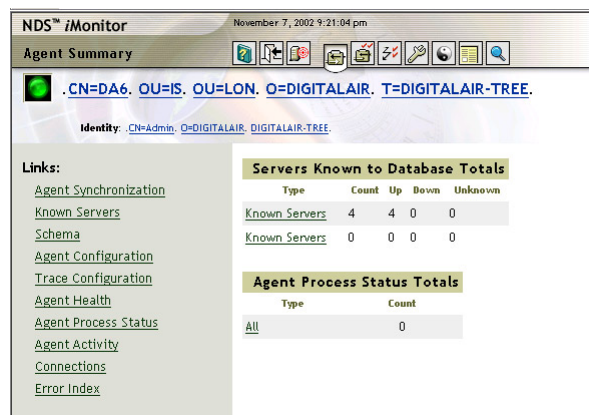
You must extend the schema to access the menu items for most ConsoleOne snap-ins.

Novell iMonitor

iMonitor provides cross-platform monitoring and diagnostic capability for all servers in your eDirectory tree.

The following shows iMonitor options:

Figure 2-2



Tell students that they use iMonitor for most administration tasks and troubleshooting in this section.

Show each of the following screens as you introduce the troubleshooting tools in iMonitor.

Tell students that NetWare Remote Manager is designed for server administration, but it can link to iMonitor for eDirectory administration as well.

You can access iMonitor from a web browser by entering **https://your server IP address:8009/nds-summary**. The address is case-sensitive.

Using iMonitor, you can monitor your servers from any location where a web browser is available.

iMonitor lets you look at the eDirectory environment in depth on a partition, replica, or server. You can also examine what processes are taking place, when they are happening, what their results are, and how long they take.

The following is a list of troubleshooting tools in iMonitor:

- Reports
- Trace
- Repair

iMonitor's diagnosing and troubleshooting capabilities replace the tools Novell developers created to debug or troubleshoot their code.

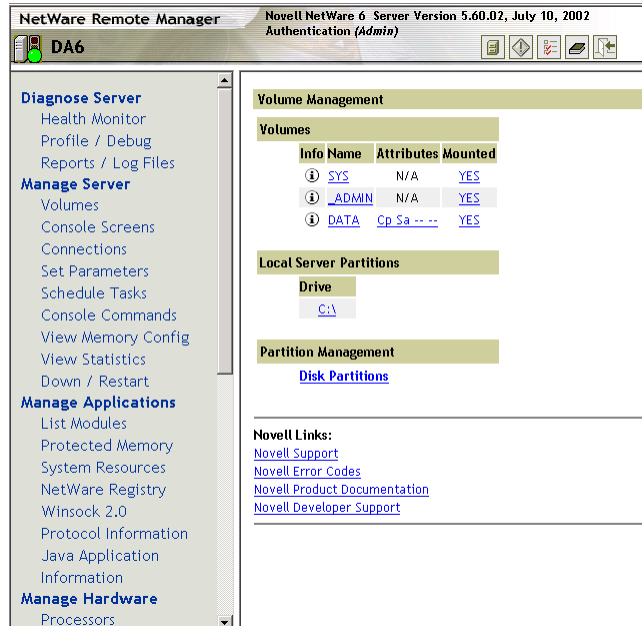
NetWare Remote Manager

NetWare Remote Manager lets you use a web browser to securely access NetWare servers from any workstation and perform specific server management tasks.

You can access NetWare Remote Manager from a web browser by entering **https://your server IP address:8009**.

The following shows some of the options:

Figure 2-3



Using Novell Remote Manager, you can monitor the health of your servers, their processes, and CPU usage. You can also perform common server management tasks such as

- Mounting and dismounting volumes
- Managing server connections
- Configuring SET parameters
- Viewing the server configuration
- Accessing files on volumes and DOS partitions
- Shutting down, restarting, and resetting your server.

Point to the troubleshooting tools in Novell Remote Manager.

Using the Console Screens feature, you can view and run all the console screens just as though you were using the keyboard at the server console.

The following is a list of Novell Remote Manager troubleshooting tools.

- Health monitor
- Profile/debug server
- Report/log files
 - Server personal log book
 - System error log file
 - Abend log file
 - Server health log file

Open iManager and review the administration options.

iManager version 1.5, which ships with eDirectory 8.7 on the Web Applications CD, includes a number of new tools.

Novell iManager

Novell iManager 1.5 is a web-based application for managing, maintaining, and monitoring eDirectory using wired and wireless devices.

Traditionally, eDirectory has been managed through utilities such as NetWare Administrator and ConsoleOne. These tools are platform specific and allow management by browsing all objects in a tree.

When an object that needs to be administered is found, it can be managed only if the object snap-in is loaded and the user has rights to administer the object.

To use iManager for management and troubleshooting, you need to know the following:

- [Role Based Management](#)
- [New iManager Administration Tools](#)
- [iManager Troubleshooting Tools](#)

You can access iManager from a web browser by entering **https://your server IP address:2200/eMFrame/iManager**. The address is case sensitive.

Role Based Management

To manage objects, administrative users must be assigned task-specific administrative roles.

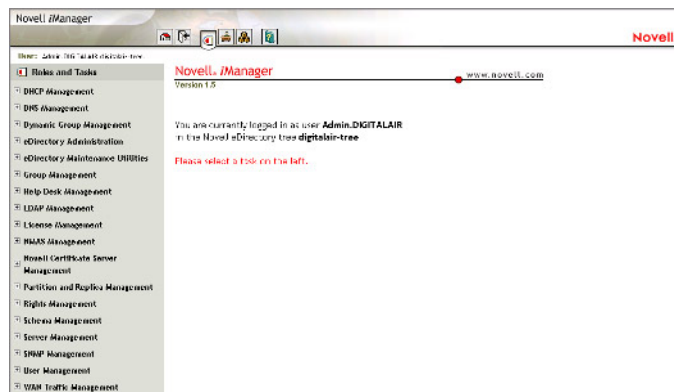
As the administrator, you can assign roles to other administrative users in your organization. If assigned a role, that user has the same rights assigned to the role enabling them to perform all tasks contained in that role.

In addition, iManager lets you organize available tasks into customized roles that suit the job roles of your organization.

New iManager Administration Tools

A number of new administration tools (as shown in the following) have been added to iManager version 1.5, which ships with eDirectory 8.7 on the Web Applications CD:

Figure 2-4



The following is a list of some of the new administration categories:

- Dynamic group management
- eDirectory administration
- eDirectory maintenance utilities
- Group management
- Novell certificate server management
- Partition and replication management
- Rights management

iManager Troubleshooting Tools

Many of the items listed as tools also contain troubleshooting capabilities, such as

- Links to iMonitor, to perform repair tasks
- A link to Novell Remote Manager, to perform server maintenance
- Rights management
- Schema management
- Server management
- WAN traffic management

Objective 3 Identify the Purpose and Function of IP/IPX Troubleshooting Tools

In this objective you learn about the following:

- [NetWare IP/IPX Troubleshooting Tools](#)
- [Client IP Troubleshooting Tools](#)
- [Protocol Analyzers](#)
- [TCP/IP Toolkits](#)
- [IP Addressing Calculators](#)

Most of these tools ship with the operating systems being discussed or are available for download.

NetWare IP/IPX Troubleshooting Tools

NetWare's configuration files and troubleshooting tools help you troubleshoot IP and IPX problems by giving you information about network configuration, status of communications, and links within an internetwork.

These tools include the following:

- **CONFIG.** Shows configuration settings, including IPX and IP addresses on this server.
- **NSLOOKUP.** Enables you to identify a DNS server by domain name or IP address.
- **HOSTS.** Contains information you enter about hosts on the IP network.
- **TCPCON.** Provides general TCP/IP stack configuration and performance statistics.
- **PING.** Provides an end-to-end connectivity test (menu-driven).
- **DEBUG.** Provides communication dumps and recording.

CONFIG

Entering CONFIG at the server console is one of the first things you should do when troubleshooting your server. This utility returns the following information:

- The file server name
- The internal network number of the file server
- The loaded LAN drivers
- The hardware settings on all network boards
- The node (station) addresses of the network boards
- The communication protocol bound to the network board
- The network number of the cabling scheme for a network board
- The frame type assigned to the board (more than one frame type can be assigned to Ethernet and Token-Ring boards)
- The board name assigned

Use CONFIG before installing network boards in the file server so that you have a current list of all hardware settings on the network boards.

NSLOOKUP

Demonstrate NSLOOKUP and point out the DNS server name and IP address.

Use NSLOOKUP at the server console to identify your DNS configuration, to diagnose DNS setup problems, or to identify DNS problems in an application.

NSLOOKUP is available only if TCP/IP is installed. Its primary function is to query DNS name servers.

Using NSLOOKUP, you can perform a forward DNS lookup (matching a domain name to an IP address) or a reverse DNS lookup (matching an IP address to a domain name).

To use NSLOOKUP, do the following:

1. Load NSLOOKUP at the server console prompt by entering **NSLOOKUP**.

The default DNS server and its IP address appear.

The prompt indicates that NSLOOKUP is active and will remain active until you enter **Exit** to close the utility.

2. View a list of NSLOOKUP commands for querying DNS by entering **?** or **HELP**.
3. Enter one or more commands.

For example, to view the SET commands on your server, enter SET ALL.

4. When you finish using NSLOOKUP, exit the NSLOOKUP utility by entering **EXIT**.

To demonstrate NSLOOKUP, enter **HELP** to show the available commands. Enter 2 or 3 commands and discuss the information shown.



For optimum performance, NSLOOKUP uses the SYS:\ETC\RESOLV.CFG file to obtain the DNS configuration information.

If this file is missing or is not configured, NSLOOKUP queries can't show the information.

You can configure the RESOLV.CFG file with the correct DNS configuration information and then exit and reload NSLOOKUP.

HOSTS and HOSTNAME Files

The SYS:ETC\HOSTS and SYS:ETC\HOSTNAME files store information you enter about the hosts on the IP network.

The HOSTS file entry has the following format:

```
IP_address host_name [alias [...]]
```

Each entry provides information about a single host and cannot extend beyond one line.

The HOSTNAME file entry has the following format:

```
IP_address host_name
```

The host_name is the name of the system associated with the internet address. The name cannot contain a space, tab, number sign (#), or end-of-line character. Each host name must be unique.

In the HOSTS file, a single host can have from one to ten aliases. The alias is another name for the same system. Typically, this is a shorter name. For example, the host Sales could have the following address and aliases:

```
139.0.9.5 sales sa
```

TCPCON

TCPCON monitors TCP/IP operations and provides detailed information on the status of network segments, protocols, routing tables, and the SNMP trap log.

TCPCON lets you view the configuration and statistics for the SNMP target only, and is an excellent troubleshooting utility for viewing errors that occur within the TCP/IP stack.

To use TCPCON, do the following:

1. At the server console, enter **TCPCON**.

From the initial screen, you see the IP Forwarded field. If any value is entered into this field, including zero, the server is configured to act as an IP router.

If this entry has **DISABLED** after the statistics, it is not set to gateway mode.

2. (Optional) To enable gateway mode, do the following:
 - a. From the server console, enter **INETCFG**.
 - b. Select **Protocols > TCP/IP**.
 - c. Enable **IP Packet Forwarding**.

The following is a brief overview of the areas and statistics available in TCPCON.

Table 2-2

Section	Purpose
SNMP Access Configuration	<ul style="list-style-type: none"> ■ Defines the SNMP target. Can use a hostname (if configured) or IP address. ■ Defines polling and timeout intervals for SNMP GET requests. ■ Defines community name for SNMP queries.
Protocol Information	<p>Examines detailed statistics on 6 primary protocols:</p> <ul style="list-style-type: none"> ■ Exterior Gateway Protocol (EGP) ■ Internet Control Message Protocol (ICMP) ■ Internet Protocol (IP) ■ Open Shortest Packet First (OSPF) ■ Transmission Control Protocol (TCP) ■ User Datagram Protocol (UDP) <p>Although similar to the Statistics section in TCPCON, this section contains more details on the performance and configuration of the protocols.</p>
IP Routing Table	Shows, refreshes, and manipulates IP routing tables.
Statistics	Shows basic statistics on EGP, ICMP, IP, OSPF, TCP, and UDP. All information in this section can be viewed through the Protocol Information options.

Table 2-2 (continued)

Section	Purpose
Interfaces	Shows interface-specific statistics, such as MAC address, MTU, speed, and bytes sent/received. Also shows unicast and non-unicast sent/received statistics.
Display Local Traps	Shows the local SNMP trap log entries, if any exist. Requires SNMPLOG.NLM on the server. Cannot be used to view remote trap log entries.

3. For help on any TCPCON option, press **F1**.
4. When you finish, press **Esc** to exit.



By default, TCPCON is configured to monitor the loopback address 127.0.0.1 with the community name PUBLIC.

PING

PING is typically the first tool to use when connectivity issues exist. If you can't access a server, try pinging it.

PING continuously transmits echo requests by default.

When you initiate a PING, the ICMP echo (ICMP type 8) and ICMP echo reply (ICMP type 0) packets are used to verify communication between the devices.

The default packet size for PING is 64 bytes.

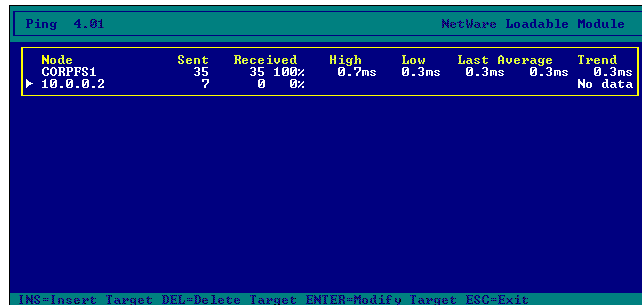
When your system is configured correctly, you can ping a device using the device name or the device's IP address. You can also PING the loopback address (127.0.0.1) to validate the local stack.

From the server console, you can ping another server by entering one of the following:

- PING *IP_address*
- PING *host name*

The following shows the results of 2 simultaneous PING tests:

Figure 2-5



The screenshot shows a NetWare console window titled "Ping 4.01" and "NetWare Loadable Module". It displays the results of two ping tests. The first test is to a device named "CORPFS1", which is successful with a 100% completion rate and a roundtrip time of 0.3ms. The second test is to the IP address "10.0.0.2", which is failing with a 0% completion rate. The console also shows high and low roundtrip times, last roundtrip time, average roundtrip time, and a trend of roundtrip times.

Node	Sent	Received	High	Low	Last	Average	Trend
CORPFS1	35	35	0.7ms	0.3ms	0.3ms	0.3ms	0.3ms
▶ 10.0.0.2	7	0					No data

The first line shows the PING test results to a device named CORPFS1. This test is successful and is supplying good responsive roundtrip times.

The second line shows that a test to a device with the IP address 10.0.0.2 is failing. The source has sent out 7 packets, but it has not received anything back. The completion rate is 0%.

This is a good indication that either the device 10.0.0.2 is down or there is no path to the device.

The PING utility provides the high and low roundtrip times, as well as the last roundtrip time, average roundtrip time, and the current trend of roundtrip times.

DEBUG

The NetWare server supports many debug screens that can help you identify and resolve TCP/IP-based problems.



Many of these debug screens are I/O intensive and can affect server performance. Use these screens with caution.

The following lists available DEBUG settings:

- **SET TCP IP DEBUG = 0/1 (default = 0).** This command shows all incoming and outgoing packets processed by TCPIP.NLM. This information is useful in debugging problems with packet translation, filtering, or connection issues.

The screen might scroll down too quickly to view. In this case, run CONLOG.NLM, change the SET TCP IP DEBUG to 1, and then unload CONLOG.

The CONSOLE.LOG file in SYS:ETC contains the output of the IP DEBUG screen.

- **SET TCP TRACE = 0-4 (Default = 0).** This command shows the following information about the current state of the TCP connection table:
 - Mode 1 shows basic information at the server console.
 - Mode 2 shows basic information at both server console and SYS:ETC\TCPxxx.LOG (where xxx begins with 0000).
 - Mode 3 shows advanced information at the server console.
 - Mode 4 shows advanced information at both server console and SYS:ETC\TCPxxx.LOG.
- **SET TCP IPCP DEBUG = 0-4 (Default = 0).** This command shows negotiated IPCP options when establishing a call over PPP with MPR. Mode 1 lists the calls. Subsequent mode numbers dump information from within each call.

- **SET TCP RIP DEBUG = 0-4 (default = 0).** This command can be used to verify RIP transmissions:
 - Mode 1 shows RIP send information at the server console.
 - Mode 2 shows RIP send information at the server console.
 - Mode 3 shows RIP send/receive information at the server console.
 - Mode 4 shows RIP send/receive information at the server console.
- **SET TCP WAN DEBUG = 0-4 (Default = 0).** This command shows IP debug information at the server console when a WAN call is made with the MPR. Mode 1 lists the calls. Subsequent modes show information from within each call.

Client IP Troubleshooting Tools

The following IP tools can be used on a client (such as a Windows 2000 workstation) to isolate problems on the network:

- **IPCONFIG**
- **PING**
- **ROUTE**
- **TRACERT (Trace Route)**
- **NSLOOKUP**
- **ARP**
- **NETSTAT**

IPCONFIG

Use IPCONFIG to identify your IP configuration for Windows 2000/XP and Windows NT workstations. This command is equivalent to WINIPCFG, which is used to configure Windows 95/98 and ME.

When you enter IPCONFIG with no parameters, IPCONFIG shows the IP address, subnet mask, and default gateway for each adapter bound to TCP/IP, as seen in the following:

Figure 2-6

```
C:\WINNT\System32\cmd.exe
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.

C:\>ipconfig

Windows 2000 IP Configuration

Ethernet adapter Local Area Connection 2:

    Media State . . . . . : Cable Disconnected

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix . : provo.novell.com
    IP Address . . . . . : 137.65.104.43
    Subnet Mask . . . . . : 255.255.252.0
    Default Gateway . . . . . : 137.65.107.254

Ethernet adapter {37BB249D-56D1-4B01-0174-93E9F466BC0E}:

    Connection-specific DNS Suffix . :
    IP Address . . . . . : 0.0.0.0
    Subnet Mask . . . . . : 0.0.0.0
    Default Gateway . . . . . :

Ethernet adapter VMware Virtual Ethernet Adapter (basic host-only support for VM
net1):

    Connection-specific DNS Suffix . :
    IP Address . . . . . : 192.168.154.1
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . :

Ethernet adapter VMware Virtual Ethernet Adapter (Network Address Translation (N
AT) for VMnet8):

    Connection-specific DNS Suffix . :
    IP Address . . . . . : 192.168.153.1
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . :

C:\>_
```

IPCONFIG /all shows all current TCP/IP configuration values, including the IP address, subnet mask, default gateway, and WINS and DNS configuration.

To use this utility, do the following:

1. From the Windows Start menu, enter CMD.

2. At the DOS command line, enter IPCONFIG with one or more of the following parameters:

Table 2-3

Parameter	Function
/?	Accesses utility help.
/all	Shows full configuration information.
/release	Releases the IP address for the specified adapter.
/renew	Renews the IP address for the specified adapter.

The syntax is

```
IPCONFIG [/? | /all | /release [adapter] | /renew [adapter]]
```

PING

Similar to the PING utility on the server, you can use this command from the Windows command line.

You use PING to query another IP device on the network to determine if it is active and how long a packet takes to get there.

This is one of the first tools you should use for any problem that appears to be caused by a lack of connectivity between network devices.

The syntax for the PING command is

```
PING [-t] [-a] [-n count] [-l size] [-f] [-i TTL] [-v TOS] [-r count] [-s count] [[-j host-list] | [-k host-list]] [-w timeout] destination-list
```

The following explains the PING parameters:

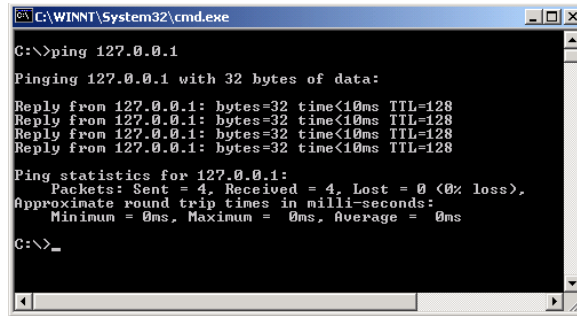
Table 2-4

Parameter	Function
-t	Pings the specified host until interrupted
-a	Resolves addresses to hostnames
-n count	Defines the number of echo requests to send
-l size	Defines the PING packet size
-f	Sets the Don't Fragment flag in packet
-i TTL	Defines the Time To Live value
-v TOS	Defines the Type Of Service value
-r count	Records the route for count hops
-s count	Defines the time stamp for count hops
-j host-list	Defines the loose source route along host list
-k host-list	Defines the strict source route along host list
-w timeout	Defines the time (in milliseconds) to wait for each reply

If a device is having problems communicating, enter PING 127.0.0.1 on that device. The address 127.0.0.1 is the loopback address.

The station will ping its own IP stack, as shown in the following:

Figure 2-7



```
C:\WINNT\System32\cmd.exe
C:\>ping 127.0.0.1
Pinging 127.0.0.1 with 32 bytes of data:
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Reply from 127.0.0.1: bytes=32 time<10ms TTL=128
Ping statistics for 127.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\>_
```

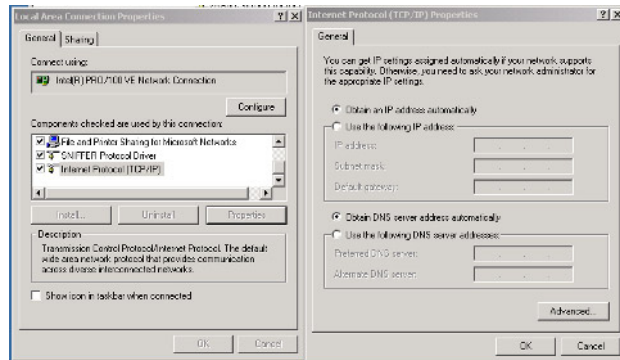
If it does not see that its own stack is active, it cannot communicate on the network.

When the workstation cannot ping its own stack, do the following:

1. Access **My Network Places > Properties > Local Area Connection > IP Protocol**.
2. Identify the following:
 - Verify the IP address information is acquired using one of the following methods:
 - The IP address is automatically assigned.
 - A dedicated IP address has been entered and the information is correct.
 - Verify the DNS settings are established using one of the following methods:
 - The DNS server address is automatically assigned.
 - A dedicated DNS server address is assigned and the information is correct.

If you are using DHCP and your TCP/IP configuration is automatic, your dialogs will look like the following:

Figure 2-8



Another tool you can use to check your communications is the routing table. The routing table defines the path a packet takes to get to its destination.

ROUTE

You use ROUTE to determine the path an IP packet is taking to reach a destination server.

When the client needs to route a packet to a remote destination, it looks at its local routing table first.

The client's routing table is dynamically learned from the network, but entries can be placed in the table manually.

For example, suppose Michelle is sending a PING packet to Drake. Michelle's workstation IP stack looks in the local routing tables to see if an entry for Drake's host address (204.10.11.5) exists.

If Michelle's workstation does not have an entry for that specific host address, the IP stack looks for a network entry (an entry for 204.10.11.0). If no entry exists for the host or the network, the IP stack looks for a default gateway setting.

In this case, Michelle's station has a default gateway (204.10.10.3) that connects to the Internet. That is not the best path to network 204.10.11.0, but Michelle's station does not know that. The workstation sends the packet to the default gateway.

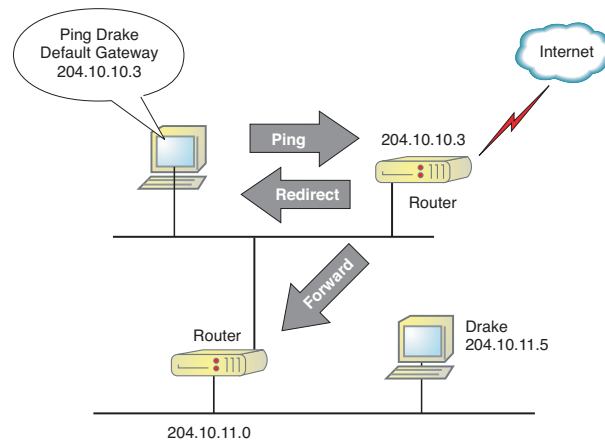
The default gateway returns an ICMP packet to the client that indicates there is a better route through router 204.10.10.4.

This packet, which is called a Redirection message, dynamically updates Michelle's workstation routing tables.

The next time Michelle sends a packet to Drake, she will find a network entry for 204.10.11.0 in her routing tables that indicates she should forward such traffic to router 204.10.10.4.

The following illustrates this example:

Figure 2-9 (slide)





If almost all of a host's packets to other devices must be rerouted, the specified default gateway is probably not the most appropriate.

The syntax for ROUTE is

```
ROUTE [-f] [command] [destination] [MASK netmask]
[gateway]
```

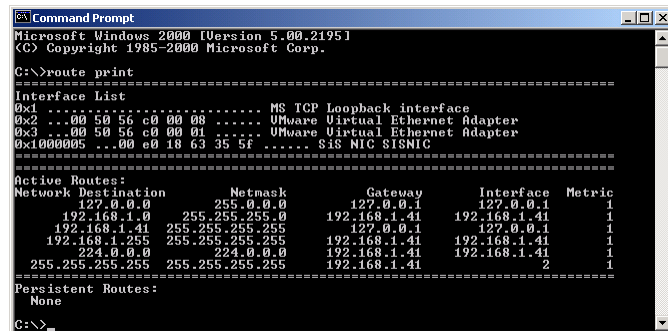
The following explains ROUTE parameters:

Table 2-5

Parameter	Function
-f	Clears the routing tables of all gateway entries. If this is used in conjunction with one of the commands, the tables are cleared prior to running the command.
Command	Specifies one of 4 commands.
PRINT	Prints/views a route.
ADD	Adds a route.
DELETE	Deletes a route.
CHANGE	Modifies an existing route.
Destination	Specifies the host to which the command will be sent.
MASK	If the MASK keyword is present, the next parameter is interpreted as the netmask parameter.
Netmask	If provided, specifies a subnet mask value to be associated with this route entry. If not specified, it defaults to 255.255.255.255.
Gateway	Specifies the gateway.

To view your routing tables, enter **ROUTE PRINT**. A screen similar to the following appears:

Figure 2-10



```

Microsoft Windows [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.

C:\>route print
=====
Interface List
0x1 ..... MS TCP Loopback interface
0x2 .. 00 50 56 c0 00 00 ..... VMware Virtual Ethernet Adapter
0x3 .. 00 50 56 c0 00 01 ..... VMware Virtual Ethernet Adapter
0x1000005 .. 00 e0 18 63 35 5f ..... SIS NIC SISNIC
=====
Active Routes:
Network Destination    Netmask          Gateway          Interface        Metric
-----
127.0.0.0                255.0.0.0        127.0.0.1        127.0.0.1         1
192.168.1.0              255.255.255.0    192.168.1.41    192.168.1.41     1
192.168.1.41            255.255.255.255  127.0.0.1        127.0.0.1         1
192.168.1.255          255.255.255.255  192.168.1.41    192.168.1.41     1
224.0.0.0                224.0.0.0        192.168.1.41    192.168.1.41     1
255.255.255.255        255.255.255.255  192.168.1.41    192.168.1.41     2
=====
Persistent Routes:
None
C:\>

```

If you want to force a workstation to rebuild the routing tables because the workstation is sending packets to the wrong router, use the **-f** parameter.

If you use the **PRINT** or **DELETE** command, wildcards can be used for the destination and gateway, or the gateway argument can be omitted.

TRACERT (Trace Route)

Use **TRACERT** to determine the possible path that a packet might take to get from one device to another (if a path exists). This tool helps define the time needed to reach the routers along the path and identify sluggish spots along that path.

When troubleshooting, **ROUTE** tells you how a packet is getting to a destination; **TRACERT** tells you where along the route the packet is failing to move forward.

TRACERT uses the Time to Live (TTL) value of the packet to obtain the routers along a path. The following shows how TRACERT uses the TTL value of a packet to obtain routers along a path:

1. Send a packet to the destination (you can use an unsupported port number, or TRACERT can send a PING packet). Set the TTL to 1.

When the local router sends a Time Exceeded in Transit message, the source has the first time stamp and first router's address.

TRACERT can repeat this process up to 3 times to get an average roundtrip time to the first router along the path.

2. Increment the TTL value in subsequent packet sets. Record all Time Exceeded in Transit messages. That is the list of routers along the path.
3. When the packet arrives at the destination device, a reply is sent (unless it was dropped at the destination). The reply provides the roundtrip time to the destination.

In some instances, a company might disable ICMP responses to TRACERT to detract hackers from exploring their network and the path to specific devices.

In the following figure, TRACERT shows one possible path to novell.com:

Figure 2-11

```
C:\> tracert novell.com
Tracing route to novell.com [192.233.80.9]
over a maximum of 30 hops:
  0  128 ms  128 ms  123 ms  pm14.san-jose.best.net [206.184.171.104]
  1  125 ms  124 ms  124 ms  e0-1.br1.snjsca02.pacific.verio.net [206.184.171.65]
  2  129 ms  126 ms  127 ms  h9-0-0.br1.atvwca.pacific.verio.net [206.86.228.117]
  3  125 ms  123 ms  123 ms  p1-0-0.cr2.atvwca.pacific.verio.net [205.149.170.66]
  4  127 ms  127 ms  123 ms  p12-0-0.br1.snjsca.pacific.verio.net [209.157.181.166]
  5  128 ms  125 ms  124 ms  f4-1-0.sjc0.verio.net [129.250.31.81]
  6  132 ms  125 ms  127 ms  107.ATM2-0-0.San-Jose9-gw.ALTER.NET [137.39.91.5]
  7  143 ms  145 ms  132 ms  118.ATM2-0.XR2.SJC1.ALTER.NET [146.188.144.142]
  8  151 ms  145 ms  147 ms  192.ATM6-0.XR2.SF01.ALTER.NET [146.188.147.97]
  9  148 ms  147 ms  135 ms  186.ATM10-0-0.CR1.SF01.ALTER.NET [146.188.148.141]
 10  223 ms  209 ms  226 ms  133.Hss15-0-0.GW1.SLT1.ALTER.NET [137.39.68.10]
 11  302 ms  226 ms  221 ms  novell-gw.customer.alter.net [157.130.162.78]
 12  246 ms  233 ms  224 ms  novell.com [192.233.80.9]
Trace complete.
```


In the following TRACERT example, the last 2 replies came from the same router (207.46.129.5). This figure also shows that TRACERT was unable to get a path to microsoft.com.

Figure 2-12

```
C:> tracert microsoft.com
Tracing route to microsoft.com [207.46.130.149]
over a maximum of 30 hops:
  0  144 ms  126 ms  123 ms  pm4.san-jose.best.net [206.184.171.104]
  1  129 ms  123 ms  124 ms  e0-1.br1.sjysca02.pacific.verio.net [206.184.171.65]
  2  130 ms  126 ms  126 ms  h9-0-0.br1.mtwca.pacific.verio.net [206.86.228.117]
  3  126 ms  122 ms  123 ms  p9-0-0.crl.mtwca.pacific.verio.net [209.157.62.201]
  4  126 ms  125 ms  135 ms  pl2-0-0.br1.plalca.pacific.verio.net [209.157.181.162]
  5  131 ms  126 ms  125 ms  g5-2-0.pao6.verio.net [129.250.15.1]
  6  146 ms  145 ms  144 ms  pao6.sea3.verio.net [129.250.3.90]
  7  155 ms  145 ms  141 ms  fe4-0-0.wes-br1.nw.verio.net [129.250.31.116]
  8  154 ms  147 ms  147 ms  microsoft-gw.nw.verio.net [204.203.0.165]
  9  151 ms  149 ms  145 ms  icpmscomc7503-a0-00-1.cp.msft.net [207.46.129.5]
 10  147 ms  147 ms  151 ms  icpmscomc7503-a0-00-1.cp.msft.net [207.46.129.5]
 11  *      *      *      Request timed out.
 12  *      *      *      Request timed out.
 13  *      *      *      Request timed out.
 14  *      *      *      Request timed out.
 15  *      *      *      Request timed out.
 16  *      *      *      Request timed out.
 17  *      *      *      Request timed out.
 18  ^C
```

The syntax of the TRACERT command is

```
TRACERT [-d] [-h maximum_hops] [-j host-list] [-w
timeout] target_name/target_ip_address
```

The parameters are explained in the following table:

Table 2-6

Parameter	Function
-d	Do not resolve addresses to hostnames.
-h maximum_hops	Sets the maximum number of hops to search for the target.
-j host-list	Sets the loose source route along host-list.
-w timeout	Sets the wait time (in milliseconds) for each reply.


TRACERT responses are not always accurate. If a network supports load balancing (as the Internet does), separate TRACERT tests can yield different paths.

The paths that are resolved might not be correct because a connection between devices in a path is only assumed.

Consider the example in the following figure. As TRACERT determines the next hop along the path, the load balancing router sends the separate path tests in different directions.

The path appears to be RA-RB-RF-RD-RH (an impossible path).

Figure 2-13 (slide)



NSLOOKUP

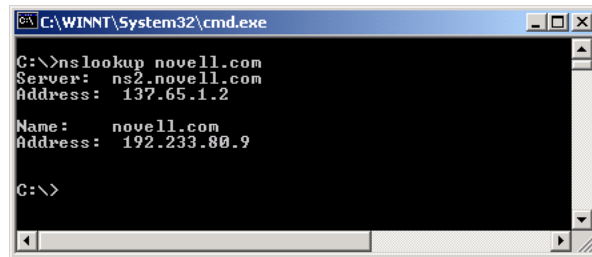
Use NSLOOKUP to query DNS name servers. NSLOOKUP is available only if TCP/IP is installed.

To use NSLOOKUP do the following:

1. From the Windows Start menu, select **Run**.
2. Enter **CMD** to access the command line.
3. Enter **NSLOOKUP** and the *IP address* or *domain name* of the server you are requesting.

The following is an example of using NSLOOKUP to access the IP address for novell.com:

Figure 2-14



```
C:\WINNT\System32\cmd.exe
C:\>nslookup novell.com
Server: ns2.novell.com
Address: 137.65.1.2

Name: novell.com
Address: 192.233.80.9

C:\>
```

ARP

You can use ARP to do the following:

- View the local device's ARP cache
- Force an ARP broadcast in an attempt to resolve an IP-to-Ethernet address
- Validate a client's local ARP cache
- Determine how entries are being acquired
- Force an entry into the tables or delete an incorrect entry

The syntax for the ARP command is

```
ARP - s inet_addr eth_addr [if_addr]
```

```
ARP - d inet_addr [if_addr]
```

```
ARP -a [inet_addr] [-N if_addr]
```

The following explains the functions of the ARP parameters:

Table 2-7

Parameter	Function
-a	Shows current ARP entries by interrogating the current protocol data. If inet_addr is specified, the IP and physical address for only the specified computer appear. If more than one network interface uses ARP, entries for each ARP table appear.
-g	Same as -a.
inet_addr	Specifies an Internet address.
-N if_addr	Shows the ARP entries for the network interface specified by if_addr.
-d	Deletes the host specified by inet_addr.
-s	Adds the host and associates the Internet address inet_addr with the physical address eth_addr. The physical address is given as 6 hexadecimal bytes separated by hyphens. The entry is permanent.
eth_addr	Specifies a physical address.
If_addr	If present, this specifies the Internet address of the interface whose address translation table should be modified. If not present, the first applicable interface is used.

NETSTAT

Use NETSTAT to view details on the current protocol operations for TCP/IP connections.

For example, if you established an FTP connection to the server and then left, you could use NETSTAT to determine whether your connection is still valid when you return.

NETSTAT shows statistics for TCP, UDP, ICMP, and IP.

The following shows the output from the NETSTAT command on a device that is connecting to an FTP server (ftp.novell.com):

Figure 2-15

```
C:> netstat
Active Connections

Proto Local Address           Foreign Address         State
TCP   chappell:1729          www.compaq.com:80      ESTABLISHED
TCP   chappell:1730          www.compaq.com:80      ESTABLISHED
TCP   chappell:1731          www.compaq.com:80      TIME_WAIT
TCP   chappell:1732          www.compaq.com:80      ESTABLISHED
TCP   chappell:1735          netscantools.com:80    CLOSE_WAIT
TCP   chappell:1736          netscantools.com:80    CLOSE_WAIT
TCP   chappell:1775          ftp.novell.com:ftp     ESTABLISHED
TCP   chappell:1776          ftp.novell.com:1909    SYN_SENT
```

The FTP connection has been established, and a secondary connection is in the process of being established.

The syntax for the NETSTAT utility is

```
NETSTAT [-a] [-e] [-n] [-s] [-p proto] [-r] [interval]
```

The following are the NETSTAT parameters:

Table 2-8

Parameter	Function
-a	Shows all connections and listening ports.
-e	Shows Ethernet statistics. This can be combined with the -s option.
-n	Shows addresses and port numbers in numerical form.
-p <i>proto</i>	Shows connections for the protocol specified by <i>proto</i> ; <i>proto</i> can be TCP or UDP. If used with the -s option to show per-protocol statistics, <i>proto</i> can be TCP, UDP, or IP.
-r	Shows the routing table.
-s	Shows per-protocol statistics. By default, statistics are shown for TCP, UDP, and IP. The -p option can be used to specify a subset of the default.
interval	Reshows selected statistics, pausing <i>interval</i> seconds between each screen. Press Ctrl + C to stop reshown statistics. If <i>interval</i> is omitted, NETSTAT prints the current configuration information once.

The NETSTAT command can be used with the -r parameter to show the routing tables. If you want to watch the NETSTAT information being dynamically updated as connections are established and terminated, enter an *interval* (in seconds) following the parameters.

For example, if you enter NETSTAT 5, the statistics are updated every 5 seconds. Press Ctrl + C to stop the display.

The following explains the different connection states:

Table 2-9

State	Function
LISTEN	Waiting for a connection request from any remote TCP device and port.
SYN-SENT	Sent connection request (SYN) as first packet of TCP handshake; waiting for a matching connection request (SYN+ACK).
SYN-RECEIVED	Sent and received connection request (SYN and SYN+ACK); waiting for a confirming connection request acknowledgment (ACK) to complete TCP handshake.
ESTABLISHED	Open connection; data received can be delivered to the user. This is the normal state for the data transfer phase of the connection.
FIN-WAIT-1 and FIN-WAIT-2	Waiting for acknowledgment (ACK) of connection termination request (FIN) from the remote TCP, or waiting for a connection termination request (FIN) from the remote TCP.
CLOSE-WAIT	Waiting for a connection termination request (FIN) from the local user.
CLOSING	Waiting for a connection termination request acknowledgment (FIN+ACK) from the remote TCP.
LAST-ACK/TIME WAIT	Waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request.
CLOSED	No connection state.

TCP/IP Troubleshooting Example

Many times, the best way to understand how to troubleshoot a situation is through example. The following is an example of troubleshooting TCP/IP protocol errors:

“In September 1999, I was going to press with an article that focused on ICMP error codes and the use of ICMP for troubleshooting. In an article sidebar, I included a list of currently assigned ICMP type numbers.

“At the last minute, the publisher asked me to double-check to see if the list was up-to-date as of that day. I logged in to the Internet and accessed www.iana.org.

“I waited, and waited, and waited. I found that I could not connect to IANA to validate the list. This was a perfect time to test some troubleshooting tools.”

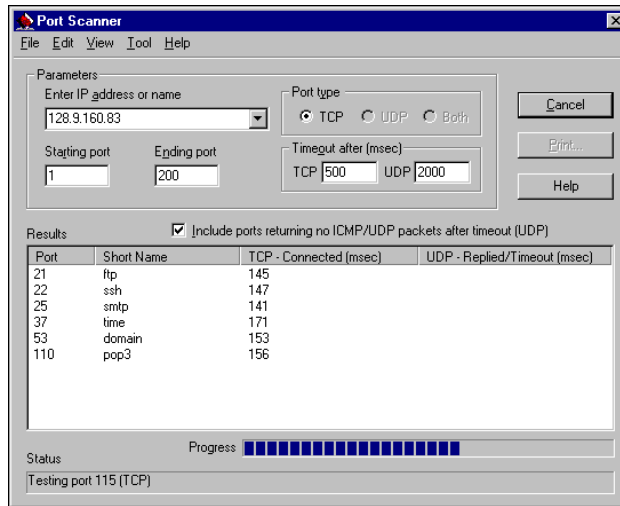
Here are the steps the author used to isolate the problem:

1. Ping www.iana.org. Results: Successful.
2. TRACERT to www.iana.org. Results: Successful.
3. Use NSLOOKUP to obtain the IP address for www.iana.org. Results: 128.9.160.83.
4. Portscan 128.9.160.83. Result: HTTP port 80 is not active on the host.

Figure 2-16 shows the result of the port scan that day. As you can see, the host is running.

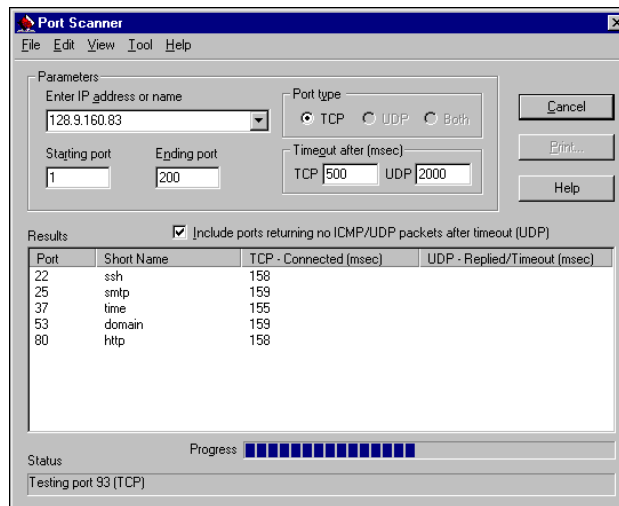
The HTTP port (used for the web server) is not functional. The Port Scanner indicates that the HTTP daemon is not loaded.

Figure 2-16



The following shows what the port scan looks like on a normal day when the web server is running:

Figure 2-17



You might need to use a variety of tools to isolate TCP/IP problems. In this case, the troubleshooter used PING, Trace Route, NSLOOKUP, and a port scan.



10 minutes

Exercise 2-2 Test Your Network

To identify your network configurations, do the following:

- From the server, enter **VERSION** to verify that your migration was successful; then record the following:
 - NetWare version:
 - Support pack version:
 - eDirectory version:
 - NDS version:

-
2. What is the difference between the eDirectory version and the NDS version?

 3. Enter **CONFIG** and record the following:
 - The IPX address:
 - The network interface board name:
 -
 -
 -
 - The frame type:
 -
 -
 -
 - The LAN protocols:
 -
 -
 -
 - IP address
 - Subnet mask
 4. Where would you typically record this information and why?

 5. Start TCPCON. What indication is there that packet forwarding is enabled?

6. If packet forwarding is not enabled, can you access novell.com?

7. What command should you use to find your workstation's IP address?

8. Using the command you specified in the previous question, record the IP address of your workstation.

9. Ping the IP address of the workstation. Was the ping test successful? (If not, notify the instructor.)

10. Using the IP address you recorded in question 3, ping the IP address for your server. Was the ping test successful? (If not, notify the instructor.)

11. What command should you use to identify the path packets might take to go from your workstation to novell.com?

12. Using the command you specified in the previous question, record the path packets would take to go from your workstation to novell.com.

13. From your server, use the IPTRACE command to identify the path from your server to the instructor server.

(End of Exercise)

Protocol Analyzers

Installing and demonstrating a protocol analyzer (such as LANalyzer) is much more effective than discussing it.

Some students might not be familiar with the term *protocol analyzer*, but they might be familiar with *sniffer*.

Protocol analyzers (also called network analyzers or sniffers) capture packets on the cabling system and show conversations and individual packets in a readable format.

Protocol analysis let you listen in on network communications to determine the health of that network.

To understand protocol analyzers, you need to know the following:

- [Analysis Session Procedure](#)
- [Use of Protocol Analysis](#)



For more on using and purchasing a protocol analyzer, see *Appendix C: Protocol Analyzers*.

Analysis Session Procedure

During a typical protocol analysis session, you perform the following tasks:


- Step 1: Access the network.
- Step 2: Capture the traffic.
- Step 3: View the captured traffic.
- Step 4: Filter out and view just the needed traffic.
- Step 5: Document your findings.

For example, to analyze the broadcast traffic from one of your servers on the network, you would do the following:

1. Connect the analyzer to the same hub the server is connected to.
2. Capture all traffic to and from the server.
3. View the traffic to identify unusual packets (such as 100 broadcasts in succession).
4. Filter out and view just the broadcast traffic from the server.
5. Document the cause of the broadcasts.

The following shows the basic flow of data through a protocol analyzer:

Figure 2-18 (slide)



You can use the display filter to view the packets you are interested in. All the captured packets are still in the trace buffer; you have just chosen to view a portion of them.

Captured packets are sent through a capture filter. You can set the capture filter to grab all packets being sent to and from the server's hardware or software address. You can then set a display filter to filter out everything but broadcast traffic.

The resulting filtered buffer would show only broadcast traffic from the server. This is one example of how to use protocol analysis to identify a possible problem.

Use of Protocol Analysis

Protocol analysis is primarily used for

- [Troubleshooting](#)
- [Optimization](#)
- [Planning and Testing](#)

Troubleshooting

Protocol analysis can help you identify problems quickly.

For example, when one device can't communicate with another on the network, the protocol analyzer can pinpoint where the problem is.

Likewise, consider a client that cannot communicate with a specific server. You can use a protocol analyzer to answer the following:

- Can the client communicate on the cabling system or is something wrong with the media?
- Can the server communicate on the cabling system or is something wrong with the media?
- Did the service discovery process work properly? Could the client find the server?
- Could the client locate the route or path to the server? Is the route or path available and within reach?
- Did the client properly authenticate to the server environment?
- Did the client make a proper request for services?
- Did the server reply to the client's request for service?
- If the server did reply to the client, was a failure or denial of service indicated in the reply?
- Are the client and server communicating using the same frame type?

Optimization

Every network can benefit from a quick review of network performance. You can use a protocol analyzer to analyze your network's performance regardless of the protocol or media access type. For example, optimization lets you

- Identify excess ICMP redirection messages and reconfigure devices to use a more appropriate default gateway.
- Identify and remove the cause of excessive broadcasts on the network, reducing the bandwidth overhead and processor drag caused by this traffic.
- Identify the cause of excessive failure replies from a server. Clients can be reconfigured to reduce these replies.
- Tune up your router configurations to reduce the routing protocol overhead caused by excessive periodic broadcasts or multicasts.
- Remove unsuccessful discovery processes that are running on the network (for example, a vendor that searches for other same vendor systems).

Planning and Testing

When you test an application or plan for growth, you can determine how much bandwidth a single user requires to run a particular application on the network.

Multiply this number by the number of users who will eventually use the application to determine if you have enough bandwidth on the segment to support that application.

TCP/IP Toolkits

If you are responsible for troubleshooting your network, you should have a comprehensive set of troubleshooting tools.

Although several tools (such as PING and TRACERT) come with the OS, these tools do not provide depth in the area of testing and information collection.

The following are examples of popular TCP/IP toolkits:

- **INetTools.** iNetTools from WildPackets is a TCP/IP troubleshooting utility.

In addition to standard features such as PING and TRACERT, INetTools provides the following:

Table 2-10

Feature	Function
PINGscan	PINGs a range of addresses
Finger	Gets user information on an email address
Whois	Queries a WHOIS server for Internet directory information
Throughput	Tests the throughput of FTP and HTTP downloads
Network statistics	Runs the NETSTAT -r -s command
ARP Cache Content (Windows Only)	Shows any locally-cached results of ARP requests
Internet port descriptions	Lists Internet port numbers and descriptions, downloaded from the IANA (Internet Assigned Numbers Authority) site



You can download a demonstration version of the toolkit at www.wildpackets.com. Click the **Demos/Buy** button to access the download screen.

- **NPS NetScanTools Pro 2002.** NetScanTools™ Pro 2002 was developed by Northwest Performance Software as a comprehensive protocol analyzer.

In addition to standard features such as PING and TRACERT, INetTools provides the following:

Table 2-11

Feature	Function
DHCP Test	The DHCP feature lets you discover DHCP servers on your local subnet. This is useful to locate servers that are not supposed to be on your network as well as to check the output of known servers.
Echo	Echo is a TCP or UDP service that echoes back all characters received on the port that it is listening on.
Finger	The Finger feature provides a client interface to a finger server. A finger server is usually located on a remote computer.
IP Packet View	The IP Packet Viewer feature launches a separate application that captures IP packets, stores them in a dBase IV compatible database, and lets you view the contents of the packets as hex bytes. This application works only on Windows 2000 or XP.
NetTopography	The NetTopography feature provides a map of routers and gateways along the various routes to a user-defined list of target hosts.
NetScanner	NetScanner is one of the most powerful and useful features in NetScanTools Pro. NetScanner sweeps a sequential IP address range and pings every IP address in that range of IP addresses.

Table 2-11 (continued)

Feature	Function
Port Probe	Port Probe is an active scanning feature to determine which ports on a target computer are active and being used by services or daemons.
Whois	Whois is a client utility that acts as a client interface to a remote server database of domain or IP address registries.
WinSock Info	This gives basic information about the active Windows Sockets (Winsock) software interface layer that is running on your computer system.



To take full advantage of the tools in NetScanTools Pro 2002, Windows 2000 is recommend. For supported platforms, see www.netscantools.com.

IP Addressing Calculators

Several tools can help you calculate subnet addressing for TCP/IP. The IP Subnet Calculator from WildPackets is a freeware calculator that offers a range of functions.



A copy of the IP Subnet Calculator can be downloaded at <http://www.wildpackets.com>.

You use subnet calculators to configure a new network addressing design and to check an existing network.

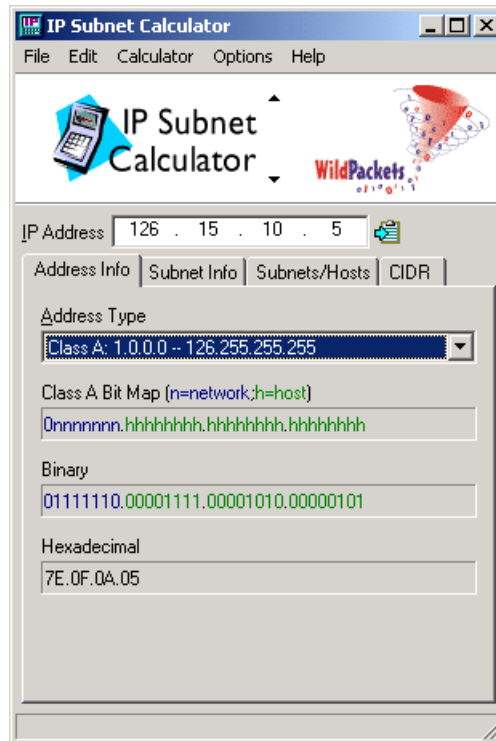
The IP Subnet Calculator computes subnet information based on the IP address and class and the number of subnet bits used.

You can use the following tabs to compute subnet information:

- **Address Info.** This tab provides information about address fields given an IP address class. Classes A, B, C, and D (multicast) addresses are supported. Using color coding, you can easily detect the role of individual bits in the address class.

The following shows the Address Info tab:

Figure 2-19



- **Subnet Info.** This tab contains user-configurable parameters to set up a subnet policy. Additional configurations on this tab include the following:
 - **Allow 1 Subnet Bit.** Use to switch between allowing a minimum of 1 or 2 subnet bits.
 - **Use Inverse Mask.** Use to set the way the Subnet Mask combo box shows masks.
 - **Subnet Bits.** Use to determine the number of subnet bits to use when creating the subnet mask using the network, subnet, and host portion of the IP address.

The following options can be used to determine the number of subnet bits to use when calculating the subnet mask:

- Mask Bits
- Subnet Mask
- Max Subnets
- Max Hosts Per Subnet

Subnet Host Address Range lists all possible addresses that hosts can own on the same subnet. The subnet ID and subnet broadcast addresses of the subnet that the current IP address belongs to appear as well.

- **Subnets/Hosts.** Use to see a table of the subnet number, the subnet ID address, host address range, and the subnet broadcast address for each possible subnet in the current IP address scheme.
- **CIDR (Classless InterDomain Routing).** Use to specify the number of bits to use in the supernet mask.

Objective 4 Identify Additional Network Troubleshooting Resources

The following are additional resources you can use when troubleshooting a Novell network:

- [Novell Web Site Resources](#)
- [Shareware and Freeware](#)

Novell Web Site Resources

The Novell web site at <http://www.novell.com> has links to troubleshooting tips and tools. The following are a few of the available resources:

- **Solutions.** You can access white papers that describe Novell business solutions that enhance and improve your existing systems, allowing you to meet new business requirements more quickly, and make sure every dollar you spend on technology translates directly into lower costs and higher revenues.
- **Products.** You can access products listed from A–Z. From this page, you can link to any product for information, FAQs, highlights, what’s new, and Novell AppNotes.
- **Training.** You can access information about Novell Education products, services, certifications, and events. You can also access Novell AppNotes.
- **Support.** You can access online support resources, including the following:
 - **Novell documentation.** You can access all product documentation online. The product documentation includes troubleshooting sections and error code information.

Access www.novell.com and show the links described. Show students your favorite links and have them explore the links available.

- ❑ **Knowledgebase.** Knowledgebase is a convenient way to search Novell Technical Information Documents (TIDs) and product manuals. Many TIDs contain links to downloadable files, patches, and drivers.

Novell Technical Support representatives write TIDs as the resolution to support calls. Although the resolution was appropriate for the customer that made that call, it might not be as appropriate for your situation.

There might also be several TIDs that resolve the same problem in different ways. You must determine which TID is most appropriate for your problem.

The following shows the Knowledgebase search screen:

Figure 2-20

Knowledgebase

Select product categories (or none to search all).

Connectivity Products - [product list](#) Excelerator - [product list](#)

Groupware - [product list](#) Management Products - [product list](#)

NetWare - [product list](#) Novell BorderManager Services - [product list](#)

Novell Directory Services - [product list](#) Other - [product list](#)

Web Services - [product list](#)

Select document sets.

Technical Information (TIDs) Manuals

Enter a word, phrase, or Technical Information Document number.

[Help](#)

or or

➤ **Natural Language Search Engine**

Try Novell's "Natural Language" search engine. You can type in a question or statement in "natural language", and the search engine will guide you to a solution in a faster, more accurate manner. [Click Here](#)

- **Cool Solutions web site.** This web site is dedicated to helping customers get the most out of Novell products. You'll find hundreds of tips, articles, Q&As, and free tools to help you.

Many articles posted are submitted by readers who have shared from their learning experience by documenting the problem and the solution. It might not fit your specific configuration, but could give you a starting point.

Shareware and Freeware

Many shareware and freeware utilities are available on the Internet and through your browser.



To access tools, demos, and evaluation copies of commercial tools that work with or enhance the functionality of many Novell products, access the Novell Cool Solutions site at <http://www.novell.com/cool solutions/tools>.

The following are examples of NSLOOKUP tools available through your browser:

- <http://www.bankes.com/nslookup.htm>
- <http://www.jimprice.com/jim-soft.htm>

Summary

The following is a summary of the objectives in this section.

Table 2-12

Objective	Summary
<p>1. Upgrade Novell Network Management Tools</p>	<p>When you upgrade NetWare or any of its components by installing a new version or support pack, you should also upgrade the Novell network administration management tools.</p> <p>The following are key Novell network management tools you should upgrade:</p> <ul style="list-style-type: none"> ■ ConsoleOne ■ iMonitor ■ Novell Remote Manager ■ iManager
<p>2. Identify the Troubleshooting Features of Novell NetWork Management Tools</p>	<p>You can use a number of Novell network management tools to administer your network, monitor processes, and troubleshoot problems.</p> <p>These tools include the following:</p> <ul style="list-style-type: none"> ■ ConsoleOne. With the release of ConsoleOne 1.3 for Windows, you can generate predefined reports to help you troubleshoot network problems. ■ Novell iMonitor. iMonitor provides cross-platform monitoring and diagnostic capability for all servers in your eDirectory tree. ■ NetWare Remote Manager. Novell Remote Manager troubleshooting tools include Health Monitor, Profile/debug server, and Report/log files. ■ Novell iManager. Many items listed as tools also contain troubleshooting capabilities such as links to iMonitor to perform repair tasks and rights management.

Table 2-12 (continued)

Objective	Summary
3. Identify the Purpose and Function of IP/IPX Troubleshooting Tools	In this objective you learned about the following: <ul style="list-style-type: none"> ■ NetWare IP/IPX Troubleshooting Tools ■ Client IP Troubleshooting Tools ■ Protocol Analyzers ■ TCP/IP Toolkits ■ IP Addressing Calculators Most of the tools in this section either ship with the OS being discussed or are available for download from the Internet.
4. Identify Additional Network Troubleshooting Resources	The following additional resources are available to you when troubleshooting a Novell network: <ul style="list-style-type: none"> ■ Novell Web Site Resources ■ Shareware and Freeware

Exercise Answers

Following are the exercise answers.

Exercise 2-2. Test Your Network

1. From the server, enter `VERSION` to verify that your migration was successful; then record the following:
 - [NetWare version:](#). Novell NetWare 6
 - [Support pack version:](#). Support Pack Revision 02
 - [eDirectory version:](#). eDirectory 8.7.0
 - [NDS version:](#). 10410.98

2. What is the difference between the eDirectory version and the NDS version?

They represent the same information. The eDirectory version is used by Novell Customer Support to help troubleshoot problems, and the NDS version is used by development to track releases.

3. Enter CONFIG and record the following:

The answers will vary.

4. Where would you typically record this information and why? Record this information in the server log file. The log file can help you to identify the following:

- Hardware settings
- Hardware upgrades
- Software setting
- Software upgrades
- System problems and solutions

5. Start TCPCON. What indication is there that packet forwarding is enabled?

IP Forwarding has a numeric value of 0 (zero) or higher. If it not enabled, it would stay disabled.

6. If packet forwarding is not enabled, can you access novell.com?

No

7. What command should you use to find your workstation's IP address?

IPCONFIG

8. Using the command you specified in the previous question, record the IP address of your workstation.

Answers vary, depending on the classroom configuration.

9. Ping the IP address of the workstation. Was the ping test successful? (If not, notify the instructor.)

The ping test should be successful.

10. Using the IP address you recorded in question 3, ping the IP address for your server. Was the ping test successful? (If not, notify the instructor.)

The ping test should be successful.

11. What command should you use to identify the path packets might take to go from your workstation to novell.com?

TRACERT

12. Using the command you specified in the previous question, record the path packets would take to go from your workstation to novell.com.

Results will vary.

13. From your server, use the IPTRACE command to identify the path from your server to the instructor server.

Results will vary.

SECTION 3 Troubleshoot and Resolve NetWare Server Issues

Duration: 3 hours

Setup: The files needed for this section should have been migrated to volume DATA from the NetWare 4.11 server.

Preparation: Perform Exercise 3-4 prior to class to verify compatibility with the network board used in the classroom.

In this section you learn how to troubleshoot and resolve problems that occur on the server.

Objectives

1. Identify Server Hardware and Operating System Components
2. Troubleshoot and Resolve NetWare Server Issues
3. Troubleshoot and Resolve Critical Server Abends
4. Troubleshoot and Resolve Server Communication Issues

Introduction

Now that you have upgraded servers and learned about Novell and IP management and troubleshooting tools, you are ready to learn about troubleshooting problems that occur on servers, resolve critical server abends, and resolve server communication issues.

These topics are covered in this section.

Objective 1 Identify Server Hardware and Operating System Components

The topics in this objective provide students the terms and concepts needed to troubleshoot server issues.

In this objective, you learn how to identify server hardware and NetWare OS components that may be causing problems in your network.

Many things can negatively affect a server's performance. These include environmental problems (natural disasters, electromagnetic fields, and so on), operator error, design and reconfiguration issues.

Ultimately, the problems that occur on a server are either hardware or software related.

The following provide a foundation to help you troubleshoot hardware and software related issues:

- Identify Server Hardware
- Identify Operating System Components

Identify Server Hardware

You need to become familiar with the following server components to effectively troubleshoot problems as they occur:

- Bus Types
- Mass Storage
- Processor Capacity
- Memory Capacity
- Scalability
- Failure Recovery



The Server+ certification from CompTIA covers server hardware in detail. See <http://www.comptia.org/certification/serverplus/index.htm>.

Bus Types

Because the bus connects devices on a computer system, it has a direct impact on performance. If there is not enough bandwidth for I/O traffic, the bus can become a bottleneck.

An analogy is a freeway with not enough lanes to handle rush hour traffic. Traffic quickly begins to slow and even stop at some points.

Similarly, a server must have a wide enough bus to meet high usage peaks.

Buses on servers usually take advantage of

- **Bus mastering.** This lets data transfer between devices on the bus without going through the CPU. Performance is increased and no additional load is placed on the CPU.

For example, the disk controller can write to the hard drive without involving the CPU.

- **Hot Plug/Hot Swap.** This lets you remove an expansion board or device while the system is running and the OS automatically recognizes the change.

Hot Swap drives are very helpful when troubleshooting or replacing failed hardware because hardware can be replaced without having to make the system unavailable to users.


Peripheral Component Interconnect (PCI), developed by Intel Corporation, is the de facto industry standard for system buses.

PCI is a 64-bit bus and has clock speeds of 33 or 66 MHz. At 64 bits and 66 MHz, the throughput rate is 528 MBps.

Not all PCI boards are 64-bit and 66 MHz, so it is important to check the specifications before using it on a server.

The following shows a server's system board that has a variety of PCI slots:

Figure 3-1 (slide)



Some legacy bus types include International Standard Architecture (ISA), Micro Channel Architecture (MCA), and Extended Industry Standard Architecture (EISA).

Mass Storage

Servers usually have a large amount of disk storage to meet the storage needs of multiple users. For example, an iFolder server allocates 200 MB per user by default.

The type of disk channel used is also important. You should be familiar with the following:

- Fibre Channel
- Small Computer System Interface (SCSI)
- Integrated Drive Electronics/Advanced Technology Attachment (IDE/ATA)

You usually find Fibre Channel or SCSI disk channels in a server, and IDE/ATA in a desktop.

Fibre Channel is an emerging standard that promises the ability to connect storage systems up to 10 kilometers (6 miles) apart. Fibre Channel is compatible with SCSI and targeted to high-performance storage systems.

On single drive systems, IDE/ATA is less expensive and performs well. In a multidrive environment, Fibre Channel and SCSI are more scalable.

IDE/ATA channels are limited to 2 devices, but SCSI can handle 8 devices with Narrow SCSI, 16 devices with Wide SCSI, or 32 devices with Very Wide SCSI.

SCSI drives also support *command queuing* (up to 256 commands per device); IDE/ATA drives do not. The performance gains are realized in a multidrive server environment.

You discuss RAID later in the course.

Servers typically support either software or hardware RAID. This is discussed in [“Troubleshoot Software RAID and Mirroring in NSS” on 6-19](#).

Students might want to know how many processors are commonly used. Before class, you might want to review the product information and processor capacity currently offered by server vendors.

Processor Capacity

The processing demands placed on a server can be great. As a result, it is often necessary to have multiple processors running on a server to meet the demand.

NetWare 6 supports up to 32 processors, with a minimum requirement of a Pentium II or AMD™ K7 processor.

The Xeon, Pentium 4, and Pentium III processors are commonly used in servers.

During installation, NetWare detects multiple processors by reading the multiprocessor (MP) configuration table in BIOS and then determining which of the available NetWare Platform Support Modules (PSMs) matches the MP hardware platform.

(A PSM is a loadable hardware abstraction layer for processor and interrupt support. A PSM is specific to a particular hardware platform.)

You can then load the PSM or run NetWare on Processor 0 only. The installation program will modify the STARTUP.NCF file to load the PSM whenever the server is started.

Novell provides MPS14.PSM, which supports any hardware platform that complies with the Intel Multiprocessor Specification 1.1 and 1.4.

Compaq and other vendors also provide a PSM for their system requirements.



For more on the Intel MPS specification, see <http://developer.intel.com/design/intarch/MANUALS/242016.htm>.

Memory Capacity

A server's performance is directly related to the amount of memory.

Each application or service running on the server uses memory from the available memory pool. When the amount of available memory becomes low or runs out, the server is negatively impacted.

NetWare 6 requires a minimum of 256 MB of RAM (512 MB recommended) and can address up to 64 GB. Up to 4 GB can be allocated to cache memory. Above 4 GB is allocated to virtual memory.

To maintain performance, memory should be added for each additional service or application running on the server.

Scalability

After a server is in place, it is difficult to replace the entire system without an interruption of service. Therefore, as a company's server needs change and grow, the server must scale to meet the demand.

You should be able to add the following to the server:

- Processors
- Memory
- Disk storage
- Clustering

Failure Recovery

Servers typically have redundant (and in some cases, hot swappable) hardware for the following:

- Hard drives
- Power supplies
- Processors
- Cooling fans
- Network boards
- Uninterruptible power supplies (UPS) / battery backup
- Memory



The *Yes, Tested and Approved* program provides a list of hardware compatible with Novell software, such as NetWare 6. See <http://developer.novell.com/nss/category.html#hardware>.

Identify Operating System Components

Drivers, applications, and services can cause a server to malfunction or stop functioning when they do not interact properly with the OS kernel.

The following terms and concepts will help you while troubleshooting a software issue on the server.

- [NetWare Load Order](#)
- [Kernel](#)
- [Threads](#)
- [The Run Queue](#)
- [Multithreading](#)

- [Multitasking](#)
- [Multiprocessing](#)
- [Processor Load Balancing](#)
- [Pre-Emption](#)

NetWare Load Order

It is helpful to understand the NetWare load order when troubleshooting problems that occur during server startup.

SERVER.EXE loads first from C:\NWSERVER on the DOS partition.

SERVER.EXE contains files, such as LOADER.EXE, SERVER.NLM, and other NLMs (that are bound in), such as

- PMLODR.NLM
- PVER500.NLM
- XLDR.NLM

The first file loaded is LOADER.EXE, which sets up the initial hardware interfaces (screen, keyboard, memory management, and interrupt handling).

SERVER.NLM is loaded next, which contains most of the server routines.

The server has a load template that defines what NLMs will be loaded at each *stage* of the boot process.

Modules are then loaded in the following order:

- LOADSTAGE 0
- STARTUP.NCF
- LOADSTAGE 1

- LOADSTAGE 2
- LOADSTAGE 3
- LOADSTAGE 4
- AUTOEXEC.NCF
- LOADSTAGE 5

Demonstrate the LIST STAGE command.

You can view the modules that are loaded in each stage by entering LIST STAGE at the server console prompt. This can be helpful when troubleshooting problems that occur when the server starts.

Module Load Color

Demonstrate module colors by typing M at the console of DA1 and discuss the blue, red, and white color codes and load locations.

After the server boots, NLMs can be viewed with the MODULES or M command. You will notice they are color coded, which is useful when troubleshooting.

The following provides a description of each color:

Table 3-1

Color	Description
Cyan (light blue)	Loaded NLM was bound into SERVER.EXE and loaded from there. For example, PVER500.NLM, is bound into SERVER.EXE.
Red	Loaded NLM was bound into SERVER.EXE; however, the NLM was also found in the C:\NWSERVER directory and loaded from there.
White	Loaded from the Novell configuration file, from any NCF files (such as, the AUTOEXEC.NCF or STARTUP.NCF files), or from the server console prompt.
Purple	Auto-loaded by another module.

In NetWare 5.x, many system NLMs were bound into SERVER.EXE and it was important to know if the loaded version was from the bound-in list or from somewhere else.

In NetWare 6, fewer NLMs are bound into SERVER.EXE.

Module Load Status

During the NetWare 6 boot process, the module will load and indicate the success of the load.

Demonstrate the Logger Screen.

The Logger Screen can be used to review the load status of modules.

Students might want to know what PUB EXISTS indicates when they view this screen. It is short for public symbols.

Pressing F1 on the Logger Screen shows the navigation keys. F2 saves the output as a text file to C:\NWSERVER\LOGGER.TXT.

This message is informational (yellow) and indicates that the public symbols required for that NLM have already been loaded.

Kernel

The kernel is the core of a network OS. It provides fundamental operating OS, such as handling interrupts and the I/O system.

In NetWare 4, there were 2 kernels, one for uniprocessing and one for multiprocessing. Starting with NetWare 5, the 2 kernels were integrated into one.

NetWare 6 is based on an integrated kernel. The multiprocessing kernel (MPK) is completely multithreaded, supports pre-emption, and runs on both multiprocessor and uniprocessor systems.

Threads

The NetWare 6 kernel manages multiple processes, called *threads*, and schedules processor resources to handle multiple threads concurrently.

A software thread is made up of 2 pieces:

- The first piece defines what the computer must do to execute a software program.
- The second piece performs the action.

A software thread works for an application, meaning that the application allocates threads just like it allocates memory.


When the application finishes processing the thread, the thread is returned to the system kernel to be called the next time it is needed.

Because a software thread is allocated CPU time, a processor can only execute one thread at a time.

The Run Queue

The kernel maintains a data structure called the *run queue*, which contains threads that are ready to be executed by the processor, as in the following:

Figure 3-2 (slide)



In a uniprocessor system, there is only one run queue from which the processor can pick up threads for execution.


In a multiprocessing system, the NetWare kernel uses a per-processor run queue. A processor picks up threads for execution only from its local run queue.

Multithreading

Multithreading is the simultaneous execution of more than one thread within a process or application.

In a uniprocessor environment, multithreaded code lets several threads run *concurrently*. However, as shown in the following, only one thread can be processed at a time:

Figure 3-3 (slide)




For example, thread 1 executes code, then it yields and waits while thread 2 executes code. Thread 2 then yields to thread 1 to process another string of code. This process continues until both threads completely process.

Threads A and B are perceived as executing simultaneously because processors are very fast and actual execution time is very small.

In a multiprocessor environment, multiple threads run at the same time on different processors:

Figure 3-4 (slide)



When threads execute simultaneously on multiple processors, they are running in *parallel*. Multithreaded code allows more efficient processor use by exploiting parallelism.

For example, thread 1 can execute on processor 0 at the same exact time that thread 2 is executing on processor 1.

With NetWare 6, applications can be written to exploit the parallelism available in multiprocessor (MP) hardware and the support for parallelism in the server operating system.

Server applications, such as GroupWise®, provide performance gains and scaling by using parallelism. However, if not developed correctly, an application can cause problems on the server.

A software developer identifies tasks that can be performed concurrently, that are not dependent on being performed in a fixed sequence, and provides the mechanisms for assigning tasks to multiple threads and for appropriate synchronization to protect data shared by the threads.

Multitasking

NetWare 6 is a multitasking operating system. This means that each single-processor NetWare 6 server can run 2 or more programs at the same time.

For example, while one program is waiting for input, instructions in another program running on the same server are executed. During the milliseconds one program waits for data to be read from a disk, millions of instructions in another program can be executed.

As a result, NetWare 6 is an efficient network OS because it constantly services client needs.

Multitasking is based upon a single processor model. By adding multiple processors, you can expect a corresponding increase in server performance.

Multiprocessing

Multiprocessor-enabled programs are written in such a way that their threads can safely execute simultaneously on multiple processors.

This means that NetWare 6 goes beyond the illusion of simultaneous processing provided by multithreading on a uniprocessor system and makes simultaneous processing of multiple threads a reality.

Analogy: The scheduler performs actions similar to a police officer directing traffic.

Aside from the processors, the component that drives the functionality of multiple processor servers is called the *scheduler*. The scheduler determines how to distribute threads.

The scheduler makes this determination based upon the type of threads being processed (ensuring that they are multiprocessor safe), or upon the processors being used.

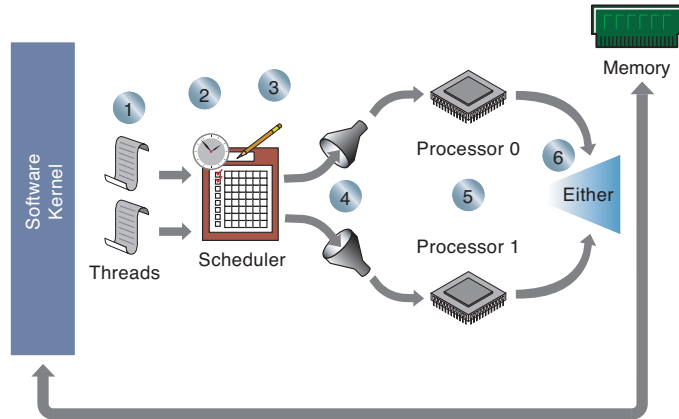
If the program or thread is not multiprocessor safe, the scheduler forces or funnels the program and its threads to be run on Processor 0. If the program is multiprocessor safe, the scheduler uses all available processors to complete the transaction.

To ensure optimal multiprocessor server performance, NetWare 6 sends threads to the processor on which they last ran. This lets the processors use their cache to process threads faster and avoid bottlenecks.

Of course, using processor cache can potentially overload one processor while underusing another.

To solve this problem, the NetWare scheduler periodically calculates utilization of each processor and determines if it is overused or underused.

Figure 3-5 (slide)



The following describes the steps in Figure 3-5:

1. One or more threads present themselves to the scheduler.
2. The scheduler determines if the transaction is multiprocessor safe.
3. The scheduler determines the workload of available processors.
4. The scheduler sends the first thread to Processor 0 and the second thread to Processor 1 until all processors are busy.
5. When possible, the threads stay on the processor they start on to avoid cache contention.



The scheduler won't use processor cache if the transaction isn't multiprocessor safe or if the scheduler is resolving load balancing.

6. The processed thread is either released into memory or returned to the software kernel for future processing.

Processor Load Balancing

To address the problem of a processor being overused or underused, NetWare uses a sophisticated load balancing algorithm.

Two important requirements of any load balancing scheme are stability (not overreacting to small load imbalances) and the ability to distribute the processing load quickly.

The NetWare scheduler handles the stability requirement by using a threshold. The threshold determines how much load imbalance is permitted in the system before the load balancing mechanism activates.

Without the threshold, threads would constantly move from one processor to another, compromising the productivity of the system.

Although the threshold is configurable through NetWare Remote Manager, it is strongly recommended that you retain the preset optimum value.

The threshold value can be viewed by selecting Set Parameters and then Multiprocessor.

Pre-Emption

NetWare allows for pre-emption of threads, within constraints. New NetWare modules can be written to be pre-emptible.

Earlier versions of NetWare implemented a non-pre-emptible round-robin (first-in, first-out) scheduling policy where threads were scheduled to run in the order that they entered the run queue.

For an application to exploit pre-emption, it must be explicitly developed to be pre-emptible. However, the Java environment utilizes pre-emption by default.

By default, if an application thread is running, it will not be pre-empted until the following conditions are met:

- The code where the thread is running must be pre-emptible. This is indicated by a flag set in the module's NLM file format. When the code is loaded into memory, the memory pages are flagged as pre-emptible.
- The thread cannot be in a critical section of the code.
- The thread has run long enough to qualify for pre-emption. The scheduler checks the elapsed time with every tick.

Support for pre-emption provides

- An execution environment that allows simplified application development. Developers can rely on the scheduler to handle preemption.
- A mechanism that prevents ill-behaved modules from monopolizing the processor.

The kernel itself is not pre-emptible.



10 Minutes

Exercise 3-1 Determine Hardware and Operating System Components

The Marketing department employees are going to be added as users on your server. You expect an increase in usage and would like to review the server hardware to establish some baseline statistics before the users are added.

In this exercise you use NetWare Remote Manager to identify server hardware by doing the following:

- [Part I: Review Disk Controller Statistics and Establish a Usage Baseline.](#)
- [Part II: Review the Network Board Driver Statistics and Version.](#)
- [Part III: View Processor Information](#)

Part I: Review Disk Controller Statistics and Establish a Usage Baseline.

Do the following:

1. From NetWare Remote Manager, select **Disk / LAN Adaptors** from the Manage Hardware section.
2. Select **Storage Statistics**.
3. Review the **Current** and **Peak** statistics and record the values in the following:

Table 3-2

	Current	Peak
IO Request Rate:		
Data Transfer Rate:		

4. Select **IO Request Rate** to view a graph.
5. Close the graph.

In a production environment, this process would be repeated to establish a baseline of usage over a period of time. After the users are added, statistics would be taken to determine if the server hardware is appropriately handling the added demand.

Part II: Review the Network Board Driver Statistics and Version.

Do the following:

1. From NetWare Remote Manager, under **Manage Hardware** select **Disk / LAN Adaptors**.
2. Select **Network Statistics** and view the statistics.
3. Select **Packets per Second Graph** to monitor the current traffic.
4. Close the graph.

5. Select **Disk / LAN Adaptors**.
 6. Next to the Network Adaptor select the Info icon.
 7. Review the information.
 8. Check the **Adaptor Resets** statistic for a non-zero value.
If the value is something other than zero, the adaptor should be tested to see if it needs to be replaced.
 9. Record the driver filename and version.
-
10. Open another browser window and access the manufacturer's web site to find out the latest version of the network board driver.
 11. If an updated driver is available, download and install it.
 12. Check the Disk Controller upgrade the drivers to the latest versions.

Part III: View Processor Information

Do the following:

1. From NetWare Remote Manager, under Manage Hardware, select **Processors**.
2. View the information for Processor 0.

If your server has multiple processors, they will also be listed.

(End of Exercise)

Objective 2 **Troubleshoot and Resolve NetWare Server Issues**

In this objective, you learn how to do the following:

- [Identify the Top Novell Technical Support Server Issues and How to Resolve Them](#)
- [Identify Problems after Installation](#)
- [Resolve Console Lock Ups](#)
- [Resolve Hard Disk Errors and Access Problems](#)
- [Resolve Application Monopolizing Server CPU](#)
- [Resolve Server Memory Problems](#)
- [Resolve Slow Server Response](#)
- [Identify Multiprocessing Problems](#)
- [Find Tools for Managing Servers](#)

Identify the Top Novell Technical Support Server Issues and How to Resolve Them

It is important to be aware of current support issues that can affect your server environment.

The Novell Technical Support Knowledgebase is a great resource to locate top support solutions.

You can search Knowledgebase by product categories or select Product Specific Online Support to view the most recent solutions.



Novell Technical Support Knowledgebase is available at http://support.novell.com/search/kb_index.jsp.

Identify Problems after Installation

If the server is not communicating properly after installation, it is usually an indication that the network was installed incorrectly or incompletely.

The following lists a few things to check:

Table 3-3

Symptom	Solution
The server is not communicating.	<ul style="list-style-type: none">■ Check all network boards for conflicting address and I/O settings.■ Verify that the network board driver is loaded and bound to the correct protocol.■ Make sure all network boards are seated properly and are initializing.■ Verify that all cables are fastened securely to all network boards and network connectors.■ Verify the correct IP address and subnet mask are being used.■ Test IP connectivity by pinging the server from a workstation and ping the workstation from the server.■ If the workstation can ping the server, but the server cannot ping the workstation, check the default route on the server.■ Verify that the server has not run out of packet receive buffers.

Resolve Console Lock Ups

If the server console locks up so that you cannot enter commands, but there is no abend message on the System Console or Logger screen, follow these steps to troubleshoot the problem:

1. Verify whether you can move between console screens. If so, the problem might be caused by high server utilization.
2. Verify whether the server console locks up when you unload a specific NLM.

If so, the NLM is probably the source of the problem. Contact the NLM vendor.

3. Make sure you are using the latest disk and LAN drivers, BIOS, and firmware. If not, update the disk and LAN drivers, BIOS and firmware.
4. Verify whether the server console locks up after you mounted the last volume. If so, the network board might not be seated correctly or might not be configured correctly. Check the board and its configuration and correct any problems.
5. Verify whether you can you access the NetWare Internal Debugger by selecting Shift + Alt on the right side of the keyboard and Shift + Esc on the left side of the system console keyboard. From the Debugger, press G to go back to the system console.

Users can't access the server while the NetWare Internal Debugger is active.

6. If the console is locked, if you can't switch among screens, and if you can't enter the Debugger, contact Novell Technical Support or your computer vendor to learn how to generate a nonmaskable interrupt (NMI) to shut down the server.

Resolve Hard Disk Errors and Access Problems

To diagnose hard disk access problems, determine whether any of the following conditions exist:

Table 3-4

Symptom	Solution
The disk driver has not been loaded.	<ul style="list-style-type: none"> ■ Make sure both the HAM and CDM driver are loaded. ■ Check the logger screen to see if the modules loaded correctly or try loading the modules from the server console prompt. ■ Verify that the Disk Adaptors list in NetWare Remote Manager matches actual hardware. Click the Disk / LAN Adaptors link in the navigation frame; then click the Info icon for each disk controller in your server. ■ Try loading the driver from the server console prompt on the server.
A hard disk is not installed or cabled correctly.	<ul style="list-style-type: none"> ■ Check the cables between the hard disks and the controller boards. Make sure that Pin 1 of each cable is attached to Pin 1 of each connector. <p>A general rule is that the red stripe on the cable is closest to the power connector on the drive.</p>
The communication channel between the controller interface board, the disk coprocessor board, and the hard disk is not functioning.	<ul style="list-style-type: none"> ■ Check the power cables and make sure they are seated correctly in the power sockets on the hard disks. ■ Verify that no interrupt conflicts exist. ■ Repair or replace the faulty hardware. ■ Verify that the driver is current.

Table 3-4 *(continued)*

Symptom	Solution
The hard disk controller board is not terminated or addressed correctly.	<ul style="list-style-type: none"> ■ Use the SCSI configuration utility to correct addressing issues and add proper termination.

Point out that the solutions discussed in this section are intended to be a starting point and are not intended to be inclusive.



The solutions listed in the table are intended as general troubleshooting guidance to lead you through some common problems and help build the skills necessary to diagnose and resolve technical server issues. They are not inclusive.

Resolve Application Monopolizing Server CPU

A significant server problem are applications that monopolize the server CPU processing time. Slow server processing speed can indicate a poorly-designed application or a conflict between applications.

A program or NLM that monopolizes the server CPU is sometimes referred to as a CPU hog.

The following are suggestions for resolving the problem:

Table 3-5

Symptom	Solution
A program or NLM is monopolizing the CPU.	<p>To solve this problem, locate the application using NetWare Remote Manager.</p> <p>Once located, the solution to the problem depends on the situation. You can</p> <ul style="list-style-type: none"> ■ Unload the module ■ Replace the module with a newer version that has been fixed ■ Use NetWare Remote Manager to debug and troubleshoot the problem down to the disassembled code level ■ Unload and reload the module and troubleshoot if the problem re-occurs ■ Contact the application vendor for support

Resolve Server Memory Problems

The following types of memory problems can occur on a server:

- [Resolve Server Not Recognizing All Memory](#)
- [Resolve Server Memory Leaks](#)
- [Free Server Memory Temporarily](#)
- [Resolve Server Memory Error Messages](#)

Resolve Server Not Recognizing All Memory

After installing a new server or adding memory, the following generally can resolve the problem of the server not recognizing all memory:

Table 3-6

Symptom	Solution
The server doesn't recognize all the memory installed on the system board.	<ul style="list-style-type: none"> ■ Remove external memory managers from CONFIG.SYS and let NetWare register the memory. ■ Remove DOS=HIGH ■ Remove any DOS devices ■ Remove memory managers such as HIMEM.SYS or EMM386.EXE. ■ Remove memory managers from loading in AUTOEXEC.BAT. ■ Make sure Windows 95 system files are not being used to boot the server because a memory manager is autoloading. ■ Make sure the server BIOS is current and reports the amount of memory correctly. ■ Make sure the memory is inserted into the system board correctly.

Resolve Server Memory Leaks

Because memory is critical to the performance of a server, you need to understand what to do if an application uses up the available memory.

For example, loading NWASPI.CDM 3.22 (from NW6SP1.EXE) creates a situation where NWPA.NLM uses large amounts of RAM. (For details see TID 10069351.)

A memory leak means an NLM or set of NLMs has requested memory from the server but has not returned the memory when finished with it.

Over time, the amount of available memory decreases until eventually the server generates memory error messages.

The memory leak might be slow or fast depending on the amount of memory requested and not released each time.

If you reboot the server, the memory is returned to the memory pool, and the low memory error messages (such as Short Term Alloc Memory errors) stop temporarily until the memory leak ties up enough memory to generate the error messages again.

The following are some troubleshooting suggestions for resolving memory leaks:

Table 3-7

Symptom	Solution
<p>An application is using memory and is not releasing it correctly.</p>	<ul style="list-style-type: none"> ■ Check the system console for error messages and identify the application name. ■ From NetWare Remote Manager Diagnose Server, view the Available Memory status. If the status is not Good, a problem exists. Select Available Memory, NLM Memory, Alloc Memory to locate the problem application. The NLMs using the most memory will be shown first. It is normal for SERVER.NLM, NSS.NLM, and DS.NLM to be the top 3. The offending module can be tracked down by watching the amount of allocated memory. ■ From NetWare Remote Manager, use options such as Diagnose Server, Health Monitor, Health Statistics Trend Graphs, Available Memory (duration), Draw Selected Graphs. The graph can help you determine if you have a memory leak.

Free Server Memory Temporarily

If the server runs out of memory, a quick fix is to free up memory until more can be added.

Table 3-8

Symptoms	Solutions
The server is out of memory.	<ul style="list-style-type: none"> ■ Unload any NLM (such as NWCONFIG, MONITOR, backup software, virus scanners), stop the GUI, and unload other services not needed. ■ On NSS volumes, monitor and adjust the cache statistics. ■ Check the status of the available cache buffers. If the cache buffers are fewer than 20%, add more RAM to your server. <p>If you are using the traditional file system, do the following:</p> <ul style="list-style-type: none"> ■ Try converting volumes from the traditional file system to NSS. It requires less memory to mount NSS volumes. ■ Dismount volumes that are not being used and reduce the size or number of volumes that the server supports. ■ Remove name space support if it is not being used. ■ Delete unused files and directories and purge them. ■ Streamline the directory structure. Every subdirectory takes at least one directory block (by default, a 4 KB block of memory). Therefore, subdirectories with only one file require as much memory as directories with 32 files.

Resolve Server Memory Error Messages

The Novell Technical Support Knowledgebase is a good place to start when you receive error messages on the server console. Error codes and messages are also described in general in the online documentation and the Novell Developer Kit.



The Novell Developer Kit documentation is available at <http://developer.novell.com/ndk/doc.htm>.

A Novell error code is a hexadecimal or decimal number that is usually shown within an error message for an application, such as 0x8996 SERVER OUT OF MEMORY:

- 0x89xx series numbers are returned by the server.
- 0x88xx series numbers are returned by the client.

Resolve Slow Server Response

Slow server response can be caused by a variety of things. To diagnose slow server response, you must identify the conditions.

The following are suggestions for resolving slow server response issues:

Table 3-9

Symptom	Solution
<p>The server or workstation network board is slow or faulty.</p>	<ul style="list-style-type: none"> ■ If a workstation or the server seems slow, insert a new network board into the slow computer to check performance. <p>If the speed is still below normal, reinstall the original network board and then replace the cable attaching the workstation or server to the network.</p> <ul style="list-style-type: none"> ■ Check the status of packet receive buffers and service processes on the server. <p>Compare their values to the maximum allowable values.</p> <p>To check the health of these values, use NetWare Remote Manager options such as Health Monitor, Allocated Server Processes, Available Server Processes, or Packet Receive Buffers.</p> <p>Packet receive buffers are used to transmit and receive packets.</p> <p>If the number of Pack Receive Buffers is increasing, the server operating system will be sluggish.</p>

Table 3-9 (continued)

Symptom	Solution
<p>The server or workstation network board is slow or faulty. (continued)</p>	<ul style="list-style-type: none"> ■ If the number of Packet Receive Buffers reaches the maximum and no ECBs are available, the system will become very sluggish and might not recover. <p>If the current server process are approaching the maximum, you should consider increasing the Maximum Server Processes SET parameter value.</p> <p>If you have only a few available server processes, your server is probably very busy. You might consider increasing the Minimum and Maximum Server Process SET Parameter.</p> <p>To change the values for these parameters, access NetWare Remote Manager; then select Set Parameter and adjust the following:</p> <ul style="list-style-type: none"> ■ The minimum and maximum Packet Receive Buffers (Communications) ■ The minimum and maximum Services Processes (Miscellaneous) <ul style="list-style-type: none"> ■ It might be necessary to analyze a LAN packet trace for slow login and slow response problems.
<p>The Network cabling is faulty.</p>	<ul style="list-style-type: none"> ■ Check the cable with a cable tester and replace faulty cabling.
<p>The server hard disk is slow or faulty.</p>	<ul style="list-style-type: none"> ■ Check the Hot Fix status of all hard disks that use the traditional file system. <p>To view the status, load MONITOR.NLM, and from Available Options, select Storage Devices, Hot Fixed Partition.</p> <p>A non-zero count for Used Hot Fix Blocks indicates problems on the drive.</p>
<p>The server is low on memory.</p>	<ul style="list-style-type: none"> ■ Free memory and then add more if required.

Table 3-9 (continued)

Symptom	Solution
The volume has too many deleted files that have not been purged.	<ul style="list-style-type: none"> ■ Purge deleted files by doing the following: From NetWare Remote Manager, select Volumes, the Volume Information icon next the volume you want to delete files on, Purge Deleted Files. ■ You can also set the Purge attribute on files you want to be purged.
Network traffic is extremely high.	<ul style="list-style-type: none"> ■ Check the LAN driver statistics by doing the following: <ul style="list-style-type: none"> ■ From NetWare Remote Manager, select Disk / LAN Adaptors, Network Adaptor Info for each network board on the Hardware Adaptors page. ■ If you are using more than one network board in the server, compare the boards' Total Packets Transmitted statistics. ■ If one board is receiving most of the traffic, recable the network so that the boards have equal loads.
The cabling system is experiencing too much interference.	<ul style="list-style-type: none"> ■ Check the cabling for interference from fluorescent lights, microwaves, radar, X-rays, and copy machines. Either move the cable or shield it from the source of interference.
Insufficient directory buffers, cache buffers, or packet receive buffers have been allocated.	<ul style="list-style-type: none"> ■ Use NetWare Remote Manager to monitor and adjust the settings.

Identify Multiprocessing Problems

A multiprocessing environment is naturally more complex than uniprocessing. It might be necessary to troubleshoot a problem by stopping and starting processors.

In this section, you learn how to

- [Show Processor Information](#)
- [Stop Processors](#)
- [Start Processors](#)

Show Processor Information

In this section you learn how to show processor status and start and stop individual processors in a multiprocessing server. Only secondary processors can be handled this way.

To view the status of all processors on the server, enter `DISPLAY PROCESSORS` or select Processors from NetWare Remote Manager.

When the Platform Support Module (PSM) is loaded in the `STARTUP.NCF` file and the Auto Start Processors `SET` parameter value (Multiprocessor category) is set to On, NetWare can start the secondary processors automatically.



Server console commands for processors affect only secondary processors. Processor 0 cannot be taken offline while the server is running.

You can start or stop secondary processors at any time while the server is running by using the console commands `START PROCESSORS` and `STOP PROCESSORS` or by clicking the Start Processor or Stop Process link on the Processor Information page in NetWare Remote Manager.

If you choose to start secondary processors manually, you can change the value for the Auto Start Processors `SET` parameter to Off.

When any secondary processor is stopped, the associated threads are automatically switched to another processor.

Before starting or stopping a process, you should show the processor information to determine the status of the processor.

Stop Processors

Processor 0 is the boot processor and cannot be taken offline.

When a secondary processor is taken offline, the threads running on that processor are switched to another processor automatically.

You can stop processors using NetWare Remote Manager by selecting Processors, or you can use the server console command STOP PROCESSORS.

Start Processors

The default is to start secondary processors when the server boots. If a secondary processor is offline, you can bring it online using NetWare Remote Manager or server console commands.

Find Tools for Managing Servers

Many resources are available to help you maintain servers, to troubleshoot, and to prevent server problems.

A few valuable resources are

- [NetWare Cool Solutions](#)
- [AppNotes](#)
- [Novell Technical Support Tools](#)
- [Third-Party NetWare Support Tools](#)

NetWare Cool Solutions

The Novell Cool Solutions Communities are dedicated to helping customers get the most out of Novell products.

You will find hundreds of tips, articles, Q&A, and free tools.



Novell Cool Solutions Communities are available at <http://www.novell.com/cool solutions/>.

AppNotes

Novell AppNotes is published monthly by Novell. The material in AppNotes is based on actual field experience and technical research performed by Novell personnel, covering topics in these main areas:

- Network design and optimization strategies
- Network management tactics
- Novell product internals and theory of operations
- Novell product implementation guidelines
- Integration solutions for third-party products
- Network applications development and tools

This information benefits Novell's technical audience: system engineers, support engineers, consultants, programmers, network managers, and information systems personnel.



AppNotes is accessible online at <http://www.novell.com/appnotes/>. A hard copy subscription is also available.

Novell Technical Support Tools

Support tools are useful for preventing and troubleshooting server problems. These support tools can be downloaded free of charge.

The following are a few of the many tools available:

- **Enhanced ToolBox Utility.** TOOLBOX.NLM lets various utility functions be executed on the server console or via NCF files without involving any clients.

You can download this tool at

<http://www.novell.com/cool solutions/tools/1490.html>.

- **Server Configuration Information Tool.** The Server Configuration Information program creates CONFIG.TXT that contains server configuration information.

CONFIG.NLM works on all versions of NetWare; however, NetWare Remote Manager provides the same functionality in NetWare 5.x and 6.

You can download this tool at

<http://www.novell.com/cool solutions/tools/1506.html>.

- **NetWare Config Reader.** The Config Reader (2.67) is a Windows 95/98/NT client utility that analyzes the CONFIG.TXT file produced by CONFIG.NLM.

You can download this tool from

<http://www.novell.com/cool solutions/tools/1500.html>.

- **Other tools.** There are many tools available for download that address a wide range of administrative needs.

Support tools from Novell Technical Support are available at <http://support.novell.com/tools-files.html>.

Third-Party NetWare Support Tools

Many support tools available from third-party vendors can assist with troubleshooting. A few are mentioned here to get you familiar with the types of software available for troubleshooting on a NetWare platform.

- **Alexander SPK NetWare (System Protection Kit).** This kit provides tools to handle server problems and help prevent or identify problems before they occur.

When an abend occurs, the abend call is intercepted by EDNA.NLM (Emergency Diagnostics for NetWare Administrator), part of the Alexander SPK, and takes control of the system. EDNA tries to prevent a crash by suspending the errant NLM.

If successful, a hard crash is prevented, a log file is generated, and administrators can be notified through SNMP traps. The SPK Crash Report immediately pinpoints the module that caused the crash. Alexander SPK v4.1 for NetWare runs on NetWare 6 - 3x.



Alexander SPK v4.1 for NetWare runs on NetWare 3.x through 6. For more information, see <http://www.alexander.com/>.

- **Storage Manager for Novell NetWare.** This is server image and recovery software that supports NetWare 6 and NSS.

Storage Manager minimizes the management, setup, installation, and reconfiguration time for NetWare servers. It offers data and disaster recovery.



For more on Storage Manager for Novell NetWare, see <http://www.portlocksoftware.com/products/stormgr/index.htm>.

- **Dave's Novell Shareware.** This is a free service for NetWare enthusiasts and authors of NetWare-based software, shareware, and freeware.

It is intended to provide a free resource location for software and utilities.



For more on Dave's Novell Shareware, see <http://www.novellshareware.com/server-management.shtml>.

- **NetWarefiles.com.** This web site is a repository for third-party utilities for NetWare and other Novell products.



For more on NetWarefiles.com, see <http://www.netwarefiles.com/>.



30 Minutes

Verify that PROBLEMS.NLM was migrated in volume DATA (EXERCISE\SECTION3) to each student's NetWare 6 server.

Exercise 3-2 Resolve Server Problems

After Digital Airlines upgraded to NetWare 6, the help desk received some support calls. As the network administrator, you also noticed a few problems. The issues include performance, communication, and file access.

Troubleshoot and resolve the following:

- **Part I: Resolve a Server Memory Leak**
- **Part II: Locate a Server CPU Hog**
- **Part III: Troubleshoot Memory Corruption Problem**
- **Part IV: Resolve Open File Access Problem**
- **Part V: Resolve a Server Process Hog Problem**

To prepare for this exercise you need to do the following:

1. From your workstation, use Windows Explorer to locate **DAx\DATA\EXERCISE\SECTION\PROBLEMS.NLM**.
(If you cannot locate PROBLEMS.NLM on volume DATA, the file is also available on your Enhanced Learning CD in EXERCISES\SECTION 3.)
2. Copy **PROBLEMS.NLM** to **SYS:SYSTEM** on your **DAx** server.

Part I: Resolve a Server Memory Leak

To simulate and troubleshoot a server memory leak, do the following:

1. Launch NetWare Remote Manager.
2. From the Diagnose Server list, select **Health Monitor**.
3. View the Available Memory status.
4. Select **Available Memory**.
5. View the current memory usage graph.
6. Record the percentage of **NLM Memory**:

7. Load **PROBLEMS.NLM** on your server console***.
8. Select **Memory Hog** from the menu.
9. Leave the Trace Portal window open (*do not* press Enter).
10. From the workstation, select Health Monitor icon in the upper left side of the window.

The status should now indicate Bad (red) for Available Memory.
11. Select **Available Memory**.
12. View the current memory usage graph.

13. Record the percentage of **NLM Memory**:

14. Compare this percentage with the percentage in Step 6.
Notice the large increase, indicating an NLM using an unusual amount of memory.
15. Select **NLM Memory** to show currently loaded NLMs.
16. Select **Alloc Memory** to sort the NLMs by allocated memory.
Notice that **PROBLEMS.NLM** is listed as the NLM using the most memory.
17. Select the **Alloc Memory** link for **PROBLEMS.NLM**; then view the **Memory Allocation Summary**.
18. The errant NLM has now been identified and you decide to unload it:
 - a. From **Manage Applications, List Modules**, select **PROBLEMS.NLM**; then select **Unload**.
 - b. Select **OK**.
19. From the **Health Monitor**, view the status of **Available Memory**.
It might be necessary to refresh the screen until the status returns to **Good** (green).

In this part of the exercise, you were able to locate an NLM using too much memory and unload it remotely through NetWare Remote Manager.

Part II: Locate a Server CPU Hog

To simulate a high CPU utilization problem, do the following:

1. Load **PROBLEMS.NLM** on your server console.
2. Select **Utilization Hog** from the menu.

3. Leave the Trace Portal window open (*do not* press Enter).
4. Locate the source of the problem:
 - a. Select **Health Monitor**.
 - b. Select **CPU Utilization** and view the graph.
5. Notice that the CPU Utilization status is **Bad** and the utilization is staying at about 98-100%.
6. Select **Diagnose Server > Profile / Debug** and note that **PROBLEMS.NLM** is causing the problem.
7. Unload the NLM:
 - a. Select **PROBLEMS.NLM** from the list.
 - b. Select **UNLOAD > OK**.
8. Select **Health Monitor**.
9. CPU Utilization status should now indicate **Good**.

In this part of the exercise, you were able to locate and resolve a server CPU hog using NetWare Remote Manager.

Part III: Troubleshoot Memory Corruption Problem

To simulate a memory corruption problem, do the following:

1. Load **PROBLEMS.NLM** on your server console.
2. Select **Memory Corruption** from the menu.
3. Leave the Trace Portal window open (*do not* press Enter).

It is now the day after the upgrade and you decide to check the console error log of your upgraded server:

1. From NetWare Remote Manager, select **Health Monitor** and notice that the list of items monitored do not report a problem.
2. Select **Reports / Log Files**.
3. Select **System Error Log File**.

While scrolling through the log file, you notice the following error near the end of the file:

```
"Free detected corrupt trailing redzone for node  
0xCCB72140, node size 24"
```

From an NTS Knowledgebase query you determine that the cause is a portion of memory has been improperly overwritten by an errant program.

The specific module mentioned in the article is not running on the server, so you decide to troubleshoot the problem further.

You suspect that it might be an application named PROBLEMS.NLM that was installed prior to the upgrade.

4. Select **List Modules**.
5. Locate PROBLEMS.NLM in the list.
6. Select the **Alloc Memory** number for that NLM.
7. From the PROBLEMS.NLM Allocation Summary, select **Display Memory Allocation Information by Size**.
8. View the Corruption Count column for any non-zero counts and select that item.

A corruption count indicates that the header or footer of the allocated memory has been corrupted.

9. Select the Address link to further view the data.
10. Switch to the PROBLEMS.NLM screen on the server.
11. Select Enter to end the Memory Corruption simulation.

You now know that there is a corruption problem, you located the application causing the problem, and you know the size of memory corrupted and what data written to it.

You can further resolve the problem when you contact the vendor and provide them with this information.

Part IV: Resolve Open File Access Problem

To simulate a file access problem, do the following:

1. Load **PROBLEMS.NLM** on your server console.
2. Select Lockup File **SYS:TESTDB.DAT** from the menu.
3. Leave the Trace Portal window open (*do not* press Enter).

The Digital Air help desk reports that the file **SYS:TESTDB.DAT** on your server cannot be updated because it is being held open by someone. Find out who has the file open and alert the person to close the file.

To resolve the problem, do the following:

1. From NetWare Remote Manger, select **Volumes** under Manage Server.
2. Select volume **SYS**.
3. At the left of the **TESTDB.DAT** file select the **question mark icon**.
4. From the Global Lock Information section, notice that the Use Count is 4.
5. From **File Lock Information by Connection**, record the connection numbers of the users listed:

6. Select **Connections**.
7. Select the user name for one of the connections that have the file locked.
8. From the Connection Information screen, you can see that the user has the file open.
9. In the Send Message box, enter a message asking the user to close the file.

10. There was no response from the user, so you decide to clear the connection so that the file lock is removed:
 - a. Select **Connections**.
 - b. Next to the user that you sent the message to select **Clear Connection**.
11. Select **Volumes**.
12. Select volume **SYS**.
13. At the left of TESTDB.DAT select the **question mark icon**.
14. From the Global Lock Information section, notice the Use Count is now 3.

In this part of the exercise, you were able to find out who had a file open, communicate with that user, and free the file for others to use by clearing the connection of the user that had the file locked.

Part V: Resolve a Server Process Hog Problem

To simulate a server process hog problem, do the following:

1. Load **PROBLEMS.NLM** on your server console.
2. Select **Server Process Hog** from the menu.
3. Leave the Trace Portal window open (*do not* press Enter).

To resolve the problem, do the following:

1. From NetWare Remote Manger, select **Health Monitor**.

The state of Available Server Processes should be Suspect or Bad.

2. Select **Available Server Processes**.

Notice the number of available server processes has dropped.

3. Select **Profile / Debug** to show the Profiling and Debug Information screen and notice there are several threads with a Thread State of Delayed.
4. Select **Thread Information**.
5. Scroll down to the server process threads named Server 00, Server 01 and so forth and notice that the Thread State is Delayed for all server processes.
6. Select any delayed server process thread names to show the Thread Information screen and notice the following:
 - The Suspend Reason is Delayed.
 - The Active Work To Do Information shows the Work Owner as PROBLEMS.NLM.

After selecting a few more server process threads, you determine the process hog is PROBLEMS.NLM.

7. From the Console Screens window, select Console Screens.
8. Select the PROBLEMS.NLM screen if necessary.
9. Select Enter on the Trace Portal window of PROBLEMS.NLM.
10. Select the **Available Server Processes** window from the Health Monitor screen and notice the number of available server processes has returned to normal.
11. Close the **Available Server Processes** window from the Health Monitor screen.
12. Select **Profile / Debug** to show the Profiling and Debug Information screen.
13. Select Thread Information.
14. Scroll down to the server process threads named Server 00, Server 01, and so forth, and notice that the Thread State is Waiting for work for all server processes.

In this part of the exercise, you were able to locate PROBLEMS.NLM as a server process hog.

Using NetWare Remote Manager, you determined which NLM was monopolizing the Available Server Processes by viewing the thread's work owner and thread state.

(End of Exercise)

Objective 3 Troubleshoot and Resolve Critical Server Abends

In this objective you learn how to troubleshoot and resolve critical server abends caused by hardware or software.

The NetWare OS is very resilient, but errors can occur. It is crucial to be able to resolve critical server errors quickly and effectively.

These types of problems are referred to as crashes, hangs, or abends. The first 2 terms are general and typically used by nontechnical people to describe why the server is not performing or available. The term abend is typically used by administrators.

You learn the following:

- [What an Abend Is](#)
- [What Types of Abends Occur](#)
- [What an ABEND.LOG File Is](#)
- [What a Core Dump Is](#)
- [How to Respond to an Abend](#)
- [How to Create and Submit a Core Dump for Analysis](#)

What an Abend Is

The term *abend* is an acronym for ABnormal END. An abend is a serious software failure that halts (brings down) the server.

When either NetWare or the CPU experiences an unexpected critical error, the NetWare fault handler starts, processing stops, and the server shows an abend message.

Because processing stopped, the server is left in a state where tasks are not completed or are only partially completed. This can result in database corruption or inconsistent data.

Starting with NetWare 4.11, automatic abend recovery features were implemented. These features are enabled by default when the server is installed. Prior versions require manual intervention.

After an abend occurs and the server abruptly halts, disk I/O requests and other processes might be in an incomplete state.

If the server is rebooted, the pending requests might be lost and there is the possibility of file system corruption.

Automatic abend recovery allows the server to be shut down gracefully, saving user data and completing disk I/O requests.

A server that has abended will have the number of times the server has abended added to the server name on the console. For example,

```
DA1 <1> :
```

indicates that server DA1 has experienced 1 abend.

What Types of Abends Occur

Abends can be divided into 2 categories:

- **Operating System (Software) Detected**
- **Processor (Hardware) Detected**

Operating System (Software) Detected

When NetWare detects a software exception, it calls the abend handling routine.

The abend message from a software exception is a text message that tries to describe the problem.

Processor (Hardware) Detected

An error condition, or fault, that is detected by the processor is called a *processor exception*.

All hardware-detected abends have the words “processor exception” in the abend message.

Some common processor-detected abends include the following:

- **Page faults.** These occur when a process tries to address memory outside its registered space or in protected memory.
- **General protection processor exceptions (GPPEs).** These occur when a process tries to address memory above the physical memory limit.
- **Non Maskable Interrupts (NMIs) and Machine Check Exceptions.** These are almost always memory related and produced by parity errors. For example, faulty/failed memory chips produce NMI errors.

These memory-related errors can occur in main memory on the system board, in add-in memory boards, and in shared-memory areas of I/O cards.

Machine checks are produced by the Intel Pentium chip when an internal hardware error is detected.

- **Invalid Opcodes.** These occur when the processor detects an instruction that is inconsistent with the processor's instruction set.

What an ABEND.LOG File Is

A log file, ABEND.LOG, is created at the time of an abend. It contains the abend message, along with additional information, and is saved on the DOS partition.

As soon as the server is restarted, the ABEND.LOG file is moved to SYS:SYSTEM.

Analyzing a core dump can be time consuming, so Novell engineers have created a process to expedite analyzing server abends.

The Abend Analysis System lets you submit an ABEND.LOG to Novell. The log file is compared with a database of known solutions and if a matching error is found, you receive suggestions on how to resolve the abend.



Novell's Abend Analysis System is located at <http://abend.novell.com/>.

The ABEND.LOG file contains the following information:

- File server name
- Date and time of abend
- Abend message
- Registers

- Abended NLM
- Running process
- Stack limit and pointer
- Stack trace
- Modules list

What a Core Dump Is

Define the term *core dump*.

If an abend continues to occur, it might be necessary for Novell Technical Support to analyze a capture of the server's memory (a *core dump*).

The term *core dump* comes from the mainframe environment where RAM was referred to as core memory. The process of saving (*dumping*) a snapshot of a server's memory to disk is referred to as creating a core dump.

A core dump is a byte-for-byte image or snapshot of a NetWare server's memory at the time an abend occurs. Memory is not refreshed when the server is in an abended state.

The core dump contains the following information that can be used to analyze the problem:

- **Processes.** All processes on the server at the time of the abend are included in the core dump. The state of these processes can be running, waiting to run, or not in use. A history of what the process has done (call stack) is also preserved in the core dump.
- **Loaded modules.** This includes module information, code, and data of all NLMs.
- **Allocated memory.** This includes the memory allocated by processes included in the core dump.
- **Cache memory.** Memory that is available for allocation by modules or processes can also be included in the core dump.

Most of the time it is not necessary to include cache memory in a core dump and the latest utilities allow it to be excluded from the core dump.

- **Screen shots.** Console screens are preserved in the core dump. This includes helpful information such as the abend message, server name, and application error messages.

How to Respond to an Abend

When the server abends, it shows an abend message similar to the following:

```
ABEND: SERVER-6.xx-message_number message_string
ADDITIONAL INFORMATION: message
```

The Additional Information section states the probable cause of the abend. It indicates where the problem occurred and gives the name of any NLM associated with the abend. This information helps you determine how to resolve the abend.

You can respond to the abend manually or have the server respond automatically.

When you respond manually, the server determines the nature of the abend and shows the appropriate response option on the screen, along with additional options for bringing down the server or executing a core dump. You must execute an option to respond to the abend.

When the server responds automatically, it executes the appropriate response without intervention.



Sometimes an abend (or a faulty NLM) can cause the server console to stop functioning. In this case, the abend message is not shown and you cannot enter commands at the server console prompt.

After a server failure, it is recommend that the computer be turned off and restarted. This process helps ensure nothing is retained in memory.

Respond to an ABEND Automatically

Demonstrate how to set the parameters in NetWare Remote Manager.

The default parameters on the server are set to respond to an abend automatically. The server automatically recovers from most abends and continues functioning normally.

Users can save their files before the server is restarted and file system corruption can be avoided because volumes can be properly dismounted.

Three SET parameters control how the server responds:

Table 3-10

Parameter	Value
Auto Restart After Abend = 1 (values 0 – 3)	<p>0. Do not try to recover from the abend. The server is left in a halted state. This option is discussed more in “Respond to an Abend Manually” on 3-54.</p> <p>1 (Default). For software abends, NMIs, and Machine Check Exceptions: attempt to recover from the problem, bring down the server in the configured amount of time, and then restart the operating system.</p> <p>For other exception abends, suspend the faulting process and leave the server up.</p> <p>2. For all software and hardware abends, attempt to recover from the problem, bring down the server in the configured amount of time, and then restart the operating system.</p> <p>3. For all software and hardware abends, immediately restart the server.</p>

Table 3-10 (continued)

Parameter	Value
Auto Restart After abend Delay Time = 2 (Range: 2 to 60 minutes)	This setting indicates how many minutes the server will wait after an abend occurs before going down and restarting itself. In most cases, the server can recover but is in a critical state that requires a restart. The purpose is to prevent data loss by giving users time to save files and log out before the server is restarted.
Auto Restart Down Timeout = 180 (Range: 0 to 600 seconds)	If there is a problem bringing down a server after an abend, this setting specifies the amount of time to wait for a server to go down before the server restarts.

To set the parameter values, use NetWare Remote Manager and select the Error Handling category. You can also use MONITOR or the console SET command.

Because the server responds to the abend automatically, you might not know when an abend has occurred.

Therefore, periodically check the ABEND.LOG file from the Profiling and Debug Information screen in NetWare Remote Manager (look for Suspended next to Abend Recovery status).

Respond to an Abend Manually

To respond manually to abends, change the following SET parameter (Error Handling category) to the value shown:

```
AUTO RESTART AFTER ABEND = 0
```

When an abend occurs, the server shows a short list of options appropriate to the nature of the abend.

To respond to the abend, you must execute one of the options by entering the first letter of the option.

The following options can appear:



Several options have the same first letter (such as R, S, or X). In a given abend situation, the option list includes only one option for any given first letter.

Table 3-11

Command	Action
S=Suspend the running process, update ABEND.LOG, and attempt to bring down the server. (OS Detected)	<p>This option appears if the abend was detected by NetWare.</p> <p>The server sends a message to users that the server is going down and advising them to save their files and log out.</p> <p>The amount of time before the server shuts down and restarts is determined by the SET parameter AUTO RESTART AFTER ABEND DELAY TIME.</p> <p>Review the ABEND.LOG file to help determine the source of the problem.</p>

Table 3-11 (continued)

Command	Action
<p>S=Suspend the running process and update ABEND.LOG. (Processor Exception)</p>	<p>This option appears if the abend was detected by the processor.</p> <p>When you execute this option,</p> <ol style="list-style-type: none"> 1. The server suspends the current process and updates the ABEND.LOG file. 2. The server does not shut down the server. Server performance might be poor, because a loaded NLM is probably malfunctioning. 3. Read the Additional Information part of the abend message to learn which NLM might be causing the problem. 4. At a convenient time, shut down the server and restart it. 5. Examine the ABEND.LOG file for more information about the source of the problem. 6. Consider submitting the ABEND.LOG file to the abend log database to see if there is a known solution.
<p>R=Resume the running process, update ABEND.LOG, and attempt to bring down the server. (Processor Exception)</p>	<p>This option appears if the abend has detected a hardware problem.</p> <p>After the server is shut down, you must</p> <ol style="list-style-type: none"> 1. Fix the hardware 2. Run diagnostics to verify new hardware 3. Boot the server <p>If needed, contact the hardware manufacturer for additional assistance.</p>
<p>S=Return the running process to a safe state and update the ABEND.LOG file. (Processor Exception)</p>	<p>When you execute this option,</p> <ol style="list-style-type: none"> 1. The server returns the running process to a safe state 2. The ABEND.LOG file is created. 3. The server is not shut down. <p>In most cases, the server completely recovers and no further action is necessary.</p>

Table 3-11 (continued)

Command	Action
Y=Copy the diagnostic image to disk.	Execute this option to perform a core dump that can be examined to determine the cause of an abend.
X=Restart the server. (DOS removed)	This option appears only if DOS has been removed. Execute this option if you want to restart the server. If DOS has been removed, the server will not create or update an ABEND.LOG file.
X=Update ABEND.LOG and then exit.	Execute this option if you want to shut down the server and exit to DOS. If you turn off the server without first executing one of the S or R options to resolve the abend, the server will not update the ABEND.LOG file.
Turn off and back on to restart.	If the console has been secured, you must turn the server off and then back on to restart the server.

How to Create and Submit a Core Dump for Analysis

In this section, you learn how to

- [Determine the Size of a Core Dump](#)
- [Determine Where to Save the Core Dump](#)
- [Create a Core Dump](#)
- [Send a Core Dump to a Local Drive](#)
- [Send a Core Dump to a Remote Drive](#)
- [Validate a Core Dump](#)
- [Submit a Core Dump to Novell for Analysis](#)

Demonstrate how to view the memory configuration from NetWare Remote Manager (from Manage Server, select View Memory Config).

Determine the Size of a Core Dump

Because a core dump is a byte-for-byte image of a NetWare server's memory, core dumps can be quite large, depending on how much memory is installed on the server.

You can use NetWare Remote Manager to view the memory configuration. It will show total memory, cache memory, and how much memory is used by the file system, swap files, NLMs, and virtual memory.

There are several options when creating a core dump that can affect the size of the file. You can include all memory or exclude cache memory. You can also use compression, which reduces the disk space needed.

The size of the full image file is approximately equal to the total RAM installed in the server.

For a cacheless core dump, the size of the image file is approximately equal to total RAM minus the amount of file cache (disk cache memory).



DIAG500.NLM 2.1.0 or later also excludes NSS cache when the exclude cache option is selected. This results in a smaller core dump image.

Determine Where to Save the Core Dump

Based on your approximation of the size of the core dump, you must determine where to save the file.

The default method is to save the core dump locally to the server's DOS partition as C:\COREDUMP.IMG.

Using an add-on utility called DBNET, you can also save the core dump to a remote server or workstation. DBNET is included in the NetWare 4, 5, and 6 support packs.

If there is not enough room on the DOS partition, an additional drive can be added to the server with a DOS partition.

Other higher capacity drives with removable storage can also be added to the server. Removable storage devices include Zip, Jaz, and SyJet. Some optical or CD-RW drives that support DOS drivers can also be used.

Create a Core Dump

Novell Technical Support might request a core dump when a server experiences a lockup or abend and other troubleshooting has failed to resolve the problem.

Core dumps can be analyzed by Novell engineers and are often the key to finding software bugs.

Your ABEND.LOG file could be equally valuable for diagnosis.



Novell Technical Support requires all necessary patches to be installed before they will analyze a core dump.

You can perform the following types of core dumps:

- A **full core dump** copies all server memory to a local drive or device.
- A **full core dump minus cache** copies all server memory except file cache (disk cache) to a local drive or device.

The cacheless core dump is smaller and in most cases provides as much useful information as a full core dump.

The core dump must still point to the cache memory pages that have been excluded. These are referred to as *phantom* pointers.

The page size is based on the memory chip, such as 4K.

A core dump can be started in 2 ways:

- By responding to the core dump choices shown by NetWare after an abend has occurred.

The Type? prompt lets you specify a full core dump or a cacheless core dump, as explained above.

The Device? prompt lets you specify a local drive or a DOSwriteable device.

- By forcing a core dump by entering .C in the NetWare Internal Debugger.

If the server is not completely locked up, you can you access the Debugger by doing the following:

1. On the right side of the system console keyboard select **Shift + Alt**; on the left side of the system console keyboard select **Shift + Esc**.
2. Enter .C to show core dump creation choices.
3. When the core dump is finished, do one of the following:
 - Enter **G** to exit the Debugger and return to the server console prompt.
 - Enter **Q** to exit to DOS.



Users can't access the server while the NetWare Internal Debugger is active.

Send a Core Dump to a Local Drive

After you start the core dump, you are asked to specify the DOS drive letter and file path that the memory image file will be written to.

The default path and name of the image file is
C:\COREDUMP.IMG.

The drive can be any writable DOS device that contains enough storage space. The device must be set up not only before the server abends but before the server is booted.

After the file is on the hard disk, it can be compressed, copied to disk, backed up to tape, or sent by FTP to ftp.novell.com (if you have opened a support incident).

After the server is running, the image file can also be copied to a workstation or network drive using NetWare Remote Manager.

For servers not running NetWare Remote Manager, this can be done using IMGCOPY.NLM or any other third-party NLM that provides this functionality.



IMGCOPY.NLM is available at <http://support.novell.com/>.

Send a Core Dump to a Remote Drive

DBNET is a set of utilities that lets you send NetWare 4, 5, and 6 core dumps to a remote location on your network.

The destination can be another NetWare server or a Windows workstation.

This method is useful when the DOS partition is not large enough to hold the core dump. This method also greatly decreases the time that the server must remain in a downed state for diagnostics.

DBNET also lets you open a diagnostic connection to a NetWare server and enter the Debugger to diagnose the server remotely.

DBNET is comprised of the following components:

- **DBNETx.NLM.** This is the main program that provides the console interface to manage DBNET.
The *x* represents the version of NetWare (4, 5, or 6). For example, DBNET6.NLM runs on NetWare 6 servers.
- **DBNET.CFG.** This configuration file contains the information DBNET uses to establish and maintain remote connections.
- **IMGHOST.NLM.** This is a NetWare version of the image host program designed to accept a core dump from a remote server.
- **IMGHOST.EXE.** This is a Windows version of the image host program designed to accept a core dump from a remote server.
- **DIAGxxx.NLM.** This is a NetWare utility used to send core dumps over a network connection to the image host program running from a remote location.

The following shows DIAGxxx.NLM and corresponding platform.

Table 3-12

DIAGxxx.NLM	Platform
DIAG411.NLM	NetWare 4.11, 4.2
DIAG500.NLM	NetWare 5x Integrated into SERVER.EXE
DIAG500.NLM 2.00c or later	NetWare 6

- **RDBHOST.NLM.** NetWare utility designed to provide a remote debugging interface.

Validate a Core Dump

After a core dump is created, the next step is to validate it before sending it to Novell for analysis.

Before you send the file for analysis, make sure that your core dump

- Contains useful information:
 - For abends, verify Auto Restart After Abend = 0. This helps ensure the core dump contains information right after the abend occurs.
 - For performance problems, verify the core dump is taken during the time the problem is occurring.
- Can be opened with the NetWare Virtual Debugger.



The NetWare Virtual Debugger is available at <http://developer.novell.com/ndk/vdb.htm>.

Submit a Core Dump to Novell for Analysis

If you cannot resolve a problem, it might be necessary to have Novell engineers analyze the core dump.

Before sending the memory image to Novell, you must have an open support incident.



For a list of available support options, see <http://support.novell.com/>.

A senior Support Engineer will analyze the memory image file and recommend a solution.

The Technical Support Engineer will make arrangements to receive the image either through parcel delivery or through the Internet and will advise you on the best media format to use.

To submit the file via FTP, do the following:

1. Compress the core dump image file using maximum compression in ZIP format.
2. Name the file *incident_number.ZIP*. For example, if your support incident number is 1234567, then name the file 1234567.ZIP.
3. Upload the file to the FTP server for your region unless your support technician tells you otherwise.

For example, you can log in to ftp.novell.com as user anonymous with your email address as the password and upload the file to the /incoming directory using binary transfer mode.



For additional details, see TID 10020665, “Submitting a Core Dump to Novell Technical Services (NTS) to Be Read and Analyzed.” It is especially important that you apply all the latest patches prior to submitting a core dump.



10 Minutes

Exercise 3-3 Submit an ABEND.LOG File for Analysis

As the network administrator, you noticed one of the servers had experienced an abend. As a first step, you decided to submit the ABEND.LOG file to Novell’s Abend Analysis System to see if it is a known issue.

An NLM will be loaded on the server to simulate the abend.

Do the following:

1. On your server, use NetWare Remote Manager to load **PROBLEMS.NLM**:
 - a. From the Manager Server options select **Console Screens**.
 - b. Select **Console Screens**.
 - c. Load **PROBLEMS.NLM**.
2. From the menu select **Page Fault**.

Internet access is required for this exercise.

3. From the server, select the **System Console** screen.
4. Press **Enter** and notice the server console prompt now indicates <1> abend has occurred.
5. Review the abend message.
The offending process has been suspended.
6. Clear the abended state by entering **RESTART SERVER**.
7. After the server restarts, launch a browser.
8. Browse to <http://abend.novell.com>.
9. Select **Submit an Abend.log**.
10. Enter your *email address*.
Your email address should preferably be from an account that you can access from a web browser.
If you do not have an email account, you can set up a free account at www.myrealbox.com. MyRealBox runs Novell NetMail (<http://www.novell.com/products/netmail/>).
11. For the Novell Technical Support incident number enter **Novell Education 3005**.
12. Browse to `SYS:\SYSTEM\ABEND.LOG` and select **Open**.
If you do not see the LOG extension, change the setting on Windows Explorer so known extensions are not hidden.
13. Select **Analyze Abend**.
14. Check your email account for the results.
The results will indicate a solution if one is found. This process might take several minutes to complete.
15. During the wait, open the ABEND.LOG file and identify the following sections:
 - File server name
 - Date and time of abend

- ❑ Abend message
- ❑ Registers
- ❑ Abended NLM
- ❑ Running process
- ❑ Stack limit and pointer
- ❑ Stack trace
- ❑ Modules list

You obtained an ABEND.LOG and submitted it to Novell's online Abend Log Analyzer to obtain a solution.

(End of Exercise)



30 Minutes

Exercise 3-4 Create a Core Dump

When the core dump is started, you might want to suggest the students take a break for a few minutes.

You might want to take a core dump prior to class to estimate the time required based on the classroom hardware being used.

As the Digital Airlines network administrator, you have been unable to resolve an abend on your server after installing a new application.

You have opened an incident with Novell Technical Support to help resolve the problem. The Support Engineer helping you has requested that you send a core dump image for analysis as soon as the server abends.

Complete the following:

- [Part I: Save a Core Dump Image to a Local Drive](#)
- [Part II: Save a Core Dump Image to a Remote Drive](#)
- [Part III: Validate a Core Dump Image](#)

You load an NLM to simulate the problem.

Part I: Save a Core Dump Image to a Local Drive

Do the following:

1. From NetWare Remote Manager, change **Auto Restart After Abend** = **0** to keep the server from automatically restarting.

This lets you obtain a core dump of the memory after the abend occurred. Do the following:

- a. Select **Set Parameters**.
 - b. Select **Error Handling**.
 - c. Select the value for **Auto Restart After Abend**.
 - d. Change the value to **0**.
 - e. Select **OK**.
2. Load **PROBLEMS.NLM** on your server console.
 3. From the menu select **Page Fault**.
 4. Review the abend message.
 5. Copy the diagnostic image to disk (core dump) by selecting **Y**.
 6. As the Coredump Type select **2 - Full W/o Cache (All Server Memory Except File Cache)**.
 7. For Compress Coredump select **1 - Yes**.
 8. Accept the default path (C:\COREDUMP.IMG) by selecting **Enter**.
 9. Allow the core dump process to complete.

This process might take several minutes to complete. After the core dump completes, you are prompted to create another core dump.
 10. *Do not* press **Y**.
 11. When prompted, turn off the server.
 12. Turn on the server to reload the OS.

13. Create a directory named **COREDUMP** off the root of drive C on the workstation.
14. From NetWare Remote Manager, select **Manage Server > Volumes > C:**.
You should see COREDUMP.IMG listed.
15. Select COREDUMP.IMG.
16. Select **Save**.
17. Save to C:\COREDUMP\ on the workstation.

Part II: Save a Core Dump Image to a Remote Drive

Do the following:

1. From **DAx\DATA\EXERCISE\SECTION3**, copy **DBNET6.NLM** and **IMGHOST.NLM** to **SYS:SYSTEM**.
2. From **DAx\DATA\EXERCISE\SECTION3**, copy **IMGHOST.EXE** to **C:\COREDUMP**.

(These files are also available in EXERCISES\SECTION 3 on your Enhanced Learning CD.)

3. On your workstation launch **IMGHOST.EXE**.
4. On your server console load **DBNET6.NLM**.



IMGHOST should indicate the core dump directory as C:\Coredump, and the IP Address will appear as 0.0.0.0. The address showing as 0.0.0.0 will not affect the ability to receive a core dump image and has to do with the Win32 Winsock library.

5. Load **PROBLEMS.NLM** on your server console.
6. Select **Page Fault** from the menu.
7. Switch to the **System Console** screen.

8. Press **Enter** and notice the server console prompt now indicates **<1>** abend has occurred.
9. Review the abend message and notice that the offending process is suspended.
10. Enter the NetWare Internal Debugger by selecting **Right-Shift + Alt** and then **Left-Shift + Esc**.
11. At the **#** prompt, enter **.C** to start the core dump.
12. As the Coredump Type select **2 - Full W/o Cache (All Server Memory Except File Cache)**.
13. For Compress Coredump select **1 - Yes**.
14. In response to Where Should Diagnostic Coredump Be Sent, select **2 - NETWORK -- Dump Across Network to Remote Host**.
15. Enter the **IP address** of the workstation where IMGHOST.EXE is loaded.
16. Allow the core dump process to complete.
17. Record the filename of the core dump image for later reference:

A "1" might be added to the beginning of the filename shown on the screen.

18. Exit the NetWare Internal Debugger and return to the system console by typing **G**.
19. Reboot the server.
20. From the workstation, launch Windows Explorer and navigate to **C:\COREDUMP*<yourservername>***.
21. Verify that the file created in Step 17 is in the directory.

Part III: Validate a Core Dump Image

1. From your administration workstation, install the NetWare Virtual Debugger from `\\DAX\DATA\EXERCISE\SECTION3\VDB.EXE`.
2. Create a shortcut to `C:\Novell\NDK\nwSDK\tools\VDB560.EXE` and name it **NetWare Virtual Debugger**.
3. Launch `VDB560.EXE`.
4. Enter `C:\COREDUMP\<yourservername>\` and the name of the core dump image created in step 17 of Part I.
The core dump image file is valid if it opens correctly.
5. Enter **Q** to quit.
6. Select **Y** to exit to DOS.

(End of Exercise)

Objective 4 Troubleshoot and Resolve Server Communication Issues

In this objective, you become familiar with the following methods to resolve communication issues:

- [Resolve Server-to-Server Communication Problems](#)
- [Resolve Workstation-to-Server Communication Problems](#)
- [Identify Preventative Maintenance Tasks](#)

Resolve Server-to-Server Communication Problems

Communication problems among servers can generate messages such as Unable to Communicate with Server or messages that a Directory partition operation can't be completed.

If DSTRACE is on, -625 errors appear in the DSTRACE output because the Directory cannot be synchronized among the servers.

Do the following to troubleshoot server-to-server communication problems:

- Check the server to make sure it is up and that an abend has not occurred.
- For IP networks, load PING.NLM and verify that each server can ping the other. (You need to know the IP address of each server.)
- For IPX networks, load IPXPING.NLM and verify that each server can ping the other. (You need to know the network and node of each server.)
- If the initial ping is unsuccessful, narrow the problem by pinging all devices in the path to the server, such as routers.
- Verify that time is synchronized on the network. Enter TIME at the system console or use DSREPAIR.NLM.
- Verify that all IPX internal network numbers and IP addresses are unique.
- If all servers in a replica ring cannot communicate with one server in the ring, it might be an address conflict. Enter CONFIG at the console prompt for IP and IPX numbers bound to the NIC.
- Make sure all network board drivers are loaded and protocols are bound. Enter CONFIG at the console prompt.
- For IPX networks, reset the server's routing table by entering RESET ROUTER at the console prompt.

- Check the server to make sure the CPU use is not remaining at or close to 100%.
- Verify that DSREPAIR.NLM is not loaded and locking the eDirectory database.

Resolve Workstation-to-Server Communication Problems

Prior to NetWare Client 32, a common error message when trying to log in to a server was “File server not found.”

With Client 32 and Windows 95/98, the login window appears only if a file server is located. If you have `NWEnableLogging=True`, then NIOS.LOG will show `A Server Could Not Be Found`.

With Client 32 and Windows NT/2000, the message “The tree or server cannot be found. Choose a different tree or server” appears after a timeout period.

Do the following to troubleshoot a *file server not found* problem:

- For IP networks:
 - Check the server IP address with CONFIG.
 - Check the workstation IP address with IPCONFIG or WINIPCFG.
 - Ping the server from the workstation.
 - Ping the workstation from the server.
 - Use TRACERT.
 - Ping devices between workstation and server, from closest to farthest.
- Verify there are enough user licenses available for login.
- Check network cabling.

Identify Preventative Maintenance Tasks

You can do the following to prevent problems before they occur:

- [Monitor Servers](#)
- [Eliminate a Potential Bottleneck](#)
- [Document the Network](#)
- [Perform Proactive Tasks](#)

Monitor Servers

It is important to know the normal behavior of your servers and network so when a problem occurs you have a baseline to help you troubleshoot the problem.

Understanding the normal behavior of your servers also lets you optimize the servers for your environment.

Eliminate a Potential Bottleneck

A system is only as stable as its weakest point. You can prevent problems by isolating a potential bottleneck and taking actions to fix it.

If a network board is a potential bottleneck, a solution could be to upgrade it to one with a wider bus.



For more information, see *Isolating the Real Bottleneck in a System* at <http://developer.novell.com/research/appnotes/1996/january/04/index.htm>.

Document the Network

Documenting the network is essential in helping solve future problems.

Use worksheets or applications that let you define your network topology to document network components.

Also, keep a network history to help you identify previous problems and normal network operation statistics.

Keeping records of network layout, hardware and software inventory, configuration, repairs, and backup schedules will save you time and work when problems occur.

For example, if you need to rebuild or replace parts of the network, justify new equipment, or restore the network after a disaster, documentation can be an invaluable resource.

You also need much of this information if you call Novell Technical Support. You might want to keep documentation in a notebook or online in a database.

Demonstrate accessing the Server Personal Log Book.

Using NetWare Remote Manager, you can keep a Server Personal Log Book at SYS:SYSTEM/NRMUSERS.LOG.

Select Reports / Log Files > Server Personal Log Book to enter and track changes made to the server or to log information you want to keep to track server performance or history.



For more information, see Section 1, Objective 3 Document Your Network of *Novell Network Management - Course 3004*.

Perform Proactive Tasks

You can often avoid problems by doing the following.

- Make sure the server has enough memory for peak usage.
- Replace hardware on a regular basis.

For example, power supplies can cause a lot of damage when they go out. Monitoring your power supplies and replacing them as needed can prevent greater damage and down time.

- Prevent static electricity problems.

An electrostatic discharge (ESD) must equal about 3000 volts before you can feel it, but sensitive electronic components such as microchips and circuit boards can be damaged by ESDs of as little as 20 or 30 volts.

These small discharges might not cause a component to fail immediately, but can cause the component to degrade over time and fail at a later date.

- Perform the following maintenance tasks on a regular basis:
 - Back up server data.
 - Check server error logs.
 - Check disk drives and controllers and monitor statistics.
 - Review server cache statistics.
 - Check for sufficient free space on server disks and volumes.
 - Test uninterruptible power supplies.
 - Update your network documentation after any change.

**20 Minutes**

Exercise 3-5 **Resolve Communication Problems**

Complete this exercise as a group.

Set up the problem at the beginning of each part.

Present the problem for discussion and let the students come to a conclusion. Guide them as needed and then discuss a solution.

Setup: Simulate a -625 error by removing the network cable from DA2 and show the console screen on the overhead.

Server-to-server communication is not working.

See “Resolve Server-to-Server Communication Problems” on 3-71.

Solution: In this case, a disconnected network cable on the server. Reconnect the cable on DA2.

In this exercise, the instructor will introduce a problem and you will need to troubleshoot the following as a group:

- [Part I: The Server Is Showing Errors](#)
- [Part II: The Workstation Cannot Login](#)
- [Part III: The Admin Account Has Been Deleted](#)
- [Part IV: Workstation Connections Are Dropping Periodically](#)
- [Part V: You Cannot Map Drives](#)

Part I: The Server Is Showing Errors

Your server is showing an error message Unable to Communicate with Server .DA2.DIGITALAIR.

When you run a DSREPAIR you are getting -625 errors for server DA2.

1. What does this error message indicate?
2. List the items you can check to locate the source of the problem.
3. What caused the error?

Setup: Simulate a login problem by using the PROBLEMS.NLM to dismount volume SYS on DA2.

Guide the students using standard troubleshooting methods.

Solution: Volume SYS is not mounted mount the volume to resolve the drive mapping issue.

Solution: Students should try searching Knowledgebase and online documentation.

Because this is an exclusive container administrator object, they need to restore eDirectory from backup or open a support call with Novell Technical Support (NTS).

NTS has a utility that will create another Admin object in the container.

Part II: The Workstation Cannot Login

Users are reporting that they cannot log in to server DA2. They can attach to it, but drive mappings fail.

1. Add the drive mapping MAP K:=\\DA2\SYS\ to your administrator login script.
2. Log in to your admin account.
3. What is causing the problem?

Part III: The Admin Account Has Been Deleted

The exclusive container administrator for .MARKETING.DEL.DIGITALAIR deleted the Admin object by accident during routine user object maintenance.

The administrator can no longer log in or administer objects in the container.

1. How will you resolve the problem?

Part IV: Workstation Connections Are Dropping Periodically

Scenario: Workstations connected to server DA2 are randomly losing connections to the server.

This NetWare server has been in use for over 3 years.

You have checked the router and cabling and everything seems to be functioning properly.

Guide the students using standard troubleshooting methods. If needed, provide them with hints that will lead them to the suspect, an outdated network board driver.

Solution: Update the LAN driver to the latest version by downloading it from the manufacturer's web site.

1. What is the most likely source of the problem?
2. What can you do to correct the problem?

Part V: You Cannot Map Drives

Users cannot map a drive to \\DA2\DATA\.

1. What troubleshooting steps would you take to resolve the problem?

Setup: Dismount the volume to simulate the problem. Guide the students as necessary.

Solution: Mount the volume.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	Summary
1. Identify Server Hardware and Operating System Components	<p>Quick and effective troubleshooting requires familiarity with server hardware and software.</p> <ul style="list-style-type: none">■ Identify Server Hardware<ul style="list-style-type: none">■ Bus types■ Mass storage■ Processor capacity■ Memory capacity■ Scalability■ Failure recovery■ Identify Operating System Components<ul style="list-style-type: none">■ NetWare load order■ Kernel■ Threads■ The run queue■ Multithreading■ Multitasking■ Multiprocessing■ Processor load balancing■ Pre-emption

Objective	Summary
2. Troubleshoot and Resolve NetWare Server Issues	<p>Problems and solutions vary from one server environment to the next.</p> <ul style="list-style-type: none">■ Identify the Top Novell Technical Support Server Issues and How to Resolve Them using the Novell Technical Support Knowledgebase. <p>Become familiar with symptoms and solutions for the following:</p> <ul style="list-style-type: none">■ Identify Problems after Installation■ Resolve Console Lock Ups■ Resolve Hard Disk Errors and Access Problems■ Resolve Application Monopolizing Server CPU■ Resolve Server Memory Problems■ Resolve Slow Server Response■ Identify Multiprocessing Problems■ Find Tools for Managing Servers

Objective	Summary
3. Troubleshoot and Resolve Critical Server Abends	<p data-bbox="954 285 1455 342">Critical server abends can be resolved using the tools and services provided by Novell.</p> <ul data-bbox="954 363 1455 951" style="list-style-type: none"><li data-bbox="954 363 1455 457">■ What an Abend Is An ABnormal END to a software routine that halts the server.<li data-bbox="954 468 1455 562">■ What an ABEND.LOG File Is A text file created at the time of an abend that contains information about the abend.<li data-bbox="954 573 1455 688">■ What a Core Dump Is A snapshot of a server's memory saved to a file. It is used to analyze the cause of an abend.<li data-bbox="954 699 1455 793">■ How to Respond to an Abend You can respond automatically or manually to an abend based on set parameters.<li data-bbox="954 804 1455 951">■ How to Create and Submit a Core Dump for Analysis Use the NetWare Debugger, NETDB, and Virtual Debugger to create and validate a core dump.
4. Troubleshoot and Resolve Server Communication Issues	<p data-bbox="954 978 1455 1014">Troubleshooting methods are used to</p> <ul data-bbox="954 1024 1455 1186" style="list-style-type: none"><li data-bbox="954 1024 1455 1081">■ Resolve Server-to-Server Communication Problems<li data-bbox="954 1092 1455 1148">■ Resolve Workstation-to-Server Communication Problems<li data-bbox="954 1159 1455 1186">■ Identify Preventative Maintenance Tasks

Exercise Answers

Following are the exercise answers.

Exercise 3-1. Determine Hardware and Operating System Components

Part I: Review Disk Controller Statistics and Establish a Usage Baseline.

3. Review the Current and Peak statistics and record the values in the following:

Answers will vary.

Part II: Review the Network Board Driver Statistics and Version.

9. Record the driver filename and version.

Answers will vary.

Exercise 3-2. Resolve Server Problems

Part I: Resolve a Server Memory Leak

6. Record the percentage of NLM Memory:

Answers will vary.

13. Record the percentage of NLM Memory:

Answers will vary; however, this percentage should be greater than the amount recorded in step 6.

Part IV: Resolve Open File Access Problem

5. From File Lock Information by Connection, record the connection numbers of the users listed:

Answers will vary. There should be 4 connections. PROBLEMS.NLM will create 4 connections for Admin and Supervisor.

Exercise 3-3. Create a Core Dump

Part II: Save a Core Dump Image to a Remote Drive

17. Record the filename of the core dump image for later reference:

Answers will vary. The format of the filename will be a hexadecimal value with an IMG extension (for example, 1021213AA.IMG).

Exercise 3-5. Resolve Communication Problems

Part I: The Server Is Showing Errors

1. What does this error message indicate?

Server-to-server communication is not working.

2. List the items you can check to locate the source of the problem.

See “Resolve Server-to-Server Communication Problems” on 3-71.

3. What caused the error?

In this case, a disconnected network cable on the server. Reconnect the cable to re-establish communication.

Part II: The Workstation Cannot Login

3. What is causing the problem?

Volume SYS is not mounted. Mount the volume to resolve the drive mapping issue.

Part III: The Admin Account Has Been Deleted

1. How will you resolve the problem?

Search the Knowledgebase and online documentation.

Because this is an exclusive container administrator object, you need to restore eDirectory from backup or open a support call with Novell Technical Support (NTS). NTS has a utility that can create another admin object in the container.

Part IV: Workstation Connections Are Dropping Periodically

1. What is the most likely source of the problem?

An outdated network board driver.

2. What can you do to correct the problem?

Update the LAN driver to the latest version by downloading it from the manufacturer's web site.

Part V: You Cannot Map Drives

1. What troubleshooting steps would you take to resolve the problem?

Mount the volume.

SECTION 4 Monitor and Troubleshoot eDirectory

Duration: 3 hours

In this section you learn how to monitor and troubleshoot eDirectory.

Objectives

1. Identify eDirectory Databases and Processes
2. Identify eDirectory Troubleshooting Steps
3. Identify Partition and Replication Placement Design
4. Use iMonitor Reports to Obtain Server and eDirectory Information
5. Perform Health Checks

Introduction

When you make changes to eDirectory, the changes are replicated throughout your eDirectory tree. The size of your eDirectory tree, the number of servers, and the number and replicas you have, determine the time it takes to distribute these changes.

Monitoring eDirectory regularly will help you identify when problems begin rather than when they might cause a system failure.

With eDirectory, you can prevent system failures if you troubleshoot the problem at the onset.

You should perform regular status reports and health checks on your system to ensure it is trouble free.

Objective 1 Identify eDirectory Databases and Processes

To help you monitor and troubleshoot eDirectory, you need to understand the following:


- [eDirectory 8.7 Databases](#)
- [eDirectory Processes](#)
- [Post-Migration or Upgrade Issues That Affect eDirectory Databases and Processes](#)

eDirectory 8.7 Databases

eDirectory uses a highly-scalable indexed database, called the FLAIM database, instead of a fixed-length record data store. It uses log files to back out and roll forward transactions in the event of a system failure.

The files that comprise the eDirectory database are shown in the following:

Figure 4-1 (slide)



The following describes each eDirectory database file:

- NDS.DB is the control file for the database. This file contains the roll-back log, used to abort incomplete transactions.
- NDS*.LOG tracks transactions that have not completed. eDirectory uses this file as a roll-forward log to reapply completed transactions that might not have been fully written to disk because of a system interruption.
- NDS.01 contains all records and indexes found on the server. When this file reaches 2 GB, NDS.02 is created for the remaining data. New files are created as necessary to keep database files from growing beyond 2 GB.

Limiting NDS.x files to 2 GB allows the database to remain scalable yet quickly accessible.

A number of indexes are maintained in the NDS.01 file to enhance performance:

- Attribute substring indexes for the CN and uniqueID fields
- Attribute indexes for the Object Class and dc fields
- Attribute indexes for positioning that include strings beginning with CN, uniqueID, Given Name, and Surname
- Streams files are named with hexadecimal characters (0–9, A–F) and hold information such as print job configurations and login scripts. Stream files have an NDS extension.

eDirectory Processes

When you make changes to eDirectory, you should make sure that the following processes are complete before you make additional changes:

- **Time synchronization.** Time synchronization is very important to eDirectory. All servers in a tree must be synchronized to the same time source. If they are not, collisions will occur when objects are being synchronized in replicas.

Synchronizing time across the network lets you maintain consistent time stamps.

The most common time stamp problem is with synthetic time. Synthetic time occurs when an eDirectory object has a modification time stamp ahead of current network time.

If the period between current time and the synthetic time is small this problem will correct itself. However, if the period is large you might need to resolve the problem manually.

- **Schema synchronization.** Schema synchronization ensures that the schema is consistent across the partitions in the eDirectory tree and that all schema changes are updated across the network.

In Section 1 you used a DSTRACE command, **SET DSTRACE=+SCHEMA**, to force schema synchronization to occur:

```
SET DSTRACE=ON
SET DSTRACE=+SCHEMA
SET DSTRACE=*S
```

By default, this process runs once every 4 hours.

Schema synchronization is required whenever a change is made to the schema. For example, when you add a new class, modify an attribute definition, or delete a schema definition, this information must be distributed.



Schema changes don't happen frequently unless you are constantly manipulating your eDirectory tree.

Servers might not receive schema updates as they occur due to communication problems or time stamp synchronization issues. When this occurs, a server will report that it did not receive schema updates properly.

- **Replica synchronization.** Replica synchronization refers to the process of copying data among the replicas of a partition. A partition is synchronized if all its replicas contain the same information.

If one replica has a more current version of a piece of data than the other replicas, it propagates this data to the other replicas.

Receiving "0" errors for replica synchronization indicates that your Directory is healthy.

Post-Migration or Upgrade Issues That Affect eDirectory Databases and Processes

You should monitor the following issues after you migrate or perform an upgrade:

- **Replica Placement.** Depending on the location of new servers in the tree, you should evaluate replica placement.

As a general guideline, you should have 3 - 5 replicas of a given partition as a minimum. The key to replica placement is based on accessibility: replicas should be placed where they will be accessed most.

As a word of caution, it possible to have too many replicas. Each time you make a change to an object, that change needs to be distributed to all replicas.

- **Unknown Objects.** eDirectory 8.7 includes new object and schema extensions. For example, new schema definitions in eDirectory 8.7 include ndsStatus Repair, ndsAgentPassword, and replication filter.

When you upgrade to eDirectory 8.7, this might cause unknown objects to appear in eDirectory.

There are 2 types of unknown objects. Each one shows a different icon in eDirectory:

- An unknown object with a question mark in a circle (the most common) indicates that the object is not recognized by eDirectory.
- An unknown object with a question mark in a square indicates that the object is recognized, but eDirectory does not have the tools to handle it.

The main reason an existing object will become unknown is that the object no longer matches its schema definition.

Unknown objects can happen for several reasons:

- The appropriate ConsoleOne snap-ins have not been installed
- The object has attribute values that are not defined by the class definition in the schema.
- A mandatory attribute is not present.
- The object is present under a container object of a class not allowed in its containment list.
- The object's naming attribute is missing (even if it is optional).

An object will appear as unknown when an older version of NDS is receiving an object that has an auxiliary class added to it.



An auxiliary class is a set of properties (attributes) added to particular eDirectory object rather than to an entire class of objects.

For example, an email application could extend the schema of your eDirectory tree to include an E-Mail Properties auxiliary class and then extend individual objects with those properties as needed.

This happens because older versions didn't allow additional superclasses to be added, which define additional attributes to be part of the definition.

Thus an object that has auxiliary classes added and has those extra attributes populated now has attribute values present that are not allowed by the more limited definition on an older server.

Objective 2 Identify eDirectory Troubleshooting Steps

Ask students to discuss problems they have seen with eDirectory.

In addition to identifying eDirectory databases and processes, a primary means of managing eDirectory is to identify steps you should take to troubleshoot eDirectory problems.

However, unlike LAN and server problems, the key to troubleshooting eDirectory is to be patient.

eDirectory uses its database processes to verify, validate, and distribute data. Most often, if given time, eDirectory will correct itself.

Ask students if they use similar troubleshooting steps to resolve the problems they experience.

If eDirectory does not correct itself, then you should use most, if not all, of the following eDirectory troubleshooting steps:

- Step 1: Identify the Scope of the Problem
- Step 2: Determine the Cause of the Problem
- Step 3: List Possible Solutions to the Problem
- Step 4: Assess Possible Solutions
- Step 5: Implement a Solution
- Step 6: Verify That the Problem Is Resolved
- Step 7: Document the Resolution to the Problem
- Step 8: Avoid Repeating the Problem

Although these are effective troubleshooting steps, keep in mind that this is only one way to troubleshoot. You might use another set of steps. The key is that you have a procedure that you consistently follow.

Step 1: Identify the Scope of the Problem

Before you can resolve a problem, you need to know the extent of the problem.

Symptoms of eDirectory problems can show themselves in many ways. They can appear on their own as error or warning messages on the server console or within utilities, or they can appear when a user attempts to log in but cannot authenticate.

The following are examples of eDirectory problems:

- Time synchronization issues
- Synchronization issues
- eDirectory version problems
- Communication problems
- Improperly moved/removed servers
- Inconsistent object/database
- Agent process errors
- Performance issues

eDirectory problems can appear as error codes when you attempt to perform an eDirectory operation or while performing a proactive eDirectory health check.

To help you assess the scope of an eDirectory problem, record the following information:

- The symptom (what happened that tells you there might be a problem)
- The eDirectory error number
- The partition, replica, server, or object with the error
- The servers holding the partition, replica, server, or object with the error

Show students the Error Messages chapter in the NetWare 6 documentation.

Don't demonstrate using the server log at this time. Discuss it during the troubleshooting steps.

If time permits, show the HEALTH.LOG file on the instructor server. Tell students that the default settings are acceptable but can be changed using Remote Manager.

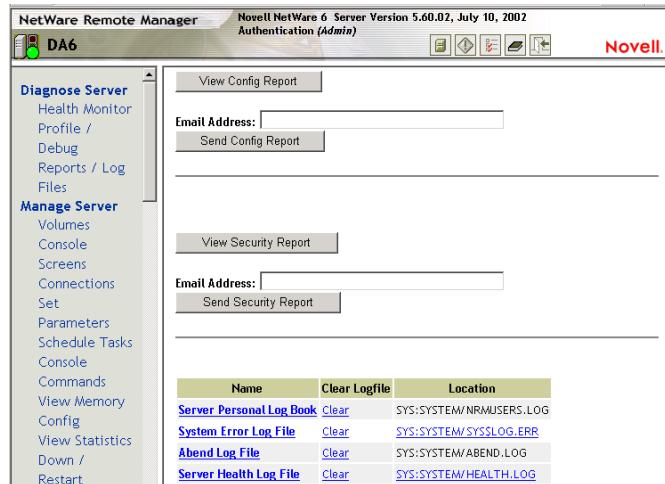
- Identify additional errors found in the following log files:
 - **Server Personal Log SYS:SYSTEM/NRMUSERS.LOG (new to NetWare 6).** Use to enter and track changes made to the server or to log any other information to track server performance.
 - **System Error Log SYS:SYSTEM/SYS\$LOG.ERR.** Use to view alert messages and system operation information sent to the System Console and Logger screens.
 - **Abend Log SYS:SYSTEM/ABEND.LOG.** Use to view information about all server abends.
 - **Server Health Log SYS:SYSTEM/HEALTH.LOG (new to NetWare 6).** Use to view all changes in the health status of the server. Using this report along with several of the health statistics and trend reports can be useful in diagnosing server problems.

Each of these log files starts automatically when you start your server, and can help you address problems.

Demonstrate accessing these log files using Novell Remote Manager > Diagnose Server.

Figure 4-2

The following shows how Novell Remote Manager provides easy access to the logs:



Step 2: Determine the Cause of the Problem

After identifying the scope of the problem, you need to determine the cause of the problem.

Usually eDirectory reports errors or has problems because a condition exists that prevents an eDirectory process from completing properly.

Not all eDirectory issues are obvious, and an eDirectory error code might only be a symptom of the real issue.

To determine the cause of the problem, do the following:

- Determine whether the problem is an eDirectory problem or something else, such as an unattached cable.
- Analyze the information gathered about your eDirectory problem carefully.

- Determine what eDirectory process is having a problem and why.
- Use eDirectory research resources to provide insight into what might be the real problem.

After you determine what process is having the problem and why, you can formulate a solution to the problem.

Step 3: List Possible Solutions to the Problem

There are often several ways to resolve an eDirectory problem. You might even find multiple support documents recommending different solutions to the same eDirectory error. eDirectory utilities also might suggest multiple ways to solve an eDirectory error.

Before taking any action to resolve an eDirectory error, you should

- Gather and list possible solutions.

Access Knowledgebase at <http://support.novell.com> to gather information regarding your problem and possible solutions to the problem.

- List the repercussions of each action.

Listing the repercussions of each action is critical because the problem can often be made worse if the wrong solution is applied, or a solution can generate other problems.

For example, if the first solution you try is to remove and then reinstall eDirectory, you will probably lose some data.



Removing and reinstalling eDirectory should be used as a last resort, after all other possible solutions have been evaluated.

Step 4: Assess Possible Solutions

After you list possible solutions to your eDirectory problem, assess the solutions based on the following:

- The likelihood that it will solve your problem
- How easy or hard the solution is to implement
- What effect the implementation process will have on users
- Whether the solution will have a destructive impact on the eDirectory tree

To assess possible solutions, you must have a clear understanding of eDirectory, the problem you are facing, the solution, and the ramifications of applying the solution.

Assessing the possible solutions might involve deciding between contradicting solutions from various sources.

Ask coworkers and others what actions they would take. If possible, test the solutions in a lab environment before implementing them.

Step 5: Implement a Solution

After determining the solution, create backups and implement the solution.

The tools used to implement a solution are often the same tools used to diagnose the problem.

Allow enough time for your actions to resolve the problem. eDirectory takes time to synchronize changes throughout the network. So even after you fix a problem, you might still see symptoms of a problem until eDirectory synchronizes.

Step 6: Verify That the Problem Is Resolved

After the actions you took to resolve an eDirectory issue have been processed, verify that the problem is resolved.

To verify that an eDirectory problem has been resolved, you should

- Use the eDirectory diagnostic tools to check the status of eDirectory.
- Attempt to repeat the actions that revealed the eDirectory problem.
 - If the eDirectory problem revealed itself while you were attempting an eDirectory operation, attempt that operation again to see if it can be successfully done.
 - If the problem was revealed through an eDirectory utility, run the same utility and see if the problem still exists.
- Continue to monitor eDirectory.

Step 7: Document the Resolution to the Problem

After you identify the eDirectory problem and solution, document them in your server maintenance logs. If you don't use a server maintenance log, now is a good time to start.

Novell Remote Manager has a user friendly interface to the NRMUSERS.LOG file, which is the server log.

NRMUSERS.LOG allows you to

- Prevent the same problem in the future
- Find a resolution to the same problem quickly in the future
- Provide insight into other problems your network might have

The idea behind documenting the resolution to a problem is that you should solve a problem only once.

Demonstrate adding information to the server log as a documentation procedure.

Although steps 7 & 8 are similar, the purpose of step 8 is to reinforce the benefits of documenting the problem.

Step 8: Avoid Repeating the Problem

The final step in the process is to avoid repeating the problem.

You might find that an eDirectory problem occurs by an inadvertent or inappropriate action or that no proper procedure has been defined to perform the action properly.

Whatever the cause of the problem, you want to make sure it does not get repeated.

You should

- Document the problem and the solution.
- Establish procedures and policies to ensure that people who administer or use the eDirectory tree will do so in a consistent and established manner.
- Take precautions, such as restricting access to servers.

Objective 3 Identify Partition and Replication Placement Design

Assigning the placement of eDirectory partitions and replicas is crucial to maintaining a healthy and stable network environment.

In the Introduction section of this course, the scenario describes the Digital Airlines network as one that continues to grow by adding and upgrading servers and applications.

In a production environment, we strongly recommend that you maintain log files identifying the following:

- Server hardware specifications
- Server software specifications
- eDirectory partitions and replica placement

- Software upgrade dates, times, and issues
- Hardware upgrade dates, times, and issues
- Problems and resolutions

As you review these log files, you question if partitions and replica placement meet the needs of the merged organizations.

The following shows the master partitions on DA1. A partition has been created for each city container, and the master replica for the partition resides on DA1.

Figure 4-3

Agent Synchronization Summary					
Replica Type	Partitions	Errors	Oldest Successful Sync	Max. File Delta	
Master	8	8	201-10-36	2:18:00:06	

Partition Synchronization Status					
Partition	Errors	Last Successful Sync.	Maximum File Delta	Replica's Periodic Data Delta	
.DIGITALAIR-TREE	1	194-49:59	2:18:00:06	0:11:26	Replica Synchronization , Agent Health , Change Cache , Continuity
.arg1.DIGITALAIR-TREE	1	168-48:53	1:68:45:51	1:58:26:34	Replica Synchronization , Agent Health , Change Cache , Continuity
.DEL.DIGITALAIR.DIGITALAIR-TREE	1	201-01:19	2:00:59:44	1:72:45:21	Replica Synchronization , Agent Health , Change Cache , Continuity
.TYO.DIGITALAIR.DIGITALAIR-TREE	1	200-32:19	2:00:29:00	0:22:59	Replica Synchronization , Agent Health , Change Cache , Continuity
.TOL.DIGITALAIR.DIGITALAIR-TREE	1	200-43:27	2:00:42:53	0:36:44	Replica Synchronization , Agent Health , Change Cache , Continuity
.SYD.DIGITALAIR.DIGITALAIR-TREE	1	200-56:54	2:00:53:41	2:16:53	Replica Synchronization , Agent Health , Change Cache , Continuity
.LON.DIGITALAIR.DIGITALAIR-TREE	1	200-59:41	2:00:59:04	38:51:26	Replica Synchronization , Agent Health , Change Cache , Continuity
.LGA.DIGITALAIR.DIGITALAIR-TREE	1	201-10:36	2:01:09:38	2:15:51	Replica Synchronization , Agent Health , Change Cache , Continuity

As network administrator for Digital Airlines, you've decided that you would like to move these master replicas to the local servers at the branch offices.

A new tool in iManger lets you manage partition and replica placement.



15 minutes

Exercise 4-1 Adding Replicas with iManager

In this exercise, you create a replica and change the replica type to better suit your needs.

Complete the following:

1. From the browser, enter the IP address for your server **https://192.168.1.x:2200** (where *x* = *your server number*)
2. From NetWare Web Manager, select **DAx** under the **iManager** heading.
3. Use the following to verify the context for the container where your server resides.

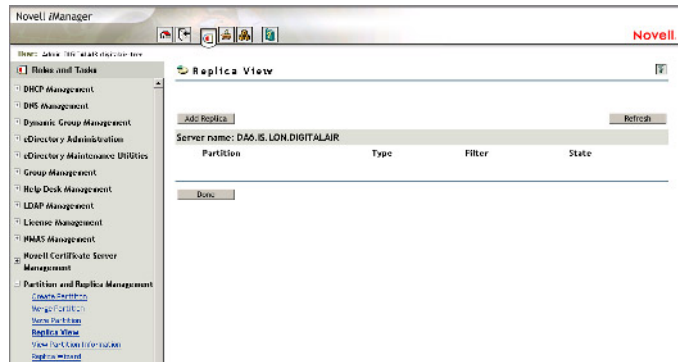
Table 4-1

Server Name	Location
DA4	IS.DEL.DIGITALAIR
DA5	IS.LGA.DIGITALAIR
DA6	IS.LON.DIGITALAIR
DA7	IS.SYD.DIGITALAIR
DA8	IS.TXL.DIGITALAIR
DA9	IS.TYO.DIGITALAIR

4. Enter your userid and password; then select **Login**.
5. From the list of Roles and Tasks, expand **Partition and Replication Management**.
6. Select **Replica View**.
7. Using the **browse** button, locate and select your server in the tree.
For server name and location see table 4-1.
8. Select **OK**.

The following appears:

Figure 4-4

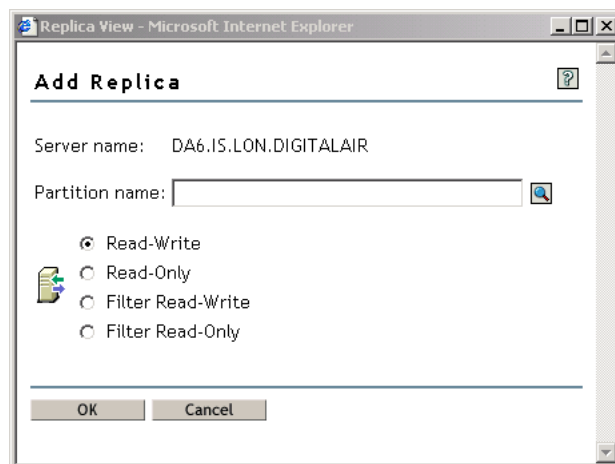


Notice that your server does not contain a replica.

9. Select **Add Replica**.

The following appears:

Figure 4-5

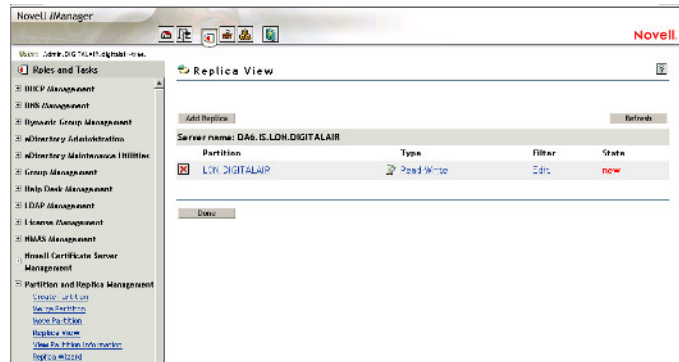


10. Verify your server name; then use the **browse** button to locate and select your location container (such as DEL.DIGITALAIR).

11. Specify that this will be a **Read/Write** replica; then select **OK**.
12. Select **Refresh** to show the replica status.

Your screen should look similar to the following:

Figure 4-6



Now that you have a replica at your location, consider the replica type. Do you want a Read/Write replica at your location?

Depending on the type of administration done at your site, a Read/Write replica might be sufficient. In this case, continue with steps 13–16 to change the Read/Write replica to a Master replica.

13. Select **Refresh** a second time to verify that the replica state is on.
14. Select the **Read/Write** replica you created.

The replica view dialog appears, with the Master replica type now available.
15. Select **Master**; then select **OK**.
16. Select **Refresh** to view the replica type change.
17. Verify the replicas assigned to your location container:
 - a. From the menu items (on the left), select **Partition and Replica Management > Replica View**.

- b. On the right, enter the context for your location container (such as **DEL.DIGITALAIR**); then select **OK**.

A list of the replicas assigned to your container appears.

- c. When you finish viewing the list, select **Done** (below the list).

(End of Exercise)

Objective 4 Use iMonitor Reports to Obtain Server and eDirectory Information

In Course 3004, reports were introduced but not explained in detail.

Select **Reports > Report Configuration** to introduce the preconfigured reports.

Tell students they will use some of these reports to verify their system in an exercise.

iMonitor reporting is a powerful tool for determining the condition of a given server, agent process, or tree.

To access iMonitor, launch your browser and in the Address field, enter **https://your server IP address:8009/nds-summary**.

From the iMonitor page, you can select the Reports option.

In the Agent Summary screen, the status of your server is shown by the icon at the left of the server address. The server status is represented by the color of the signal light: green means the server is functioning properly, yellow indicates a problem, and red signifies that communication cannot happen.

You can also run reports to determine the status of the server and eDirectory.

To use these reports effectively, you need to understand the following:

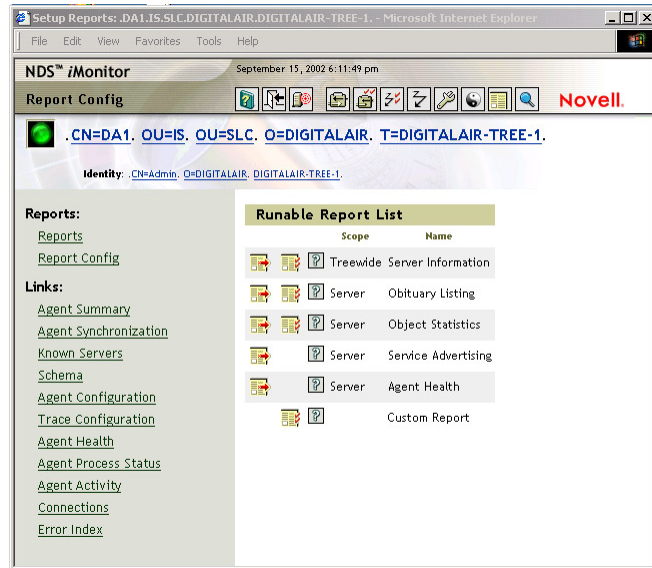
- [How to Review Report Options](#)
- [How to Run a Report](#)
- [How to View Saved Reports](#)

How to Review Report Options

Demonstrate running reports. Tell students that they will run a report in the next exercise.

When you select the Report icon you will notice that there are 2 Report options (Reports and Report Config), as shown in the following:

Figure 4-7



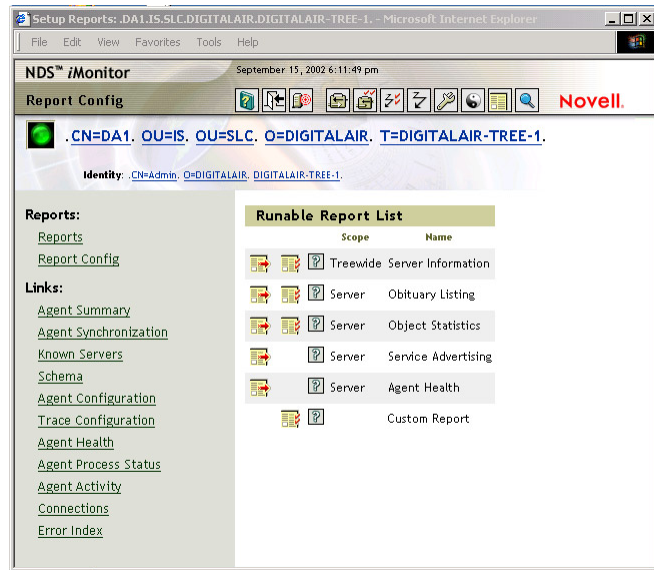
The following describes these 2 options:

- **Reports.** Shows the reports you have generated. This is the default setting.

The first time you use the Report feature, you see a message that no reports have been run for this server. After you generate your first report, each time you select the Reports option, the reports you have run will appear.

- **Report Config.** Shows the report list, as viewed in the following:

Figure 4-8



The following preconfigured reports are listed:

- **Server Information.** This report searches the entire tree, communicates with every Netware Core Protocol™ (NCP™) server it can find, and reports errors it finds.
 You can use this report to diagnose time synchronization and limber problems, to verify communication with all other servers from this server's perspective, or determine if a server has been improperly removed.
 If you select this report in the Configuration page, the server can also generate DS Agent Health information for every server in the tree.
- **Obituary Listing.** This report lists all obituaries on this server.

- ❑ **Object Statistics.** This report evaluates the objects in a given scope and then generates lists of objects matching the requested criteria.

The criteria can include such things as future time, unknown objects, renamed objects, counts of base classes, containers, aliases, and external references.

- ❑ **Service Advertising.** This report lists all directories and servers known to this server through SLP or SAP.
- ❑ **Agent Health.** This report gathers health information for this server.
- ❑ **Custom Report.** This report can create a customized report and scheduled events.

Scheduled events are reports that run when you configure them to run.

Your user identity and security settings are used to gather report data. The report data is then stored on the server from which the report was run.

When you schedule a report to run, it does not use your identity and security settings. It uses the identity of Public.

How to Run a Report

Complete the following:

1. From iMonitor, select the **Report** icon.
2. Select the **Report Configuration** link.

The list of preconfigured reports appears.

Use one of the following icons to access report options:

- ❑ **Red arrow.** Use to run the report.
- ❑ **Check boxes.** Use to configure report options.
- ❑ **Question mark.** Use to access report help.

Move your cursor over an icon and point out the various icons on the reports screen.

3. Select **Report Configuration** for the report you want to run.
In the following figure, the Obituary Listing report is selected.

Figure 4-9

The screenshot shows two configuration forms. The top form, titled "Configure Obituaries Report", includes a "Report Options" section with checkboxes for "Restored (primary)", "Dead (primary)", "Moved (primary)", "Inhibit move", "Backlink", "Old RDN", "New RDN", "Unknown", "Tree old RDN", "Tree new RDN", "Moved tree", "Moved from", "Used by", and "Object version". It also has a "Scope Options" section with a "Start at:" field set to "DIGITALAIR-TREE", radio buttons for "Base", "Subordinate", "Partition", and "Subtree", and a "Saved reports:" field set to "5". The bottom form, titled "Schedule Report", includes a "Scheduling Options" section with radio buttons for "One Time", "Hourly", "Daily", "Weekly", and "Monthly", and fields for "Start Time" (21:40), "Start Day" (Friday), and a month/year selector (November 8).

This report lists the obituaries that match the specified criteria.

4. Select **Run Report**.

The report appears as follows:

Figure 4-10

Report

Obituary: [\[+\]](#)

Obituary Info						
	Full Name	Entry ID	Flags	Modification Time	Creation Time	Obit Types
<input type="checkbox"/>	.D6415.DEL.DIGITALAIR.DIGITALAIR-TREE.	0000233F	Not Present	10-31-02 10:04:24 pm 1:9	10-30-02 9:02:18 pm 1:1	Dead (primary), Backlink, Used by
<input type="checkbox"/>	.D6415.IS.DEL.DIGITALAIR.DIGITALAIR-TREE.	00002341	Not Present	10-31-02 11:40:23 pm 2:71	10-30-02 9:02:18 pm 1:42	Dead (primary), Backlink, Used by
<input type="checkbox"/>	.D44_5YS.IS.DEL.DIGITALAIR.DIGITALAIR-TREE.	00002344	Not Present	10-31-02 11:40:23 pm 2:101	10-30-02 9:02:29 pm 2:1747	Dead (primary), Backlink, Used by
<input type="checkbox"/>	.ADMIN_DMS.IS.DEL.DIGITALAIR.DIGITALAIR-TREE.	00002345	Not Present	10-31-02 11:40:23 pm 2:118	10-30-02 9:02:29 pm 2:1763	Dead (primary), Backlink, Used by
<input type="checkbox"/>	.D64_0AT4.S.DEL.DIGITALAIR.DIGITALAIR-TREE.	00002346	Not Present	10-31-02 11:40:23 pm 2:125	10-30-02 9:02:29 pm 2:1819	Dead (primary), Backlink, Used by
<input type="checkbox"/>	.D44_5YS.PROL.KDEL.DIGITALAIR.DIGITALAIR-TREE.	00002348	Not Present	10-31-02 11:40:23 pm	10-30-02 9:02:29 pm	Dead (primary),

In this example, a number of obituaries appear.

5. Scroll to the right of the screen to view the obituary status.
6. (Conditional) For additional details, select the Entry ID link.

How to View Saved Reports

When you run a report, you can configure iMonitor to save a specified number of the same report. For example, the default setting for the Obituary Listing report is to keep the last 5 occurrences of the report.

To view the details of the report, do the following:

1. From iMonitor, select the **Report** icon.

The Report List shows the reports that have been generated and who initiated the report.

2. Select a report; then select **View Report** (the magnifying glass icon).
3. Review the results on your screen.

Depending on the report, the report might open to a results screen or you might see the link to a specific screen.



15 minutes

Discuss the questions in this exercise as a group to facilitate students comparing results and helping each other determine causes and solutions.

Exercise 4-2 Verify eDirectory Status Using Reports

As the system administrator for your Digital Airlines office, you have just upgraded a server from NetWare 4.11 (with NDS 6.11) to NetWare 6 (with eDirectory 8.7).

You need to confirm with Mark Bassil, the vice president of IT, that the upgrade is complete.

To verify that the schema has been properly updated and distributed to all partitions and replicas, run the Server Information report.

Do the following:

1. Access iMonitor using the IP address for your server **https://192.168.1.x:8009** (where *x* = your server number).
2. Enter your userid and password; then select **Login**.
3. Select the **NDS iMonitor** link.
4. Select the **Report** icon.
If this is the first report, no reports are shown.
5. Select **Report Config**.
6. From the list of runnable reports, select the **Configure Report** icon for **Server Information**.
7. Accept the default report options by selecting **Run Report**.

The Report results appear. This report is divided into 2 areas:

- Servers with warnings
- Servers that are functioning properly

8. Which servers are showing errors?

9. Describe the errors.

10. Which servers have recommended actions?

11. Why are these recommendations being made?

12. Select **Report Config**.
13. From the list of runnable reports, select the **Configure Report** icon for **Server Information**.
14. Select the **Try: IPX** option; then run the report a second time.
15. What are the current errors for any IPX servers in your network (such as DA3)? Why have they changed?

(End of Exercise)

Objective 5 Perform Health Checks

You should perform health checks on a regular basis and any time you make a change to NetWare or eDirectory.

To effectively perform an eDirectory health check, you need to understand the following:

- [Health Check Items](#)
- [iMonitor Health Check Features](#)
- [iMonitor Health Check Procedure](#)
- [How to Run the Agent Health Report](#)
- [How to Perform a Trace with iMonitor](#)
- [How to Perform Directory Service Repair](#)

Health Check Items

A complete health check includes checking for the following:

- **eDirectory revision.** If your version (or revision level) of eDirectory is outdated, download the latest software patch from <http://support.novell.com>.
- **Time synchronization.** Time stamps are assigned to each eDirectory object and property. They ensure the correct order for object and property updates.

Using time stamps, eDirectory determines which replicas need to be synchronized. For synchronization to happen properly, all eDirectory servers must maintain accurate time.

- **Partition continuity.** Partition continuity ensures that all replicas for a partition can be updated with Directory changes.
- **Background processes.** eDirectory changes are replicated in background processes. The following background process should be checked:

- **Schema synchronization status.** This process identifies the current condition of schema synchronization.
Schema synchronization ensures that the schema is consistent across the partitions in the eDirectory tree and that all schema changes are updated across the network.
- **Obituaries.** This process is based upon eDirectory ID numbers rather than object names. This process ensures that name collisions do not occur during certain operations.
Obituaries are attributes applied to an object. There are 11 obituary types. For example, there is an obituary for move, one for rename, and one for delete.
- **External references/distributed reference links (DRL).**
External references are place holders in eDirectory that contain information about entries the server does not hold.
For example, when a user browses the eDirectory tree and requests information about an entry that is not stored locally, eDirectory creates an external reference to the entry.
- **Limber status.** This process ensures that all server information (such as IP address and server name) is correct.
You can initiate the process in 2 ways:
 - To initiate the Limber request in iMonitor select **Agent Configuration > Agent Triggers > Limber > Submit.**
 - To initiate this process with DSTRACE, at the server enter
SET DSTRACE=ON
SET DSTRACE=+LIMBER
SET DSTRACE=*L

iMonitor Health Check Features

The reporting capabilities in iMonitor let you check the health of your server and eDirectory. In addition to running reports to verify the condition of your system, you can manually perform health checks.

In Section 1, you performed a health check prior to the upgrade using the DSTrace and DSRepair commands at the server. After the upgrade, you used these same tools to initiate the limber and backlinker processes to take place.

You can use iMonitor to perform the same tasks you might have performed using the following tools: DSTrace, DSRepair, DSBrowse, and NDS Manager.

As a general guideline, if your system is constantly changing, use iMonitor to verify the health status weekly. If your system has not had changes, you should verify the health status monthly.

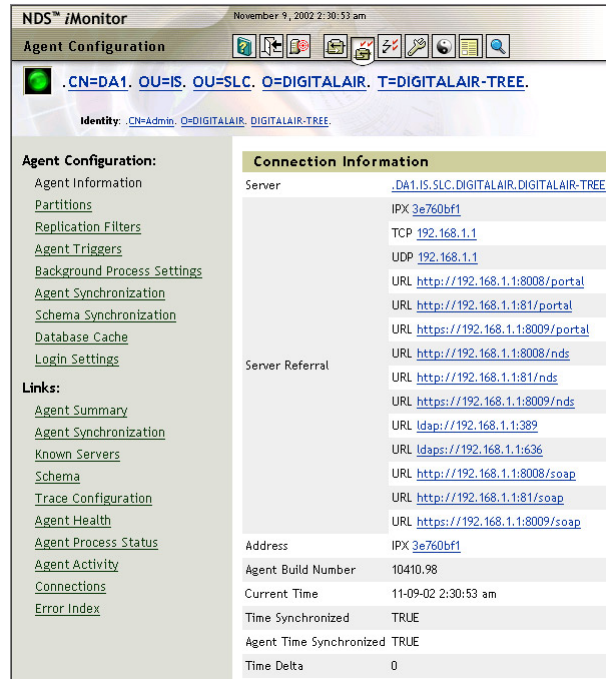
iMonitor Health Check Procedure

To use iMonitor to perform a health check, do the following:

1. From iMonitor, access the **Agent Configuration** link to obtain the following health check information:
 - Agent (DS) Version
 - Time Synchronization

You screen should look like the following:

Figure 4-11



Notice the agent build number and the time synchronization information.

2. From the list of links on the left, select **Agent Synchronization**; then verify partition synchronization information.

The following appears:

Figure 4-12

Agent Synchronization Summary					
Replica Type	Partitions	Errors	Oldest Successful Sync	Max. Ring Delta	
Master	8	8	201-10-36	218:00:06	

Partition Synchronization Status					
Partition	Errors	Last Successful Sync.	Maximum Ring Delta	Replica's Permissible Data Delta	
.DIGITALAIR-TREE	1	194:49:59	218:00:06	0:11:26	Replica Synchronization , Agent Health , Change Cache , Continuity
.org1.DIGITALAIR-TREE	1	168:48:53	168:48:51	198:24:34	Replica Synchronization , Agent Health , Change Cache , Continuity
.DEL.DIGITALAIR.DIGITALAIR-TREE	1	201:01:19	200:59:44	172:45:21	Replica Synchronization , Agent Health , Change Cache , Continuity
.TLO.DIGITALAIR.DIGITALAIR-TREE	1	200:32:19	200:29:00	0:22:59	Replica Synchronization , Agent Health , Change Cache , Continuity
.TLO.DIGITALAIR.DIGITALAIR-TREE	1	200:43:27	200:42:53	0:36:44	Replica Synchronization , Agent Health , Change Cache , Continuity
.SYD.DIGITALAIR.DIGITALAIR-TREE	1	200:56:54	200:53:41	2:16:53	Replica Synchronization , Agent Health , Change Cache , Continuity
.LON.DIGITALAIR.DIGITALAIR-TREE	1	200:59:41	200:59:04	38:51:26	Replica Synchronization , Agent Health , Change Cache , Continuity
.LOA.DIGITALAIR.DIGITALAIR-TREE	1	201:10:36	201:09:38	2:15:51	Replica Synchronization , Agent Health , Change Cache , Continuity

The following information is available:

- ❑ **Errors.** Shows the number of errors during synchronization.
- ❑ **Last Successful Sync.** Lists the amount of time since all replicas of the partition were successfully synchronized from this server.
- ❑ **Maximum Ring Delta.** Shows the amount of data that might not be successfully synchronized to all replicas in the ring.

For example, if Mark has changed his login script within the past 30 minutes, and the maximum ring delta has a 45-minute allocation, Mark's login might not be successfully synchronized, and he might get the previous login script when he attempts to log in.

However, if Mark changed his login script more than 45 minutes ago, he should get the new login script consistently from all replicas.

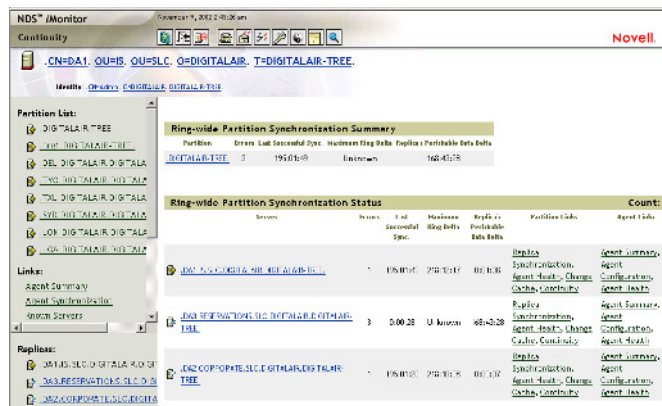
- **Replica’s Perishable Data Delta.** Lists the amount of time since this server has last synchronized that partition.

At the right of the partition information, notice the links to Replica Synchronization, Agent Health, Change Cache, and Continuity.

3. Select **Continuity**. This link lists every server that holds a replica of the partition.

The following appears:

Figure 4-13

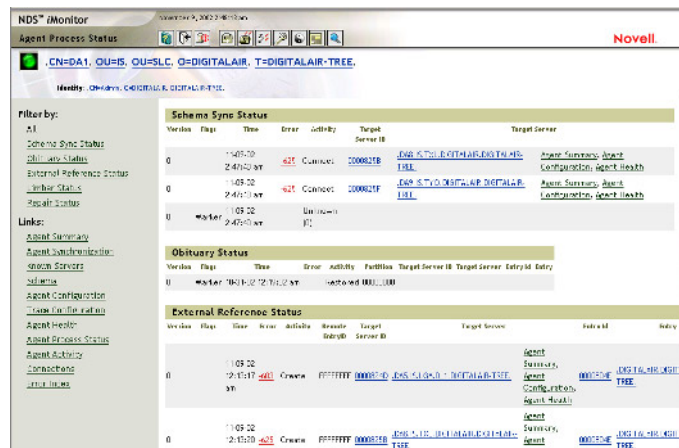


If a server’s replica ring differs from the replica ring maintained by other servers in this list, or if the server cannot participate in the replica synchronization process, an error message appears.

4. From the list of links on the left, select **Agent Process Status** (you might need to use the scroll bar to see the link).

The following appears:

Figure 4-14



This link lists the following background processes and their status:

- Schema synchronization
- Obituary processing
- External references/distributed reference links (DRL)
- Limber

Problems with these processes appear as an error code. To further identify the problem, you can run the Health report option.

How to Run the Agent Health Report

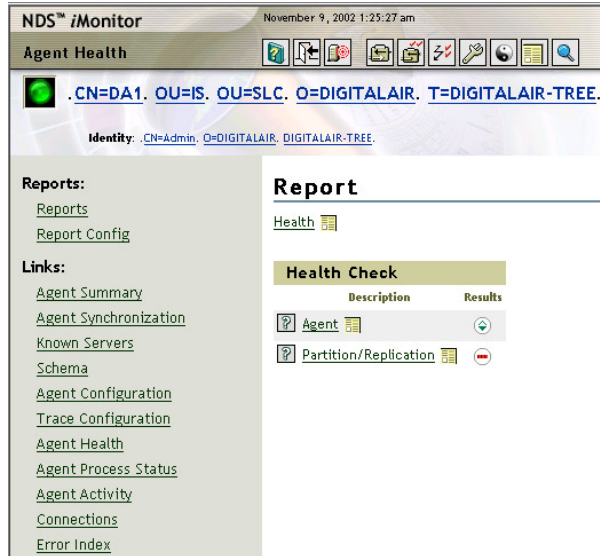
This option provides quicker access to problems with eDirectory. Do the following:

1. Select the **Report** icon.
2. Select **Report Config**.

3. From the list of runnable reports, select the option to run the **Agent Health** report.

The report results appear similar to the following:

Figure 4-15



In this example, you can immediately see the report results show problems with both the Agent and Partition/Replication status.

The Agent results icon indicates that there are marginal problems, and the Partition/Replication processes indicate a warning.

- When problems exist, you can select the **Health Check > Agent** link to see the reason for the marginal status, as seen in the following:

Figure 4-16



The Health Check Agent shows processes that are experiencing problems. In this example, 2 error codes appear:

- Agent Reference Check shows error 603.
- Agent Process Schema shows error 625.

Selecting the error code shows the help screen with a description of the error and solutions to resolve the errors.

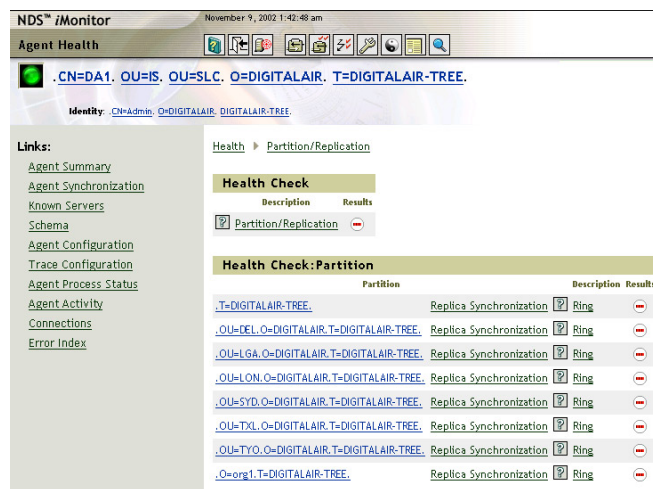
The 603 error lets you know that an attribute is not defined. The 625 error tells you that the communication process for the schema is not happening, most likely due to a LAN problem.

These errors are not critical at this time, but if left unattended, they could cause problems. To identify the full scope of the problem, check the partition and replication error you saw.

5. Use the **Back** button to return to the report status.
6. Select the Partition/Replication link.

If there are no errors, the Result area will show a green indicator. Because there are errors in this example, the following appears:

Figure 4-17



In this screen, each replica is showing a problem.

7. Select the Ring link.

This opens the details for each replica. You can see that replication is not happening, as shown in the following:

Figure 4-18

Replica	Description	Results	Current
	On		
	Last Successful Sync		168:13:20
	Last Attempt Sync		168:13:20
	Send Delta		168:13:14
	Receive Delta		
Local Replica	Issued Future Time		-3:03:05
	Purge Time		168:13:20
	Perishable Data		157:51:58
	On		
	Last Successful Sync		168:13:20
	Last Attempt Sync		0:00:05
	Send Delta		
	Receive Delta		168:13:11
	On		
	Last Successful Sync		168:13:20
	Last Attempt Sync		157:52:03
	Send Delta		168:11:14
	Receive Delta		157:51:58

8. Select the replica where the error is located.
9. Scroll down to the partition to view the error, as seen in the following:

Figure 4-19

Timestamp	Flags	Type	Synchronization Time	Error	Server	Entry
10-31-02 11:43:39 pm	Present, 1:1	Replica	10-31-02 11:43:37 pm	OK		DIGITALAIR.TREE
11-09-02 2:13:44 am	Present, 1:1	Replica	11-09-02 2:13:43 am		DS2.CORPORATE.SLC.DIGITALAIR.DIGITALAIR.TREE	DIGITALAIR.TREE
11-09-02 2:14:40 am	Present, 1:1	Replica	11-09-02 2:14:40 am	-694	DS2.RESERVATIONS.SLC.DIGITALAIR.DIGITALAIR.TREE	RESERVATIONS.SLC.DIGITALAIR.TREE

10. Select the error code again to view the details.

In this example, 694 is a lost entry. The eDirectory object being updated is using an eDirectory background process that has not been received.

In the problems shown in the previous figures, server DA3 is not accessible by DA1 and updates cannot be processed. To resolve this problem you must access DA3 and perform troubleshooting steps on that server.

Resolving problems is discussed later in this section.

iMonitor shows errors as the DSAgent runs agent processes. These processes, which run in the background at regular intervals, distribute the changes you make to eDirectory.

You can force a process to run ahead of its schedule interval to make these changes immediately.

To initiate a process ahead of schedule and to view the process as it occurs, you can use the Trace option.

How to Perform a Trace with iMonitor

The following are the high-level steps to perform a trace with iMonitor:

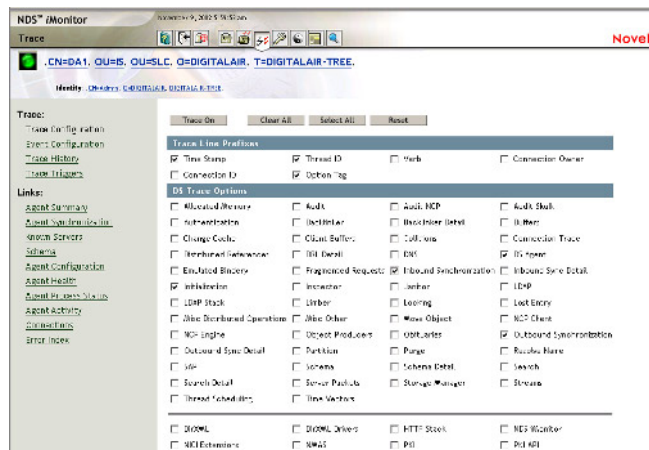
- Configure Trace Options
- Turn Trace On
- Select the Process to perform the trace

Use the following to specify the amount of information you want to obtain from a specific trace:

1. **Select Trace Configuration.**

The following appears:

Figure 4-20



Demonstrate running Trace.

In most cases, the default trace line prefixes will meet your needs.

2. Specify the DS Trace Options; then scroll to the bottom of the screen and select **Trace On**.

These settings are stored for this trace run.

After you specify the options, you must tell the agent what process to perform these options on.

3. Select **Agent Configuration** using the navigation button at the left of the Trace Configuration button.

After you select Trace On and select Agent Configuration, a new icon appears on the navigation bar. This is the Trace button, which you use to access all Trace options.

4. Select **Agent Triggers**; then specify the background process you want to trace.
5. Select **Submit** to begin the trace.
6. Select the **Trace** icon.



The line items that appear might not be sequential, making diagnosing messages very difficult. The Trace screen is a developer debugging tool and is not very intuitive. The iMonitor screens are more understandable than reading the Trace results.

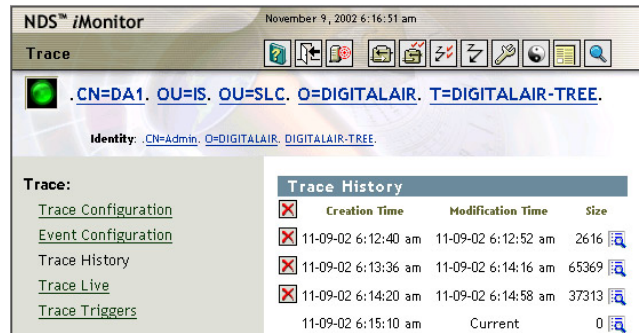
8. Select Trace Configuration > Trace Off.

Trace processing can generate a great deal of network traffic. After you complete your trace, turn the trace off.

A server must have a replica to show eDirectory trace information (except for the schema). The schema still replicates to all servers in the tree.

Each instance of Trace is stored in a log file. To review the results of a previous trace, you can select Trace History. The reports can be viewed, as seen in the following:

Figure 4-22



The X is a delete option. You can delete individual reports or delete all reports by selecting the delete icon at the top of the Trace History screen.

When you find problems in your eDirectory tree, you can initiate the Repair process to correct them.

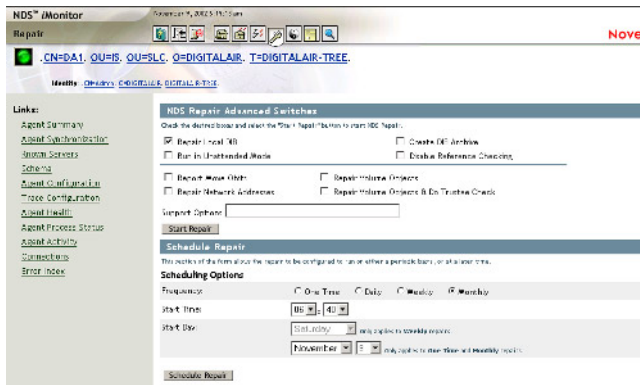
How to Perform Directory Service Repair

One of the most critical procedures you can perform to correct Directory problems is a Directory service repair.

As you learned in Course 3004, you can run Repair on the server where you are running iMonitor.

The Repair screen is shown in the following:

Figure 4-23



You must be the equivalent of Administrator of the server or a console operator on the server where you are accessing the Repair feature.

To repair a single object, select the object before selecting Repair.

Selecting a replica or partition root before you access Repair lets you repair all objects with that partition or that single object.



If Repair is loaded at the server, the Repair option in iMonitor cannot run.

The following shows the Repair log file, REPAIR.HTM:

Figure 4-24

```
Repairing volume object for volume DATA
Directory services volume object ID: 0000801C
ERROR: The volume ID cannot be resolved in the database
New volume object DN: CN=DA1_DATA.OU=IS.OU=SLC.O=DIGITALAIR.T=DIGITALAIR-TREE
Contacted a replica on server: CN=DA1.OU=IS.OU=SLC.O=DIGITALAIR.T=DIGITALAIR-TREE
The volume object has been created for this volume, ID: 00008053
The volume has been attached to the volume object
Volume: DATA, object ID: 00008053, CN=DA1_DATA.OU=IS.OU=SLC.O=DIGITALAIR.T=DIGITALAIR-TREE
Checking trustee on volume DATA
000080CD 00008017
Number of unique Trustee ID's found on this volume: 4
ERROR: Purging the invalid Trustees from the volume, total: 2

Repairing volume object for volume EDIR_87_WINDFW
Directory services volume object ID: 00000000
ERROR: The volume has never been installed
New volume object DN: CN=DA1_EDIR_87_WINDFW.OU=IS.OU=SLC.O=DIGITALAIR.T=DIGITALAIR-TREE
Contacted a replica on server: CN=DA1.OU=IS.OU=SLC.O=DIGITALAIR.T=DIGITALAIR-TREE
The volume object has been created for this volume, ID: 00008427
The volume has been attached to the volume object
Volume: EDIR_87_WINDFW, object ID: 00008427,
CN=DA1_EDIR_87_WINDFW.OU=IS.OU=SLC.O=DIGITALAIR.T=DIGITALAIR-TREE
Volumes checked: 3

** Automated Repair Mode **
Fixed: Sunday, November 9, 2002 6:20:09 am Local Time
Total repair time: 0.0043
```

As you review the log file, notice that the databases are checked, errors are identified, and are then corrected.



20 minutes

Exercise 4-3 **Verify Network Health**

In this exercise you manually perform a health check, and then run an Agent Health Report.

The results of the manual health check are identical to the results you receive when you run the Agent Health Report.

You perform the following Health Check procedures:

- [Part I: Check Schema Synchronization](#)
- [Part II: Check Agent Status](#)
- [Part III: Perform Health Check on Subsequent Servers](#)
- [Part IV: Run an Agent Health Report Check](#)

Discuss the questions in Part I as a group to facilitate student understanding of the results.

Part I: Check Schema Synchronization

eDirectory 8.7 contains a number of new objects and new schema extensions. After you complete an upgrade, you must verify that the schema is synchronized throughout the tree.

You should check schema synchronization on the server containing the master of root. After checking this server, check the other servers in the replica ring.

To check schema synchronization, do the following:

1. Launch your browser and in the Address field, enter **http://your server IP address:8009/nds-summary**.
2. Enter your userid and password: then select **Login**.
3. From the navigation toolbar, select **Trace Configuration**.
4. From Trace Line Prefixes, verify that the following prefixes are selected:
 - Time Stamp**
 - Thread ID**
 - Option Tag**
5. From DS Trace Options, select **Schema** and **Inbound Synchronization** (in addition to the options already selected).
6. From the bottom of the dialog, select **Trace On**.
7. From the navigation toolbar, select **Agent Configuration**.
8. From the list of links on the left, select **Agent Triggers**.
9. Select **Schema Synchronization**.
10. Select **Submit**.
11. View the trace by returning to Trace Configuration.
12. From the list of trace links on the left, select **Trace Live**.
13. Select **Refresh On**.

In a classroom setting, discuss the following as a group:

14. What information do you see in the Live Trace screen?

15. What indication is given when all trace information completed?

16. Why should you turn Trace off?

17. What steps must you take to turn Trace off?

Part II: Check Agent Status

To check the health of your system, do the following.

1. Select **Agent Configuration**.
2. Record the following information:
 - Agent Build Number:
 - Time Synchronized:
 - Root Most Master:

3. From the list of links on the left, select **Agent Synchronization**.
4. Are any errors listed?
5. At the right of the .DIGITALAIR-TREE master partition, select **Continuity**; then verify that there are no errors.
6. Select **Agent Process Status**.
7. What status information is available in the Agent Process status screen?

8. Select **Agent Synchronization**.
9. From the .DIGITALAIR-TREE master replica, select **Replica Synchronization**.
10. Verify the replica state and the last successful synchronization for each replica on the server.
11. What does the **On** replica state mean?

12. From the Partition Synchronization Status, select **Agent Health**.
13. List and describe the 2 health check categories and their function:
 -

 -

14. From the health check you just performed, are you satisfied that eDirectory on your server functioning properly?



You might see a 628 error for DA3 because your server can't synchronize to that server.

Part III: Perform Health Check on Subsequent Servers

You can use one of the following to ensure that all servers in your tree are functioning properly:

- Perform DSTrace and DSRepair at the server console.
- Change the URL in iMonitor to pull the information from a specific server.
- Use the links in Replica synchronization to view each server.

Using Replica Synchronization links in iMonitor, do the following:

1. Select **Agent Synchronization**.
2. From Partition Synchronization Status, select **Replica Synchronization** on the right.
3. Select a **Replica** link on the bottom left to access another server.

Notice the reference at the top left of the screen shows that you are now accessing information for that server.

4. Select **Agent Summary**.
5. Using the steps in this exercise, you can now verify the health on this server.

Part IV: Run an Agent Health Report Check

Do the following:

1. On the left under the **Links** heading, select **Agent Health**.
On the right, a Health Check list appears.
2. Select an option (such as **Partitions/Replica**) to run a report.
3. What additional information is shown?

4. When would you use the Agent Health Report check feature instead of performing your own health check with the options in Parts I - III?

(End of Exercise)



20 minutes

Exercise 4-4 Evaluate an eDirectory Problem

In a classroom setting, discuss as a group or class the following scenario and exercise questions.

In the Agent Summary screen in iMonitor, you notice that the master replica shows errors. Using the troubleshooting steps you learned in this section, complete the following:

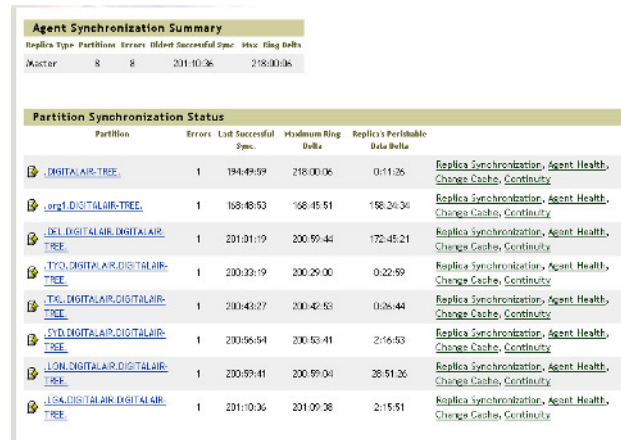
- [Part I: Identify the Scope of the Problem](#)
- [Part II: Determine the Cause of the Problem](#)
- [Part III: List Possible Solutions to the Problem](#)
- [Part IV: Assess Possible Solutions](#)
- [Part V: Implement a Solution](#)

- Part VI: Verify that the Problem Is Resolved
- Part VII: Document the Resolution to the Problem
- Part VIII: Avoid Repeating the Problem

Part I: Identify the Scope of the Problem

Using the figures provided, answer the following questions:

Figure 4-25

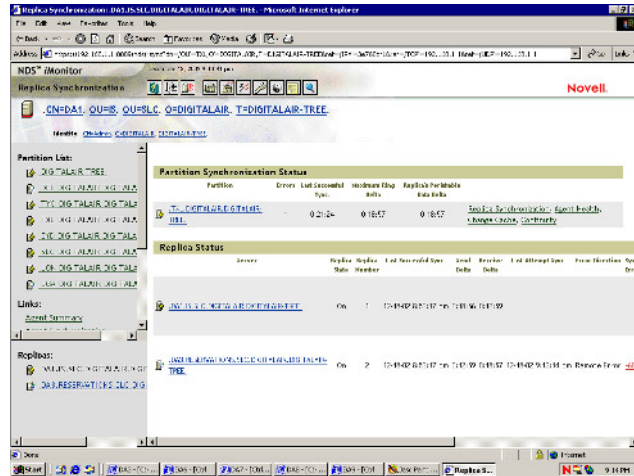


1. In the Agent Synchronization Summary screen, what indications do you have that there is a problem?

2. From Partition Synchronization Status, list all partitions experiencing the problem:

The following shows the synchronization status.

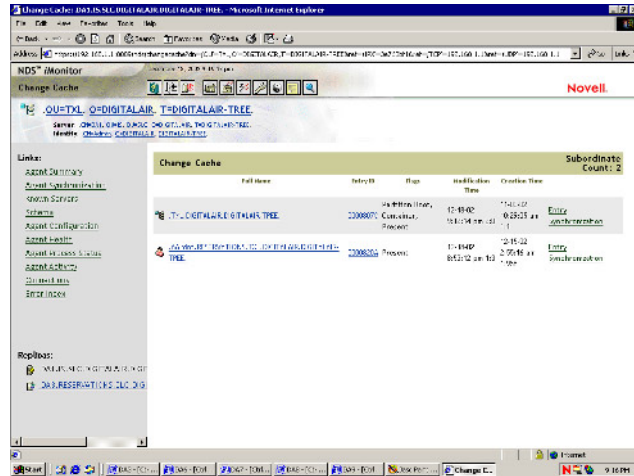
Figure 4-26



- 3. What error code is shown?
- 4. What is the source of the error?
- 5. Select Change Cache.

The following appears:

Figure 4-27



6. What information does Change Cache provide?

Part II: Determine the Cause of the Problem

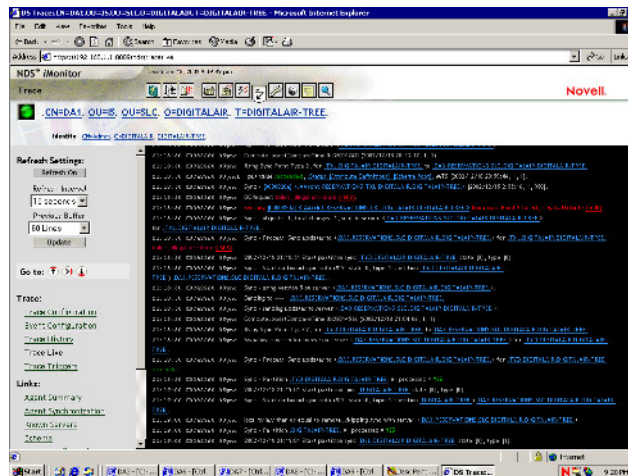
To determine the cause of the problem, identify the following:

1. Is the problem a LAN, server, or eDirectory problem?

2. What sources should you use to identify the problem?

To help identify additional detail regarding the error, you could perform a Trace. The following shows the error using Trace commands.

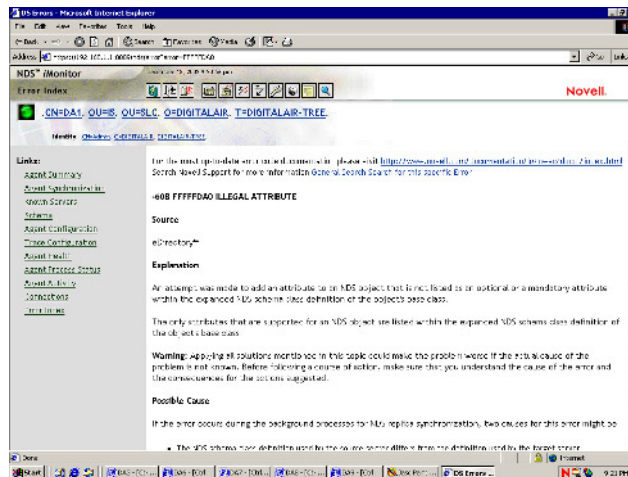
Figure 4-28



Before taking action on the user's record, try to determine what the illegal attribute is for this user's record.

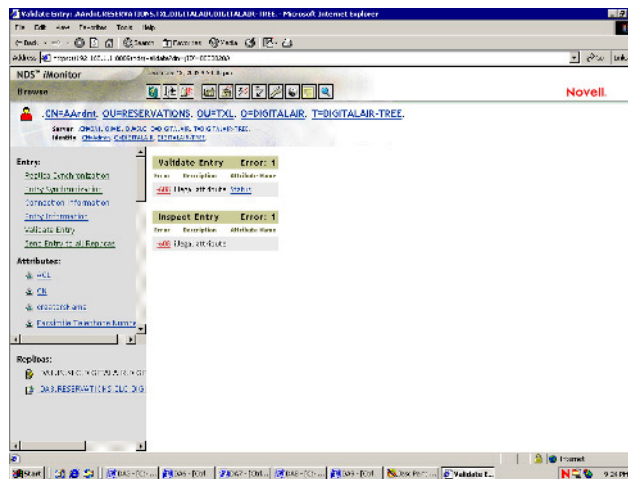
From the Trace Live screen, you can select the Red text to be linked to the Error section in the Help file, as seen in the following:

Figure 4-29



A quick way to verify a problem is to select the object and then use the Validate Entry option, as seen in the following:

Figure 4-30



This shows the error letting you know that the attribute **STATUS** has been changed and does not match the schema.

3. From your browser, access <http://www.novell.com>; then select **Support > Knowledgebase**.
4. Perform a search using 608 Illegal Attribute as your search criteria.

A number of TID references are listed, relating to the problem and possible solutions.

5. Using the Error Codes help screen and Knowledgebase, what could have caused the problem?

Table 4-2

Possible Causes

1.

2.

3.

Part III: List Possible Solutions to the Problem

1. From your browser, access <http://www.novell.com>; then select **Support > Knowledgebase**.
2. Perform a search using 608 Illegal Attribute as your search criteria.

A number of TID references are listed, relating to the problem and possible solutions.

3. Using the Error Codes help screen and Knowledgebase, what solutions can you try?

Table 4-3

Solutions

1.

2.

3.

Part IV: Assess Possible Solutions

Based on the solutions you have discovered, you must evaluate which of them will correct the problem without causing an adverse effect elsewhere in the tree.

1. Which solution seems most likely to resolve the problem with minimal difficulty?

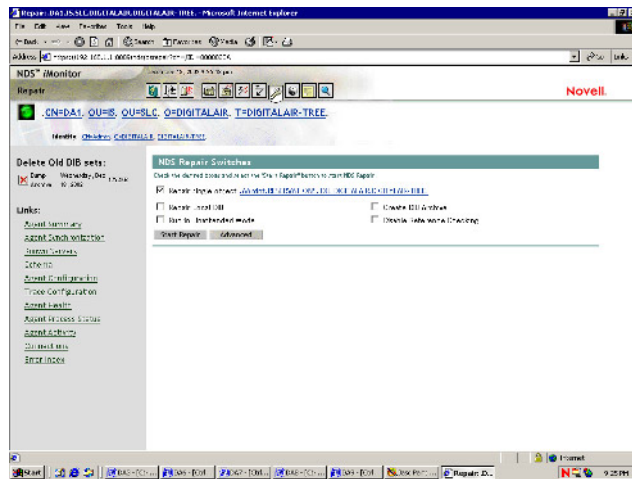
2. Talk to the members of your class and determine the best solution.

Part V: Implement a Solution

Now that you've identified possible solutions and determined the best solution for your situation, implement the solution.

In this scenario, it was determined that you should repair the user's record. The following shows the repair dialog:

Figure 4-31



The Repair Single Object option is selected and the user's record appears. This happens because the user's record was being viewed prior to selecting Repair.

After Repair is run, you can view the results by accessing **dsrepair.htm**, as seen in the following:

Figure 4-32

```
Repairing volume object for volume DATA
Directory services volume object ID: 0000801C
ERROR: The volume ID cannot be resolved in this database
New volume object DN: CN=DA1_DATA,OU=IS,OU=SLC,O=DIGITALAIR,T=DIGITALAIR-TREE
Contacted a replica on server: CN=DA1,OU=IS,OU=SLC,O=DIGITALAIR,T=DIGITALAIR-TREE
The volume object has been created for this volume, ID: 00008053
The volume has been attached to the volume object
Volume: DATA, object ID: 00008053, CN=DA1_DATA,OU=IS,OU=SLC,O=DIGITALAIR,T=DIGITALAIR-TREE
Checking trustees on volume: DATA
000080D 00008017
Number of unique Trustee IDs found on this volume: 4
ERROR: Purging the invalid Trustees from the volume, total: 2

Repairing volume object for volume EDIR_87_WINNW
Directory services volume object ID: 00000000
ERROR: The volume has never been installed
New volume object DN: CN=DA1_EDIR_87_WINNW,OU=IS,OU=SLC,O=DIGITALAIR,T=DIGITALAIR-TREE
Contacted a replica on server: CN=DA1,OU=IS,OU=SLC,O=DIGITALAIR,T=DIGITALAIR-TREE
The volume object has been created for this volume, ID: 00008427
The volume has been attached to the volume object
Volume: EDIR_87_WINNW, object ID: 00008427,
CN=DA1_EDIR_87_WINNW,OU=IS,OU=SLC,O=DIGITALAIR,T=DIGITALAIR-TREE
Volumes checked: 3

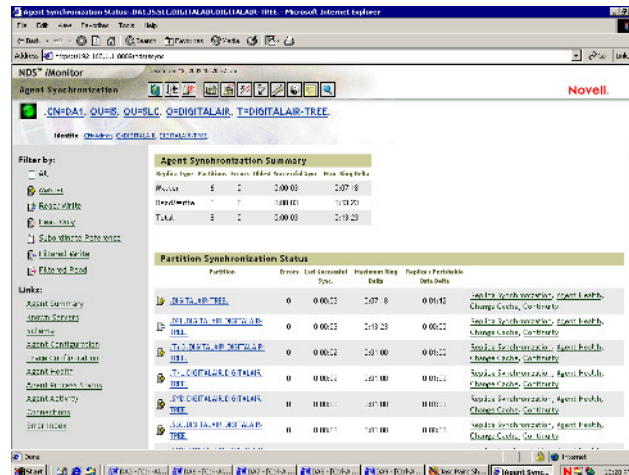
** Automated Repair Mode **
Finish: Saturday, November 9, 2002 6:20:09 am Local Time
Total repair time: 0.0043
```

1. What action did Repair perform?

Part VI: Verify that the Problem Is Resolved

After you implement your solution, verify that the problem is resolved. The following shows the **Agent Summary** link:

Figure 4-33



1. What indications are there that the problem is resolved?

Part VII: Document the Resolution to the Problem

Now that eDirectory is functioning properly, record the problem and the solution you used. This record can be accessed if the problem presents itself in the future.

1. In a second browser window, open **Novell Remote Manager**.
2. Access the **Server Personal Log**.

3. What was the problem (including error codes)?

4. How did the problem occur? (You might not know how the problem was introduced to your system, but if you do, record it so it can be prevented.)

5. What steps did you follow to resolve the problem?

6. What additional actions did you take to prevent the problem from occurring again?

For example, if another application was generating the problem, you could configure the application to stop initiating the problem. Or if an attribute is required, you can extend the schema to accept the attribute.

Part VIII: Avoid Repeating the Problem

The problem you just experienced and resolved can occur again in the future for a number of reasons.

With the problem and solution documented, you can save yourself and your peers a great deal of time if this problem occurs again. Consider how you would respond to the following:

1. In addition to documenting the problem, what can you do to avoid duplication of problems?

Summary

The following is a summary of the objectives in this section:

Objective	Summary
1. Identify eDirectory Databases and Processes	<p>Knowing the databases and the processes involved in eDirectory can help resolve problems as they arise.</p> <p>eDirectory 8.7 Databases include the following:</p> <ul style="list-style-type: none">■ NDS.DB■ NDS*.LOG■ NDS.01■ Stream files <p>eDirectory Processes include the following:</p> <ul style="list-style-type: none">■ Time synchronization■ Schema synchronization■ Replica synchronization

Objective	Summary
2. Identify eDirectory Troubleshooting Steps	<p>The key to troubleshooting eDirectory is to have a process in place. The following identifies a suggested process:</p> <ul style="list-style-type: none">■ Step 1: Identify the Scope of the Problem■ Step 2: Determine the Cause of the Problem■ Step 3: List Possible Solutions to the Problem■ Step 4: Assess Possible Solutions■ Step 5: Implement a Solution■ Step 6: Verify That the Problem Is Resolved■ Step 7: Document the Resolution to the Problem■ Step 8: Avoid Repeating the Problem
3. Identify Partition and Replication Placement Design	<p>Placement of partitions and replicas depends on the needs of your organization. The following identifies areas of concern:</p> <ul style="list-style-type: none">■ Place replicas where they are needed.■ Configure a minimum of 3 replicas.■ Do not add replicas to areas where they are not needed. This could cause unnecessary synchronization processes.

Objective	Summary
4. Use iMonitor Reports to Obtain Server and eDirectory Information	<p>iMonitor reporting is a powerful tool for determining the condition of a given server, agent process, or tree.</p> <p>The following reports are available:</p> <ul style="list-style-type: none">■ Server Information. This report searches the entire tree, communicates with every Netware Core Protocol™ (NCP™) server it can find, and reports errors it finds.■ Obituary Listing. This report lists all obituaries on this server.■ Object Statistics. This report evaluates the objects in a given scope and then generates lists of objects matching the requested criteria.■ Service Advertising. This report lists all directories and servers known to this server through SLP or SAP.■ Agent Health. This report gathers health information for this server.■ Custom Report. This report can create a customized report and scheduled events.

Objective	Summary
5. Perform Health Checks	<p data-bbox="954 281 1479 373">To ensure that eDirectory is functioning properly or to identify problems, you should check the following:</p> <ul data-bbox="954 394 1312 684" style="list-style-type: none"><li data-bbox="954 394 1179 422">■ eDirectory revision<li data-bbox="954 432 1203 459">■ Time synchronization<li data-bbox="954 470 1175 497">■ Partition continuity<li data-bbox="954 508 1312 684">■ Background processes, such as<ul data-bbox="992 543 1276 684" style="list-style-type: none"><li data-bbox="992 543 1276 571">■ Schema synchronization<li data-bbox="992 581 1133 609">■ Obituaries<li data-bbox="992 619 1224 646">■ External references<li data-bbox="992 657 1166 684">■ Limber status <p data-bbox="954 699 1479 791">Historically, performing Health Checks required you to use several tools. Using iMonitor, you can perform all your checks using one tool.</p> <p data-bbox="954 806 1479 863">iMonitor includes an Agent Health Report that quickly indicates any problems with eDirectory.</p> <p data-bbox="954 877 1479 949">After problems are identified, you can perform a Repair to correct the inconsistencies.</p>

Exercise Answers

Exercise 4-2. Verify eDirectory Status Using Reports

8. Which servers are showing errors?

DA3.

9. Describe the errors.

The Connection status says that the remote server is down, and the error message says that iMonitor can't connect to the specified server.

10. Which servers have recommended actions?

Results will vary.

11. Why are these recommendations being made?

Results will vary.

15. What are the current errors for any IPX servers in your network (such as DA3)? Why have they changed?

The remote server is up because the NetWare 4.11 server uses IPX. By using the IPX protocol, iMonitor can communicate with this server.

Exercise 4-3. Verify Network Health

Part I: Check Schema Synchronization

14. What information do you see in the Live Trace screen?

The date and time

Begin schema sync time

Schema sync upto “?”

All processed = YES.

15. What indication is given when all trace information completed?

All processed = Yes

16. Why should you turn Trace off?

Trace increases network traffic and can be turned off when not performing system checks.

17. What steps must you take to turn Trace off?

1. Select **Trace Configuration**.
2. Select **Trace Off**.

Part II: Check Agent Status**2. Record the following information:**

- Agent Build Number: 10410.98
- Time Synchronized: True
- Root Most Master: Yes

4. Are any errors listed?

Yes

7. What status information is available in the Agent Process status screen?

Schema Synchronization, Obituaries, External References, Limber, and server status

11. What does the On replica state mean?

The **On** state indicates that there are no active processes.

13. List and describe the 2 health check categories and their function:

- Agent. To view information about time synchronization and the state of eDirectory on the server.
- Partition. To view health information, such as replica synchronization and replica ring detail, for each partition on the server.

14. From the health check you just performed, are you satisfied that eDirectory on your server functioning properly?

Yes

Part IV: Run an Agent Health Report Check

3. What additional information is shown?

Answers will vary

4. When would you use the Agent Health Report check feature instead of performing your own health check with the options in Parts I - III?

Answers will vary

Exercise 4-4: Evaluate an eDirectory Problem

Part I: Identify the Scope of the Problem

1. In the Agent Synchronization Summary screen, what indications do you have that there is a problem?

There is an error under the Agent Synchronization Summary from the Master replica.

3. What error code is shown?

608

4. What is the source of the error?

NDS or eDirectory

6. What information does Change Cache provide?

Change cache reflects what objects were changed, but does not show errors associated with this change.

Part II: Determine the Cause of the Problem

1. Is the problem a LAN, server, or eDirectory problem?
eDirectory
2. What sources should you use to identify the problem?
iMonitor Reports, Trace, Error logs
5. Using the Error Codes help screen and Knowledgebase, what could have caused the problem?

Possible Cause

1. If the error occurs during the background processes for NDS or eDirectory replica synchronization, 2 causes might be
 - The NDS or eDirectory schema class definition used by the source server differs from the definition used by the target server. Additionally, the error indicates that the NDS or eDirectory schema on the source server contains additional information.
 - The database on the source server is damaged.
 2. If this error occurs while attempting to add an attribute to an object, an unsupported attributed might have been used.
 3. If this error occurs while adding an attribute to an NDS or eDirectory alias object, NDS or eDirectory will not accept this task because this task is not supported by NDS or eDirectory.
-

Part III: List Possible Solutions to the Problem

3. Using the Error Codes help screen and Knowledgebase, what solutions can you try?

Solutions

-
1. If you suspect that the database on the source server is the problem, try repairing it using REPAIR.

 2. Only use attributes for an NDS or eDirectory object that are supported in the expanded NDS or eDirectory schema class definition of the object's base class.

 3. NDS or eDirectory does not support adding attributes to alias objects. Therefore, this task cannot be performed.
-

Part IV: Assess Possible Solutions

1. Which solution seems most likely to resolve the problem with minimal difficulty?

To run Repair.

Part V: Implement a Solution

1. What action did Repair perform?

The attribute is purged from the users record.

Part VI: Verify that the Problem Is Resolved

1. What indications are there that the problem is resolved?

iMonitor is no longer showing errors on the Agent Summary screen.

Part VII: Document the Resolution to the Problem

3. What was the problem (including error codes)?.

eDirectory displayed a 608 error. This error identifies that there was an illegal attribute on a users record.

4. How did the problem occur? (You might not know how the problem was introduced to your system, but if you do, record it so it can be prevented.)

- This error occurred when an attribute was changed on the users record illegally.

5. What steps did you follow to resolve the problem?

The following troubleshooting steps were used:

- Step 1: Identify the Scope of the Problem
- Step 2: Determine the Cause of the Problem
- Step 3: List Possible Solutions to the Problem
- Step 4: Assess Possible Solutions
- Step 5: Implement a Solution
- Step 6: Verify That the Problem Is Resolved
- Step 7: Document the Resolution to the Problem
- Step 8: Avoid Repeating the Problem

6. What additional actions did you take to prevent the problem from occurring again? For example, if another application was generating the problem, you could configure the application to stop initiating the problem. Or if an attribute is required, you can extend the schema to accept the attribute.

In this case, the error was a forced change. If the error was generated by an application or by another system administrator, document the cause and what process you would put in place to correct the problem.

Part VIII: Avoid Repeating the Problem

In addition to documenting the problem, what can you do to avoid duplication of problems?

Discuss your experience with other administrators in your organization.

Implement any necessary procedures or processes to prevent the problem from happening again.

MODULE 3

Demonstrate Advanced Novell Network Storage Management Skills

Section 5 Perform Advanced Novell Storage Services Tasks

Section 6 Configure and Troubleshoot a RAID Solution Using NSS

Section 7 Perform Advanced iFolder Tasks and Troubleshooting

SECTION 5 Perform Advanced Novell Storage Services Tasks

Duration: 1 hour 30 minutes

In this section you learn about advanced Novell Storage Services (NSS) management tasks.

Objectives

1. Expand an NSS Storage Space
2. Configure NSS Volume Attributes
3. Mount a DOS Partition as an NSS Volume
4. Use VCU to Create an NSS Volume from a Traditional Netware Volume
5. Resolve Common NSS Errors
6. Restore a Deleted Logical Volume
7. Describe Storage Area Networks and Network Attached Storage

Introduction

As data storage needs on the network grow, satisfying those needs becomes more complex. Novell Storage Services (NSS) is a very scalable and flexible file system that can meet these needs, but it can require advanced NSS management skills.

In this section you learn about these skills and 2 other storage architectures and technologies: Storage Area Networks (SANs) and Network Attached Storage (NAS).

The Scenario

At Digital Airlines, you are the network administrator for one of the branch offices and have installed NetWare 6 on your servers.

You want to take maximum advantage of NSS to resolve file storage requests and issues such as improving access to data. You are considering using a SAN as a file storage solution.

Objective 1 Expand an NSS Storage Space

As storage demands grow, you will need to expand the NSS storage space on your servers. There are 3 ways in which NSS lets you expand the amount of space available:

- [Increase the Number of Logical Volumes in a Storage Pool](#)
- [Overbook the Storage Pool](#)
- [Increase the Size of a Storage Pool](#)

Increase the Number of Logical Volumes in a Storage Pool

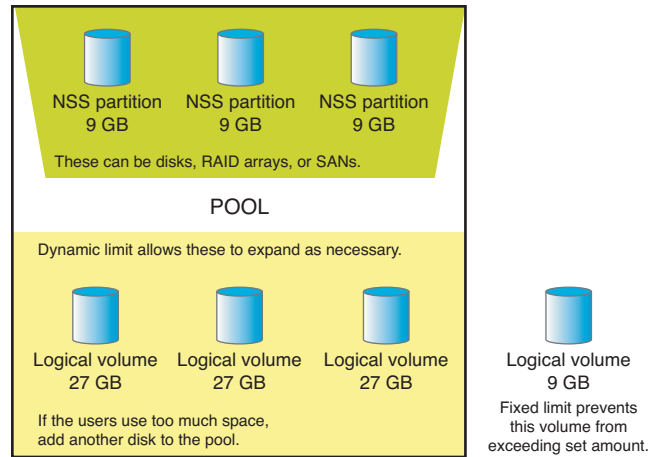
To increase the storage space on a server, you can increase the number of logical volumes in an NSS storage pool.

In Course 3004, you learned that a *storage pool* is a specified amount of space from various storage devices used to contain multiple NSS logical volumes.

You might want to demonstrate how to create a volume as a review of Course 3004.

The following shows logical volumes in a pool:

Figure 5-1



A pool can contain many partitions, but a partition can only be included in one pool.

You can create up to 255 logical volumes in the storage pool. Remember though, that clients can only map up to 26 drives (AZ).

Overbook the Storage Pool

You might want to demonstrate how to overbook a pool as a review of Course 3004.

Another method for increasing the amount of storage space on a server is to overbook the storage pool.

In Course 3004 you learned that although the size of an individual logical volume cannot exceed the size of a storage pool, the sum of multiple logical volumes in the pool can exceed the pool size.

This feature, called *overbooking*, can be an efficient way to manage your file system, without having to add more disk space.

You can limit some volumes to a certain size and allow others to grow as necessary in the pool.

For example, you might have students at a university constantly pushing the limits of Volume A, while the faculty of a department has plenty of space for their needs on Volume B.

Using NSS, the sum of Volume A and Volume B can actually exceed the size of the storage pool.

With overbooking, the users of Volume A can, essentially, borrow space from Volume B as long as Volume B is not filled to the limit. If too much space is used, the pool will report “Out of space.”

To view pool information, at the server console, enter **NSS /SPACE**. Information similar to the following appears:

Figure 5-2

```

001:nss /space
-----
Pool/Volume Name      Total Space/Quota  Used Space  Available Space
NSS Pools/Volumes:
  DATA                390.00 MB          8.85 MB    381.14 MB
  DATA                390.00 MB @       560.00 KB  381.14 MB *
  SYS                  3.41 GB           655.50 MB  2.77 GB
  SYS                  3.41 GB @       591.03 MB  2.77 GB *
* -- Available size is limited by the free space available on the pool.
@ -- Volume has no quota sat and can grow to the pool size.
001:~

```

Increase the Size of a Storage Pool

If a storage pool is not large enough to accommodate your storage needs, you can increase the size of the pool by adding more storage devices, without having to recreate the pool or its volumes.



While you can increase the size of a pool, you cannot reduce the size.

To increase the size, use one of the following tools:

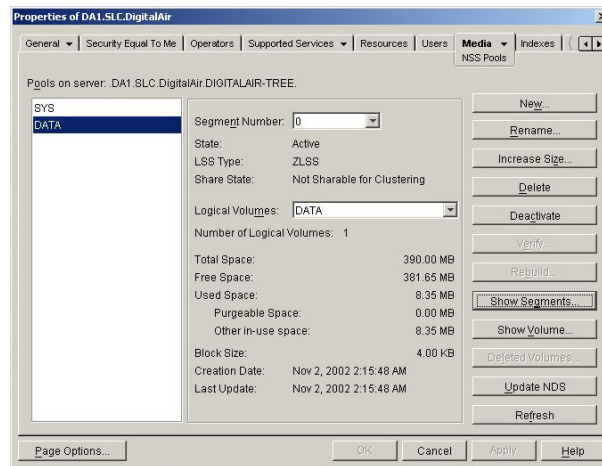
- [ConsoleOne](#)
- [Remote Manager](#)
- [NSSMU](#)

ConsoleOne

ConsoleOne lets you complete NSS management tasks from the Media tab in a Server Object Properties dialog:

Demonstrate how to use ConsoleOne to increase the size of a storage pool.

Figure 5-3



To increase the size of a pool, do the following:

1. Add storage devices to your server.
2. Create NSS partitions on the new storage devices.

3. Access the Media Properties page for your server object:
 - a. From ConsoleOne, browse to a **server object**.
 - b. Right-click the server; then select **Properties**.
 - c. Select **Media > Pools**.
4. Note the size of the pool:
 - a. Select a **pool**.
 - b. Note the amount of space in the pool.
5. Increase the size of the pool:
 - a. Select **Show Pool**.
 - b. Select **Increase Size**.
 - c. Select the partition or free space you want to add.
 - d. Check **Used**.
 - e. Select **Finish**.
 - f. Select **Yes**.
 - g. Note the amount of space in the pool.

Dynamic NSS volumes can grow to the size of the pool, but fixed-size volumes can't grow, even by adding space to the pool.

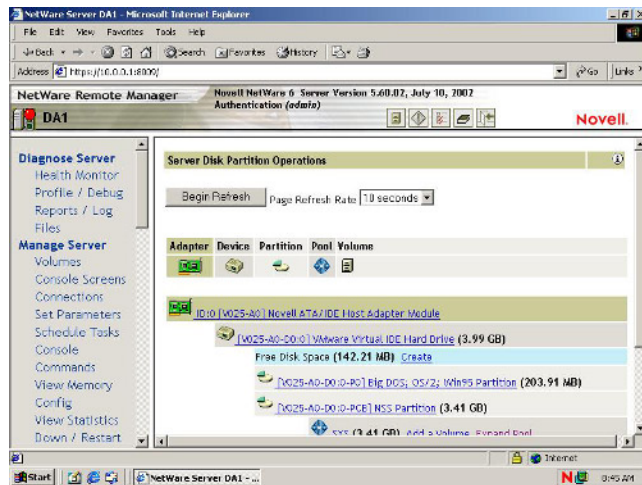
Remote Manager

Demonstrate how to use Remote Manager to increase the size of a storage pool.

Remote Manager lets you complete NSS tasks from anywhere you have a web browser and Internet access.

The following shows the partition screen in Remote Manager from which you can perform these tasks:

Figure 5-4



Remote Manager is particularly convenient when you are at home or away from the office somewhere and need to manage NSS.

To increase the size of a pool using Remote Manager, do the following:

1. From a workstation, open an Internet browser.
2. Browse to **Remote Manager**.
3. Select **OK** at the security alert.
4. Select the **Volumes** icon.
5. Under Partition Management, select **Disk Partitions**.
6. Browse to the pool you want to increase.
7. Select **Expand Pool**.
8. Select **Free Disk Space**.
9. Enter the amount.

10. Select **Expand**.
11. Select **OK**.
12. Note the amount of space in the pool.

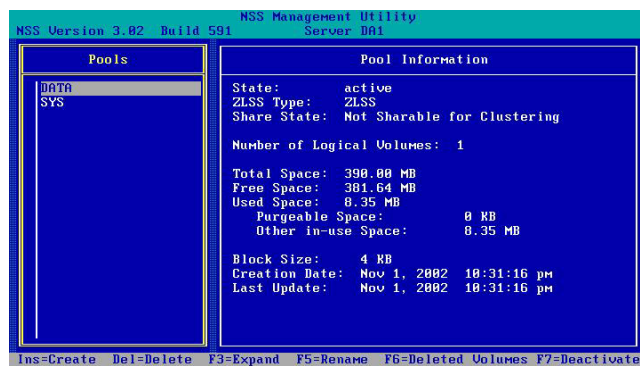
NSSMU

Demonstrate how to use NSSMU to increase the size of a storage pool.

NSS Management Utility (NSSMU) lets you complete the same tasks at the server console that you can complete using ConsoleOne or Remote Manager at a workstation.

The following shows the Pool Information screen in NSSMU:

Figure 5-5



NSSMU is convenient to use when you are working on the server and not near a workstation.

To increase the size of a pool, use NSSMU to do the following:

1. At the server console, enter **LOAD NSSMU**.
2. Select **Pools**.
3. Press **F3**.
4. Press **Enter**.
5. Enter a number.

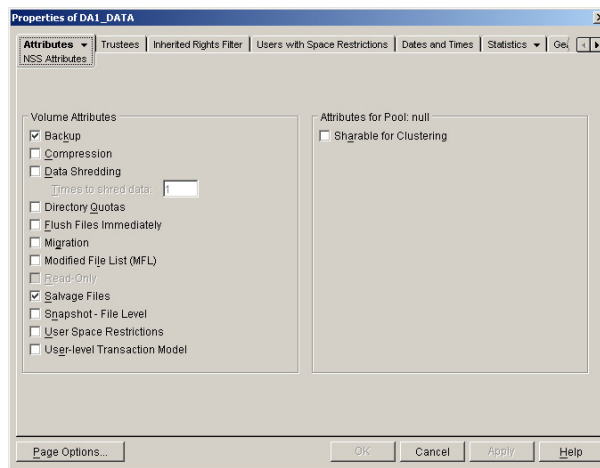
6. Press **Enter**.
7. Note the amount of space in the pool.

Objective 2 Configure NSS Volume Attributes

With NSS, you can configure several attributes of your NSS volumes to increase their manageability.

You can set NSS volume attributes from ConsoleOne on the Attributes tab, as shown in the following:

Figure 5-6



The following describes the NSS volume attributes:



Do not enable features you do not need. Many of these features will affect the performance of NSS.

- **Backup.** Enables backup of the volume. This does not automatically enable or launch third-party backup utilities. You need to configure third-party backup utilities as directed.

- **Compression.** Enables compression for the entire volume. After you select Compression, you cannot turn it off for the volume.
- **Data Shredding.** Activates the Data Shredding security feature. This feature scrambles any data you delete to prevent anyone from accessing the information on a disk reader.

This might be useful to ensure confidentiality of information when preparing to surplus storage devices.

The Data Shredding feature uses random hex characters to write to the blocks where deleted files used to reside.

Enter the number of times you want to apply Data Shredding to your deleted files. The range is from 1 to 7 times.
- **Directory Quotas.** Lets you assign the maximum quota of space that a directory can have.
- **Flush Files Immediately.** Flushes all file data to disk immediately when you close a file.
- **Migration.** Activates the Data Migration feature for the files on this volume. Migration moves old data to an M, N, or O drive.
- **Modified File List (MFL).** Maintains a list of all files modified since the last backup. Your file system maintains this list; however, a third-party vendor must implement this feature for you to use it.
- **Salvage Files.** The file system keeps all deleted files in an allocated space until that space is needed for other data. The Salvage Files feature tracks the files and lets you retrieve the deleted data for a time until the space is needed for other data.
- **Snapshot-File Level.** Activates the Snapshot feature at the file level. This feature lets the backup utility capture a snapshot of the last closed version of a file.

For example, if your system backs up or crashes while you have a file open, this feature will save a copy of the file before you opened it. You might lose some new information, but you will retain all the previous information.

When you enable File Snapshot on a volume, the backup utility copies information about the individual file, such as owner, creation date and time, and modification date and time.

During write requests, only the data that is being written is copied. NSS recognizes how to return the appropriate data when read requests come in for the real file or the snapshot file.

After you set up File Snapshot on your volumes, you must deactivate and then reactivate and remount the volumes. This ensures that there are no open files without a snapshot.

- **User Space Restrictions.** Lets you assign space usage quotas for users on a volume. Setting up user space restrictions is beneficial for systems that have a large number of users, such as students or contractors.
- **User-Level Transaction Model.** Protects database applications from corruption by backing out incomplete transactions that result from a failure in a network component.

The following describes the pool attribute:

- **Sharable for Clustering.** Lets you configure the pool (and accompanying volumes) to reside on a device that is sharable for clustering. This feature cannot be modified at the pool level on the device level.

Objective 3 **Mount a DOS Partition as an NSS Volume**

NSS lets you mount DOS partitions as NSS volumes. For example, if you are having problems with server abends and need to create a core dump, you can attach a removable storage device with a DOS-formatted disk, such as an Iomega Jaz drive, to your server.

After loading the appropriate hardware drivers, you can mount the removable disk as an NSS volume and create a core dump on the disk. Because it is a DOS-formatted disk, the file can then be read from the disk on any Windows computer.

Mounting a DOS partition as an NSS volume also lets you apply patch files to or test new files on the DOS partition.

However, NSS-mounted DOS partitions should only be used for maintenance or troubleshooting purposes, not for user volumes.

To mount a DOS partition as an NSS volume, you must understand the following:

- [How DOSFAT.NSS Works](#)
- [How to Mount a DOS Partition Using DOSFAT.NSS](#)
- [How to Verify That the DOS Partition Is Mounted](#)

How DOSFAT.NSS Works

DOSFAT.NSS is an NSS module that mounts DOS (FAT16) partitions as NSS volumes. When you load the module, DOS FAT partitions on the server are mounted as logical NSS volumes. NSS will not mount FAT32 or Compaq system partitions.

The following are features of the DOSFAT.NSS module:

- **Support for long filenames.** The module uses the Windows 95/98 method for placing long filenames onto a DOS FAT partition.
- **Support for NetWare trustee rights.** You can grant other users access to the volume using eDirectory rights assignments, but typically, only the network administrator should have access to the volume.

- **Support for NetWare utilities.** Any client or server utility that accesses or manages NetWare volumes can do the same with volumes created with the DOSFAT.NSS module.
- **MAP command support.** NetWare clients can map to the volume and use it the same as any other NetWare drive mapping.

How to Mount a DOS Partition Using DOSFAT.NSS

Demonstrate how to mount a DOS partition.

To mount existing DOS partitions on your NetWare 6 server as NSS volumes, do the following:

1. At the server console prompt, enter **SET AUTO RESTART AFTER ABEND = 0.**



If you do not turn off automatic abend recovery, you risk corrupting the DOS drive's FAT tables.

If you do not set **AUTO RESTART AFTER ABEND = 0**, you will see the following:

```
WARNING: The 'Auto Restart After Abend' settable
parameter has a value of 1. Are you sure you want to
load DOSFAT.NSS (y/n)?
```

When the server abends, it writes detailed information into the ABEND.LOG file on the DOS bootable partition.

This write operation bypasses the internal DOSFAT LSS cache buffers, and it might corrupt the DOS drive's FAT tables.

To prevent this, set **Auto Restart After Abend = 0.**

2. Enter **LOAD DOSFAT.NSS.**

After DOSFAT.NSS is loaded, DOS partitions on the server are mounted and made available as logical volumes, as shown in the following:

Figure 5-7

```
DAI:load dosfat.nss
Loading Module DOSFAT.NSS
FAT LSS: Long name support enabled.
Moving DOS FAT volume 'DOSFAT_C' to active state
  Volume DOSFAT_C set to the DEACTIVE state.
Activating volume "DOSFAT_C"...
Loading FAT for volume DOSFAT_C into memory (52224 bytes)
  Volume DOSFAT_C set to the ACTIVE state.
Mounting Volume DOSFAT_C
** DOSFAT_C mounted successfully
DOS FAT Support Module loaded successfully.
Upgrading NetWare Configuration File from 5.5 to 5.6
DAI:~
```

To dismount DOSFAT.NSS, enter **UNLOAD DOSFAT.NSS**.

How to Verify That the DOS Partition Is Mounted

To verify that the DOS partition was mounted properly, look for the following:

- After loading DOSFAT.NSS, you should see a message stating **** DOSFAT_C mounted successfully** on the server console.
- At the console prompt, enter **VOLUMES**.

If the module loaded correctly, the volume appears as DOSFAT_*x*, where *x* is the drive letter, such as DOSFAT_C.

If NSS cannot determine the drive letter, the volume appears as DOSFAT_0 or DOSFAT_1.

Objective 4 Use VCU to Create an NSS Volume from a Traditional Netware Volume

You can copy traditional volumes to logical NSS volumes using the Volume Conversion Utility (VCU). To use VCU, you must understand the following:

- [How VCU Works](#)
- [How to Copy a Traditional Volume to an NSS Volume](#)
- [The Correct VCU Syntax](#)

How VCU Works

The VCU utility creates a new NSS volume and then copies data (keeping the same file structure) from the source traditional volume to the new NSS logical volume. VCU can copy volumes with long name space applied.

Because VCU creates a new volume, you must have adequate space on your server.

For example, if you want to copy a 2 GB traditional volume, you need to have an additional 2 GB of available disk space on the server (more if your traditional volume has compression turned on).



You must have enough space for both the traditional volume and the new logical volume.

The original volume is renamed **VOLUMENAME_OLD**. The new logical volume keeps the original volume name and maintains trustee assignments.

Keep in mind that copying traditional volumes to NSS volumes requires significant processing. Using VCU will affect server performance during the copy process.

Use VCU only when server demands are low (such as after working hours), after you disconnect all other users and disable login, and after you create a backup of the volume.

After you copy the traditional volume to a logical volume, restart the server to ensure the volume copied properly and then remove the traditional volume.



After you copy traditional volume data to a logical volume in NetWare 6, you cannot access the new logical volume using prior versions of NetWare.

Create a small traditional volume; then demonstrate how to copy an NSS volume from the traditional volume.

How to Copy a Traditional Volume to an NSS Volume

To copy a volume, load VCU.NLM, specify the *volume* to copy, and then specify the *NSS pool* where you want to create the new NSS volume. Use the following syntax:

VCU *volume pool*

For example, to copy a traditional volume named APPS to an NSS volume in a pool named POOLONE, enter the following at the server console prompt:

VCU APPS POOLONE

Don't be alarmed if this process returns an error about a file that cannot be copied. The traditional volume has a hidden system file that will not be copied. As a result, VCU will return an error but completes the process.

The Correct VCU Syntax

The syntax for VCU.NLM is as follows:

**VCU /P/L/I/D/R ORIGINALVOLUME NSSPOOL
[DS_CONTAINER [DS_VOLNAME]]**

The following describe each VCU command parameter:

Table 5-1

VCU Parameter	Description
/P	Do not print directory filenames.
/L	Do not write errors to log file (DST_VOL:ERROR.OUT).
/I	Keep file COMPRESS_FILE _IMMEDIATELY_BIT.
/D	Delete the original volume if the copy process is successful. If you delete the traditional volume, the new volume retains the name of that volume.
/R	Remove the new NSS volume name and restore the traditional volume name (use to keep the original volume name for the new logical volume name).
DS_CONTAINER	Designate the original volume of the eDirectory container.
DS_VOLUMENAME	Specify so VCU uses this name to rename or delete the original volume's eDirectory object. Otherwise, VCU.NLM uses SERVERNAME_ORIGINALVOLNAME as the default eDirectory name.



15 minutes

Exercise 5-1 Perform Advanced NSS Storage Management Tasks

In this exercise you do the following:

- [Part I: Create an NSS Volume from a Traditional NetWare Volume](#)
- [Part II: Expand an NSS Volume](#)
- [Part III: Mount a DOS Partition as an NSS Volume](#)

Part I: Create an NSS Volume from a Traditional NetWare Volume

Do the following:

1. Create a new pool called NEWPOOL:
 - a. From ConsoleOne, browse to *your server*.
 - b. Right-click *your server*; then select **Properties**.
 - c. Select **Media > NSS Pools**.
 - d. Select **New**.
 - e. In the Name field, enter **NEWPOOL**; then select **Next**.
 - f. Mark the *unpartitioned space*.
 - g. In the Used field, enter a *size* (in MB) equal to or greater than the size of volume DATA; press **Enter** or **Tab**; then select **Next**.
 - h. Select **Finish**.
 - i. If prompted with a warning concerning the Hot Fix size, select **Yes**.
2. Verify that NEWPOOL was created:
 - a. At the server console, enter **NSS /SPACE**.
 - b. Verify that NEWPOOL appears in the list.

3. Convert the traditional volume DATA to an NSS volume in NEWPOOL:
 - a. At the server console, enter **VCU /D DATA NEWPOOL**.
 - b. When prompted to rename volumes, type **Y**.
 - c. When prompted to delete old volumes, type **Y**.
 - d. Wait while files are copied.
 - e. When prompted, press *any key* to complete the process.
 - f. Restart the server.

Part II: Expand an NSS Volume

Do the following:

1. Access the Media Properties page for your server object:
 - a. From ConsoleOne, browse to *your server*.
 - b. Right-click *your server*, then select **Properties**.
 - c. Select **Media > NSS Pools**.
2. Note the size of the NEWPOOL pool:
 - a. Select the **NEWPOOL** pool.
 - b. Record the total amount of space in the pool:
3. Increase the size of the pool:
 - a. Select **Increase Size**.
 - b. Select the *unpartitioned space*.
 - c. In the Used field, delete the current number and type **50**; then press **Enter**.
 - d. Continue by selecting **Finish**.
 - e. If prompted with a warning concerning the Hot Fix size, select **Yes**.
 - f. Note the amount of space available.

Part III: Mount a DOS Partition as an NSS Volume

Do the following:

1. Mount your DOS partition as an NSS volume:
 - a. At the server console enter **SET AUTO RESTART AFTER ABEND = 0**.
 - b. Enter **LOAD DOSFAT.NSS**.
2. At the server console, enter **VOLUMES**.
3. Verify that **DOSFAT_C** mounted successfully by mapping a drive to the volume and viewing the files.

(End of Exercise)

Objective 5 Resolve Common NSS Errors

To resolve common NSS errors, you need to do the following:

- Determine the Cause of the Problem and List Possible Solutions
- Use **VERIFY** to Determine the Integrity of an NSS Pool
- Assess Possible Solutions
- Use **REBUILD** as a Last Resort Solution

Determine the Cause of the Problem and List Possible Solutions

When a storage problem exists on the network, NSS prompts you with an error code. An NSS error code can lead you to a resolution, but it might only be a symptom of the real issue.

To determine the cause of the problem, identify whether the problem is actually an NSS problem or an external problem, such as a unattached cable or a low memory situation on the server.

Sometimes problems are produced when NSS defaults are changed without a clear understanding of the consequences.

For example, NSS cache parameters are best left at their default value. If you change them without consulting with Novell, you could lose large amounts of data.

Also, SYS should always be the only volume in its pool. Create other pools for other volumes.

The following are errors you might experience while working with NSS:

Table 5-2

Problem	Explanation	Solution
NSS does not recognize a device.	NSS can only use what the Media Manager recognizes.	At the server console, enter SCAN FOR NEW DEVICES; then enter LIST DEVICES. If the device is not listed or appears as an unbound object, most likely the device is malfunctioning or the appropriate driver is not loaded. Make sure the correct board driver (*.HAM) and device driver (*.CDM) are loaded.

Table 5-2 (continued)

Problem	Explanation	Solution
NSS does not let you create a storage pool or a logical volume.	<ul style="list-style-type: none"> ■ Your storage devices might not have enough free space to create more storage pools or logical volumes. ■ All logical volumes might not be part of the same storage pool. ■ NSS might not own the free space you want to use for a storage pool. 	<ul style="list-style-type: none"> ■ Make sure you have enough free space to create another storage pool or logical volume. ■ Before you create a logical volume, create a storage pool. ■ Create an NSS partition for your storage pools and logical volumes.
You cannot configure a logical volume.	NSS might not have enough free space for another logical volume.	<ul style="list-style-type: none"> ■ Add another storage device. ■ Delete a logical or traditional volume to free up space for a storage pool.
You cannot compress a file.	You didn't choose the file compression option when you created a logical volume.	<p>Apply the file compression option to an existing logical volume.</p> <p>From ConsoleOne, select Media > NSS Logical Volumes > Properties > NSS Attributes; then select Compression.</p>



You can access a list of all NSS error codes and possible solutions at <http://www.novell.com/documentation/lg/nwec/index.html?treetitl.html>

Demonstrate how to use VERIFY.

Use VERIFY to Determine the Integrity of an NSS Pool

Whenever you need to check the integrity of an NSS file system of a pool, you can use VERIFY.

VERIFY is a read-only utility that checks the file system integrity for an NSS pool by searching for inconsistent data blocks or other errors. This utility checks to see if there are problems with the file system.

VERIFY performs a read-only assessment of a pool. It dismounts the volumes, so you'll have to mount them when finished.

Run VERIFY before running a REBUILD process. To run VERIFY, do the following:

1. At the server console prompt, enter NSS **/POOLVERIFY=poolname**.
2. When prompted that volumes will be dismounted, select **Yes**.
3. Press **F1** for a list of errors.



If you encounter an unaccounted block error, don't do anything about it. The unaccounted block error will work itself out.

4. Press **F6** for a list of conflicts.
5. When the process is complete, at the server console enter **MOUNT ALL**.

Assess Possible Solutions

After you gather a list of possible solutions to your NSS problem, assess the solutions based on the following:

- The likelihood that it will solve your problem
- How easy or hard the solution is to implement

- What effect the implementation process will have on users
- Whether the solution will have a negative impact on storage access
- The validity of backup copies
- The amount of time it will take to restore the backup

Ask coworkers and others what actions they would take. If possible, test the solutions in a lab environment before implementing them.

Assessing the possible solutions often involves deciding between contradicting solutions from various sources.

Use REBUILD as a Last Resort Solution

If other solutions don't work, you might have to use REBUILD to fix the problems with the NSS pool.

Demonstrate how to use REBUILD.



Only use REBUILD as a last resort to recover the file system because it could cause loss of data.

Before using REBUILD, try restoring the pool from a tape backup first. If this does not work, contact Novell for help in using REBUILD.

REBUILD copies errors and transactions into a file named *volume_name.RLF* at the root of SYS. Every time you rebuild an NSS volume, its previous error file is overwritten.

Do not confuse REBUILD with VREPAIR, which is used with traditional volumes.

Before you run REBUILD, deactivate pools and volumes in the pools. Users must not have access to volumes you are rebuilding.

To run REBUILD, do the following:

1. At the server console enter **NSS /POOLREBUILD=poolname**.
2. When prompted that volumes will be dismounted, select **Yes**.
3. When the process is complete, at the server console, enter **MOUNT ALL**.

The REBUILD and VERIFY utilities both generate a log file at the root of the DOS drive.

Use Third-Party Software or Services to Recover Data

If you are unsuccessful using REBUILD, you might need to use third-party tools or services to recover the data. Third-party services, such as those provided by OnTrack Data Recovery, let you send in the disk to be recovered.

These solutions can be expensive, but they are usually less costly than losing the data. The best way to avoid this situation is to consistently back up the data.

After you implement a solution, test the system to make sure the issue is resolved. Continue to monitor the situation to ensure the problem does not recur.



10 minutes

Exercise 5-2 Resolve NSS Error Codes

In this exercise, you encounter an NSS error, research the solution, and fix the problem.

You thought you had enough storage space on the system, but one day you received the following error: 20103 zERR OUT OF SPACE.

Do the following to resolve the issue:

- [Part I: Research an Error Code and Find a Possible Solution](#)
- [Part II: Implement the Solution](#)

Part I: Research an Error Code and Find a Possible Solution

Do the following:

1. From a web browser, access <http://www.novell.com/documentation/lg/nwec/index.html?treeitl.html>.
2. Under **Novell Storage Services (NSS) Error Codes**, select **List of Codes**.
3. Select **20103 zERR OUT OF SPACE**.
4. Record the action recommended in the list:

Part II: Implement the Solution

Do the following:

1. Access the Media Properties page for your NSS logical volumes:
 - a. From ConsoleOne, browse to *your server*.
 - b. Right-click *your server*, then select **Disk Management > NSS Logical Volumes**.
2. Access the NSS Attributes page for volume DATA and determine whether the Salvage Files property is enabled:
 - a. Select **DATA**.

- b. Select **Properties**.
 - c. Verify that **Salvage Files** is selected.
3. Delete unnecessary salvaged files:
 - a. From your workstation, map a drive to volume **DATA**.
 - b. In Windows Explorer, right-click **DATA**; then select **Purge Files**.
 - c. Select **Purge Subdirectories**; then select **Yes**.
 - d. Check volume **DATA** again.

Notice that your data is still there.

(End of Exercise)

Objective 6 Restore a Deleted Logical Volume

There are times when you need to restore a volume that has been accidentally or maliciously deleted.

If you delete a volume, it is removed from the storage pool. However, for a specified amount of time, called the *purge delay time*, you can review and even restore the contents of the volume you deleted.

You must retrieve the volume before the delay time elapses; otherwise, the volume is purged from the system, and you can no longer restore it.



If you delete a storage pool, you delete all volumes in that pool. Volumes deleted in this manner cannot be restored.

The default setting for the purge delay time is 2 days. After this time expires, NSS purges the volume. You can change the purge delay time to extend or reduce the time for the automatic purging cycle.

To change the purge delay time, at the server console enter

NSS /LOGICALVOLUMEPURGEDELAY=*delay time in seconds*



The purge delay change command is not persistent. The parameter is lost if the server is restarted. To make the change permanent, add the command to the server's AUTOEXEC.NCF file.

You can also manually purge deleted volumes. To restore or purge a deleted volume, do the following:

1. Start **ConsoleOne**.
2. Authenticate as **Admin**.
3. Browse to the *server object*.
4. Right-click the *server object*; then select **Properties**.
5. Select **Media > NSS Pools**.
6. Select **Deleted Volumes**.
7. Select one of the following:
 - ❑ **Purge**. Use to immediately purge all deleted volumes.
 - ❑ **Prevent Purge/Allow Purge**. Use to stop the volume purging process.
 - ❑ **Salvage**. Use to restore the deleted volume.
 - ❑ **Refresh**. Use to rescan the volumes that have been deleted and update the list in the panel.
8. Select **Close**.

Objective 7 Describe Storage Area Networks and Network Attached Storage

Storage Area Networks (SANs) and Network Attached Storage (NAS) are newer storage technologies that let you optimize and centralize your data storage as well as use redundant array of independent disks (RAID) to provide efficient file access to your users.

They are quickly replacing traditional server-attached storage configurations, providing better speed and flexibility.

The acronyms SAN and NAS are sometimes confused, but you must know the capabilities of each and the differences between the 2 for later sections in this course where you will implement a SAN for clustering.

To use SAN and NAS, you need to be able to


- [Identify How a SAN Works](#)
- [Identify How NAS Works](#)
- [List SAN and NAS Design and Implementation Considerations](#)

Identify How a SAN Works

Just as a LAN or WAN is a type of network, a SAN is also a type of network, or network architecture. It is designed for high-volume, block-oriented data storage and retrieval.

A SAN is a separate network for storage and is located between a LAN and storage devices as shown in the following:

Figure 5-8



Because the SAN is not part of the LAN, it is not slowed by the normal network traffic on the LAN.

Because of this, a SAN is often used to resolve bandwidth problems as well as to consolidate storage. You can locate your SAN in a remote location, up to 150 kilometers away.

A SAN must be viewed in a different light than a traditional network. While a traditional network uses TCP/IP and tolerates a degree of re-transmission of packets, a SAN uses Fibre Channel and Small Computer Systems Interface (SCSI) protocols to maintain the highest data integrity.

The components of a SAN are

- Fibre Channel or SCSI connection to the LAN
- Master backup servers
- Media backup servers
- RAID devices
- Tape libraries

Backup and transfer of data within the SAN is very fast and efficient because of the high-speed, large-block transfers that Fibre Channel allows.

Vendors for SAN equipment include the following:

- EMC
- Compaq
- Dell
- Net Appliance

Benefits of using a SAN include the following:

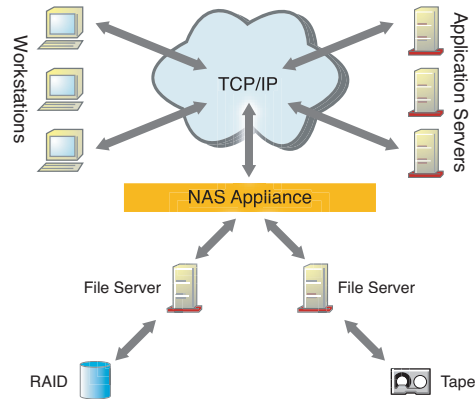
- Centralized storage management
- Dynamic and scalable storage from a pool
- High level of fault tolerance
- Ability to add storage without downtime or disruption
- High-volume data transfers

Identify How NAS Works

While a SAN is type of network, NAS is a product that allows clients to access files directly. Because it is dedicated to storage it allows for much quicker file access than from a general-purpose server.

The following shows the NAS architecture:

Figure 5-9



The NAS is a component, often referred to as a *NAS appliance*, and is usually located between your application servers and your file system. Unlike a SAN it uses typical network protocols such as TCP/IP.

Files are saved or retrieved directly from the NAS appliance. The complexities of the file system are hidden from the user. Because of this, a NAS appliance is often a great place for hosting home directories for users.

Keep in mind that while a NAS appliance stores and retrieves files more quickly than a general-purpose server, network traffic does impact the performance.

Benefits of using NAS include the following:

- NAS is dedicated solely to storage.
- NAS is self-contained, and it is easy to add to an existing LAN.
- NAS is scalable and can support a large number of clients, platforms, and networks.
- NAS supports long-distance data transfers.

List SAN and NAS Design and Implementation Considerations

Consider the following when implementing a SAN or NAS:

- SAN and NAS are not mutually exclusive. You can use both together.
- Because of the complexity of implementing a SAN, you should consider having a dedicated storage management team.
- The line between the capabilities of SAN and NAS are beginning to blur. Many SAN implementations now employ NAS as part of the solution.
- SANs can use a new protocol called Internet SCSI (iSCSI). iSCSI is used to send SCSI information over IP networks and is mainly used to transfer storage information over long distances. This allows for location-independent data storage and management.
- SAN and NAS significantly reduce the overall cost of data storage compared to traditional server-centric storage.
- Don't mix NetWare 5.x and NetWare 6 servers on the same SAN.
- Don't give access to clustered partitions.

Summary

The following is a summary of the objectives in this section:

Table 5-3

Objective	What You Learned
<p>1. Expand an NSS Storage Space</p>	<ul style="list-style-type: none"> ■ Increase the Number of Volumes in a Storage Pool. You can place as many logical volumes in the storage pool as you need. ■ Overbook the Storage Pool. The sum of multiple logical volumes can exceed the pool size. ■ Increase the Size of a Storage Pool. While you can increase the size of a pool, you cannot reduce the size.
<p>2. Configure NSS Volume Attributes</p>	<ul style="list-style-type: none"> ■ Because adjusting these attributes can affect NSS performance, only use the attributes you know you need.
<p>3. Mount a DOS Partition as an NSS Volume</p>	<ul style="list-style-type: none"> ■ Understand DOSFAT.NSS. When you load the module, all DOS FAT partitions on the server are mounted as logical NSS volumes. ■ Mounting a DOS Partition Using DOSFAT.NSS. Make sure you SET AUTO RESTART AFTER ABEND = 0.
<p>4. Use VCU to Create an NSS Volume from a Traditional Netware Volume</p>	<ul style="list-style-type: none"> ■ Understand VCU. Remember that when using VCU.NLM, you must have enough space for both the traditional volume and the new logical volume. You can delete the traditional volume after you know the new logical volume works well. ■ Using Correct VCU Syntax. The /R: variable lets you keep the original volume name of the traditional volume for the new logical volume name.

Table 5-3 (continued)

Objective	What You Learned
5. Resolve Common NSS Errors	<ul style="list-style-type: none"> ■ Use VERIFY to Determine the Integrity of an NSS Pool. The VERIFY utility checks the file system integrity for an NSS pool by searching for inconsistent data blocks or other errors. ■ Use REBUILD as a Last Resort Solution. Only use REBUILD as a last resort to recover the file system. If you use it to recover from data corruption, you could lose data.
5. Restore a Deleted Logical Volume	<ul style="list-style-type: none"> ■ Restore a Deleted Logical Volume. NSS lets you restore a deleted volume within 2 days (default). You can also adjust this interval. Remember, when you delete a pool, you also delete the volumes in the pool and these cannot be restored.
7. Describe Storage Area Networks and Network Attached Storage	<ul style="list-style-type: none"> ■ Describe Differences between a SAN and NAS. A SAN is a type of network; NAS is a product or network component dedicated to storage. ■ List SAN and NAS Design and Implementation Considerations. Remember that SAN and NAS are not mutually exclusive. In fact, they work quite well together.

SECTION 6 Configure and Troubleshoot a RAID Solution Using NSS

Duration: 1 hour

In this section you learn to configure and troubleshoot a RAID solution with NSS.

Objectives

1. Implement RAID 0 with NSS
2. Configure Partition Mirroring and Duplexing in NSS
3. Troubleshoot Software RAID and Mirroring in NSS

Introduction

You are the network administrator at a Digital Airlines branch office. With your servers upgraded to Netware 6, you want to quickly and effectively resolve some server performance and file storage issues.

Specifically, you notice that your servers are experiencing heavy traffic loads and you want to implement Redundant Array of Independent Disks (RAID) Level 0 and stripe data across multiple disk drives to improve server performance.

You are also responsible for preserving and protecting sensitive financial data, the loss of which would cost your job. You learn that you can use partition mirroring to provide a degree of protection from data loss by creating redundant disk sets.

Objective 1 **Implement RAID 0 with NSS**

In this section you learn the following about RAID:

- [What RAID Is](#)
- [Hardware RAID versus Software RAID](#)
- [Hardware and Software RAID Levels](#)
- [How to Configure Software RAID 0 in NSS](#)
- [How to Use NSSMU to Create RAID Arrays](#)

What RAID Is

RAID combines several inexpensive disks into an *array* or *disk posture*.

By creating an array, you can achieve performance and reliability results that exceed that of one large drive.

Hard disks are mechanical devices and slowly wear out with use. Every hard disk has an associated Mean Time Before Failure (MTBF).

To guard against data loss from a disk failure in your server, you can implement RAID.


Implementing RAID spreads information across multiple disk drives to improve performance and reliability.

Hardware RAID versus Software RAID

RAID can be implemented using hardware or software. Hardware RAID makes a disk array look like a single disk or multiple disks to the operating system.

Figure 6-1

The operating system only “sees” what the controller represents:



Software RAID is a set of kernel modules that work together with management utilities to implement RAID solely using software. The modules are layered between the low-level disk drivers and the file system that resides above it.

Examples of a software RAID solution are the NSS RAID modules found in NetWare 6. The NSS RAID modules are independent. Their performance is dependent on server CPU performance and load.

The following compares the advantages and disadvantages of hardware and software RAID:

Table 6-1

Hardware RAID	Software RAID
<ul style="list-style-type: none"> ■ More expensive ■ Operates on entire disk drives ■ Simplified management ■ More RAID level options available ■ Much faster 	<ul style="list-style-type: none"> ■ Less expensive ■ Operates on a partition-by-partition basis ■ More complicated management ■ Fewer RAID level options available ■ Uses processing time

Hardware and Software RAID Levels

There are many ways to implement a RAID solution. The industry has defined several standard implementations called *levels*.

RAID levels differ in performance, redundancy, reliability, cost, and storage capacity.

Some RAID levels offer disk redundancy to protect against data loss. Other RAID levels read and write using multiple disks at the same time to increase performance. Others offer both advantages.

The following provides an overview of some common RAID levels and their functions:

Table 6-2

Level	Technology	Description
0	Disk striping	<ul style="list-style-type: none"> ■ Writes data to multiple disks ■ Enhances performance ■ No fault tolerance
1	Disk mirroring and duplexing (Both duplicate data: mirroring does it while connected to one disk controller; duplexing requires disks connected to 2 disk controllers.)	<ul style="list-style-type: none"> ■ Maintains duplicate copies of all data on 2 drives ■ Fault tolerant
2	Hamming error-correcting code (ECC)	<ul style="list-style-type: none"> ■ Writes error-correcting information to a separate disk drive ■ Ensures data integrity ■ Rarely implemented
3	Parallel transfer with shared parity	<ul style="list-style-type: none"> ■ Stripes data at the byte level across 2 or more drives ■ Stores parity information on a third drive ■ Fault tolerant
4	Independent data disks with shared parity	<ul style="list-style-type: none"> ■ Identical to RAID 3 except data is striped at block level

Table 6-2 (continued)

Level	Technology	Description
5	Independent data disks with distributed parity	<ul style="list-style-type: none"> ■ Stripes data and parity across 3 or more drives ■ Fault tolerant
6	Independent disks with 2-dimensional parity	<ul style="list-style-type: none"> ■ Stripes data and 2 complete copies of parity information across 3 or more drives ■ Fault tolerant ■ Vendor specific
7	Asynchronous RAID	<ul style="list-style-type: none"> ■ Hardware solution consisting of striped data array and a separate parity drive ■ Dedicated operating system coordinates disk storage activities ■ Vendor specific
10	Striping of mirrored disks	<ul style="list-style-type: none"> ■ Combines RAID 0 and RAID 1 ■ Stripes data across mirrored pairs of disks ■ Enhances performance ■ Fault tolerant
53	Striped array of arrays	<ul style="list-style-type: none"> ■ Stripes data across multiple RAID 5 arrays ■ More enhanced performance ■ Fault tolerant ■ Vendor specific
0+1	Mirroring of striped disks	<ul style="list-style-type: none"> ■ Combines RAID 0 and RAID 1 ■ Mirrors data stored on identical striped disk arrays

NSS provides software RAID 0 and software RAID 1.

Most implementations use hardware RAID.

NSS also fully supports RAID arrays created using hardware RAID adapters.

How to Configure Software RAID 0 in NSS

NSS lets you create a RAID 0 device by striping data across multiple drives on your system. This option is a software configuration that emulates an actual hardware RAID 0 system.

A RAID device is set up by securing space from all of your disk drives and then putting segments on the combined space. Data is then sequentially placed or striped on the RAID disks.

The RAID stripe size is the amount of data the file system places on a disk before moving to the next disk. The stripe size ranges between 4 KB and 256 KB in increments of 2 KB.

The size of the stripe units depends on the application for which the array is used. For example, if the system will store large files, such as graphics or digital video, the stripes are generally small, around 512 bytes.

The small size of the stripes ensures that a single file spans as many disks as possible. This ensures that the files can be manipulated quickly because modifying the file will require reading and writing to all disks in the array at the same time.

This configuration occurs at the software level.



Remember that RAID 0 improves and enhances performance but does not provide fault tolerance.



Each segment in the RAID 0 configuration should come from a different device. NSS will let you obtain RAID elements from the same device, but this will severely impede the performance of your file system.

Emphasize that each segment in the RAID 0 configuration should come from a different device or performance suffers.

You can use RAID 0 for both logical and traditional volumes.

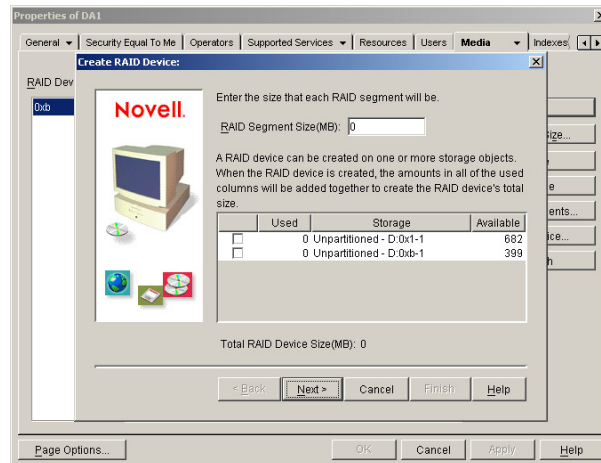
Create a Software RAID 0 Array

To create a software RAID 0 array, do the following:

1. Start ConsoleOne
2. Authenticate as your *admin user*.
3. Browse to your server object.
4. Right-click the *server object*; then select **Properties**.
5. Select **Media > Raid Devices**; then select **New**.

The following appears:

Figure 6-2



6. In RAID Segment Size (MB), enter the amount of *space* you want to secure from each storage device in megabytes.
7. Mark the *devices* to be used in the array; then select **Next**.

The following appears:

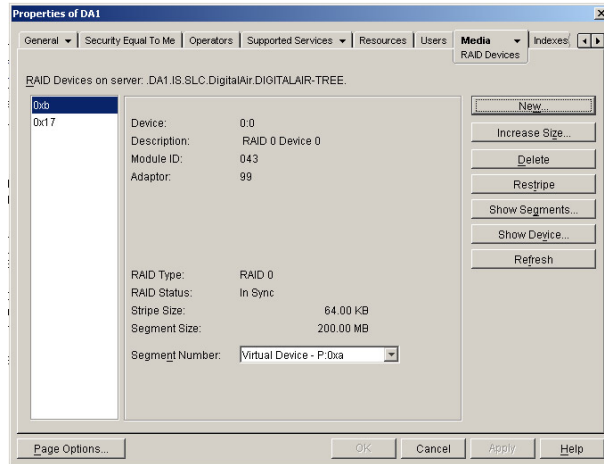
Figure 6-3



8. Select the *stripe size* and *RAID type*.
9. Select **Finish**.
10. When the Server Properties window reappears, select **Media > Raid Devices**.
11. Select the *RAID device* whose size you want to increase.

The following appears:

Figure 6-4



12. Select **Increase Size**; then select the *device* you want to add to the RAID configuration.
13. Select **Finish**.
14. Select **Restripe**.

After you add another physical storage device to your RAID device, restripe the array. This creates stripes on the new device and redistributes the data across all devices.

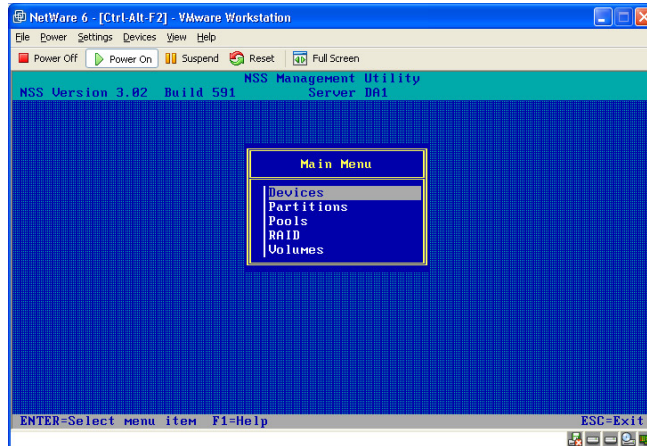


The restriping process takes time to complete, depending on the size of the disk devices involved. As a result, file system performance is impacted during the restriping process.

How to Use NSSMU to Create RAID Arrays

You can also use NSS Management Utility (NSSMU) to create NSS RAID arrays. This is done from the server console, as shown in the following:

Figure 6-5



You can also use Remote Manager (NRM) to configure software RAID but this course uses ConsoleOne.



20 minutes

Exercise 6-1 Configure a Software RAID Solution

This exercise works with 1 or 2 channels.

In this exercise, you how to configure a RAID 0 solution by doing the following:

- Part I: Create a Software RAID 0 Array
- Part II: Create a Storage Pool
- Part III: Create a Logical Volume

Before starting the exercise, make sure both IDE controllers are recognized by the computer BIOS (if you are using 2 controllers) and that a second driver is loaded for the second hard drive.

During this exercise, you use 2 hard drives installed on your NetWare 6 server. Before you start, do the following:

- If there are 2 IDE controllers for the hard drives, check the computer BIOS settings to make sure the computer recognizes both hard drive IDE controllers.
- Using NWCONFIG, make sure there is a driver for each hard drive. You might need to load a driver for the second hard drive.

If you are in a classroom setting, your instructor might have performed these tasks for you. Check with your instructor for additional information.

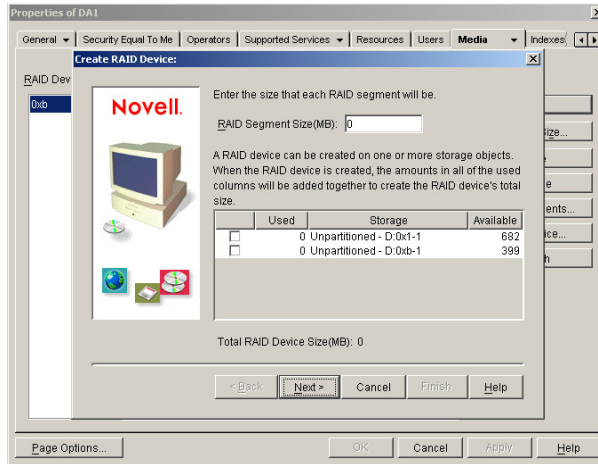
Part I: Create a Software RAID 0 Array

In this part you create a RAID 0 device to stripe data across 2 or more disk drives:

1. Authenticate as **admin** to *your server*.
2. Start **ConsoleOne**.
3. Browse to and right-click *your server object*; then select **Properties**.
4. Select **Media > Raid Devices**; then select **New**.

The following appears:

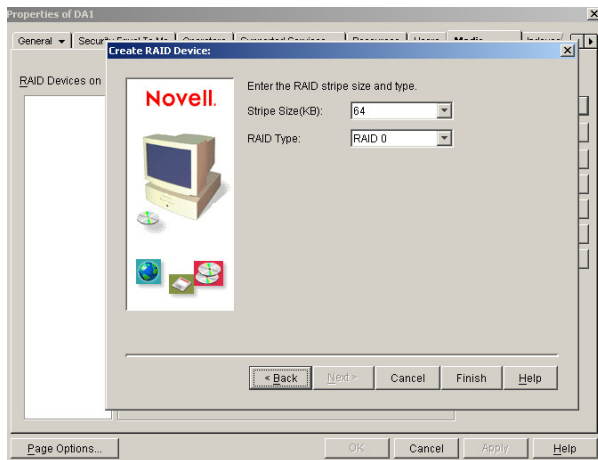
Figure 6-6



5. In the RAID Segment Size (MB) window enter **200**.
6. From the list select your *first hard drive* (storage device) to be used in the array; then select **Next**.

The following appears:

Figure 6-7



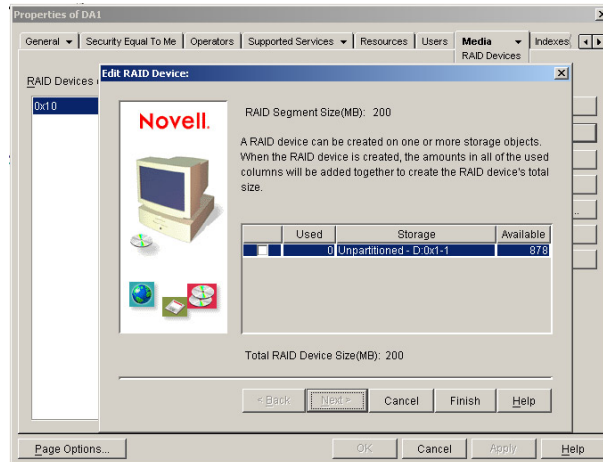
7. Make sure the stripe size is **64 (KB)** and the RAID type is **RAID 0**; then select **Finish**.

The new RAID device is listed with settings shown to the right.

8. Select **Increase Size**.

A screen similar to the following appears:

Figure 6-8



9. Add your *second hard drive* to the RAID configuration by selecting the drive: then select **Finish**.

After you add another physical storage device to your RAID device, restripe the array. This creates stripes on the new device and redistributes the data across all devices.

10. Select **Restripe**.
11. To verify that you have 2 raid device segments, select **Show Segments**.
12. On the right under RAID information, find the **Device ID** and record it to use later in the exercise:

13. Select **Close**.



The restriping process takes some time to complete, depending on the size of the disk devices involved. As a result, file system performance is impacted during the restriping process.

Part II: Create a Storage Pool

In this part you create a storage pool:

1. From the Server Properties window, select **Media > NSS Pools**.
2. Select **New**.
3. For the storage pool name enter **POOL1**; then select **Next**.
4. Select the **RAID device** you just created.

Notice that instead of seeing 2 separate segments, you only see a single unit of unpartitioned space. This space is equal to the sum of the 2 devices together.

5. In the **Used** column, enter **200**; then select **Next**.
6. Make sure **Activate on Creation** is selected.

This option activates your pool and any logical volumes when you create a pool.

7. Select **Finish**.
8. (Conditional) If a warning window appears, select **Yes**.

This warning appears because you are in a classroom environment.

Part III: Create a Logical Volume

Do the following:

1. From the Server Properties window, select **Media > NSS Logical Volumes**.
2. Select **New**.
3. For the volume name enter **VOLUME1** for the volume; then select **Next**.
4. Select **POOL1**.

You can also select unpartitioned space at this point. If you do, NSS creates an NSS partition and a storage pool for your volume.
5. Select **Allow volume quota to grow to the pool size**; then select **Next**.
6. Select **Finish**.
7. From the **DAx** server console enter **VOLUMES** and verify that **VOLUME1** was mounted.

(End of Exercise)

Objective 2 **Configure Partition Mirroring and Duplexing in NSS**

Now that you have learned to configure a software RAID array in ConsoleOne, you are ready to learn how to

- [Configure Partition Mirroring and Duplexing in NSS](#)
- [Troubleshoot Software RAID and Mirroring in NSS](#)

Configure Partition Mirroring and Duplexing in NSS

Mirroring your partitions lets you protect critical data by storing the same data on 2 separate disks using the same disk controller. If one disk goes down, the system uses the other disk.

Duplexing your partitions is essentially the same as mirroring with one main difference: with duplexing, data is stored on 2 disks that are controlled by 2 separate disk controllers.

Using ConsoleOne, you can mirror or duplex both traditional and NSS partitions.

The following is a list of requirements for mirroring partitions:

- Mirrored partitions must have the same partition type as you mirror the partition to. In other words, you can only mirror NSS partitions to other NSS partitions. The same applies to traditional partitions.
- You can only mirror partitions. However, you might want to mirror an entire storage pool. Because a partition can only be a member of one storage pool, the only way to mirror that storage pool is to mirror all partitions the pool resides on.
- To mirror partitions, select an option that makes the partitions compatible for mirroring when you create them—you cannot change that mirroring option after you create a partition.

You can mirror to an existing group or create a new mirror group for the partition. You cannot combine mirror groups (existing groups with multiple mirrored partitions).

- Mirrored partitions must have compatible data area size. This means the new partition must be at least the same size or slightly larger than the other partitions in the group.

For example, the physical size (data and hot fix size combined) shouldn't be more than 2 MB in size difference.

- The file system adjusts the hot fix size to the allowable ranges to make the data area identical to the other partitions in the mirror group.
- Mirrored partitions must both be marked for sharable for clustering.

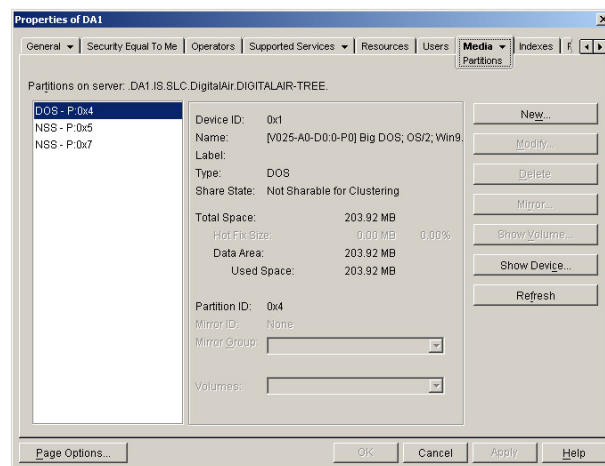
The partitions you add to a group cannot be part of an existing group—they must be individual mirrored objects.

The first task for setting up NSS mirroring or duplexing in NetWare 6 is to create partitions on your storage devices to be mirrored. Do the following:

1. Start ConsoleOne.
2. Right-click the *server* object and select **Properties**.
3. Select **Media > Partitions**.

The following appears:

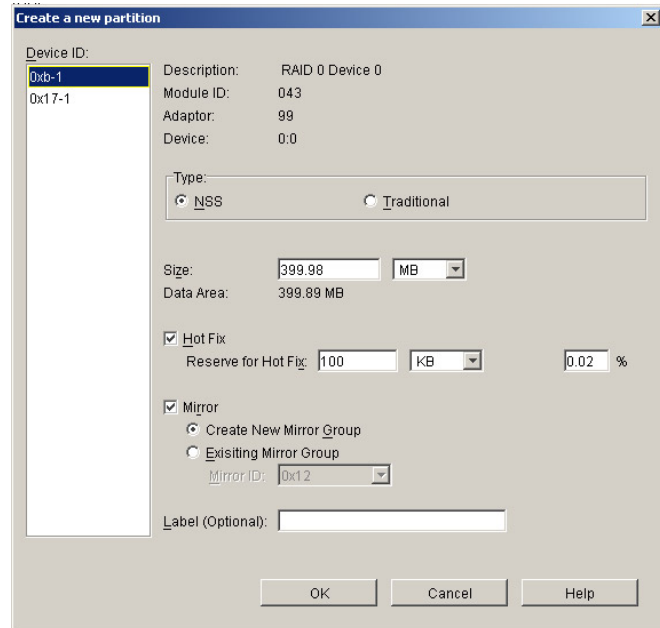
Figure 6-9



4. Select **New**.

The following appears:

Figure 6-10



5. Select a *device* to create the partition on.
6. Select the *type* of partition you want to create as NSS.
7. Enter the *size* of the partition in bytes (B), kilobytes (KB), megabytes (MB), or gigabytes (GB).
8. To reserve space for the hot fix error correction feature, select **Hot Fix** and enter the amount of *space* or *percentage* you want to reserve.

This is required for mirroring to work. If the existing partitions don't have a hot fix area, they can't be mirrored.

Mirrored partitions must be compatible in data area size. This means the new partition must be at least the same size or slightly larger than the other partitions in the group.

The physical size (combined data and hot fix size) of the partition must be at least 100 KB, but no more than 120 MB larger than the data size of the existing partitions in the mirror group.

9. To mirror the partition, select **Mirror**; then select one of the following options:
 - **Create New Mirror.** This option allows the partition to be part of a mirror group. You do not actually create the group until you add another mirrored partition to the partition you are creating.
 - **Existing Mirror Group.** This option shows a list of existing mirrored groups that are compatible in data area size. This option lets you add this new partition to one of the mirrored groups in the list. If you select this option, also select the ID of the mirrored partition.
10. (Optional) Enter a *label* for the partition.
11. Select **OK**.
12. Highlight a *mirrored partition*; then select **Mirror**.

This shows the partitions in the mirror group as well as the status of the mirror group. At this point allow the partitions to complete the mirror.

Troubleshoot Software RAID and Mirroring in NSS

When you troubleshoot mirroring in NSS, remember the following:

- Mirrored partitions must be the same size or within 2 MBs of the same size.
- When creating a partition, you must create a partition with Hot Fix selected. If you create the partition without Hot Fix selected, the partition does not mirror.

**10 minutes**

You should have 2 devices per machine. If that is not possible, demonstrate the procedure for the students.

Exercise 6-2 Mirror an NSS partition

In this exercise, you learn to provide fault tolerance for sensitive data by mirroring an NSS partition. Do the following:

1. From your workstation start **ConsoleOne**.
2. Right-click **DAX**; then select **Properties**.
3. Select **Media > Partitions > New**.
4. Select a device to create a partition on (either the hard drive or RAID).
5. Select the type of partition as **NSS**.
6. For the volume size enter **300 MB**.
7. Make sure **Hot Fix** is selected with a setting of **100 KB**.
Remember that this is required for mirroring to work. If the existing partitions don't have a hot fix area, they can't be mirrored.
8. Make sure **Create New Mirror Group** is selected.
9. In the Label field enter **MIRROR1**.
This is important because it will help you identify the mirrored or duplexed partition later.
10. Select **OK**.
11. On the left, locate and select the partition you just created (look for the **MIRROR1** label on the right).
This partition is probably located near the bottom of the list.
12. Record the **Mirror ID** number:
13. Select **New**.
14. For the size, select **MB** from the drop-down list; then enter **300**.
15. Under Mirror, select **Existing Mirror Group**.
16. From the drop-down list, select your *mirror ID number*.

17. In the Label field, enter **MIRROR2**; then select **OK**.
18. From server **DAx**, switch to the console prompt and notice that the partitions were not synchronized; when remirroring occurred, they were synchronized.
19. Return to the Properties Window and select **Media > NSS Logical Volumes**.
20. Select **New**.
21. In the Name field enter **MIRROR**; then select **Next**.
22. Select the *mirrored space* you just created using the mirror ID number as a guide.
23. Select **Allow volume quota to grow to the pool size**; then select **Next**.
24. In the Name field enter **POOL2**; then select **OK**.
25. Select **Finish**.
26. From your server console prompt list the MIRROR volume by entering **VOLUMES**.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	Summary
1. Implement RAID 0 with NSS	<ul style="list-style-type: none">■ NSS lets you create a RAID 0 device by striping data across multiple drives on your system. This option is a software configuration that emulates an actual hardware RAID 0 system.■ RAID levels differ in performance, redundancy, reliability, cost, and storage capacity.
2. Configure Partition Mirroring and Duplexing in NSS	<ul style="list-style-type: none">■ Mirroring partitions lets you protect critical data by storing the same data on 2 separate disks using the same disk controller. If one disk goes down, you can use the other one.■ Duplexing your partitions is essentially the same as mirroring with one main difference: with duplexing, data is stored on disks that are controlled by separate disk controllers.■ Mirrored partitions must be virtually the same size.■ To mirror a partition, Hot Fix must be selected.

SECTION 7 Perform Advanced iFolder Tasks and Troubleshooting

Duration: 4 hours

In this section you learn advanced iFolder administration tasks and troubleshooting steps.

Objectives

1. [Describe iFolder Configuration Files](#)
2. [Perform iFolder Management Tasks](#)
3. [Maintain and Troubleshoot the iFolder Client](#)
4. [Maintain and Troubleshoot the iFolder Server](#)

Objective 1 Describe iFolder Configuration Files

iFolder is configured each time it launches by reading the following Apache configuration files:

SYS:Apache\iFolder\Server\HTTPD.CONF

SYS:Apache\iFolder\Server\HTTPD_ADDITIONS_NW.CONF

You can use these files to help troubleshoot certain iFolder problems. To change these files, you need to know the following:

- [How to Edit iFolder Configuration Files](#)
- [When to Edit iFolder Configuration Files](#)

How to Edit iFolder Configuration Files

iFolder configuration files are text files. To edit them you use a text editor, make changes, and save the file in the correct location.

iFolder accesses these files when it is launched. So if you make a change, you must stop and start iFolder for changes to take affect.

Before you edit the contents of these files it is useful to know how they are structured.

Two files are used to configure Apache for iFolder:

- [HTTPD.CONF](#)
- [HTTPD_ADDITIONS_NW.CONF](#)

HTTPD.CONF

HTTPD.CONF begins with the following about the file structure:

```
#
# Based upon the NCSA server configuration files
# originally by Rob McCool.
#
# This is the main Apache server configuration file.
# It contains the configuration directives that give
# the server its instructions. See
# <URL: http://www.apache.org/docs/> for detailed
# information about the directives.
#
# Do NOT simply read the instructions in here without
# understanding what they do. They're here only as
# hints or reminders. If you are unsure consult the
# online docs. You have been warned.
#
# After this file is processed, the server will look
# for and process sys:/apache/conf/srm.conf and then
# sys:/apache/conf/access.conf unless you have
# overridden these with ResourceConfig and/or
# AccessConfig directives here.
#
```



```
# The configuration directives are grouped into three
# basic sections:
# 1.Directives that control the operation of the
# Apache server process as a whole (the 'global
# environment').
# 2.Directives that define the parameters of the
# 'main' or 'default' server, which responds to
# requests that aren't handled by a virtual
# host. These directives also provide default
# values for the settings of all virtual hosts.
# 3.Settings for virtual hosts, which allow Web
# requests to be sent to different IP addresses
# or hostnames and have them handled by the same
# Apache server process.
#
# Configuration and log file names: If the filenames
# you specify for many of the server's control files
# begin with "/" (or "drive:/" for Win32 and sys:/
# for NetWare), the server will use that explicit
# path. If the filenames do *not* begin with "/",
# the value of ServerRoot is prepended -- so
# "logs/foo.log" with ServerRoot set to
# "/usr/local/apache" will be interpreted by the
# server as "/usr/local/apache/logs/foo.log".
```

This is the default Apache configuration file. This is a long text file that does not need to be discussed in its entirety to understand how to edit iFolder configuration settings.

HTTPD.CONF is organized into 3 sections:

- Global environment
- Main server
- Virtual hosts

For troubleshooting iFolder, you need to know that each section contains an entry that can affect iFolder:

- In the global environment section, the entry is

```
Listen IP address:80
```

If you look at this file on your DAx server, it looks like this:

```
Listen 192.168.1.x:80
```

- In the main server section, the entry is

```
ServerName IP address
```

If you look at this file on your DAx server, it looks like this:

```
ServerName 192.168.1.x
```

- The virtual hosts section contains the following:

```
<IfModule mod_tls.c>
    SecureListen IP address:443 "SSL
    CertificateIP"
</IfModule>
```

If you look at this file on your DAx server, this entry looks like this:

```
<IfModule mod_tls.c>
    SecureListen 192.168.1.x:443 "SSL
    CertificateIP"
</IfModule>
```



The term *virtual host* refers to the practice of maintaining more than one server on one machine.

HTTPD_ADDITIONS_NW.CONF

The HTTPD.CONF file creates an Apache web server exclusively for iFolder to use. At the end of HTTPD.CONF, you see the following entry:

```
include SYS:\apache\iFolder\Server\
httpd_additions_nw.conf
```

This entry causes the Apache web server configuration to include specific iFolder server configuration information found in HTTPD_ADDITIONS_NW.CONF.

This file is as follows:

```
#
# iFolder Server LoadModule
#
LoadModule ifolderserver_module
"iFolder/Server/iFolder.nlm"

#
# Edit the xxx.xxx.xxx.xxx to your IP address
# =====
<VirtualHost 192.168.1.1:80>
    DocumentRoot "SYS:\apache\iFolder\DocumentRoot"
<Directory "SYS:\apache\iFolder\DocumentRoot">
    Options Indexes FollowSymLinks MultiViews
    AllowOverride None
    Order allow,deny
    Allow from all
</Directory>
<location /iFolderServer>
    SetHandler ifolderserver-form-handler
</location>

#
# iFolder Server LDAP Settings
#
# Edit the LdapHost and LdapSecondaryHost
# if the SSL LDAP port of 636 is used, you must
# enter the path to the LdapRootCert.
# =====
LdapHost dal.digitalair.com
LdapPort 636
LdapLoginDnContext "O=DigitalAir"
# -or -
# LdapLoginDnContext ",ou=xxxx,o=xxxx"

LdapRootCert"SYS:\apache\iFolder\server\
RootCert.der"
```

```

# Suggestion: Enter IP address of the iFolder
# server so that
# you can use NWAdmin to add "shared" accounts

# LdapSecondaryHost %LdapSecondaryHost%
# LdapSecondaryPort %LdapSecondaryPort%
# LdapSecondaryLoginDnContext
# "%LdapSecondaryLoginDnContext%"

# -or -
# LdapSecondaryLoginDnContext ",ou=xxxx,o=xxxx"

# LdapSecondaryRootCert "%LdapSecondaryRootCert%"

#
# iFolder Volume \ directory for user files
#
# Edit the iFolderServerRoot
# =====
iFolderServerRoot SYS:\iFolder

#
# iFolder Admin Settings for Server Management
# Console
#
# Edit the iFolderAdminName
# =====

iFolderAdminName admin

#
# iFolder Server Secure Port
#
# Edit the ServerSecurePort
# =====

ServerSecurePort 443

</VirtualHost>

# Virtual Host for SSL Port 443

#
# Edit the xxx.xxx.xxx.xxx to your IP address
# =====
<VirtualHost 192.168.1.1:443>

    DocumentRoot "SYS:\apache\iFolder\DocumentRoot"

```

```

<Directory "SYS:\apache\iFolder\DocumentRoot">
    Options Indexes FollowSymLinks MultiViews
    AllowOverride None
    Order allow,deny
    Allow from all

</Directory>

<location /iFolderServer>
    SetHandler ifolderserver-form-handler
</location>

#
# iFolder Server LDAP Settings
#
# Edit the LdapHost and LdapSecondaryHost
# if the SSL LDAP port of 636 is used, you must
# enter the path to the LdapRootCert.
# =====
LdapHost dal.digitalair.com
LdapPort 636
LdapLoginDnContext "O=DigitalAir"
# -or -
# LdapLoginDnContext ",ou=xxxx,o=xxxx"
LdapRootCert "SYS:\apache\iFolder\server\
RootCert.der"

# Suggestion: Enter IP address of the iFolder
# server so that you can use NWAdmin to add
# "shared" accounts

# LdapSecondaryHost %LdapSecondaryHost%
# LdapSecondaryPort %LdapSecondaryPort%
# LdapSecondaryLoginDnContext
# "%LdapSecondaryLoginDnContext%"

# -or -
# LdapSecondaryLoginDnContext ",ou=xxxx,o=xxxx"

# LdapSecondaryRootCert "%LdapSecondaryRootCert%"

```

```

#
# iFolder Volume \ directory for user files
#
# Edit the iFolderServerRoot
# =====
iFolderServerRoot SYS:\iFolder

#
# iFolder Admin Settings for Server Management
# Console
#
# Edit the iFolderAdmin lines
# =====

iFolderAdminName admin

#
# iFolder Server Secure Port
#
# Edit the ServerSecurePort
# =====
ServerSecurePort 443

</VirtualHost>

```

The file is divided into 2 sections: the first section gives the configuration of the nonsecure virtual host; the second section gives the configuration of the secure virtual host.

These sections create 2 virtual hosts that are run by the same Apache server. The sample code shown above configures a nonsecure virtual host to use port 80 and a secure virtual host to use port 443, both on the same server.

Aside from the different port numbers, the configurations for each virtual host are identical. For iFolder to function, each virtual host must have the same login contexts, server name, admin name, and server root location.

In most cases you edit this file by using the iFolder server management console, which writes changes to this file. In some cases you must make the changes manually to the file.

When to Edit iFolder Configuration Files

You edit the Apache iFolder configuration files in the following cases:

- Your organization grows and new organization or organizational units are added to your eDirectory tree. You must add login contexts to accommodate the new contexts.
- You did not configure sufficient login contexts when you installed iFolder.
- You must authorize additional admin users to accommodate the increased size of your organization.
- Your iFolder server needs its own IP address so you can use ports 80 and 443 without causing port conflicts with other web-based services.
- Your organization has changed its IP addressing scheme.
- iFolder was installed without changing the default data location from volume SYS to a more suitable volume.

You have 2 options for making changes to the configuration files:

- [Use the iFolder Server Management Console](#)
- [Make Manual Changes to the Configuration Files](#)

Use the iFolder Server Management Console

You use the iFolder server management console to make changes to the following:

- **Login Contexts** indicate the containers that iFolder users are located in. The sample code shown above allows users in the O=DIGITALAIR container to log in to iFolder and have accounts created for them. No other users can log in.

There are 2 ways to add more user contexts in the iFolder server management console. You can individually specify each container or you can use the Search Subcontainer option.

You add login contexts from the LDAP configuration page in the iFolder server management console.

When you individually specify each container, you enter the contexts in the DN field, as shown in the following:

Figure 7-1

LDAP	
Primary LDAP Host	da4.digitalairlines.com
Port	636
DN	OU=DEL,O=DIGITALAIR,O=DIGITALAIR

The syntax you use for entering contexts must follow LDAP conventions, which are different from eDirectory conventions.

eDirectory requires that contexts be written with period (.) separators, like this:

```
OU=IS.OU=SLC.O=DIGITALAIR
```

LDAP requires that contexts be written with comma (,) separators, like this:

```
OU=IS,OU=SLC,O=DIGITALAIR
```

If you enter more than one context in the DN field, separate each context with a semicolon (;) and no spaces, as shown in the following example:

```
O=DIGITALAIR;OU=SLC,O=DIGITALAIR;OU=IS,OU=SLC,O=DIGITALAIR
```

If you only have 1 or 2 contexts to configure for iFolder users, manually specifying them is the easiest way to accomplish this.

If you want to configure all contexts that fall under a certain container, you can save time by using the Search Subcontainer option.

You select this option by selecting the Search Subcontainer box, as shown in the following:

Figure 7-2

LDAP	
Primary LDAP Host:	dsa4.digitalairlines.com
Port:	636
DN:	OU=DEL,O=DIGITALAIR,O=DIGITALAIR
Search subcontainers:	<input checked="" type="checkbox"/>
Trusted Root Certificate:	SYS:\apache\iFolder\server\RootCert.der

In iFolder 1.01, selecting this option only applies to the first context listed in the DN field.

The iFolder server management console writes this change in the HTTPD_ADDITIONS_NW.CONF file by placing a wildcard asterisk (*) in front of the first context in the DN field.

For example, if you enter your contexts in the following order:

```
OU=SLC , O=DIGITALAIR ; O=DIGITALAIR
```

the Search Subcontainer option applies to all contexts below the SLC container.

The change that is written to both virtual server sections of HTTPD_ADDITIONS_NW.CONF looks like this:

```
LdapLoginDnContext
"*OU=SLC,O=DigitalAir;O=DigitalAir"
```

Using the Search Subcontainer option requires that [Public] be given inheritable Compare and Read rights to the CN property of the first container in the DN field.

- **Admin names** indicate the names of users that can log in to the iFolder server management console.

You add authorized administrators from the Admin Sessions page in the iFolder server management console, as shown in the following:

Figure 7-3



To indicate additional admin names you enter them in the Authorized Admins field. They must be separated by a semicolon (;) and no spaces.

Changes you make to this field are written to HTTPD_ADDITIONS_NW.CONF.

For example, if you added CVALDEZ to the Authorized Admins field, the change to both virtual server sections would look like this:

```
iFolderAdminName admin;cvaldez
```

Make Manual Changes to the Configuration Files

You must make changes to HTTPD_ADDITIONS_NW.CONF manually if you need to change the following:

- **Server name** changes must be made to both virtual server sections of HTTPD_ADDITIONS_NW.CONF.

For example, to change the iFolder server name from 192.168.1.1 to 192.168.1.11, you would change the following entries:

```
<VirtualHost 192.168.1.1:80>  
<VirtualHost 192.168.1.1:443>
```

to

```
<VirtualHost 192.168.1.11:80>  
<VirtualHost 192.168.1.11:443>
```

You do not change the LDAP host information because your LDAP server is not changing names or IP addresses.

You must also change the following entries in HTTPD.CONF:

```
Listen 192.168.1.1:80  
ServerName 192.168.1.1  
SecureListen 192.168.1.1:443 "SSL CertificateIP"
```

to

```
Listen 192.168.1.11:80  
ServerName 192.168.11.1  
SecureListen 192.168.11.1:443 "SSL CertificateIP"
```

Additional steps for changing the server name are shown in the exercise that follows.

- **Server root location** indicates where the iFolder server saves user data. Changes must be made to both virtual server sections of HTTPD_ADDITIONS_NW.CONF.

For example, to change the location of iFolder user data from volume SYS to volume DATA, you must change the following:

```
iFolderServerRoot SYS:\iFolder  
  
to  
  
iFolderServerRoot DATA:\iFolder
```

This entry exists in both the nonsecure and secure virtual host sections of HTTPD_ADDITIONS_NW.CONF. You must make the change to both entries.

You must also manually move the SYS:\IFOLDER directory to DATA:\IFOLDER.

Objective 2 Perform iFolder Management Tasks

To manage iFolder you must be able to perform the following tasks:

- [Stop and Start the iFolder Server](#)
- [Set iFolder Client and Server Policies](#)
- [Change the Location of iFolder User Data](#)
- [Add Contexts](#)
- [Add Additional Administrators](#)
- [Change the iFolder Server IP Address](#)

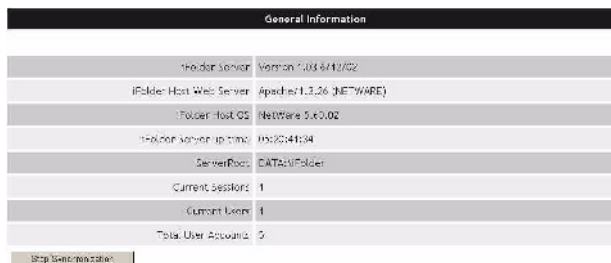
Stop and Start the iFolder Server

After making changes to HTTPD.CONF or HTTPD_ADDITIONS_NW.CONF you must stop and restart iFolder before the changes can take affect.

Before you stop iFolder you should stop the iFolder server from synchronizing. This task is performed from the iFolder server management console.

The Stop Synchronization button is found on the General Information page, as shown:

Figure 7-4



After you stop synchronization, you stop the iFolder server at the server console by entering the following command, which immediately terminates the iFolder service:

STOPIFOLDER

If you use the STOPIFOLDER command without stopping synchronization first the server won't close all connections with the clients before it stops.

Doing the stop synchronization first insures that all iFolder data files are closed when the server goes down. It also makes bringing the server back up much faster.

If you do not stop synchronization first, the next time iFolder is started the Apache console screen reports an error that iFolder was not shut down properly.

The server must then perform a data integrity check of all iFolder data on the server. On an iFolder server that hosts a lot of data, this check can take 45 minutes or longer to perform.

STOPIFOLDER runs the STOPIFOLDER.NCF file that is located in SYS:\APACHE\iFOLDER\SERVER.

To start the iFolder server, at the server console enter

STARTIFOLDER

This also runs an NCF file located in
SYS:\APACHE\IFOLDER\SERVER.

The iFolder installation places the STARTIFOLDER command in AUTOEXEC.NCF so that iFolder is launched each time the server is started.

When STARTIFOLDER is run the Apache for NetWare screen is added to the list of server console screens. You can view this screen to verify that iFolder initialized successfully.

You can also view this screen to view the log of iFolder events that have taken place during the current iFolder session.

(If you have problems starting iFolder, enter NVXADM DN at the server console prompt, to down the Apache server; then enter NVXADMUP to start the Apache server. After the Apache server is running again, enter STARTIFOLDER.)

Set iFolder Client and Server Policies

The iFolder server management console lets you set several client and server policies. These policies affect the way users interact with iFolder and the way servers perform.

iFolder Client Policies

The Client Policies page, found in the iFolder server management console, is shown in the following:

Figure 7-5

Client Policies						
Encryption	<input checked="" type="checkbox"/> On	<input checked="" type="checkbox"/> Enforced	<input checked="" type="checkbox"/> Hidden			
Save Password	<input checked="" type="checkbox"/> On	<input checked="" type="checkbox"/> Enforced	<input checked="" type="checkbox"/> Hidden			
Save Pass Phrase	<input checked="" type="checkbox"/> On	<input checked="" type="checkbox"/> Enforced	<input checked="" type="checkbox"/> Hidden			
Automatic Sync	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Enforced	<input type="checkbox"/> Hidden			
Sync To Server Delay	<input type="text" value="5"/> sec	<input type="checkbox"/> Enforced	<input type="checkbox"/> Hidden	Min: <input type="text" value="1"/> sec	Max: <input type="text" value="60000"/> sec	
Sync From Server Interval	<input type="text" value="30"/> sec	<input type="checkbox"/> Enforced	<input type="checkbox"/> Hidden	Min: <input type="text" value="1"/> sec	Max: <input type="text" value="60000"/> sec	
Conflict In Space	<input type="text" value="25"/> MB	<input type="checkbox"/> Enforced	<input type="checkbox"/> Hidden	Min: <input type="text" value="1"/> MB	Max: <input type="text" value="100"/> MB	

iFolder Server Policies

The Server Policies page, found in the iFolder server management console, is shown in the following:

Figure 7-6

Server Policies	
Initial Client Quota	<input type="text" value="300"/> MB
Session Timeout	<input type="text" value="60"/> minutes
Debug Output	<input type="checkbox"/>

Change the Location of iFolder User Data

The default location for iFolder user data is volume SYS, but it is rarely advisable to use volume SYS to store user data. iFolder users can potentially fill up the volume and bring down the server.

To change the location, do the following:

1. Open SYS:\APACHE\IFOLDER\SERVER\HTTPD_ADDITIONS_NW.CONF.

2. For both virtual host sections, edit the iFolderServerRoot entry to indicate the new location for user data.

For example, if your new location for user data is volume DATA, the entries should look like the following:

```
iFolderServerRoot: DATA:\iFolder
```

3. After you save the file, stop iFolder; then move the iFolder directory from the root of SYS to the root of DATA.
4. Start iFolder; then verify that your changes have taken effect.

Add Contexts

You add user login contexts through the iFolder server management console by doing the following:

1. Access the iFolder server management console through **https://iFolder server IP address** or **https://DNS name/iFolderServer/Admin**.

This URL is case sensitive.

2. Log in as an authorized admin; then select **LDAP**.
3. Make changes to the DN field; then select **Update**.

Changes made to the DN field are reflected in HTTPD_ADDITIONS_NW.CONF. The configuration remarks out the old lines that configured DN contexts and creates a new line for the new configuration.

For example, if your original DN context configuration only included the context O=DIGITALAIR, the line in the configuration file looks like the following:

```
LdapLoginDnContext "O=DigitalAir"
```


When you use the iFolder server management console to add OU=SLC,O=DIGITALAIR as a DN context, the configuration file is changed to look like the following:

```
#LdapLoginDnContext "O=DigitalAir"  
LdapLoginDnContext "O=DigitalAir;OU=SLC,O=DigitalAir"
```

Each time you make a change, the previous configuration is remarked out and a new line is added. If you make many changes, HTTPD_ADDITIONS_NW.CONF can become so long that iFolder will not launch.

You can avoid this by making a backup copy of the original that you can use to replace the overgrown file. You can also open the file and delete the extra remarked out lines, but you should have a backup copy in case you make a mistake.

Add Additional Administrators

You can add additional administrators through the iFolder server management console by doing the following:

1. Access the iFolder server management console through **https://iFolder server IP address** or **https://DNS name/iFolderServer/Admin**.

This URL is case sensitive.

2. Log in as an authorized admin; then select **ADMIN SESSIONS**.
3. Make changes to the Authorized Admins field; then select **Update**.

These changes are also reflected in HTTPD_ADDITIONS_NW.CONF. Old lines are remarked out and new lines are added with each change.



Admin user names and passwords cannot contain special characters. All characters must belong to the UTF 8 character set.

Change the iFolder Server IP Address

If you installed the iFolder service to use the server's primary IP address, you might decide later that it would be better if iFolder had its own IP address.

This might happen if port conflicts occur between iFolder and other services that must use the primary IP address. iFolder can use a secondary IP address so it is logical to change the IP address that iFolder is using.

The iFolder server IP address cannot be changed through the iFolder server management console. All necessary changes must be made manually in the Apache iFolder configuration files.

`SYS:\APACHE\IFOLDER\DOCUMENTROOT\IFOLDERCLIENT.EXE` must be modified to deliver the new IP address when new users download it.

Changes also need to be made to `AUTOEXEC.NCF`. If you are using DNS names but do not have a DNS server configured and running, changes also need to be made to `SYS:\ETC\HOSTS`.

Do the following:

1. Change the necessary entries in `HTTPD.CONF`.
2. Change the necessary entries in `HTTPD_ADDITIONS_NW.CONF`.

3. Modify iFolderClient.exe to use the new IP address:
 - a. At the server console, enter **FIXUP IP ADDRESS [FILE PATH]**.

For example, if you are changing the IP address to 192.168.1.11, enter the following:

```
FIXUP 192.168.1.11 SYS:\APACHE\IFOLDER\DOCUMENTROOT\IFOLDERCLIENT.EXE
```
 - b. View the Logger screen to verify that FIXUP loaded and ran without errors.
4. Configure the new IP address as a secondary IP address on the server that is running the iFolder service:
 - a. At the server console enter **ADD SECONDARY IPADDRESS IP address**
 - b. Edit AUTOEXEC.NCF by placing the ADD SECONDARY IPADDRESS command after the BIND IP command (or after the INITSYS.NCF command).

Placing the ADD SECONDARY IPADDRESS command in AUTOEXEC.NCF ensures that the secondary IP address will be configured each time the server is restarted.



The ADD SECONDARY IP ADDRESS command can also cause a “race” condition on some systems. If this occurs add a PAUSE command after ADD SECONDARY IP ADDRESS.

5. (Conditional) If DNS is not configured on your network and if you are using a DNS name for your iFolder server, place an entry in the SYS:\ETC\HOSTS file to resolve the DNS name to the new IP address.

**2 hours**

This exercise must be completed successfully on all servers that will be cluster enabled.

Exercise 7-1 Perform Advanced iFolder Management Tasks

In this exercise you perform the following tasks:

- Part I: Install iFolder
- Part II: Reapply Support Pack 2
- Part III: Access the iFolder Server Management Console
- Part IV: Stop and Start the iFolder Server
- Part V: Download the iFolder Client and Create an Account for the Admin User
- Part VI: Change the iFolder Server Data Location
- Part VII: Add User Account Contexts
- Part VIII: Enable Subcontainer Search
- Part IX: Assign Additional iFolder Administrators
- Part X: Change the iFolder Server IP Address and Port Number
- Part XI: Change the DNS Name in iFolderClient.exe
- Part XII: Configure Client and Server Policies

Part I: Install iFolder

Do the following:

1. Mount the **NetWare 6 CD** on the server.
2. From the graphical console, select **Novell > Install**.
3. Select **Add**.
4. Browse to the NetWare 6 volume and highlight **PRODUCT.NI**; then select **OK**.
5. Select **OK** again to run the product installation program.
6. From the Components screen, select **Clear All**.

7. Select **Novell iFolder Storage Services**; then select **Next**.
8. Authenticate as **admin**.
9. From the Configure IP-Based Services window, leave **Single IP Address** selected; then select **Next**.
10. From the LDAP Configuration window, select **Next**.
11. From the iFolder Server Options window, select **Next**.
12. From the Summary window, make sure iFolder is in the list of products to be installed; then select **Finish**.
13. Complete the installation by selecting **Close**.
14. Restart the server.

Part II: Reapply Support Pack 2

You have installed iFolder 1.01 from the NetWare 6 OS CD. Support Pack 2 updates installed products and updates iFolder to version 1.03.

If you install a product from the original OS CD after applying a support pack, you must reapply the support pack.

Do the following:

1. Mount the Support Pack CD as a NetWare volume on DAX.
2. At the server console, enter **NWCONFIG**.
3. In Configuration Options, select **Product Options**.
4. In Other Installation Actions, select **Install a Product Not Listed**.
5. Continue by pressing **Esc**.
6. To specify the directory path, press **F3**.
7. In Specify a Directory Path, change the default setting to **NW6SP2:** (include the colon).

8. Press **Enter**.
9. In the Novell Terms and Conditions screen, press **Esc** to continue.
10. Accept the license agreement by selecting **Yes**.
11. In the License Agreement for JReport Runtime JInfonet software, press **Esc** to continue.
12. Accept the license agreement for JReport Runtime by selecting **Yes**.
13. Install NetWare Support Pack 6.0.2 by pressing **Enter**.
14. On the Backup Files Replaced by NetWare Support Pack screen, select **No**.
15. On the Do You Want to Update the Storage/LAN/PSM/WAN Drivers Currently in Use screen, select **Yes**.
16. To reboot your server after the file copy, select **Yes**.
17. In the Warning screen, press **Enter** to continue.
18. (Conditional) If prompted, authenticate as **admin** using your full context and password; then allow files to copy and your server to reboot.
19. When asked, do not press any key to exit.

Part III: Access the iFolder Server Management Console

Do the following:

1. At the workstation, use a supported browser to access **<https://DAX.DIGITALAIRLINES.COM:52443/iFolderServer/Admin>**.
The iFolderServer/Admin part of this URL is case sensitive.
2. From the Security Alert window, select **Yes**.
3. Authenticate as **admin**.

The admin username is case sensitive. Be sure to use all lowercase characters.

4. Familiarize yourself with all management options in the menu on the left, but don't make any configuration changes.
5. When you are done exploring, return to the **General Info** page.

Part IV: Stop and Start the iFolder Server

Do the following:

1. From the iFolder server management console, at the bottom of the General Info page, select **Stop Synchronization**.
2. Close the browser.
3. To stop iFolder, at the DAx server console prompt enter **STOPIFOLDER**.
4. Verify that the iFolder server is no longer running:
 - a. Press **Ctrl + Esc**.
 - b. Note that the Apache for NetWare screen is not listed.
 - c. To return to the server console, enter **1**.
5. Verify that the iFolder web site is not functioning:
 - a. At the workstation, launch **Internet Explorer**.
 - b. For the URL enter **http://DAx.DIGITALAIRLINES.COM:52080**.
You might receive a message that the page cannot be displayed or nothing at all appears.
6. To start iFolder again, at the DA1 server console enter **STARTIFOLDER**.

Note that Apache for NetWare screen states that iFolder initialized successfully.

7. Verify that the iFolder web site is functioning:
 - a. At the workstation in Internet Explorer, enter **http://DAX.DIGITALAIRLINES.COM:52080**.
You should see the iFolder page from which you can download the iFolder client.
 - b. Remain at this point to do the next part of this exercise.

Part V: Download the iFolder Client and Create an Account for the Admin User

Do the following:

1. Download the iFolder client:
 - a. From the iFolder web page, select **Download**.
 - b. Select **Open**.
 - c. Select **Run This Program from its Current Location**; then select **OK**.
 - d. From the Security Warning window, select **Yes**.
 - e. From the Welcome screen, select **Next**.
 - f. Choose your language, then select **Next**.
 - g. Accept the license agreement by closing the browser window and selecting **Yes**.
 - h. From the Choose Destination Location window, select **Next**.
 - i. When prompted, select **Finish**.
 - j. Close the **ReadMe** file.
 - k. Restart the workstation by selecting **Finish**.
2. Create an account for Admin:
 - a. When the workstation restarts, authenticate to eDirectory.
 - b. When the iFolder window appears, select **Continue**.
 - c. In the Login window, enter **admin** for the username and **novell** as the password; then select **Login**.

- d. Make sure the server name is **DAx.DIGITALAIRLINES.COM:52080**; then select **Login**.
 - e. Select both **Enable Auto Start** and **Encrypt Files**; then select **OK**.
 - f. In the next window, enter **novell** in both Pass Phrase fields.
 - g. Select **Remember Pass Phrase**; then select **OK**.
- Note that an admin Home shortcut folder is placed on the desktop.

Part VI: Change the iFolder Server Data Location

Do the following:

1. From the workstation, map a drive to the root of your server's volume SYS.
2. Use Notepad to browse to and open **SYS:\APACHE\IFOLDER\SERVER\HTTPD_ADDITIONS_NW.CONF**.
3. Change the data location from SYS to DATA:
 - a. In the nonsecure virtual host configuration portion of this file, find iFolderServerRoot SYS:\iFolder.
 - b. Change this entry to **iFolderServerRoot DATA:\iFolder**.
 - c. Repeat this change for the same entry in the secure virtual host configuration portion of this file.
 - d. Save the file and exit Notepad.
4. Create data files in Admin's iFolder home directory:
 - a. Using the shortcut on the desktop, open Admin's iFolder home directory.
 - b. Inside the folder, right-click; then select **New > Bitmap Image**.
 - c. Repeat Steps **a** and **b** to create any other files.

- d. Right-click the **iFolder trayapp icon**; then select **Sync Now**.
- e. Log out of iFolder.
5. At the DAX server console, stop iFolder.
6. Move the iFolder directory from SYS to DATA:
 - a. Use Windows Explorer to browse to **DAX_SYS:\iFolder**.
 - b. Right-click **iFolder**; then select **Copy**.
 - c. Browse to **DAX_DATA**; then right-click and select **Paste**.
7. At the DAX server console start iFolder.
8. View the Apache for NetWare screen and note the following entry:
iFolder Server Root: DATA:\iFolder
9. At the workstation, log in to iFolder again, synchronize your files, and verify that the data files you created are still there.

Part VII: Add User Account Contexts

Do the following:

1. View the current iFolder user account contexts:
 - a. At the workstation, use Internet Explorer to access the iFolder server management console at **https://DAX.DIGITALAIRLINES.COM:52443/iFolderServer/Admin**.
This URL is case sensitive.
 - b. Log in with the username **admin** and the password **novell**.
The username is case sensitive.
 - c. From the left column, select **LDAP**.
2. Attempt to log in as a user from xxx.Digitalair:
 - a. Right-click the **iFolder trayapp icon**; then select **Logout**.

- b. Right-click the **iFolder trayapp icon** again; then select **Login**.
- c. In the User ID field, enter the *username* (see Table 7-1); then use the same steps and password you used when you created Admin's iFolder account.

Table 7-1

Server	Username	Context
DA4	KSINGH	DEL
DA5	KWILDE	LGA
DA6	LMORGAN	LON
DA7	EKING	SYD
DA8	WMOZART	TXL
DA9	KHIRATA	TYO

When you attempt to log in you will get an invalid password error.

- d. Close the invalid password message by selecting **OK**; then close the Login dialog.
3. Add your context to the iFolder user account context:
 - a. In the LDAP configuration DN field in the iFolder server management console, place your cursor after **O=DIGITALAIR** without entering a space.
 - b. Enter **;OU=*your_context*,O=DIGITALAIR** without entering any spaces (see Table 7-1 for your context).
 - c. Select **Update**.
 - d. At the top of the page, look for the Are You Sure You Want to Update Your LDAP Settings prompt; then select **Yes**.
 - e. Do not close the iFolder server management console.
 4. Attempt to log in as the user from *xxx.Digitalair* again.
This attempt now succeeds.

Part VIII: Enable Subcontainer Search

You want to enable all users in *xxx.Digitalair* and its subcontainers to create iFolder accounts, but you do not want to manually enter all the contexts.

Do the following:

1. Prove that you cannot create an account for a user in one of your subcontainers by attempting to log in as a user from the CUSTRSVC container (see Table 7-2 for your username).

Table 7-2

Server	Username
DA4	JKURMI
DA5	SDEES
DA6	KCHILDS
DA7	BFULLER
DA8	HWAARLE
DA9	KCHUNG

This attempt should fail.

2. Enable subcontainer search for all contexts below *xxx.DIGITALAIR*:
 - a. Because the Search Subcontainer option only applies to the first context in the list, change the order the contexts appear in the DN field to the following:
**OU=*your_context*,O=DIGITALAIR;OU=IS,
 OU=*your_context*,O=DIGITALAIR**
 - b. In the iFolder server management console, select **Search Subcontainers**.
 - c. Select **Update**; then select **Yes**.
 - d. Do not close the iFolder server management console.

Students might miss this step. If they do they will experience problems with the rest of the exercise.

After the first student accomplishes this task, no one else needs to do it.

3. Attempt to log in to iFolder as the user from CUSTSVC again.
This attempt also fails because using Subcontainer Search requires you to give PUBLIC inheritable CN property rights to O=DIGITALAIR.
4. (Only one student should perform this step.) Give PUBLIC inheritable CN property rights to O=DIGITALAIR:
 - a. In ConsoleOne, right-click **DIGITALAIR**; then select **Trustees of This Object**.
 - b. Select **Add Trustee**.
 - c. Select **[PUBLIC]**; then select **OK**.
 - d. Select **Add Property**.
 - e. Select **Show All Properties**.
 - f. Select **CN**; then select **OK**.
 - g. Select **Inheritable**; then select **OK**.
 - h. Select **Apply**; then select **Close**.
5. Attempt to log in to iFolder as the user from CUSTSVC again.

Part IX: Assign Additional iFolder Administrators

Do the following:

1. Add the user from Table 7-1 as an authorized iFolder administrator:
 - a. From the left column in the iFolder server management console, select **Admin Sessions**.
 - b. In the Authorized Admins field, place your cursor after **admin**.
 - c. Without entering any spaces, enter **;username**.
For example, for username KSINGH you would enter **;KSINGH**. The semicolon is required.



Remember that the case you use here is the case you must use when logging in to the iFolder server management console.

- d. Select **Update**.
- 2. Verify that the user you added can access the iFolder server management console:
 - a. From the left column, select **Login/Logout**; then from the right select **Logout**.
 - b. From the left column, select **Login/Logout**.
 - c. Enter *username* and use **novell** as the password.
 - d. Select **Admin Sessions**.
 - e. Note that the Current Admin is *username*.
 - f. Log out of the iFolder server management console.

Part X: Change the iFolder Server IP Address and Port Number

Do the following:

- 1. At the workstation, log out of iFolder.
- 2. Add a secondary IP address that the iFolder server can use:
 - a. At the DA x server console, enter **ADD SECONDARY IPADDRESS 192.168.1.x** (see Table 7-3 for the IP address).

Table 7-3

Server	Secondary IP Address
DA4	192.168.1.34
DA5	192.168.1.35
DA6	192.168.1.36
DA7	192.168.1.37

Table 7-3 (continued)

Server	Secondary IP Address
DA8	192.168.1.38
DA9	192.168.1.39

- b. Add this command to AUTOEXEC.NCF so the secondary IP address is added each time the server is restarted. Place it between the BIND IP and MOUNT ALL statements.
3. Change the IP address and DNS name that are configured in HTTPD.CONF:
 - a. At the workstation, open **SYS:\APACHE\IFOLDER\SERVER\HTTPD.CONF** and search for **192.168.1.x**.
For example, if you are using server DA4, search for 192.168.1.4. You will find 2 instances: the first indicates port 52080 and the second indicates port 52443.
 - b. Change the first instance you find to **192.168.1.x:80**.
For example, if you are using 192.168.1.34 as your secondary IP address, change the IP address and port to 192.168.1.34:80
 - c. Change the second instance you find to **192.168.1.x:443**.
 - d. Find the following lines
ServerAdmin admin@DAx.DIGITALAIRLINES.COM
ServerName DAx.DIGITALAIRLINES.COM
and change them to
ServerAdmin admin@iFolderx.DIGITALAIRLINES.COM
ServerName iFolderx.DIGITALAIRLINES.COM.
 - e. Save the file.
 4. Change the port number that are configured in HTTPD_ADDITIONS_NW.CONF:

- a. Open **SYS:\APACHE\IFOLDER\SERVER\HTTPD_ADDITIONS_NW.CONF**.
 - b. Find the following lines
<VirtualHost DAx.DIGITALAIRLINES.COM:52080>
<VirtualHost DAx.DIGITALAIRLINES.COM:52443>
and change them to
<VirtualHost iFolderx.DIGITALAIRLINES.COM:80>
<VirtualHost iFolderx.DIGITALAIRLINES.COM:443>
 - c. Find both instances of
ServerName DAx.DIGITALAIRLINES.COM
and change them to
ServerName iFolderx.DIGITALAIRLINES.COM
 - d. Find both instances of
ServerSecurePort 52443
and change them to
ServerSecurePort 443
 - e. Save the file.
5. Change the DNS name that is configured in **IFOLDER_NAV.HTML**:
- a. Use Notepad to open **SYS:\APACHE\IFOLDER\DOCUMENTROOT\HTML\IFOLDER_NAV.HTML**.
 - b. Find the line that contains
“https://DAx.digitalairlines.com:52443/applet/java.htm”
and change it to
“https://iFolderx.digitalairlines.com:443/applet/java.htm”
 - c. Save the file.

If the iFolder service starts with the error that it cannot resolve the iFolder server's hostname, make sure RESOLVE.CFG is not empty.

If RESOLVE.CFG is empty have students copy the contents from another server where the file is not empty.

6. Stop and start the iFolder service.
7. Log in to iFolder as any of the users you have already logged in as:
 - a. At the workstation, right-click the **iFolder Trayapp icon**; then select **Login**.
 - b. In the Novell iFolder Login window change the server name to **iFolderx.DIGITALAIRLINES.COM** (remove the port number); then select **Login**.

iFolder is now running on the new IP address and is using the standard ports: 80 and 443.

Part XI: Change the DNS Name in iFolderClient.exe

Do the following:

1. Use FIXUP.NLM to assign a new DNS name to iFolderClient.exe:
 - a. At the DA_x server console enter **FIXUP iFolderx.DIGITALAIRLINES.COM SYS:\APACHE\IFOLDER\DOCUMENTROOT\IFOLDERCLIENT.EXE**.
 - b. View the Logger screen to verify that FIXUP loaded and reported no errors.
2. Remove iFolder from the workstation so you can verify that iFolderClient is using the new IP address:
 - a. Log out of iFolder.
 - b. Delete all iFolder home directories and their shortcuts.
 - c. Select **Start > Settings > Control Panel > Add/Remove Programs**.
 - d. Select **Novell iFolder 1.03**.
 - e. Select **Change/Remove**; then select **Yes**.
 - f. After uninstall completes, restart the workstation.

3. Access iFolder using the new IP address:
 - a. After the workstation restarts, use Internet Explorer to go to **http://iFolderx.DIGITALAIRLINES.COM**.
 - b. Download and install the iFolder client from this web site.
 - c. Restart the workstation.
 - d. When prompted select **Continue**.
Note that the Server field contains the new iFolder DNS name.
 - e. Log in as **admin** with the password and pass phrase **novell**.

Part XII: Configure Client and Server Policies

Do the following:

1. Access the iFolder server management console at **http://iFolderx.DIGITALAIRLINES.COM/iFolderServer/Admin**.
2. Log in as **admin**.
3. Select **Client Policies**.
4. For Encryption, select the following:
 - On**
 - Enforced**
 - Hidden**
5. For Save Password, select the following:
 - On**
 - Enforced**
6. For Save Pass Phrase, select the following:
 - On**
 - Enforced**
7. Select **Update Policy**; then select **Refresh**.

8. Log out of iFolder.
9. Log in to iFolder with a user from FLIGHTOPS in xxx.Digitalair (see Table 7-4 for your username).

Table 7-4

Server	Username
DA4	RBHAT
DA5	KFULLMER
DA6	ASANDERS
DA7	DCROCKETT
DA8	IBLUNCK
DA9	HNAGAI



If you experience difficulty logging in, make sure you have configured subcontainer search properly. See the steps to [Part VIII: Enable Subcontainer Search](#).

10. As you log in, note of the following:
 - The Enable Automatic Login at Startup option (Save Password) is shown but it is selected and grayed so you cannot change it.
This option is On and Enforced, but not Hidden.
 - The Encryption option is not shown.
This option is On, Enforced, and Hidden.
 - You are asked to provide a Pass Phrase, and the Remember Pass Phrase option (Save Pass Phrase) is selected and grayed so you cannot change it.
This option is On and Enforced, but not Hidden.

11. While still logged in as your user from FLIGHTOPS, check the current client quota:
 - a. Right-click the **iFolder icon** in the System tray.
 - b. Select **Account Information**.
 - c. Select the **Account Information** tab.
 - d. In the Server Information box, note that this user has 200 MB total space on the server.
 - e. Close the Account Information window and log out of iFolder.
12. Return to the iFolder server management console; then select **Server Policies**.
13. Change the Initial Client Quota to **300 MB**.
14. Select **Update**; then select **Refresh**.
15. Log in to iFolder as a user from **MARKETING** in `xxx.Digitalair` (see Table 7-5).

Table 7-5

Server	Username
DA4	MJAIN
DA5	DDECKER
DA6	RHAYMOND
DA7	JASTIN
DA8	KBAHR
DA9	MYAMADA

16. Check the new client quota:
 - a. Right-click the **iFolder icon** in the System tray.
 - b. Select **Account Information**.
 - c. Select the **Account Information** tab.

- d. In the Server Information box, note that this user has 300 MB total space on the server.
 - e. Close the Account Information window; then log out of iFolder.
- 17.** Check the client quota for a user that already has an account:
- a. Log in again as the user from FLIGHTOPS and check that user's client quota.
Note that it still says 200 MB. This has not changed because the setting you changed only affects the initial client quota.
 - b. Close the Account Information window and log out of iFolder.
- 18.** Change an individual user's client quota:
- a. Return to the **iFolder server management console**.
 - b. Select **User Accounts**.
 - c. Select the user from FLIGHTOPS.
 - d. Change the Disk Quota to **300 MB**.
 - e. Select **Change**; then select **Refresh**.
 - f. Log in again as the user from FLIGHTOPS and check the user's client quota.
Note that it now says 300 MB.
 - g. Close the Account Information window and log out of iFolder.

(End of Exercise)

Objective 3 Maintain and Troubleshoot the iFolder Client

As a network administrator you maintain and troubleshoot the iFolder client. To do so it is helpful to understand the following:

- [The Benefit of the iFolder Client](#)
- [How the iFolder Client Works](#)
- [Common Issues Involving the Client](#)

The Benefit of the iFolder Client

The main benefit of using the iFolder client is automatic backup of data.

When the user is logged in, the iFolder client watches for new files and for changes made to existing files in the iFolder home directory and transmits those changes to the server.

This is done in the background without user intervention. With the web browser, the user must remember to manually upload new files and modified files to the server.

If the original file is lost from the user's computer, there is always an up-to-date backup of the file on the iFolder server.

How the iFolder Client Works

The iFolder client goes through the following steps as it keeps data synchronized with the iFolder server:

1. When a user enters a username and password on a workstation, the iFolder client sends them to the iFolder server encrypted with RSA Encryption.

2. The iFolder server takes the user name and password and performs an LDAP bind to the LDAP servers.



LDAP must be running in the iFolder environment for iFolder users to authenticate. On a NetWare server, NLDAP.NLM must running.

After an LDAP bind is successful, LDAP verifies that the user is connected to the correct iFolder server. If the user is connected to the wrong server, it redirects the client to the correct server.

3. When the connection is in place, the iFolder client reconciles the files in the local iFolder with the iFolder server and determines if there have been changes since the last login.

The client compares the sync index, file maps, and dirmaps on the client with the master copies of those same files that are held on the server.

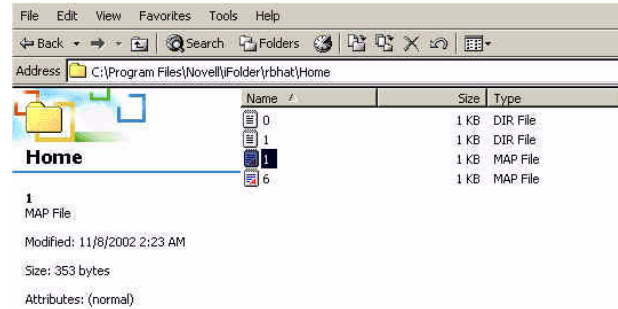
- **The Sync index** is a 4 kilobit file that increments each time a synchronization takes place.
- **File maps** contain metadata that describes the files in your iFolder home directory.
- **Dirmaps** contain metadata that describes the folders you create in your iFolder home directory.

You can view file maps and dirmaps on client computers in

**C:\PROGRAM FILES\NOVELL\IFOLDER\USERNAME
\HOME**

An example of these files is shown in the following:

Figure 7-7



4. If the client discovers through the comparison that there are new files or changes to files on the server or client, the new data on the server is downloaded, and any new data on the client is uploaded.



The maximum size for a file placed in an iFolder home directory is 4 GB. Files larger than 4 GB will not synchronize.

The uploads and downloads consist of only the changed portion of an existing file on a 4 kilobit block level. This minimizes the bandwidth that iFolder synchronization uses.

When uploading or downloading data, iFolder always uses port 80. iFolder's IP packets are never encrypted.

However, if the user has chosen to encrypt data, the client encrypts the data (using the pass phrase as the key) before it is placed in the packet.

When the server receives the packet, the data is stored on the server in its encrypted form.

If the user does not choose to encrypt data, the client transmits the data in plain text and the server saves it in plain text.

5. When the iFolder server receives the new data it increments the sync index.

The sync index indicates the current state of file system. As long as the sync index on the client matches the sync index on the server, there has been no change. The sync index is only 4 bytes and is a very small operation on the server.

6. When the same user connects to the iFolder server using another computer (such as at home or on a laptop), iFolder compares indexes. If it sees that the sync index on the server is different from the sync index on the client, it downloads the changes.

Now the files in both computer's iFolder home directories are the same.

7. While logged in to iFolder, the user creates a document or modifies an existing file. The iFolder client watches for file system changes and is notified of the new or modified file.

The new file or changes to the existing file are then uploaded to the iFolder server.

8. When the iFolder server receives the data it increments the sync index.

The server is ready for the next client login.

Common Issues Involving the Client

You might encounter the following client-related issues:

- [New Files Are Not Synchronizing from the Client to the Server](#)
- [Files Are Not Synchronizing from the Server to the Client](#)
- [The Conflict Bin Does Not Behave as Expected](#)

New Files Are Not Synchronizing from the Client to the Server

Users might discover that new files they thought were being synchronized from the client to the server are actually not on the server, while at the same time changes to existing files are being synchronized.

Users discover this condition when they try to access a file while using a different computer than the one they created the file on. The cause of this problem is corrupted file maps and dirmaps.

The solution is to do the following:

1. Delete all the file maps and dirmaps from the home directory on the computer where the nonsynchronizing files are originating.
2. Have the user log in to iFolder from the originating computer.
The iFolder client recognizes that the file maps and dirmaps are not present and will download them from the master copies on the server.
3. Open the iFolder client Account Information window and view the client's synchronizing activity. You should see that the missing files are now being synchronized.

Files Are Not Synchronizing from the Server to the Client

Users might discover that files that have been synchronized to the server from one computer are not being synchronized when they log in from another computer.

The most likely cause is that the second computer does not have sufficient disk space for the synchronization process to complete.

As files are downloaded from the server they are placed in the user's working home directory before being copied to their iFolder home directory.



The user's working home directory is found in C:\PROGRAM FILES\NOVELL\IFOLDER\USERNAME\HOME.

The user's iFolder home directory is found in C:\DOCUMENTS AND SETTINGS\ADMINISTRATOR\MY DOCUMENTS\IFOLDER\USERNAME\HOME.

The iFolder synchronization process requires that a certain amount of disk space always be available. If the disk is getting too full additional space must be made available.

The Conflict Bin Does Not Behave as Expected

The purpose of the conflict bin is to save files that have been deleted or overwritten by the synchronization process. If a user expects to find a file in the conflict bin but it is not there, you should look for the following possible causes:

- **The conflict bin is too small.** The default setting for the conflict bin is 25 MB. If the conflict bin is full or the deleted files are more than 25 MB, they will not be placed in the conflict bin.

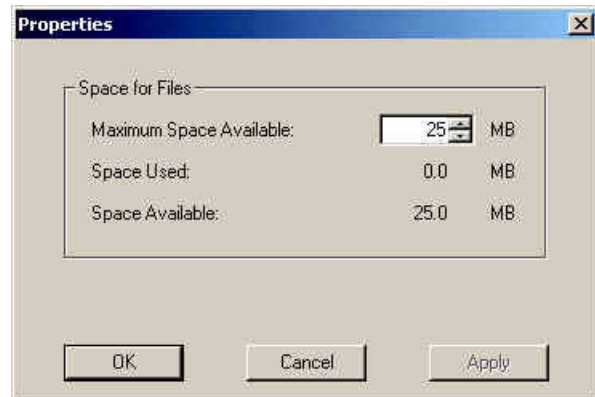
There is no way to solve this problem after it happens, but it can be prevented by allocating more space to the conflict bin. This is a client policy setting that you can change at any time.

Users can also increase the size of their conflict bin by doing the following:

1. Right-click the **iFolder icon**.
2. Select **View Conflict Bin**.
3. Select **File > Properties**.

The following appears:

Figure 7-8



4. Enter the *desired value* in the Maximum Space Available field; then select **Apply**.
- **The file has been deleted on the originating computer.** When a file is deleted or overwritten, iFolder only deems it a conflict when it takes place on a computer other than the one the file originated on.

For example, if a file originated on computer A and then was later deleted from computer B, it is placed in the conflict bin. If it is deleted on computer A it is not placed in the conflict bin.

Objective 4 Maintain and Troubleshoot the iFolder Server

Following are some common maintenance and troubleshooting issues related to administering the iFolder server:

- [Adjust the Number of Threads per Child](#)
- [The Admin Cannot Access the Server Management Console](#)

- [Port Conflict with iPrint Secure Port](#)
- [Restoring User Accounts When Pass Phrases Are Forgotten](#)
- [LDAP Incorrectly Configured for Non-secure Port](#)

Adjust the Number of Threads per Child

When many users are accessing iFolder through the browser client, you might need to increase the number of threads per child.

This setting is configured in the HTTPD.CONF file and looks like this:

```
ThreadsPerChild 150
```

The default setting for threads per child is 150. The maximum setting is 2048. Threads are used to maintain user connections to the iFolder server. The iFolder client does not use persistent connections.

For client connections a thread is used for each 32 kilobits of data transmitted; then the connection is released. However, the iFolder browser requires persistent connections.

If most of your users are using the client, you do not need to have 1 thread per user. 50 threads can service 1,200 concurrent iFolder client connections.

However, if most of your users are connecting through the web browser and you have many concurrent connections, you need to adjust the ThreadsPerChild setting in HTTPD.CONF.

The maximum number of threads correlates with the bandwidth of your server's network board. With a 100 Mb board, you can set your threads to 312. With a 1 Gb network board, you can set them to 2048.

The Admin Cannot Access the Server Management Console

If the server certificates that provide security for the admin to log in become corrupted, the server management console is not available.

Users can still log in using the client because the client uses port 80. But browser-based access is not available because the secure port is used.

You can give yourself admin access by making a temporary modification to the HTTPD_ADDITIONS_NW.CONF file. This modification creates only one virtual host, the nonsecure host, for the Apache server.

To make this modification, do the following:

1. Make a backup copy of the current HTTPD_ADDITIONS_NW.CONF file by saving it under a different name.
2. Open HTTPD_ADDITIONS_NW.CONF and remove the entire secure virtual host portion of the file.
3. From the nonsecure virtual host portion of the file, change the ServerSecurePort line to
ServerSecurePort 80
4. Save the file; then stop and restart iFolder.
You cannot perform the stop synchronization step in this case.
5. Access the server management console as you would normally.
6. After the problems with the certificates are resolved, rename the original HTTPD_ADDITIONS_NW.CONF file to its correct name.

Port Conflict with iPrint Secure Port

A port conflict can arise because iFolder and iPrint both use port 443 by default as their secure port. The port resolver should prevent this problem.

If the port resolver detects the conflict when iFolder is installed, it will force iFolder to use port 52443 instead. But if the port resolver does not detect the conflict during installation you have 2 choices:

- **Give iFolder a new IP address.** This process is outlined in an earlier objective (see [“Perform iFolder Management Tasks” on 7-14](#)).
- **Move iPrint to a new IP address.** iPrint uses the HTTP stack that is bound to the primary IP address on the iPrint server. The primary IP address is the first one bound in AUTOEXEC.NCF or INITSYS.NCF.

This is determined during the installation of iPrint so the only way to change it is to issue a command that moves the HTTP stack to a new IP address. Do the following:

1. Find out which SSL certificate the HTTP stack is using.
2. Determine the new IP address that you will use.

You have the following options:

- Use an IP address that is bound to another network board on the server.
 - Use a secondary IP address that is bound to the same board.
3. Enter the following command at the console prompt; then add it to AUTOEXEC.NCF:

```
HTTPBIND new_IP_address /keyfile:"SSL Certificate"
```

4. Unload and reload iPrint by entering the following at the server console:

```
UNLOAD NDPSM Manager_Name  
UNLOAD BROKER Broker_Name (if applicable)  
BROKER Broker_Name (if applicable)  
NDPSM Manager_Name
```

(Unloading and loading a print broker is only applicable if a broker is running on the server.)

Restoring User Accounts When Pass Phrases Are Forgotten

When users forget their pass phrase they can no longer get into their iFolder account. Access to their accounts can be restored by deleting all user data from the server and then allowing the user to log in again and resynchronize with the server.

Do the following:

1. Make sure the user has a complete local copy of their data.
2. Make sure the user also has a connection to the server with adequate bandwidth for completing the resynchronization.
3. Find out which folder the user's data is saved in on the iFolder server:
 - a. Find the user's account in iFolder server management console on the User Account page.
 - b. Place your mouse on the user's name.
 - c. Note that in the status line at the bottom of the browser window there is a long hexadecimal number.

Example:

```
User=9E226380764BAA07696D656DF6F7B191.htm
```


This is the name of the directory where all user data is stored. An example of the status line is this:

Figure 7-9



4. Determine where the user’s data is stored on the iFolder server:

- a. Browse to the location of the iFolder data.

For example, you have configured your iFolder data to be saved on

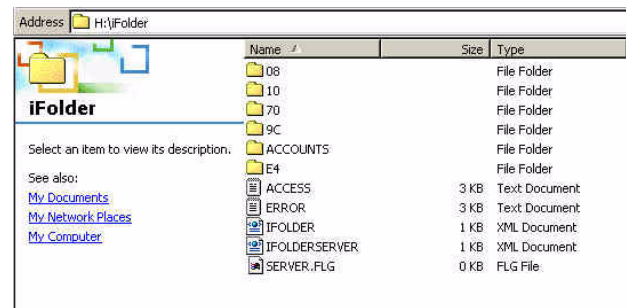
\\DAx\DATA\iFolder

- b. Open the iFolder directory.

The iFolder directory contains a series of subfolders with 2-digit hexadecimal names. User directories are categorized based on the first 2 digits of their directory’s name.

The first 2 digits are the directory identifiers. An example of subfolders is shown in the following:

Figure 7-10



- c. Find the subfolder that contains the user’s directory.


To find the correct subfolder you need to know that each of the subfolders contains a range of 4 sequential directory identifiers. The range begins the name of the sub-folder.

For example, the 00 subfolder contains the 00, 01,02, and 03 directory identifiers. All user directories that begin with any of these identifiers are in the 00 subfolder.

For another example, the 9C subfolder contains the 9C, 9D, 9E, and 9F directory identifiers. All user directories that begin with any of these identifiers are in the 9C subfolder.

The 9E226380764BAA07696D656DF6F7B191 user directory is in the 9C subfolder because it begins with 9E, as shown in the following:

Figure 7-11



- d. Open the user's directory and verify that this is the correct directory.

Within each user's directory is a CONTROL.DAT file.

Open the CONTROL.DAT file with Notepad and you will see the user's name in the last characters of the first line.

This is the name of the user that this directory belongs to.

5. Delete the user's directory.

Delete the entire directory that belongs to the user that forgot the pass phrase. For example, for the user associated with 9E226380764BAA07696D656DF6F7B191, you would delete the 9E226380764BAA07696D656DF6F7B191 folder.

6. Have the user login to iFolder.

iFolder will recognize that this user doesn't have a directory on the server and it will create one for the user.

The user enters a new pass phrase and the iFolder client initiates the synchronization process that places all the user's data back on the server. This time the data is encrypted with the new pass phrase as the key.

(iFolder 2.0 has an added feature that allows the administrator to restore a user's forgotten pass phrase.)

LDAP Incorrectly Configured for Non-secure Port

If iFolder is configured at installation to use port 389 (rather than port 636) LDAP must be configured to allow clear text passwords.


Using port 389 is a legitimate option when the LDAP server and the iFolder server are running on the same physical computer. In that scenario no communications take place over the wire so no encryption is required.

If you do not configure LDAP during installation to allow clear text passwords you must configure it from ConsoleOne by doing the following:

1. In ConsoleOne, navigate to the container that holds the server running LDAP.
2. Open the **LDAP Group** object.
3. Deselect **Require TLS for simple binds with password**.

This option is shown in the following:

Figure 7-12



4. Select **OK**.

If you do not have LDAP configured to allow clear text passwords you will receive an error message that confidentiality is required.

Summary

The following is a summary of the objectives in this section:

Table 7-6

Objective	What You Learned
1. Describe iFolder Configuration Files	<p>iFolder reads 2 configuration files to find its configuration information. These files can be used to help troubleshoot certain iFolder problems.</p> <p>The names of these files are</p> <ul style="list-style-type: none"> ■ HTTPD.CONF This is the default Apache configuration file. ■ HTTPD_ADDITIONS_NW.CONF The HTTPD.CONF file creates an Apache web server exclusively for iFolder to use.
2. Perform iFolder Management Tasks	<p>To manage iFolder you must be able to perform the following tasks:</p> <ul style="list-style-type: none"> ■ Stop and Start the iFolder Server ■ Set iFolder Client and Server Policies ■ Change the Location of iFolder User Data ■ Add Contexts ■ Add Additional Administrators ■ Change the iFolder Server IP Address

Table 7-6 (continued)

Objective	What You Learned
3. Maintain and Troubleshoot the iFolder Client	<p>As network administrator you need to know the following about the iFolder client:</p> <ul style="list-style-type: none"> ■ The Benefit of the iFolder Client ■ How the iFolder Client Works ■ Common Issues Involving the Client <p>You might encounter any or all of the following client-related issues:</p> <ul style="list-style-type: none"> ■ New Files Are Not Synchronizing from the Client to the Server ■ Files Are Not Synchronizing from the Server to the Client ■ The Conflict Bin Does Not Behave as Expected
4. Maintain and Troubleshoot the iFolder Server	<p>Following are common maintenance and troubleshooting issues related to administering the iFolder server:</p> <ul style="list-style-type: none"> ■ Adjust the Number of Threads per Child ■ The Admin Cannot Access the Server Management Console ■ Port Conflict with iPrint Secure Port ■ Restoring User Accounts When Pass Phrases Are Forgotten ■ LDAP Incorrectly Configured for Non-secure Port

MODULE 4

Deliver High Availability Services with Novell Cluster Services

Section 8 Design and Set Up an NCS Cluster Configuration

Section 9 Install and Test NCS on a 2-Node Cluster

Section 10 Configure and Test High Availability File Access

Section 11 Configure and Test High Availability Services

SECTION 8 Design and Set Up an NCS Cluster Configuration

Duration: 2 hours

In this section, you learn the purpose and advantages of clustering data and applications and how to design and set up a 2-node NCS cluster.

Objectives

1. [Identify the Purpose and Advantages of Implementing an NCS Solution](#)
2. [Design and Set Up an NCS Cluster Configuration](#)

Introduction

Consider introducing this section by demonstrating the migration of a cluster-enabled volume while playing the video used in Exercise 10-2.

As part of the demo, have students run the video (located on the DA1/DA2 cluster-enabled SCSI hard drive) from their workstations while migrating the cluster-enabled volume

Your employees, customers, and partners need access to data, applications, web sites, and other services 24 hours a day, 7 days a week, 365 days a year.

Keeping that data online, along with critical applications that depend on that data, requires an intelligent approach to system design that includes clustering services.

Novell Cluster Services (NCS) 1.6 is a multinode clustering system for NetWare 6, and is enabled for eDirectory. NCS ensures availability and manageability of volumes, applications, server licenses, and services.

Enlist the assistance of students who have experience with NCS 1.5 or 1.6 clustering installations to share their experience.

A license for a 2-node NCS 1.6 cluster is included with NetWare 6. Mixing prior NCS cluster license versions with NetWare 6 clusters is not supported.

Scenario

Because of the increase in company data storage and service availability requirements, the company CIO has given you the task of researching the use of clustering as a possible solution.

You know that NetWare 6 provides a 2-node NCS clustering license, but you're not sure what advantages NCS provides.

You want to set up a 2-node cluster in your lab for testing and demonstration purposes but funds are not available to purchase a commercial clustering hardware solution.

Objective 1 Identify the Purpose and Advantages of Implementing an NCS Solution

In this objective you identify the purpose and advantages of implementing an NCS solution to maintain high availability of data and services.

To do this, you need to understanding the following:

- [High Availability Terms](#)
- [High Availability Definition](#)
- [Computer System Outage Factors](#)
- [Benefits and Features of an NCS High Availability Solution](#)

High Availability Terms

These terms are used later in this objective when introducing the purpose and benefits of an NCS high availability solution.

They are introduced here to provide students with a basic vocabulary to help you determine what students know about high availability solutions, and to prevent student interruptions in this objective.

These terms give students the ability to explain the benefits of high availability to management, and the ability to discuss high availability with clustering consultants.

You should understand the following about high availability:

- **Resource.** Any service or data that can be migrated from one server to another in a cluster

For example, you cannot migrate a physical printer from one server to another, but you can migrate a service such as iPrint that provides access to the printer.
- **Service.** A resource that is available to a customer or employee from a server

A server is a host for services. Network administrators care about servers. Users care about services; they do not care about servers until a service goes down.

Examples of services include printing, file access, web services, and email.
- **Availability.** The percentage of total system time that a service is accessible for normal use
- **Uptime.** The duration of time a service is functioning
- **Outage.** The loss of a computer service
- **Downtime.** The duration of an outage (planned or unplanned)

When determining high availability, downtime duration can be a critical factor. For example, a system that sustains 10 outages of 10 seconds duration each has a higher availability than a system that has one 10 minute outage.
- **Reliability.** The amount of time before a system is expected to fail
- **Mean time between failures (MTBF).** The average time (usually listed in hours) that a device or system works without failure

You can calculate the MTBF by dividing the total number of operating hours by the total number of failures.

A true measure of high availability is not the number of times a system fails, but the MTBF time between failure and recovery.

- **Mean time to recovery (MTTR).** The average time that a device takes to recover from a nonterminal failure

MTTR is often part of a maintenance contract, where you would pay more for a system whose MTTR is 24 hours than for a system with an MTTR of 7 days.

This means that the supplier guarantees to have the system running again within 24 hours (or 7 days) of being notified of the failure.

Some devices have an MTTR of zero, which means that they have components that can take over the instant the primary components fail.

Availability is often expressed as a percentage of total uptime. You can calculate availability using the following formula:

$$\% \text{ Availability} = (\text{MTBF}) / (\text{MTBF} + \text{MTTR}).$$

Notice that if MTTR is 0 then % Availability is 100%. However, because there is always time associated with recovery (even if it is a few seconds), you can never achieve 100% uptime.

High Availability Definition

High availability generally means one or more of the following to your company management, employees, partners, and customers:

- **24x7x365.** This represents the ability to access resources 24 hours a day, 7 days a week, 365 days a year.

The term *24x7* is often used to represent 24x7x365, with 365 days a year assumed.

- **24x7x365 at 100%.** This represents 100% availability of resources all the time, and is what most organizations and management personnel define as high availability.

Emphasize that high availability should always be viewed from the standpoint of the customer or employee.

Although you can work towards 24x7x365 at 100%, this level of high availability is impossible to achieve. Even with a minimal MTTR, unforeseen natural and manmade disasters can impact 100% availability.

- **6-6 or 6-11 (A.M. to P.M.) at 100%.** This represents 100% availability of resources during certain hours of the day.
- **Five 9s guaranteed.** When discussing high availability, clustering consultants often talk about the five 9s of high availability.

The five 9s represent access to resources 99.999% of the time 24x7x365. This means only 5.2 minutes of downtime during the year.

The definition for high availability can differ from organization to organization. When you talk with managers and customers about high availability, they might want 24x7x365 at 100%. However, that level of availability is impractical in terms of equipment and maintenance.

For example, a call center organization might only care about availability during business hours and have no requirements outside that time. On the other hand, a commercial web site really requires 24x7x365 availability.

You need to determine the level of availability required for each service on your cluster, and then work with others (such as your power vendor, application retailer, and a clustering consultant) to make sure you have the service contracts and equipment to support that level.

The following gives you 24x7x365 downtime figures for high availability services percentages:

Table 8-1

High Availability Access	Percentage	Yearly Downtime
Five 9s	99.999%	5.2 minutes

Table 8-1 *(continued)*

High Availability Access	Percentage	Yearly Downtime
Four 9s	99.99%	52.5 minutes
Three 9s	99.9%	8.7 hours
Two 9s	99.0%	87.6 hours
	98.0%	175.2 hours
	96.0%	350.4 hours

When computing this in terms of hourly outage costs and the yearly loss at five 9s availability, see the following (as compiled by Stratus Technologies):

Table 8-2

Service	Hourly Outage Cost	99.999% Yearly Loss
Brokerage	\$5.6 – \$7.3 million	\$485 – \$633 thousand
Credit card	\$2.2 – \$3.1 million	\$191 – \$269 thousand
Pay-per-view	\$67 – \$233 thousand	\$6 – \$20 thousand
Home TV shopping	\$87 – \$140 thousand	\$8 – \$12 thousand
Catalog sales	\$60 – \$120 thousand	\$5 – \$10 thousand
Airline reservations	\$67 – \$112 thousand	\$6 – \$10 thousand
Teleticket sales	\$56 – \$82 thousand	\$5 – \$7 thousand
ATM fees	\$12 – \$17 thousand	\$1 – \$2 thousand

When defining high availability, remember that customer and employee needs and perception should be the focus.

Even if a system or application is technically available, slow response times could mean that it is not usable.

Factors that can contribute to slow response times include hardware (CPU, memory, application overload, and network failures), misconfiguration, and lack of recovery processes.

After defining high availability, you can design the clustering solution necessary to deliver high availability services.

Computer System Outage Factors

Ask students how often their servers fail or how often they bring down their servers and for what reasons (such as maintenance or software upgrades).

Make sure students understand that NCS and NetWare are normally the most reliable components in a high availability clustering solution.

This is especially true when configuring a 2-node SCSI cluster with a SAN.

The SCSI hard drive and adaptor cards are not designed specifically for clustering, and must be carefully configured before clustering software (such as NCS) can properly use the SAN.

Although NCS provides all software features necessary to manage and configure a high availability clustering solution, other factors contribute significantly to making services highly available.

These factors often cause computer system outages and include the following:

- **Physical.** Physical faults or hardware failures
- **Design.** Design errors in both the hardware and software you want to cluster-enable
- **Operations.** Errors caused by operations personnel or users
- **Environmental.** Power or cooling system failures, failures of external network connections, natural disasters, and so on
- **Reconfiguration.** Scheduled maintenance, upgrades, or configuration changes

In addition, there are single points of failure, such as one power source or one hub. The more single points of failure, the greater the risk for maintaining high availability.

Although many factors can cause an interruption, NCS solves the problem of unavailable services due to an abended server.

As soon as a server in a cluster abends, resources are migrated to other servers in the cluster, with little or no interruption in service.

After reviewing the features of NCS, you might want to discuss with students the benefits of implementing NCS clustering in their own network environment.

Benefits and Features of an NCS High Availability Solution

NCS 1.6 includes the following to help you ensure high availability:

- **Multinode all-active cluster (up to 32 nodes).** NCS lets you configure up to 32 NetWare servers (nodes) into a high-availability cluster, where resources can be dynamically switched or moved to any server in the cluster.

Services can be assigned across the cluster to different servers. Any NetWare server in the cluster can restart resources from a failed server in the cluster.

- **Multiprocessor and multithreading enabled.** NCS 1.6 is more efficient than ever because NetWare 6 is not just multiprocessor enabled, it's also multithreaded.

Each processor can be maximized to execute commands faster and more efficiently, providing faster network throughput that delivers 24 hours a day, every day of the year.

- **Consolidation of applications and operations.** NCS lets you tailor a cluster to the specific applications and hardware infrastructure that fit your organization.

You can also reduce unplanned and planned outages for software and hardware maintenance and upgrades.

In addition, you can lower costs by consolidating applications and operations onto a cluster. Customers find they can reduce the number of servers used to provide services by 50% or more.

- **Flexible resource management.** You can configure resources to automatically switch or be moved when a server fails, or you can move them manually to troubleshoot hardware or balance the workload.
- **Shared storage support.** NCS provides support for shared SCSI devices or Fibre Channel SANs.

SAN technologies deliver a lower total cost of operation (TCO) by providing higher levels of availability for centralized storage and server resources.

And with NCS you can dynamically assign and reassign resources on the shared storage as needed.

Shared disk fault tolerance can be obtained by implementing RAID Level 5 on the shared disk system.



According to the Gartner Group, it costs at least 40% more to manage storage as it does to purchase it. In many cases, this discrepancy can be much higher — 300% to 400%.

After being implemented, SANs and NASs require 75% less personnel to maintain than traditional direct attached storage.

- **Single point of control.** You can manage a cluster from a single point of control and adjust resources to meet changing workload requirements (thus, manually load-balance the cluster).

You can manage and configure NCS through the browser-based NetWare Remote Manager or through ConsoleOne cluster configuration and monitoring.

The browser-based NetWare Remote Manager lets you remotely manage your cluster.

- **Fan-out failover.** You can configure migration and load balancing of resources to other nodes during a failover based on factors such as node traffic and availability of installed applications

In NCS clustering, the clustering consultant or network administrator configures fan-out failure for the cluster.

- **Cluster event and state notification.** You can configure NCS to notify administrators through email of cluster events and cluster state changes.

Objective 2 Design and Set Up an NCS Cluster Configuration

The clustering terms provide students with a basic clustering vocabulary and help you determine what students know about clustering.

Students do not need to thoroughly understand each term at this point. You discuss the terms again during the review of NCS configuration and when summarizing exercises.

While discussing NCS clustering, students might become confused as you begin to use the words *node* and *server* interchangeably.

A node in a cluster is a physical server that is cluster-enabled.

However, clustering consultants and administrators often use both words to mean the same thing when referring to a cluster-enabled server.

Now that you understand the purpose and benefits of an NCS high availability solution, you can design and set up a simple 2-node NCS cluster configuration.

The following help you prepare for this task:

- [Basic Clustering System Terms](#)
- [NCS Cluster Components](#)
- [Typical NCS Shared Disk System Cluster Configurations](#)
- [NCS System Terms](#)
- [Rules for Managing an NCS SCSI SAN](#)
- [Troubleshooting a 2-Node NCS SCSI SAN](#)

Basic Clustering System Terms

You should understand the following terms when discussing a clustering solution:

- **Cluster.** A group of servers linked together in a dedicated network to minimize the loss of service by reducing or managing failures and minimizing downtime
- **Node.** A server in a cluster
- **Cluster resource.** A server resource, application, or network service with a dynamic location managed by clustering software

In NCS, a cluster resource can only be assigned to one node at a time.

- **Shared storage device.** A device (such as external hard drives, disk arrays, and Fibre Channel disks) in a cluster used to store shared cluster resources

Emphasize that using the word *failover* to describe migrating a resource can communicate the idea that NCS and NetWare 6 have failed when both are still running.

- **Storage area network (SAN).** A dedicated network (such as a cluster) connecting servers and shared storage devices
- **Migration.** The process of moving resources to other nodes in your cluster without waiting for a server to fail
Migration is useful for tasks such as load balancing of resources and upgrading servers in the cluster.
- **Failover.** The process of restarting a failed node's resources on one or more of the surviving nodes
Failover normally results from a server hardware or power source problem.
Network administrators often refer to migrating resources as "failing over" resources. This can cause confusion and misconceptions.
Failover happens when a server (node) in a cluster fails (usually due to a power or server hardware failure) and is unplanned. Anything else is a migration.
For example, a migration can be a conscious choice to move resources for maintenance and load balancing. NCS also migrates resources when a cluster node fails.
- **Failback.** The process of returning the failed node's resources back to the way they were before the failover
- **Fibre Channel.** The Fibre Channel Standard (FCS) defines a high-speed data transfer interface to connect workstations, mainframes, supercomputers, storage devices, and displays
Optical and electrical media are supported, transferring data from 260 megabits/second (copper wire) up to 4 GB/second (fiber optics), up to a distance of 10 km (with fiber optics).

NCS Cluster Components

The following components make up an NCS 1.6 cluster:

- From 2 to 32 NetWare 6 servers configured to use IP, each with at least one local disk device (used for a local volume SYS)
- NCS 1.6 and NetWare 6 installed on each server in the cluster
- A shared disk system connected to all servers in the cluster (recommended for most configurations)

These versions of NCS and NetWare must be loaded on each server for the cluster nodes to access NSS pools and volumes.

- NetWare services that do not require a shared disk system include licensing, LDAP server, and DHCP.
- **One** of the following:
 - High-speed Fibre Channel cards and cables to connect the servers to the shared disk system
 - SCSI cards and cables used to connect the servers to a shared disk system
 - A dedicated SCSI hardware system (cluster in a box) that includes the servers and shared disk system



If you use SCSI cards and a SCSI hard drive to configure a 2-node cluster, make sure each card and the hard drive are assigned a different SCSI ID number.

Also, make sure you purchase enclosures and cables with the correct pin-outs.

Typical NCS Shared Disk System Cluster Configurations

Some students might be confused about the spelling and use of the term *fib*re.

Fibre Channel refers to a protocol that uses fiber cable or copper wire to transmit data from the nodes to the shared storage device.

Typical cluster configurations include a shared disk system connected to all servers in the cluster. If a server fails, another server is assigned the resources. This gives users continuous access to resources such as data, applications, and services.

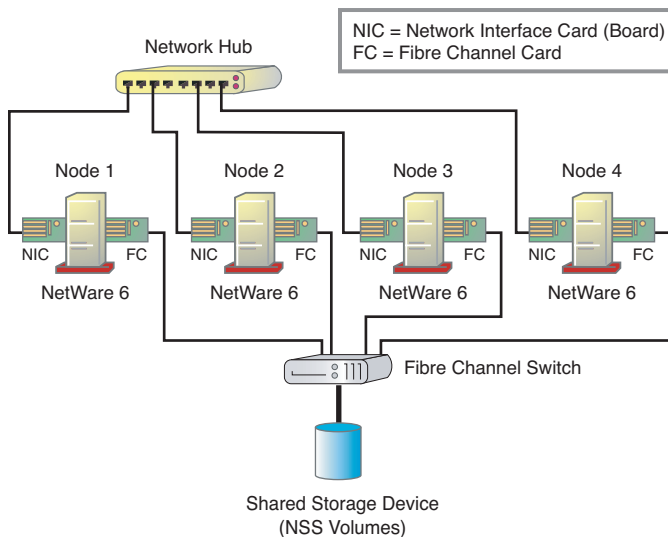
You can use 2 basic shared disk system configurations:

- [Fibre Channel Cluster Configuration](#)
- [SCSI Hard Drive Cluster Configuration](#)

Fibre Channel Cluster Configuration

The following shows a Fibre Channel cluster configuration. (Fibre Channel cards are also called Host Bus Adaptors (HBAs).)

Figure 8-1 (slide)



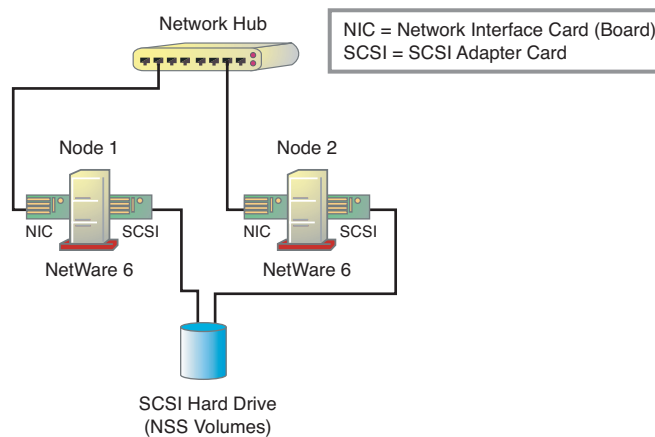
In this figure, the nodes represent cluster-enabled servers on the network; the SAN can consist of a single hard drive or an array of storage devices.

SCSI Hard Drive Cluster Configuration

Although Fibre Channel is recommended, you can use a dedicated SCSI hardware solution. For testing and demonstrations, you can configure your cluster to use an external shared SCSI hard drive.

The following shows a SCSI hard drive cluster configuration:

Figure 8-2 (slide)



The purpose of building this type of SCSI configuration (using SCSI adaptor cards in NetWare 6 servers and a shared SCSI hard drive) is for testing and demonstration.

Setting up a SCSI configuration in a lab at home or work helps you take advantage of the 2-node license available with NetWare 6, and can help you demonstrate the benefits and viability of cluster-enabling a company network.

However, external SCSI hard drives and SCSI adaptor cards are not designed specifically for clustering.

If you want to use SCSI as a less expensive alternative to Fibre Channel, you can purchase a dedicated SCSI hardware solution (such as one provided by Compaq or Dell) created specifically for clustering.



If the SCSI adaptor cards and drivers are multi-initiator enabled, they can share a SCSI hard drive, but be careful about configuration issues (see [Troubleshooting a 2-Node NCS SCSI SAN](#) in this section for details).

You might want to use the Cluster State view (ConsoleOne) or the Cluster Status view (NetWare Remote Manager) to identify and reinforce some of these terms.

NCS System Terms

You need to know the following NCS system terms:

- [Master Node](#)
- [Cluster-Enabled Volumes and Pools](#)
- [Shared Storage Device](#)
- [Cluster Resource](#)
- [Heartbeats, Tics, Poison Pills, and the Split Brain Detector \(SBD\)](#)
- [Fan-Out Failover](#)

Master Node

The first server that comes up in an NCS cluster is assigned the cluster IP address and becomes the master node. (All other nodes in the cluster are often referred to as slave nodes.)

The master node updates information transmitted between the cluster and eDirectory, and monitors the health of the cluster nodes.

If the master node fails, NCS migrates the cluster IP address to another server in the cluster for you, and that server becomes the master node.

Cluster-Enabled Volumes and Pools

A cluster-enabled volume in NetWare 6 is an NSS volume that gives users continuous read/write file access on the shared storage device.

Volumes are associated with NSS pools that provide a unique secondary IP address (through a virtual server object) for locating the volumes on the cluster's shared storage device.

In other words, you migrate or failover pools instead of volumes. This means that you can migrate or failover more than one volume at a time by assigning the volumes to a pool.

Shared Storage Device

The shared storage device in an NCS cluster is where customers and employees access files on the SAN.

For example, if employees need access to an EMAIL volume, you can mount EMAIL as a volume on the shared storage device and instruct employees to map a drive to that copy of EMAIL.

By placing EMAIL on the shared storage device, you ensure that public services associated with the volume are always available, no matter which node fails or which node you take offline.

Other examples of placing files on the shared disk include web sites, print drivers, and iFolder user data files.

By moving all customer and employee file access to the shared storage device, you can reduce the number of servers needed in your network and reserve files on the local hard drive of a node (such as volume SYS) for network administrators.

Make sure students understand that if employees or customers need direct read/write access to data and files, you should store the files in a cluster-enabled volume on the shared storage device.

However, if the service or application maintains the employee or customer data (such as synchronized iFolder data), you should create a cluster resource that accesses the volume for the service or application.

Make sure students understand that resources in a cluster are no longer dedicated to a particular server. They become server independent in a cluster.

Although a resource is initially assigned to a server, that resource can be migrated at any time to another server.

Cluster Resource

A cluster resource is an object in eDirectory that represents an application or other type of service (such as DHCP or the master IP address) that you can migrate or failover from one node to another in an NCS cluster.

The cluster resource object includes scripts for unloading the service from one node and loading it on another node.

In most cases, make sure the service is installed on all nodes in the cluster that will host the service.

Heartbeats, Tics, Poison Pills, and the Split Brain Detector (SBD)

NCS uses heartbeats on the LAN, tics on the SAN, and a split brain detector (SBD) on the shared storage device, and poison pills to keep all services highly available on the cluster when a node fails:

- A *heartbeat* is a small IP packet sent periodically over the LAN (not the SAN) by the master node and the slave nodes in the NCS cluster.

The master node sends out a multicast heartbeat to all slave nodes. Each nonmaster node (*slave node*) sends out a unicast heartbeat to the master node.

The nodes monitor the heartbeat of other nodes in the cluster at a tolerance rate of 8 seconds (default setting).

The *tolerance rate* is the amount of time a node waits for a heartbeat from another node before taking action that results in casting off (abending) the failed node.

- A *tic* (Transport Independent Checking) is a type of heartbeat sent over the SAN by a node. The tic writes an epoch number to the node's sector in an SBD partition on the shared storage device.

The epoch number increases by 1 each time a node leaves or joins the cluster.

- *SBD* information is stored in an *SBD partition* on the shared storage device. Each node in the cluster stores and maintains its own *SBD* information, including an epoch number, in a separate sector.
- A *poison pill* is a voluntary abend by a node that has been cast off by the cluster by other nodes.

Whenever a node completely fails (such as a power outage), no heartbeat or tic is sent over the LAN or SAN.

If a heartbeat is not detected within 8 seconds (the default tolerance rate), the master node notes the failure, a new cluster view is created that does not include the failed node, and the failed node resources are migrated to other designated nodes in the cluster.

The failed node can only join the cluster again when you reboot the server and the node starts running the cluster protocol.

However, if a node is still active, but fails to send a heartbeat over the LAN, a condition called *split brain* occurs where all other nodes update their epoch numbers by 1, but the node without a heartbeat becomes isolated and cast off by the cluster.

The following helps you understand what happens to an isolated node that fails to send a heartbeat over the LAN:

- [How NCS Casts Off an Isolated Slave Node](#)
- [How NCS Casts Off an Isolated Master Node](#)
- [An Isolated Slave Node Example](#)



For details on heartbeats, split brains, and poison pills, see TID 10053882.

How NCS Casts Off an Isolated Slave Node

NCS determines when a slave node fails and casts off the failed node through the following process:

1. The master node monitors the heartbeats of all other nodes in the cluster to determine if they are still “alive.” The master node also reads the epoch numbers for all nodes in the cluster.
2. If a heartbeat is not received from a slave node within 8 seconds (the default tolerance rate), the master node and remaining slave nodes create a new cluster membership view without the failed node.
3. Each node in the new membership (including the master node) updates its epoch number by 1 in the SBD partition.

This causes a *split brain*, because the epoch number for the isolated node is one less than for the rest of the nodes.

4. NCS uses this information in the SBD to vote between the 2 cluster memberships.

The cluster membership that has the most nodes wins. If there are equal nodes in both views, the side with the membership that contains the previous master node wins.

In the special case of a 2-node cluster, if one of the nodes can still communicate over the LAN and the other node can't, the node with the good connectivity wins.

5. The nodes in the surviving cluster membership write a special token to the sector in the SBD partition for the losing node.
In this case, the token is written to the sector for the failed slave node.
6. The losing node reads the special token, and then abends by taking a *poison pill*. The poison pill causes a self-inflicted abend that stops all processes on the node.

Abending ensures that nodes on the losing side cannot corrupt the new, healthy cluster.

7. The new cluster (minus the failed node) migrates the resources (volumes and services) assigned to the failed node to other nodes in the cluster and services continue uninterrupted for customers and employees.

How NCS Casts Off an Isolated Master Node

Each slave node in a cluster continuously monitors the heartbeat of the master node. If the master node fails to send a heartbeat over the LAN within 8 seconds (the default), the following occurs:

1. A new cluster membership view is created that includes only the slave nodes.
2. Each node in the new cluster membership view updates its epoch number by 1 in the SBD partition.
3. NCS uses the information in the SBD partition to vote between the 2 cluster memberships.
4. Because the master node is the only node with a different epoch number, the new cluster membership view with the slave nodes wins.
5. The nodes in the new cluster membership view write a special token to the sector in the SBD for the master node.
6. The slave nodes use an algorithm to vote on which node becomes the new master node.
7. The failed master node reads the special token, takes a poison pill, and abends.
8. Any resources assigned to the failed master node are migrated to other nodes in the cluster.

Working together, the heartbeats, tics, epoch numbers, and SBD partition allow NCS to monitor and quickly respond to a failed server.



There are 4 counters in network board drivers that NCS monitors to determine which nodes can send heartbeats. Some drivers do not support these counters. If you run a 2-node cluster and the master node fails to send a heartbeat over the LAN, NCS might cast off the slave node instead of the master node.

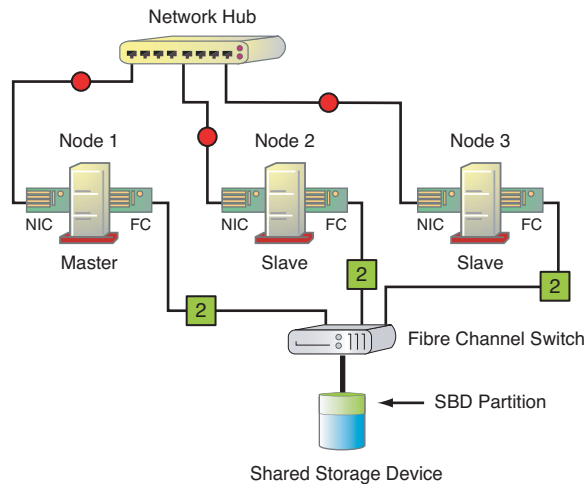
This only happens in a 2-node cluster where both master and slave have a vote and NCS can't determine which node is communicating. In a 3-node cluster, the slaves have 2 votes and will cast off the master node. For details on these counters, see TID 10057437.

An Isolated Slave Node Example

Although you work with a 2-node cluster in this section to demonstrate clustering, NCS clusters normally include 3 or more nodes to ensure high availability.

For example, the following NCS cluster is configured with 3 nodes and a shared storage device:

Figure 8-3 (slide)



Each node sends a heartbeat over the LAN and writes an epoch number in a tic to the SBD partition. Node 1 is the master node and the current epoch number for all nodes is 2.

Suppose the network board for node 3 fails. The following occurs to maintain the health of the cluster:

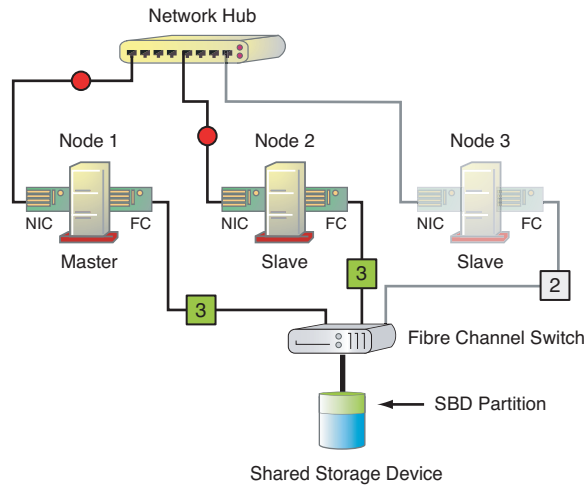
1. The master node (node 1) listens for a heartbeat over the LAN from nodes 2 and 3.
2. After 8 seconds (the default), the master node notifies node 2 over the LAN that there is a new cluster membership view that does not include node 3.

Node 3 continues to maintain the old cluster membership view and the old epoch number (2).

3. The master node and node 2 write a new epoch number (3) to the SBD partition.
4. NCS uses the information in the SBD partition to vote between the 2 cluster membership views.
5. Because there are 2 nodes in the new cluster membership view with a new epoch number, node 3 loses.
6. The master node and node 2 write a special token to the sector in the SBD sector for node 3.
7. Node 3 reads the special token, takes a poison pill, and abends.
8. NCS migrates the resources on node 3 to the master node and node 2.

The results look like the following:

Figure 8-4 (slide)



Fan-Out Failover

When a node fails in an NCS cluster, the cluster-enabled volumes and resources assigned to that node are migrated to other nodes in the cluster.

Although this migration happens automatically, you must design and configure where each volume and resource migrates during failover.

Emphasize that the clustering administrator and consultant are responsible for making sure that resources are configured to fan out correctly across the cluster when a node fails.

NCS lets you configure this resource distribution, but does not detect the network traffic, data, or services load of each node nor does it adjust the fanning out of resources.


Figure 8-5 (slide)

You will probably want to distribute, or fan out the volumes and resources to several nodes based on factors such as load balancing and the availability of installed applications.

For example, suppose you configure a 3-node NCS cluster as follows:

- DHCP server and iFolder resources assigned to node 2
- 2 web site resources assigned to node 2
- 2 email resources assigned to node 3.

The following illustrates how this setup might look:

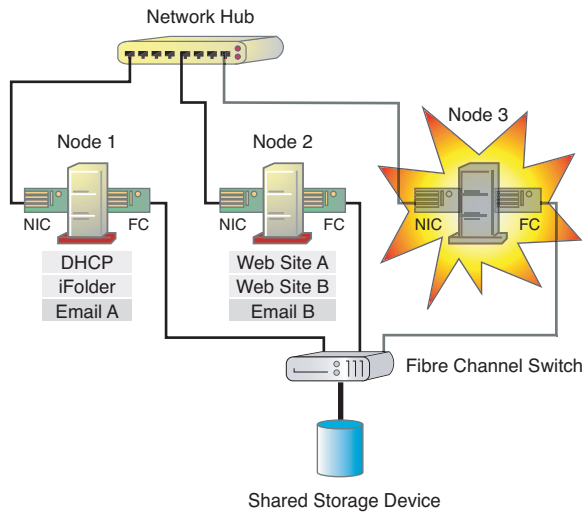


During normal cluster operation, each node remains in constant communication with the other nodes in the cluster by sending out a periodic heartbeat over the LAN.

However, suppose node 3 experiences hardware problems and fails. As a result, the users depending on node 3 for email access might lose their connections.

The following shows how the email resources migrate when node 3 fails (based on the network administrator configuration):

Figure 8-6 (slide)



When node 3 failed, NCS did the following:

1. Detected a failure of node 3 through the heartbeat channel on the LAN and confirmed the failure by checking the epoch numbers in the SBD partition and casting off node 3
2. Reassigned the shared data volumes for the email services (located on the shared disk system) to nodes 1 and 2
3. Restarted the email servers that were running on node 3 on nodes 1 and 2
4. Transferred the email resources (including the resource IP addresses) to node 1 and node 2

NCS knew where to migrate these resources because, when configuring the NCS cluster, the network administrator decided where the email resources assigned to node 3 should be reassigned if a failure occurred.

Because Email A is used by most employees, the network administrator made sure that Email A migrated to node 1, which has a faster processing chip and can handle the additional load.

Because Email B is used by a limited number of employees, the network administrator made sure that Email B migrated to node 2.

Although node 2 has a slower processing chip than node 3, employees using Email B will probably not notice the difference because the web sites on node 2 do not generate a significant amount of traffic.

In this example, the migration of resources happens quickly, and employees regain access to email within seconds and, in most cases, without having to log in again.

When the problems with node 3 are resolved and the network administrator starts running node 3, Email A and Email B remain running on nodes 1 and 2 because the resources are not configured to automatically failback to node 3.

Instead, the network administrator migrates Email A back to node 3, and then tests the service to make sure it works properly on the node.

Because Email B seems to be working well on node 2, the network administrator can leave the email service running on that node to allow for expansion and faster processing of Email A on node 3.

Selectively migrating resources from one node to another allows the network administrator to not only load-balance services, but to upgrade email servers and other cluster services during normal working hours without bringing down a server and temporarily disconnecting employees from vital services.



By default, NCS resources are configured with failback disabled. This configuration gives you the opportunity to bring a failed node back into the cluster and test its viability without resources automatically loading on the node.

However, you can set resources to automatically failback to a particular node when that node joins the cluster.

These rules focus on the importance of cluster-enabling all nodes in a cluster. If a node is not clustered, serious data and volume loss occurs.

Rules for Managing an NCS SCSI SAN

To avoid serious problems with migrating and failing over services, you must follow several rules when managing an NCS SCSI SAN.

When you install NCS with shared storage, each shared storage device is assigned a global unique ID (GUID), and a flag is set on the device to help Netware 6 distinguish between local server storage devices and shared storage devices.

This identification of shared storage helps prevent problems such as assigning a local volume and a shared volume to the same pool. If the server with the local volume fails, the shared volume is assigned another node, but the local volume remains inaccessible to users.

When working with shared storage, you must observe the following rules or risk data corruption or volume loss:

- Don't attach a noncluster server to the shared storage device unless you isolate the storage so the noncluster server has access only to its own volumes.

All servers attached to the shared storage device (whether in the cluster or not) have access to all volumes on the shared storage device unless you specifically prevent such access.

NCS manages access to shared volumes for all cluster nodes but cannot protect shared volumes from being corrupted by noncluster servers.

- Don't install NetWare 6 on a server that is attached to shared storage. The NetWare 6 installation deletes all NetWare partitions it finds on local and shared storage devices.
You must disconnect the shared device from the server before installing NetWare. After installation is complete, you can connect the server to the shared device.
- Don't perform NSS cluster volume operations (such as deleting, resizing, and renaming) from noncluster nodes (NetWare 6 servers running the NCS NLMs).
- If an application or users will read or write to data or files, store the volume containing the data or files on the shared storage device. Otherwise, keep the volume on the local hard drive.

Troubleshooting a 2-Node NCS SCSI SAN

Most problems in setting up a SAN result from errors in preparing and connecting devices on the SCSI bus. If you have problems with your SAN, use the following questions to check your configuration:

- Are the SCSI adaptor cards identical?
Although not required, you can reduce potential failures and troubleshooting by making sure that the SCSI adaptor cards are the same model and version.
- Is the SCSI adaptor card and driver multi-initiator enabled?
To find out, refer to the card and driver specifications. This information is normally listed on the manufacturer's web site. If not, call the manufacturer for the information.



Novell does not maintain a list of cards and drivers that are multi-initiator enabled.

- Do the SCSI cables have the same impedance?

Try to use quality SCSI cables with the same specifications. Use the same type and length from the same manufacturer.

- Are all SCSI devices turned on and all SCSI cables and power cables properly connected?
- Is the SCSI adaptor card seated and secure in the slot?
- Are all SCSI devices (hard drive and adaptor cards) assigned unique SCSI IDs?

SCSI ID numbers are assigned to SCSI devices and adaptor cards by the manufacturer. The lower the ID number, the higher the priority of the SCSI device or card.

SCSI hard drives are normally assigned an ID number of 0, giving them the highest priority. SCSI cards are given a lower priority number such as 7.

Make sure each SCSI hard drive and adaptor card have a different ID number. Also, make sure the SCSI hard drive has a lower ID number than the SCSI adaptor cards.

When configuring a 2-node SCSI cluster for NCS, consider using 7 for one card and 6 for the other card. Even if they are available, avoid using higher ID numbers (8–15).



For external SCSI devices such as a hard drive, the SCSI ID usually is set with a switch on the back of the device.

However, you might need to set the number by configuring a jumper on the hard drive board if the enclosure and the hard drive were not pre-assembled by the manufacturer.

- Are SCSI hard drive and SCSI adaptor cards terminated properly?

To ensure reliable communication on the SCSI bus, the ends of the SCSI bus must be properly terminated.

For a 2-node NCS SCSI cluster, the hard drive must not be terminated. The 2 adaptor cards are at the ends of the SCSI bus and should be terminated.

However, if an adaptor card fails, termination also fails. You can avoid this situation by using an external terminator at each end of the SCSI bus.

(Most newer SCSI devices use autotermination.)

- Have you done a low-level format on the SCSI hard drive?

Every SCSI hard disk must be physically low-level formatted, partitioned, and logically formatted before it can be used to store data.

Most SCSI drives are pre-formatted at the factory. However, if you connect a used SCSI hard drive to the NCS nodes for clustering, you must perform a low-level format before you can use the drive.

Because a low-level format destroys all data on the drive, make sure you back up the data before performing a low-level format.



SCSI hardware manufacturers (such as Adaptec) indicate that you must perform a low-level format if the drive was previously connected to a different SCSI card.

In addition, you might want to use the SCSI utility to configure the following BIOS settings of your SCSI adaptor cards:

Table 8-3

BIOS Setting	Suggested Configuration
Maximum Sync Transfer Rate	Set both SCSI adaptor cards to a common transfer speed. If one card is transferring data at a faster rate than the other, you might experience problems when running the cluster.

Table 8-3 *(continued)*

BIOS Setting	Suggested Configuration
Advanced Speed Increase Options	Disable advanced options designed to increase the speed of transferring data. These include Wide, Ultra, and Send Start Unit Command.
Extended BIOS Translation for DOS Drive	Turn off (or disable) drive translation. This function can be fatal to NetWare, and possibly cause an abend.
Reset SCSI bus	Disable to avoid SCSI Bus Reset by Third Party messages.

Make sure the BIOS settings for both SCSI adaptor cards are identical. You can then adjust the settings to resolve problems with the SAN.



For more on configuring SCSI devices, see www.paralan.com.



For tips on configuring SCSI adaptor cards for a NCS 2-node SCSI SAN, see <http://developer.novell.com/research/sections/netsupport/abend/2001/april/spv.htm> and <http://developer.novell.com/research/sections/netmanage/tips/2001/November/t011101.pdf>.

Exercise 8-1 Design and Set Up a 2-Node SCSI Clustering Configuration



40 minutes

Perform Part I as a group exercise. You might want to have one student draw a design for their 2-node cluster on the whiteboard.

Make sure you give students the name of the SCSI adaptor card.

Now that you know more about clustering with NCS, you are ready to set up a 2-node NCS cluster in your lab using an external SCSI hard drive for a SAN.

In this exercise, you design the 2-node SCSI cluster and set up the SAN for the installation of NCS.

Specifically you do the following:

- **Part I: Create a Design for a 2-Node SCSI Cluster**
- **Part II: Connect the SCSI Hard Drive to the SAN**
- **Part III: Check the Drivers for the SCSI Card and Hard Drive**
- **Part IV: Initialize the SCSI Hard Drive**

Part I: Create a Design for a 2-Node SCSI Cluster

The following provides cluster component names and IP addresses for a 2-node cluster for Digital Airlines:

Table 8-4

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.DEL.DIGITALAIR	192.168.11.14
Cluster nodes	
DA4.IS.DEL.DIGITALAIR (<i>first server</i>)	192.168.1.4
DA5.IS.LGA.DIGITALAIR (<i>second server</i>)	192.168.1.5
Cluster name (DN)	
DACluster.IS.LON.DIGITALAIR	192.168.1.16
Cluster nodes	
DA6.IS.LON.DIGITALAIR (<i>first server</i>)	192.168.1.6
DA7.IS.SYD.DIGITALAIR (<i>second server</i>)	192.168.1.7

Table 8-4 (continued)

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.TXL.DIGITALAIR	192.168.1.18
Cluster nodes	
DA8.IS.TXL.DIGITALAIR (<i>first server</i>)	192.168.1.8
DA9.IS.TYO.DIGITALAIR (<i>second server</i>)	192.168.1.9



In a classroom setting, students can work together to design, configure, and set up the 2-node SCSI SAN.

Note that the cluster name (DACluster) is the same for all clusters, but the context (such as IS.DEL.DIGITALAIR) and IP address are different for each cluster. *First server* and *second server* are references used in the exercise.

Emphasize that designing a cluster for a production environment is critical to the success of the cluster.

Using Figure 8-2 as reference and the information in Table 8-4, draw a 2-node cluster configuration that uses the following:

- A LAN that connects 2 NetWare 6 servers with a hub
- A SCSI SAN with a shared SCSI hard drive
- Two Windows 2000 workstations connected to the LAN through the hub

Use cluster names, node names, and IP addresses from the table.

Use the following space to draw your 2-node clustering configuration:

Part II: Connect the SCSI Hard Drive to the SAN

Now set up the SAN. In addition to connecting the SCSI hard drive to your 2 servers (such as DA4 and DA5), you must also remove all partitions from the SCSI hard drive in preparation for installing the SBD partition.

To do this, you need the following:

- NetWare 6 installed and running on both servers
- A SCSI adaptor card installed in your servers with a SCSI cable connected to each card
- A SCSI hard drive in a SCSI enclosure with 2 connectors
- Your 2-node SCSI configuration design from Part I

Do the following:

1. Bring down NetWare 6 on *your servers*; then turn off both servers.

You can bring down NetWare 6 on a server by pressing **Ctrl + Esc**, entering **1** for System Console, and then entering **DOWN** at the console prompt.
2. Plug in the SCSI hard drive to a power source; then make sure the drive is off.
3. Connect both servers to the SCSI hard drive enclosure using the SCSI cables.
4. Turn on the SCSI hard drive.
5. Turn on *your servers*.
6. As each server boots, look for the name of the SCSI adaptor card (such as Adaptec 2940) and the keystrokes for accessing the adaptor card BIOS (such as Ctrl+A).
7. Access the BIOS configuration utility for the SCSI adaptor card on each server.
8. Select the option that shows the SCSI ID numbers for the SCSI components associated with the server (the hard drive and the adaptor card).

Each adaptor card and the SCSI hard drive should have a different SCSI ID number.

For example, your SCSI hard drive might be assigned a SCSI ID number of 0 (the same on each server) while the SCSI adaptor card should be assigned a different SCSI ID (such as 6 or 7).
9. (Conditional) If 2 of the components use the same SCSI ID number, change the number for one of the components; then reboot the server.

Part III: Check the Drivers for the SCSI Card and Hard Drive

Now that the SCSI hard drive is connected to both servers and running, and you have checked the SCSI ID numbers, you can configure the SAN for NCS. Do the following:

1. After NetWare 6 loads (and the GUI interface appears) on both servers, make sure a driver for the SCSI adaptor card and a driver for the SCSI hard drive are installed on your servers by doing the following on each server:
 - a. Start the NetWare Configuration utility by entering **NWCONFIG** at the console prompt.



If a dialog appears indicating that new hardware has been detected, continue by selecting **F3**. NWCONFIG loads the drivers and exits.

Continue the exercise by starting NWCONFIG again.

- b. In the Configuration Options dialog, select **Driver Options**.
- c. In the Driver Options dialog, select **Configure disk and storage device drivers**.
- d. Press **Tab** and scroll through the list of selected drivers.
- e. Choose from the following:
 - If you see **SCSIHD** listed for the hard drive and a driver listed for your SCSI card, exit NWCONFIG by pressing **Esc** until an Exit the Install message appears; then select **Yes** and start at step 2.
 - If you do not see **SCSIHD** or a driver listed for your SCSI card, press **Tab** and select **Select an additional driver**; select the SCSIHD driver from the list; copy the driver by selecting **Yes** and select the default path.

Select **Yes** or **No** (depending on your need to select another driver); then press **Tab** and scroll through the disk drivers list.

You can also have students select the option to discover a driver; then use that driver instead of SCSIHD for the SCSI card.

You should see **SCSIHD** and the driver for your SCSI card listed with a **Currently Loaded** status.

- f. Exit NWCONFIG by pressing **Esc** until an Exit the Install message appears; then select **Yes**.
2. At the console prompt for each server, enter **SCAN FOR NEW DEVICES**; then enter **LIST DEVICES**.

After a few moments, you see the SCSI hard drive listed with a device ID number (such as 0x0005 or 0x000C) and a name (such as FUJITSU MAE3091LP).

The device ID number might be different on each server, but the hard drive name should be the same.



If you see **Unbound Device** instead of the SCSI hard drive name, the driver for the SCSI hard drive has not been properly installed.

To correct this situation, copy the **SCSIHD.CDM** and **SCSIHD.DDI** files from C:\NWSERVER\DRIVERS to C:\NWSERVER and load the driver from the command prompt; at the console prompt enter **SCAN FOR NEW DEVICES**; then enter **LIST DEVICES**.

The SCSI hard drive name is listed instead of **Unbound Device**.

3. Record the name of the SCSI hard drive and the device ID number for each server:

Table 8-5

Hard Drive Name	Device ID Numbers
	DAx:
	DAx:

The device ID number is normally the same on each server; however, some servers show different ID numbers.

Part IV: Initialize the SCSI Hard Drive

When you connect a SCSI hard drive for clustering that contains partitions and data, you must initialize the drive. This prepares the drive for installing the SBD clustering partition and for creating pools and volumes.

(Make sure you back up valuable data from the hard drive before initializing it.)



Initializing the hard drive removes the partitions and rewrites the MBR. This works most of the time in class for this exercise.

However, you might need to do a low-level format if you experience problems (such as a server abending when you create an NSS volume) with other exercises in sections 9, 10, and 11.

A low-level format can be performed with a SCSI utility or a DOS-based low-level format utility.

Do the following:

1. Start ConsoleOne from *your second server*; then log in as **admin** to *your second server*.

For example, if you are clustering DA4 and DA5, *your second server* is **DA5**.

2. Right-click the *second server* object (in **IS.xxx.DIGITALAIR** where *xxx* = your location container); then select **Properties**.

For example if you are clustering DA4 and DA5, right-click **DA5** in **IS.LGA.DIGITALAIR**.

3. List the devices connected to the server by selecting **Media > Devices**.

4. Select the *device ID number* of your SCSI hard drive.

The name of the SCSI hard drive appears in the Description field.

5. Select Initialize Hard Disk.

A message indicates that all partitions will be deleted.

6. Initialize the SCSI hard drive and delete all partitions by selecting Yes.**7. (Conditional) If you see a message indicating that the hard drive contains cluster management software select OK or Yes to remove the software and continue.****8. (Conditional) If you see a warning indicating a problem removing NSS pools or traditional volumes, select OK.**

For the exercise, don't investigate or fix the problem.

9. When the process is complete, check settings in the Media > Devices list for the SCSI hard drive device.

The capacity, free space, and unpartitioned space are all equal and include most space on the hard drive. In addition, no partitions are listed in the Partitions drop-down list.

10. Close the Properties dialog by selecting Cancel.**11. Close ConsoleOne by selecting File > Exit.**

With the SCSI hard drive connected to servers and initialized, and the SCSI drivers installed on both servers, your SAN is ready for installing and configuring the NCS software.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	What You Learned
<p>1. Identify the Purpose and Advantages of Implementing an NCS Solution</p>	<p>You learned the following:</p> <ul style="list-style-type: none"> ■ High Availability Terms: <ul style="list-style-type: none"> ■ Resource ■ Service ■ Availability ■ Uptime ■ Outage ■ Downtime ■ Reliability ■ Mean time between failures (MTBF) ■ Mean time to recovery (MTTR) ■ High Availability Definition: <ul style="list-style-type: none"> ■ 24x7x365 ■ 24x7x365 at 100% ■ 6-6 or 6-11 (am to pm) at 100% ■ Five 9s guaranteed ■ Computer System Outage Factors <p>These include physical (physical faults or hardware failures), design (hardware and software), operations (errors caused by operations personnel or users), environmental, and reconfiguration.</p> ■ Benefits and Features of an NCS High Availability Solution <p>These include items such as multinode all-active cluster (up to 32 nodes), consolidation of applications and operations, fan-out failover, and shared storage support.</p>

Objective	What You Learned
<p>2. Design and Set Up an NCS Cluster Configuration</p>	<p>Knowing the following helps you prepare to design and set up a simple 2-node NCS cluster configuration:</p> <ul style="list-style-type: none"> ■ Basic Clustering System Terms These terms include cluster, node, cluster resource, shared storage device, storage area network (SAN), migration, failover, and Fibre Channel ■ NCS Cluster Components <ul style="list-style-type: none"> ■ From 2 to 32 NetWare 6 servers configured to use IP ■ NCS 1.6 and NetWare 6 installed on each server in the cluster ■ A shared disk system connected to all servers in the cluster ■ High-speed Fibre Channel or SCSI shared disk system ■ Typical NCS Shared Disk System Cluster Configurations These include Fibre Channel and SCSI hard drive configurations. ■ NCS System Terms These terms include master node, cluster-enabled volumes and pools, shared storage device, cluster resource, heartbeats, tics, epoch numbers, split brain detector (SBD), and fan-out failover. ■ Rules for Managing an NCS SCSI SAN These rules cover tasks such as attaching a noncluster server to the shared storage device, and installing NetWare 6 on a server attached to a shared storage device. ■ Troubleshooting a 2-Node NCS SCSI SAN If you have problems with your SAN, use the series of questions in this objective to check your configuration.

Exercise Answers


The following are the exercise answers.

Exercise 8-1. Design and Set Up a 2-Node SCSI Clustering Configuration

Part I: Create a Design for a 2-Node SCSI Cluster

The following is the suggested design for your 2-node SCSI cluster:

Figure 8-7



SECTION 9 Install and Test NCS on a 2-Node Cluster

Duration: 2 hours 30 minutes

In this section you learn how to install, monitor, and test a 2-node NCS cluster for use in a non-production lab environment.

Objectives

1. [Verify NCS System Requirements](#)
2. [Create a Cluster by Installing NCS](#)
3. [Check Cluster Configuration Settings](#)
4. [Test and Monitor the Cluster](#)

Introduction

After designing a 2-node cluster with a SCSI shared disk drive and setting up and configuring the SAN, you can install and test NCS on the 2-node cluster.

In this section you prepare to cluster-enable data and network services by installing NCS and testing and monitoring the health of the cluster.



If you are upgrading NCS from a previous version, see “Installing or Upgrading Novell Cluster Services” at <http://www.novell.com/documentation/lg/ncs6p/index.html> or see Novell AppNotes at <http://developer.novell.com/research/sections/netsupport/abend/2001/October/x011001.htm> for steps on performing the upgrade.

Objective 1 **Verify NCS System Requirements**

Before installing NCS, your system must meet the following requirements:

- [Hardware Requirements](#)
- [Software Requirements](#)
- [License Requirements](#)
- [Shared Disk System Requirements](#)

Hardware Requirements

The following lists minimum hardware requirements for installing NCS 1.6 on a 2-node NCS cluster:

- Two NetWare 6 servers
- 256 MB of memory on all servers in the cluster (512 MB is recommended for failing multiple applications to the same server)
- At least one local disk device (not shared) for SYS on each server

If you are configuring a 2-node SCSI cluster, you must have a SCSI adaptor card installed in each server. In addition, the external SCSI hard drive and each SCSI adaptor card must have a unique SCSI ID.

You might want to use these requirements as a checklist and have students confirm that their hardware and software meet the requirements.

For example, as you discuss memory requirements for the servers, have students check each server to make sure the requirement is met.

Additional hardware might be necessary depending on how you use NCS.

Software Requirements

- NetWare 6 running on each cluster server (with the same service pack installed)
- All servers in the cluster configured for IP and on the same IP subnet (NCS is not IPX compatible.)
- An additional IP address reserved for the cluster and each cluster resource and cluster-enabled volume

The IP address assigned to a cluster resource is a secondary IP address that NCS uses to find and migrate the resource from one node to another in the cluster.

If a service (such as a web site) already has an IP address, that address is preserved as part of the service when you cluster-enable the service and is not used as a resource secondary IP address for clustering.

- All servers in the cluster in the same eDirectory tree

License Requirements

NCS requires a Cluster Server License for each server that is part of the cluster. The Cluster Server License allows a server to join a cluster. Cluster Server License objects are created in the same eDirectory context as the Cluster object.

Cluster Server Licenses for a 2-node cluster are provided with NetWare 6 and are added during NCS Services installation. You only need additional cluster server licenses if you have a 3-node or larger cluster.

Additional Cluster Server Licenses can be obtained from Novell or from your Novell Authorized Reseller.

Shared Disk System Requirements

A shared disk system is required for each cluster for data to be highly available. If you use a shared disk system, ensure the following:

- At least 15 MB of free disk space is available on the shared disk system for creating the SBD partition

NCS installation allocates one cylinder on one drive of the shared disk system for the special cluster partition.

- The shared disk system is properly set up and functions according to the manufacturer's instructions

Prior to installation, verify that all drives in your shared disk system are recognized by NetWare by entering LIST DEVICES on each server you will add to your cluster.

If any drives in the shared disk system do not show up in the list, consult your NetWare documentation or the shared disk system documentation for troubleshooting information.

- The disks contained in the shared disk system configured in a mirroring or RAID 5 configuration to add fault tolerance to the shared disk system



If the disks in the shared disk system are not configured to use mirroring or RAID 5, a single disk error can cause a volume failure. NCS does not protect against such faults.

Objective 2 Create a Cluster by Installing NCS

You must run the NCS installation when you do the following:

- Create a cluster
- Add nodes to an existing cluster
- Upgrade NCS software in an existing cluster

The installation does the following:

- Creates a cluster object in eDirectory
- Installs NCS software on servers you specify for your cluster

After running the installation the first time to create a cluster, run the installation again to add servers to your cluster or to upgrade NCS software on a cluster.

(Licenses for a 2-node cluster are provided with NetWare 6. The license file is located on the NetWare 6 CD, and the path to the license file is provided during NetWare 6 installation.)



For information on removing clustering services and the eDirectory clustering objects, see TID 10015339.


To install NCS for a cluster, do the following:

1. Insert the *NetWare 6* CD in a NetWare administrator workstation and allow NetWare Deployment Manager to launch.

(Or, run **NWDEPLOY.EXE** from the root of the CD to launch NetWare Deployment Manager.)

The following appears:

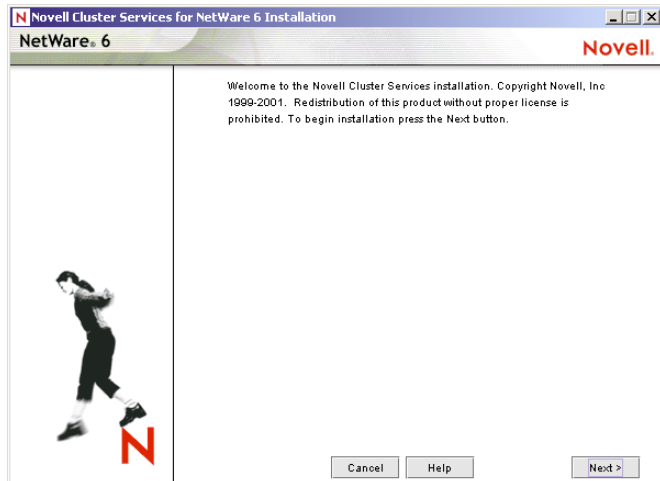
Figure 9-1



2. In Deployment Manager, open the **Post-Installation Tasks** folder.
3. Start the NCS installation by selecting **Install or Upgrade a NetWare Cluster**.

The Welcome screen appears:

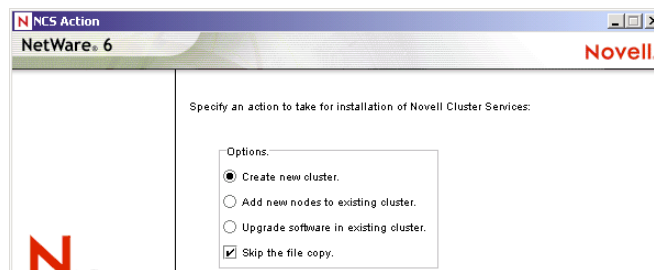
Figure 9-2



4. Continue by selecting **Next**.

The NCS Action screen appears:

Figure 9-3



Use the options available in the NCS Action screen to

- ❑ Create a cluster
- ❑ Add nodes to an existing cluster
- ❑ Upgrade NCS software in an existing cluster

The Skip the File Copy option is useful if NCS files have been copied to cluster nodes and you want to save time.

For example, the NetWare 6 installation program copies NCS files to every NetWare 6 server. If you have already installed NetWare 6 on the cluster servers, leave the Skip the File Copy option selected.

Even though NCS files might exist on each NetWare 6 server, you still need to run the NCS installation program to configure and set up cluster nodes.

Make sure students understand the purpose of the Skip the File Copy option.

When you install NetWare 6 on a server, the NLMs for clustering are copied to the server and do not need to be copied during a cluster installation.

However, if you are upgrading NCS software, you will probably want to deselect the option to let NCS installation copy new clustering files to the servers in the cluster.

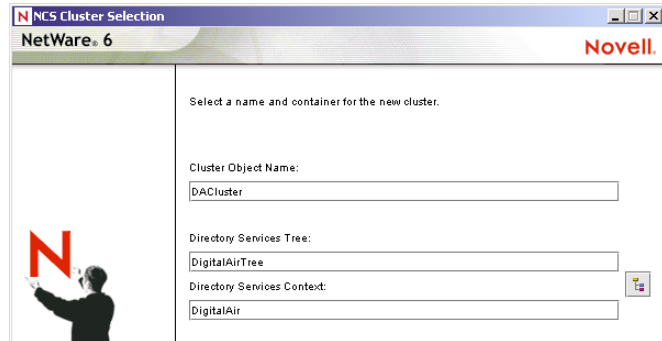


If the Skip the File Copy option is not selected, existing NCS files on the servers will be replaced. However, this will not otherwise affect the installation.

5. Select **Create a New Cluster**; then select **Next**.

The NCS Cluster Selection screen appears:

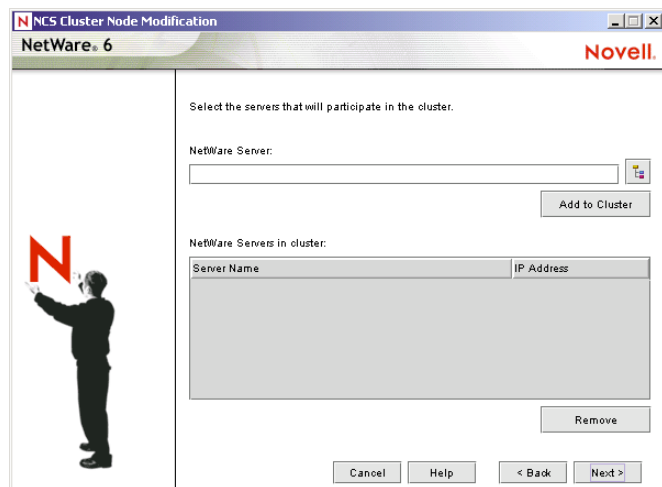
Figure 9-4



6. Enter the name for the cluster object you are creating and specify the eDirectory tree and context where you want it created; then select Next.

The NCS Cluster Node Modification screen appears:

Figure 9-5



7. Add the servers you want in the cluster to the NetWare Servers in Cluster list by doing one of the following:
 - Enter the name of the server in the NetWare Servers box; then select **Add to Cluster**.
 - Select the browse button, find and select a server, and select **Add**. Repeat this for each server you want in the cluster. When you finish, select **OK**.

The installation program detects each server and then adds the server name and IP address to the list.

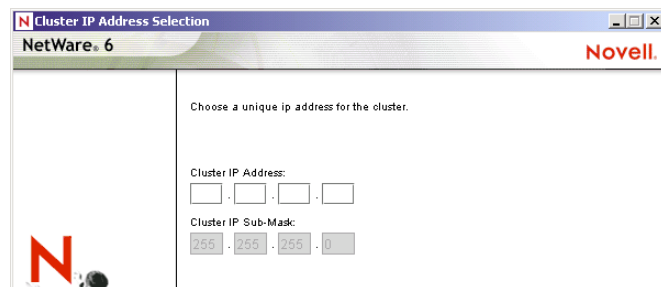
You can remove a server you added to the list by selecting the server and selecting **Remove**.

If the server you are adding has more than one IP address, you are prompted to select the IP address you want NCS to use.

8. When you finish adding servers to the list, select **Next** to continue.

The Cluster IP Address Selection screen appears:

Figure 9-6



9. Enter a unique IP address for the cluster.

The cluster IP address is separate from each server IP address in the cluster and is required for external programs like ManageWise® to get cluster status alerts.

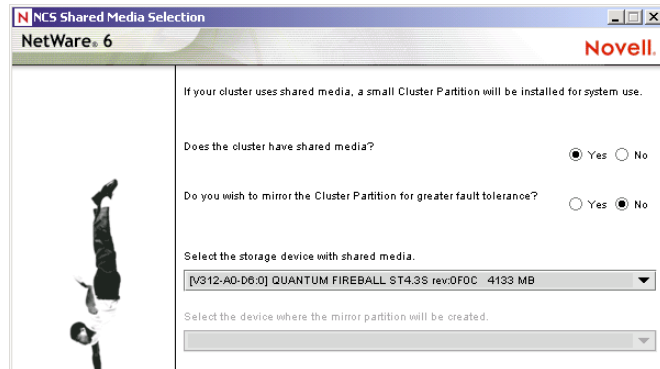
NetWare Remote Manager and ConsoleOne also require a cluster IP address.

The cluster IP address is assigned to the master node and remains with the master node regardless of which server is assigned as the master node.

10. Continue by selecting **Next**.

The NCS Shared Media Selection screen appears:

Figure 9-7



11. Specify whether your cluster has a shared disk system; if so, select the drive where you want the special cluster partition created.

NCS requires a special cluster partition on the shared disk system. You are also given the option of mirroring the partition for greater fault tolerance.




To create the SBD partition you must have at least 10 MB of free space that is not part of an NSS partition on one of the shared disk drives. If no free space is available, the shared disk drives can't be used by NCS.

12. Continue by selecting **Next**.

The Start Clustering screen appears:

Figure 9-8



13. Choose whether you want the servers you are adding to your cluster to start NCS software after the installation.

If you choose to not start NCS software on each server that you upgrade or add to your cluster, you must manually start the server after the installation or you must reboot cluster servers to automatically start the server.

You can manually start NCS by entering **LDNCS** at the server console on each cluster server.

14. Continue by selecting **Next**.
15. (Conditional) Specify the location of the license files or browse and select a path; then select **Add**.

This screen appears only if you are installing or upgrading a 3-node or larger cluster.

You can install without licenses. If you install without licenses and you have a cluster with more than 2 nodes, you must manually install the licenses later using Novell iManager.



NCS will not function without the proper licenses in place.

After the installation program recognizes the license, the Summary screen appears:

Figure 9-9



Novell Cluster Services should be listed as the product to be installed.

16. Begin the NCS installation by selecting **Finish.**

The installation program creates a cluster object in eDirectory and installs NCS on the servers you specified to be nodes in your cluster.

During installation, several clustering NLMs are loaded on each server (such as NISP.NLM, NCSPROXY.NLM, and SBDLIB.NLM).

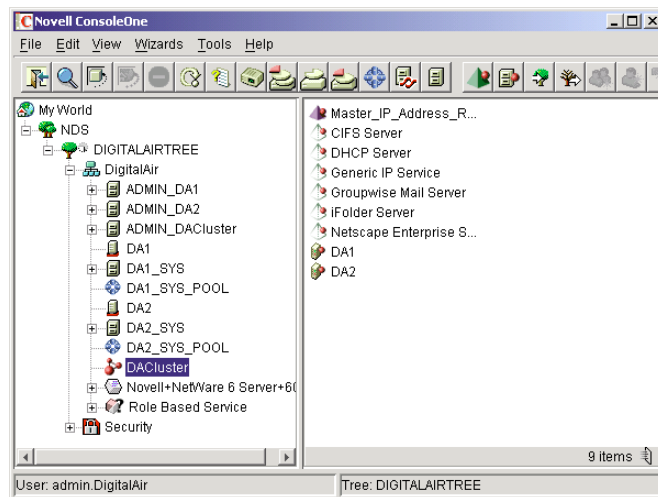
When installation is complete, the Cluster Membership Monitor appears on each server, listing each server node with an UP status.

Objective 3 Check Cluster Configuration Settings

In this section you learn how to access basic cluster object configuration settings and what the settings mean.

When you install NCS, several cluster objects are created in eDirectory, as shown in the following:

Figure 9-10



Students do not need to have a thorough understanding of each setting. However, they should understand the flexibility of configuring an NCS cluster to meet specific cluster requirements.

Encourage students to follow along as you show the various configuration settings.

These include the following:

- [Cluster ADMIN Object](#)
- [Cluster Object](#)
- [Master IP Address Resource Object](#)
- [Cluster Server Node Objects](#)

You can view and edit cluster object settings in ConsoleOne or NetWare Remote Manager.

The following provides steps for accessing the cluster object settings in ConsoleOne. Some steps for using NetWare Remote Manager are included in the exercises. (For more information, see <http://www.novell.com/documentation/lg/ncs6p/index.html>.)

Cluster ADMIN Object

An ADMIN object (such as ADMIN_DACLuster) is created for the cluster, and lets you set attributes, trustees, and rights for cluster objects and operations. You can also show creation facts about objects in the cluster.

For setting up and testing clustering on a 2-node SCSI configuration with an external drive, you do not need to know how to configure the properties of the cluster ADMIN object.

Cluster Object

The cluster object contains several objects necessary for configuring and running the cluster. The cluster objects include the following:

- **Master IP address resource.** When you install a cluster, you assign an IP address to the cluster. The IP address and the scripts for loading and unloading the IP address are stored in this object.

The master IP address resource object is new to NCS 1.6. The address is always assigned to the master node and allows the cluster to advertise on the LAN as though it were a virtual server.

- **Cluster server nodes.** Configuration settings for each server in the cluster are stored in a node object. These settings include the server IP address.

Some students might be confused about services and servers in a cluster. This is especially true of the role of the master node.

Help students understand that the master node is determined by where the Master IP Address resource is running, and not by a specific node configured as the master node in the cluster.

- **Resource templates.** These templates let you quickly create resources for a variety of services including DHCP, iFolder, and GroupWise.

The following are basic properties of an NCS cluster object:

- [Timeout and Quorum Membership](#)
- [Cluster Protocol](#)
- [Cluster IP Address and Port](#)
- [Resource Priority](#)
- [Cluster Email Notification](#)

Timeout and Quorum Membership


When you first start the cluster, NCS waits for a specific amount of time and for a specific number of nodes to join the cluster before starting. The specific number of nodes is called a *quorum*.

To check quorum properties in ConsoleOne, do the following:

1. Right-click the cluster object; then select **Properties**.
2. In the Properties dialog, select the **Quorum** tab.

The Cluster Quorum Trigger settings appear:

Figure 9-11



You can configure the following:

- **Timeout.** Use to specify the amount of time to wait for the number of servers defined in the Membership field to run.

If the timeout period elapses before the quorum membership reaches its specified number, resources load on the servers that are running in the cluster.

For example, if you specify a Membership value of 4 and a Timeout value equal to 30 seconds, and after 30 seconds only 2 servers are running in the cluster, resources begin loading on the 2 servers that are running.

- **Quorum Membership.** Use to specify the number of nodes that must be running in the cluster before resources start to load.

When you start the nodes in your cluster, NCS reads the number specified in the Membership field and waits until that number of servers is running in the cluster before it loads resources.

Set the Membership value to a number greater than 1 so resources don't load on the first server brought up in the cluster.

For example, if you set the Membership value to 4, 4 servers must be up in the cluster (within the timeout period) before resources load and start.

By ensuring an adequate number of nodes in the cluster before clustering starts, you make sure that all resources are not loaded on the first server to join the cluster.

3. Save the configuration settings by selecting **OK**.

Cluster Protocol

You can use the Cluster Protocol tab pages to view or edit the transmit frequency and tolerance settings for all nodes in the cluster, including the master node.

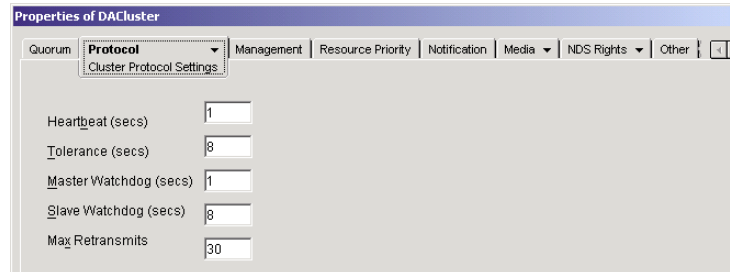
The master node is generally the first node brought online in the cluster. However, if that node fails, any of the other nodes in the cluster can become the master.

To check Cluster Protocol properties in ConsoleOne, do the following:

1. Right-click the cluster object; then select **Properties**.
2. In the Properties dialog, select the **Protocol** tab.

The following appears:

Figure 9-12



This tab has 2 pages: Cluster Protocol Settings and Cluster Protocol Internals.

The Internals page lets you view the script used to configure cluster protocol settings. However, you cannot edit the script.

3. Use the Settings page to make changes to the following cluster protocol properties:
 - **Heartbeat.** You can set the heartbeat to indicate the amount of time between LAN transmits for all nodes (except the master node) in the cluster.
For example, if you set this value to 1, nonmaster nodes in the cluster send a signal that they are alive to the master node every second.
 - **Tolerance.** You can set the tolerance to specify the amount of time the master node gives all other nodes in the cluster to signal that they are alive.

For example, if you set this value to 8 and the master node does not receive an “I’m alive” signal from a node in the cluster within 8 seconds, that node is cast off from the cluster.

You might want to increase the tolerance value if there is significant traffic on the LAN and you want to ensure that the master node waits long enough before initiating the cast-off process.

You might want to decrease the tolerance value if you feel the master node is not responding quickly enough to meet the design specifications of the cluster.

However, unless there is some compelling reason to change this setting, keep the tolerance value at 8 seconds. This setting is optimal for most installations of NCS cluster.

- **Master Watchdog.** You can set Master Watchdog to specify the amount of time between transmits for the master node in the cluster.

For example, if you set this value to 1, the master node in the cluster transmits an “I’m alive” signal to all other nodes in the cluster every second.

- **Slave Watchdog.** You can set Slave Watchdog to specify the amount of time the master node has to signal that it is alive.

For example, if you set this value to 8 and the nonmaster nodes in the cluster do not receive an “I’m alive” signal from the master within 8 seconds, the master node is cast off from the cluster and one of the other nodes becomes the master node.

- **Max Retransmits.** You can set this option to the number of times the master node waits for a heartbeat from another node before casting it off from the cluster.

4. Save the configuration settings by selecting **OK**.
5. Restart the cluster.



You should not make any changes to the configuration settings unless you check with Novell Technical Support or a qualified NCS clustering consultant.

Cluster IP Address and Port

When you install NCS, you assign an IP address to the cluster. The cluster IP address normally does not need to be changed, but you can change it if needed.

The default cluster port number is 7023. It is assigned when the cluster is created. The cluster port number does not need to be changed unless a conflict is created by another resource using the same port number.

If there is a port number conflict, change the port number to any other value that doesn't cause a conflict.

To check the cluster IP address and port number using ConsoleOne, do the following:

1. Right-click the cluster object; then select **Properties**.
2. On the Cluster Object property page, select the **Management** tab.
3. Make necessary changes; then save the changes by selecting **OK**.

Resource Priority

Some students might have questions about the colors associated with the resource priority list.

The colors have no specific significance in relation to the order or placement of resources in the list.

You can use the Resource Priority configuration settings to control the order in which multiple resources start on a given node when the cluster is brought up or during a failover or failback.

For example, if a node fails and 2 resources fail over to another node, the resource priority determines which resource loads first.

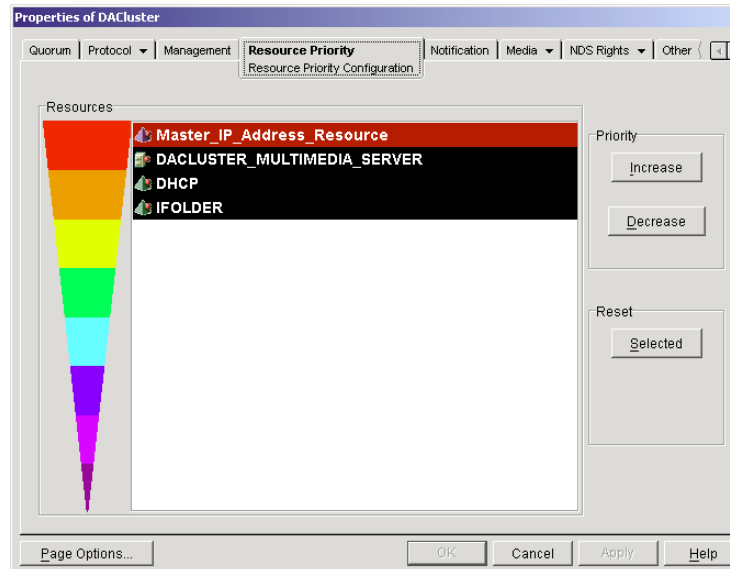
This is useful for ensuring that the most critical resources load first and are available to users before less critical resources.

To check resource priorities using ConsoleOne, do the following:

1. Right-click the cluster object; then select **Properties**.
2. In the Properties dialog, select the **Resource Priority** tab.

The following appears:

Figure 9-13



3. To change the priority for a resource, select the resource in the list and then select the **Increase** or **Decrease** button to move the resource up or down in the list.

This lets you change the load order of the resource relative to other cluster resources on the same node.

You can also select a resource and then click the **Selected** button to reset the resource to its default load order.

4. Save changes made to resource priorities by selecting **Apply**.

Cluster Email Notification

You can automatically send email messages for cluster events such as cluster and resource state changes or nodes joining or leaving the cluster with cluster email notification.

This feature lets you keep yourself and other administrators informed about changes to the status of the cluster without showing the Cluster State view (ConsoleOne), the Cluster Status view (NetWare Remote Manager), or the Cluster Membership Monitor (NetWare 6 server).

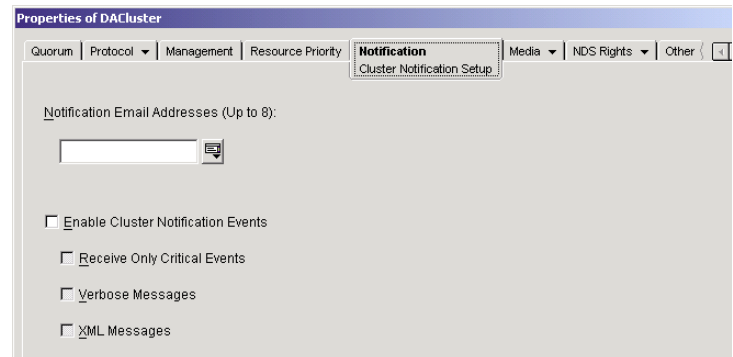
You can enable or disable email notification for the cluster and specify up to 8 administrator email addresses for cluster notification.

To enable cluster email notification using ConsoleOne, do the following:

1. Right-click the cluster object; then select **Properties**.
2. In the Properties dialog, select the **Notification** tab.

The following appears:

Figure 9-14



3. Enable the email notification feature and create the notification list:
 - a. Select the **Enable Cluster Notification Events** box.
 - b. Enter an email address in the field provided.
 - c. Select the button next to the field to add the address to the list.
 - d. Repeat steps **b** and **c** for each address.
4. Select the type of cluster events you want administrators to receive messages for:
 - Receive Only Critical Events.** Use to receive notification of critical events such as a node failure or a resource going comatose.
 - Verbose Messages.** Use to receive notification of all cluster state changes including critical events and resource state changes and nodes joining and leaving the cluster.
 - XML Messages.** Use to receive notification of all cluster state changes in XML format.
XML format messages can be interpreted and formatted with a parser that lets you customize the message information for your specific needs.

5. Save changes made to email notification by selecting **Apply**.

Master IP Address Resource Object

You can view information such as loading and unloading scripts for the cluster; start, failover, and failback node settings; and nodes associated with the cluster by showing the Master_IP_Address_Resource object properties.

This object is assigned to the master node in the cluster.

To view the settings in ConsoleOne, do the following:

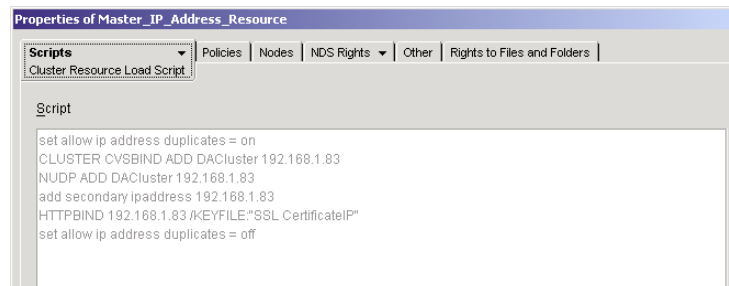
1. Select the cluster object.
2. On the right side of the ConsoleOne display screen, right-click the **Master_IP_Address_Resource** object.
3. Select **Properties**.
4. In the Properties dialog, select the **Scripts** tab.

The Properties dialog appears:

Make sure students understand that the Master_IP_Address Properties dialog is for viewing properties only.

NCS updates configuration information in the dialog.

Figure 9-15



You can view (but not edit) the following:

- **(Node) Cluster resource load and unload scripts.** Like all other cluster resources, a load script and an unload script are provided for taking offline, bringing online, or migrating the cluster resource.

Notice that the commands in the master IP address unload script are in the opposite order of those in the load script.

You cannot change the load and unload scripts for the master IP address resource, but you can change the IP address in the Cluster Object Properties dialog.

- **Policies.** You can view the current start, failover, and failback settings for the master IP address.
- **Nodes.** You can view the nodes assigned to the master IP address resource. Because the master IP address is assigned to the master node, the first node in the list is always the master node.

Because NCS assigns this resource to the master node, you can only view the assigned list.

5. When you finish viewing the information, close the Properties dialog by selecting **Cancel**.

Cluster Server Node Objects

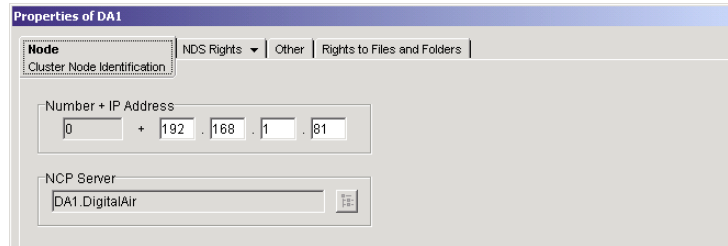
You can view or edit the cluster server node number or IP address of the selected node. You can also view the context for the NetWare Server object.

To check cluster server node properties in ConsoleOne, do the following:

1. Select the cluster object.
2. On the right side of the ConsoleOne display screen, right-click a cluster server node.
3. Select **Properties**.
4. In the Properties dialog, select the **Node** tab.

The Properties dialog appears:

Figure 9-16



You can configure the following:

- **(Node) Number + IP Address.** Use to specify the cluster node number and IP address for the selected node.
If the cluster node number or IP address changes for the selected node, the new information is not automatically updated in eDirectory.
- **NCP Server.** Use to view the context for the NetWare Server object. This field cannot be edited.

5. Save the configuration changes by selecting **Apply**.



40 minutes

Exercise 9-1 *Install and Check NCS on a 2-Node Cluster*

With the SCSI SAN connected to the NetWare 6 servers, you are ready to set up the clustering software.

In this exercise you install NCS for the 2-node SCSI clustering solution you designed and set up in section 8.

You need the following information to install the NCS software:

Table 9-1

For the sake of the scenario, some students may question the validity of clustering 2 nodes across a WAN.

If the issue arises, emphasize that the clustering exercises are meant to be replicated in a lab (not a production) environment.

Also emphasize the ability to cluster-enable servers located anywhere in the same tree.

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.DEL.DIGITALAIR	192.168.1.14
Cluster nodes	
DA4.IS.DEL.DIGITALAIR (<i>first server</i>)	192.168.1.4
DA5.IS.LGA.DIGITALAIR (<i>second server</i>)	192.168.1.5
Cluster name (DN)	
DACluster.IS.LON.DIGITALAIR	192.168.1.16
Cluster nodes	
DA6.IS.LON.DIGITALAIR (<i>first server</i>)	192.168.1.6
DA7.IS.SYD.DIGITALAIR (<i>second server</i>)	192.168.1.7
Cluster name (DN)	
DACluster.IS.TXL.DIGITALAIR	192.168.1.18
Cluster nodes	
DA8.IS.TXL.DIGITALAIR (<i>first server</i>)	192.168.1.8
DA9.IS.TYO.DIGITALAIR (<i>second server</i>)	192.168.1.9

In this exercise, you do the following:

- [Part I: Install NCS on the 2-Node SCSI Cluster](#)
- [Part II: Verify That NCS Is Loaded and Configured on Each Cluster Server](#)
- [Part III: Check the Cluster Configuration Settings in ConsoleOne](#)

Part I: Install NCS on the 2-Node SCSI Cluster

To install NCS, do the following:

1. Make sure you are logged in to *both servers* as **admin** from *your second workstation*.

2. Insert the **NetWare 6 OS** CD in the workstation and allow NetWare Deployment Manager to launch.
If the program does not launch, run **NWDEPLOY.EXE** from the root of the CD.
3. From the NetWare Deployment Manger Select Product screen, open the **Post-Installation Tasks** folder.
4. Start the NCS installation by selecting **Install or Upgrade a Novell Cluster**.
The Welcome screen appears.
5. Continue by selecting **Next**.
The NCS Action screen appears.
6. Make sure **Create new cluster** and **Skip the file copy** are selected; then continue by selecting **Next**.
The NCS Cluster Selection screen appears.
7. Enter the following:
 - Cluster object name: **DACluster**
 - Directory Services tree: **DigitalAir-Tree**
 - Directory Services context: **IS.xxx.DigitalAir**
For example, if you are clustering DA4 and DA5, your context is **IS.DEL.DIGITALAIR**.
8. Continue by selecting **Next**.
The installation program authenticates to the eDirectory tree you listed in the NCS Cluster Selection screen.
The NCS Cluster Node Modification screen appears.
9. Add *your first server* and *your second server* to the list of Netware servers in the cluster:
 - a. Select the browse button at the right of the NetWare Server box.

The installation program scans for objects and then shows a Browser dialog.

- b. In the left pane of the dialog, select the **IS.xxx.DIGITALAIR** container (where *xxx* = your location container) for *your first server*.

For example, if you are clustering DA4 and DA5, your first server is **DA4** and the container is **IS.DEL.DIGITALAIR**.

- c. In the right pane of the dialog, select *your first server* object; then select **Add**.

Your first server is added to the Selected Items list at the bottom of the dialog.

- d. In the left pane of the dialog, select the **IS.xxx.DIGITALAIR** container for *your second server*.

For example, if you are clustering DA4 and DA5, your second server is **DA5** and the container is **IS.LGA.DIGITALAIR**.

- e. In the right pane of the dialog, select *your second server* object; then select **Add**.

Your second server is added to the Selected Items list at the bottom of the dialog.

- f. From both servers show the logger screen by pressing **Ctrl + Esc** and selecting **2**.

- g. From your workstation add the names and IP addresses of your servers to the **NetWare Servers in cluster** list by selecting **OK**.



You can remove a server you added to the list by selecting the server and selecting **Remove**.

The installation program accesses each server and then adds the server name and IP address to the list.

- h. Watch the logger screen on both servers for any messages.

As the installation program accesses each server, it loads several cluster modules. Any error messages are also shown in the logger screen.

10. When the installation program finishes accessing and listing the servers, continue by selecting **Next**.

The Cluster IP Address Selection screen appears.

11. For the cluster IP address enter the *cluster IP address* indicated in Table 9-1; then select **Next**.

For example if you are clustering DA4 and DA5, the cluster IP address is **192.168.1.14**.

The NCS Shared Media Selection screen appears.

The shared SCSI hard drive should be listed with **Yes** selected for “Does the cluster have shared media?” and **No** selected for “Do you wish to mirror the Cluster Partition for greater default tolerance?”

For this exercise (or for testing purposes) you do not need to mirror the SCSI hard drive.

12. Accept the default settings by selecting **Next**.

The Start Clustering screen appears.

13. Make sure the **Start clustering on newly added or upgraded servers after installation** option is selected; then select **Next**.

The Summary screen appears. NetWare Cluster Services should be listed as the product to be installed.

14. Start the NCS installation by selecting **Finish**.

The installation program creates a cluster object in eDirectory using the name you specified, and loads NCS on the servers you selected to include in the cluster.



If you receive a replica error while installing NCS, create a replica on each server that holds a replica of the other server.

For example, if you are using DA6 and DA7 in your cluster, make sure DA6 holds a replica of DA7 (and vice versa); then start the installation again.

15. (Optional) View the Readme file by selecting **View** from the Installation Complete dialog.
16. End the installation by selecting **Close** from the Installation Complete dialog.
17. Exit NetWare Deployment Manager by selecting **Cancel**; then select **Yes**.
18. Remove the **NetWare 6 OS** CD from the workstation.

Part II: Verify That NCS Is Loaded and Configured on Each Cluster Server

During installation, NCS is configured and loaded on each server in the cluster. You can verify this by doing the following on each server:

1. Make sure that the **Cluster Membership Monitor (CMON)** appears.

You see nodes 01 and 02 listed with a blinking Up status. One of the Up messages is yellow, indicating the master node.

Notice that node 1 is *your first server* and node 2 is *your second server* (check the top of the membership monitor screen).

Now that both servers have joined the cluster, they are referred to as nodes instead of servers.

2. Press **Ctrl + Esc**.

CLUSTER RESOURCE SCREEN and CMON Screen should be listed as options.

3. Select System Console by entering **1**.
4. View the contents of AUTOEXEC.NCF by entering **EDIT AUTOEXEC.NCF** at the console prompt.
5. Scroll to the bottom of the AUTOEXEC.NCF file.

You see an LDNCS.NCF line. This command starts NCS on the server.

6. View the contents of LDNCS.NCF by pressing **Esc**; then enter **LDNCS.NCF** at the text editor prompt.

Several clustering NLMs are included in the file. Notice that the last command in the file instructs the server to join the cluster.

7. Exit the NetWare Text Editor by pressing **Esc** twice; then select **Yes**.

Part III: Check the Cluster Configuration Settings in ConsoleOne

Review the property values students record for the cluster and the nodes to make sure students understand the impact of those values on the operation of the cluster.

When you finish installing a cluster, check the basic cluster and node properties to make sure the correct settings are included in the cluster object.

You might also want to record these settings for future reference.

For example, to check and record the settings for the DACluster object, do the following:

1. On both workstations, make sure a shortcut exists for ConsoleOne.
2. Start ConsoleOne using the **ConsoleOne** shortcut on *your first server*; then open the IS container where *your first server* is located.

3. Right-click the **DACluster** object in the container; then select **Properties**.
4. Record the cluster configuration settings:

Table 9-2

Properties Tab	Settings	Values
Quorum	Timeout	
	Membership	
Protocol	Heartbeat	
	Tolerance	
	Master WatchDog	
	Slave WatchDog	
Management	Max Retransmits	
	IP Address	
	Port	

5. When you finish, select **Cancel**.
6. Select the **DACluster** object, right-click the *your first server* node object in the right panel; then select **Properties**.
7. Record the node settings:

Table 9-3

Properties Tab	Settings	Values
Node	Node Number	
	IP Address	
	NCP Server	

Notice that the node number is 0 even though your first server is listed as node 1 in the Cluster Membership Monitor. When identifying nodes internally, NCS starts counting nodes from 0.

8. When you finish, select **Cancel**.

9. Right-click the *your second server* node object in **DACluster**; then select **Properties**.
10. Record the node settings:

Table 9-4

Properties Tab	Settings	Values
Node	Node Number	
	IP Address	
	NCP Server	

11. When you finish, select **Cancel**.
12. Right-click the **Master_IP_Address_Resource** object in **DACluster**; then select **Properties**.
13. Select **Scripts** and view the load and unload scripts for the master IP address resource.

Whenever the master node fails, NCS uses the load and unload scripts to migrate the resource to a healthy node.

Notice that the IP address for the cluster is deleted and added each time the resource migrates from one node to another.

14. Select **Nodes**.
Your first and second servers appear in the assigned list. If your first server (currently the master node) fails, the master IP address resource is migrated to your second server and the cluster IP address continues to be broadcast on the network.
15. Close the Properties dialog by selecting **Cancel**.
16. Exit ConsoleOne by selecting **File > Exit**.

(End of Exercise)

Objective 4 Test and Monitor the Cluster

To test and monitor the cluster state, you need to learn about the following:

- [Cluster State and Cluster Status Views](#)
- [Console Prompt Commands](#)

Cluster State and Cluster Status Views

View and discuss the Cluster State view in ConsoleOne and the Cluster Status view in NetWare Remote Manager.

ConsoleOne and NetWare Remote Manager provide a special status screen for cluster objects that you can use to perform tasks such as view the cluster state, migrate resources, check an events log, and print or save an HTML report (in ConsoleOne only) on the cluster state.

To view the status screen from ConsoleOne, do the following:

1. Select the *cluster object* for your cluster.
2. Select **View > Cluster State View**.

To view the status screen from NetWare Remote Manager, do the following:

1. Open a web browser (such as Internet Explorer) on your workstation.
2. Enter a URL that includes the IP address to access a server in the cluster, plus port **8008** or **8009**.

For example, if the IP address is 192.168.1.1, enter **HTTPS://192.168.1.81:8009**.

3. When requested, log in to NetWare Remote Manager by entering the appropriate ID and password.

NetWare Remote Manager appears in the web browser.

4. In the left column under the **Clustering** section (near the bottom), select **Cluster Management**.

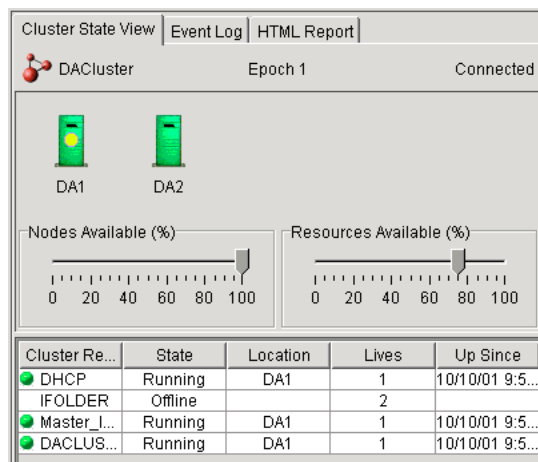
The following tabs or links are provided in the status screen:

- [Cluster State or Status View](#)
- [Event Log](#)
- [HTML Report \(ConsoleOne Only\)](#)

Cluster State or Status View

Selecting Cluster State View (ConsoleOne) or Cluster Status (NetWare Remote Manger) shows a detailed report of the state of your cluster:

Figure 9-17



The cluster object name and the epoch number appear at the top of the view. The epoch number indicates the number of times the cluster state has changed. The cluster state changes every time a node joins or leaves the cluster.

An icon for each cluster server appears in the middle; the resources running in the cluster appear in a list at the bottom. Gauges indicate the percentage of nodes and resources available in the cluster.

The cluster node and resource icons display in different colors, depending on their operating state:

- **Green.** When the icons are green, nodes and resources are in a normal operating condition.
- **Red.** When a node icon is red with a break in the icon, the node has failed.

When a resource icon is red, the resource is waiting for administrator intervention.

- **Gray.** When a node icon is gray with no break in the icon, that node is not a member of the cluster, or its state is unknown.
- **Blank (or no color).** When a resource is blank or has no colored icon, it is unassigned, offline, changing state, or in the process of loading or unloading.

In addition, the yellow ball in the middle of a node icon designates the master node in the cluster.

Event Log

Selecting Event Log gives you a detailed history of your cluster:

Figure 9-18

Cluster State View Event Log HTML Report			
Timestamp	Node	Resource	Cluster Event
10/11/01 7:45...		IFOLDER	NDS Sync
10/11/01 7:45...		IFOLDER	Offline
10/11/01 7:45...		IFOLDER	Administrative ...
10/10/01 10:45...		IFOLDER	Offline
10/10/01 10:45...		IFOLDER	Offline
10/10/01 10:45...		IFOLDER	NDS Sync
10/10/01 10:25...		DACLUSTER_...	Running on DA1
10/10/01 10:24...		DHCP	Running on DA1
10/10/01 10:24...		Master_IP_Add...	Running on DA1
10/10/01 10:24...		DACLUSTER_...	Running on DA1
10/10/01 10:24...		DHCP	Running on DA1
10/10/01 10:24...		Master_IP_Add...	Running on DA1
10/10/01 10:24...	DA2		Joined at epoc...

Every time the cluster state changes, a new event is added to the event log.

You can perform the following while viewing the event log:

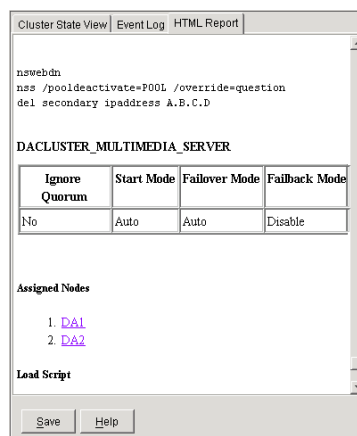
- Sort events in the log by selecting the column headings of the table
- Reverse the sort order by pressing the Shift key while selecting a column heading
- Save the event log to a file

The event log is stored on the SBD cluster partition on the shared storage device, which ensures that the event log is always available.

HTML Report (ConsoleOne Only)

Selecting the HTML Report tab for the cluster object in ConsoleOne shows a more detailed report of the state of your cluster:

Figure 9-19



You can save this report to an HTML file for printing or viewing with a web browser.

Console Prompt Commands

To view a list of commands, enter **HELP CLUSTER** or **HELP SBD** at a server console prompt, or enter **HELP** and a command (such as **HELP CLUSTER VIEW**) to see a description and example.

You can use the following to view the status of the cluster and SBD partition, and to load and unload clustering:

Table 9-5

Command	Status Information
CLUSTER VIEW	<ul style="list-style-type: none"> ■ Number of the current node ■ Cluster epoch number ■ Master node number ■ List of nodes that are members of the cluster
CLUSTER POOLS	<ul style="list-style-type: none"> ■ List of cluster pools ■ Node assigned to each pool
CLUSTER RESOURCES	<ul style="list-style-type: none"> ■ List of cluster resources ■ The state, node assignment, and lives of each resource
CLUSTER STATS DISPLAY	<ul style="list-style-type: none"> ■ Node number and name ■ Heartbeat information
CLUSTER STATS CLEAR	Clears the status information from the screen
CLUSTER STATUS <i>resource</i>	<ul style="list-style-type: none"> ■ Number of lives ■ State ■ Assigned node
SBD VIEW	Shows information such as heartbeat ID, node state, and group epoch number stored in the SBD partition

Table 9-5 (continued)

Command	Status Information
SBD VIEW ALL	Shows nodes in the SBD partition with information about each node, such as node state and epoch number
LDNCS	Loads NCS on a cluster node
ULDNCS	Unloads NCS on a cluster node
CLUSTER DOWN	Removes all cluster nodes from the cluster It has the same effect as executing the CLUSTER LEAVE command on every server in the cluster.
CLUSTER JOIN	Adds the node where the command is executed to the cluster and makes the node visible to other servers in the cluster NCS must already be installed on a node for it to join the cluster.
CLUSTER LEAVE	Removes the node where the command is executed from the cluster The node is not visible to other servers in the cluster.

**30 minutes****Exercise 9-2 Test the SBD Partition and Heartbeats**

After installing NCS, you decide to run some tests to check the health and reliability of the cluster.

In this exercise you test the SBD partition and heartbeat monitoring on the 2-node SCSI clustering system you have configured by observing a node abend.

The following lists information you might need when testing the partition:

Table 9-6

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.DEL.DIGITALAIR	192.168.1.14
Cluster nodes	
DA4.IS.DEL.DIGITALAIR (<i>first server</i>)	192.168.1.4
DA5.IS.LGA.DIGITALAIR (<i>second server</i>)	192.168.1.5
SCSI hard drive ID	
SCSI Hard Drive ID on DA4:	
SCSI Hard Drive ID on DA5:	
Cluster name (DN)	
DACluster.IS.LON.DIGITALAIR	192.168.1.16
Cluster nodes	
DA6.IS.LON.DIGITALAIR (<i>first server</i>)	192.168.1.6
DA7.IS.SYD.DIGITALAIR (<i>second server</i>)	192.168.1.7
SCSI hard drive ID	
SCSI Hard Drive ID on DA6:	
SCSI Hard Drive ID on DA7:	
Cluster name (DN)	
DACluster.IS.TXL.DIGITALAIR	192.168.1.18
Cluster nodes	
DA8.IS.TXL.DIGITALAIR (<i>first server</i>)	192.168.1.8
DA9.IS.TYO.DIGITALAIR (<i>second server</i>)	192.168.1.9
SCSI hard drive ID	
SCSI Hard Drive ID on DA8:	
SCSI Hard Drive ID on DA9:	

To prepare for the test, do the following:

1. At the console prompt of both servers, enter **LIST DEVICES**.

2. Record the SCSI hard drive device ID number (such as 0x1) in Table 9-6 and note that the hard drive ID might be different on each node.

To run the test, do the following:

1. Verify that an SBD clustering partition exists on the SCSI hard drive:
 - a. From one (or both) workstation desktops start ConsoleOne.
 - b. In the left pane, right-click *your first server* object; then select **Properties**.
 - c. Select **Media > Devices**.

A list of device IDs (such as 0x1) appears at the left for the server you are logged into.

When you select a device ID in the list, information about the device appears at the right, including a description and partitioning information.
 - d. Select the device ID for the shared SCSI hard drive.

This is the device ID you listed in Table 9-6 at the beginning of the exercise.

Notice the following at the bottom of the information page:

 - A shared clustering partition (such as “Clustering - P:0xf”) appears. This is the SBD partition created during NCS installation.
 - The Sharable for Clustering option is selected. This indicates that the NCS cluster can use the SCSI hard drive for storing and sharing cluster data.
 - e. View additional partition information by selecting **Show Partition**.

Notice that the total space and used space are about 16 MB (depending on the size of your SCSI hard drive).
 - f. Close the Properties dialog by selecting **Cancel**.

2. Check the node information in the SBD partition on the SCSI hard drive by entering **SBD VIEW ALL** from one (or both) console prompts.

Both nodes are alive and have an epoch number of 1.
3. Prepare to view the results of split brain processing:
 - a. Switch to the logger screen at each node by pressing **Ctrl + Esc** and entering **2** for Logger Screen.
 - b. View the Cluster State view in ConsoleOne on one (or both) workstations by selecting the **DACluster** object in the left pane; then select **View > Cluster State View**.

Notice that both server icons in the Cluster State view are green (in a normal operating condition).

One node in your cluster is serving as the master node (indicated by a yellow ball). Because it is the master node, the master IP address resource is running on it.
4. View and record the effects of removing the Ethernet cable from *your first server*:
 - a. On *your first server*, remove the Ethernet cable from the network board.
 - b. Check the logger screens.

It can take 30 seconds or longer for the split brain process to complete and for messages to show in the logger screens.
 - c. Check the Cluster State view in ConsoleOne.

Notice that the server icon for your first server changes to red (with a break through it) or gray to indicate that the node has failed.

If the master IP address resource was running on your first server, it is now running on your second server with a yellow dot indicating that your second server is now the master node.

The epoch number updates to epoch 2 to indicate that your first server has left the cluster.



If your second server is cast off (abends) instead of your first server, NCS has had a problem reading the statistics on the network board drivers and can't determine which node is still communicating over the LAN.

Because there are only 2 nodes with 1 vote each, NCS breaks the tie by keeping the master node (your first server) alive and casting off the slave node (your second server).

- d. Read the cluster warning messages at the bottom of the logger screen on *your second server* monitor.

The messages indicate that the cluster lost communication with your first server and a poison pill was processed by the server to ensure cluster stability.

The message at the bottom of the screen indicates that the secondary IP address for the cluster (such as 192.168.1.14) has been added to your second server, making it the master node (if it was not already the master).

- e. Record the abend message (in blue) on *your first server* monitor in the following:

Node 1 (your first server) Abend Message

- f. Switch to the Cluster Membership Monitor (**CMON Screen**) on *your second server*.

Notice that the status for node 1 (*your first server*) is FAIL.

- g. Reconnect the Ethernet cable for *your first server*.

Table 9-7

- h. On *your first server* update the ABEND.LOG file and exit the Abend screen by entering **X**.
- 5. Restart NetWare 6 on *your first server* by entering **SERVER** at the DOS prompt.
- 6. After NetWare 6 loads, check the Cluster State view in ConsoleOne.

Notice that the server icon for your first server changes to green, indicating that the server is operating normally. The epoch number has updated again (to epoch 3), indicating your first server has joined the cluster.

- 7. Check the Cluster Membership Monitor on *your second server*. Notice that the status for node 1 is UP.
- 8. At *your first server* console prompt, enter **TIME** and record the date and time (your local time) in the following:

Table 9-8

Date	Time

- 9. At *your first server* console prompt enter **EDIT ABEND.LOG**.
- 10. Find the entry that is closest to the date and time you recorded.
- 11. Compare the abend message in the log entry to the abend message you recorded for *your first server* earlier in the exercise (step 3e).
The messages are the same.
- 12. Exit the Abend log by pressing **Esc** twice and selecting **Yes**.
- 13. Switch to the logger screen on *your first server*.
- 14. View the effects of turning off a server in the cluster:
 - a. Turn off *your second server*.
 - b. Check the cluster warning messages in the logger screen on *your first server*.

To ensure cluster stability, NCS created a new cluster view that only included your first server, and sent a token (poison pill) to the SBD sector for your second server.

A message indicates that the secondary IP address for the cluster (such as 192.168.1.14) has been added to your first server, making it the master node (if it was not already).

- c. Check the Cluster State view in ConsoleOne.

Notice that the server icon for your second server changes to red (with a break) or gray to indicate the node has failed.

Also notice that the master IP address resource is now running on your first server. A yellow dot indicates that your first server is the master node.

When a node fails in the cluster, NCS immediately migrates the cluster IP address resource to make sure the cluster is still available on the network.

It doesn't matter if your first or second server is the master node, as long as the cluster IP address is available on the network.

15. Switch to the Cluster Membership Monitor (CMON Screen) on *your first server*.

16. Turn on *your second server* and wait for NetWare 6 to load.

The status in the Cluster Membership Monitor on your first server for your second server changes from FAIL to UP.

17. As the server comes up, notice the change in color and state of the icon in the Cluster State view for *your second server*.

The icon changes from red or gray to green, indicating that the node has joined the cluster and is operating normally (sending heartbeats over the LAN and writing to the SBD partition on the SCSI hard drive).

18. On the workstations, close ConsoleOne by selecting **File > Exit**.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	What You Learned
1. Verify NCS System Requirements	<p>Before installing NCS, your system must meet the following requirements:</p> <ul style="list-style-type: none">■ Hardware Requirements Minimum hardware for installing NCS 1.6 on a 2-node NCS cluster include 2 NetWare 6 servers, 256 MB of memory each server (512 MB recommended), and 1 local disk device (not shared) for SYS on each server.■ Software Requirements Software includes all nodes in the same eDirectory tree and running NetWare 6, all servers configure for IP and on the same IP subnet, and an IP address for the cluster and each resource or volume.■ License Requirements NCS requires a cluster server license for each server that is part of the cluster. The license allows a server to join a cluster.■ Shared Disk System Requirements:<ul style="list-style-type: none">■ At least 15 MB of free disk space available on the shared disk system for creating the SBD partition■ The shared disk system set up and functioning according to the manufacturer's instructions■ The disks in the disk system configured in a mirroring or RAID 5 configuration to add fault tolerance to the system

Objective	What You Learned
2. Create a Cluster by Installing NCS	<p>You must run the NCS installation when you do the following:</p> <ul style="list-style-type: none">■ Create a cluster■ Add nodes to an existing cluster■ Upgrade NCS software in an existing cluster <p>The installation does the following:</p> <ul style="list-style-type: none">■ Creates a cluster object in eDirectory■ Installs NCS software on servers you specify for your cluster <p>After running the installation the first time to create a cluster, run the installation again to add servers to your cluster or to upgrade NCS software on a cluster.</p>

Objective	What You Learned
3. Check Cluster Configuration Settings	<p>When you install NCS, the following cluster objects are created in eDirectory:</p> <ul style="list-style-type: none">■ Cluster ADMIN Object<p>An ADMIN object (such as ADMIN_DACLuster) is created for the cluster. This object lets you set attributes, trustees, and rights for cluster objects and operations. You can also show creation facts about objects in the cluster.</p>■ Cluster Object<p>The cluster object contains the following objects for configuring and running the cluster:</p><ul style="list-style-type: none">■ Master IP address resource■ Cluster server nodes■ Resource templates<p>The following are basic properties of an NCS cluster object:</p><ul style="list-style-type: none">■ Timeout and Quorum Membership■ Cluster Protocol■ Cluster IP Address and Port■ Resource Priority■ Cluster Email Notification■ Master IP Address Resource Object<p>You can view information such as loading and unloading scripts for the cluster; start, failover, and failback node settings; and nodes associated with the cluster by showing the Master_IP_Address_Resource object properties.</p>■ Cluster Server Node Objects<p>You can view or edit the cluster server node number or IP address of the selected node. You can also view the context for the NetWare Server object.</p> <p>You can view and edit cluster object settings in ConsoleOne or NetWare Remote Manager.</p>

Objective	What You Learned
4. Test and Monitor the Cluster	<p>To test and monitor the cluster state, you need to learn about the following:</p> <ul style="list-style-type: none"> <li data-bbox="954 359 1481 590">■ Cluster State and Cluster Status Views ConsoleOne and NetWare Remote Manager provide a special status screen for cluster objects that you can use to perform tasks such as view the cluster state, migrate resources, check an events log, and print or save an HTML report (in ConsoleOne only) on the cluster state. <li data-bbox="954 604 1481 791">■ Console Prompt Commands To view a list of commands, enter HELP CLUSTER or HELP SBD at a server console prompt, or enter HELP and a command (such as HELP CLUSTER VIEW) to see a description and example.

Exercise Answers

Following are the exercise answers.

Exercise 9-1. Install and Check NCS on a 2-Node Cluster

Part III: Check the Cluster Configuration Settings in ConsoleOne

4. Record the cluster configuration settings:

Table 9-9

Properties Tab	Settings	Values
Quorum	Timeout	60
	Membership	2

Table 9-9 (continued)

Properties Tab	Settings	Values
Protocol	Heartbeat	1
	Tolerance	8
	Master WatchDog	1
	Slave WatchDog	8
	Max Retransmits	30
Management	IP Address	192.168.1.3x
	Port	7023

7. Record the node settings:

Table 9-10

Properties Tab	Settings	Values
Node	Node Number	0
	IP Address	192.168.1.x
	NCP Server	DAx.IS.xxx.DigitalAir

10. Record the node settings:

Table 9-11

Properties Tab	Settings	Values
Node	Node Number	1
	IP Address	192.168.1.x
	NCP Server	DAx.IS.xxx.DigitalAir

SECTION 10 Configure and Test High Availability File Access

Duration: 1 hour 30 minutes

In this section you learn how to cluster-enable and test high availability access of data using NCS on a NetWare network.

Objectives

1. [Configure NCS for High Availability File Access](#)
2. [Manage Resources in an NCS Cluster](#)

Introduction

After creating an NCS cluster, you need to create and configure cluster resources to make them available to customers and employees.

One type of resource is a shared cluster volume on a SAN. By combining the logical volume and storage pool features of NSS with NCS, you can provide a totally scalable, always-available data-access solution.

Objective 1 **Configure NCS for High Availability File Access**

When configuring NCS for making data and files highly available to users, you perform the following tasks:

- [Create a Shared Disk Partition](#)
- [Create and Cluster-Enable an NSS Volume and Pool on a Shared Storage Device](#)
- [Cluster-Enable an Existing Pool or Volume on the Shared Disk System](#)

Create a Shared Disk Partition

If students follow while you demonstrate creating a partition, pool, and volume on the shared SCSI hard drive, make sure they remove them before starting Exercise 10-1 or they might not have enough room on the shared drive to complete the exercises.

Before creating disk partitions on shared storage devices in your SAN, you must install NCS. You should carefully plan how you want to configure your shared storage prior to installing NCS.

To create a shared disk partition on a SAN, do the following:

1. Start ConsoleOne and authenticate to the eDirectory tree where the cluster object resides.

ConsoleOne runs faster from a client than a server. From a NetWare client, access `SYS:PUBLIC\MGMT\CONSOLEONE\1.2\BIN` on a server in the cluster and run `CONSOLEONE.EXE`.

We recommend accessing ConsoleOne on the server that is your primary NetWare connection.

2. In the left pane, right-click the cluster object or the server object of a server in the cluster; then click **Properties**.
3. On the Media tab, select **Devices**; then select the device where you want to create the shared partition.
4. Make sure the **Sharable for Clustering** box is selected for the device.

Although the Sharable for Clustering option lets students configure local disk devices as shared, they should avoid doing this to keep the SAN isolated from local drives on server nodes.

The only time students need to select the Sharable for Clustering option is if NetWare does not detect a device as shared storage on the SAN.

When you add a device to the SAN, NetWare 6 detects that the device is shared storage and identifies it as sharable for clustering.

If NetWare does not detect a device as shared storage on the SAN, you need to select this option.

Device names are not changeable and might be labeled something like 0x2 or 0x1.

5. On the Media tab, select **Partitions**; then select **New**.
6. Select the device where you want to create the partition (the same device you selected in step 3).
7. Specify the size of the partition and make sure of the following:
 - **NSS** is selected as the partition type (the default)
 - **Hot Fix** and **Mirror** are selected
 - **Create New Mirror Group** is selected
8. Create the partition by selecting **OK**.

Create and Cluster-Enable an NSS Volume and Pool on a Shared Storage Device

Although storage pools must be created prior to creating volumes, you can create and cluster-enable an NSS volume and pool at the same time by using the Create a New Logical Volume option on the Media tab of the Server properties dialog.

To create and cluster-enable a volume and pool from ConsoleOne, do the following:

1. Right-click the *server object* for a node in the cluster; then select **Properties**.
2. Select **Media > NSS Logical Volumes > New**.

The Create a New Logical Volume dialog appears.

3. In the Name box, enter a name for the volume.

Each volume in the pool must have a unique name.

4. Continue by selecting **Next**.

The Storage Information dialog appears.

5. Select an NSS partition on the shared storage device for the new pool and volume you want to create.
6. Enter a quota for the volume, or select the box to allow the volume to grow to the pool size.

The quota is the maximum possible size of the volume. If you have more than one volume per pool, you might want to enter a quota for each volume rather than allowing multiple volumes to grow to the pool size.

7. Continue by selecting **Next**.

The Create a New Pool dialog appears.

8. In the Name box, enter a name for the pool; then select **OK**.

Because the partition you selected is on a shared storage device, the Create a New Pool Cluster Info dialog appears with the following options:

- Cluster Enable on Creation.** This option is selected by default and causes the pool to be activated and cluster-enabled when it is created.
- Virtual Server Name.** When you cluster-enable a pool, a virtual server object is created and given the name of the cluster object plus the cluster-enabled pool.

For example, if the cluster name is *cluster1* and the cluster-enabled pool name is *pool1*, the default virtual server name is *cluster1_pool1_server*.

The virtual server object contains the IP address for the NSS pool and is associated with the pool and volume.



If you are cluster-enabling a volume in a pool that is cluster-enabled, the virtual server object has been created, and you can't change the virtual server object name.

- **IP Address.** Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access, migration, and failover capability to the cluster-enabled pool.

The IP address remains assigned to the pool regardless of which server in the cluster is accessing the pool.

- **Advertising Protocols.** You can select NCP, CIFS, and AFP as the advertising protocols over the network for the NSS pool IP address.

NCP is the protocol used by Novell clients; CIFS is the protocol used by Microsoft clients; and AFP is the protocol used by Macintosh clients.

Selecting any of the protocols causes lines to be added to the pool resource load and unload scripts to activate the selected protocols on the cluster.

This lets you ensure that the cluster-enabled pool you create is highly available to all your clients.

If you select the CIFS box, the CIFS Server Name field becomes active. The CIFS server name is the server name CIFS clients see when they browse the network.

A default CIFS server name is listed, but you can change the name by editing the text in the field.

9. Enter an **IP address** for the NSS pool; then make any other changes to the configuration information.
10. When you finish, continue by selecting **OK**.
The Attribute Information dialog appears.
11. Review and change the attributes as necessary.

The Flush Files Immediately option flushes files in the volume from the cache as soon as the file is closed to improve file system reliability. However, the option can hamper network performance.

Select this option if you want to ensure the integrity of volume data.

12. Create the volume and pool by selecting **Finish.**

eDirectory objects for the volume, pool, and virtual server are created in the same container as the NetWare server. In addition, a cluster resource object for the NSS pool is created in the cluster object.

For example, if you create a cluster-enabled USERS pool with a MEDIA volume for DAACLUSTER, the following objects are created:

Table 10-1

Object	Description
DAACLUSTER_MEDIA	The cluster-enabled volume object
DAACLUSTER_USERS_POOL	The cluster-enabled pool object
DAACLUSTER_USERS_SERVER	The virtual server object for the cluster-enabled pool
DAACLUSTER_USERS	The resource object for the NCS cluster

Cluster-Enable an Existing Pool or Volume on the Shared Disk System

In addition to creating a cluster-enabled NSS volume and pool, you can also cluster-enable an existing NSS volume or pool on the shared disk system.

To cluster-enable an existing NSS volume and pool, you need to know the following:

- [What Happens to eDirectory Objects during Cluster-Enabling](#)
- [How to Cluster-Enable an Existing Volume and Pool](#)
- [Guidelines for Cluster-Enabling the Pool and Volume](#)

What Happens to eDirectory Objects during Cluster-Enabling

When you cluster-enable an existing NSS volume and pool, you associate the volume and pool with a new virtual server object with its own IP address. This enables the volume to be accessible even if the server fails.

During the cluster-enabling process, the old volume object is replaced with a new volume object that is associated with the pool, and the old pool object is replaced with a new pool object associated with the new virtual server object.

In addition, a volume resource object is created and is listed under Resources in the Cluster State or Cluster Status view.

How to Cluster-Enable an Existing Volume and Pool

To cluster-enable an existing volume (and pool) using ConsoleOne, do the following:

1. Browse and select the cluster object.
2. Select **File > New > Cluster > Cluster Volume**.
3. Browse and select a volume on the shared disk system to cluster-enable.
4. Enter an IP address for the volume.

This is only required for the first volume to be cluster-enabled in the pool. This IP address is assigned to the pool where the volume resides.

Selecting **Online Resource** after **Create** causes the volume to mount when it is created.

Selecting **Verify IP Address** instructs NCS to verify there are no IP address conflicts.

5. (Optional) Change the default name of the virtual server object.

When you cluster-enable a pool, a virtual server object is created and given the name of the cluster object plus the cluster-enabled pool.

For example, if the cluster name is *cluster1* and the cluster-enabled pool name is *pool1*, the default virtual server name is *cluster1_pool1_server*.

If you are cluster-enabling a volume in a pool that is cluster-enabled, the virtual server object has been created, and you can't change the virtual server object name.

6. (Optional) Change the default name of the cluster-enabled volume object.

When you cluster-enable a volume, a new object is created and given the name of the cluster object and the volume name.

For example, if the cluster name is *cluster1* and the volume name is *vol1*, then the default cluster-enabled volume object name is *cluster1_vol1*.

7. Ensure **Define Additional Properties** is selected; then select **Create**.

You can also cluster-enable a volume using NetWare Remote Manager.

The first 2 guidelines are targeted at students who are already familiar with previous versions of NCS and need to know what's new in NCS 1.6.

Guidelines for Cluster-Enabling the Pool and Volume

The following are guidelines for cluster-enabling a pool or volume in NCS 1.6:

- Cluster-enabled volumes no longer appear as cluster resources. The load and unload scripts in cluster resource objects apply to pools (not volumes).
- Each cluster-enabled NSS pool requires its own IP address for the virtual server. This means that each cluster-enabled volume does not have an associated load and unload script or an assigned IP address.
- The first volume you cluster-enable in the pool cluster-enables the pool where the volume resides.

After a pool is cluster-enabled, you must cluster-enable the other volumes in the pool if you want them to be mounted on another node during a failover.

- When a node fails, any cluster-enabled pools being accessed by that node are migrated to other nodes in the cluster.

All volumes in the pool are migrated with the pool, but only volumes that have been cluster-enabled are mounted. Any volumes in the pool that are not cluster-enabled must be mounted manually.

For this reason, volumes that aren't cluster-enabled should be in separate pools that are not cluster-enabled.

- If you want each cluster-enabled volume to be its own cluster resource, each volume must have its own pool.
- If a server application does not require NetWare client access to volumes, cluster-enabling those pools and volumes might not be necessary.
- Pools should be deactivated and volumes should be dismounted before being cluster-enabled.

Exercise 10-1 Create a Cluster-Enabled Volume for High Availability File Access



20 minutes

Although it is easier to create and cluster-enable an NSS pool and volume using NetWare Remote Manager, this exercise focuses on using ConsoleOne to perform these tasks.

NetWare Remote Manager is more task-oriented and less focused on Directory objects. By using ConsoleOne, students can observe and explore more conveniently the eDirectory objects that support NCS clustering.

When you initialized the SCSI hard drive, you removed all existing partitions. Installing NCS created an SBD partition on the hard drive to monitor clustering.

At this point, you could use the rest of the hard drive as a partition for an NSS pool. However, all data and services on the hard drive will migrate or failover to the same node if they are stored in the same partition and pool.

You decide to create at least 2 partitions on the SCSI hard drive to test migration and failover of part of the data on the hard drive from one node to the other.

In this exercise, you create a partition that uses half of the available storage space on the SCSI hard drive, and then create a MULTIMEDIA pool and VIDEO volume in that partition.

The following lists information you need during the exercise:

Table 10-2

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.DEL.DIGITALAIR	192.168.1.14
Cluster nodes	
DA4.IS.DEL.DIGITALAIR (<i>first server</i>)	192.168.1.4
DA5.IS.LGA.DIGITALAIR (<i>second server</i>)	192.168.1.5
SCSI hard drive ID	
SCSI Hard Drive ID on DA4:	
SCSI Hard Drive ID on DA5:	
SCSI hard drive partition	
Name: MediaResources	
Partition ID:	
Mirror ID:	
NSS information	
Pool Name: MULTIMEDIA	192.168.1.24
Volume Name: VIDEO	
Cluster name (DN)	
DACluster.IS.LON.DIGITALAIR	192.168.1.16
Cluster nodes	
DA6.IS.LON.DIGITALAIR (<i>first server</i>)	192.168.1.6
DA7.IS.SYD.DIGITALAIR (<i>second server</i>)	192.168.1.7
SCSI hard drive ID	
SCSI Hard Drive ID on DA6:	
SCSI Hard Drive ID on DA7:	
SCSI hard drive partition	
Name: MediaResources	
Partition ID:	
Mirror ID:	
NSS information	
Pool Name: MULTIMEDIA	192.168.1.26
Volume Name: VIDEO	

Table 10-2 (continued)

Cluster Information	IP Addresses
Cluster name (DN)	
DAcluster.IS.TXL.DIGITALAIR	192.168.1.18
Cluster nodes	
DA8.IS.TXL.DIGITALAIR (<i>first server</i>)	192.168.1.8
DA9.IS.TYO.DIGITALAIR (<i>second server</i>)	192.168.1.9
SCSI hard drive ID	
SCSI Hard Drive ID on DA8:	
SCSI Hard Drive ID on DA9:	
SCSI hard drive partition	
Name: MediaResources	
Partition ID:	
Mirror ID:	
NSS information	
Pool Name: MULTIMEDIA	192.168.1.28
Volume Name: VIDEO	

To prepare for the exercise, do the following:

1. At the console prompt of each server, enter **LIST DEVICES**.
2. In the Component Name/ID column of Table 10-2 for your cluster, record the SCSI hard drive device ID number.

The device ID might be different on each server.

To create and cluster-enable the partition, pool, and volume, do the following:

- [Part I: Create a Shared Partition on the SCSI Hard Drive](#)
- [Part II: Create a Cluster-Enabled VIDEO NSS Volume on the MediaResources Partition](#)
- [Part III: Verify That the Cluster Objects Are Created for the Volume](#)

Part I: Create a Shared Partition on the SCSI Hard Drive

Before creating the MULTIMEDIA pool on the SCSI hard drive, create a partition by doing the following:

1. Make sure that
 - You are logged in as **Admin** to *your first server* from a workstation
 - ConsoleOne for *your first server* is running on the workstation
2. From ConsoleOne on the workstation, right-click *your first server* object in **IS.xxx.DIGITALAIR** (where *xxx* = your location container); then select **Properties**.

The Properties dialog for the server appears.
3. Select **Media > Partitions > New**.
4. Select the device ID for your shared SCSI hard drive (such as 0x6-1).

The name of the SCSI hard drive and the settings for the new partition appear at the right.
5. For the partition type select **NSS**.
6. Set the partition size to one-half the size listed in the Size box.

For example, if the available space for the selected device is 2 GB, enter 1 GB in the Size box.
7. Make sure the **Hot Fix** and **Mirror** options are selected.
8. In the Label box, name the partition by entering **MediaResources**.
9. Create the partition by selecting **OK**.

The list of partitions for the node appears.
10. Select the new partition in the list.

The partition is an NSS partition with an ID such as 0x11. Look for the MediaResources label in the right panel.

11. In Table 10-2, record the partition ID number and the mirror ID number (also in the right panel) for your cluster.

Notice that the partition is sharable for clustering and that there is no used space in the partition.

Part II: Create a Cluster-Enabled VIDEO NSS Volume on the MediaResources Partition

To create a VIDEO volume in a MULTIMEDIA pool on the MediaResources partition, do the following:

1. From *your first server* properties dialog, select **Media > NSS Logical Volumes > New**.

The Create a New Logical Volume dialog appears.

2. In the Name box, name the volume by entering **VIDEO**; then select **Next**.

The VIDEO - Storage Information dialog appears.

3. Select the NSS partition you created in Part I for MediaResources (see Table 10-2 for the ID number).
4. Select **Allow volume quota to grow to the pool size**.
5. Continue by selecting **Next**.

The Create a New Pool dialog appears.

6. In the Name box, name the pool by entering **MULTIMEDIA**; then select **OK**.

Because the partition you selected is on a shared storage device, the Create a New Pool Cluster Info dialog appears.

The following parameters are already selected for you:

- Cluster Enable on Creation
- Virtual Server Name:
DACluster_MULTIMEDIA_SERVER

To cluster-enable the pool and volume, all you need to do is enter an IP address for the pool (virtual server).

7. For the pool IP address, enter *your pool IP address* (see Table 10-2).
8. Continue by selecting **OK**.

The VIDEO - Attribute Information dialog appears.

9. Make sure the following are the only attribute settings selected:
 - Backup
 - Salvage Files
 - On Creation:
Activate
Mount

10. Create and cluster-enable the volume and pool by selecting **Finish**.

The volume VIDEO appears in the list of NSS Logical Volumes on *your first server*.

11. Select the **VIDEO** logical volume and read the volume information in the right panel.

Notice that the host pool is MULTIMEDIA, the state is active and mounted, and the volume can grow to the pool size.

12. Make sure the volume is mounted on *your first server*:
 - a. At *your first server* console prompt, enter **VOLUMES**.

The volume VIDEO is listed for your first server.

Because you chose to bring the resource online while cluster-enabling the pool and volume, NCS used the load script to activate the pool and mount the volume.

- b. From *your second server* console prompt, enter **VOLUMES**.

Volume VIDEO is not listed. Although the volume is on the shared SCSI hard drive, it is assigned as a volume resource to your first server.

Part III: Verify That the Cluster Objects Are Created for the Volume

To verify that the cluster-enabled objects have been created for the MULTIMEDIA pool and volume VIDEO, do the following:

1. From one or both workstations, make sure you are logged in as **Admin** to one of the servers; then start **ConsoleOne**.
2. Find the following in **IS.xxx.DigitalAir** for *your first server*:
 - **DACluster_VIDEO** (cluster-enabled volume)
 - **DACLUSTER_MULTIMEDIA_SERVER** (virtual server)
 - **DACLUSTER_MULTIMEDIA_POOL** (cluster-enabled pool)
3. Right-click **DACluster_VIDEO**; then select **Properties**.
4. Select **General > Identification**.
5. Verify the name of the host server on the Identification tab page.
You see the **DACLUSTER_MULTIMEDIA_SERVER** virtual server listed.
6. Close the Properties dialog by selecting **Cancel**.
7. Right-click **DACLUSTER_MULTIMEDIA_POOL**; then select **Properties**.
8. Select **Other**; then expand the **Host Server** attribute.
9. Verify the name of the host server.

The pool object is associated with the same virtual server.

10. Close the Properties dialog by selecting **Cancel**.
11. Right-click **DACLUSTER_MULTIMEDIA_SERVER**; then select **Properties**.
12. Select **General > Identification**.

In the Network address field you see the IP address you entered when cluster-enabling the volume (such as 192.168.1.24).

With the pool and volume associated with a virtual server, VIDEO is always available in the cluster because the virtual server IP address continues to be broadcast as the resource migrates from node to node.
13. Close the Properties dialog by selecting **Cancel**.
14. Verify that the cluster volume resource object is created:
 - a. Make sure you are in **Console View** (ConsoleOne).
 - b. Select the **DACluster** object.
 - c. In the right window, right-click **MULTIMEDIA_SERVER**; then select **Properties**.
 - d. Select **IP Address > Cluster Resource IP Address**.
 - e. Verify the listed IP address.

You see the IP address you entered when cluster-enabling the volume (such as 192.168.1.24).
 - f. Select **Nodes > Cluster Resource Preferred Nodes**.
 - g. Verify the nodes assigned to the volume resource.

In the Assigned list you see your first and second servers. Because your first server is listed first, MULTIMEDIA_SERVER is assigned to your first server when the cluster first starts.

If your first server goes down, NCS migrates VIDEO to the next node in the list (your second server).
 - h. Select **Scripts > Cluster Resource Load Script**.

The first 2 lines of the script activate MULTIMEDIA and mount VIDEO.

The remaining lines ensure that the IP address for your pool (such as 192.168.1.24) is bound to the server and that the server is broadcasting.

- i. Close the Properties dialog by selecting **Cancel**.

(End of Exercise)

Objective 2 Manage Resources in an NCS Cluster

To manage resources in an NCS cluster, you need to know the following:

- [How to Migrate Resources](#)
- [How to Troubleshoot Resource States](#)

How to Migrate Resources

You can migrate resources to different nodes in your cluster without waiting for a failure to occur.

You might want to migrate resources to do the following:

- Lessen the load on a specific server
- Free a server so it can be brought down for scheduled maintenance
- Increase the performance of the resource by putting it on a faster machine

Migrating resources lets you balance the load and evenly distribute applications among the servers in your cluster.

You can migrate resources from the status view in ConsoleOne or NetWare Remote Manager.

To migrate resources using ConsoleOne, do the following:

1. Select the cluster object that contains the resource you want to migrate.

Resources must be in a running state to be migrated.

2. Make sure the right half of ConsoleOne shows the Cluster View State by selecting **View > Cluster State View** from the menu at the top of the screen.
3. In the Cluster Resource List, select the resource you want to migrate.

The Cluster Resource Manager screen appears, showing the server that the selected resource is running on and a list of possible servers you can migrate resources to.

4. Select a server from the list; then select **Migrate** to move the resource to the selected server.

To migrate resources using NetWare Remote Manager, do the following:

1. Open a web browser (such as Internet Explorer) from the Windows desktop of your workstation.
2. Enter **https://server IP address:8009/**.

For example, if your server IP address is 192.168.1.1, enter **https://192.168.1.1:8009/**.

One or more security alert dialogs appear.

3. Continue by selecting **Yes** or **OK**.
4. Log in to the server as the network administrator.
NetWare Remote Manager appears.
5. In the left panel under the **Clustering** heading (near the bottom), select **Cluster Management**.

A Cluster Status view similar to the Cluster State view in ConsoleOne appears.

6. Select a resource from the Resource list.
7. Select a node; then select **Migrate**.
8. Select a refresh rate (such as 2 seconds); then select **Begin Refresh**.

If you select a resource and click Offline, the resource is unloaded from the current node. It does not load on any other nodes in the cluster and remains unloaded until you load it again.

This is useful when editing resources because resources can't be edited while loaded or running on a node.

How to Troubleshoot Resource States

This topic introduces the troubleshooting table. You do not need to cover every state, description, and possible action listed.

When running or testing an NCS cluster, you can view valuable information about cluster resource states from the Cluster State view in ConsoleOne or from the Cluster Status view in NetWare Remote Manager.

The first solution to most error messages listed (such as Comatose or NDS Sync) is to take the resource offline and then bring it online again.

If this does not work, take the resource offline, check the configuration settings (especially the load and unload scripts), and then bring the resource online again.

The following identifies different resource states you see in the Cluster State view (ConsoleOne) or the Cluster Status view (NetWare Remote Manager) and gives a description and possible actions for each state:

Table 10-3

State	Description	Possible Actions
Alert	The Start, Failover, or Failback mode for the resource has been set to Manual. The resource is waiting to start, failover, or failback on the specified server.	Select the Alert status indicator. Depending on the resource state, you are prompted to start, failover, or failback the resource.
Comatose	The resource is not running and requires administrator intervention.	Select the Comatose status indicator and take the resource offline. After resource problems are resolved, the resource can be put back online (returned to the Running state).
Unloading	The resource is unloading from the server it was running on.	None.
Running	The resource is in a normal running state.	Select the Running status indicator and choose to either migrate the resource to another server in your cluster, or unload (take offline) the resource.
Loading	The resource is loading on a server.	None.

Table 10-3 (continued)

State	Description	Possible Actions
Unassigned	There isn't an assigned node up that the resource can be loaded on.	Select the Unassigned status indicator and, if desired, take the resource offline. Taking the resource offline prevents it from running on any of its preferred nodes if any of them join the cluster.
NDS_Sync	The properties of the resource have changed and the changes are still being synchronized with eDirectory.	None.
Offline	The resource is shut down or is in a dormant or inactive state.	Select the Offline status indicator and, if desired, click the Online button to load the resource on the best node possible, given the current state of the cluster and the resource's preferred nodes list.
Quorum Wait	The resource is waiting for a quorum to be established so it can begin loading.	None.

**15 minutes**

During classroom setup, you should have created a C:\MARKETING VIDEO folder on each student workstation with a video file that runs 1–2 minutes.

This folder is also available in EXERCISES\SECTION 10 on the Enhanced Learning CD.

Before starting the exercise, instruct students on how to start the video for Parts II and III.

Exercise 10-2 Test High Availability File Access on the 2-Node Cluster

Now that your cluster is running and configured with a cluster-enabled volume, you can test high availability file access.

In this exercise you copy a video file to volume VIDEO and then run the video from a Windows 2000 NetWare 6 workstation while migrating the MULTIMEDIA_SERVER resource and failing a node.

Specifically, you do the following:

- **Part I: Verify That Volume VIDEO Is Cluster-Enabled**
- **Part II: Test the Effects of Migrating Volume VIDEO while Playing a Video**
- **Part III: Test the Effects of Failing a Cluster Node while Playing a Video**

(Remember: if a program is not cluster-aware or cannot recover from file-level interruption, you can have problems when saving a file or playing a video during a migration, such as when migrating volume VIDEO during this exercise.)

Part I: Verify That Volume VIDEO Is Cluster-Enabled

You can verify that volume VIDEO is cluster-enabled by migrating the MULTIMEDIA_SERVER volume resource from one cluster node to the other.

To verify cluster-enabling, do the following:

1. From ConsoleOne on one of your workstations, make sure the Cluster State view appears for the **DACluster** object.
2. Check the information for the **MULTIMEDIA_SERVER** volume resource in the Cluster State view.

You might need to expand the width of the Cluster Resources column to see the entire MULTIMEDIA_SERVER name.

The MULTIMEDIA_SERVER volume resource is running on *your first server*.

3. Verify cluster-enabling by migrating **MULTIMEDIA_SERVER** from *your first server* to *your second server*:
 - a. In the Cluster Resource list select **MULTIMEDIA_SERVER**.
The Cluster Resource Manager dialog appears.
 - b. Select *your second server*; then select **Migrate**.
 - c. In the Cluster State view, verify what happens to the MULTIMEDIA_SERVER volume resource.
The resource is unloaded from your first server, loaded on your second server, and starts running on your second server.
4. Verify the volumes mounted on *your first server* and *your second server* by entering **VOLUMES** at the console prompt of each node.
After migrating MULTIMEDIA_SERVER to your second server, volume VIDEO is mounted on that server.
5. Migrate **MULTIMEDIA_SERVER** from *your second server* to *your first server*.
6. Use the **VOLUMES** command at the console prompt of each node to verify that you have migrated VIDEO back to its original node assignment.
7. Close ConsoleOne on the workstation.

Part II: Test the Effects of Migrating Volume VIDEO while Playing a Video

In this part you test clustering by migrating the MULTIMEDIA_SERVER resource from ConsoleOne on one workstation (called *your first workstation*) while playing a video on the other workstation (called *your second workstation*).

For example, if you are using WS4 and WS5, you can use **WS4** as *your first workstation* and **WS5** as *your second workstation*.

To perform the test, do the following:

1. View the Cluster Status view in NetWare Remote Manager on *your first workstation* by doing the following:
 - a. Open Internet Explorer from the desktop; then enter **https://192.168.1.x:2200/** (where 198.168.1.x = your first server IP address).

For example, if your cluster nodes are DA4 and DA5, the IP address of your first server (DA4) is 192.168.1.4.
 - b. (Conditional) If a Security Alert dialog appears, continue by selecting **Yes** or **OK**.

NetWare Web Manager appears.
 - c. Under the **NetWare Remote Manager** heading, select *your first server*.
 - d. (Conditional) If a Security Alert dialog appears, continue by selecting **Yes** or **OK**.
 - e. Log in by entering **Admin.IS.xxx.DIGITALAIR** for the user name and **novell** for the password; then select **OK**.
 - f. In the left panel under the **Clustering** section (near the bottom), select **Cluster Management**.

A Cluster Status view similar to the Cluster State view in ConsoleOne appears.

However, you select **Begin Refresh** to see the cluster status update when migrating resources or when a node fails.

2. From *your second workstation*, copy the **MARKETING VIDEO** folder from C:\ to the cluster-enabled VIDEO volume on the SCSI hard drive:
 - a. Right-click **My Computer** on the desktop and select **Open**; then double-click (C:).
 - b. Right-click the **MARKETING VIDEO** folder; then select **Copy**.
(If the Marketing Video folder is not available on C:\, you can copy the same folder from EXERCISES\SECTION 10 on your Enhanced Learning CD.)
 - c. From the Address drop-down list select **My Network Places**; then double-click **Novell Connections > Digitalair-Tree > DIGITALAIR > xxx > IS**.
 - d. Right-click the **DACLUSTER_VIDEO** folder; then select **Novell Map Network Drive**.
 - e. Select an available driver letter; then select the following:
 - Check to make folder appear as the top most level**
 - Check to always map this drive letter when you start Windows.**
 - f. Select **Map**.
A DACLUSTER_VIDEO folder window appears.
 - g. Right-click in the empty folder window; then select **Paste**.
The file is copied to VIDEO on the SCSI hard drive.
 - h. When the copying is complete, close the **DACLUSTER_VIDEO** folder window.
3. Test high availability file data access by migrating the **MULTIMEDIA_SERVER** resource on *your first workstation* while playing the marketing video from *your second workstation*:
 - a. On *your second workstation*, right-click **My Computer**; then select **Open**.

-
- b. Start the video file from the **MARKETING VIDEO** folder by double-clicking *the drive* you mapped to the **DACLUSTER_VIDEO** folder; then double-click **Marketing Video > video file**.
 - c. From NetWare Remote Manager on *your first workstation*, migrate the **MULTIMEDIA_SERVER** resource from *your first server* to *your second server*.
 - d. From the Page Refresh Rate drop-down list select **2 seconds**; then select **Begin Refresh**.
 - e. Verify what happens to the video.

The video pauses during migration, but does not continue playing. This results from the share-level (volume-level) reconnect feature of Windows 2000.

If you were using Windows 98, the video would momentarily pause during the migration and then continue until it finished playing. This results from the file-level reconnect feature of Windows 98.

- f. From *your second workstation*, restart the video.
How you restart the video depends on the video player you are using. For most video players, you select a play button. Other video players might have a special restart option or button.



Depending on the video player you use, error messages might appear as the video player attempts to reconnect to the file. Because the video player is not cluster-aware, error messages might appear as the video player attempts to reconnect to the video file.

Try closing the video player, opening it again, and then start the video file. Because the migration is complete, you should be able to access and play the video again.

- g. When the video finishes, exit the video player.
- h. On *your first workstation* in NetWare Remote Manager, select **Stop Refresh**.

Part III: Test the Effects of Failing a Cluster Node while Playing a Video

Migrating services from one node to another in a cluster is a convenient way to upgrade servers while continuing to provide network services. However, when a node fails, you still want to provide the same level of high availability.

To test the effects of failover when a node fails, do the following:

1. Verify the following:
 - NetWare Remote Manager on *your first workstation* is maximized.
 - In the left panel of NetWare Remote Manager, Cluster Management is selected.
 - The Cluster Status view is set to refresh every 2 seconds.
2. From *your second workstation*, start the *video file* in the **MARKETING VIDEO** folder on the SCSI hard drive.

You can access the **MARKETING VIDEO** folder from Windows Explorer by using the drive you mapped earlier.
3. After the video starts playing, pull the network cable from *your second server*.
4. Verify what happens to the video.

The video pauses while NCS detects that your second server is down and loads the MULTIMEDIA_SERVER resource on your first server.
5. After the resource is running, restart the video.
6. On *your second server* at the abend screen, return to DOS by entering **X**; then reconnect the network cable and start NetWare 6 by entering **SERVER**.

7. When the video finishes and Netware 6 loads on *your second server*, exit the video player; then close the **Marketing Video** dialog and **NetWare Remote Manager**.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	What You Learned
1. Configure NCS for High Availability File Access	<p>When configuring NCS for making data and files highly available to users, you</p> <ul style="list-style-type: none">■ Create a Shared Disk Partition. Before creating disk partitions on shared storage devices in your SAN, you must install NCS. You should carefully plan how you want to configure your shared storage prior to installing NCS.■ Create and Cluster-Enable an NSS Volume and Pool on a Shared Storage Device. Although storage pools must be created prior to creating volumes, you can create and cluster-enable an NSS volume and pool by using the Create a New Logical Volume option on the Media tab of the server properties dialog.■ Cluster-Enable an Existing Pool or Volume on the Shared Disk System. To cluster-enable an existing NSS volume and pool, you need to know the following:<ul style="list-style-type: none">■ What Happens to eDirectory Objects during Cluster-Enabling■ How to Cluster-Enable an Existing Volume and Pool■ Guidelines for Cluster-Enabling the Pool and Volume

Objective	What You Learned
2. Manage Resources in an NCS Cluster	<p data-bbox="951 281 1435 344">Before you can manage resources in an NCS cluster, you need to learn the following:</p> <ul data-bbox="951 359 1479 583" style="list-style-type: none"><li data-bbox="951 359 1479 583">■ How to Migrate Resources. You might want to migrate resources to do the following:<ul data-bbox="987 422 1479 583" style="list-style-type: none"><li data-bbox="987 422 1393 453">■ Lessen the load on a specific server<li data-bbox="987 464 1479 516">■ Free a server so it can be brought down for scheduled maintenance<li data-bbox="987 527 1479 583">■ Increase the performance of the resource by putting it on a faster machine <p data-bbox="987 594 1479 678">Migrating resources lets you balance the load and evenly distribute applications among the servers in your cluster.</p> <ul data-bbox="951 688 1479 856" style="list-style-type: none"><li data-bbox="951 688 1479 856">■ How to Troubleshoot Resource States. When running or testing an NCS cluster, you can view valuable information about cluster resource states from the Cluster State view in ConsoleOne or from the Cluster Status view in NetWare Remote Manager. <p data-bbox="987 867 1330 898">The following are resource states:</p> <ul data-bbox="987 909 1166 1251" style="list-style-type: none"><li data-bbox="987 909 1081 940">■ Alert<li data-bbox="987 951 1133 982">■ Comatose<li data-bbox="987 993 1133 1024">■ Unloading<li data-bbox="987 1035 1117 1066">■ Running<li data-bbox="987 1077 1117 1108">■ Loading<li data-bbox="987 1119 1149 1150">■ Unassigned<li data-bbox="987 1161 1149 1192">■ NDS_Sync<li data-bbox="987 1203 1101 1234">■ Offline<li data-bbox="987 1245 1166 1276">■ Quorum Wait

SECTION 11 Configure and Test High Availability Services

Duration: 2 hours

If you are running out of time to complete this section, briefly introduce each objective but focus on cluster-enabling an application and accessing load and unload scripts.

Remember that the overall objective for this module is to prepare students for a successful experience implementing a 2-node cluster on their own.

As students install their own 2-node cluster at home, they can experiment with accessing features outside the scope of the exercises.

In this section, you learn how to cluster-enable and test service solutions such as DHCP and iFolder on your 2-node NCS cluster.

Objectives

1. [Identify Cluster-Aware and Cluster-Naive Applications](#)
2. [Identify How to Cluster-Enable an Application](#)
3. [Identify How to Assign Nodes to a Resource](#)
4. [Identify How to Set Start, Failover, and Failback Modes](#)
5. [Identify How to View and Edit Load and Unload Scripts](#)
6. [Identify How to Find NCS Configuration and Troubleshooting Information](#)

Introduction

Depending on your needs and design, additional configuration is required for you to effectively provide services using NCS.

This additional configuration normally requires cluster-enabling the application or service and possibly changing the properties for the cluster object and the cluster node objects.

Objective 1 Identify Cluster-Aware and Cluster-Naive Applications

When creating a resource for an NCS cluster you need to be familiar with the following types of applications:

- **Cluster aware.** Cluster-aware applications are programmed to take advantage of NCS clustering.

When cluster-enabled, these applications know they are running on an NCS cluster and try longer and harder to reconnect to the cluster.

GroupWise is an example of a cluster-aware application.

- **Cluster naive.** You can cluster-enable any application, but it might not be programmed to recognize that it is running on a cluster. This is a cluster-naive application.

For a cluster-naive application or service, NCS does all the work to ensure that the resource is reloaded on another node if the assigned cluster node fails.

An example of a cluster-naive application is the video player you used in Exercise 10-2.

The following are cluster-aware applications for NCS 1.6:

- Apache Web Server
- AppleTalk Filing Protocol (AFP)
- BorderManager® (Proxy and VPN)
- DHCP Server
- Enterprise Web Server (LDAP and NDS)
- GroupWise 5.5 and 6 (MTA, POA, GWIA, and WebAccess)
- iFolder
- iManager
- iPrint

- NetWare 5.1 FTP Server
- NFAP Common Internet File Services (CIFS)
- NFS 3.0
- NDPS
- Novell Clients (Windows 98 and Windows 2000)
- Oracle database
- Btrieve®
- Norton AntiVirus™
- WebDAV
- ZENworks for Servers
- ZENworks for Desktops 2 and 3

Objective 2 Identify How to Cluster-Enable an Application

You cluster-enable a service such as an application by creating a cluster resource. The resource includes a unique IP address and is available for migration from the Resource list in the Cluster State view (ConsoleOne) or the Cluster Status view (NetWare Remote Manager).

Cluster resources can be created for cluster-aware or cluster-naive applications such as web sites, email servers, databases, or any other server-based applications or services you want to make available to users at all times.

You can create a cluster resource using ConsoleOne or NetWare Remote Manager.

To create a cluster resource using ConsoleOne, do the following:

1. Browse and select the *cluster object* you want to create resources for.
2. Select **File > New > Cluster > Cluster Resource**.
3. Enter a *name* for the new cluster resource.
4. Choose one of the following:
 - If a template exists for the resource you are creating, in the Inherit from Template field enter the template name, or browse and select it from the list.
 - If a template does not exist, select **Define Additional Properties**.
5. (Conditional) If you want the resource to start on the master node as soon as it is created and configured, select **Online Resource after Create**.
6. Select **Create**.

To create a cluster resource using NetWare Remote Manager, do the following:

1. From the left column under the Clustering section, select **Cluster Config**.
2. At the bottom of the screen above Create New Objects, select **New Cluster Resource**.
3. Enter a name for the new cluster resource.
4. Choose from the following:
 - If a template exists for the resource you are creating, select it from the list.

Additional resource configuration is performed by the template.

The default selection for this field is No Template.

- If you are not using a template, create the cluster resource by configuring load and unload scripts, setting failover and failback modes, and if necessary, changing the node assignments for the resource.

5. When you finish, select **Apply**.



More information on cluster-enabling GroupWise, NDPS, NetWare Enterprise Web Server and many other applications is available at Novell's documentation web site, <http://www.novell.com/documentation/lg/ncs6p/index.html>.

Objective 3 Identify How to Assign Nodes to a Resource

You assign nodes to a resource for NCS to know which nodes to migrate the resource to during a failover.

When you create a resource on a cluster or when you cluster-enable a volume, the nodes in the cluster are assigned to the resource or volume. The order of assignment is the order the nodes appear in the resource list.

You can assign or unassign nodes to the resource or volume or change the failover order.

To assign or unassign nodes, or to change node assignments using ConsoleOne, do the following:

1. From the cluster object container, right-click the cluster-enabled resource or volume object; then select **Properties**.
2. Select **Node > Cluster Resource Preferred Nodes**.
3. From the list of unassigned nodes, select the server you want the resource assigned to; then click the **Right-arrow** button to move the selected server to the Assigned Nodes list.

Repeat this step for all servers you want assigned to the resource.

You can also use the Left-arrow button to unassign servers from the resource.

4. Click the Up- and Down-arrow buttons to change the failover order of the servers assigned to the resource or volume.

To assign or unassign nodes, or to change node assignments using NetWare Remote Manager, do the following:

1. In the left column under the Clustering section, select **Cluster Config**.
2. From the list of resources, select a resource.
3. On the Resource Information screen, select **Nodes**.
4. Select or enter the nodes you want assigned to this resource.
5. Save the node assignment changes by selecting **Apply**.

Objective 4 Identify How to Set Start, Failover, and Failback Modes

You can configure the start, failover, and failback of cluster resources to happen manually or automatically:

- **Start mode (AUTO).** With the resource Start mode set to AUTO, the resource starts on a server when the cluster is brought up.
- **Start mode (MANUAL).** If the resource Start mode is set to MANUAL, you manually start the resource on a server when you want, instead of having it start when servers in the cluster are brought up.
- **Failover (AUTO).** With the resource Failover mode set to AUTO, the resource starts on the next server in the Assigned Nodes list in the event of a hardware or software failure.

- **Failover (MANUAL).** If the resource Failover mode is set to **MANUAL**, you can intervene after a failure occurs and before the resource is moved to another node.
- **Failback (DISABLE).** With the resource Failback mode set to **DISABLE**, the resource will not fail back to its most preferred node when the most preferred node rejoins the cluster.
- **Failback (AUTO).** If the resource Failback mode is set to **AUTO**, the resource fails back to its most preferred node when the most preferred node rejoins the cluster.
- **Failback (MANUAL).** If the Failback mode is set to **MANUAL**, the resource doesn't move back to its preferred node when that node is brought back online until you are ready to allow it to happen.

The preferred node is the first server in the list of the assigned nodes for the resource.

You can set Start, Failover, and Failback modes using ConsoleOne or NetWare Remote Manager.

To set resource Start, Failover, and Failback modes using ConsoleOne, do the following:

1. Right-click the resource object and select **Properties**; then select the **Policies** tab on the property page.
2. If you don't want the cluster-wide timeout period and node number limit enforced, select the **Ignore Quorum** box.

The quorum default values were set when you installed NCS. You can change the quorum default values by accessing the properties page for the cluster object.

Selecting this box ensures the resource is launched immediately on any server in the Assigned Nodes list as soon as any server in the list is brought online.

3. Select the **Start**, **Failover**, and **Failback** modes for this resource.

The default for both Start and Failover modes is AUTO; the default for Failback mode is DISABLE.

4. Make sure the resource runs only on the master node in the cluster by selecting the **Master Only** box.

If the master node in the cluster fails, the resource fails over to whichever node becomes the master.

To set resource Start, Failover, and Failback modes using NetWare Remote Manager, do the following:

1. In the left column under the Clustering section, select **Cluster Config**.
2. From the list of resources, select a resource.
3. In the Resource Information screen, select **Policies**.
4. Select or deselect the **Ignore Quorum** box.
5. Select the **Start**, **Failover**, and **Failback** modes for this resource.
6. Select or deselect the **Master Only** box.

If the Master Only box is selected, the resource runs only on the server designated as the Master node in the cluster.

7. Select **Apply**.

Objective 5 Identify How to View and Edit Load and Unload Scripts

When editing a resource, you need to know how to view and edit the load and unload scripts for the resource.

A load script is required for each resource or volume in your cluster. The load script specifies the commands to start the resource or mount the volume on a node.

An unload script is also required to unload the resource or volume from a node when you migrate resources or volumes from one node to another in the cluster.

You can use any commands in a load script that would be used in an NCF file run from the server console. If you don't know which commands to add to your load script, consult the documentation for the application or resource.

Load and unload scripts are created for disk pools when you cluster-enable them. Because of this, it might not be necessary to configure or change the scripts for a pool.

You can view or edit a load or unload script using ConsoleOne or NetWare Remote Manager. To view or edit a script using ConsoleOne, do the following:

1. On the resource property page, select the **Load Script** or **Unload Script** tab.
2. Edit or add the necessary commands to the script to load or unload the resource on a node.

Some commands might require command-line input. You can add << to a command to indicate command-line input. For example, a script command might read like this:

```
LOAD SLPDA <<Y
```

This means that when SLPDA is loaded, it receives a Y at the command line, presumably to a question that needs a yes answer.

If more inputs are required, they can be continued on subsequent lines, as follows:

```
LOAD SLPDA <<Y
<<Y
<<N
```

The string can be up to 32 characters.

3. Specify a timeout value.

The default is 600 seconds, or 10 minutes. The timeout value determines how much time the script is given to complete.

If the script does not complete within the specified time, the resource becomes comatose.

To configure a load or unload script using NetWare Remote Manager, do the following:

1. In the left column under the Clustering heading, select **Cluster Config**.
2. From the list of resources, select a resource.
3. On the Resource Information screen, select **Loading** or **Unloading**.
4. Edit or add the necessary commands to the script to load or unload the resource on a node.
5. Specify the **Load Timeout**; then save the script by selecting **Apply**.

Objective 6 Identify How to Find NCS Configuration and Troubleshooting Information

You can find NCS documentation and TIDs at www.novell.com to help you troubleshoot configuring and running a 2-node NCS cluster.

You can access NCS documentation at <http://www.novell.com/documentation/lg/ncs6p/index.html>. The documentation contains information that is valuable when troubleshooting an NCS cluster.

To access NCS clustering TIDs, do the following:

1. Access the Novell web site (www.novell.com); then select **Support > Knowledgebase**.

2. Select the **NetWare** product category; then select the **TIDs** product set.
3. For the search term enter **cluster**; then select **Search Now**.

Many TIDs reference NCS for NetWare 5.x, but the content is valid for NetWare 6. The following TIDs address common clustering problems and questions:

Table 11-1

TID	Title	Modified	Comments
10015339	How to Remove Clustering Services and the Clustering DS Objects	02 Aug 2002	You might also be able to use the NetWare 6 install program to remove the NLMs.
10016861	What Are the Issues with Backups and NCS?	22 May 2001	Contains good background information on backups and NCS.
10017340	Novell Cluster Services for NetWare 5	22 May 2001	Contains several issues to consider when working with NCS. Although the content focuses on NetWare 5, many comments are valid for NetWare 6.
10024057	Novell Clustering Services and Virus Scanners	22 May 2001	Includes a URL for listing products (such as virus scanners) that support NetWare.

Table 11-1 (continued)

TID	Title	Modified	Comments
10050099	When Migrating a Cluster Volume, the Client Loses Connection and Takes Longer than Usual to Reconnect	04 Dec 2001	You might want to try this solution with NCS 1.6 if you are having problems with the Novell Client trying to reconnect.
10053882	Novell Cluster Services: the Gory Details	14 Jan 2002	A detailed explanation of heartbeats, node failure, false node failure, split brains, false split brains, and poison pills. A support guide is also included.
10058446	Mapping to Cluster Volume Fails after Failover	11 Jun 2002	Important information about re-initializing volume mappings.
10058722	What Are the Proper Client Configuration Settings for Clustering?	25 Apr 2001	Contains some valuable information about client settings for the Windows client and the Novell client.
10061786	Procedure for Doing an eDirectory Tree Rename/Merge	16 Apr 2001	Explains why it is necessary to take down the cluster while eDirectory synchronizes itself.
10063341	Should I Turn IPX Off on a Novell Cluster?	03 Jul 2001	Explains why IPX cannot be used with NCS.

TID 10053882 is very useful for understanding the algorithm NCS uses to cast off a cluster node. It also includes several troubleshooting tips.

Table 11-1 (continued)

TID	Title	Modified	Comments
10063780	Installing and Configuring NFS 3.0 on Novell Cluster Services	04 Jan 2002	Provides procedures and tips for cluster-enabling NFS 3.0.
10063923	How to Install Novell Enterprise Web Server on Novell Cluster Services	04 Jan 2002	Provides procedures and tips for cluster-enabling Novell Enterprise Web Server.

Exercise 11-1 (Optional) Cluster-Enable and Test DHCP Server on Your 2-Node Cluster



20 minutes

For students to perform this exercise, they must install and configure DHCP on both cluster-enabled servers.

You might want to use this exercise as a demonstration for cluster-enabling a resource.

Novell's DHCP server stores DHCP lease information in eDirectory. A volume is not needed, which means you can cluster-enable the DHCP service without a SAN.

You also do not need to use a secondary IP address for the DHCP resource because DHCP uses the IP address of the server to broadcast its service over the network.

Because the DHCP service is already running on DA2 in your classroom LAN, the instructor might decide to demonstrate this exercise by cluster-enabling DHCP on DA1 and DA2.

If you have time in class (or in a self-study environment), try installing DHCP on your own servers and then use the steps that follow to cluster-enable the service.

The following lists information you might need during the exercise;

Table 11-2

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.DEL.DIGITALAIR	192.168.1.14
Cluster nodes	
DA4.IS.DEL.DIGITALAIR (<i>first server</i>)	192.168.1.4
DA5.IS.LGA.DIGITALAIR (<i>second server</i>)	192.168.1.5
Cluster name (DN)	
DACluster.IS.LON.DIGITALAIR	192.168.1.16
Cluster nodes	
DA6.IS.LON.DIGITALAIR (<i>first server</i>)	192.168.1.6
DA7.IS.SYD.DIGITALAIR (<i>second server</i>)	192.168.1.7
Cluster name (DN)	
DACluster.IS.TXL.DIGITALAIR	192.168.1.18
Cluster nodes	
DA8.IS.TXL.DIGITALAIR (<i>first server</i>)	192.168.1.8
DA9.IS.TYO.DIGITALAIR (<i>second server</i>)	192.168.1.9

You do the following:

- [Part I: Create the DHCP Cluster Resource](#)
- [Part II: Test the Effect on the Workstations of Migrating the DHCP Resource](#)



To cluster-enable DHCP in this exercise, you must install the service on both servers before beginning Part I.

Your instructor might demonstrate this exercise for you instead.

Part I: Create the DHCP Cluster Resource

To cluster-enable the DHCP server, you create a DHCP cluster resource by doing the following:

1. Unload DHCP on *your first server* by entering **UNLOAD DHCPSRV**R at the server console prompt.

When you create a cluster resource for DHCP, NCS attempts to load DHCP on the server. If DHCP is loaded, an error occurs and the resource is listed as comatose in the Cluster State view.

2. Open and authenticate to NetWare Remote Manager on both workstations using **https://192.168.1.x:8009** (where **192.168.1.x** = your first server IP address).

For example if your cluster nodes are DA4 and DA5, the IP address of your first server (DA4) is **192.168.1.4**.

3. Minimize the **NetWare Remote Manager** window on *your second workstation*.

For example, if you are using WS4 and WS5, you can minimize NetWare Remote Manager on WS5 (your second workstation).

4. From the left panel in NetWare Remote Manager on *your first workstation* under **Clustering**, select **Cluster Config**.
5. In the right panel, scroll to the bottom and select **New Cluster Resource**.
6. Select **DHCP Server** from the drop-down list; then enter **DHCP** in the Resource Name box.
DHCP Server is the name of the cluster resource template you can use to create a DHCP resource.
7. Select **Define Additional Properties**; then select **Apply**.

8. Modify the load script:

- a. View the load script by selecting **Loading**.
- b. Edit the CLUSTER command as follows:

```
CLUSTER DHCP CN=your first  
server.OU=IS.OU=xxx.O=DigitalAir.  
T=DigitalAir-Tree
```

For example, if your cluster nodes are DA4 and DA5, you would edit the CLUSTER command to look like the following:

```
CLUSTER DHCP CN=DA4.OU=IS.OU=DEL.O=DigitalAir.  
T=DigitalAir-Tree
```

- c. Save the changes by selecting **Apply**.
- d. (Conditional) If you see a security alert, select **Yes** to continue.

9. Bring the **DHCP** resource online:

- a. In the left panel, select **Cluster Management**.

The Cluster Status view appears. Notice that the DHCP resource is added to the Cluster Resource list with an Offline status message shown.

- b. Select **DHCP**; then select **Online**.

- c. From the Refresh drop-down list, select **2 seconds**; then select **Begin Refresh**.

The Offline status message is replaced by “Loading” and then “Running” messages. The DHCP service is now cluster-enabled and running on your first server.

10. Comment-out the **DHCPSRVR** command in the **AUTOEXEC.NCF** file of *your first server*:

- a. At *your first server* console prompt, enter **EDIT AUTOEXEC.NCF**.
- b. Find the **DHCPSRVR** command line; then comment-out the command by entering a # at the beginning of the line.

- c. Save the change by pressing **Esc** and selecting **Yes**; then close the editor by pressing **Esc** and selecting **Yes**.

If you leave the command active in the AUTOEXEC.NCF file, your first server starts the DHCP server when rebooting.

After clustering starts, NCS uses the DHCP resource load script and configuration to start the DHCP server again on your first server (after it is already running).

This causes the DHCP resource to move into a comatose state. By commenting-out the DHCPSRVR command in the AUTOEXEC.NCF file, you avoid this problem.

Part II: Test the Effect on the Workstations of Migrating the DHCP Resource

To test the effect on *your first workstation* of migrating the DHCP resource, do the following:

1. View the logger screen on *your first server* and *your second server* by pressing **Ctrl + Esc** and entering **2**.
2. On *your second workstation*, make sure the IP address for the workstation is configured for automatic selection:
 - a. Right-click **My Network Places**; then select **Properties**.
 - b. Right-click **Local Area Connection**; then select **Properties**.
 - c. Select **Internet Protocol (TCP/IP)**; then select **Properties**.
 - d. Make sure that **Obtain an IP address automatically** is selected; then close the open dialogs by selecting **OK**.
 - e. Close the **Network and Dial-up Connections** window.
3. On *your second workstation*, select **Start > Programs > Accessories > Command Prompt**.
4. At the command prompt enter **IPCONFIG /RELEASE**; then enter **IPCONFIG /RENEW**.

5. At the command prompt, enter **IPCONFIG /ALL**.

Record the following information:

Table 11-3

IP Address	DHCP Server

6. From NetWare Remote Manager on WS1, migrate the **DHCP** resource from *your first server* to *your second server* and begin refreshing the page every 2 seconds.
7. Check the logger screens for unload and load messages.
8. When the migration is complete, at *your second workstation* command prompt enter **IPCONFIG /RELEASE**.

The IP addresses for the workstation and DHCP server are cleared.
9. At the command prompt enter **IPCONFIG /RENEW**.

IP addresses for the workstation and DHCP server are renewed.
10. At the command prompt enter **IPCONFIG /ALL**.
11. Compare the IP addresses in the IP configuration information with those you recorded in Step 5.

Although the workstation IP address might be the same (Windows attempts to keep the same IP address for the workstation), the DHCP server address has changed from your first server IP address to your second server IP address.

For example, if your cluster nodes are DA4 and DA5, the IP address changed from **192.168.1.4** to **192.168.1.5**.
12. Close the Command Prompt window by entering **EXIT** at the command prompt.

To test the effect on *your first workstation* of migrating the DHCP resource, do the following:

1. On *your first workstation*, minimize the **NetWare Remote Manager** window.
2. On *your second workstation*, maximize the **NetWare Remote Manager** window and show the **Cluster Status** view.
3. Refresh the view by selecting **Begin Refresh**; then select **Stop Refresh**.
4. On *your first workstation*, make sure the IP address for the workstation is configured for automatic selection:
 - a. Right-click **My Network Places**; then select **Properties**.
 - b. Right-click **Local Area Connection**; then select **Properties**.
 - c. Select **Internet Protocol (TCP/IP)**; then select **Properties**.
 - d. Make sure that **Obtain an IP address automatically** is selected; then close the open dialogs by selecting **OK**.
 - e. Close the **Network and Dial-up Connections** window.
5. On your first workstation, select **Start > Programs > Accessories > Command Prompt**.
6. At the command prompt enter **IPCONFIG /RELEASE**; then enter **IPCONFIG /RENEW**.
7. At the command prompt enter **IPCONFIG /ALL**.
Record the following information:

Table 11-4

IP Address	DHCP Server

8. From NetWare Remote Manager on WS2, migrate the **DHCP** resource from *your second server* to *your first server* and begin refreshing the page every 2 seconds.
9. Check the logger screens for unload and load messages.

10. When the migration is complete, at *your first workstation* command prompt enter **IPCONFIG /RELEASE**.
The IP addresses for the workstation and DHCP server are cleared.
11. At the command prompt enter **IPCONFIG /RENEW**.
IP addresses for the workstation and DHCP server are renewed.
12. At the command prompt enter **IPCONFIG /ALL**.
13. Compare the IP addresses in the IP configuration information with those you recorded in Step 7.
14. Close the Command Prompt window by entering **EXIT** at the command prompt.
15. Close NetWare Remote Manager on both workstations.

(End of Exercise)



40 minutes

Exercise 11-2 Cluster-Enable and Test iFolder on Your 2-Node Cluster

A key reason for upgrading to NetWare 6 at Digital Airlines is to implement iFolder as a remote file access and management tool for SLC office employees.

You decide that cluster-enabling iFolder is an ideal solution for ensuring that the employees have constant access to their user data and folders.

By placing the iFolder user data on the shared storage device, you can migrate iFolder from one node to another when you need to perform maintenance on an iFolder server. You can also ensure high availability if the server hosting iFolder fails.

The following lists information you need during the exercise:

Table 11-5

Cluster Information	IP Addresses
Cluster name (DN)	
DACluster.IS.DEL.DIGITALAIR	192.168.1.14
Cluster nodes	
DA4.IS.DEL.DIGITALAIR (<i>first server</i>)	192.168.1.4
DA5.IS.LGA.DIGITALAIR (<i>second server</i>)	192.168.1.5
Current iFolder addresses	
DA4	192.168.1.34
DA5	192.168.1.35
iFolder secondary IP address	192.168.1.34
Cluster name (DN)	
DACluster.IS.LON.DIGITALAIR	192.168.1.16
Cluster nodes	
DA6.IS.LON.DIGITALAIR (<i>first server</i>)	192.168.1.6
DA7.IS.SYD.DIGITALAIR (<i>second server</i>)	192.168.1.7
Current iFolder addresses	
DA6	192.168.1.36
DA7	192.168.1.37
iFolder secondary IP address	192.168.1.36
Cluster name (DN)	
DACluster.IS.TXL.DIGITALAIR	192.168.1.18
Cluster nodes	
DA8.IS.TXL.DIGITALAIR (<i>first server</i>)	192.168.1.8
DA9.IS.TYO.DIGITALAIR (<i>second server</i>)	192.168.1.9
Current iFolder addresses	
DA8	192.168.1.38
DA9	192.168.1.39
iFolder secondary IP address	192.168.1.38

In this exercise, you do the following:

- [Part I: Cluster-Enable iFolder](#)
- [Part II: Test the IFOLDER Resource](#)

Part I: Cluster-Enable iFolder

Before cluster-enabling iFolder, iFolder needs to be installed on all server nodes that will host iFolder during migration or failover.

Your first server and *your second server* already have iFolder installed with the IP addresses listed in Table 11-5.

Cluster-enabling iFolder includes several post-installation tasks including editing AUTOEXEC.NCF and creating an iFolder resource.

To cluster-enable iFolder, do the following:

1. Comment-out the iFolder command lines in the **AUTOEXEC.NCF** file of *your first server* and *your second server*:
 - a. At the server console prompt, enter **EDIT AUTOEXEC.NCF**; then scroll through the file and find the following commands:

```
ADD SECONDARY IPADDRESS iFolder IP address
...
SEARCH ADD SYS:\APACHE\IFOLDER\SERVER
STARTIFOLDER
```

The *iFolder IP address* represents the address entered when installing iFolder (see Table 11-5).

Because these command lines are included in the load script of the iFolder resource you create, you do not need them in the AUTOEXEC.NCF file.

- b. Comment-out the command lines by entering a # at the beginning of each line.

The command lines should look like the following:

```
#ADD SECONDARY IPADDRESS iFolder IP address
...
#SEARCH ADD SYS:\APACHE\IFOLDER\SERVER
#STARTIFOLDER
```

- c. Save the changes by pressing **Esc** and selecting **Yes**.
 - d. Exit the NetWare Text Editor by pressing **Esc** and selecting **Yes**.
2. Restart *your first server* and *your second server* by entering **RESTART SERVER** at the console prompt.

NetWare 6 reloads on each server without starting iFolder, and the cluster restarts.

3. Start NetWare Remote Manager on *your first workstation* using **https://192.168.1.x:8009** (where **198.168.1.x** = your first server IP address).



For this exercise, you can select which workstation you want to use as your first (and second) workstation.

4. Using NetWare Remote Manager, on the SCSI hard drive create volume **USERDATA** in an **IFOLDER** pool:
 - a. In the left panel under the Manage Server heading, select **Volumes**.
 - b. In the right panel under Partition Management, select **Disk Partitions**.
A list of hardware adaptors with all associated devices, partitions, pools, and volumes appears.
 - c. Scroll to the SCSI adaptor for *your first server*.

Notice the MULTIMEDIA pool listed. Also notice the free disk space available for partitioning on the SCSI hard drive.

- d. Next to the largest amount of free disk space available (there might be more than one listing), select **Create**.
- e. Under the Novell Storage Services heading, select **Create a New Pool and Volume**.
- f. For the pool name enter **IFOLDER**; for the volume name enter **USERDATA**; then scroll down and select **Create**.
- g. Confirm the creation by selecting **OK**.

You are returned to the list of hardware adaptors. An IFOLDER pool appears with volume USERDATA listed.



If you do not see USERDATA, or if the size of the IFOLDER pool is 0, the pool needs to be activated.

You can activate the pool by using ConsoleOne. View the properties for *your first server*; then select **Media > NSS Pools > IFOLDER > Activate**. Refresh the NetWare Remote Manager screen to view the pool and volume.

- h. In the left panel of NetWare Remote Manager select **Volumes**; then verify the mounted status of **USERDATA**.
 - i. (Conditional) If the volume is not mounted, mount the volume by selecting **NO** in the Mounted column.
5. Create a cluster-enabled resource for iFolder:
 - a. In the left panel in NetWare Remote Manager, select **Cluster Config**; then select **New Cluster Resource**.
 - b. From the drop-down list, select **iFolder Server**; then enter **IFOLDER** in the Resource Name text box.
 - c. Select **Define Additional Properties**; then select **Apply**.
 6. Modify the load and unload scripts for the IFOLDER resource:
 - a. In the IFOLDER Resource Information page, select **Loading**.

- b. Make the following changes to the first 3 command lines:

```
nss /poolactivate=IFOLDER
mount USERDATA
add secondary ipaddress 192.168.1.3x
```

The secondary IP address for cluster-enabling iFolder is listed in Table 11-5.

For example, if your cluster includes nodes DA4 and DA5, your iFolder secondary IP address is **192.168.1.34**.

- c. Make the following changes to the last command line:

```
load address space = ifolder apache -f
sys:\apache\ifolder\server\httpd.conf
```

Enter a space (not a hard return) after -f.

- d. Select **Apply**; then select **Unloading**.

- e. Delete the following set of command lines:

```
unload apache
delay 10
unload ldapssl
unload ldapsdk
unload fpsm
```

- f. Enter the following as the first command line in the unload script:

```
unload address space = ifolder apache -f
sys:\apache\ifolder\server\httpd.conf
```

- g. Edit the **nss** and **del** command lines to reflect the following:

```
nss /pooldeactivate=IFOLDER /override=question
del secondary ipaddress 192.168.1.3x
```

The secondary IP address you want to delete is the same IP address used in your load script (see Table 11-5).

- h. Select **Apply**.

7. Copy the existing iFolder user data to USERDATA:

- a. On *your second workstation*, make sure you are logged in as **Admin** to *your first server*.

- b. From *your second workstation*, right-click **My Network Places**; then select **Open**.
 - c. Double-click **Novell Connections** > *your first server* > **DATA**.
 - d. Select the **IFOLDER** folder; then right-click and select **Copy**.
 - e. From the Address drop-down list select **My Network Places**; then double-click **Novell Connections** > **DigitalAir-Tree>DigitalAir>xxx>IS> your first server_USERDATA**.
 - f. Right-click the **USERDATA** window; then select **Paste**.
All user data is now available on the SCSI hard drive in the shared **USERDATA** volume.
You might need to select **View** > **Refresh** to see the copied files.
8. Edit the **HTTPD_ADDITIONS_NW.CONF** file on *your first server* to reflect the iFolder secondary IP address and **USERS** folder on the SCSI hard drive:
- a. From the Graphical Console on *your first server*, select **Novell** > **Utilities** > **Editor**.
 - b. Select **File** > **Open**; then find and open the **HTTPD_ADDITIONS_NW.CONF** file by double-clicking **APACHE** > **IFOLDER** > **SERVER**.
 - c. Find the **LdapHost** command (it occurs twice); then change both commands to reflect the following:

```
LdapHost 192.168.1.3x
```

The **LdapHost** IP address is the iFolder secondary IP address listed in Table 11-5.
 - d. Find the **iFolderServerRoot** command (it occurs twice); then change both commands to reflect the following:

```
iFolderServerRoot USERDATA:\IFOLDER
```

- e. When you finish editing, select **File > Save**; then exit the Editor.
9. Copy **HTTPD.CONF** and **HTTPD_ADDITIONS_NW.CONF** from *your first server* to **SYS:\APACHE\IFOLDER\SERVER** on *your second server*:
 - a. On *your second workstation* desktop, right-click **My Network Places**; then select **Open**.
 - b. Double-click **Novell Connections > DigitalAir-Tree > DIGITALAIR > xxx > IS > your first server_SYS > Apache > iFolder > Server**.

For example, if your cluster includes nodes DA4 and DA5, you would double-click **Novell Connections > DigitalAir-Tree > DIGITALAIR > DEL > IS > DA4_SYS > Apache > iFolder > Server**.
 - c. Select and copy **HTTPD.CONF** and **HTTPD_ADDITIONS_NW.CONF**.
 - d. From the Address drop-down list select *your location container* for *your second server*; then double-click **IS**.

For example, if your cluster includes nodes DA4 and DA5, you would select **LGA > IS**.
 - e. Double-click *your second server_SYS > Apache > iFolder > Server*.
 - f. Paste **HTTPD.CONF** and **HTTPD_ADDITIONS_NW.CONF** into the **SERVER** folder, replacing the existing files.
 10. Restart *your first server* and *your second server* by entering **RESTART SERVER** at the console prompt.
 11. After NetWare 6 loads on both servers, check the Cluster Status view in NetWare Remote Manager.

All resources in the Cluster Resource list are running, except for the IFOLDER resource.

Part II: Test the IFOLDER Resource

Do the following:

1. From NetWare Remote Manager, bring the IFOLDER resource online by selecting **IFOLDER** in the Cluster Resource list; then select **Online** and begin the refresh.

NCS loads and runs the IFOLDER resource on *your first server*.

2. From *your first server* console prompt, enter **VOLUMES**.

USERDATA is mounted on your first server.

3. From NetWare Remote Manager, migrate the **IFOLDER** resource to *your second server*.

4. After the IFOLDER resource is running on your second server, at *your second server* console prompt enter **VOLUMES**.

USERDATA is now mounted on your second server.

5. On *your second server*, check the load messages for Apache server by pressing **Ctrl + Esc** and selecting **Apache for Netware**.

If Apache server loaded successfully, you see an “iFolder server initialization complete” message.

6. Now that you have successfully migrated IFOLDER from one node to another, write your own plan for testing cluster-enabled iFolder:

7. Try out your test plan and share your results with the rest of the students.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	What You Learned
1. Identify Cluster-Aware and Cluster-Naive Applications	<p>When creating a resource for an NCS cluster you need to be familiar with the following types of applications:</p> <ul style="list-style-type: none">■ Cluster aware■ Cluster naive <p>Examples of cluster-aware applications for NCS 1.6 include Apache Web Server, BorderManager® (Proxy and VPN), DHCP Server, and Enterprise Web Server (LDAP and NDS).</p>
2. Identify How to Cluster-Enable an Application	<p>You cluster-enable a service such as an application by creating a cluster resource.</p> <p>The resource includes a unique IP address and is available for migration from the Resource list in the Cluster State view (ConsoleOne) or the Cluster Status view (NetWare Remote Manager).</p> <p>Cluster resources can be created for cluster-aware or cluster-naive applications such as web sites, email servers, databases, or any other server-based applications or services you want to make available to users at all times.</p> <p>You can create a cluster resource using ConsoleOne or NetWare Remote Manager.</p>

Objective	What You Learned
3. Identify How to Assign Nodes to a Resource	<p>You assign nodes to a resource for NCS to know which nodes to migrate the resource to during a failover.</p> <p>When you create a resource on a cluster or cluster-enable a volume, the nodes in the cluster are assigned to the resource or volume. The order of assignment is the order the nodes appear in the resource list.</p> <p>You can assign or unassign nodes to the resource or volume, or change the failover order.</p>
4. Identify How to Set Start, Failover, and Failback Modes	<p>You can configure the start, failover, and failback of cluster resources to happen manually or automatically.</p> <p>You can set Start, Failover, and Failback modes using ConsoleOne or NetWare Remote Manager.</p>
5. Identify How to View and Edit Load and Unload Scripts	<p>When editing a resource, you need to know how to view and edit the load and unload scripts for the resource.</p> <p>A load script is required for each resource or volume in your cluster. The load script specifies the commands to start the resource or mount the volume on a node.</p> <p>An unload script is also required to unload the resource or volume from a node when you migrate resources or volumes from one node to another in the cluster.</p> <p>You can view or edit a load or unload script using ConsoleOne or NetWare Remote Manager.</p>

Objective	What You Learned
6. Identify How to Find NCS Configuration and Troubleshooting Information	<p>You can find NCS documentation and TIDs at www.novell.com to help you troubleshoot configuring and running a 2-node NCS cluster.</p> <p>You can access NCS documentation at http://www.novell.com/documentation/lg/ncs6p/index.html.</p> <p>The documentation contains information that is valuable when troubleshooting an NCS cluster.</p>

Exercise Answers

The following are the exercise answers.

Exercise 11-2. Cluster-Enable and Test iFolder on Your 2-Node Cluster

Part II: Test the IFOLDER Resource

The following is a suggested method for testing cluster-enabled iFolder:

1. Open ConsoleOne or NetWare Remote Manager on both workstations and view the Cluster State view.
2. Make sure you are logged in to iFolder as the same user on both workstations.
3. Open a user's **Home** folder on both workstations.
4. Copy the **Marketing Video** folder on *your first workstation* to the user's **Home** folder on *your first workstation*.

The folder is copied to the Home folder on your first workstation, synchronized to USERDATA on the SCSI hard drive, and then synchronized to the Home folder on your second workstation.

5. During the copying and synchronization, migrate the **IFOLDER** resource between the cluster nodes twice.

The synchronization pauses during migration and then continues after the resource is running on the new node.

6. When the synchronization is complete, on both workstations right-click the **iFolder** icon in the system tray, select **Account Information**; then select **View Activity**.

You see messages indicating that iFolder lost connectivity with the resource (server) during migration, tried to find the resource again by using the IP address (such as 192.168.1.34), and then logged in again as the user and continued synchronizing files.

These messages indicate that iFolder is cluster-aware and able to handle a migration or failover without prompting the user for help.

MODULE 5

Troubleshoot a NetWare 6 Network

Section 12 Troubleshoot a NetWare 6 Network

SECTION 12 Troubleshoot a NetWare 6 Network

Duration: 4 hours

In this section, you troubleshoot server and communication problems on a NetWare 6 network.

Prior to beginning this section, introduce problems into the network by completing the steps in "Setup for Section 12" on Setup-50.

Objectives

1. [Create a Disaster Recovery Plan](#)
2. [Troubleshoot Network Problems](#)

Introduction

In this section, you develop a disaster recovery plan and then troubleshoot server and communication problems on the network.

Objective 1 **Create a Disaster Recovery Plan**

A critical component of any network implementation is a plan for recovering from a disaster. The services provided by your network represent a significant investment on the part of your organization.

The data or services hosted by the network are of great value to your company. Losing data or services can potentially cost your organization millions of dollars.

Don't wait until a disaster strikes to determine how you will restore service. Doing so could cost you your job. Creating a disaster plan involves the following:

- [Planning for Hardware Failures](#)
- [Planning for Calamities](#)

Planning for Hardware Failures

The first element in a disaster recovery plan is a strategy for dealing with hardware failures.

Some components in a computer system are more susceptible to failure than others. Purely electronic components, such as the system CPU, have a relatively low failure rate.

Rather than wearing out, failure of these components is usually due to physical impact or from electrostatic discharge (ESD).

However, mechanical components slowly wear as they are used. Devices such as hard disk drives have a mean time before failure (MTBF) associated with them. Eventually, they will fail.

To plan for failure of such devices, consider doing the following:

- Keep a supply of replacement parts available, especially items such as hard disk drives, CD drives, and network boards, or know where to obtain new parts within a couple of hours.

If you don't have a spare when a device fails, you might be forced to order a replacement. Most organizations can't afford to have data or services unavailable waiting for a part to arrive.
- To preserve network data, develop and abide by a rigid backup policy.

Unfortunately, many busy network administrators place a low priority on creating and verifying eDirectory and file system backups.

If a server hard disk fails and you don't have a recent backup to restore from, it could cost your job.

- To further protect data, consider implementing a RAID solution that incorporates redundant disks. If a particular disk in a RAID array fails, the redundant disk can take over, ensuring the data remains intact.
- To make sure NetWare services remain available, consider implementing a clustering solution. If a server in the cluster goes down, other servers in the cluster can take over, ensuring services remain available to end users.

Planning for Calamities

In addition to hardware failures, you need to be prepared for other types of disasters such as flood, fire, earthquake, or even urban violence.

These events can destroy an entire site, making redundant hard disks, clustered servers, and backups useless.

To prepare for calamities, consider the following:

- When backing up, make at least 2 copies. Keep one copy on site and one copy somewhere else. This helps ensure that your organization's data remains intact should a disaster strike.

If possible, select a secondary storage site in a different locale. The likelihood of a calamity striking 2 communities at the same time is relatively low.

As an alternative, you can purchase Internet-based backup services. Server backups are performed over the Internet to a remote storage system hosted by the company providing the service.

- If your organization's data and services are exceptionally valuable, consider implementing redundant servers at a remote site.

Configuring such a system is beyond the scope of this course. However, doing so ensures that a localized system outage will not hamper network services.

Objective 2 Troubleshoot Network Problems

In this objective, you troubleshoot and repair server and communication problems on the network.



3 hours 30 minutes

Exercise 12-1 Troubleshoot Network Problems

In this exercise, help students isolate each problem and then guide them to the proper solution.

However, avoid giving students too much information. Make them draw upon their experience in class to arrive at a solution.

Verify that students fix every bug you introduce into the system.

Digital Airlines is currently experiencing company-wide network failures. Users are frustrated; management and customers are angry.

In this exercise, you must troubleshoot and repair every problem on the network.

The symptoms being reported to the Digital Airlines help desk include:

- Users in all locations can't log in to DA4, DA5, DA6, DA7, DA8, or DA9.
- Three users can't log in to DA3.
- Local administrators report various time error messages on several servers.
- iFolder users on DA7 can't log in with their iFolder clients.
- iFolder users on DA8 and DA9 report that all of their files are missing.
- Administrators in Sydney report that workstations aren't receiving DHCP addresses.

(End of Exercise)

Summary

The following is a summary of the objectives in this section:

Objective	Summary
1. Create a Disaster Recovery Plan	<p>A critical component of any network implementation is a plan for recovering from a disaster. The services provided by your network represent a significant investment on the part of your organization.</p> <p>The data or services hosted by the network are of great value to your company. Losing data or services can potentially cost your organization millions of dollars.</p>
2. Troubleshoot Network Problems	<p>You successfully troubleshoot and repair every problem on the network.</p>

APPENDIX A Network Components

To troubleshoot LAN issues, you must have a solid understanding of what a network is and the components that comprise a network.

The following is an overview of the components of a network. Each could be considered as a point of failure when troubleshooting your LAN/WAN environment:

Table A-1

Components	Characteristics
Local area network (LAN)	<ul style="list-style-type: none"> ■ Small group of connected computers in one location. ■ Typically does not exceed tens of kilometers. ■ Provides data transmission services. ■ Communication links are owned and maintained by LAN owner. ■ Transmission speed is typically measured in megabits per second (mbps).
Wide area network (WAN)	<ul style="list-style-type: none"> ■ Comprised of multiple LANs. ■ Often uses telephone or satellite communications. ■ Access can be leased from a WAN service provider who is responsible for maintenance. ■ Transmission speed is typically measured in kilobits per second (kbps). ■ Also known as one of the following network types: <ul style="list-style-type: none"> ■ Enterprise network, which connects all LANs of a single organization. ■ Global network, which spans the earth, and can include networks of several organizations. The Internet is an example of a global network.

Table A-1 (continued)

Components	Characteristics
Network services	<ul style="list-style-type: none"> ■ Network services are the capabilities that networked computers share. ■ A service provider is not a computer; it is a subset of the computer's software and hardware. <ul style="list-style-type: none"> ■ <i>Servers are classified as service providers.</i> They only provide services. ■ <i>Clients are classified as service requesters.</i> They only request services. ■ <i>Peers</i> can be classified as both a service requester or provider. They provide and request services.
Network classification	<ul style="list-style-type: none"> ■ Peer-to-peer networks allow any entity to both request and provide network services. Peer-to-peer network software is designed so that peers perform the same or similar functions for each other. ■ Server-centric networks involve strictly defined roles. By definition, a server-centric network places restrictions upon which an entity can make requests or service them.
Transmission media	<ul style="list-style-type: none"> ■ A transmission media is the pathway networked entities use to contact each other. ■ A transmission media cannot guarantee that other network devices will understand a message. However, It can guarantee a message delivery path.
Protocol	<ul style="list-style-type: none"> ■ Protocols are the rules required to help entities communicate with or understand each other. ■ A protocol can be one rule or a complete set of rules and standards that allow different devices to hold conversations.

Table A-1 (continued)

Components	Characteristics
Connectivity devices	<p data-bbox="893 283 1477 346">Connectivity devices to connect separate segments of the network or internetwork.</p> <p data-bbox="893 357 1477 483">A segment is a portion of the network transmission media that is assigned a network address and provides access to network resources for all attached clients and servers.</p> <p data-bbox="893 493 1477 556">Network connectivity devices connect individual devices to a single network. Devices include</p> <ul style="list-style-type: none"> <li data-bbox="893 567 1477 682">■ Transmission media connectors. These attach directly to the transmission media and serve as the physical interface between the media and computing devices. <li data-bbox="893 693 1477 829">■ Network interface boards. A network interface board includes all the circuitry needed to create the necessary physical and logical connections between your computer, or other device, and the transmission medium. <p data-bbox="893 840 1477 903">The following terms also describe network interface boards or devices that attach to them:</p> <ul style="list-style-type: none"> <li data-bbox="893 913 1477 997">■ Transceivers. A device that can transmit as well as receive electric or electromagnetic signals on the transmission media. <li data-bbox="893 1008 1477 1228">■ Network interface card (or board). A printed circuit board called a <i>network interface card</i> (NIC, network board, or network adapter) includes the circuitry and mechanical connections to convert the computer's electric signals to the signals used on the medium. Some network boards provide more than one type of media connector. <p data-bbox="893 1239 1477 1415">A network board usually uses an internal transceiver (built into the circuit board). However, some implementations require the use of external transceivers that attach to the cable or to the media connector of the network board.</p>

Table A-1 (continued)

Components	Characteristics
Connectivity devices (continued)	<ul style="list-style-type: none"> ■ Transmission media adapter. When a network board uses a connector that is different from what is already attached to the transmission medium, a <i>transmission media adapter</i> is used. ■ Modems. A <i>modem</i> (MOdulator/DEModulator) converts a computer's digital signals to an analog transmission signal to use with telephone lines or microwave transceivers. You can use modems in the following ways: <ul style="list-style-type: none"> ■ To connect directly to a remote server or your network. ■ In some instances, to take the place of network boards in connecting a device to a network. Some small organizations use modems as their WAN interface, but this type of connection can be very slow. ■ To connect to an ISP for Internet access. Modem types: <ul style="list-style-type: none"> ■ Telephone line. ■ Cable/DSL modems. ■ Hubs. Some network implementations require a central point of connection between media segments. These central points are referred to as hubs, multiport repeaters, or concentrators. The hub organizes the cables and transmits incoming signals to the other media segments. <ul style="list-style-type: none"> ■ Active hubs. An <i>active hub</i>, which connects medium segments together, regenerates or amplifies signals. ■ Passive hubs. A <i>passive hub</i> connects medium segments together; it does not regenerate or amplify signals.

Table A-1 (continued)

Components	Characteristics
Connectivity devices (continued)	<ul style="list-style-type: none"> ■ Multiport repeaters. A multiport repeater regenerates the signal and then transmits it to all ports. ■ Switches. A switch receives a transmission and forwards the signal through the port that allows the transmission to be delivered. Using switches, you can set up a network where all transmission media segments are permanently connected, but each segment is used only when a signal is directed to a computer on that segment. This can significantly improve performance by optimizing bandwidth use. ■ Multiplexers. A <i>multiplexer</i> combines 2 or more separate signals on a transmission media segment, allowing you to efficiently use the entire transmission media bandwidth.
Internetwork Connectivity Devices	<p>In an internetwork, 2 or more networks are connected to provide access to remote resources using internetwork connectivity hardware.</p> <p>Each network in an internetwork must be assigned a unique network address.</p> <p>The following devices connect distinct networks while protecting their individuality:</p> <ul style="list-style-type: none"> ■ Routers. A router segregates networks and passes data to the network it is intended for. Routers operate using a <i>logical</i> network address assigned by the network administrator. They also use MAC addresses, which are assigned by the hardware manufacturer.

Table A-1 (continued)

Components	Characteristics
Internetwork Connectivity Devices (continued)	<ul style="list-style-type: none"> <li data-bbox="893 283 1479 315">■ Channel Service Unit/Digital Service Unit. <li data-bbox="893 315 1479 409">Organizations that provide transmission media for others might require <i>channel service units</i> (CSU) and <i>digital service units</i> (DSU) for connections. <li data-bbox="893 409 1479 493">CSU/DSU devices, also known as an Integrated Services Unit (ISU), are 2 components of a Data Communications Equipment (DCE) device. <li data-bbox="893 493 1479 640">The CSU/DSU is comparable to a modem, but provides digital-to-digital services instead of digital-to-analog. They protect you, and other public network users, from electrical noise or unsafe electric voltages. <li data-bbox="893 640 1479 766">These devices prepare digital signals for transmission using the rules specified for the network, and ensure that the transmitted signal is of the proper signal strength and format. <li data-bbox="893 766 1479 869">The CSU/DSU is usually attached to a router by a synchronous serial interface (such as a V.35 connection).

APPENDIX B The Network Communication Process

In addition to needing a firm understanding of the components that make up a network, you must understand the following to determine where in the communication process a specific problem might occur:

- [ISO Layers and the Communication Process](#)
- [IP Routing](#)

ISO Layers and the Communication Process

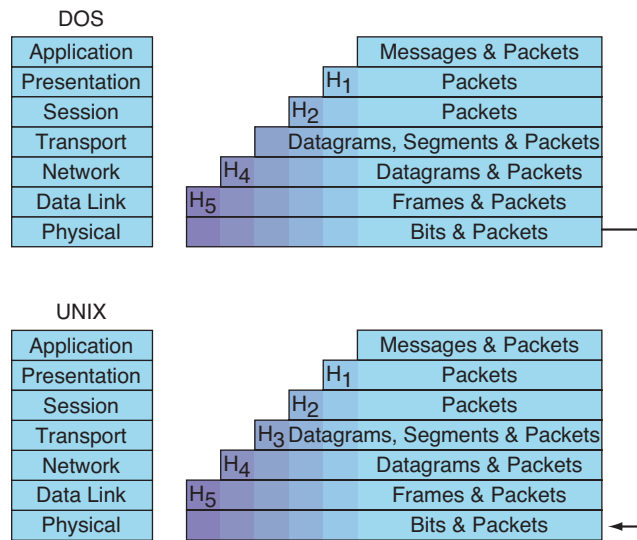
You have already learned that the International Organization for Standardization (ISO) developed a standard for communication called the OSI model.

Remember that this is a standard and not all layers in the OSI model are used in every communication.

The communication process in the OSI model begins by encapsulating the data in each layer. *Encapsulation* is the process by which layer-specific information is added to the data.

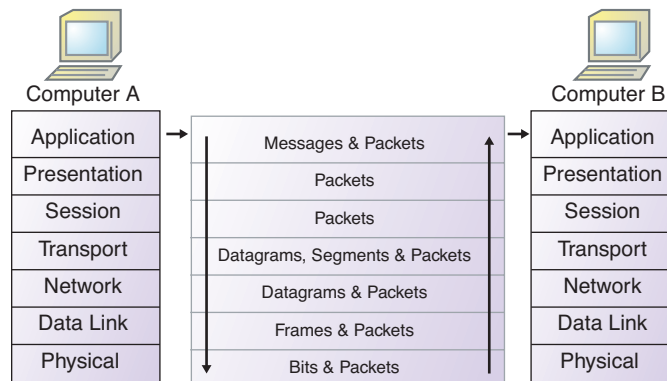
The following shows the OSI layers and the data packet reference names:

Figure B-1



The following can be used identify the communication process:

Figure B-2



To illustrate the OSI model, suppose a user wants to log in to the network from computer A. The data entered, such as the username and the password, is sent to computer B for authentication.

The following process occurs in computer A:

Table B-1

Process	OSI Layer
1. The username and password are entered into the login application.	Application
2. This data is broken into data packets. It is encrypted because it contains secure information, such as the password. Data packets are also compressed to enable quick data transfer.	Presentation
3. Computer A prepares to send the data and establish a communication link with computer B.	Session
4. Information regarding the sequence of data packets is included in the data packets.	Transport
5. The data packets are collated into logical groups called <i>datagram packets</i> . At this stage, computer A determines the network address of computer B and information about the best path for sending the packet. This information is included in the datagram packet.	Network
6. The datagram packets are broken into <i>frames</i> . Information regarding the address of computer B (destination) is included in each frame. Additional information called a <i>checksum</i> is also included. A checksum is a simple error detection method. It makes sure that each transmitted message is accompanied by a numerical value based on the number of bits in the message.	Data Link

Table B-1 (continued)

Process	OSI Layer
7. Finally, frames are converted into bits and sent to the network board of computer A. From the network board of computer A, a stream of bits is sent to computer B.	Physical
When the data reaches computer B, the following occurs:	
1. Data in the form of a stream of bits is retrieved from the transmission medium with the help of the network board in computer B.	Physical
2. The stream of bits is reorganized into logical groups called <i>frames</i> , and the checksum is verified. If the checksum is not correct, computer B sends a message to computer A asking it to retransmit the data.	Data
3. From the frames, address information is verified.	Data/MAC
4. Data integrity is checked.	Data/LLC
5. The original data packet is recreated in the proper sequence.	Presentation
6. The username and the password are passed to the application that authenticates the user.	Application

IP Routing

Every packet being processed by a TCP/IP host has a source and destination IP address. The router examines the destination address on each packet, compares it with entries in its local routing table, and decides what action to take.

There are three code paths that the packet can take:

- It can be passed up to a protocol layer above IP if the destination IP address is itself, that is, to a local application such as GroupWise or BorderManager Proxy.



This occurs when you are using a software router on NetWare, Linux, or Windows, etc.

- It can be forwarded using one of the locally attached network interface boards if the packet is destined for another known network.

This assumes that the TCP/IP host has multiple interfaces and has routing enabled.

- It can be discarded.

The route table can maintain 4 types of routes. They are listed as follows in the order that they are searched for a match:

1. Host (a route to a single, specific destination IP address)
2. Subnet (a route to a subnet)
3. Network (a route to an entire network)
4. Default (used when there is no other match)

IP compares the destination IP address of the packet that it is processing with the entries in the table. The next action is based on the following:

- **Host entry exists.** If a host entry exists that matches the destination IP address, IP forwards it to the next hop associated with that host entry.

Host entries are usually found in routing tables when ICMP has added the entry due to the pathMTU algorithm or an ICMP redirect.

Check the TCPCON > IP Routing table and verify if the protocol associated with that route is ICMP.

The maximum transfer unit (MTU) is the maximum size of data packets that can be transferred across a given physical network.

For LANs, the MTU is determined by the network hardware. For WANs that use serial lines to interconnect packet switches, the MTU is determined by the software.

The Path MTU is the smallest of all MTUs. It governs the size of the largest IP packet that can be sent across the path without fragmentation.

Servers send ICMP redirect packets to notify other IP hosts that they should not use a specific gateway to route IP traffic to certain destinations.

The server receives ICMP redirect packets from other IP hosts that detect that the server is sending packets over inappropriate gateways.

If the redirect value received is high, it might indicate that the default router you're pointing to is not the correct one, or that you have a routing problem in your network.

- **Subnet entry exists.** If a host entry is not found but a subnet entry exists matching the destination IP address, IP forwards the packet to the next hop associated with that subnet entry.

Subnet entries exist when RIP2, OSPF, or static entries have been added with a nondefault subnet mask.

- **Network entry exists.** If a subnet entry is not found, but IP finds a network entry matching this destination IP address, it forwards the packet to the next hop associated with that network entry.

Customers running in default NetWare TCP/IP mode will have network entries.

- **Default route entry exists.** If a network entry is not found, but IP finds that a default route entry exists, it forwards the packet to the next hop associated with that default entry.

The default route is most commonly inserted as a static route through INETCFG but can also be learned via RIP or OSPF.

Failure to have a default route can often lead to communication failures on the network.

- **No Match is Found.** If the match has *not* been found in the table at this stage, the packet is dropped, and an ICMP Destination Unreachable message is triggered to notify the sender that the host or network is unreachable.

Whenever a communication problem occurs, most likely a route entry doesn't exist for the network or host you are trying to communicate with.

APPENDIX C Protocol Analyzers

The following provides additional information about using and purchasing a protocol analyzer:

- [Protocol Analyzer Elements](#)
- [Protocol Analyzer Types](#)
- [Protocol Analyzer Placement](#)

Protocol Analyzer Elements

Most analyzers use the same or similar elements to analyze the network. The same element in different protocol analyzers might have different names.

For example, LANalyzer for Windows (LZFW) uses the term *capture filter*; Sniffer uses the term *pre-filter*.

The elements of protocol analyzers are described below:

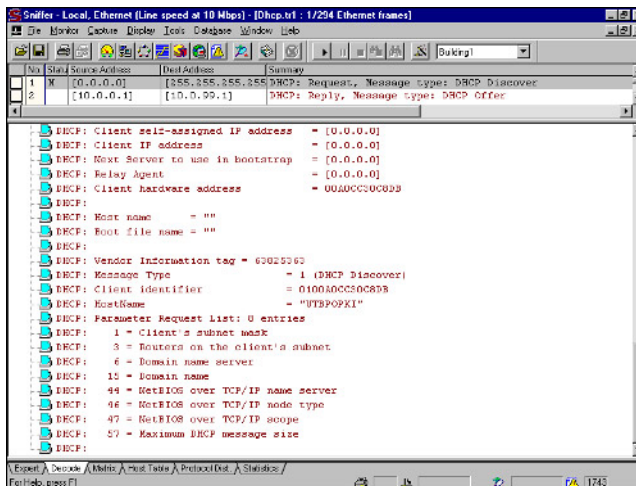
- **Ports.** The port is the connection through which packets flow into the analyzer.

The analyzer port determines the method you use to connect to the network. For example, if your analyzer has a 10Base-T port, you can only connect to a 10Base-T network, not a token ring network.
- **Decodes.** Decodes, which are wrapped around all captured packets, are the deciphered version of the bits on the wire.

The decodes control how packet contents appear on the screen.

For example, the following shows a decoded DHCP packet. If the analyzer didn't have a decode for that protocol, you would see a list of hexadecimal numbers.

Figure C-1



When an analyzer decodes a packet for you, it breaks down the packet contents and lists the individual fields, the field contents, and their meanings.

For example, the decode in the previous figure breaks down the DHCP options to list the magic cookie, message type, client identifier, hostname, and parameter request list.

If a DHCP decode was not available with LZFW, you would need to look up the field lengths and values in the DHCP specifications (RFC 2131) and manually decode the packets.

- **Capture filters.** Capture filters define the traffic that should be copied from the network into the trace buffer.

For example, you could set a capture filter to capture all TCP/IP traffic to and from a workstation, or you could set a capture filter to capture all IPX/SPX traffic.

By setting a capture filter, you reduce the number of packets actually captured. This allows you to focus on the traffic.

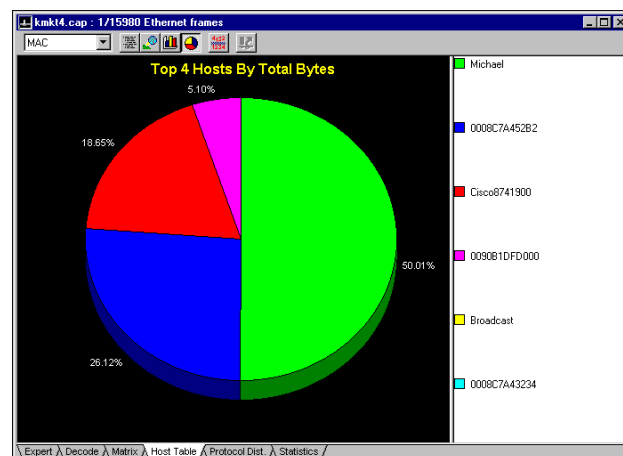
- **Display filters (post-filters).** Display filters create subsets of packets in the trace buffer and allow you to reduce the number of packets that you view.

For example, you can set a capture filter for all TCP/IP traffic. If you want to determine which types of broadcasts occurred on the network, you can apply a display filter looking only for packets addressed to the broadcast address (0xFF-FF-FF-FF-FF-FF).

- **Gauges and graphs.** Gauges and graphs show long- and short-term traffic trends and can help you by giving you a graphical view of the network health and traffic flow.

In the following figure, the Top Hosts pie chart on Sniffer indicates that 50% of the traffic is attributed to one user, Michael.

Figure C-2



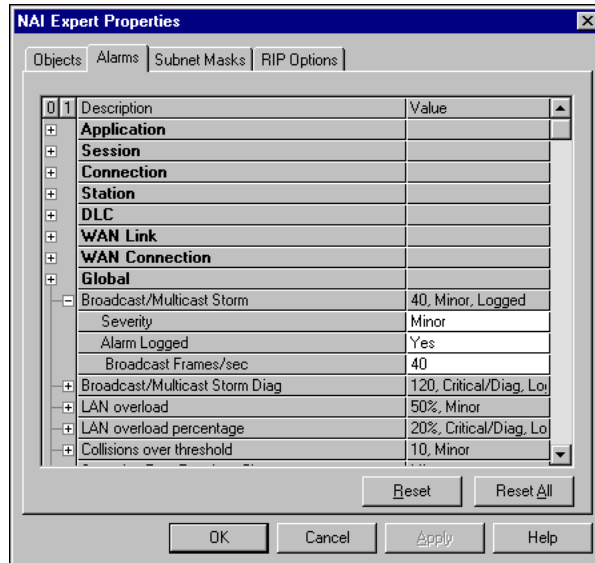
If this network experiences utilization overload, you know who or what is causing the overload. You can then make any necessary adjustments—moving Michael could help.

- **Alarms and alarm thresholds.** Alarms and alarm thresholds alert you to possible problems.

Alarms can alert you to such problems as data link errors, network congestion, too many packets per second, slow applications, application faults, routing problems, service discovery faults, and address allocation errors.

The following shows a sample of alarm threshold settings for Sniffer Pro:

Figure C-3



Some alarms are more critical than others. The following list identifies the more critical alarms:

- **Utilization Percentage.** The percentage of the bandwidth (available roadway) that is used up.

For example, on Ethernet networks, performance degrades significantly when use is above 40%. If utilization is high, you should also consider the collision/fragment error count.

- **Packets Per Second.** The number of packets seen per second on the network. This number can give you an idea of how many packets an interconnecting device (such as a router or switch) needs to process per second.
- **Broadcasts Per Second.** The number of packets addressed to 0xFF-FF-FF-FF-FF-FF on your network.
Broadcasts must be processed by all devices (regardless of their OS or protocol). Excessive broadcasts slow the network.
- **Server/Router Down.** The status of a server or router.
- **MAC-Layer Errors.** The number of errors in per-second increments. MAC layer errors, (defined as layer 1 and layer 2,) corrupt packet formats or make access to the network impossible.

- **Trace Buffer.** The trace buffer is the storage location for all packets you capture.

Typically, the trace buffer is a reserved area of memory in the local system and is limited to the amount of memory in the machine.

When you apply a display filter to the contents of the trace buffer, you do not override the contents of the original buffer. You simply view a subset of the full buffer.

Most analyzers have their trace buffers set up for cyclical operations. The analyzer keeps capturing packets even after the buffer has been completely filled.

The analyzer dumps the oldest packets captured to make room for newer ones (first in, first out).

More sophisticated analyzers have triggers that define when specific actions (such as starting or stopping the capture process or launching a supplemental application) occur.

For example, you can set a trigger to alert you via pager when a Server Down alarm is generated on the analyzer.

Protocol Analyzer Types

You should purchase and implement an analyzer solution based on your network design. A hubbed network requires a different analyzer than a heavily routed network. Switched networks have different requirements.

Protocol analyzer types include the following:

- [Standalone Analyzers](#)
- [Distributed Protocol Analyzers](#)
- [Hardware and Software Analyzers](#)

Standalone Analyzers

A standalone protocol analyzer captures packets that cross the wire the analyzer is connected to.

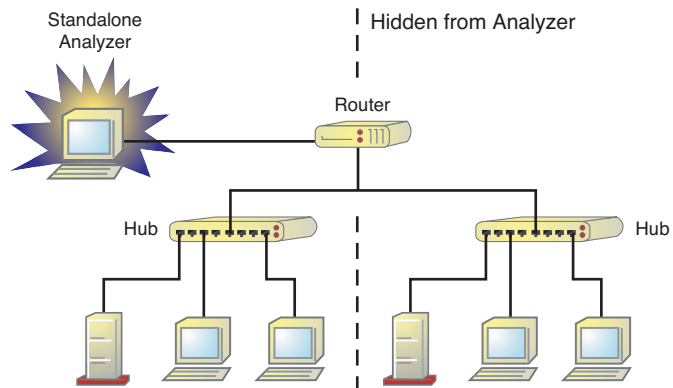
Standalone analyzers are the most popular type of analyzer. You can use a portable analyzer in various locations.

You can set up your own desktop machine as a standalone analyzer and monitor all visible traffic from the comfort of your own desk.

Standalone analyzers such as LANalyzer for Windows (LZFW) can only capture data that is seen locally.

For example, if your LZFW system is connected to a hub with 4 other devices (a router, 2 PCs, and a server), as shown in the following figure, LZFW can see all traffic that flows through that hub. LZFW cannot see traffic on the other side of the router.

Figure C-4



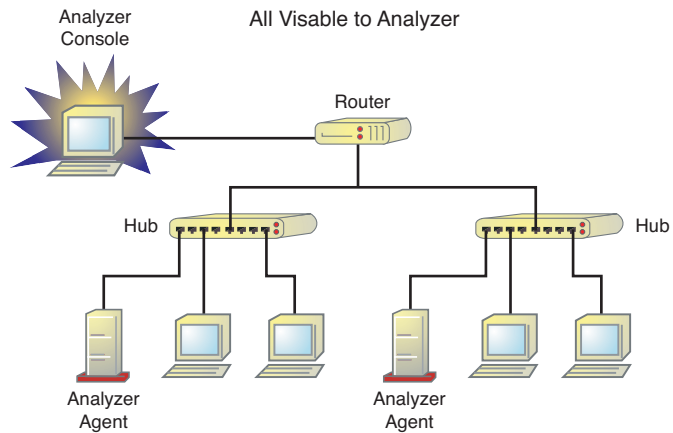
To view traffic on the other side of the router, LZFW would need to be moved to a hub port on the other side of the router.

Distributed Protocol Analyzers

A distributed protocol analyzer captures packets on remote networks. A distributed protocol analyzer uses probes or remote agents to capture the packets on another network and send them to a centralized analyzer station for viewing.

The following illustrates what a distributed analyzer can see based on its configuration:

Figure C-5



Because an analyzer agent has been placed on each server, you can capture data on both sides of the filtering device (router). In this case, you can see all traffic on each side of the router, accounting for the entire network.

On a larger network, you can use the distributed analyzer to troubleshoot multiple buildings or sites from a single location.



Be sure you examine the traffic crossing the network. Make sure you are not causing problems (such as adding extra load onto the network).

Hardware and Software Analyzers

After you make your decision regarding standalone and distributed solutions, you can decide if you want a software solution or a more powerful hardware/software solution.

Software solutions (such as LZFW, Sniffer Basic, and ManageWise®) are typically less expensive than hardware/software solutions. They are also generally easy to install.

Hardware/software solutions (such as Sniffer Pro) are more expensive, but they typically offer greater functionality.

For example, Sniffer Pro includes a specialized network board to ensure greater capability for capturing traffic at higher speeds. Sniffer Pro also includes an Expert System that diagnoses network performance based on traffic patterns.

Protocol Analyzer Placement

The design of your network determines, in part, where you put your analyzer.

The following subsections explain how interconnecting devices such as hubs, bridges, switches, and routers affect the network traffic and how each design should be analyzed:

- [Hubbed Network](#)
- [Bridged Network](#)
- [Switched Network](#)
- [Routed Network](#)
- [WAN Links](#)

Hubbed Network

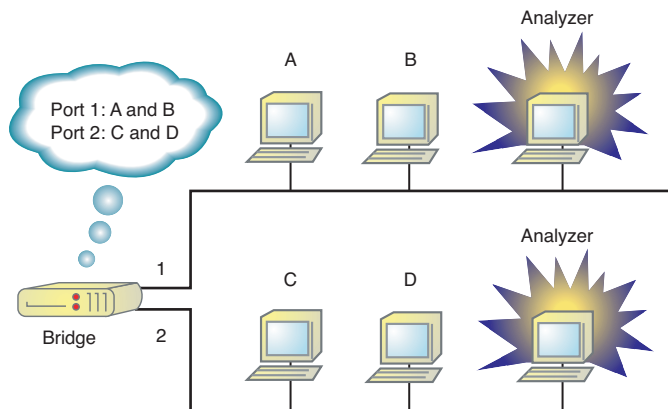
A hubbed network is the simplest network to hook up an analyzer to because all traffic goes everywhere. You can plug the analyzer into any hub to view all traffic from all devices.

Bridged Network

A bridge isolates and localizes traffic based on hardware addresses.

The following illustrates the use of a bridge:

Figure C-6



If a packet from client A is destined for client B, the bridge will not forward the packet because both devices are on the same port (port 1) of the bridge.

However, if client A sends a packet to client C, the bridge forwards the packet to client C's segment because the bridge knows that client C is located off port 2.

If client A sends a broadcast, the bridge forwards the packet because it was destined for all devices. If client A sends a packet to an unknown address, the bridge forwards the packet—bridges always forward packets that they do not have an entry for.

To analyze both sides of a bridge, you must place an analyzer on each side of the bridge.

Switched Network


Switching effectively reduces unnecessary traffic on connected ports.

If you plug your analyzer into a switch port and classify the traffic, you'll see only broadcast and multicast packets and any packets specifically addressed to your analyzer's board address.

On a switch, traffic from one device on a port flows directly to the destination device on another port.

In the following figure, the client and server share a virtual circuit between them. Their traffic flows from port 1 to port 3 and vice versa. Their traffic is not sent to any other ports.

Figure C-7




If you plug your analyzer into the switch, you won't be able to observe the communications because the switch is isolating local conversations.

This natural traffic isolation process only allows you to see broadcast and multicast traffic generated by your host or sent to your host.

You can use the following methods to analyze a switched network:

- **Hub out.** Place a hub between the device of interest (a server, for example) and the switch. Then connect your analyzer to one of the hub ports, as shown:

Figure C-8



You must either connect the switch to the hub's crossover port or use a crossover cable to connect the 2.

The problem with this solution is that you must move the hub/analyzer combination around from device to device if you want to look at multiple devices on the network.

- **Analyzer agents.** Analyzer agents are used by distributed analyzers. These agents are typically software programs loaded on switches to enable them to capture traffic from all ports and send the data to a management console.

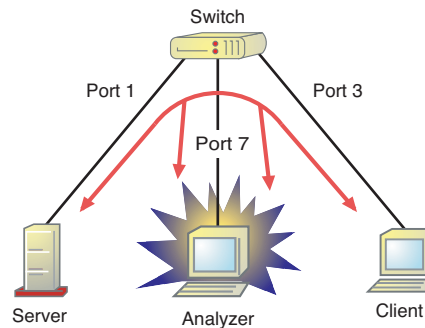
The agents allow you to manage switched traffic from a central location. Unfortunately, this type of feature might make the switch too expensive.

- **Port spanning or mirroring.** Port spanning or mirroring enables you to configure the switch to send a copy of any port's traffic down another port, specifically the port your analyzer is connected to.

This is the most effective way of dealing with switched networks, but it can only be used if the switch supports this functionality.

As shown in the following figure, the traffic from ports 1 and 3 is copied down to Port 7, where they are analyzed:

Figure C-9



Routed Network


Routers isolate traffic based on the network (software) address. If you place an analyzer on one side of a router, you should only see traffic that is sent to that network.

For example, the following figure shows 2 networks (10.1.0.0 and 10.2.0.0, subnetted 255.255.0.0). Traffic between the clients and servers on network 10.2.0.0 is not visible to the analyzer on network 10.1.0.0.

In this case, you have 3 options:

- Place a standalone analyzer on each side of the router.
- Load an analyzer agent on the router (make sure the analyzer agent is a multisegment agent that can capture packets from both connected networks).

- Load an analyzer agent on devices (such as file servers) on each side of the router:

Figure C-10

WAN Links


WAN links typically consist of routing devices at each end of the WAN and a link connecting them. You must place an analyzer or analyzer agent on each side of the WAN routers.

How you place an analyzer on the WAN link depends on the WAN link and the analyzer solution.

The following illustrate 2 possible solutions using Network Associates Sniffer Pro Internetwork Analyzer:


Solution 1: Connect the WAN analyzer to the router and CSU using a Y-cable:

Figure C-11



Solution 2: Connect the WAN analyzer to the CSU monitor port:

Figure C-12



Your WAN analyzer vendor can provide additional details on how to make the physical connection for your WAN link.

The traffic-isolating nature of the devices on your network governs where you place your analyzer. Place the analyzer (or analyzer agent) on each side of any device that might filter traffic.

Index

A

adaptor Setup-30, 3-20, 3-33, 8-7, 8-14–8-15, 8-28–8-32, 8-34–8-36, 9-2, 11-23

address 2-22, 2-30, 7-21, 7-49, 8-16, 10-15, 10-17, 11-33, C-3

administration 2-5, 2-16

administrator Intro-1

agent 4-30
 process status 4-33

alias 2-19

Apache 1-4, 2-7, 7-1–7-9, 7-15–7-17, 7-20–7-21, 7-25, 7-27–7-28, 7-33–7-35, 7-48, 7-54, 11-2, 11-22–11-23, 11-25–11-28, 11-30

asynchronous 6-5

B

background 4-28, 4-34, 4-38–4-40, 4-64, 4-68, 7-40, 11-11

backup 5-24

bandwidth 2-51, 3-3, 5-30, 7-42, 7-47, 7-50, A-5, C-4

binary 3-64

bindery 1-36

block 3-29

bootable Setup-16, 5-13

BorderManager 11-2

bound Setup-18–Setup-19, 1-6, 1-9–1-10, 1-49, 2-18, 2-26, 3-9–3-11, 3-22, 3-71, 7-49, 10-18

C

cable Setup-46, Setup-51, 3-24, 3-31–3-33, 3-76, 3-83, 4-11, 5-21, 8-13, 8-34, 9-42–9-43, 10-28, A-3–A-4, C-12, C-15

cache 3-58–3-59, 3-67, 3-69

casting off 8-17, 8-25, 9-43

class Setup-49, Setup-51, Intro-4, 2-54–2-55, 3-1, 3-6, 3-66, 4-4–4-5, 4-7, 4-49, 4-56, 4-68–4-69, 8-38, 11-13, 12-4

client 2-25

cluster Intro-3–Intro-4, Intro-8–Intro-9, Intro-11, 1-17, 1-42, 7-22, 8-1–8-3, 8-5, 8-7–8-30, 8-32–8-33, 8-39–9-49, 10-1–10-12, 10-14–11-16, 11-18–11-25, 11-27–11-33, 12-3

cluster-enable 8-7, 9-1, 9-3, 9-26, 10-1, 10-3–10-4, 10-6–10-10, 10-12, 10-15, 10-29, 11-1–11-3, 11-5, 11-9, 11-13–11-15, 11-20, 11-22, 11-30–11-31

clustering Setup-5, Setup-46, Setup-52, 1-17, 3-7, 5-11, 5-29, 6-17, 8-1–8-3, 8-5, 8-7–8-10, 8-15, 8-21, 8-24, 8-30, 8-32, 8-34, 8-38, 9-3, 9-5, 9-7, 9-11–9-12, 9-14, 9-16, 9-19, 9-25–9-29, 9-31, 9-34, 9-38–9-39, 9-41, 10-2–10-3, 10-10, 10-14, 10-19, 10-25, 11-2, 11-4, 11-6, 11-8, 11-10–11-12, 11-15, 11-17, 12-3

cluster-naive 11-2–11-3, 11-30

compatibility Setup-51, 1-66, 3-1

- component Setup-10, Setup-25, Setup-28, Setup-31, Setup-34, Setup-48, 1-62–1-63, 1-66, 1-69, 2-2–2-3, 2-8, 2-60, 3-2, 3-4, 3-8, 3-15, 3-18, 3-62, 3-74–3-75, 5-11, 5-30, 5-32, 5-35, 7-22, 8-4, 8-7, 8-12, 8-32, 8-35, 10-12, 12-1–12-2, 12-5, A-1, A-6–B-1
- compressed 1-8, 3-61, B-3
- compression 1-45, 3-58, 3-64, 5-10, 5-15, 5-22
- concurrent 7-47
- configuration Setup-17, Setup-22, Setup-25–Setup-26, Setup-28, Setup-31, Setup-34, Setup-38–Setup-39, Setup-44–Setup-45, Intro-3–Intro-4, 1-29–1-31, 1-33–1-34, 1-36, 1-48, 1-55–1-56, 1-60, 1-62, 1-68, 2-2, 2-4–2-5, 2-13, 2-17–2-21, 2-26–2-27, 2-30, 2-40, 2-59, 2-63, 3-6, 3-10, 3-23, 3-25, 3-37, 3-58, 3-62, 3-74, 4-20, 4-22–4-24, 4-29–4-30, 4-39–4-40, 4-42, 4-45–4-46, 4-66, 6-6–6-7, 6-9, 6-13, 6-22–7-3, 7-5, 7-8–7-10, 7-12, 7-18–7-20, 7-23, 7-25, 7-27, 7-29, 7-54, 8-1, 8-7, 8-9–8-10, 8-13–8-15, 8-25, 8-27–8-28, 8-30, 8-32–8-36, 8-41, 9-4, 9-13–9-14, 9-16, 9-18–9-20, 9-23, 9-25, 9-31–9-32, 9-46, 10-5, 10-20, 11-1, 11-4, 11-10, 11-12, 11-17–11-18, 11-20, C-8
- configure Setup-9, Setup-11, Setup-16–Setup-17, Setup-32, Setup-36–Setup-39, Setup-44–Setup-45, Intro-3, 1-9, 1-43–1-47, 2-19, 2-26, 2-54, 4-23, 4-25–4-27, 4-39, 4-60, 4-62, 5-9, 5-11, 5-22, 6-1, 6-6, 6-10, 6-15–6-16, 7-2, 7-9–7-11, 7-21, 7-23, 7-36, 7-53, 8-7–8-9, 8-12, 8-14, 8-23–8-24, 8-30, 8-33, 8-36, 9-7, 9-14–9-15, 9-17, 9-25, 9-46, 10-1–10-3, 10-29, 11-1, 11-6, 11-9–11-10, 11-13, 11-31, C-12
- connection 2-29
- ConsoleOne 2-4
- container 1-64, 4-18, 4-20, 5-17, 7-30, 9-28
- context 8-33
- conversion utility 5-15
- core dum 3-50–3-52, 3-57–3-64, 3-66–3-70, 3-81, 5-11–5-12
- create Setup-8, Setup-12–Setup-14, Setup-16, Setup-20, Setup-26, Setup-30, Setup-37–Setup-40, Setup-52, 1-2, 1-8, 1-19, 1-22–1-23, 1-39, 1-43–1-45, 1-53–1-54, 3-57, 3-59, 3-66–3-68, 3-70, 3-77, 3-81, 3-83–3-84, 4-13, 4-17, 4-23, 5-2–5-3, 5-5, 5-11–5-12, 5-15–5-16, 5-18, 5-21–5-22, 6-6–6-7, 6-10–6-11, 6-14–6-20, 6-22, 7-8, 7-26–7-27, 7-30, 7-41, 7-52, 8-16, 8-19, 8-32, 8-38, 9-5, 9-7, 9-10, 9-15, 9-22, 9-27, 9-30, 9-47, 10-1–10-6, 10-8, 10-10, 10-12–10-15, 10-29, 11-3–11-5, 11-15, 11-22–11-24, 11-30–11-31, 12-1, A-3, C-3
- cursor 4-23, 7-29, 7-31
- ## D
- deactivate 5-11, 5-24
- debug 1-9, 2-12, 2-14, 2-17, 2-24–2-25, 2-60, 3-26, 3-42, 3-46, 3-54

decompression 1-8
 device 8-10, 8-12, 8-27, 9-2, 9-46, C-5, C-8
 DHCP Setup-9, Setup-31,
 Setup-36–Setup-39, Setup-49,
 Setup-52, Intro-8–Intro-9, 2-5, 2-30,
 2-53, 8-12, 8-17, 8-24, 9-15,
 11-1–11-2, 11-13–11-20, 11-30,
 12-4, C-2
 diagnose 2-18, 3-24–3-25, 3-28, 3-30, 3-40,
 3-42, 3-62, 4-11, 4-13, 4-22
 directory Setup-12–Setup-14, Setup-17,
 Setup-23, Setup-26–Setup-28,
 Setup-33–Setup-34, Setup-42,
 Setup-52–Intro-1, 1-4–1-8, 1-21,
 1-23, 1-25–1-26, 1-33, 1-40, 1-42,
 1-45, 1-51, 1-57, 1-59–1-63, 1-65,
 2-52, 3-10, 3-29, 3-33, 3-64,
 3-68–3-69, 3-71, 4-5, 4-28, 4-43,
 5-10, 5-17, 7-5–7-8, 7-14, 7-18,
 7-23, 7-27–7-28, 7-40–7-42,
 7-44–7-45, 7-51–7-52, 9-27, 10-10
 dismounting 2-13
 DNS Setup-31, Setup-36–Setup-42, Setup-47,
 Intro-8–Intro-9, 2-5, 2-17–2-19,
 2-26, 2-29, 2-36, 7-18–7-21,
 7-33–7-36
 Domain
 Name Service 1-46
 driver 5-21

E

encrypted 7-40, 7-42, 7-53, B-3
 entry 2-31
 epoch 8-17–8-20, 8-22, 8-25, 8-41, 9-35,
 9-38–9-39, 9-42, 9-44
 export Setup-42
 external Setup-2, Setup-44, Setup-46,
 1-39–1-40, 1-59, 1-63, 1-69, 3-27,
 4-23, 4-29, 4-34, 4-64, 4-66, 5-21,
 8-7, 8-10, 8-14–8-15, 8-29–8-30,
 8-32, 9-2, 9-9, 9-14, A-3

F

failback 11-7
 failover 11-6–11-7
 FAT 5-12–5-13, 5-34
 FCS 8-11
 file
 system Setup-21, 1-4–1-8, 1-18, 1-27,
 1-46, 1-55, 3-29, 3-32, 3-48,
 3-53, 3-58, 5-1, 5-3, 5-10,
 5-23–5-24, 5-32, 5-35, 6-3,
 6-6–6-7, 6-9, 6-14, 6-17, 7-43,
 10-6, 12-2
 flush 1-41, 5-10, 10-6

G

generate 2-9–2-11, 2-60, 3-23, 3-28, 3-71,
 4-12, 4-21–4-22, 4-42, 5-25, 8-26
 global Setup-38–Setup-39, 3-44–3-45, 7-3,
 8-27, A-1
 unique ID 8-27
 GUI Setup-17, 1-59, 3-29, 8-36
 GUID 8-27

H

HAM 1-65, 3-24, 5-21
 hardware Setup-1–Setup-3, Setup-46, Intro-1,
 1-1–1-2, 1-4, 1-7, 1-18, 1-25, 1-42,
 2-2, 2-18, 2-49, 2-62, 3-2–3-3,
 3-5–3-6, 3-8–3-9, 3-14, 3-18–3-20,
 3-24, 3-33, 3-47, 3-49–3-50, 3-53,
 3-56, 3-66, 3-74–3-75, 3-79,
 4-15–4-16, 5-12, 6-2–6-3, 6-5–6-6,
 6-22, 8-2, 8-7–8-8, 8-11–8-12,
 8-14–8-15, 8-24, 8-30, 8-36, 8-40,
 9-2–9-3, 9-46, 11-6, 11-23–11-24,
 12-2–12-3, A-2, A-5, B-6, C-8–C-10
 header 3-43

- health Setup-23, Setup-27, Intro-11, 1-33, 1-57, 2-13–2-14, 2-47, 2-60, 3-28, 3-31, 3-40–3-42, 3-45–3-46, 4-2, 4-9–4-10, 4-22–4-23, 4-28, 4-30, 4-33–4-36, 4-44, 4-46–4-49, 4-64, 4-66, 8-15, 8-22, 9-1, 9-39, C-3
- check Setup-27, 4-9, 4-28, 4-30, 4-36, 4-44, 4-47–4-49
- hexadecimal 2-38, 3-30, 3-83, 4-4, C-2
- high availability Intro-1, Intro-3, Intro-8, 8-2–8-8, 8-10, 8-21, 10-1–10-2, 10-10, 10-23, 10-26, 10-28, 11-1, 11-20
- hostname 2-21
- HTTP
port 2-43
- I**
- iFolder Setup-25–Setup-26, Setup-31, Setup-52, Intro-3, Intro-8–Intro-9, Intro-11, 1-4, 1-42, 1-67, 1-69, 2-9, 3-5, 7-1–7-55, 8-16, 8-24, 9-15, 11-1–11-2, 11-20–11-29, 11-32–11-33, 12-4
- iManager Setup-26, Setup-28–Setup-29, Setup-34–Setup-36, 1-4, 1-10, 1-13, 1-64, 1-67, 2-5, 2-7–2-8, 2-14–2-16, 2-60, 4-17, 9-11, 11-2
- Internet
protocol Setup-12, 2-21, 11-17, 11-19
- interrupt 3-23
- interval 2-39–2-40, 4-39, 5-35
- IP
address 7-21, 8-16, 10-15, 10-17, 11-33
- J**
- Java Setup-29, Setup-34, 2-7
- JReport Setup-16, Setup-26, Setup-28, Setup-33, 1-44, 1-60, 1-62, 7-24
- L**
- LAN Setup-27, Setup-44–Setup-45, Setup-50–Intro-1, Intro-4, Intro-8–Intro-9, Intro-11, 1-2, 1-6, 1-24, 1-61, 1-65, 1-69, 2-1–2-2, 2-18, 2-45, 3-19–3-20, 3-23–3-24, 3-32–3-33, 3-78, 3-84, 4-8, 4-37, 4-52, 5-29–5-30, 5-32, 7-24, 8-17–8-22, 8-24–8-25, 8-33, 9-14, 9-17–9-18, 9-43, 9-45, 11-13, A-1
- LBURP Setup-43
- LDIF Setup-3, Setup-42–Setup-43
- limber 4-29
- list 5-10, 10-17
- LOAD Setup-13, Setup-18, Setup-20, Setup-27, Setup-33, Setup-36, Setup-41, Setup-51–Setup-52, 1-7, 1-29–1-30, 1-45, 1-48–1-49, 1-61, 1-64–1-65, 2-19, 2-35–2-36, 3-3, 3-6, 3-9–3-11, 3-16–3-17, 3-32, 3-40–3-42, 3-44–3-45, 3-64, 3-66–3-68, 3-71, 3-79, 5-8, 5-12–5-13, 5-16, 5-20, 5-34, 6-3, 6-11, 8-9, 8-11, 8-24, 8-26, 8-37, 9-16, 9-20–9-21, 9-23–9-24, 9-33, 9-38, 9-45, 10-5, 10-9, 10-15, 10-17–10-18, 10-20, 10-22, 10-30–11-1, 11-5, 11-8–11-10, 11-15–11-19, 11-22, 11-24–11-25, 11-28, 11-31, C-1, C-8, C-13–C-14
- load
balancing 2-35
location Setup-11, Setup-49, 7-5, 7-7
log 2-24, 4-10
log file 5-17

- logical 5-2-5-3, 5-12, 5-14-5-17, 5-22,
5-26-5-27, 5-34-5-35, 6-7,
6-14-6-15, 6-21, 7-20, 10-1, 10-3,
10-14-10-15, 10-29, A-3, A-5,
B-3-B-4
volumes 5-22, 10-3, 10-14
- login 1-26
- long name Setup-20, 1-7, 1-49, 5-15
- M**
- management Setup-7, Setup-27, Setup-33,
Setup-36-Setup-40, Intro-1, Intro-3,
1-59, 1-64, 1-66, 2-2-2-5, 2-8-2-9,
2-12-2-16, 2-60, 3-1, 3-9, 3-36,
3-38-3-39, 3-74, 4-17, 4-19, 5-1,
5-5, 5-7-5-8, 5-18, 5-26, 5-31, 5-33,
6-3, 6-10, 7-6, 7-8-7-12, 7-14,
7-16-7-20, 7-22, 7-24-7-25,
7-28-7-32, 7-36, 7-38-7-39, 7-48,
7-50, 8-3-8-4, 8-8, 8-39, 9-19, 9-32,
9-34, 10-19, 10-25, 10-28, 11-16,
11-20, 11-23, 12-4, C-12
- master Setup-52, Intro-9, 1-17, 1-37, 4-16,
4-19, 4-45-4-47, 4-49, 4-67, 5-30,
7-41, 7-44, 8-15-8-22, 8-41, 9-10,
9-14, 9-16-9-18, 9-23-9-24, 9-30,
9-32-9-33, 9-36, 9-38, 9-42-9-43,
9-45, 9-48, 11-4, 11-8
- Mean Time Between Failures 8-3, 8-40
- Mean Time To Recovery 8-4, 8-40
- media 5-6, 5-18-5-19, 5-22, 5-28, 6-7-6-8,
6-11, 6-14-6-15, 6-17, 6-20-6-21,
8-38-8-39, 9-41, 10-3, 10-13-10-14,
11-24
- memory 3-7-3-9, 3-12, 3-16, 3-18,
3-26-3-30, 3-32, 3-40-3-43,
3-49-3-53, 3-58-3-61, 3-63, 3-67,
3-69, 3-75, 3-79, 3-81, 5-21, 8-7, 9-2,
9-46, C-5
- MFL 5-10
- migrate Setup-6-Setup-7, Setup-13, Setup-17,
Setup-21-Setup-24, Setup-51,
Intro-2, 1-1, 1-7, 1-9, 1-18, 1-23,
1-25, 1-27, 1-29, 1-31, 1-35, 1-40,
1-42-1-43, 1-55-1-58, 1-68, 4-6,
8-3, 8-16-8-17, 8-25, 9-3,
9-33-9-34, 9-49, 10-10,
10-18-10-21, 10-24, 10-27, 10-30,
11-5, 11-9, 11-18-11-20, 11-28,
11-31, 11-33
- migrating Setup-24, Intro-10, 1-6-1-9, 1-20,
1-25, 1-39, 1-42, 1-58, 1-67, 2-5, 8-1,
8-11, 8-26-8-27, 9-23, 10-18,
10-23-10-26, 10-28, 10-30, 11-12,
11-17, 11-19
- Modified File List 5-10
- monitor Setup-32, Intro-3, 2-8, 2-12-2-14,
2-22, 2-60, 3-19, 3-28-3-29,
3-31-3-33, 3-40-3-42, 3-45-3-46,
3-54, 3-73, 3-75, 4-1-4-2, 4-6, 4-14,
5-25, 8-17, 8-20, 9-1, 9-12, 9-21,
9-30, 9-32, 9-34, 9-43-9-45, 9-49,
10-10, C-6, C-15
- MTBF 6-2, 8-3-8-4, 8-40, 12-2
- MTTR 8-4-8-5, 8-40
- N**
- navigation Setup-36, 3-11, 3-24, 4-40, 4-45
- NCS Intro-1, Intro-3, Intro-8, Intro-11, 1-42,
8-1-8-3, 8-7-8-13, 8-15-8-17,
8-19-8-25, 8-27-8-32, 8-36, 8-39,
8-41, 9-1-9-16, 9-18-9-19,
9-23-9-27, 9-29-9-33, 9-39, 9-41,
9-43, 9-45-9-48, 10-1-10-2, 10-6,
10-8-10-10, 10-15, 10-17-10-18,
10-20, 10-28-11-2, 11-5, 11-7,
11-10-11-12, 11-15, 11-17, 11-28,
11-30-11-32

- NDS Setup-10, Setup-15, Setup-17, Setup-20, Setup-22–Setup-25, Setup-30–Setup-31, Setup-42, Setup-48, 1-21, 1-31, 1-35, 1-37, 1-41, 1-47, 1-49–1-54, 1-56–1-58, 1-63, 2-12, 2-44–2-45, 2-62, 4-3–4-4, 4-7, 4-20, 4-26, 4-30, 4-45, 4-61, 4-67–4-69, 10-20, 10-22, 10-30, 11-2, 11-30
 - NetWare 3-38, 3-62
 - NetWare 6 Intro-1, 1-1, 3-38, 4-26
 - NetWare Migration Wizard Setup-20, Setup-24, 1-19, 1-53, 1-58
 - NetWare Remote Manager 3-54, 3-58
 - NetWare Web Manager Setup-25, 2-9, 4-17, 10-25
 - network Setup-50, 2-51, 3-69, 4-27, 8-11, 8-41, A-1, B-5, C-13
 - NFAP 11-3
 - NFS 11-3, 11-13
 - NLM Setup-41, Setup-52, 1-26, 1-30, 1-33, 1-56, 1-67, 2-22, 2-24, 3-9–3-11, 3-18, 3-23, 3-25–3-29, 3-32, 3-37–3-47, 3-51–3-52, 3-56, 3-58, 3-61–3-62, 3-64, 3-66–3-68, 3-71–3-72, 3-77, 3-83, 5-16–5-17, 5-34, 7-5, 7-35, 7-41, 9-12
 - node Intro-3–Intro-4, Intro-11, 2-18, 3-43, 3-71, 8-1–8-2, 8-7–8-34, 8-40–9-3, 9-5, 9-7–9-8, 9-10–9-12, 9-14–9-18, 9-20–9-27, 9-30–9-33, 9-35–9-36, 9-38–9-39, 9-41–9-48, 10-3, 10-9–10-10, 10-13, 10-17–10-18, 10-20, 10-22–10-25, 10-28, 11-1–11-2, 11-4–11-10, 11-12–11-13, 11-15–11-16, 11-18, 11-20, 11-22, 11-25, 11-27, 11-29, 11-31–11-33
 - Notes 8-18
 - Novell
 - Novell Cluster Services Intro-3, Intro-8, 1-42, 8-1, 9-2, 9-12, 11-11–11-13
 - Novell Modular Authentication Service Setup-28, Setup-34, 1-62
 - Novell Storage Services Intro-8, 1-42, 5-1, 5-26
 - Novell Technical Support 3-77, 3-84
 - NSS Intro-3, Intro-8, 1-6, 1-39–1-40, 1-42, 1-69, 3-8, 3-28–3-29, 3-38, 3-58, 5-1–5-2, 5-4–5-9, 5-11–5-28, 5-34–5-35, 6-1–6-3, 6-6–6-7, 6-10, 6-14–6-22, 8-12, 8-16, 8-28, 8-38–8-39, 9-10, 10-1, 10-3–10-7, 10-9–10-15, 10-29, 11-24–11-25
 - NWPA 3-27
- O**
- object 1-16, 8-38, 9-14, 9-48
 - operating system 3-49
 - options 4-21
 - outage 8-3
- P**
- packets 2-48
 - parameters Setup-32
 - partition 9-41, 10-3
 - path 3-67
 - physical 2-38, 3-49, 6-9, 6-13, 6-16, 6-19, 7-53, 8-3, 8-7, 8-10, 8-40, 12-2, A-3, B-4, B-6, C-15
 - ping 2-27
 - platform support module 3-34
 - pool 5-11, 10-9, 10-15–10-16, 10-18, 11-24
 - port Setup-42, 2-4–2-5, 2-9, 2-34, 2-40–2-44, 2-52–2-54, 7-5–7-9, 7-20, 7-32–7-33, 7-35, 7-42, 7-48–7-49, 7-53, 9-19, 9-32, 9-34, A-5, C-1, C-7, C-10–C-13, C-15

- post-installation 1-41, 1-65–1-66, 1-69, 9-6, 9-27, 11-22
 - post-migration 1-39, 4-6
 - pre-migration Setup-6, Setup-15–Setup-17, 1-9, 1-43–1-44
 - print
 - job 4-4
 - printer Setup-40, 1-12, 1-16, 8-3
 - processes 3-32
 - processor 2-51, 3-6, 3-11–3-12, 3-14–3-18, 3-20, 3-34–3-35, 3-49–3-50, 3-56, 3-79, 8-8
 - property Setup-12, 1-45, 2-29, 4-7, 5-5–5-6, 5-18–5-19, 5-22, 5-26–5-28, 6-7–6-8, 6-11, 6-14–6-15, 6-17, 6-20–6-21, 7-31, 7-45, 8-38–8-39, 9-14–9-15, 9-17, 9-19–9-21, 9-23–9-25, 9-31–9-33, 9-41, 9-48, 10-2–10-3, 10-8, 10-13–10-14, 10-16–10-18, 10-22, 10-29, 11-1, 11-4–11-5, 11-7, 11-15, 11-17, 11-19, 11-24
 - protocol Setup-12, Setup-45, 1-2, 1-6, 1-46, 2-20–2-21, 2-45, 3-71, 4-22, 5-30, 5-32, 10-5, 11-2, 11-17, 11-19, A-2
 - PSM Setup-27, 1-61, 1-65, 3-6, 3-34, 7-24
 - purge 3-29, 3-33, 5-27–5-28
- R**
- RAID Intro-3, 3-5, 5-29–5-30, 6-1–6-15, 6-19–6-20, 6-22, 8-9, 9-4, 9-46, 12-3
 - RAM 3-7
 - recovery 8-4, 8-40
 - Remote Manager 3-54, 3-58, 4-11, 6-10
 - replica
 - ring Setup-13, 3-71, 4-33, 4-45, 4-66
 - synchronization 4-5, 4-33, 4-47–4-48, 4-61, 4-66, 4-68
 - reports 2-4, 3-42, 3-74, 4-20
 - requirements 8-9
 - resource Setup-40, 3-21, 3-39, 3-74, 8-3, 8-8, 8-10–8-11, 8-16–8-17, 8-20, 8-23–8-25, 8-40–8-41, 9-3, 9-14–9-15, 9-19–9-24, 9-31, 9-33, 9-36, 9-38, 9-42, 9-45–9-46, 9-48, 10-1, 10-5–10-9, 10-15–10-28, 10-30, 11-2–11-10, 11-13, 11-15–11-19, 11-22, 11-24, 11-27–11-28, 11-30–11-31, 11-33
 - revision 2-61, 4-28, 4-64
 - router 2-35, C-13
- S**
- SBD 8-17–8-20, 8-22, 8-25, 8-34, 8-38, 8-41, 9-4, 9-10, 9-37–9-39, 9-41–9-42, 9-45–9-46, 9-49, 10-10
 - SCSI Setup-2, Setup-5, Setup-16, Setup-30, Setup-46, Setup-52, 1-43–1-44, 3-5, 3-25, 5-30, 5-33, 8-1, 8-7–8-8, 8-12, 8-14–8-15, 8-27–8-39, 8-41–9-2, 9-14, 9-25–9-26, 9-29, 9-39–9-42, 9-45, 10-2, 10-10–10-13, 10-16, 10-26, 10-28, 11-23–11-24, 11-26, 11-32
 - ID Setup-52, 8-12, 8-29, 8-35–8-36, 9-2
 - secondary Setup-52, 1-9, 1-46, 2-39, 3-34–3-35, 7-20–7-21, 7-32–7-33, 7-49, 8-16, 9-3, 9-43, 9-45, 11-13, 11-21–11-23, 11-25–11-26, 12-3
 - security Setup-7, 2-8, 2-10, 4-23, 5-7, 5-10, 7-24, 7-26, 7-48, 10-19, 10-25, 11-16
 - server Setup-2–Setup-3, Setup-5–Setup-6, Setup-8–Setup-9, Setup-11–Setup-23, Setup-25–Setup-34, Setup-36–Setup-44, Setup-46–Setup-47, Setup-49, Setup-51–Setup-52, Intro-3, Intro-9–Intro-11, 1-2–1-16, 1-18, 1-21, 1-23–1-27, 1-29–1-37, 1-39–1-68, 2-1–2-2, 2-4–2-9,

- 2-12-2-25, 2-27, 2-29-2-30, 2-36, 2-39, 2-43-2-44, 2-46-2-54, 2-60, 2-62, 3-1-3-11, 3-14-3-15, 3-18-3-35, 3-37-3-48, 3-50-3-62, 3-64-3-72, 3-74-3-81, 3-83, 4-3, 4-5, 4-7-4-11, 4-14-4-15, 4-17-4-18, 4-20-4-23, 4-26-4-27, 4-29-4-30, 4-32-4-33, 4-39, 4-41-4-43, 4-45, 4-47-4-48, 4-52, 4-59, 4-65-4-66, 4-68-4-69, 5-2-5-6, 5-8, 5-11-5-16, 5-18-5-21, 5-23, 5-25-5-26, 5-28-5-29, 5-31-5-34, 6-1-6-3, 6-7-6-8, 6-10-6-11, 6-14-6-15, 6-17, 6-21, 7-1-7-25, 7-27-7-44, 7-46-7-55, 8-1, 8-3, 8-7-8-13, 8-15-8-18, 8-20, 8-24, 8-26-8-28, 8-32-8-33, 8-35-8-38, 8-41, 9-2-9-4, 9-7, 9-9-9-12, 9-14, 9-16, 9-21, 9-24-9-35, 9-38-9-46, 9-48-9-49, 10-2-10-9, 10-11-10-19, 10-21, 10-23-10-30, 11-2-11-3, 11-5-11-9, 11-13-11-28, 11-30, 11-33, 12-1, 12-3-12-4, A-2, A-4, B-6, C-5, C-7-C-8, C-11-C-12
- management Setup-37-Setup-39, 2-12-2-13, 2-16, 7-6, 7-8-7-12, 7-14, 7-16-7-20, 7-24-7-25, 7-28-7-32, 7-36, 7-38-7-39, 7-48, 7-50
- node 9-12, 9-24, 9-32-9-33, 9-48
- object 8-38
- server-to-server 3-71, 3-76, 3-83
- service Intro-10, 1-10, 8-17, 9-3
- settings Setup-12, 1-2, 7-35, 10-20
- setup Setup-2
- shared
 - disk partition 10-2
 - storage device 8-27
- size 5-18, 6-8, 6-12, 6-16, 6-19
- slave 8-15, 8-17, 8-19-8-21, 9-18, 9-32, 9-43
- SLP 1-6, 4-23
- SMDR Setup-21, 1-55, 1-66
- SMS 1-18, 1-66, 1-69
- SNMP 1-62, 2-20-2-22, 3-38
- software Setup-1, Setup-3-Setup-4, Setup-6, Setup-18, Setup-26, Setup-28, Setup-33, 1-2, 1-10, 1-20, 1-53, 1-60, 1-62, 2-2, 2-49, 2-53-2-54, 2-62, 3-2, 3-5, 3-8, 3-12, 3-14, 3-16, 3-29, 3-38-3-39, 3-47-3-49, 3-53, 3-59, 3-74, 3-79, 3-81, 4-15-4-16, 4-28, 5-25, 6-2-6-3, 6-6-6-7, 6-10-6-11, 6-15, 6-19, 6-22, 7-24, 8-7-8-8, 8-10, 8-39-8-40, 9-2-9-3, 9-5, 9-7, 9-11, 9-25-9-26, 9-46-9-47, 11-6, A-2, B-5-B-6, C-8-C-9, C-12-C-13
- source
 - server Setup-23, 1-57
- space Setup-13, 11-25
- SSL Setup-42, 7-5-7-7, 7-49
- standalone C-6, C-8, C-13
- start Setup-4, Setup-11-Setup-13, Setup-15, Setup-18, Setup-20-Setup-21, Setup-24, Setup-29-Setup-30, Setup-34-Setup-35, Setup-38-Setup-39, Setup-42, 1-19, 1-50, 1-52-1-53, 1-58, 2-6, 2-8, 2-26, 2-36, 2-45, 3-16, 3-30, 3-34-3-35, 3-61, 3-69, 4-10, 4-14, 5-28, 6-7, 6-11, 6-17, 6-20, 7-2, 7-14, 7-16, 7-18, 7-25, 7-27-7-28, 7-35, 8-31, 8-36, 8-38, 9-6, 9-11, 9-15-9-16, 9-20, 9-23-9-24, 9-27, 9-29-9-31, 9-41, 9-48, 10-2, 10-16, 10-21, 10-23, 10-26-10-28, 11-4, 11-6-11-8, 11-17, 11-19, 11-23, 11-31
- Start Mode 11-6
- startup directory 1-65

- storage Setup-13, Setup-25, Setup-27, Intro-1, Intro-3, Intro-8, 1-42, 1-59, 1-61, 1-66, 3-5, 3-7, 3-19, 3-32, 3-38, 3-59, 3-61, 3-79, 5-1-5-6, 5-8, 5-10-5-11, 5-18, 5-20, 5-22, 5-24-5-27, 5-29-5-35, 6-1, 6-4-6-5, 6-8-6-9, 6-12-6-17, 6-22, 7-23-7-24, 8-2, 8-8-8-11, 8-13-8-14, 8-16-8-18, 8-21, 8-27-8-28, 8-36, 8-40-8-41, 9-37, 10-1-10-4, 10-10, 10-14, 10-29, 11-20, 11-24, 12-3, C-5
- area network 8-11, 8-41
 - device 8-27
- Storage Management Services 1-66
- stripe 3-24, 6-1, 6-6, 6-8, 6-11, 6-13
- subcontainer 7-10-7-11, 7-30-7-31, 7-37
- subnet Setup-9, Setup-12, Setup-17, Setup-31, Setup-37-Setup-38, Setup-51, 1-46, 2-26, 2-32, 2-45, 2-53-2-56, 3-22, 9-3, 9-46, B-5-B-6
- Support Pack Setup-1, Setup-3, Setup-8, Setup-26, Setup-32, 1-6-1-7, 1-16, 1-60-1-61, 2-3, 2-44, 2-60-2-61, 7-23-7-24
- synchronization 4-45, 4-64
- SYS Setup-17, Setup-21, Setup-24, Setup-29, Setup-34-Setup-35, Setup-42, Setup-52, 1-24-1-26, 1-30, 1-40-1-41, 1-45, 1-54, 1-56, 2-6-2-7, 2-9, 2-19, 2-24, 3-27, 3-40, 3-44-3-45, 3-50, 3-65, 3-68, 3-74, 3-77, 3-84, 4-10, 5-21, 5-24, 7-1-7-9, 7-13-7-18, 7-20-7-21, 7-27-7-28, 7-33-7-35, 8-12, 8-16, 9-2, 9-46, 10-2, 11-22-11-23, 11-25, 11-27
- system Setup-19, Setup-21, Setup-29, Setup-34-Setup-35, Setup-45, Setup-52-Intro-1, Intro-9-Intro-10, 1-2, 1-4-1-8, 1-18, 1-20, 1-24, 1-27, 1-41, 1-46, 1-55, 2-14, 2-20, 2-22, 2-47, 2-50, 2-54, 2-62, 3-2-3-8, 3-11-3-12, 3-14-3-15, 3-17-3-18, 3-23, 3-27-3-29, 3-31-3-33, 3-36, 3-38, 3-40, 3-42, 3-48-3-50, 3-53, 3-58, 3-60, 3-64-3-65, 3-68-3-69, 3-71, 3-73-3-74, 4-1-4-3, 4-10, 4-20, 4-26, 4-30, 4-46, 4-60, 4-66, 4-70, 5-1, 5-3, 5-10-5-12, 5-16, 5-23-5-25, 5-27, 5-32, 5-35, 6-2-6-3, 6-5-6-7, 6-9, 6-14, 6-16-6-17, 6-22, 7-38, 7-43, 8-1, 8-3-8-4, 8-6-8-7, 8-9-8-10, 8-12-8-13, 8-15, 8-25, 8-35, 8-41, 9-2, 9-4, 9-10, 9-31, 9-39, 9-46, 10-6-10-7, 11-33, 12-2-12-4, C-5, C-7, C-9
- ## T
- target
- server 4-68
- TID 1-6, 2-58, 3-27, 3-64, 4-55, 8-18, 8-21, 9-5, 11-11-11-12
- time Setup-1-Setup-2, Setup-6, Setup-9, Setup-18, Setup-32, Setup-43, Setup-50, Setup-52, Intro-10, 1-7-1-10, 1-12, 1-23, 1-31, 1-35, 1-37, 1-46, 1-48-1-49, 2-1, 2-3, 2-23, 2-28, 2-31, 2-33-2-35, 2-41-2-42, 3-12-3-14, 3-18-3-19, 3-25, 3-28, 3-34, 3-38, 3-50-3-56, 3-62-3-63, 3-65-3-66, 3-71, 3-74-3-75, 3-81, 4-1, 4-4-4-6, 4-8-4-10, 4-13-4-14, 4-19,

4-21-4-23, 4-27-4-28, 4-30-4-33,
 4-37, 4-41, 4-45-4-46, 4-61,
 4-64-4-66, 5-10-5-11, 5-24,
 5-27-5-28, 6-2-6-4, 6-6, 6-9, 6-14,
 7-1, 7-11, 7-15-7-16, 7-19, 7-21,
 7-33, 7-41, 7-44-7-45, 7-53,
 8-3-8-5, 8-10, 8-16-8-18, 8-38,
 8-40, 9-5, 9-7, 9-15, 9-17-9-18,
 9-33, 9-35, 9-37, 9-44, 9-47, 10-3,
 11-1, 11-10, 11-13, 12-2-12-4

server Setup-52, 1-35, 1-46

Time Synchronization 4-64

TIMESYNC type Setup-52, 1-9

traditional

 volume 5-15, 5-24, 6-7, 8-39

transaction 3-15-3-16, 5-11

transmission 2-21, 5-30, A-1-A-6, B-4

tree Setup-31, 10-26, 11-27

TSA 1-66

tune 2-51

type Setup-52, 1-9, 1-40, 10-3

U

UAL 1-12-1-16, 1-68

unattended Setup-6, Setup-17-Setup-18, 1-7,
 1-41-1-42, 1-64, 4-37

UNIX 2-3

upgrade Setup-9, Setup-27, Setup-33, Intro-8,
 Intro-10, 1-1-1-4, 1-18, 1-37,
 1-40-1-42, 1-61, 1-63, 1-65, 1-69,
 2-2-2-3, 2-5-2-6, 2-60, 3-20,
 3-42-3-43, 3-73, 4-6, 4-16, 4-26,
 4-30, 4-45, 8-26, 9-2, 9-5-9-7, 9-11,
 9-27, 9-47, 10-28

upload 3-64, 7-40

user

 account 7-28-7-29, 7-50

utilities Setup-29, Setup-34, 11-26

V

value 3-34

version 4-28

video 8-1, 10-14-10-16, 10-27

view 8-15, 9-21, 9-23, 9-34-9-35, 9-42,
 10-17, 10-19-10-20, 11-3, 11-26,
 11-30

volume 1-2, 3-68, 5-22, 7-6, 7-8, 10-3, 10-7,
 10-14, 10-17

 conversion utility 5-15

 name 5-17

W

web

 server 1-4, 2-43-2-44, 7-4-7-5, 7-54,
 11-2, 11-5, 11-13, 11-30

 services Intro-4, 8-3

X

XML 2-9

Z

zone Setup-39-Setup-40, 1-35, 1-46