The IBM TotalStorage NAS Integration Guide

- Learn all about the IBM TotalStorage NAS 100, 200, and 300
- Implement high speed cross platform storage solutions
- Configure NAS using this hands-on guide

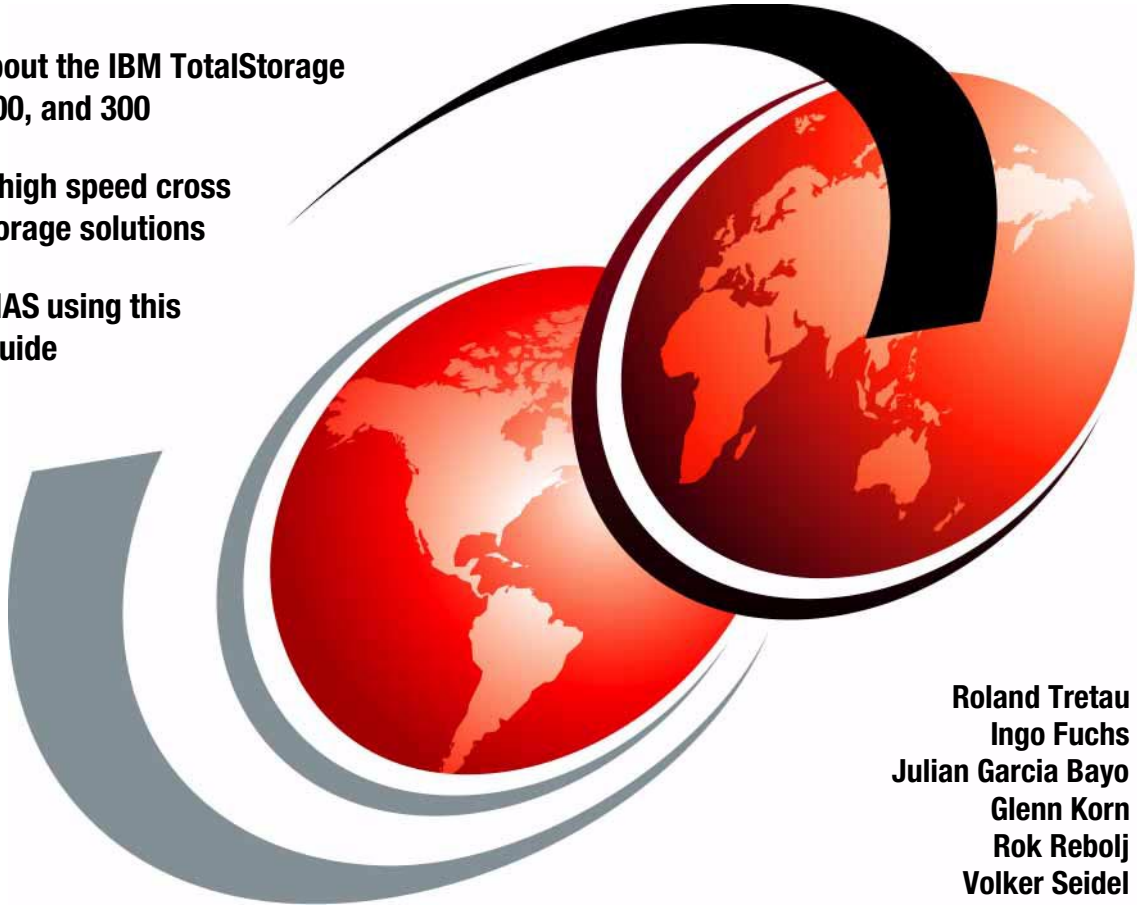

Roland Tretau
Ingo Fuchs
Julian Garcia Bayo
Glenn Korn
Rok Rebolj
Volker Seidel

Redbooks

ibm.com/redbooks

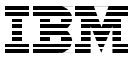

International Technical Support Organization

The IBM TotalStorage NAS Integration Guide

March 2003

SG24-6505-02

Note: Before using this information and the product it supports, read the information in “Notices” on page xxxi.

Third Edition (March 2003)

This edition applies to the IBM TotalStorage Network Attached Storage 100, 200, and 300 running the Windows Powered OS.

Contents

Figures

Tables

Examples

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AFS®
AIX®
AIX 5L™
Approach®
Balance®
DB2®
DB2 Universal Database™
DFS™
Domino™
Enterprise Storage Server™
ESCON®
IBM®
IBM eServer™
Lotus®
Lotus Notes®
Metaphor®
Micro Channel®
Netfinity®
NetView®
Notes®
PAL®
Perform™
PowerPC®
Predictive Failure Analysis®
PS/2®
RACF®
RAMAC®
Redbooks™
Redbooks (logo)™ 
RMF™
SANergy™
Sequent®
ServeRAID™
ServicePac®
SP2®
TCS®
Tivoli®
Tivoli Enterprise™
TotalStorage™
Word Pro®
xSeries™

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

SAS: Portions of this document (Section 14.5 and Appendix E) are Copyright (C) 2001, SAS Institute Inc., Cary, NC, USA. All rights reserved. Reproduced with permission of SAS Institute Inc., Cary, NC, USA.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. (R) indicates USA registration.

Other company, product, and service names may be trademarks or service marks of others.

Preface

This IBM Redbook describes how to integrate, install, and configure the very latest IBM TotalStorage Network Attached Storage 100, 200, and 300 Release 2.5 in heterogeneous environments. Concrete scenarios for IBM DB2 UDB, Lotus Domino Server, Microsoft SQL Server, Microsoft Exchange Server, Oracle, and SAS illustrate the integration process for storage solutions.

The NAS 100, 200, and 300 are innovative Network Attached Storage (NAS) appliances that connect clients and servers on an IP network to storage. Their value is enhanced by their support of multiple protocols, allowing seamless file sharing across dissimilar platforms. They provide excellent Microsoft Windows performance that enhances client productivity while simultaneously protecting a customer's data and business continuity. This book shows how to integrate the units and explains how a company may benefit by utilizing these innovative solutions.

This easy-to-follow guide describes the market segments that may benefit from the NAS 100, 200, and 300, and explains NAS installation, ease-of-use, remote management, expansion capabilities, Microsoft Active Directory integration, high availability (clustering), and backup and recovery techniques. Other concepts, such as cross platform storage and methodologies for common data sharing for Linux/UNIX and Windows NT/2000 environments, are also covered.

This book makes use of the IBM NAS initiative in the marketplace and defines its position and value-add. Also discussed is how the reliability, availability, scalability, and security of the NAS 100, 200, and 300 have the potential to be at the heart of an enterprise's data storage system.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

The NAS 200 and 300 team, from left to right: Julian, Glenn, Rok, Roland, Ingo

Roland Tretau is a Project Leader with the IBM International Technical Support Organization, San Jose Center. Before joining the ITSO in April 2001, Roland worked in Germany as an IT Architect for Cross Platform Solutions and Microsoft Technologies. He holds a Masters degree in Electrical Engineering with a focus in telecommunications. He is a Microsoft Certified Systems Engineer (MCSE) and he also holds a Masters Certificate in Project Management from the George Washington University, School of Business and Public Management.

Ingo Fuchs is a Chief IT Architect and Alliance Manager at Memorex Systemhaus in Frankfurt, Germany. His areas of expertise include Storage Networking (SAN/NAS), Intel-based server systems, and high-availability solutions. Ingo holds a Masters degree in Information Technology with focus in Project Engineering and has a teaching share at the University of Cooperative Education in Mannheim, Germany.

Julian Garcia Bayo is a Storage expert and consultant with IBM Spain in the SSG Group. He holds a Masters degree in Telecommunications Engineering with focus in Electronic Equipment. His areas of expertise include Storage Networking Solutions (NAS/SAN), Microsoft and Linux Operating Systems, and disk subsystems. He is a Microsoft Certified Systems Engineer (MCSE) and an IBM Certified Specialist for xSeries and Enterprise Disk Solutions.

Glenn Korn is a member of the Technical Support Marketing group for the IBM SSG group where he specializes in SAN and NAS solutions. Glenn has been working with computer hardware and networks for over 12 years. He has extensive experience in IT support of both hardware and software. He holds a B.S. degree in Industrial Technology and is a Microsoft Certified Systems Engineer (MCSE) and Novell CNE.

Rok Rebolj is a Systems Engineer and Instructor in Slovenia. He has 8 years of experience in the IT field. He holds a degree in Electronics Engineering from the University in Ljubljana. His areas of expertise include IBM Netfinity and xSeries servers, storage networking, and systems management. He is a Microsoft Certified Systems Engineer (MCSE) and an IBM Certified Expert for xSeries.

Volker Seidel is a consultant for cross platform solutions in Germany. He has four years of experience in planning and implementing IT solutions for companies focusing on Linux/UNIX. He holds a bachelor-of-arts degree in business administration from the Hamburg Academy of Business Management. His areas of expertise include cross platform connectivity for SOLARIS, HP-UX, AIX, and RedHat Linux with Windows based networks. He is also a certified HP 9000 Technical Professional.

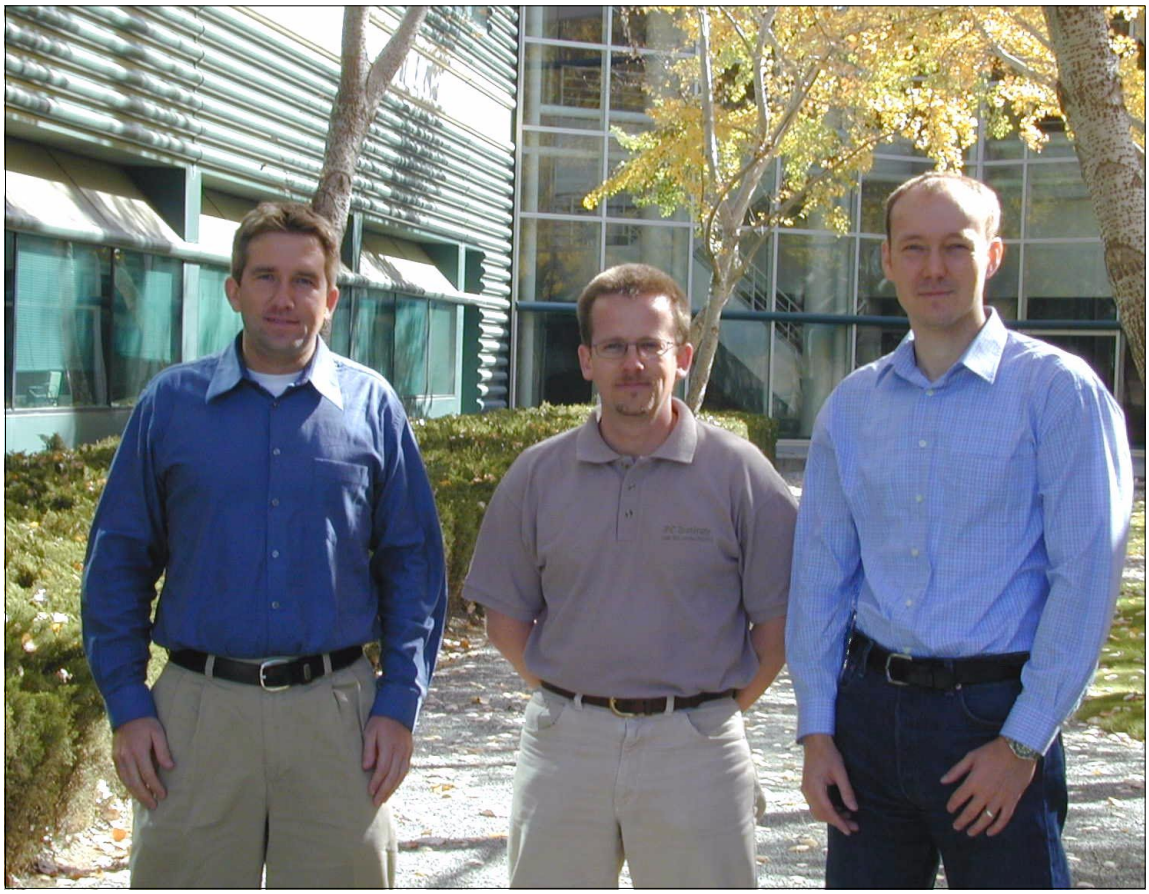

The NAS 100 team, from left to right: Volker, Rok, Roland

Thanks to the following people for their contributions to this project:

Yvonne Lyon, Deanna Polm, Emma Jacobs
International Technical Support Organization, San Jose Center

Scott Hovey, Jeff Ottman, Thomas Daniels, Kevin Goldsmith, Richard Kisley, Douglas Dewey, Larry Gooch, Frank Tutone, Scott Hovey, David Hutton, John B Hartman, Robert C. Davis, Florian Kandefer, James Cybrynski, Daniel Heffley, Bobby Griggs, Jeff Reedy
IBM US

David Worley
LSI Logic Storage Systems, Inc.

Notice

This publication is intended to help network or storage administrators to install and configure the IBM TotalStorage NAS 100, 200, and 300. The information in this publication is not intended as the specification of any programming interfaces that are provided by IBM TotalStorage NAS 200 and 300. See the PUBLICATIONS section of the IBM Programming Announcement for the IBM TotalStorage NAS 100, 200 and 300 for more information about what publications are considered to be product documentation.

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to the address on page ii.

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition may also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-6505-02
for The IBM TotalStorage NAS Integration Guide
as created or updated on April 23, 2003.

March 2003, Third Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

- Information regarding the IBM TotalStorage NAS 100 has been added.

Changed information

- The book has been restructured for easier use.
- The descriptions of how to implement the IBM TotalStorage NAS 100, 200, and 300 are now covered in separate chapters.

Part 1

Network Attached Storage concepts and hardware

In Part 1 of this book we cover these topics:

- First we introduce our basic concepts (Chapter 1, “The main concept behind Network Attached Storage” on page 3).
- Next we cover NAS concepts and introduce the hardware we used in this book, including the IBM TotalStorage NAS 100, 200, and 300 (Chapter 2, “IBM TotalStorage NAS products” on page 25).

If you are already familiar with these concepts, or if you are anxious to get to the main points of this book, you can jump ahead to Part 2, “Implementing the NAS appliance in your network” on page 43.

1

The main concept behind Network Attached Storage

Given the expansive growth in both storage and network technology, it is not surprising that an easy-to-implement and scalable solution has been developed to meet the various storage needs.

Network Attached Storage (NAS) exploits the existing intermediate speed messaging network with a very easy-to-integrate storage solution.

In this book, we focus on NAS as a storage networking solution. Reading this book should adequately equip you to implement a NAS solution using one or more of the products we describe to meet your networked storage requirements.

This introductory chapter covers the following topics:

- How this book is organized
- Local Area Networks
- Open Systems Interconnection (OSI) model
- File systems and I/O
- Network Attached Storage (NAS)
- Industry standards

1.1 How this book is organized

Basically, here is how the material in this book is presented:

- First we provide the concepts and technical knowledge needed (Chapter 1, “The main concept behind Network Attached Storage” on page 3).
- Next we offer a brief overview of the IBM products we used (Chapter 2, “IBM TotalStorage NAS products” on page 25).
- After that we show the first steps for getting connected with the NAS appliance (Chapter 3, “Connecting to the NAS appliance” on page 45).
- Then we describe how to integrate the NAS appliances into your storage network. Each IBM TotalStorage NAS appliance has a special chapter (Chapter 4, “Implementing the IBM TotalStorage NAS 100” on page 57, Chapter 5, “Implementing the IBM TotalStorage NAS 200” on page 93, Chapter 6, “Implementing the IBM TotalStorage NAS 300” on page 113).
- Next we take a look at the basic user and security management features (Chapter 7, “User and security management” on page 199).
- After that we take a look at the advanced configuration of the NAS appliance systems (Chapter 8, “Advanced NAS configuration” on page 207).
- After we have finished with the configuration part, we then show how to manage the NAS appliance using IBM systems management tools (Chapter 9, “Systems management for NAS” on page 285).
- Next we explain how to set up the NAS appliance as a cross platform storage solution (Chapter 10, “Cross platform storage” on page 353).
- Then we talk about some backup and restore considerations, and we integrate the Persistent Storage Manager (PSM) into this picture to show how it can be used with Tivoli Storage Manager (Chapter 11, “Backup considerations” on page 431).
- Next we show how to rebuild the NAS appliance from scratch in case of a disaster (Chapter 12, “NAS recovery procedures” on page 453).
- After this we show how to integrate IBM TotalStorage NAS and Microsoft Active Directory (Chapter 13, “Microsoft Active Directory and IBM TotalStorage NAS” on page 467).
- Then we show how you can use the NAS appliance with Windows based applications (Chapter 14, “Solutions for Windows based applications” on page 497).
- Finally, after discussing Windows based applications, we switch to UNIX based applications (Chapter 15, “Solutions for UNIX based applications” on page 543).

Most of this book is a hands-on guide to implementing the NAS appliance as part of a storage networking solution, but before we can leap into the how-to section, it is important that you understand a few of the basic concepts about networks and storage.

Note: If you are a seasoned storage networking professional and are already very familiar with this subject, feel free to skip ahead to Part 2, “Implementing the NAS appliance in your network” on page 43. However, if you would like a quick primer, you will need to read these first two chapters. They provide the background information you need to understand, not only how to proceed with the integration, but also what you stand to gain from doing so.

1.2 Local Area Networks

A Local Area Network (LAN) is simply the connection of two or more computers (nodes) to facilitate data and resource sharing. They proliferated from the mid-1980s to address the problem of “islands of information” which occurred with standalone computers within departments and enterprises. LANs typically reside in a single or multiple buildings confined to a limited geographic area which is spanned by connecting two or more LANs together to form a Wide Area Network (WAN).

LAN designs are based typically on open systems networking concepts, as described in the network model of the Open Systems Interconnection (OSI) standards of the International Standards Organization (ISO). The OSI model is shown in detail in Figure 1-4, “Comparing the Internet protocol suite with the OSI reference model” on page 8.

LAN types are defined by their topology, which is simply how nodes on the network are physically connected together. A LAN may rely on a single topology throughout the entire network but typically has a combination of topologies connected using additional hardware. The primary topologies defined for Local Area Networks are:

Bus topology

In a bus topology, all nodes are connected to a central cable, called the bus or backbone. Bus networks are relatively inexpensive and easy to install. Ethernet systems use a bus topology (Figure 1-1).

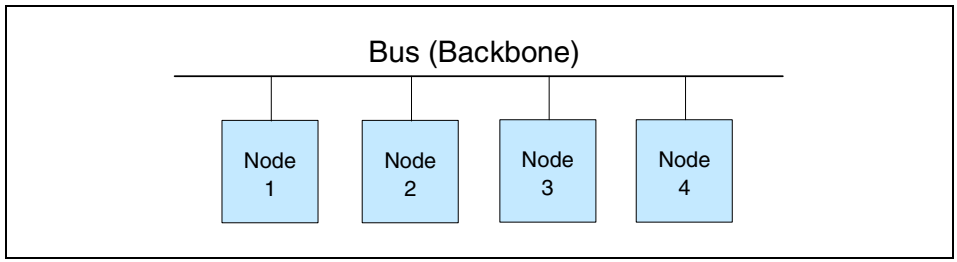

Figure 1-1 Bus topology

Ring topology

Nodes in a ring topology are connected via a closed loop such that each node has two other nodes connected directly to either side of it. Ring topologies are more costly and can be difficult to install. The IBM Token Ring uses a ring topology (Figure 1-2).

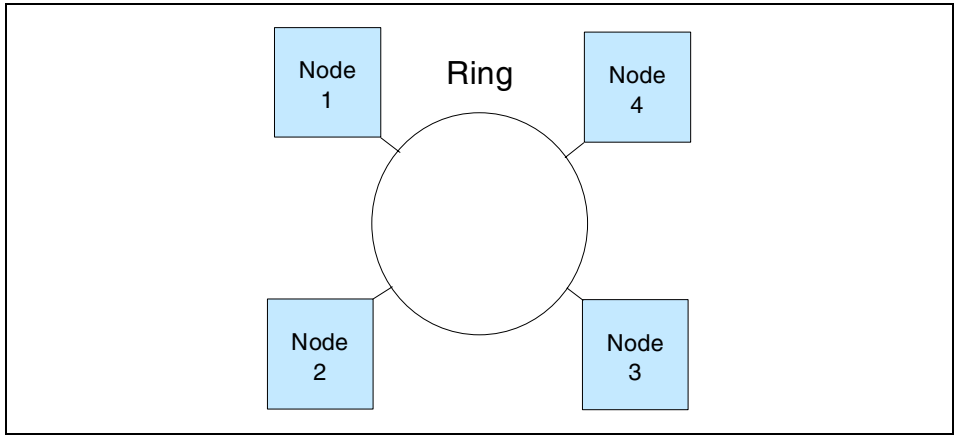

Figure 1-2 Ring topology

Star topology

A star topology uses a centralized hub to connect the nodes in the network together. Star networks are easy to install and manage. However, bottlenecks occur since all of the network traffic travels through the hub. Ethernet systems also use a star topology (Figure 1-3).

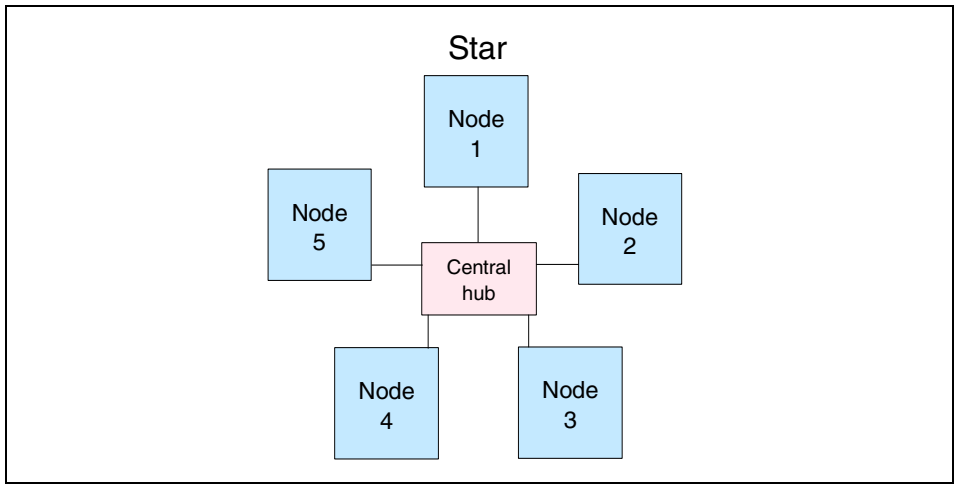

Figure 1-3 Star topology

Today, Ethernet topologies are predominant. International Data Corporation (IDC) estimates more than 85% of all installed network connections worldwide are Ethernet. It is popular due to its simplicity, affordability, scalability, and manageability. Ethernet includes definitions of protocols for addressing, formatting and sequencing of data transmissions across the network and also describes the physical media (cables) used for the network.

1.3 Open Systems Interconnection (OSI) model

The Open Systems Interconnection (OSI) model describes the layers in the network required for communication between computers. OSI is a seven layered model illustrated with the Internet protocol suite (or stack) in Figure 1-4. Each layer is responsible for a certain set of tasks associated with moving data across the network. Most Ethernet networks (including ours) communicate using the TCP/IP protocol. In this section, we discuss TCP/IP and how it relates to the OSI model, since it is the default communication protocol for the NAS appliance.

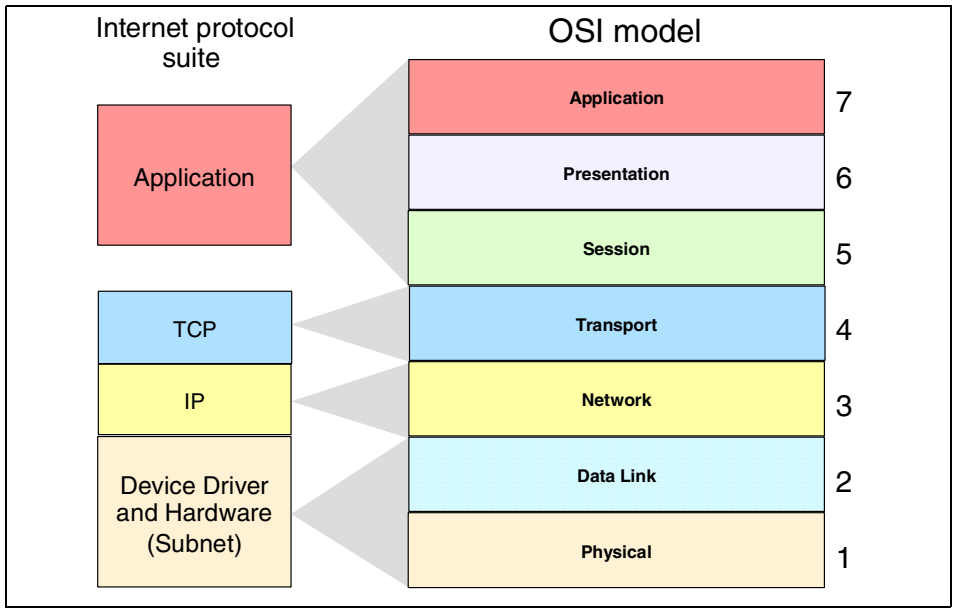

Figure 1-4 Comparing the Internet protocol suite with the OSI reference model

1.3.1 Device driver and hardware layer

Also called the Subnet layer, the device driver and hardware layer comprises both the physical and data link layers of the OSI model. It is considered the hardware that is part of each node on the network. The hardware handles the electrical and mechanical aspects of data transfers, moving the bits across a physical link. The data link layer packages packets of data into frames, ensures that they arrive safely to the target destination, and encompasses error detection and correction.

1.3.2 Internet Protocol layer

In the OSI model, the Network layer finds the best route through the network to the target destination. It has little to do in a single discrete LAN; but in a larger network with subnets, or access to WANs, the Network layer works with the various routers, bridges, switches, gateways, and software, to find the best route for data packets.

The Internet Protocol (IP) layer in the Internet protocol suite performs the functions of the network layer. It is the common thread running through the Internet and most LAN technologies, including Ethernet. It is responsible for moving data from one host to another, using various “routing” algorithms. Layers above the network layer break a data stream into chunks of a predetermined size, known as packets or datagrams. The datagrams are then sequentially passed to the IP layer.

The job of the IP layer is to route these packets to the target destination. IP packets consist of an IP header, together with the higher level TCP protocol and the application datagram. IP knows nothing about the TCP and datagram contents. Prior to transmitting data, the network layer might further subdivide it into smaller packets for ease of transmission. When all the pieces finally reach the destination, they are reassembled by the network layer into the original datagram.

IP connectionless service

The IP is the standard that defines the manner in which the network layers of two hosts interact. These hosts may be on the same network, or reside on physically remote heterogeneous networks. IP was designed with inter-networking in mind. It provides a connectionless, best-effort packet delivery service. Its service is called connectionless because it is like the postal service rather than the telephone system.

IP packets, like telegrams or mail, are treated independently. Each packet is stamped with the addresses of the receiver and the sender. Routing decisions are made on a packet-by-packet basis. On the other hand, connection-oriented, circuit switched telephone systems explicitly establish a connection between two users before any conversation takes place. They also maintain the connection for the entire duration of conversation.

A best-effort delivery service means that packets might be discarded during transmission, but not without a good reason. Erratic packet delivery is normally caused by the exhaustion of resources, or a failure at the data link or physical layer. In a highly reliable physical system s(p)-16 r.8(i)-16.6ad6 r.8(i5( )24.4(a)7.s(y)0.34(a)7. In)7.5E

The IP packet

All IP packets or datagrams consist of a header section and a data section (payload). The payload may be traditional computer data, or it may, commonly today, be digitized voice or video traffic. Using the postal service analogy again, the “header” of the IP packet can be compared with the envelope and the “payload” with the letter inside it. Just as the envelope holds the address and information necessary to direct the letter to the desired destination, the header helps in the routing of IP packets.

The payload has a maximum size limit of 65,536 bytes per packet. It contains error and/or control protocols, like the Internet Control Message Protocol (ICMP). To illustrate control protocols, suppose that the postal service fails to find the destination on your letter. It would be necessary to send you a message indicating that the recipient's address was incorrect. This message would reach you through the same postal system that tried to deliver your letter. ICMP works the same way: It packs control and error messages inside IP packets.

IP addressing

An IP packet contains a source and a destination address. The source address designates the originating node's interface to the network, and the destination address specifies the interface for an intended recipient or multiple recipients (for broadcasting).

Every host and router on the wider network has an address that uniquely identifies it. It also denotes the sub-network on which it resides. No two machines can have the same IP address. To avoid addressing conflicts, the network numbers are assigned by an independent body.

The network part of the address is common for all machines on a local network. It is similar to a postal code, or zip code, that is used by a post office to route letters to a general area. The rest of the address on the letter (i.e., the street and house number) are relevant only within that area. It is only used by the local post office to deliver the letter to its final destination.

The host part of the IP address performs a similar function. The host part of an IP address can further be split into a sub-network address and a host address.

Time to Live (TTL)

The IP packet header also includes Time to Live (TTL) information that is used to limit the life of the packet on the network. It includes a counter that is decremented each time the packet arrives at a routing step. If the counter reaches zero, the packet is discarded.

1.3.3 TCP layer

The transport layer is responsible for ensuring delivery of the data to the target destination, in the correct format in which it was sent. In the event of problems on the network, the Transport layer finds alternative routes. It is also responsible for delivering the sequence of packets in the correct order. In the Internet protocol suite, the protocol operating in the transport layer is the Transmission Control Program (TCP).

The application data has no meaning to the Transport layer. On the source node, the transport layer receives data from the application layer and splits it into data packets or chunks. The chunks are then passed to the network layer. At the destination node, the transport layer receives these data packets and reassembles them before passing them to the appropriate process or application.

The Transport layer is the first end-to-end layer of the TCP/IP stack. This characteristic means that the transport layer of the source host can communicate directly with its peer on the destination host, without concern about 'how' data is moved between them. These matters are handled by the network layer. The layers below the transport layer understand and carry information required for moving data across links and subnetworks.

In contrast, at the transport layer or above, one node can specify details that are only relevant to its peer layer on another node. For example, it is the job of the transport layer to identify the exact application to which data is to be handed over at the remote end. This detail is irrelevant for any intermediate router. But it is essential information for the transport layers at both the ends.

1.3.4 Application layer

The functions of the Session, Presentation, and Application layers of the OSI model are all combined in the Application layer of the Internet protocol suite. It encompasses initial logon, security, final termination of the session, interpretation services (compression, encryption, or formatting), and delivery of the network messages to the end user program.

The Application layer is the layer with which end users normally interact. It is responsible for formatting the data so that its peers can understand it. Whereas the lower three layers are usually implemented as a part of the OS, the application layer is a user process. Some application-level protocols that are included in most TCP/IP implementations, include:

- Telnet for remote login
- File Transfer Protocol (FTP) for file transfer
- Simple Mail Transfer Protocol (SMTP) for mail transfer

1.3.5 Protocol suites

A protocol suite (or protocol stack), as we saw in the Internet protocol suite, is organized so that the highest level of abstraction resides at the top layer. For example, the highest layer may deal with streaming audio or video frames, whereas the lowest layer deals with raw voltages or radio signals. Every layer in a suite builds upon the services provided by the layer immediately below it.

Note: You may see the different terms Internet protocol suite, *TCP/IP suite*, or *TCP/IP stack*. These are simply names for the same thing, the group of network layers to describe how two nodes on the Internet communicate.

The terms protocol and service are often confused. A *protocol* defines the exchange that takes place between identical layers of two hosts. For example, in the IP suite, the transport layer of one host talks to the transport layer of another host using the TCP protocol. A *service*, on the other hand, is the set of functions that a layer delivers to the layer above it. For example, the TCP layer provides a reliable byte-stream service to the application layer above it.

Each layer adds a header containing layer-specific information to the data packet. A header for the network layer might include information such as source and destination addresses. The process of appending headers to the data is called encapsulation. Figure 1-5 shows how data is encapsulated by various headers. During de-encapsulation the reverse occurs; the layers of the receiving stack extract layer-specific information and process the encapsulated data accordingly. The process of encapsulation and de-encapsulation increases the overhead involved in transmitting data.

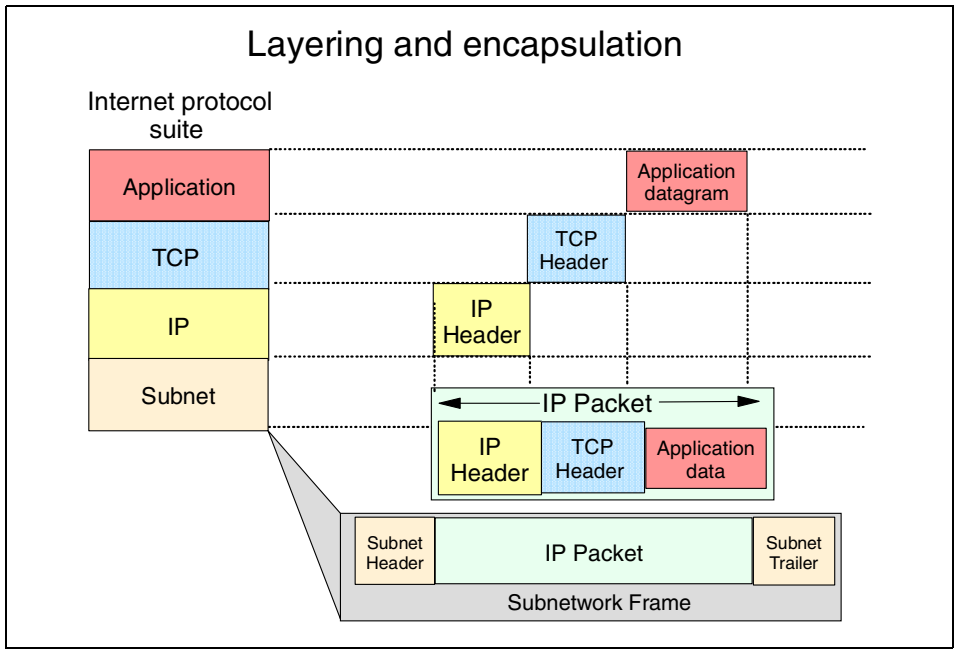

Figure 1-5 Layering and encapsulation

1.4 File systems and I/O

In this section we describe the most common file level protocols and attempt to untangle the confusion surrounding the various I/O concepts.

1.4.1 Network file system protocols

The two most common file level protocols used to share files across networks are Network File System (NFS) for UNIX and Common Internet File System (CIFS) for Windows. Both are network based client/server protocols which enable hosts to share resources across a network using TCP/IP. Users manipulate shared files, directories, and devices such as printers, as if they were locally on or attached to the user's own computer. The NAS appliances are preconcerted to support both NFS and CIFS.

Network File System (NFS)

NFS servers make their file systems available to other systems in the network by *exporting* directories and files over the network. Once exported, an NFS client can then “mount” a remote file system from the exported directory location. NFS controls access by giving client-system level user authorization based on the assumption that a user who is authorized to the system must be trustworthy. Although this type of security is adequate for some environments, it is open to abuse by anyone who can access a UNIX system via the network.

For directory and file level security, NFS uses the UNIX concept of file permissions with *User* (the owner’s ID), *Group* (a set of users sharing a common ID), and *Other* (meaning all other user IDs). For every NFS request, the IDs are verified against the UNIX file permissions.

NFS is a *stateless* service. Therefore, any failure in the link will be transparent to both client and server. When the session is re-established the two can immediately continue to work together again.

NFS handles file locking by providing an *advisory lock* to subsequent applications to inform them that the file is in use by another application. The ensuing applications can decide if they want to abide by the lock request or not. This has the advantage of allowing any UNIX application to access any file at any time, even if it is in use. The system relies on “good neighbor” responsibility which, though often convenient, clearly is not foolproof. This is avoided by using the optional Network Lock Manager (NLM). It provides file locking support to prevent multiple instances of open files.

Common Internet File System (CIFS)

Another method used to share resources across a network uses CIFS, which is a protocol based on Microsoft’s Server Message Block (SMB) protocol. Using CIFS, servers create *file shares* which are accessible by authorized clients. Clients subsequently connect to the server’s shares to gain access to the resource.

Security is controlled at both the user and share level. Client authentication information is sent to the server before the server will grant access. CIFS uses access control lists that are associated with the shares, directories, and files, and authentication is required for access.

A *session* in CIFS is oriented and *stateful*. This means that both client and server share a history of what is happening during a session, and they are aware of the activities occurring. If there is a problem, and the session has to be re-initiated, a new authentication process must be completed.

CIFS employs opportunistic locks *(oplocks)* to control file access. Depending on the type of locking mechanism required by the client, CIFS offers nodes the ability to cache read or write data from the file being accessed to improve network performance. Exclusive rights to the file prevents other nodes on the network from gaining access to that file until it is closed. During a CIFS session the lock manager has historical information concerning which client has opened the file, for what purpose, and in which sequence.

1.4.2 Understanding I/O

A major source of confusion regarding NAS is the concept of *File I/O* versus *Block I/O*. We try to shed a little light on this subject here. Understanding the difference between these two forms of data access is crucial to realizing the potential benefits of any SAN-based or NAS-based solution.

When a partition on a hard drive is under the control of an operating system (OS), the OS will format it. Formatting of the partition occurs when the OS lays a file system structure on the partition. This file system is what enables the OS to keep track of where it stores data. The file system is an addressing scheme the OS uses to map data on the partition. Now, when you want to get to a piece of data on that partition, you must request the data from the OS that controls it.

For example, suppose that Windows 2000 formats a partition (or drive) and maps that partition to your system. Every time you request to open data on that partition, your request is processed by Windows 2000. Since there is a file system on the partition, it is accessed via File I/O. Additionally, you cannot request access to just the last 10 KB of a file. You must open the entire file, which is another reason that this method is referred to as File I/O.

Block I/O (raw disk) is handled differently: There is no OS format done to lay out a file system on the partition. The addressing scheme that keeps up with where data is stored is provided by the application using the partition. An example of this would be DB2 using its tables to keep track of where data is located rather than letting the OS do that job. That is not to say that DB2 cannot use the OS to keep track of where files are stored. It is just more efficient, for the database to bypass the cost of requesting the OS to do that work.

Using File I/O is like using an accountant. Accountants are good at keeping up with your money for you, but they charge you for that service. For your personal checkbook, you probably want to avoid that cost. On the other hand, for a corporation where many different kinds of requests are made, an accountant is a good idea. That way, checks are not written when they should not be.

When sharing files across a network, something needs to control when writes can be done. The operating system fills this role. It does not allow multiple writes at the same time, even though many write requests are made. Databases are able to control this writing function on their own so in general they run faster by skipping the OS although this depends on the efficiency of the implementation of file system and database.

For a more in-depth study of these topics, refer to the redbook, *IP Storage Networking: IBM NAS and iSCSI Solutions*, SG24-6240.

1.5 Network Attached Storage (NAS)

Storage devices which optimize the concept of file sharing across the network have come to be known as Network Attached Storage (NAS). NAS solutions utilize the mature Ethernet IP network technology of the LAN. Data is sent to and from NAS devices over the LAN using TCP/IP.

By making storage devices LAN addressable, the storage is freed from its direct attachment to a specific server and any-to-any connectivity is facilitated using the LAN fabric. In principle, any user running any operating system can access files on the remote storage device. This is done by means of a common network access protocol, for example, NFS for UNIX servers, and CIFS for Windows servers.

A storage device cannot just attach to a LAN. It needs intelligence to manage the transfer and the organization of data on the device. The intelligence is provided by a dedicated server to which the common storage is attached. It is important to understand this concept. NAS comprises a server, an operating system, plus storage which is shared across the network by many other servers and clients. So NAS is a *device*, rather than a *network infrastructure*, and shared storage is either internal to the NAS device or attached to it.

1.5.1 File servers

Early NAS implementations in the late 1980s used a standard UNIX or NT server with NFS or CIFS software to operate as a remote file server. In such implementations, clients and other application servers access the files stored on the remote file server, as though the files are located on their local disks. The location of the file is transparent to the user.

Several hundred users could work on information stored on the file server, each one unaware that the data is located on another system. The file server has to manage I/O requests accurately, queuing as necessary, fulfilling the request and returning the information to the correct initiator. The NAS server handles all aspects of security and lock management. If one user has the file open for updating, no-one else can update the file until it is released. The file server keeps track of connected clients by means of their network IDs, addresses, and so on.

1.5.2 Network appliances

More recent developments use application specific, specialized, “thin server” configurations with customized operating systems, usually comprising a stripped down UNIX kernel, reduced Linux OS, or a specialized Windows 2000 kernel, as with the IBM TotalStorage NAS appliances. In these reduced operating systems, many of the server operating system functions are not supported. The objective is to improve performance and reduce costs by eliminating unnecessary functions normally found in the standard hardware and software. Some NAS implementations also employ specialized data mover engines and separate interface processors in efforts to further boost performance.

These specialized file servers with a reduced OS are typically known as appliances, describing the concept of an application specific system. The term “appliance” borrows from household electrical devices the idea of a specialized application specific tool, such as a coffee maker or a toaster. NAS appliances, like the IBM TotalStorage NAS appliances, typically come with pre-configured software and hardware, and with no monitor or keyboard for user access. This is commonly termed a “headless” system. A storage administrator accesses the appliance and manages the disk resources from a remote console.

One of the typical characteristics of a NAS appliance is its ability to be installed rapidly using minimal time and effort to configure the system. It is integrated seamlessly into the network as shown in Figure 1-6. This approach makes NAS appliances especially attractive when lack of time and skills are elements in the decision process.

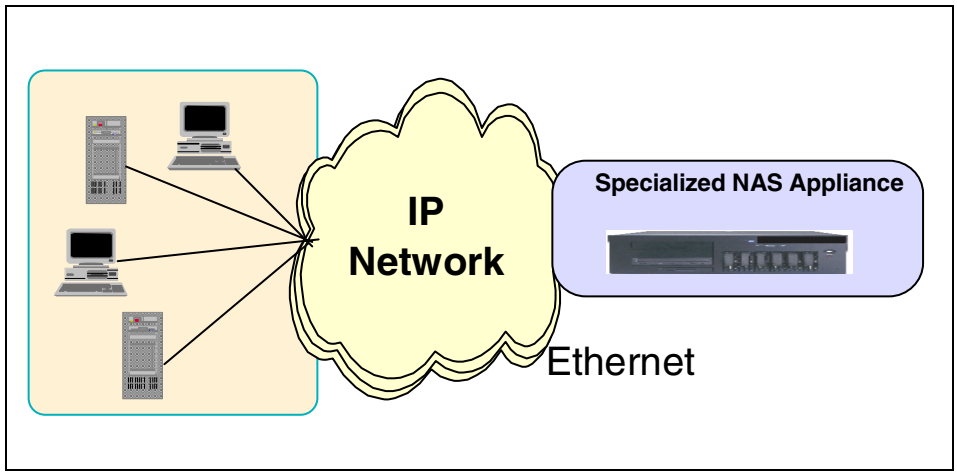

Figure 1-6 The role of the NAS appliance in your storage network

So, a NAS appliance is an easy-to-use device, which is designed for a specific function, such as serving files to be shared among multiple clients. It performs this task very well. It is important to recognize this fact when selecting a NAS solution. The NAS appliance is not a general purpose server, and should not be used (indeed, due to its reduced OS, probably cannot be used) for general purpose server tasks. However, it does provide a good solution for appropriately selected shared storage applications.

1.5.3 NAS uses File I/O

One of the key differences of a NAS disk device, compared to direct access storage (DAS) is that all I/O operations use file level I/O protocols. File I/O is a high level type of request that, in essence, specifies only the file to be accessed, but does not directly address the storage device. This is done later by other operating system functions in the remote NAS appliance.

A File I/O request specifies the file and the offset into the file. For instance, the I/O may specify “Go to byte ‘1000’ in the file (as if the file was a set of contiguous bytes), and read the next 256 bytes beginning at that position”. Unlike Block I/O, there is no awareness of a disk volume or disk sectors in a File I/O request. Inside the NAS appliance, the operating system keeps track of where files are located on disk. The OS issues a Block I/O request to the disks to fulfill the File I/O read and write requests it receives.

Network access methods, NFS and CIFS, can only handle File I/O requests to the remote file system. I/O requests are packaged by the node initiating the I/O request into packets to move across the network. The remote NAS file system converts the request to Block I/O and reads or writes the data to the NAS disk storage. To return data to the requesting client application, the NAS appliance software re-packages the data in TCP/IP protocols to move it back across the network. This is illustrated in Figure 1-7.

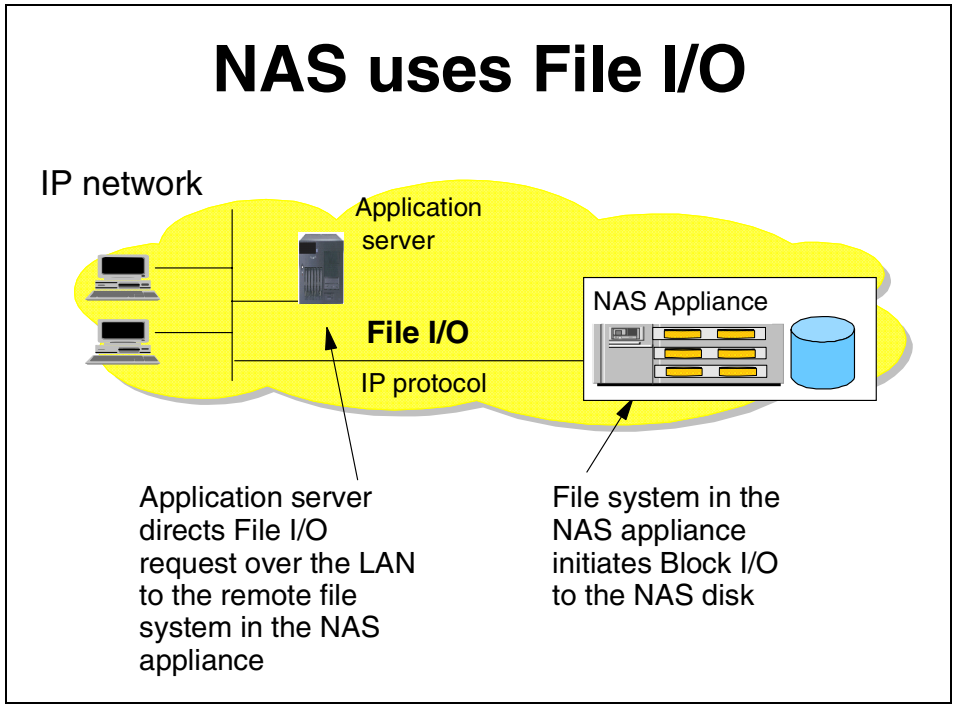

Figure 1-7 NAS devices use File I/O

1.5.4 NAS benefits

NAS offers a number of benefits that address some of the limitations of directly attached storage devices, and that overcome some of the complexities associated with SANs.

Resource pooling

A NAS appliance enables disk storage capacity to be consolidated and pooled on a shared network resource, at great distances from the clients and servers which will share it. Thus a NAS device can be configured as one or more file systems, each residing on specified disk volumes. All users accessing the same file system are assigned space within it on demand. This contrasts with individual DAS storage, when some users may have too little storage, and others may have too much.

Consolidation of files onto a centralized NAS device can minimize the need to have multiple copies of files spread on distributed clients. Thus overall hardware costs can be reduced.

NAS pooling can reduce the need to physically reassign capacity among users. The results can be lower overall costs through better utilization of the storage, lower management costs, increased flexibility, and increased control.

Exploits existing infrastructure

Because NAS utilizes the existing LAN infrastructure, there are minimal costs of implementation. Introducing a new network infrastructure, such as a Fibre Channel SAN, can incur significant hardware costs. In addition, new skills must be acquired, and a project of any size will need careful planning and monitoring to bring it to completion.

Simple to implement

Because NAS devices attach to mature, standard LAN implementations, and have standard LAN addresses, they are typically extremely easy to install, operate, and administer. This plug-and-play operation results in lower risk, ease of use, and fewer operator errors, all of which contributes to lower costs of ownership.

Enhanced choice

The storage decision is separated from the server decision, thus enabling the buyer to exercise more choice in selecting equipment to meet the business needs.

Connectivity

LAN implementation allows any-to-any connectivity across the network. NAS appliances may allow for concurrent attachment to multiple networks, thus supporting many users.

Scalability

NAS appliances can scale in capacity and performance within the allowed configuration limits of the individual appliance. However, this may be restricted by considerations such as LAN bandwidth constraints, and the need to avoid restricting other LAN traffic.

Heterogeneous file sharing

Remote file sharing is one of the basic functions of any NAS appliance. Multiple client systems can have access to the same file. Access control is serialized by NFS or CIFS. Heterogeneous file sharing may be enabled by the provision of translation facilities between NFS and CIFS, as with the NAS 100, 200 and 300.

Improved manageability

By providing consolidated storage, which supports multiple application systems, storage management is centralized. This enables a storage administrator to manage more capacity on a NAS appliance than typically would be possible for distributed, directly attached storage.

Enhanced backup

NAS appliance backup is a common feature of most popular backup software packages. For instance, the IBM TotalStorage NAS appliances all provide TSM client software support. Some NAS appliances have some integrated, automated backup facility to tape, enhanced by the availability of advanced functions such as the IBM NAS appliance facility called Persistent Storage Manager (PSM). This enables multiple point-in-time copies of files to be created on disk, which can be used to make backup copies to tape in the background. This is similar in concept to features such as IBM's Snapshot function on the IBM RAMAC Virtual Array (RVA).

1.5.5 Other NAS considerations

On the converse side of the storage network decision, you need to take into consideration the following factors regarding NAS solutions.

Proliferation of NAS devices

Pooling of NAS resources can only occur within the capacity of the individual NAS appliance. As a result, in order to scale for capacity and performance, there is a tendency to grow the number of individual NAS appliances over time, which can increase hardware and management costs.

Software overhead impacts performance

As we explained earlier, TCP/IP is designed to bring data integrity to Ethernet-based networks by guaranteeing data movement from one place to another. The trade-off for reliability is a software intensive network design which requires significant processing overheads, which can consume more than 50% of available processor cycles when handling Ethernet connections. This is not normally an issue for applications such as Web-browsing, but it is a drawback for performance intensive storage applications.

Consumption of LAN bandwidth

Ethernet LANs are tuned to favor short burst transmissions for rapid response to messaging requests, rather than large continuous data transmissions. Significant overhead can be imposed to move large blocks of data over the LAN. The maximum packet size for Ethernet is 1518 bytes. A 10 MB file has to be segmented into more than 7000 individual packets. Each packet is sent separately to the NAS device by the Ethernet collision detect access method. As a result, network congestion may lead to reduced or variable performance.

Data integrity

The Ethernet protocols are designed for messaging applications, so data integrity is not of the highest priority. Data packets may be dropped without warning in a busy network, and have to be resent. Since it is up to the receiver to detect that a data packet has not arrived, and to request that it be resent, this can cause additional network traffic.

With NFS file sharing there are some potential risks. Security controls can fairly easily be by-passed. This may be a concern for certain applications. Also the NFS file locking mechanism is not foolproof, so that multiple concurrent updates could occur in some situations.

Impact of backup/restore applications

One of the potential downsides of NAS is the consumption of substantial amounts of LAN bandwidth during backup and restore operations, which may impact other user applications. NAS devices may not suit applications which require very high bandwidth. To overcome this limitation, some users implement a dedicated IP network for high data volume applications, in addition to the messaging IP network. This can add significantly to the cost of the NAS solution.

1.5.6 Total cost of ownership

Because it makes use of both existing LAN network infrastructures and network administration skills already employed in many organizations, NAS costs may be substantially lower than for directly attached or SAN-attached storage. Specifically, NAS-based solutions offer the following cost-reducing benefits:

- They reduce administrative staff requirements.
- They improve reliability and availability.
- They bridge the gap between UNIX and Windows environments.

Reduced administrative staff requirements

Implementing single or clustered NAS appliances to manage your networked storage concentrates the administrative tasks and thereby reduces the number of people required to maintain the network. Since the NAS appliance is a headless system, administration is usually performed via a Web-based GUI interface accessible from anywhere on the network. In addition, more capacity can be managed per administrator, thus resulting in a lower cost of ownership.

Improved reliability and availability

In today's business world, it has become the de facto standard to provide customers with access to information 24 hours per day, 7 days per week, allowing very little time available for unplanned outages. Some IBM NAS appliances offer the ability to provide great availability with options for clustered models.

Bridges the gap between UNIX and Windows environments

Most companies today contain heterogeneous operating environments. A NAS solution offers customers the ability for true cross-platform file sharing between Windows and UNIX clients by offering support for CIFS and NFS. This becomes increasingly important when application data becomes more common across platforms.

1.6 Industry standards

There is a clear customer need for standardization within the storage networking industry to allow users to freely select equipment and solutions, knowing that they are not tying themselves to a proprietary or short term investment. To this end, there are extensive efforts among the major vendors in the storage networking industry to cooperate in the early agreement, development, and

adoption of standards. A number of industry associations, standards bodies, and company groupings are involved in developing and publishing storage networking standards. The most important of these are the Storage Networking Industry Association (SNIA) and the Internet Engineering Task Force (IETF).

In addition, IBM, IBM Business Partners, and other major vendors in the industry, have invested heavily in inter-operability laboratories. The IBM laboratories in Gaithersburg (Maryland, USA), Mainz (Germany), and Tokyo (Japan) are actively testing equipment from IBM and many other vendors, to facilitate the early confirmation of compatibility between multiple vendors servers, storage, and network hardware and software components.

1.6.1 Storage Networking Industry Association

The Storage Networking Industry Association (SNIA) is an international computer industry forum of developers, integrators, and IT professionals who evolve and promote storage networking technology and solutions. SNIA was formed to ensure that storage networks become efficient, complete, and trusted solutions across the IT community.

SNIA is accepted as the primary organization for the development of SAN and NAS standards, with over 150 companies and individuals as its members, including all the major server, storage, and fabric component vendors. SNIA is committed to delivering architectures, education, and services that will propel storage networking solutions into a broader market.

IBM is one of the founding members of SNIA, and has senior representatives participating on the board and in technical groups. For additional information on the various activities of SNIA, see its Web site at:

http://www.snia.org

1.6.2 Internet Engineering Task Force

The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. The actual technical work of the IETF is done in its working groups, which are organized by topic into several areas (for example, routing, transport, and security).

For more information on the IETF and its work groups, refer to:

http://www.ietf.org

2

IBM TotalStorage NAS products

In this chapter we provide a brief overview of the IBM TotalStorage NAS products used during the development of this IBM Redbook. NAS appliances like the IBM TotalStorage Network Attached Storage 100, 200 and 300 are fully integrated and dedicated storage solutions that can be quickly and easily attached to an IP network. Their storage will then become immediately and transparently available as a network file-serving resource to all clients. These specialized appliances are also independent of their client platforms and operating systems. Thus, NAS appliances appear to the client application as just another server.

We cover these topics:

- IBM TotalStorage NAS features and benefits
- IBM TotalStorage Network Attached Storage 100
- IBM TotalStorage Network Attached Storage 200
- IBM TotalStorage Network Attached Storage 300
- IBM TotalStorage NAS Release 2.5
- IBM TotalStorage NAS Release 2.7

2.1 IBM TotalStorage NAS features and benefits

The NAS 100, 200 and 300 are appliances that are designed to work in heterogeneous environments right out of the box. Figure 2-1 on page 26 visually demonstrates how the built-in features of the NAS 100, 200 or 300 allow it to plug into almost any environment.

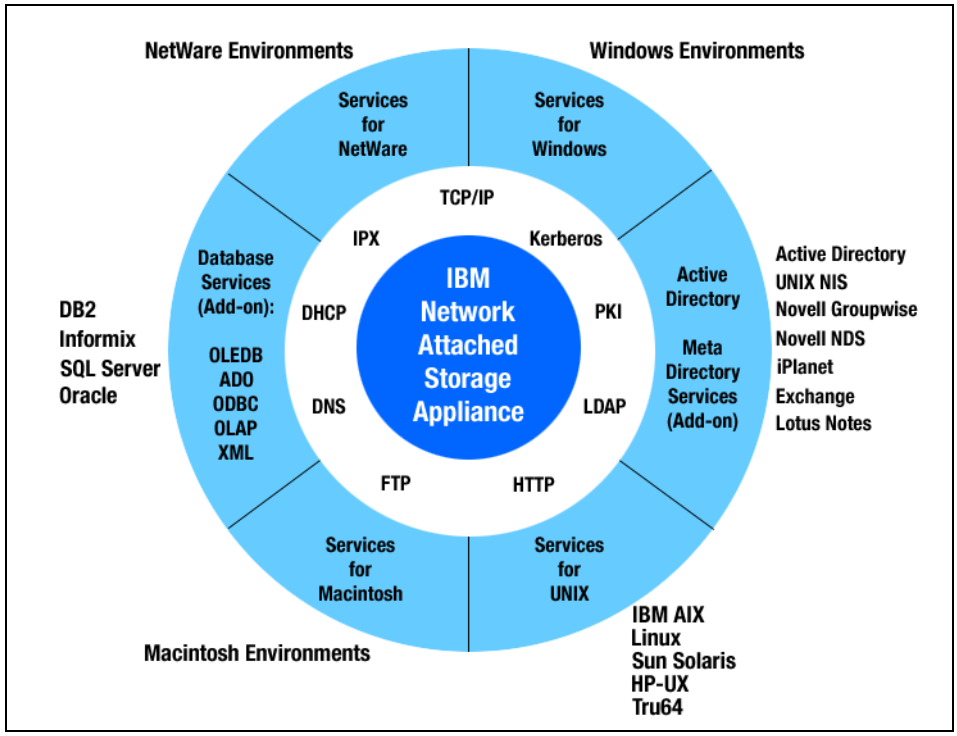

Figure 2-1 Visualization of interoperability features

2.1.1 Included software

The software in Table 2-1 is included in the NAS 200 and 300.

Table 2-1 NAS 200 and 300 software

Software	IBM TotalStorage NAS 200 and 300
Operating system	Windows Powered OS
Data Protection	Columbia Data Products Persistent Storage Manager enables 250 point-in-time True Image data views.

Software	IBM TotalStorage NAS 200 and 300
Backup/Restore management	NT/2000 Backup
Backup/Restore management	Tivoli Storage Manager Client(V4.2.1.20)
Systems management	IBM Director 3.1 agent
Remote Administration	Web-based GUI Microsoft Terminal Services

2.1.2 Preloaded and optional software

Each NAS 200 or 300 is preloaded at the factory with its base operating system, installation and administration software. The code is loaded to the system's hard disk with a backup copy provided on an emergency recovery CD-ROM. The operating system and NAS application code have been specifically tuned to enable the NAS 200 and 300 as high performance NAS server appliances. The difference between the NAS 200 and 300 is that the NAS 300 comes with additional software for clustering.

In addition to the operating system and application software, each unit contains tools which simplify remote configuration and administration tasks. Additionally, included network management agents provide options for managing the units.

Specifically, the units come preconfigured with the following functions:

- Windows Powered OS:
  - Windows 2000 Advanced Server code optimized for the IBM TotalStorage NAS 200 and 300 Models
- Multiple file systems support:
  - CIFS
  - NFS
  - Netware
  - Macintosh
- Multiple file transfer services:
  - FTP
  - HTTP
- Remote NAS system administration:
  - Administrative tasks can be performed in the Web-based GUI
  - IBM Advanced Appliance Configuration Utility
  - Alternate administrative task performed using Windows Terminal Service
  - Advanced management functions available via Windows Terminal Service
  - Backup and restores using NT Backup or Tivoli Storage Manager

  - NAS Backup Assistant MMC Snap-in Web page
- UNIX services:
  - Pre-configured NFS support
  - Web-based GUI for performing administrative tasks
  - Microsoft Services for UNIX V2.3
  - NFS V3.0 (IETF RFC 1830)
- Setup tutorial:
  - NAS Setup Navigator will help you in the setup process
- Automatic disaster recovery of operating system:
  - The IBM Snap command-line utility creates a point-in-time persistent True Image of the specified drive letter, and then backs up that persistent image
  - Supports full and incremental backups
  - The NAS Backup Assistant is a GUI front end to IBM Snap that generates a batch file that invokes NT Backup from settings configured by the user in the GUI
  - Original factory CD-ROM reload of operating system
  - Prevention of accidental reloads via reload enablement diskette
- IBM Fibre Management utility:
  - IBM Fibre Stand-Alone Management utility
  - MMC snap-in that launches the utility
  - Users can use terminal services to remotely monitor the fibre adapter configuration
- Advanced Aggregate Management:
  - IBM Director Agent
- Persistent Storage Manager (PSM) for IBM NAS:

  Persistent Storage Manager (PSM) creates True Images (these are multiple point-in-time persistent images of any or all system and data volumes). All persistent images survive system power loss or a planned or unplanned reboot. Each instance of PSM seamlessly handles 250 concurrent images of up to 255 independent volumes for a total of 63,750 independent data images. Any image can be easily managed through the Microsoft Web user interface, and accessed the same as any other active volume.

  In case of data corruption or loss, any persistent image can be used for manual retrieval of individual files (by the administrator or end users), or more importantly, for instant restoration (by a PSM function initiated by the administrator, in the Web user interface) of the entire volume from image, which can substantially reduce the amount of system down time.

- Persistent Storage Manager creates and keeps multiple point-in-time persistent images (maximum of 250 concurrent images of up to 255 independent volumes)
- All images for a volume are mounted under a single directory (in the root directory of the volume), with each image under its own mount point
- User-level access can be granted to one or more of the images, to allow users to restore their own files from the images (users automatically have the same access privileges to individual files and directories in the images that they would have on the actual volume)
- Images can be read-only, or read-write (with ability to reset (undo changes to) read-write images)
- Images can be assigned retention levels (if an image needs to be automatically deleted by PSM, the highest priority images can be kept)
- Any image can be used to restore an entire volume instantly (for data volumes, typically within seconds; for system volume, system reboot is required)
- Flexible, configurable image access and administration via Web-based user interface
- Schedule images (for each schedule entry (image group), specify interval, number to keep, image name, properties (read-only or read-write, retention level))
- Create a new image immediately (specify name and properties)
- Delete images
- View and change properties of images (also reset read-write images)
- Restore a volume (from any image of that volume)
- Configure advanced parameters:
  - Maximum number of images to keep concurrently
  - Name of image root directory
  - Quiescent period and quiescent period wait time-out
  - Size of the image cache file (per volume)
- Image cache file usage warning and automatic image deletion thresholds (per volume)

► IBM Director with Universal Manageability (UM) Services V3.1:

The IBM TotalStorage NAS 200 and 300 contain a IBM Director agent and can be managed by this powerful, highly-integrated, systems management software solution that is built upon industry standards and designed for ease-of-use. Using its intuitive Java-based GUI, an administrator can

centrally manage individual or large groups of IBM and non-IBM PC-based servers. IT administrators can view the hardware configuration of remote systems in detail and monitor the usage and performance of crucial components, such as processors, disks, and memory.

The following functions for NAS have been added in V3.1:

- Learn detailed inventory information about the deployed NAS, including operating system, memory, network card and hardware.
- Track all managed NAS boxes operatively with features such as power management, event log and system monitor capabilities.
- Upwardly integrate with Tivoli Enterprise, Tivoli NetView, Computer Associated Unicenter, HP Openview, Microsoft SMS and Intel LANDesk Management Suite

IBM Director with UM Services V3.1 is the latest update to IBM world-class systems manageability solutions. V3.1 replaces all earlier versions of NF Director and UM Services.

2.1.3 Limitations of the Windows Powered OS

There are some limitations in the Windows Powered OS running on the NAS units. The operating system in tuned for optimal performance but the following functions cannot be done with an IBM NAS appliance:

- Windows Domain Controller (Included in Domainlet feature for 300/300G)
- DHCP Server
- DNS Server (Included in Domainlet feature for 300/300G)
- WINS Server

2.1.4 IBM Advanced Appliance Configuration Utility Tool

The IBM Advanced Appliance Configuration Utility tool helps you set or reconfigure the network configuration for one or many appliance servers. This software consists of an agent on each appliance server and a Java application residing on a Windows-based client workstation acting as a configuration station.

You can use this configuration station to do the following:

- Discover appliance servers
- Set up and manage server network configurations
- Launch the comprehensive Web-based server management console in a separate browser window

Network administrators not currently running DHCP servers will find the advanced appliance configuration utility particularly useful for automatically configuring network settings for newly added IBM TotalStorage NAS appliances. Even administrators with networks using DHCP servers can benefit from the advanced appliance configuration utility, by permanently assigning IP addresses and host names automatically and launching Web-based management.

2.2 IBM TotalStorage Network Attached Storage 100

The IBM TotalStorage NAS 100, 5190 Model R12, is the newest member of IBM's growing family of Network Attached Storage (NAS) products, joining the TotalStorage NAS 200, NAS 300, and NAS 300G as the entry-level family member. It is a low profile (1U) table top NAS solution designed for central management of your remote and branch locations' data storage needs.

As a member of the IBM NAS family of products, the NAS 100 is characterized by many of the same outstanding features and attributes as those of other NAS family members. Those features include preloaded operating systems and application code, integrated software functions that help ease configuration and use, ease of deployment in the network, as well as built-in tools to facilitate remote management and systems management. The NAS 100 supports the major industry-standard file protocols: CIFS, NFS, NetWare, Apple File Protocol, HTTP, and FTP.

The IBM TotalStorage NAS 100 is a high-performance network storage device designed for a variety of Ethernet LAN storage applications. The NAS 100 is intended primarily for mid-size and large enterprise customers with many remote branch locations or distributed offices who desire centralized management of their IT infrastructure.

The NAS 100 supports applications such as home page directories and e-mail archival. Other applications include desktop publishing, illustration, photo imaging, and archival of document images such as legal documents, job tickets, or shipping manifests.

In comparison, the NAS 200 is designed to support storage applications for departments, workgroups, general/medium sized business customers, and service providers who require the advantages that network attached storage can provide and who need the added scalability and capacity provided by the NAS 200.

The NAS 100 is also well suited for many of the same applications as the NAS 200. The main differences between the NAS 100 and the NAS 200 are in the areas of HDD type (ATA versus SCSI), memory size, scalability and maximum

storage capacity, support of the Remote Supervisor Adapter (RSA) to augment the function of the IBM Director management functions, and the rack space requirement (1U versus 3U). Both use the same Microsoft operating system.

The 100 series products provide these important advantages:

- High reliability via redundant, hot swap hard disk drives so business operations can continue in the event of a sub-system failure.
- 250 persistent True Image data views provide quick data protection, enable backups without slowing the system and allow files and volumes to be restored quickly and accurately.
- Low-cost entry into network attached storage for customers who are reducing their use of general purpose servers.
- Multi-protocol support for CIFS (Windows), NFS (Unix), FTP, HTTP, Apple Talk and Novell file systems enables the sharing of files between and within any of these environments, eliminating the need for separate, dedicated file servers for each protocol.
- Excellent performance to enhance productivity for users in both Windows and mixed Windows/UNIX environments.
- Fully integrated, pre-loaded software suite that minimizes setup time — a minimum amount of dedicated IT resource is required for setup due to simplified installation and integration into the IP network via an easy-to-use Web browser.
- On-board Advanced System Management processor with Light Path Diagnostics, Predictive Failure Analysis and Remote Connect capabilities to help ensure system up time.

These are the highlights for the IBM TotalStorage NAS 100:

- The NAS 100 is a 1U-high, table top or 19-inch rack-mountable system (using the optional feature, Rack Mounting Kit) with a 1.26 GHz Pentium III processor, 512 MB ECC memory.
- The NAS 100 includes four 120 GB slim-high, 3.5-inch ATA HDDs. The rotational speed of the drives is 7200 rpm. They are all hot swappable; they can be removed or inserted with the power on. Operation of the NAS 100 is capable of continuing even with one of the drives removed.
- Two 10/100/1000 Mbps RJ-45 ports for Ethernet connections are located on the back of the NAS 100. These Ethernet connections are full-duplex and auto-sensing.
- The NAS 100 supports Software RAID 0, 1, and 5.
- The hot swap HDD support and support for RAID auto-rebuilding are not available on other solutions on the market.

More details can be found at:

http://www.storage.ibm.com/snetwork/nas/100/index.html

Figure 2-3 shows a picture of the IBM TotalStorage NAS 100 rack model.

Figure 2-2 The IBM Network Attached Storage 100 model R12

2.3 IBM TotalStorage Network Attached Storage 200

The IBM Network Attached Storage 200 series products, available in both tower and rack configurations, are well suited for both workgroup and departmental applications. These storage products are great choices for distributed environments such as retail, banking or insurance branch offices. In these environments they can be used for file serving and client backups, or other storage needs resulting from server consolidation. Combined with their ease of installation and easy-to-use features, the low cost per megabyte of the IBM NAS 200 series also makes them an excellent fit for NetGen environments where large amounts of inexpensive storage are required.

The tower version of the IBM NAS 200 series is scalable from 109.2 GB to 1.69 TB and is well-suited for use in workgroup environments. The rack version of the NAS 200 series scales from 109.2 GB to 3.52 TB and is suitable for support of larger workgroups or departments.

The 200 series products provide these important advantages:

- High reliability via redundant, hot swap power supplies and hard disk drives so business operations can continue in the event of a sub-system failure.
- 250 persistent True Image data views provide quick data protection, enable backups without slowing the system and allow files and volumes to be restored quickly and accurately.
- Low-cost entry into network attached storage for customers who are reducing their use of general purpose servers.
- Both affordable and scalable as storage requirements grow.
- Multi-protocol support for CIFS (Windows), NFS (Unix), FTP, HTTP, Apple Talk and Novell file systems enables the sharing of files between and within any of these environments, eliminating the need for separate, dedicated file servers for each protocol.
- Excellent performance to enhance productivity for users in both Windows and mixed Windows/UNIX environments.
- Fully integrated, pre-loaded software suite that minimizes setup time - a minimum amount of dedicated IT resource is required for setup due to simplified installation and integration into the IP network via an easy-to-use Web browser.
- On-board Advanced System Management processor with Light Path Diagnostics, Predictive Failure Analysis and Remote Connect capabilities to help ensure system up time.

All actual details can be found at:

http://www.storage.ibm.com/snetwork/nas/200/index.html

Figure 2-3 shows a picture of the IBM TotalStorage NAS 200 rack model.

Figure 2-3 The IBM 5194-225 NAS 200 rack model

2.4 IBM TotalStorage Network Attached Storage 300

The IBM Network Attached Storage 300 was designed to meet your storage requirements across many demanding environments. With rack configuration flexibility, the IBM NAS 300 is an excellent choice for use in departmental and small enterprise applications and it provides an affordable but robust solution for the storage needs of either a large branch office or department, or a small enterprise. In these environments the IBM NAS 300 is well suited for support of mission critical tasks such as accounts receivable processing, payroll support or customer service support.

The 300 offers all the same features and benefits as the 200 series products, but with these additional advantages:

- A second high-performance server that provides:
- increased levels of system reliability via a second engine, for mission critical processes, and enhanced performance so additional users in a department or small enterprise can maintain high productivity levels.
- Highly scalable storage capacity ranging from 109.2 GB to 6.61 TB providing flexibility and control over the initial purchase investment.
- Instant volume restore capability utilizing any one of a possible 250 True Images to recover a corrupted volume.

- Both affordable and scalable as storage requirements grow.
- Employs high-performance Fibre Channel communication between the engines and hard disk drives, as well as Fibre Channel technology for the hard disks themselves.

All actual details can be found at:

http://www.storage.ibm.com/snetwork/nas/300/index.html

Figure 2-4 shows a photograph of the IBM TotalStorage NAS 300 with optional extra storage.

Figure 2-4 The IBM TotalStorage NAS 300

2.5 IBM TotalStorage NAS Release 2.5

We are in the fastest-changing business ever, and we are constantly trying to improve our products for greater customer satisfaction and higher performance. Therefore, IBM announced the IBM TotalStorage NAS Version 2.5 in May 2002.

In the following sections, we describe the new features. However, do be aware that, due to the rapidly-changing nature of this business, there may already be another update available. Therefore, we recommend that you check the following Web link for the most current specifications on the IBM TotalStorage NAS product family:

http://www.storage.ibm.com/snetwork/index.html

In Release 2.5, we have all of the previous release functionality, plus these features:

- Sanergy V2.2.3 (Included with 300G only)
- WQuinn enhanced disk quota management V1.01
- ServeRAID Manager V4.84
- Storage Manager V8.2 (Included with NAS 300 only)
- TSM Agent 4.2.1.20
- Columbia Datastore Products Persistent Storage Manager (PSM) V2.2
- SNMP protocol stack
- IBM Director Services agent Version 3.1
- Domainlet feature for clustering (Included with 300/300G only)
- Diskpart Utility
- Microsoft SAK V2.01 (NAS GUI)

2.5.1 IBM TotalStorage NAS 200 Version 2.5

These are the highlights for the IBM TotalStorage NAS 200:

- New engines with 1.13 GHz Pentium III processors provide an estimated performance increase of 15 to 30%.
- 73.4 GB HDD support is provided with the same footprint, and the new platforms will continue to support 36.4 GB HDD as well.
- Quad 10/100 Megabit Ethernet Adapters, for elimination of a potential network constraint that single port adapters created — storage is more accessible and there is higher system throughput.
- Partially populated JBOD allows greater control and flexibility so storage can keep pace with growth without creating excess capacity.

These are the improvements on the NAS 200 Tower Model:

- Single Node
- Up to two 1.13 GHz Pentium III processors
- Storage 109.2 to 440.4 GB
- Maximum memory 2.5 GB.

These are the improvements on the NAS 200 Rack Model:

- Single Node
- Up to two 1.13 GHz Pentium III processors
- Storage 109.2 to 3.52 TB

2.5.2 IBM TotalStorage NAS 300 Version 2.5

These are the highlights for the IBM TotalStorage NAS 300:

- New engines with 1.13 GHz Pentium III processors provide an estimated performance increase of 15 to 30%.
- 73.4 GB HDD support is provided with the same footprint, and the new platforms will continue to support 36.4 GB HDD as well.
- Mirror OS HDD feature now permits hardware-based mirroring of operating system HDD when configured with optional, second 18 GB HDD.
- Quad 10/100 Megabit Ethernet Adapters for elimination of a potential network constraint that having a single port adapter created — storage is more accessible and there is higher system throughput.
- Partially populated JBOD allows greater control and flexibility so storage can keep pace with growth without creating excess capacity.

These are the improvements on the NAS 300 Model:

- Dual Node
- Dual 1.13 GHz Pentium III Processors
- Storage 109.2 to 6.61 TB
- Domainlet Feature for Clustering
- Storage Manager 8.0

2.5.3 IBM TotalStorage NAS Version 2.5 at a glance

Table 2-2 provides a short overview of the new highlights of the IBM TotalStorage NAS 100, 200 and 300 comparisons.

Table 2-2 Comparison of features — Release 2.5 for NAS 100, 200 and 300

Enhancement	NAS 100 Model R12	NAS 200 Model 201	NAS 300 Model 226
Processor	1.26 GHz	1.33 GHz	1.33 GHz
Dual Processor Option	no	yes	yes
Operating System	Windows Powered OS	Windows Powered OS	Windows Powered OS
Memory	512 MB	512 MB to 2.5 GB	1 GB to 3 GB
Capacity	480 GB	109 GB to 1.68 TB	109 GB to 3.51 TB
Disk Technology	ATA	SCSI	SCSI
Slots	1 PCI	5 PCI	PCI
IBM Director	yes	yes	yes
Remote Supervisor Adapter Option	no	yes	yes

2.5.4 IBM TotalStorage NAS 200 Model 25T (5194-25T) Version 2.5

These are the highlights for the IBM TotalStorage NAS 200:

- New engines with 2.4 GHz XEON processors
- Two redundant hot plugable power supplies
- 6 hot-swap drive bays
- Disk options:
  - 73 GB, Ultra 160 SCSI, ‘10.000 RPM’
  - 146 GB (Dec GA), Ultra 160 SCSI, ‘10.000 RPM’

These are the improvements on the NAS 200 Tower Model:

- Single Node
- Up to two 2.4 XEON processors
- Storage 202 GB to 7.04 TB
- Maximum memory 4.5 GB
- 10/100/1000 onboard Ethernet controller

These are the improvements on the NAS 200 Rack Model:

- Single Node with 5U height
- Up to two 2.4 XEON processors
- Storage 202 GB to 7.04 TB
- Maximum memory 4.5 GB.
- 10/100/1000 onboard Ethernet controller

2.5.5 IBM TotalStorage NAS Gateway 300 HW Version 2.5

These are the highlights for the IBM TotalStorage NAS Gateway 300:

- Dual node cluster
- New engines with 2.4 GHz XEON processors provide an performance increase
- More options in configuring Fibre Channel
- More options in configuring Ethernet

These are the improvements on the NAS 300 Model:

- Dual Node
- 5U height and 6 PCI expansion slots available per node
- Dual 2.4 GHz XEON Processors
- Embedded system management processor
- Two 560W (180-256V) hot-swap power supplies
- Storage up to 22 TB

2.6 IBM TotalStorage NAS Release 2.7

We are in the fastest-changing business ever, and we are constantly trying to improve our products for greater customer satisfaction and higher performance. Therefore, IBM announced the IBM TotalStorage NAS Version 2.7 in November 2002.

In the following sections, we describe the new features. However, do be aware that, due to the rapidly-changing nature of this business, there may already be another update available. Therefore, we recommend that you check the following Web link for the most current specifications on the IBM TotalStorage NAS product family:

http://www.storage.ibm.com/snetwork/index.html

In Release 2.7, we have all of the previous release functionality, plus these features:

- Windows 2000 service pack 2 [build 2195] + SRP1
- Microsoft Internet Explorer V5.5 + SP2
- SAK / MS GUI V2.01 Build 2204
- SFU V2.2 + QFE Q320175 + QFE Q327314
- IAACU V2.1
- PSM V2.2/Build 2262
- IBM Director Agent & UMS V3.11
- SANergy V2.2.3 (NAS Gateway 300 only)
- TSM Client V4.2.2 + PTFs
- TSM Agent V4.2.2 + PTFs
- Storage Manager for SAK V2 (WQuinn SRM V1.01/Build 163)
- UPS support

Configuration utilities:

- ServeRAID V5.11.01
- IBM Cluster Utility V1.1 (NAS Gateway only)
- IBM WWN Utility V1.1 + fix for new HBA
- NAS NT Backup V1.1
- Intel Proset II V5.3.42.0
- Alacritech SLICUSER V5.15 W/Driver V5.34
- IBM NAS Setup Tutorial V1.0

Diagnostic tools:

- Qlogic FAStT MSJ V2.0.0 Release 32
- E-gatherer V0.8.2 (c:\IBM)
- NAS Support Center (c:\Tools\Level_3)

Software ship group:

- Enablement Diskette V2.7
- Supplementary CD V2.7
- Documentation CD V2.7
- Basic LUN Expansion (Diskpart) V5.1.3553
- Domain Controller (NAS Gateway 300 only)

Part 2

Implementing the NAS appliance in your network

In Part 2 of this book we start the step-by-step walkthrough we promised:

- We explain the first steps that are necessary to get connected to the NAS appliance (Chapter 3, “Connecting to the NAS appliance” on page 45).
- Then we describe how to integrate the NAS appliances into your storage network, covering each IBM TotalStorage NAS appliance separately: (Chapter 4, “Implementing the IBM TotalStorage NAS 100” on page 57; Chapter 5, “Implementing the IBM TotalStorage NAS 200” on page 93; Chapter 6, “Implementing the IBM TotalStorage NAS 300” on page 113).
- Finally, we take a closer look at basic user and security management: (Chapter 7, “User and security management” on page 199).

3

Connecting to the NAS appliance

In this chapter we discuss the first thing that must be done after unpacking the device and plugging it into the network: you need to connect into the NAS device in order to start managing it.

We cover these topics:

- First steps
- Assigning IP address to the appliance with IAACU
- Connecting to the appliance
- NAS Setup Navigator overview

3.1 First steps

For the first time you connect to a IBM TotalStorage NAS device, you only can do it in two ways:

- **Locally:** Using a keyboard, mouse, and monitor
- **Remotely:** Using IBM Advanced Appliance Configuration Utility (IAACU) to set up the networking parameters

Using a keyboard, mouse, and monitor

The NAS 200 and NAS 300 units are designed as headless appliances. However, you can connect a keyboard, mouse, and monitor directly to the system. This is sometimes more convenient than using the remote utility to configure the NAS system. For logging into the system, use the user: *Administrator* and the initial password: *password*. For security purposes, we highly recommend changing the user and password immediately.

Using IBM Advanced Appliance Configuration Utility (IAACU)

You can install the IBM Advanced Appliance Configuration Utility from the IBM TotalStorage Network Attached Storage System Supplementary CD, which is delivered with the appliance.

This tool should be installed on a workstation connected to the same network segment with the NAS unit. While using this tool, you are able to discover any NAS appliances in the same network and see their characteristics and specific data.

3.2 Assigning IP address to the appliance with IAACU

The objective is to assign the network settings to the appliance that allow you to continue configuring the NAS device with Terminal Services or a Web Browser.

The process for providing the network configuration for the first time is made using groups.

In Figure 3-1 you can see that the NAS appliances are listed under the unassigned appliances.

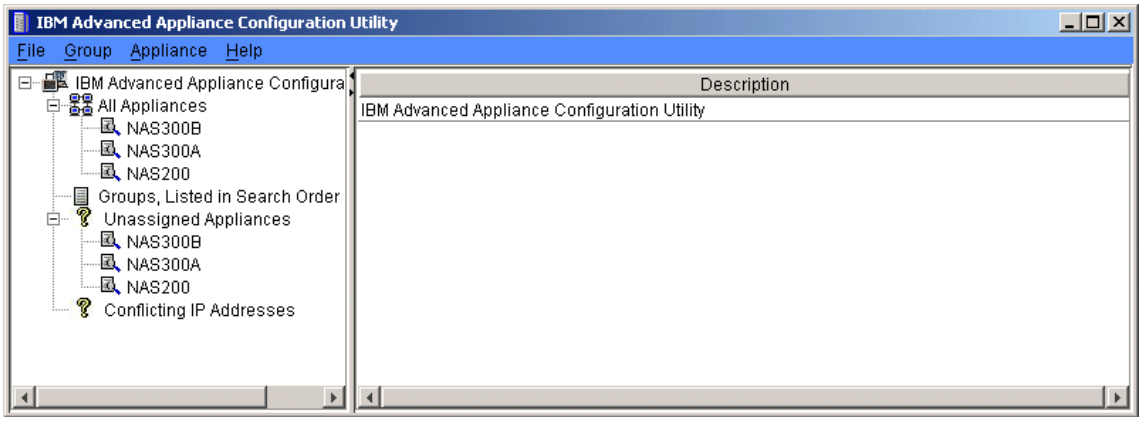

Figure 3-1 IAACU main screen

These are the steps you need to follow:

1. To create a new group, click **Group > Create Group**.
2. In the Group Type Setup dialog box, supply the name of the group, the type of the appliances to be included (IBM NAS devices in our case), and for the Network Setting question, “Allow this Group to assign the network settings?”, select the **Yes** radio button (Figure 3-2).
3. Click **Next** to continue.

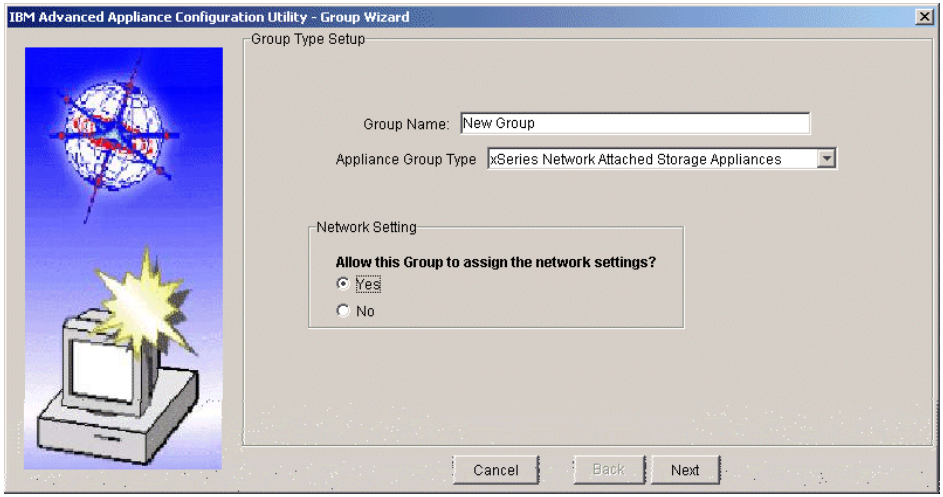

Figure 3-2 Group Type Setup dialog box

4. Assign the TCPIP settings for the group, including the group IP address range, as shown in Figure 3-3.

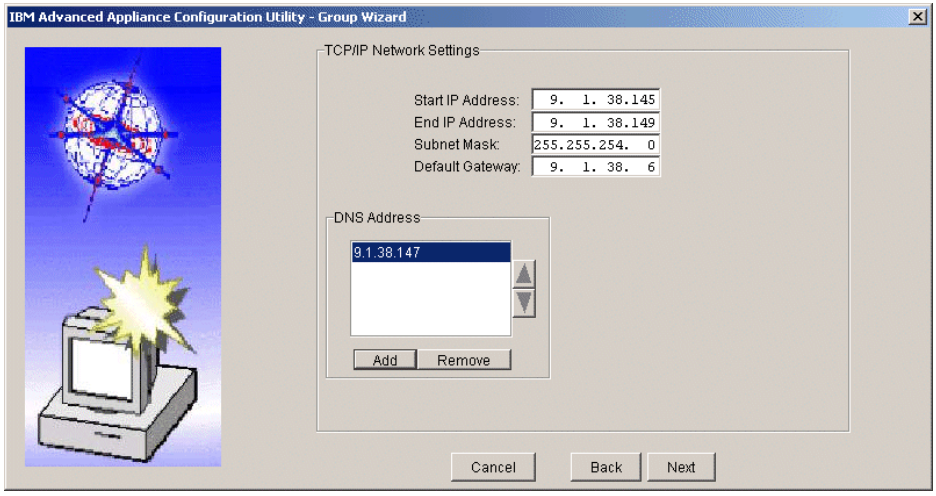

Figure 3-3 Setting the IP address

5. Click **Next** to continue.
6. In the Host Name Assignment Type screen (Figure 3-4) you can define the policy for assigning names:
   - Use Current Host name.
   - Use the “i” prefix and model-serial number.
   - Use a predefined prefix and model-serial number.
7. Provide the domain name for the NAS Device.

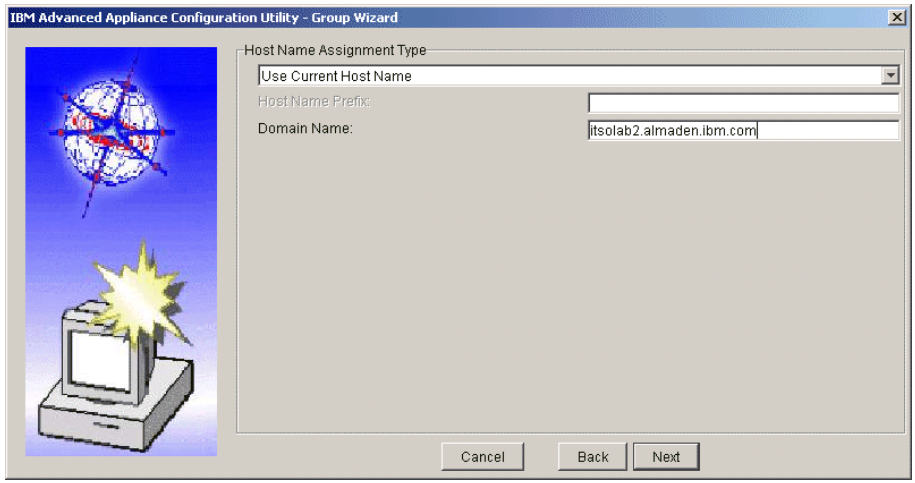

Figure 3-4 Defining the name policy

8. The next screen in the wizard (Figure 3-5) asks you if you want to enable reprovisioning. This means that is possible to maintain an FTP site with OS images of the appliances to replace the current OS in the appliance.

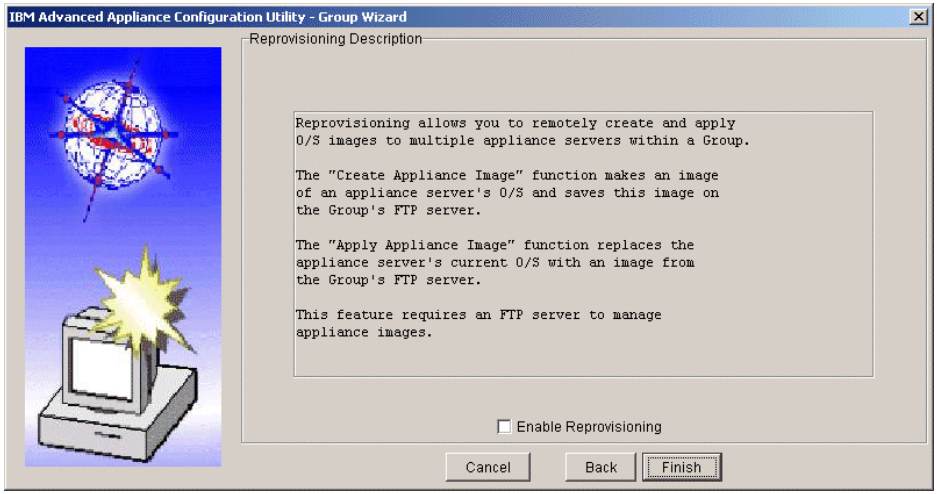

Figure 3-5 Enabling provisioning

9. Click **Finish** to create the group. You can see that the NAS 200 box has been moved to the New Group and has the first free IP address specified in the IP range for the group (Figure 3-6).

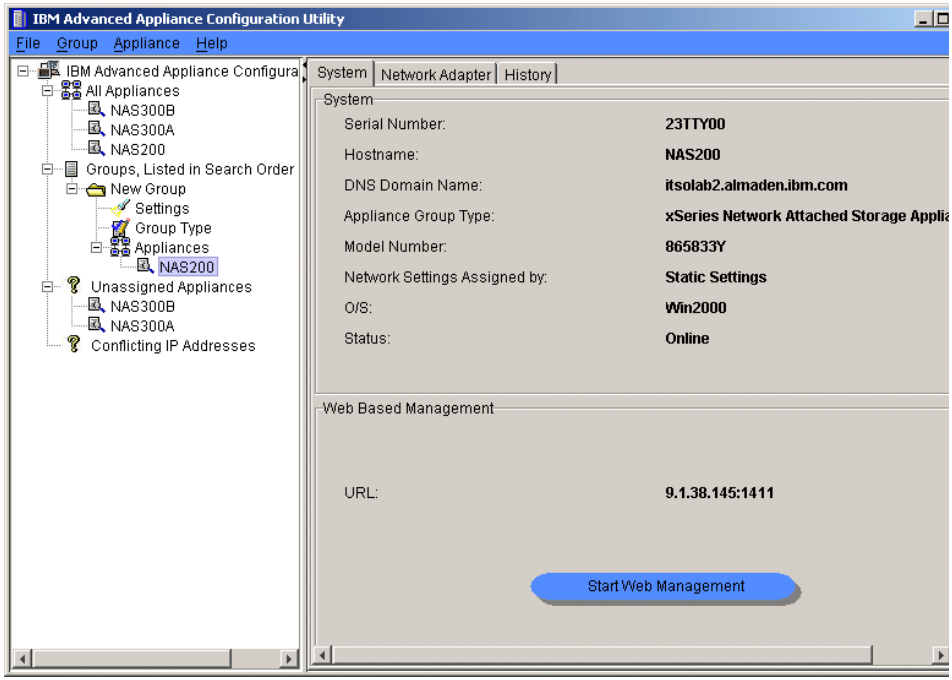

Figure 3-6 NAS Appliance with IP configured

Now you can start to manage the devices by selecting the appliance that you want to manage and clicking the **Start Web Manager** button. This opens a connection to the NAS device through the TCP/IP port 1411 as described in 3.3, “Connecting to the appliance” on page 52.

Note: We recommend that you connect the on-planar Ethernet port to the network that you use for the administration of the NAS device. This procedure might not work when the on-planar Ethernet port is not connected.

3.2.1 Universal Manageability Services

Universal Manageability Services (UM Services) is a Windows application that acts as both a stand-alone management tool for the system it is installed on, and a client to IBM Director. As a Director client, it receives and sends information to the Director Server as controlled from the IBM Director console. As a stand-alone tool, it provides a Web-browser based interface and a Microsoft Management Console (MMC) interface, where you can view the system status, perform certain management tasks, and configure alerts.

UM Services is a lightweight client that resides on each managed computer system. It comes preinstalled on all IBM NAS appliances.

Note: At some points in the documentation, UM Services is referred to as the IBM Director Agent. For more information, refer to 9.1, “Description of IBM Director” on page 286.

You can start UM Services to manage the NAS appliance either from the IAACU by clicking the **Start Web Management** button (Figure 3-6 on page 50), or directly by opening the browser with the appliance’s IP address and port 1411 specified.

As a prerequisite for your Web browser interface, you must use Microsoft Internet Explorer Version 5.0 or higher. You also need to install Java Swing/JFC and XML support on your system (Figure 3-7).

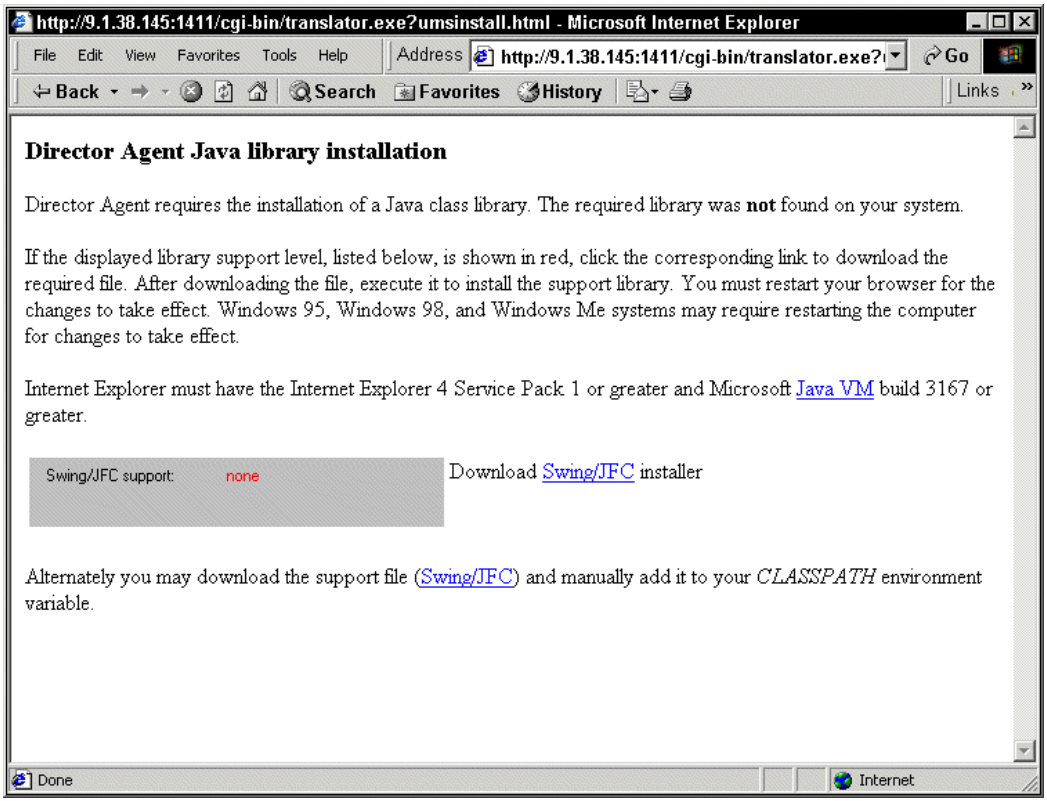

Figure 3-7 UM services Java Library installation

After installing the components, close your Web browser and start Web Manager again. Now you should be able to explore all IBM appliances and their characteristics, as shown in Figure 3-8.

Remember to use the administrative account to get access to the NAS appliance.

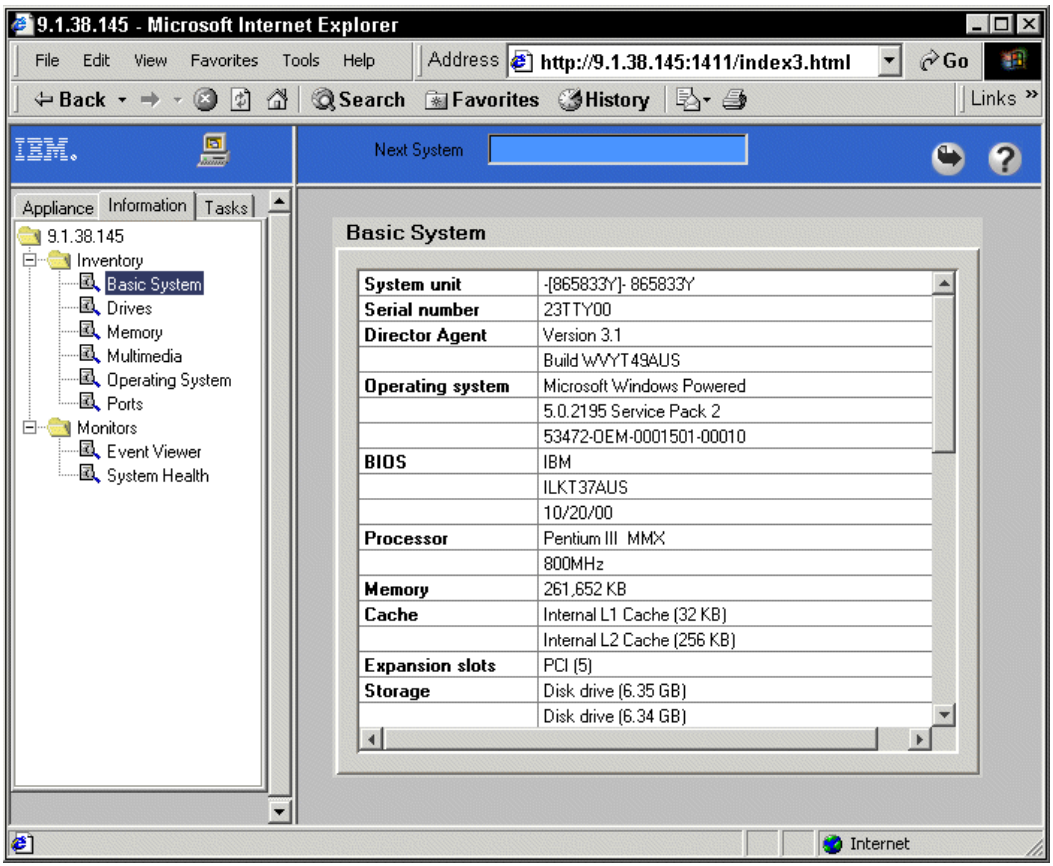

Figure 3-8 NAS characteristics

3.3 Connecting to the appliance

After you set up the network parameters, either with the IAACU Utility or directly attached to the NAS 300 device, you can use other tools to manage the NAS device:

- Terminal Service Client
- Web browser Interface (ports 1411 and 8099)

3.3.1 Using Terminal Services

The Windows Powered OS that is pre-installed on the NAS appliance is also running Windows Terminal Services. You can use a Terminal Services Web Connection from your Web interface by connecting to:

```
http://YourIpAdressAppliance:1411
```

Then click in Terminal Services or the Terminal Service Client to get access to the NAS unit.

The Terminal Services Client is included on the TotalStorage Network Attached Storage System Supplementary CD. A sample connection is shown in Figure 3-9.

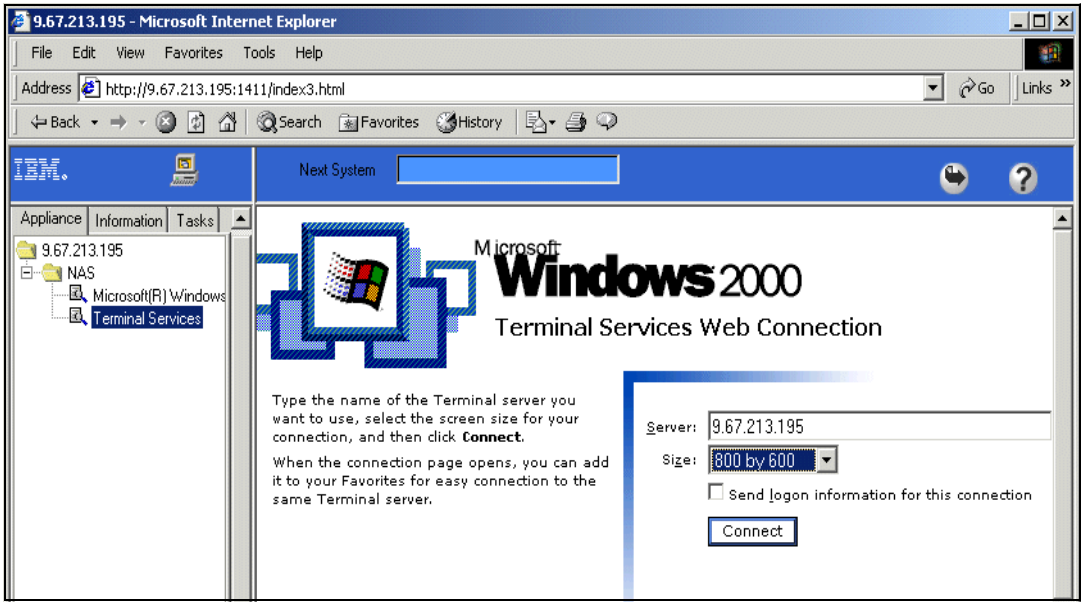

Figure 3-9 Using the http://YourIpAdressAppliance:1411 connection

3.3.2 Getting access with Internet Explorer

As we mentioned before, all administration tasks only work using Microsoft Internet Explorer 5.0 or higher. In order to have access to a NAS appliance, you can type this URL:

```
http://YourIpAddressAppliance:8099
```

Use an IP address that matches your NAS unit for “YourIpAddressAppliance”. The TCP/IP socket that is used for NAS is port 8099. After executing the URL, the system will prompt you for the username and password. See the example in Figure 3-10.

Figure 3-10 NAS login screen

After a successful login, you will see the Web interface shown in Figure 3-11.

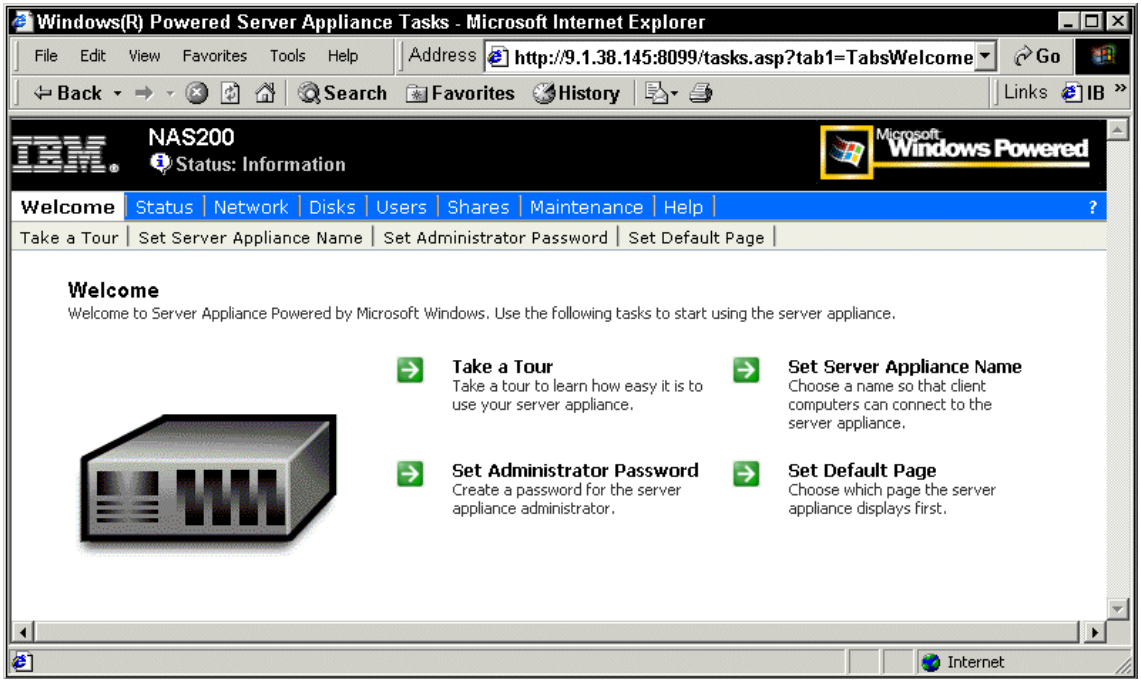

Figure 3-11 NAS access using Microsoft Internet Explorer

3.4 NAS Setup Navigator overview

The NAS Setup Navigator was created to assist you in setting up your IBM TotalStorage NAS appliance and is new to Release V2.5. It appears as an icon on the desktop of all the NAS appliances. The easiest way to use this tool is to follow the screens in the order presented. Should a question arise during the installation of your NAS appliance, you can refer to other sections in this chapter which describe in detail how to perform a specific task. Also, at the bottom of many of the navigator screens, there are links with additional information.

To use the NAS Setup Navigator, you need to either connect to the appropriate NAS appliance locally (with keyboard, mouse, and video), through Microsoft's Terminal Services, or you can use the Web interface and then Terminal Services.

You can use the IAACU tool described in 3.2, "Assigning IP address to the appliance with IAACU" on page 46 to supply an IP address to the NAS device for remote connectivity.

Depending on your system, the NAS Setup Navigator can assist you in accomplishing many tasks. These are the navigator screens:

- Information and Setup Options
- Configuration Prerequisites
- Clustering Overview (NAS300)
- System Language
- Administrator Password
- Date and Time
- Network Identification
- Public LAN Settings
- Configure Pooled Storage
- Redefining RAID Arrays and Logical Drives (NAS200)
- Writing Disk Signatures
- Creating Partitions
- Verifying Disk Health
- Services for UNIX
- Setting up Server for NFS
- Setting up Gateway for NFS
- Server for PCNFS
- User Name Mapping
- User and Security Management
- Setting Up Windows Users and Groups
- Share Pooled Storage
- File Share for Windows Clients
- File Share for UNIX Clients
- Cluster Setup (NAS300)

- Private LAN settings (NAS300)
- Active Directory Domain Controller Setup (NAS300 Domainlet feature)
- Administering the Cluster (NAS300)
- Joining the Cluster (NAS300)
- Configure Cluster Properties (NAS300)
- Creating Clustered Windows File Shares (NAS300)
- Creating Clustered UNIX File Shares (NAS300)
- Cluster Resource Balancing (NAS300)
- Failover and Failback Setup (NAS300)
- FAStT200 (RAID) Setup (NAS300)

As you can see, the NAS Setup Navigator can aid in the setup of your NAS appliance. In the following sections we show you how to use the Navigator to configure an IBM TotalStorage NAS 200 and an IBM TotalStorage NAS 300 clustered system.

4

Implementing the IBM TotalStorage NAS 100

In this chapter we explain how to configure the IBM TotalStorage NAS 100 from the very beginning.

We cover these topics:

- Initial configuration
- Using the Navigator to set up the NAS 100
- Ethernet adapter teaming

4.1 Initial configuration

The NAS 100 device is designed as a headless appliance. It is not possible to connect the monitor directly to the system, because the device doesn't have a video card installed, and no PS2 keyboard and mouse ports are integrated. However, the system supports a USB connected keyboard. So the only access method for configuration is through the network.

There is also a separate IBM Redbook available regarding the NAS 100 appliance: *The IBM TotalStorage NAS 100 Integration Guide*, SG24-6913. See the URL:
http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246913.html?Open

To be able to work with the NAS 100 device, you have to ensure that network recognizes the new appliance. The device has two networking interfaces (LAN ports) integrated. LAN port 1 is preconfigured for dynamic address assignment (DHCP) and LAN port 2 has a static address preassigned. Both networking ports are positioned on the back side of the appliance. You can use any port for the initial configuration.

Important: The new appliance's default host name is IBM5190-*xxxxxxx*, where *xxxxxxx* is the serial number, located on the back panel of the NAS 100 device.

Accessing the NAS 100 using DHCP address on LAN port 1

If the NAS 100 device is connected to a network where a DHCP Server exists and serves out available dynamic addresses, the new appliance will pick up an available IP address (as well as subnet mask, DNS Server address, and gateway address) and will be ready to connect to. If there is no DHCP Server available, the IAACU utility can be installed on another machine in the same network. When IAACU receives a DHCP request from the new NAS 100 appliance, it will respond to it and supply an IP address.

Accessing the NAS 100 using static address on LAN port 2

There is a static IP address of 192.168.0.1 with a subnet mask of 255.255.255.0 and a default gateway address of 192.168.0.254 predefined on the LAN port 2. You can access the NAS 100 appliance from a workstation in the same network segment, or a router must be correctly configured.

4.1.1 Methods for setting up the NAS 100 device

There are two methods to configure the basic settings:

- **Easy Setup, using a browser:** This is recommended for initial configuration.
- **IBM Advanced Appliance Configuration Utility (IAACU):** This is an alternative method to set up the networking parameters.

Using Easy Setup for initial configuration

Easy Setup is part of a complete configuration and management tool called Windows Powered Server Appliance Tasks, which comes preloaded with the NAS 100 appliance. Easy Setup is intended to be used for setting or changing the administrator password, configuring the network settings, and defining the host name and share. To access it, use a Microsoft Internet Explorer browser.

Note: To be able to use different configuration and management tools to manage the NAS 100 appliance, it is recommended that you set Microsoft Internet Explorer 5.x (or later) as your default browser.

1. Open the browser and enter the system name, then continue by clicking **Next**, as shown in Figure 4-1:
   - For a secured port, type:

     `https://IBM5190-`*`xxxxxxx`*`:8098`
   - For an unsecured port, type:

     `http://IBM5190-`*`xxxxxxx`*`:8099`

   In each of these expressions, *xxxxxxx* is the serial number on the back of the NAS 100 appliance.

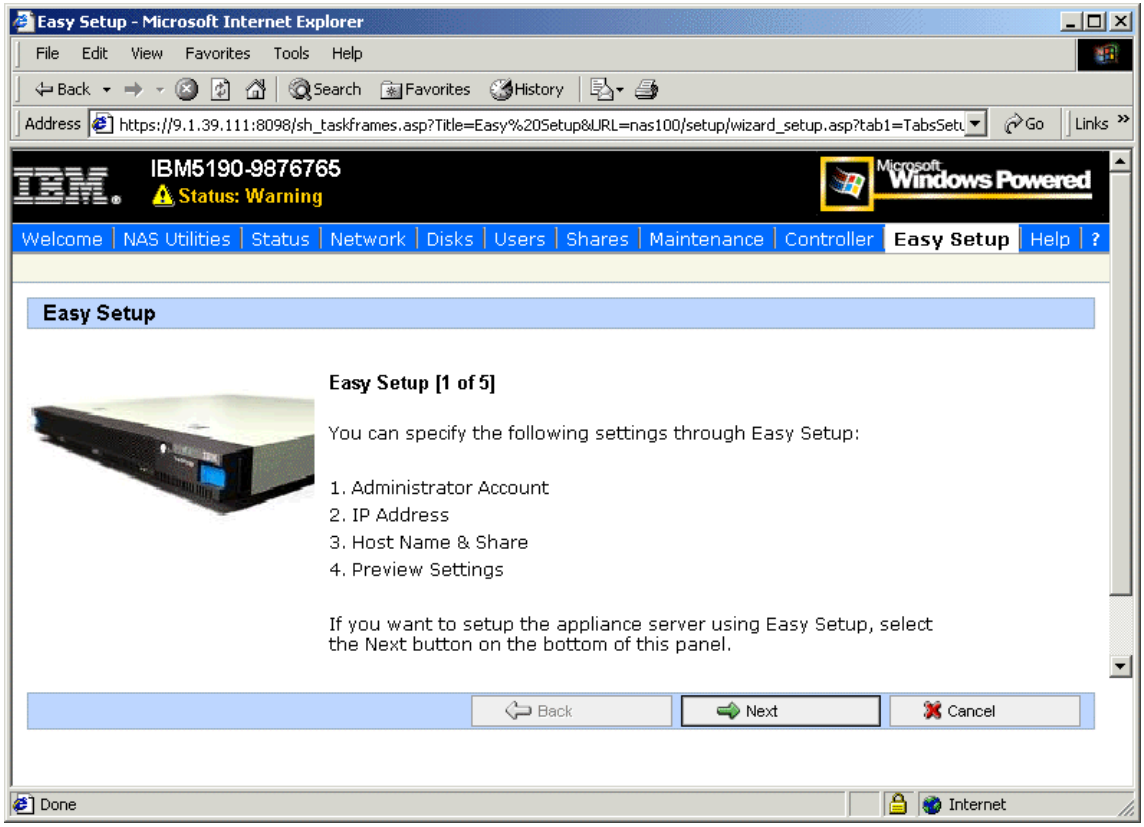

Figure 4-1 Starting Easy setup

2. Enter the administrator user name and password.

Note: The default user/password combination for the NAS 1000 system is: *Administrator* and *password*. For security purposes, it is recommended to change the initial password immediately.

3. After you change it, continue by clicking **Next**, as shown in Figure 4-2.

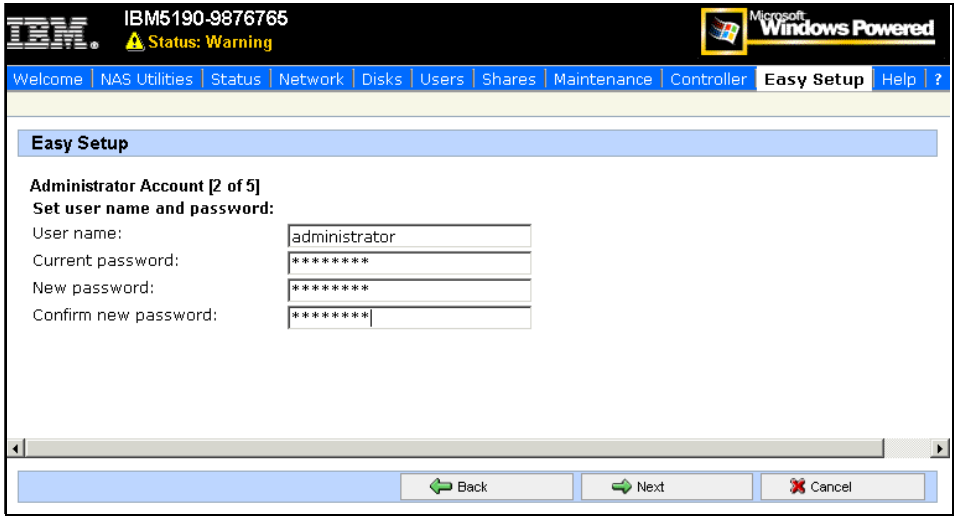

Figure 4-2 Administrator Account

4. The Easy Setup IP Address panel appears. Here you can set the configuration mode (DHCP or static) for both LAN ports and configure the IP address, subnet mask, default gateway, and the DNS server address. Notice that the Easy Setup senses whether the network cable is connected to the LAN port and displays this under the heading Status (Figure 4-3).

Note: If the IP address for a particular port is static, the DNS server address must also be specified statically.

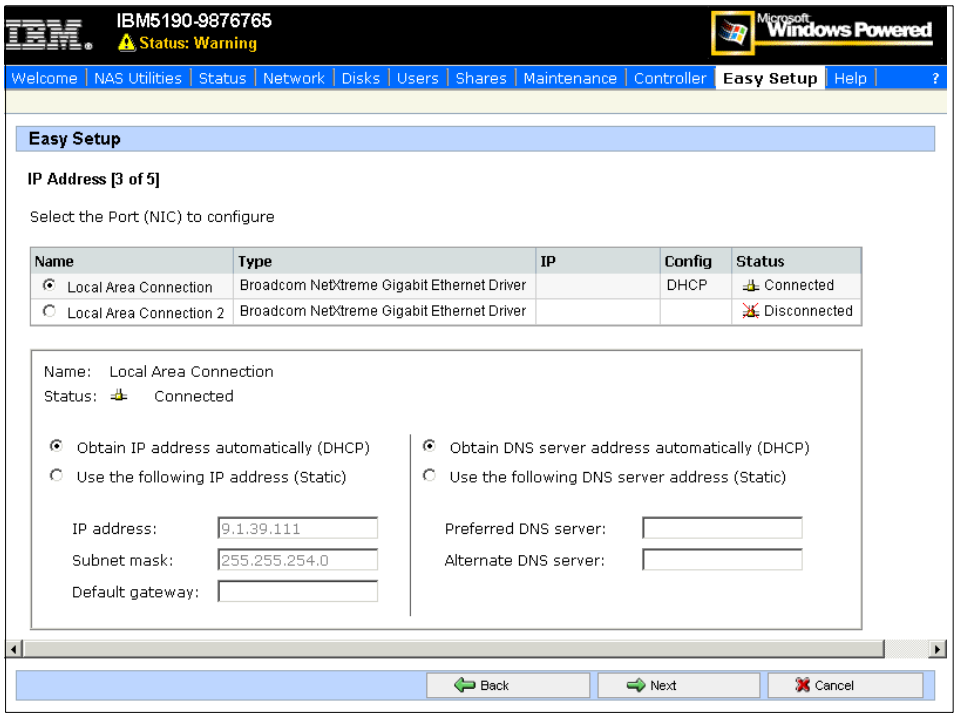

Figure 4-3 Setting IP addresses

5. After clicking **Next**, the Host Name & Share panel is presented. Here you can change the default host name. This is not recommended, though, as you will not be able to use IAACU. In the bottom part of this panel you can configure a share. Enter the Share name, path, and type, as shown in Figure 4-4. If you need to create and manage more than one share, use the Shares tab at the top of this panel. When you are done, click **Next**.

Note: If you select UNIX (NFS) checkbox under the Share type, you need to map the UNIX user to the Windows user using Services for UNIX. If you don't do this, your user will only have limited privileges. For more information, see Chapter 10, "Cross platform storage" on page 353.

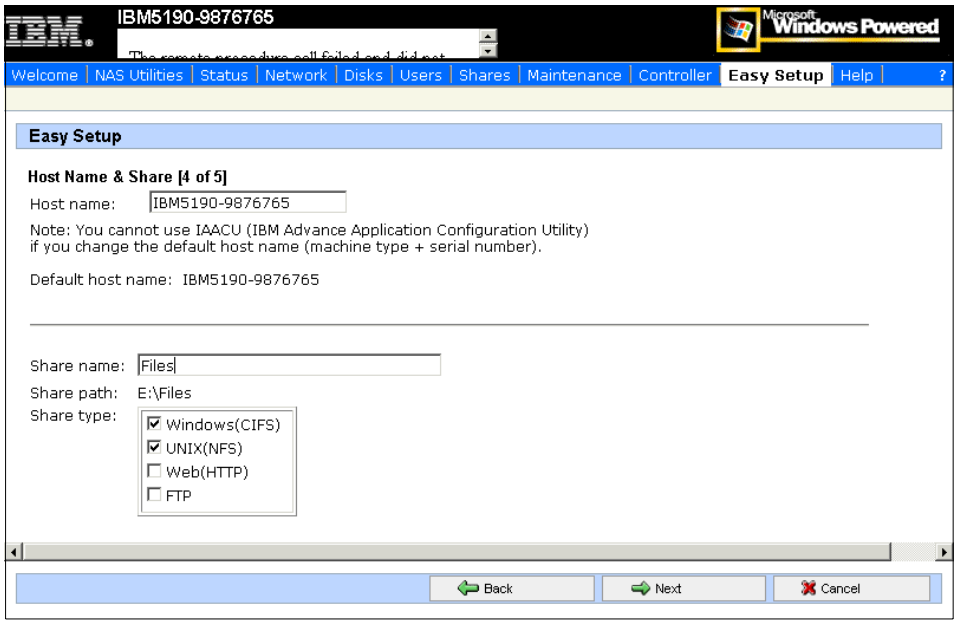

Figure 4-4 Changing a host name and creating a Share

6. In the Summary panel you can preview the settings. The ones which were changed are shown in shaded boxes. If you are satisfied with them, you can accept them by clicking **Finish**, as shown in Figure 4-5.

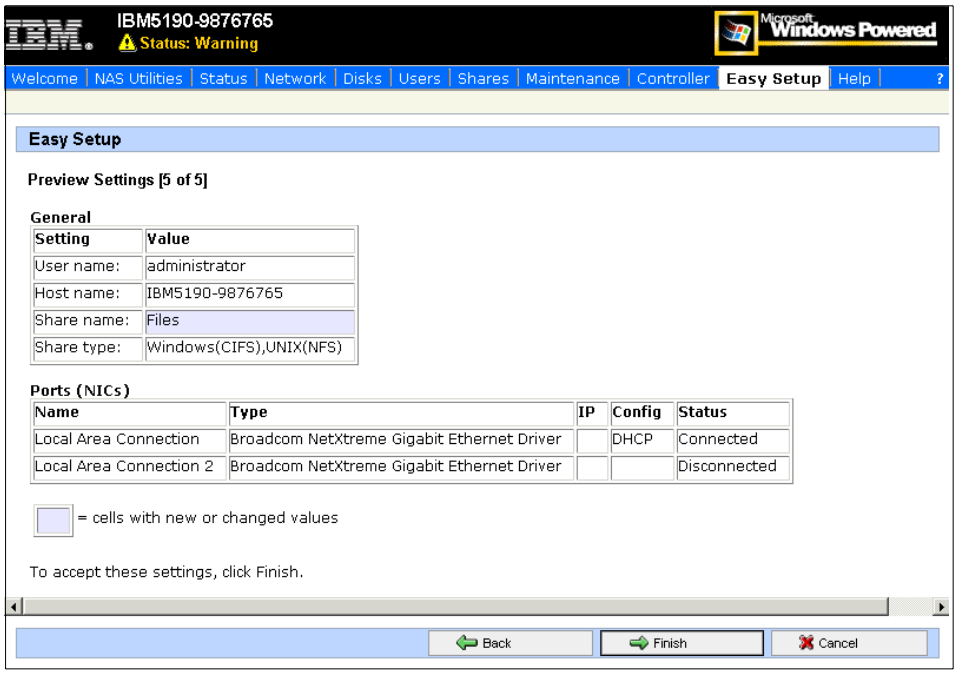

Figure 4-5 Preview and confirm settings

If you want to make a correction to any settings, you can always click **Back** to return to the appropriate panel and change the settings. After you are done making changes, you save them by clicking **Finish**.

Using IBM Advanced Appliance Configuration Utility (IAACU)

Another useful tool for initial configuration of NAS 100 appliance is the IBM Advanced Appliance Configuration Utility. It consists of Console and Agent code. The Advanced Appliance Configuration Agent is preinstalled as a Windows Powered OS service on the NAS 100 appliance. The Console can be installed on a workstation connected to the same network segment as the NAS 100 unit. The Console code is available on the Supplementary CD 1, which is delivered together with the appliance.

You can use the IAACU to automatically discover any NAS appliances in the same network and modify their specific settings. This tool comes very handy when there is no DHCP Server available in the network and you don't know if there was a static address configured for one of the LAN ports of the NAS 100 appliance.

You can also use the IAACU to start UM Services on the appliance, enabling you to perform several systems-management tasks.

Note: The IAACU console can display only one NAS 100 appliance attached to the network. Therefore, when configuring multiple NAS 100 appliances using IAACU, you need to attach and configure one appliance to the network at a time. To configure the remaining appliances, remove the appliance from the network or power off each appliance after it has been configured.

Assigning IP Address to the appliance with IAACU

Assigning the network settings to the NAS appliance will allow you to continue configuring the NAS device with Terminal Services or a Web Browser.

IAACU enables you to group all discovered appliances into function-specific families. Appliances are added to a family based on the appliance type. Appliances that perform the same function appear in the same family. Additionally, you can divide appliances into different groups:

1. After IAACU is installed on a local workstation, you can start it by clicking **Start —> Programs —> IBM Advanced Appliance Configuration Utility**. The IAACU main screen will open and all discovered NAS appliances will be listed under Unassigned appliances, as you can see in Figure 4-6.

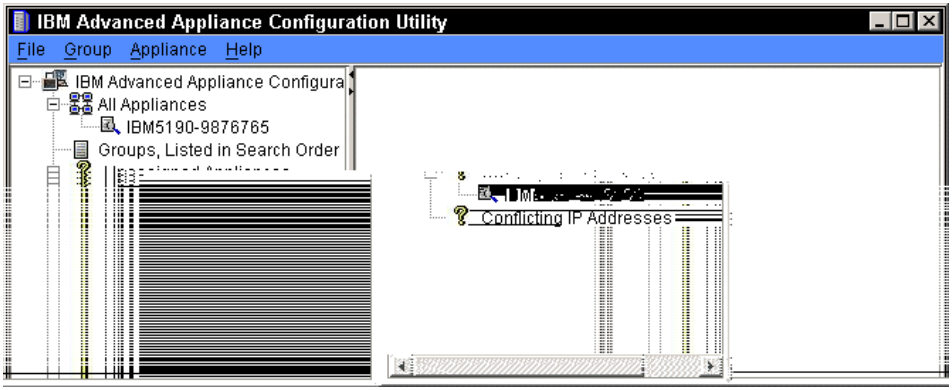

Figure 4-6 IAACU main screen

2. To create a new group, click **Group —> Create Group**.
3. In the Group Type Setup dialog box, supply the name of the group, then the type of the appliances to be included (IBM NAS devices in our case), and for the Network Settings question, “Allow this Group to assign the network settings?”, select the **Yes** radio button (Figure 4-7). Click **Next** to continue.

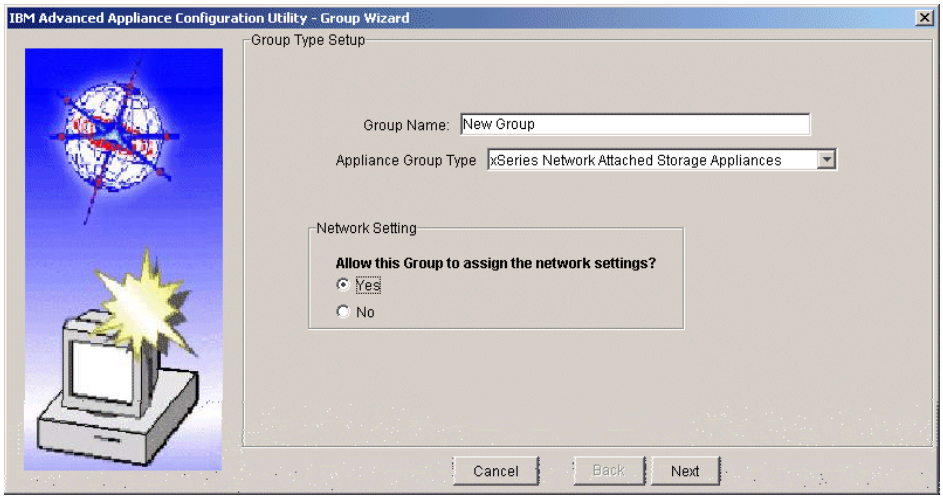

Figure 4-7 Group Type Setup dialog box

4. Assign the TCPIP settings for the group, including the group IP address range and click **Next** to continue, as shown in Figure 4-8.

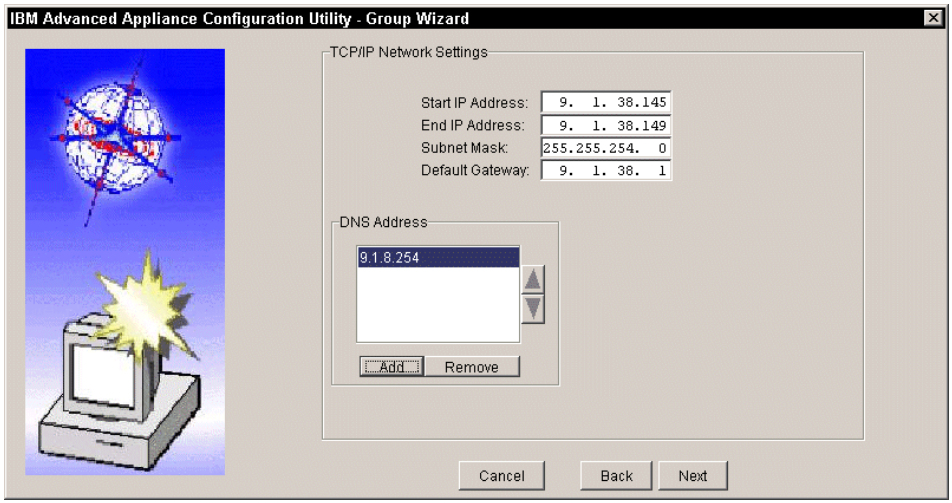

Figure 4-8 Setting the IP address

5. In the Host Name Assignment Type screen (Figure 4-9) you can define the policy for assigning names:
   - Use Current Host name.
   - Use the “i” prefix and model-serial number.
   - Use a predefined prefix and model-serial number.

- In the bottom part of the window you can provide the domain name into which the new NAS device will be placed. When done, continue by clicking **Next**.

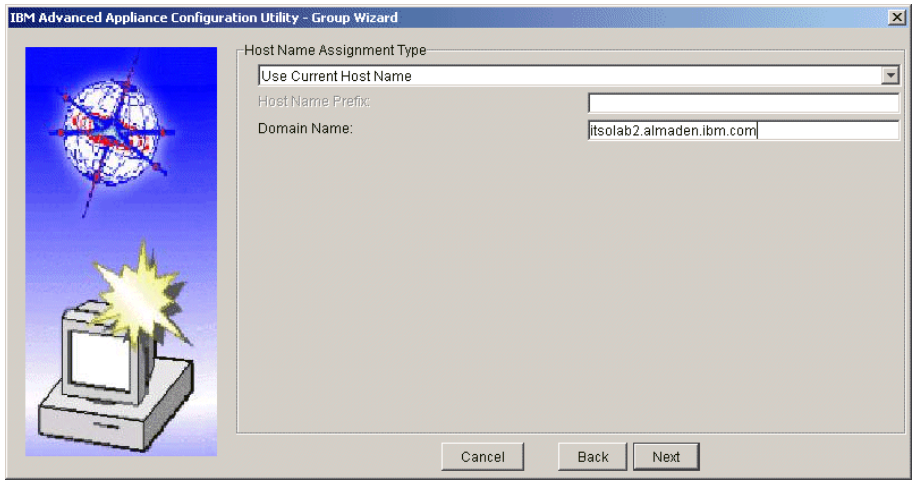

Figure 4-9 Defining the name policy and Domain name

6. Another of IAACU's options is Reprovisioning. It allows you to maintain an FTP server with appliance OS images. If appliance OS is damaged, these images can be used to replace the appliance's current OS. To enable this function, select the Enable Reprovisioning checkbox and continue by clicking **Next** (Figure 4-10).

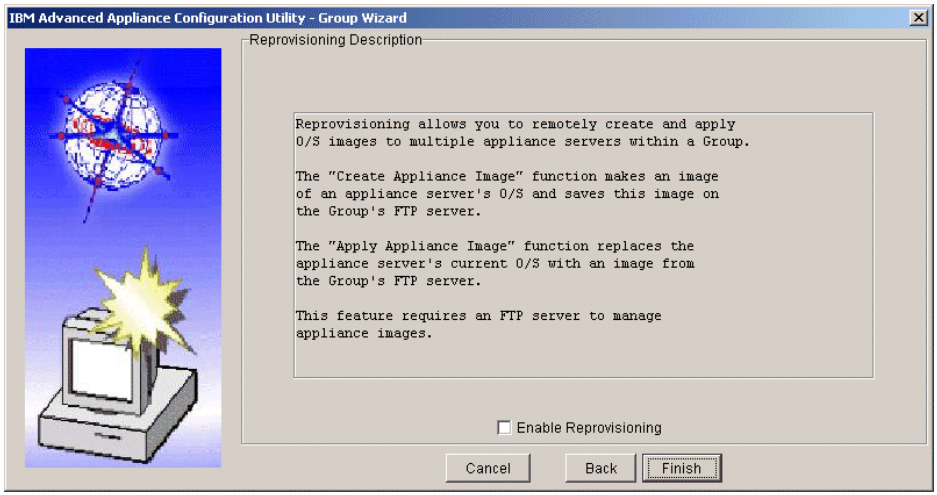

Figure 4-10 Enabling reprovisioning

7. Click **Finish** to create the group. You can see that the NAS 100 device has been moved to the New Group. The first available IP address of the specified range for this group has been assigned to it (Figure 4-11).

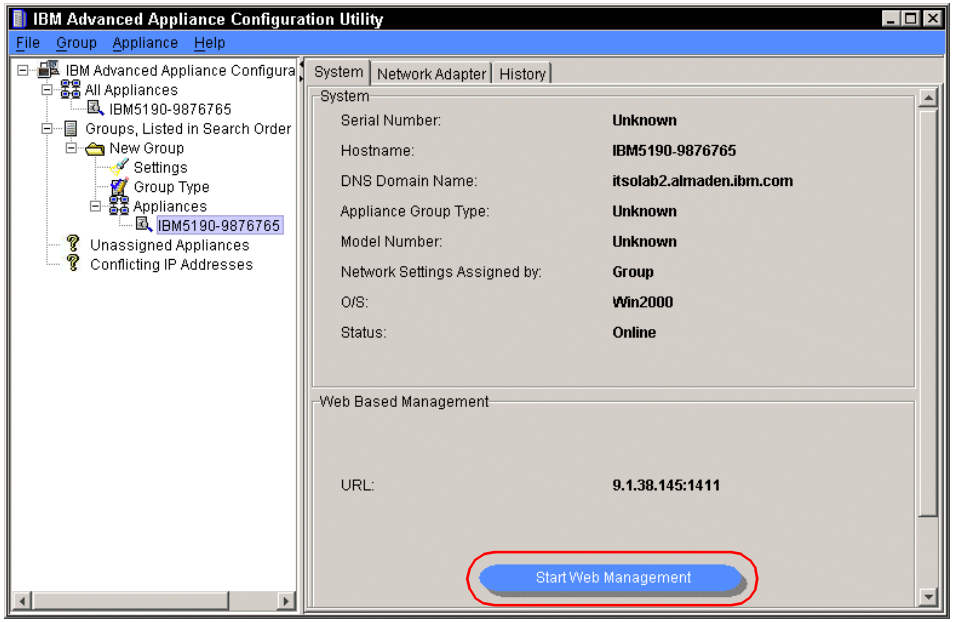

Figure 4-11 NAS Appliance with network settings applied

4.2 Using the Navigator to set up the NAS 100

In this section we demonstrate how to use the NAS Setup Navigator to set up the IBM TotalStorage NAS 100. As mentioned before, the wizard guides you through each step. Many of the NAS Setup Navigator screens have links in them that open additional configuration screens to make the process easier. To set up the NAS 100, follow these steps.

4.2.1 Basic configuration

These are the steps for basic configuration:

1. Connect to the NAS 100 appliance using any of the Terminal Services methods explained in 3.3.1, “Using Terminal Services” on page 53 and logon with an administrative account (default is **Administrator** and **password**).
2. Open the NAS Setup Navigator. The icon is located on the desktop of the IBM TotalStorage NAS appliance. The Getting Started screen appears first and explains the navigator. To advance, click the **Forward** button, which brings you to the Information and Setup Options screen, as shown in Figure 4-12.

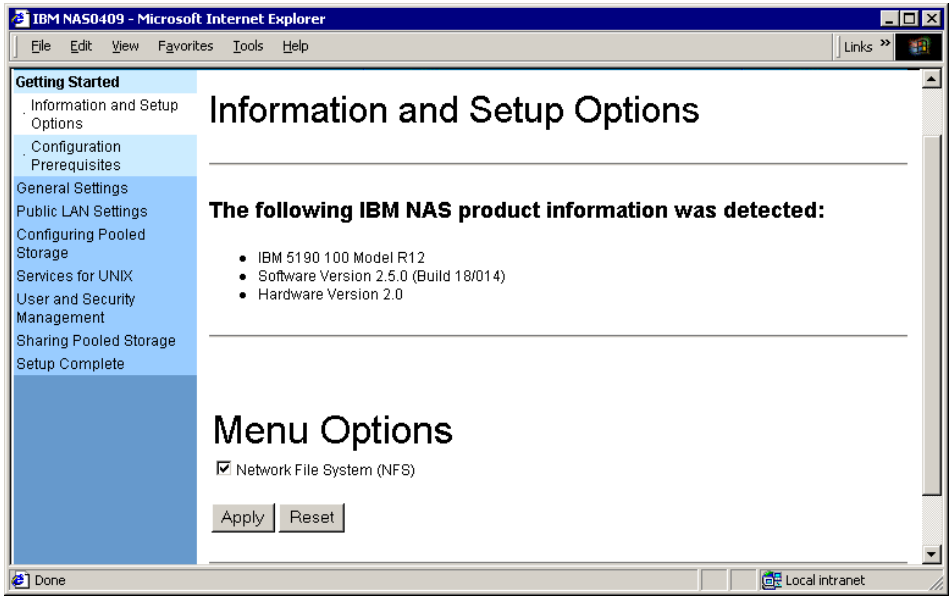

Figure 4-12 NAS Setup Navigator — Information and Setup Options screen

Note: Checking the NFS box will result in bringing up screens later on in the wizard process for installing support for NFS.

3. Select the Network File System box if support for NFS is required; otherwise, leave it unchecked and click **Apply**.

   Clicking the **Apply** button will refresh the screen and bring you back to the top of the Information and Setup Options screen.

4. Click the **Forward** button. This will bring up the Configuration Prerequisites page. This is an information-only screen and is here to help you gather and record information concerning the NAS appliance. When you are finished, you can click the **Forward** button.

Note: If the navigator screen has a link (which will be in a different color and underlined) and directs you to click it, a file download screen will come up as shown in Figure 4-13 and ask you to either open the file from the current location or save to disk. Always select the radio button, **Open this file from its current location**, and click **OK**. This will bring up a new Windows screen.

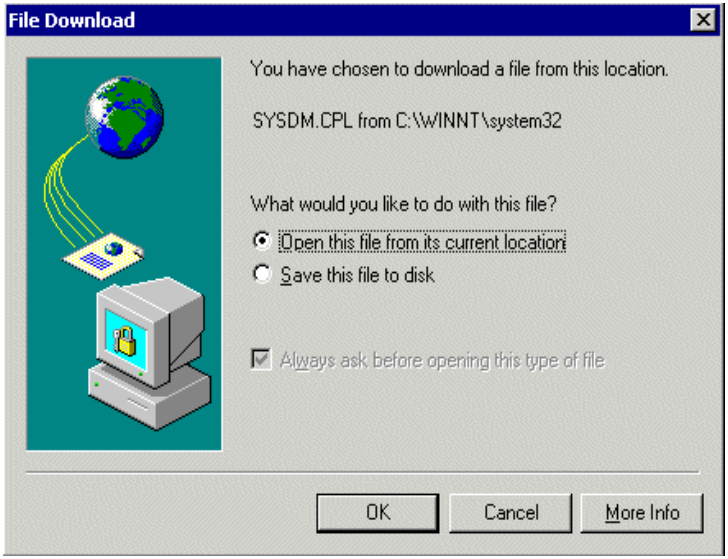

Figure 4-13 NAS Setup Navigator — additional screen

5. After the Configuration Prerequisites screen, the next set of NAS Setup Navigator screens will guide you through setting up some General Settings for your NAS appliance. Follow the steps and screens and configure the NAS appliance with the appropriate information for System Language, Administrator Password, and Date and Time. When you are done with each screen, click the **Forward** button to advance. A typical NAS Setup Navigator screen is shown in Figure 4-14.

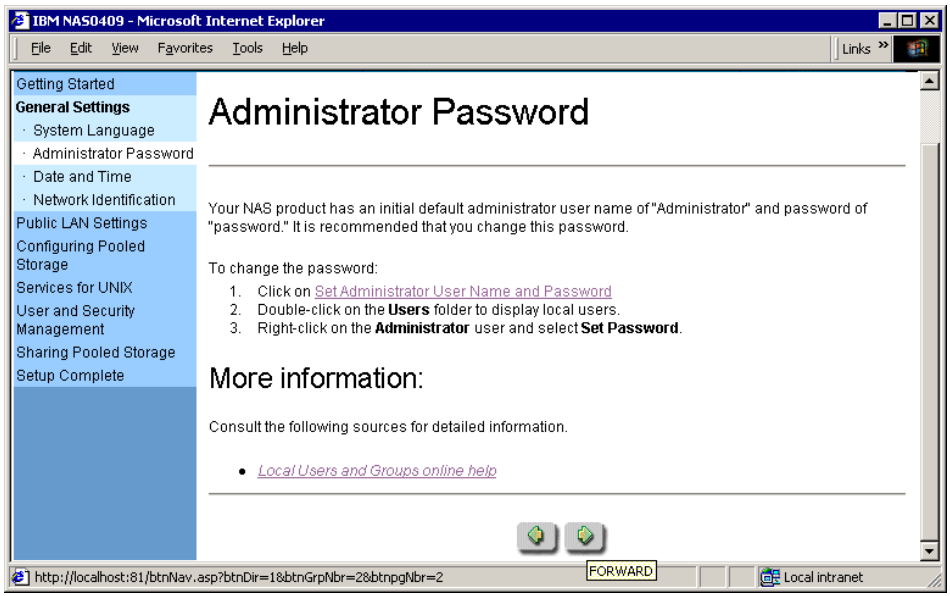

Figure 4-14 NAS Setup Navigator — Administrator Password screen

Note: As you can see in Figure 4-14, the link **Set Administrator User Name and Password** is a different color. Clicking it will bring up the file download screen. After you have told the file download screen to **open this file from its current location**, another screen will open (in this case the Local Users and Groups screen, shown in Figure 4-15) which will allow you to alter the configuration specified by the link. After making the necessary changes, close that screen (not the NAS Setup Navigator) and continue on with the Navigator.

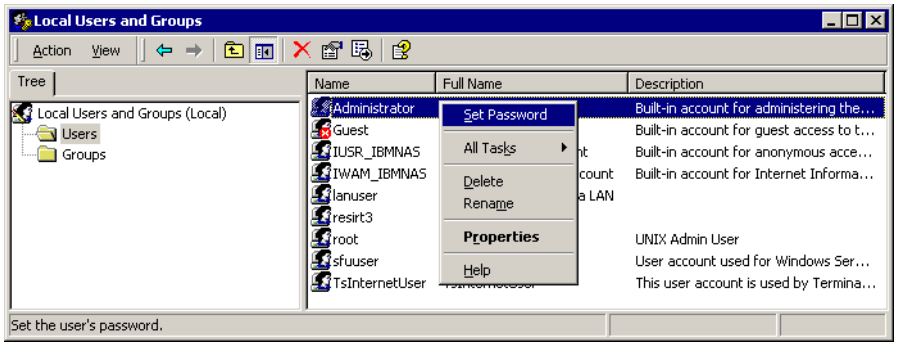

Figure 4-15 Local Users and Groups screen

6. Configure Network Identification on your NAS 100 by clicking the Network Identification link. This will bring up the system properties window. Click the **Network Identification** tab and then **Properties**. Now you have the option to join the NAS 100 appliance to your environment, as shown in Figure 4-16.

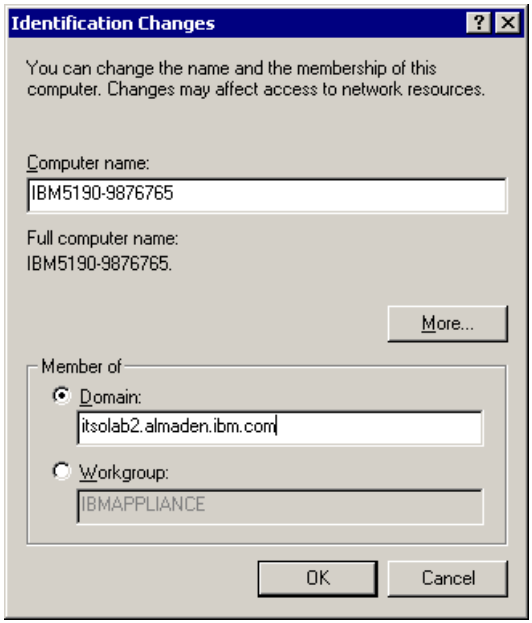

Figure 4-16 Network identification — Domain membership

7. If you want the NAS 100 to be a storage device for your workgroup, check **Workgroup**, and type in the name of your workgroup. Be aware that all security management within such a workgroup is local. This means you have to administer all user accounts on the NAS 100. Also be aware that the default behavior when sharing a network drive with Windows is to grant all users full access to the data.
8. If you want the NAS 100 to be part of a domain, then check **domain** and enter the name of the domain you would like to join. A popup window will ask for the username and password of a domain administrator or equivalent. Either way, when you are finished, click the **OK** button. A reboot is required.
9. You will now continue with the NAS Setup Navigator by configuring Public LAN settings. Configure your IBM TotalStorage NAS appliance with the appropriate information for Public LAN Settings by clicking the Network and Dial-up Connections link. This will bring up the Network and Dial-up Connections window as shown in Figure 4-17.
10. Right-click the connection you want to work with and choose **Properties**.

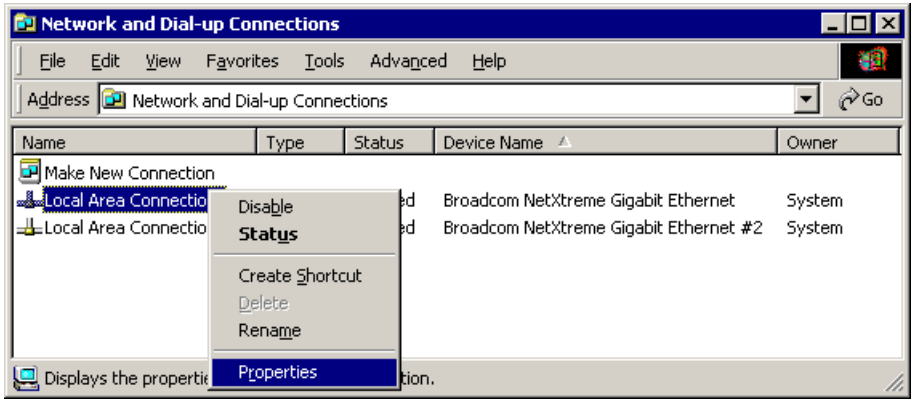

Figure 4-17 Network configuration window

11. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 4-18).

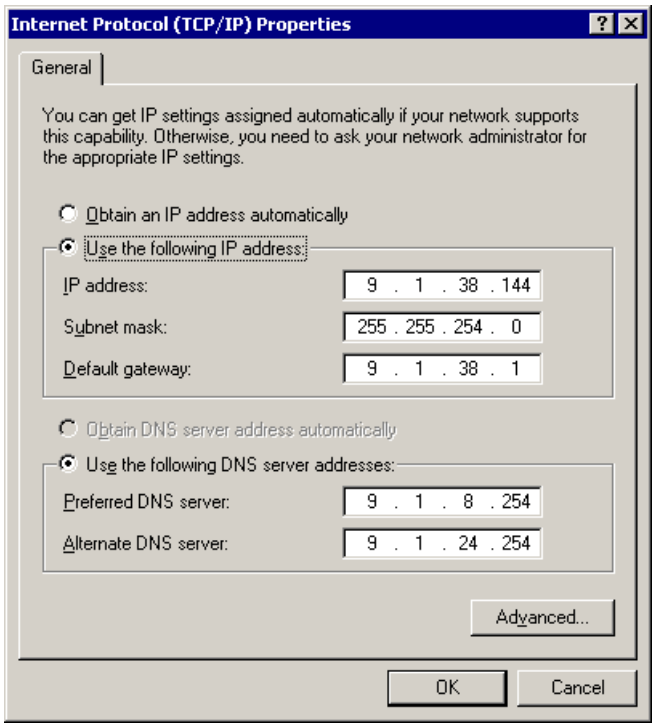

Figure 4-18 IP Properties screen

12. Our example shows an IP address of 9.1.38.144. Configure the IP address and settings for your network. You should get the information from your network administrator:
    - IP address
    - Subnet Mask
    - Default Gateway
    - Preferred DNS Server
    - DNS Suffix

The DNS Suffix and WINS, etc., can be found under the **Advanced** tab. When you are finished configuring your network adapters, click the **OK** button and exit back to the NAS Setup Navigator. Click the **Forward** button to configure your drives.

Note: The NAS 100 comes preconfigured, and you may be able to skip these steps if the factory default settings meet your business needs.

4.2.2 Storage configuration and management

The NAS 100 comes with storage space already configured. Four hard drives are connected to the on board ATA controller and the NAS 100 is using the Windows 2000 integrated software RAID function. Hard drive storage space is configured as follows (See Figure 4-19):

- A mirrored volume on HDD1 and HDD2 for the System volume (drive letter will be C:)
- A mirrored volume on HDD3 and HDD4, for the Recovery system volume (drive letter will be D:)
- A RAID-level 5 volume, comprised of the remaining storage on HDD1, HDD2, HDD3, and HDD4 (drive letter will be E:)

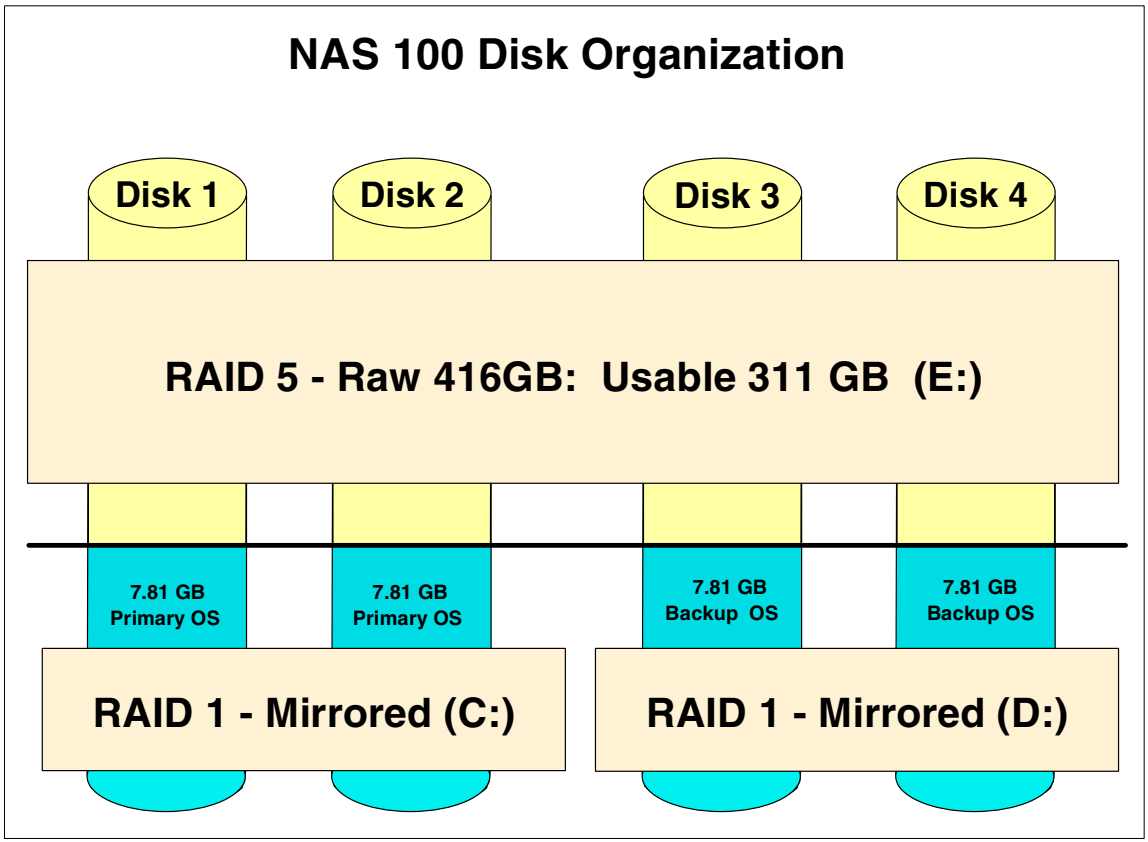

Figure 4-19 NAS 100 disk organization

Important: Although usable disk space is available on the System and Recovery volumes, it is not recommended to repartition them. NAS 100 System Recovery will not function if the System volume configuration is changed.

Data volume (drive letter E:) comes preconfigured and ready to use. If, for any reason, you want to reconfigure it, you can do so by following directions in the Configure Pooled Storage screen of the NAS Setup Navigator. By clicking the Disk Management link in the Creating Partitions panel (shown in Figure 4-20), the Disk Management snap-in will start and you have the possibility to delete, create new partitions, or verify existing partitions.

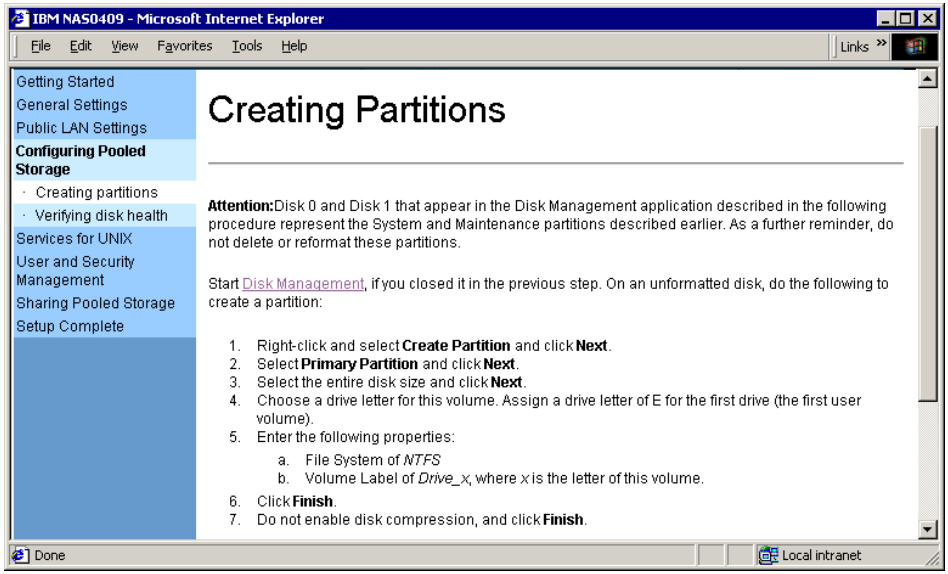

Figure 4-20 NAS Setup Navigator — Creating Partitions screen

By clicking **Forward** in the NAS Setup Navigator, you move to the Verifying disk health screen. By clicking the Disk Management link, the same plug-in is started and the health status is indicated for each volume, as shown in Figure 4-21.

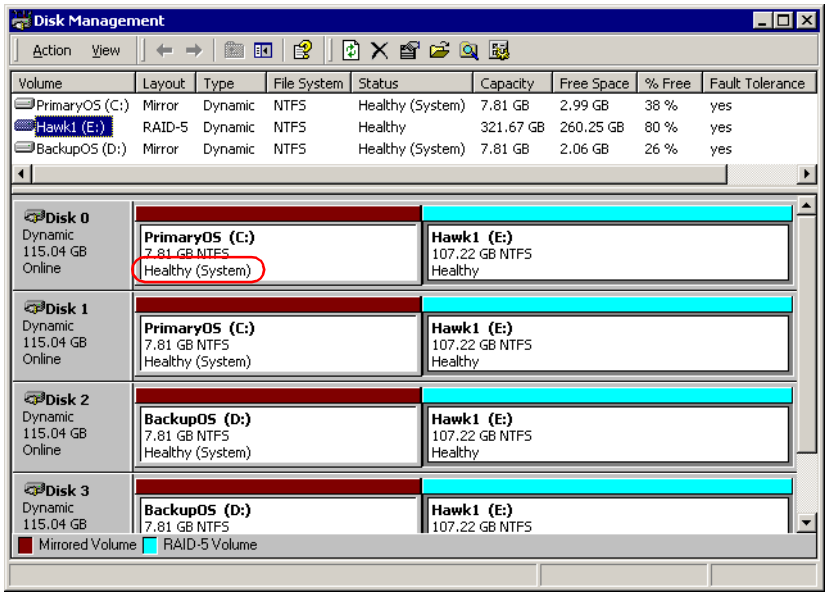

Figure 4-21 Disk Management plug-in

Disk subsystems management

Unlike the NAS 200 and NAS 300 appliances, the NAS 100 device doesn't use a hardware based RAID controller and there are no specific management applications for the disk subsystem. NAS 100 employs industry-standard management tools included in Windows Powered for NAS operating system, meaning the Microsoft Management Console. Specifically, the Disk Management plug-in can and should be used for all disk-related management tasks.

4.2.3 Microsoft Services for UNIX

As part of the initial preload, the Microsoft Services for UNIX V2.2 provide file access to UNIX and UNIX-based clients and servers using Network File System (NFS) protocol. In this section you can:

- Set up the Server for NFS.
- Set up the Gateway for NFS.
- Configure the Server for PCNFS.
- Configure User Name Mapping.

Each of these tasks can be started separately by choosing the corresponding link in the left pane, or they can be accomplished one after another by clicking the **Forward** button. Selecting the links in the center of this screen will open a comprehensive Microsoft Services for the UNIX online help library (Figure 4-22).

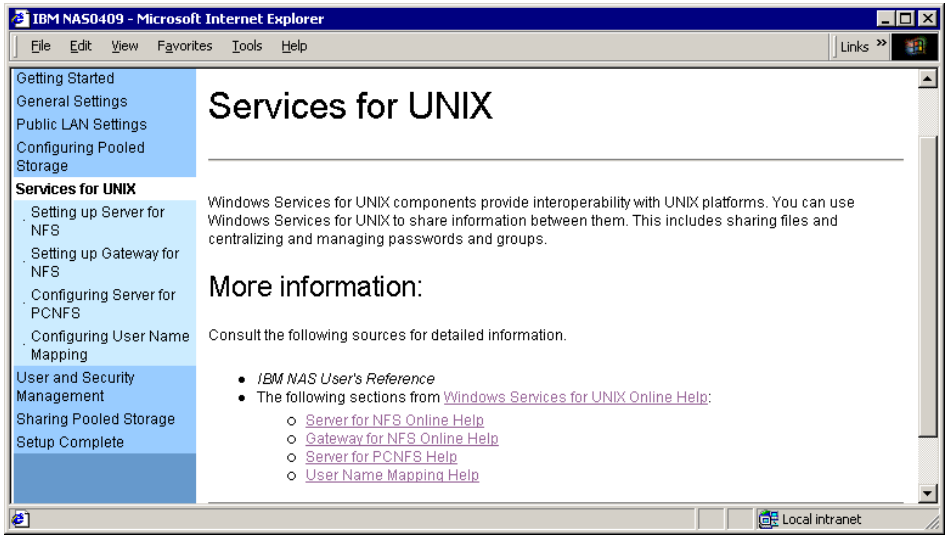

Figure 4-22 NAS Setup Navigator — Services for UNIX

UNIX Network Information System (NIS) integration

The UNIX Network Information System (NIS) services work just like using the yellow pages. While the NAS 100's feature set includes support for NIS, the security standard of NIS is not very high. Therefore we do not recommend the use of NIS with this product.

Within Microsoft Services for UNIX V2.2 there is a NIS migration wizard. This tool allows you to migrate a NIS. The tool takes your NIS source files and migrates them into Active Directory.

The Server for NIS feature allows a Windows Domain controller to be an NIS master server or an NIS subordinate (slave) by integrating NIS into Active Directory. When using the NIS server as a slave, the NIS master server must be a Windows 2000 Server.

For more information, you can check the following Web site:

`http://www.microsoft.com/WINDOWS2000/sfu/default.asp`

Password synchronization

Another tool that is included within the Microsoft Services for UNIX 2.2 is a password synchronization tool (2-way). It allows you to synchronize password changes between Windows NT or Windows 2000 and UNIX. Pre-compiled single sign-on daemons are available for:

- IBM AIX 4.3+
- Linux (Red Hat 5.2, 6.0, 7.x, 8.0)
- Sun Solaris 2.6+
- HP-UX 10.3+
- Compaq Tru64 UNIX

Note: Even if your UNIX version is not on the list — it may still work. Microsoft provides the source code for the password synchronization tool.

For a detailed overview and functional explanation of Services for UNIX, see Chapter 10, "Cross platform storage" on page 353.

4.2.4 User and security management

This section describes integrating the NAS 100 into a secure environment. The NAS 100 appliance is designed to plug right into your existing user and security management system.

Active Directory, NT 4 Domains, and Workgroups

The NAS 100 will integrate with all of the Microsoft Operating System versions that you may have in your current network environment. It will work with any existing user and security management for those systems, including:

- Windows Workgroup Computing
- Windows NT 4 Domains
- Windows 2000 Active Directory (mixed and native mode)

Fully describing user and security management for Windows is beyond the scope of this chapter, so we will just provide you with a quick overview. A more detailed look at Active Directory is provided within Chapter 13, “Microsoft Active Directory and IBM TotalStorage NAS” on page 467.

Also, you can refer to the literature regarding Microsoft Operating Systems. Some examples are listed in “Related publications” on page 637.

To manage local users on your NAS 100, from the NAS Setup Navigator left pane, select the **User and Security Management** section (Figure 4-23).

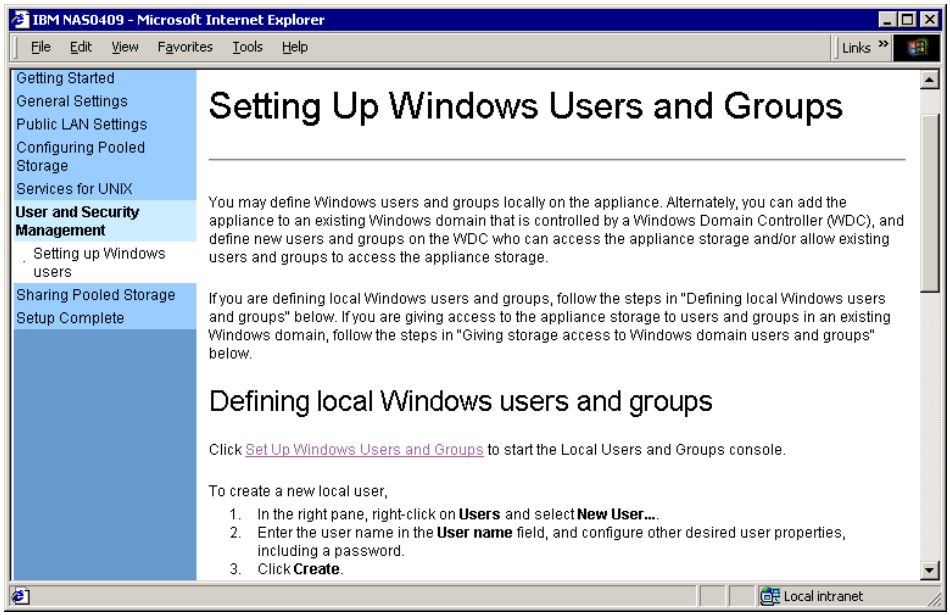

Figure 4-23 NAS Setup Navigator — Setting Up Windows Users and Groups

On the screen, Setting Up Windows Users and Groups, click the link, **Set Up Windows Users and Groups**, and the Local Users and Groups plug-in will start in a separate window. To create a new user, click **Action —> New User**, and enter the User name and password for the new user (Figure 4-24).

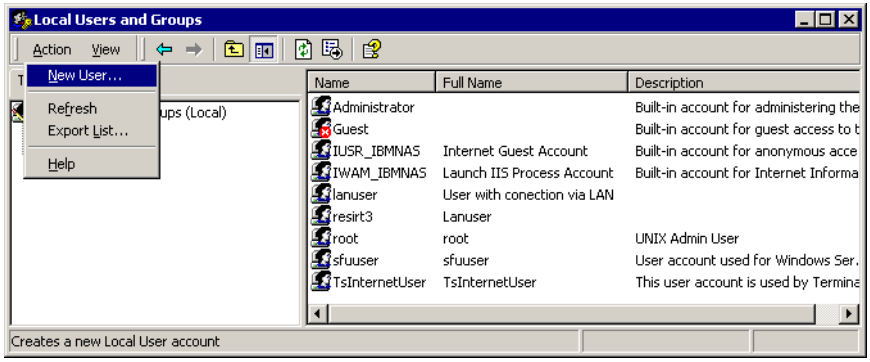

Figure 4-24 Local Users and Groups plug-in

The same tool can be used for creating a new local group on the NAS 100 appliance.

To give users from any domain access to the storage resources on the NAS 100 appliance, you have to move the device into the domain first. This can be done as part of the Network Identification procedure. It is shown in Figure 4-16 on page 72.

When you want to join an existing Windows NT 4.0 Domain, check the **Domain** box and type in the name of your domain. Click **OK** to confirm your choice and you will soon be a happy member of an existing Windows NT 4.0 Domain.

Tip: When joining an existing NT 4 Domain across subnets or via routed paths, define the PDC as the primary WINS server, even if the WINS service is not running on the PDC. This way, the joining client will find the PDC easily.

Joining an Active Directory tree is almost identical to joining an NT 4 Domain. Check **Domain** and type in the name of your Active Directory tree.

Important: When joining an Active Directory, it is essential that your TCP/IP configuration and DNS name resolution be working properly. Make sure both machines can ping each other using the IP address and the fully qualified domain name before joining the domain. For example, type:

```
ping lochness.itsolab2.almaden.ibm.com
```

4.2.5 Sharing pooled storage

As the last step of the NAS Setup Navigator, we have to share the NAS 100 storage space to the users in the network. IBM TotalStorage NAS 100 appliance supports multiple network protocols used in different client operating systems, including:

- CIFS (NT LM 0.12) for Windows clients
- NFS (V2.0, V3.0) for UNIX clients
- FTP
- HTTP (V1.1)
- Apple File Protocol
- Netware (Novell 4,5,6 via Windows Services for Netware V5.0)

Sharing for Windows clients using the CIFS protocol

To get information about sharing folders on the NAS 100 device, read the directions in the File Sharing for Windows clients screen of the NAS Setup Navigator. Then click the **Windows Explorer** link, as shown in Figure 4-25.

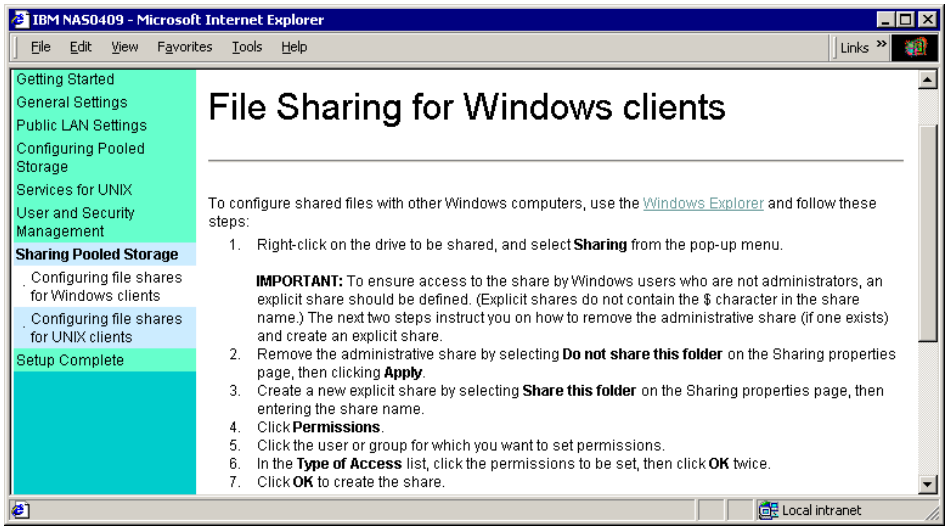

Figure 4-25 NAS Setup Navigator — File Sharing for Windows Clients

Windows Explorer opens. To share a certain folder on the NAS 100 appliance, right-click it and select **Sharing**. In the Folder properties window that opens, you can change the Share name, enter a description of the shared folder, and define the maximum number of connected users, as shown in Figure 4-26.

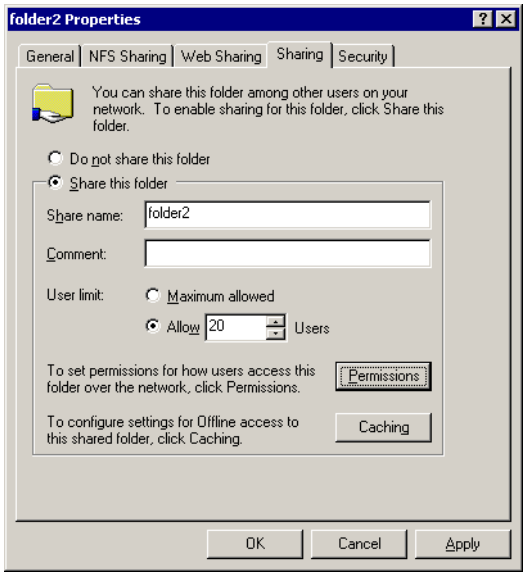

Figure 4-26 Shared Folder properties

To change Permissions for different users, click **Permissions** and change them according to the security policy in place.

Sharing for UNIX clients using the NFS protocol

To share folders for the UNIX clients (Figure 4-27), get the needed guidance on the *Sharing for UNIX clients* screen. Again, Windows Explorer will be used for defining shares.

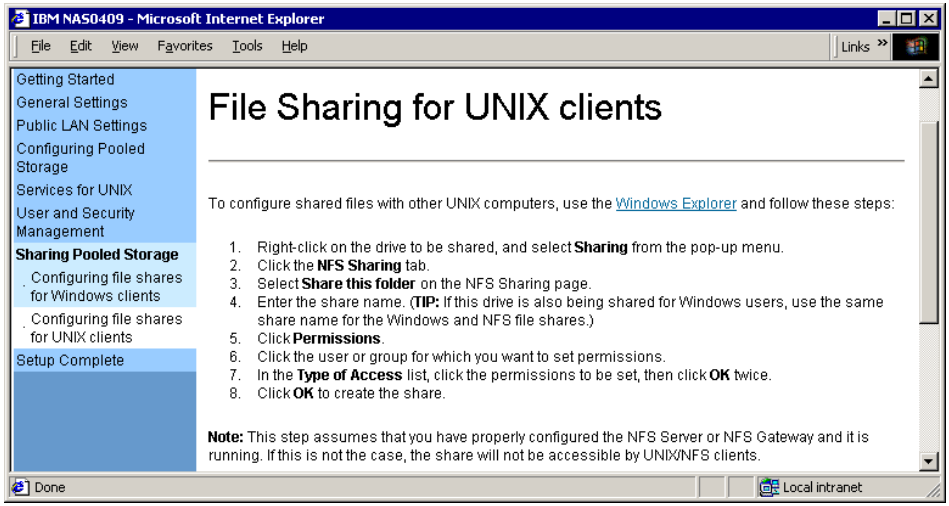

Figure 4-27 NAS Setup Navigator — File Sharing for UNIX clients

4.2.6 Completing setup

This concludes the setup of the NAS 100 appliance using the NAS Setup Navigator. You can return to any of the available sections at any time to view or change some of the settings mentioned (Figure 4-28).

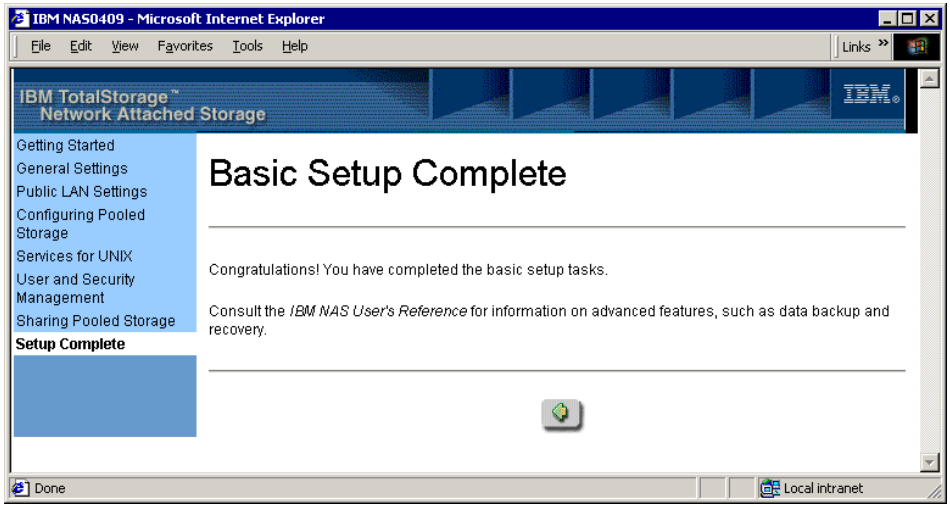

Figure 4-28 NAS Setup Navigator — Setup Complete

4.3 Ethernet adapter teaming

This section describes how to enable adapter teaming on the Ethernet adapters of the NAS 100.

4.3.1 Overview of adapter teaming

Teaming provides traffic load balancing and redundant adapter operation in the event that a network connection fails. NAS 100 has two Gigabit Ethernet adapters that can be grouped into teams. If traffic is not identified on any of the adapter team members' connections due to failure of the adapter, cable, switch port, or switch (where the teamed adapters are attached to separate switches), the load distribution is reevaluated and reassigned among the remaining team members. In the event that all the primary adapters are down, the hot standby adapter becomes active. Existing sessions are maintained, causing no user impact. NAS 100 supports three schemes of load balancing:

- Smart load balancing
- Link aggregation (802.3ad)
- Generic link aggregation (Trunking)

Smart load balancing

The implementation of load balancing based on IP flow. This feature supports balancing IP traffic across multiple adapters (that is, team members). In this mode, all adapters in the team have separate MAC addresses. It provides automatic fault detection and dynamic failover to another team member or to a hot standby member; this is done independently of layer 3 protocol (IP, IPX, NetBEUI). It works with existing layer 2 and 3 switches.

Link aggregation

This mode supports link aggregation through static configuration and conforms to the IEEE 802.3ad specification. Configuration software allows you to statically configure which adapters you want to participate in a given team. If the link partner is not correctly configured for 802.3ad link configuration, you will receive errors. With this mode, all adapters in the team are configured to receive packets for the same MAC address.

Generic link aggregation (Trunking)

This mode is very similar to 802.3ad in that all adapters in the team need to be configured to receive packets for the same MAC address. This mode supports a variety of environments where the link partners for the NICs are statically configured to support a proprietary trunking mechanism.

Basically, this mode is a *light* version of the 802.3ad link aggregation. This approach is much simpler because there is not a formalized link aggregation control protocol. As with the other modes, the creation of teams, and the allocation of physical adapters to various teams, is done statically with user configuration software.

Trunking supports load balancing and failover for both outbound and inbound traffic.

4.3.2 Load balancing for the configuration

Load Balance provides an easy way to configure the load balancing and redundant adapter function by grouping multiple adapters into teams.

The Load Balance panel allows you to configure advanced features. Teaming is a method of grouping multiple adapters into a virtual adapter (bundling multiple adapters to look like a single adapter). The benefit of this approach is load balancing.

To access the Ethernet Adapter Teaming tool, use the following sequence:

1. Open Control Panel via Remote Control or Terminal Services Session, **Start** -> **Settings** -> **Control Panel**.
2. Double-click the **Broadcom NetXtreme Gigabit** icon, to start the tool (Figure 4-29).

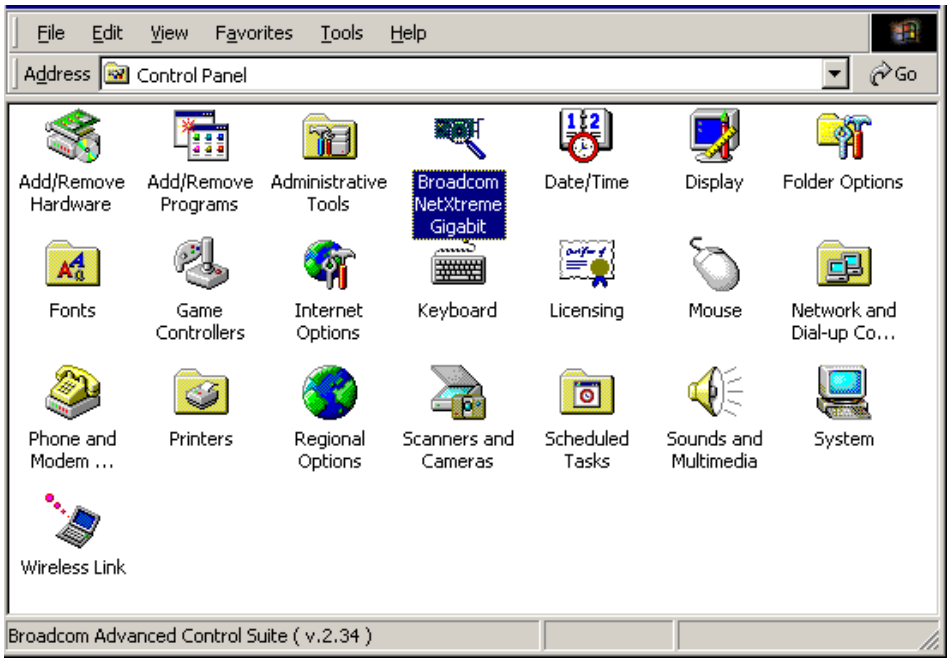

Figure 4-29 Control Panel on NAS 100 with Broadcom NetXtreme Gigabit icon

3. The initial configuration panel will open (Figure 4-30).

Important: You must have a working DHCP Server in your network infrastructure. This is necessary because the Load Balance tool will disconnect the TCP/IP connection while configuring the new team (Figure 4-34 on page 90)! It is not possible to configure the IP-address before restarting the network! Make sure that the new ethernet team will get an IP address and is registered by the DNS server or Windows domain controller. You can connect to the NAS100 via terminal services or remote control after network has been restarted. Otherwise, you may lose connectivity to your NAS 100 system.

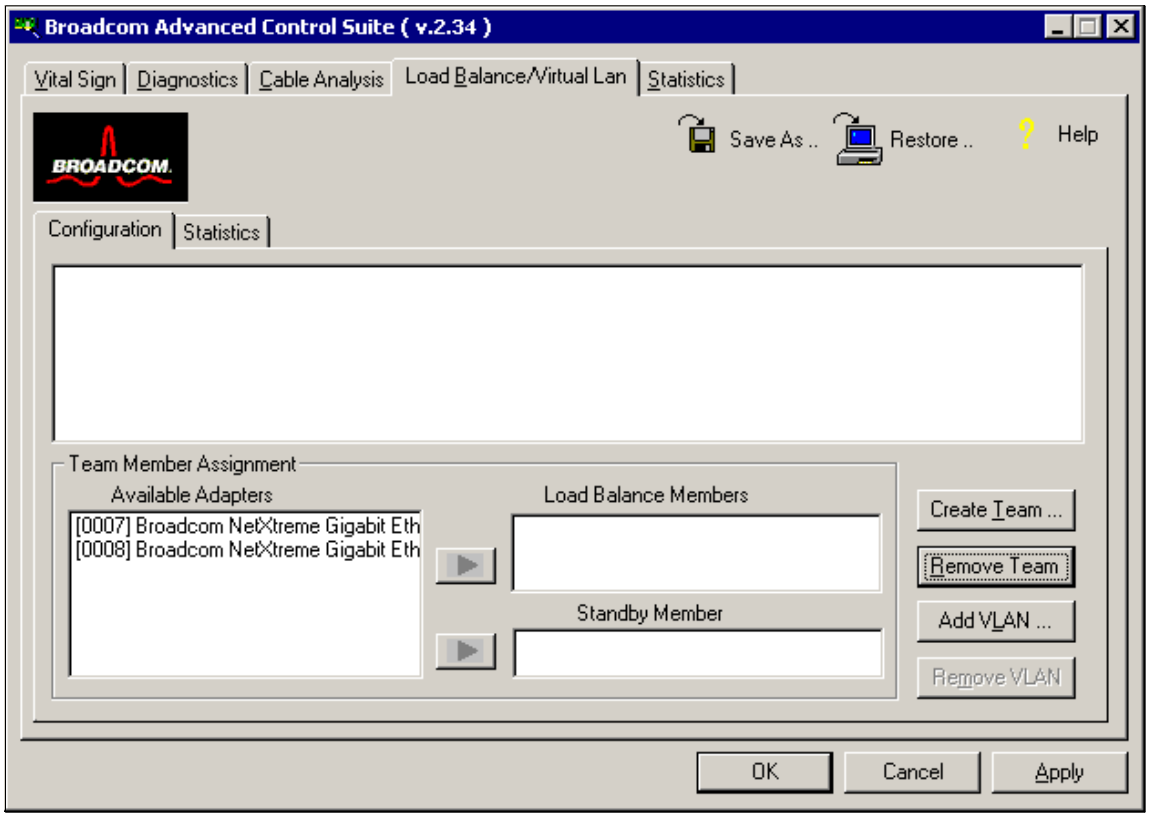

Figure 4-30 Initial configuration panel

4. From the Load Balance window, click **Create Team**. This displays the Add New Team window shown in Figure 4-31.

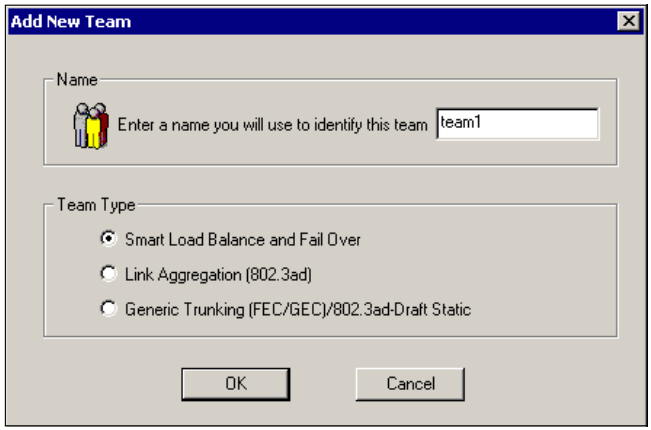

Figure 4-31 Add New Team screen

5. Type a team name, select the team type, and click **OK**.

   The minimum number of characters that can be used in a team name is 1. The maximum number of characters that can be used in a team name is 39. A team name cannot begin with spaces, nor can it contain the character “&”. If you attempt to use an identical team name, an error message window is displayed indicating that the name already exists.

6. In the Available Adapters list in the **Load Balance/Virtual Lan** panel (Figure 4-32), select the available adapter or adapters that you want to add to the team created in the previous step. Move the selected adapters to the Load Balance Members list box using the arrows.

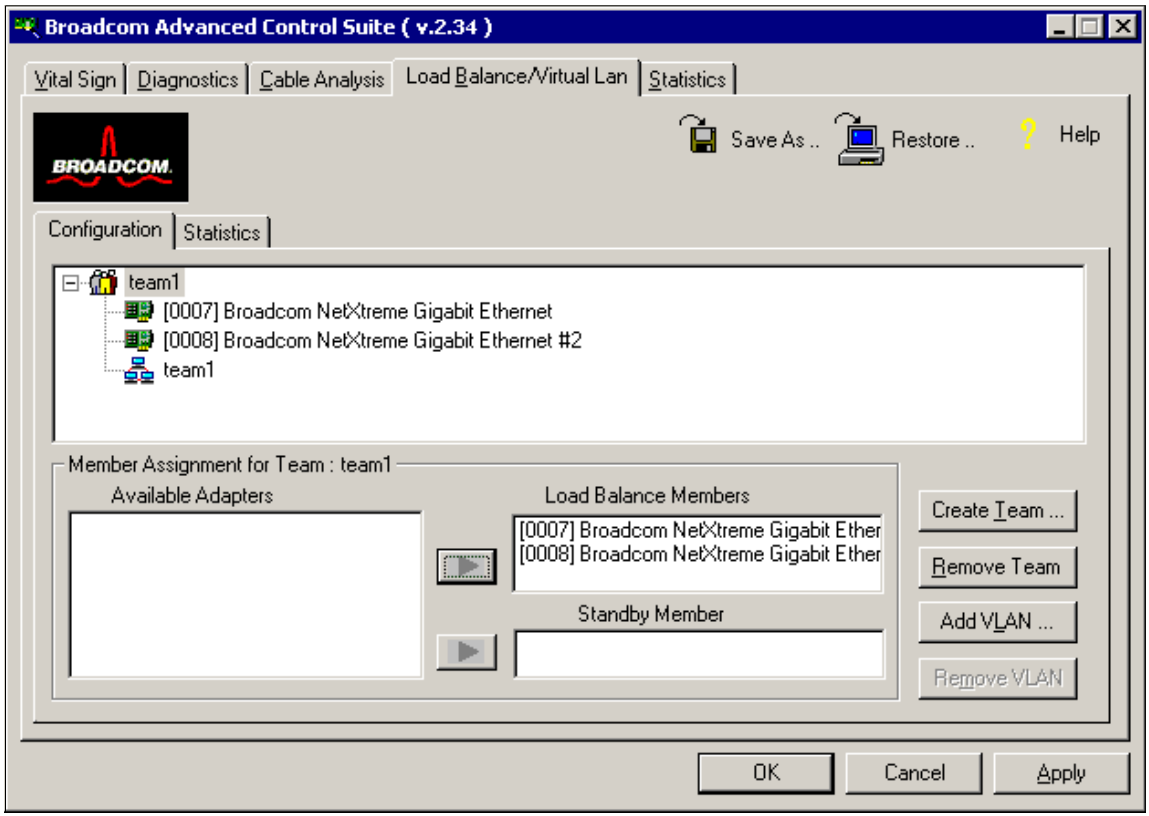

Figure 4-32 Available adapters added to team

7. When you have finished configuring failover teams, click **OK** or **Apply** to accept the changes.

Note: At least one adapter must be displayed in the Load Balance list box

8. A Microsoft Digital Signature message will appear (Figure 4-33). The BASP Virtual Adapter was not yet certified with a digital signature from Microsoft. Clic **Yes** to continue the installation process.

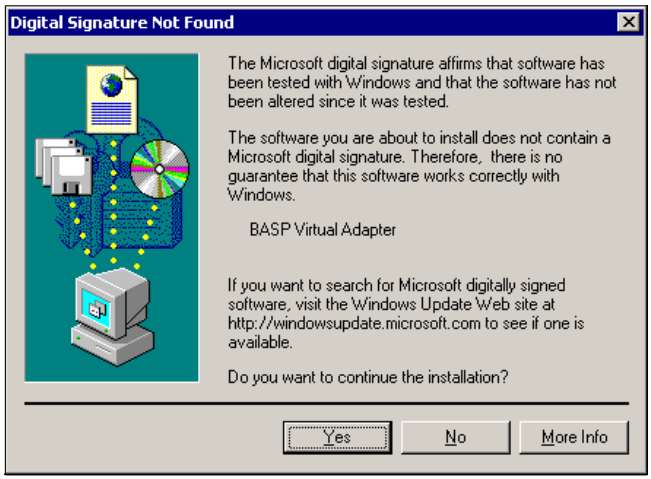

Figure 4-33 Microsoft Digital Signature message

As mentioned earlier, you must have a working DHCP Server in your network infrastructure. This is necessary because the Load Balance tool will disconnect the TCP/IP connection while configuring the new team (Figure 4-34).

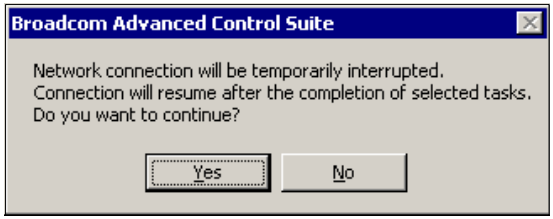

Figure 4-34 Network connection interruption

9. When team configuration has been correctly performed, one Virtual Team adapter driver will be created for each configured team.

   When you create a generic trunking team, you cannot select a Standby Member. Standby Members work only with Smart Load Balance and Failover Teams.

10. Configure the Team IP address if necessary. If other adapters in your system use TCP/IP bindings, the TCP/IP Properties window will open.

11. To access the Internet Protocol Properties window in Windows 2000, right-click the **My Network Places** icon and select **Properties** to view the panel shown in Figure 4-35.

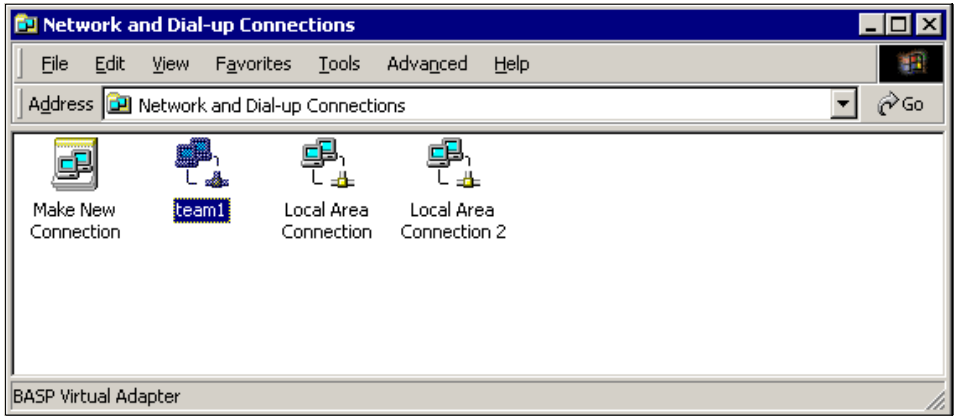

Figure 4-35 Network and dial-up connections

12. When the Network and Dial-up Connections window opens, right-click any network adapter. This displays the Internet Protocol (TCP/IP) Properties window shown in Figure 4-36.

    Use this window to set an adapter's IP address. Configure the IP address and any other necessary TCP/IP configuration parameters for the team and click **OK** when finished.

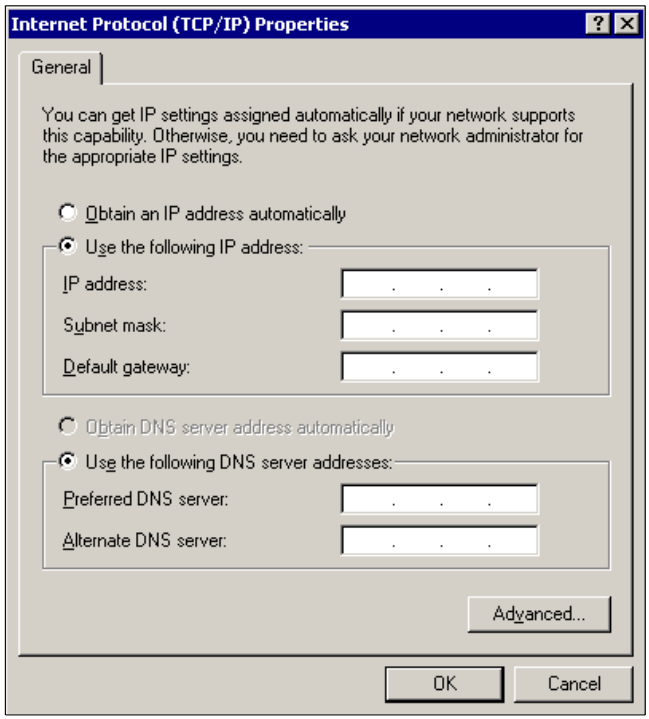

Figure 4-36 TCP/IP configuration properties

5

Implementing the IBM TotalStorage NAS 200

In this chapter we explain how to configure the IBM TotalStorage NAS 200 and NAS 300 from the very beginning.

We cover these topics:

- Using the Navigator to set up the NAS 200
- Managing disks using ServeRAID Manager

5.1 Using the Navigator to set up the NAS 200

In this section we demonstrate how to use the NAS Setup Navigator to set up the IBM TotalStorage NAS 200 from the factory. As mentioned before, the wizard will guide you through each step. Many of the NAS Setup Navigator screens contain links that will open additional configuration screens to make the process easier. To set up the NAS 200 out-of-the-box, follow these steps.

5.1.1 Basic configuration

These are the steps for basic configuration:

1. Connect to the NAS 200 appliance and logon with and administrative account (default is **Administrator** and **password**).
2. Open the NAS Setup Navigator. The icon is located on the desktop of the IBM TotalStorage NAS appliance. The Getting Started screen appears first and explains the navigator. To advance, click the **Forward** button, which brings you to the Information and Setup Options screen, as shown in Figure 5-1.

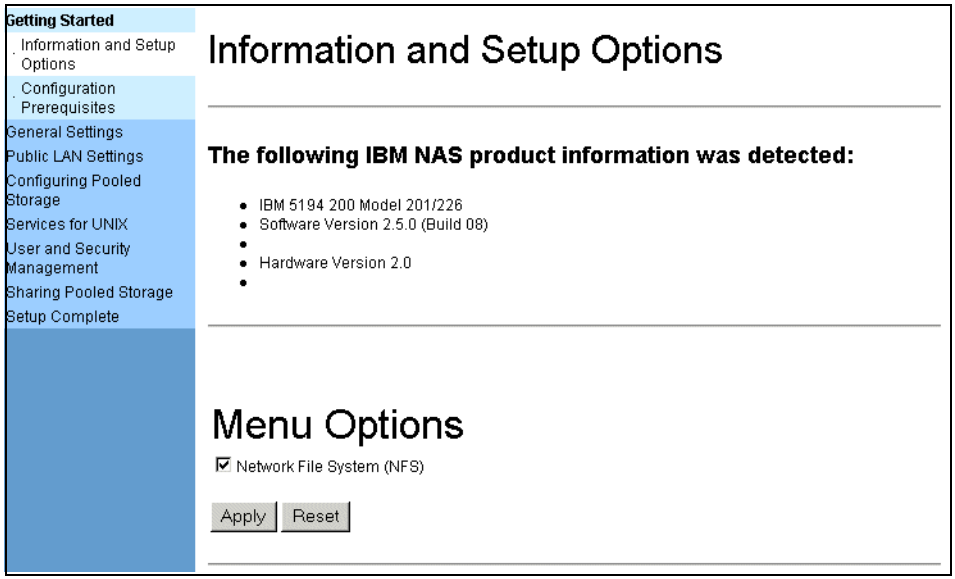

Figure 5-1 NAS Setup Navigator — Information and Setup Options screen

Note: Checking the NFS box will result in bringing up screens later on in the wizard process for installing support for NFS.

3. Select the Network File System box if support for NFS is required; otherwise, leave it unchecked and click **Apply**.

   Clicking the **Apply** button will refresh the screen and bring you back to the top of the Information and Setup Options screen.

4. Click the **Forward** button. This will bring up the Configuration Prerequisites page. This is an information only screen and is here to help you gather and record information concerning the NAS appliance. When you are finished, you can click the **Forward** button.

Note: If the navigator screen has a link (which will be in a different color and underlined) and directs you to click it, a file download screen will come up as shown in Figure 5-2 and ask you to either open the file from the current location or save to disk. Always select the radio button, **Open this file from its current location**, and click **OK**. This will bring up a new Windows screen.

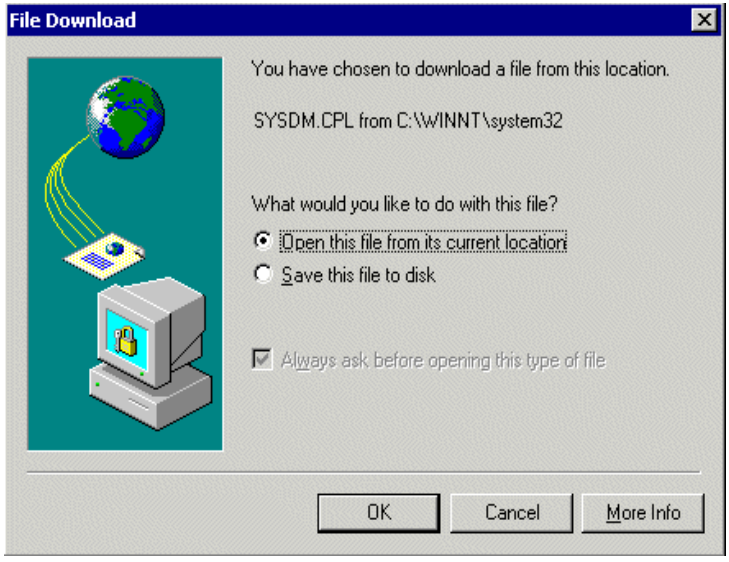

Figure 5-2 NAS Setup Navigator — additional screen

A typical NAS Setup Navigator screen is shown in Figure 5-3.

Note: As you can see in Figure 5-3, the link **Set Administrator User Name and Password** is a different color. Clicking it will bring up the file download screen.

After you have told the file download screen to **open this file from its current location**, another screen will open (in this case the Local Users and Groups Screen) which will allow you to alter the configuration specified by the link. After making the necessary changes, close that screen (not the NAS Setup Navigator) and continue on with the NAS Setup Navigator.

Figure 5-3 NAS Setup Navigator — Administrator Password screen

5. After the Configuration Prerequisites screen, the next set of NAS Setup Navigator screens will guide you through setting up some General Settings for your NAS appliance. Follow the steps and screens and configure the NAS appliance, and when you are done with each screen, click the **Forward** button to advance.
6. Configure your IBM TotalStorage NAS appliance with the appropriate information for System Language, Administrator Password, and Date and Time.
7. Configure Network Identification on your NAS 200 by clicking the Network Identification link. This will bring up the system properties window. Click the **Network Identification** tab and then **Properties**.

Now you have the option to join the NAS 200 to your environment, as shown in Figure 5-4.

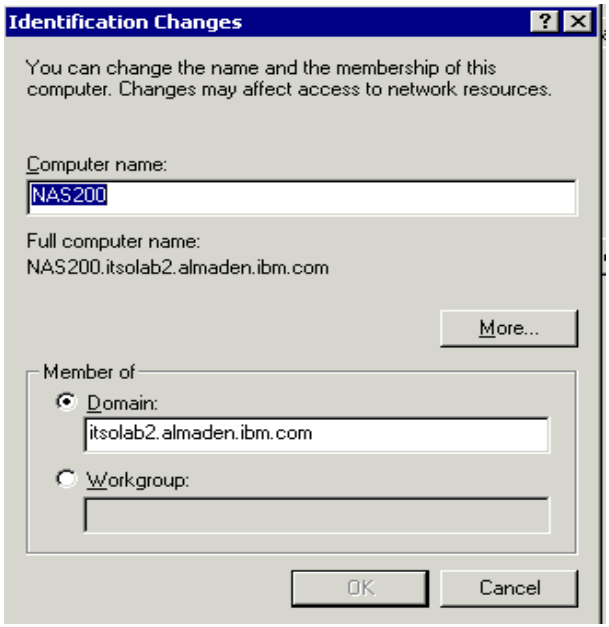

Figure 5-4 Network identification — computer name

- If you want the NAS 200 to be a storage device for your workgroup, check **Workgroup**, and type in the name of your workgroup. Be aware that all security management within such a workgroup is local. This means you have to administer all user accounts on the NAS 200. Also be aware that the default behavior when sharing a network drive with Windows is to grant all users full access to the data.
- If you want the NAS 200 to be part of a domain, then check **Domain** and enter the name of the domain you would like to join. A popup window will ask for the username and password of a domain administrator or equivalent. Either way, when you are finished, click the **OK** button. A reboot is required.

8. You will now continue with the NAS Setup Navigator by configuring Public LAN settings. Configure your IBM TotalStorage NAS appliance with the appropriate information for Public LAN Settings by clicking the Network and Dial-up Connections link. This will bring up the Network and Dial-up Connections window as shown in Figure 5-5.
9. Select the connection you want to work with. Right-click it, and choose **Properties**.

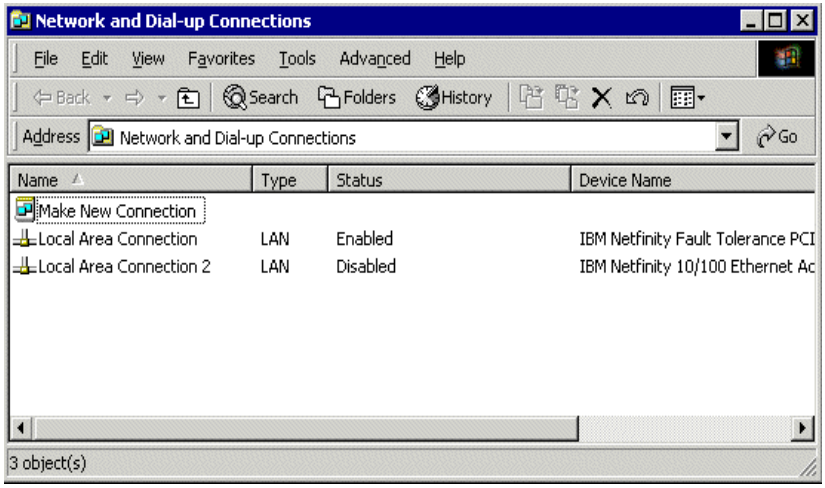

Figure 5-5 Network configuration window

10. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 5-6).

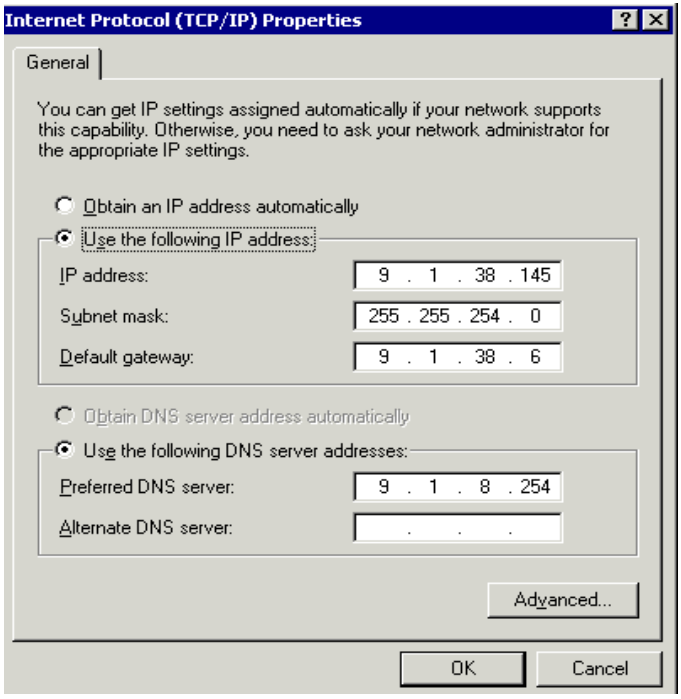

Figure 5-6 IP properties screen

11. Our example shows an IP address of 9.1.38.145. Configure the IP address and settings for your network. You should get the information from your network administrator:
    - IP address
    - Subnet Mask
    - Default Gateway
    - Preferred DNS Server
    - DNS Suffix

The DNS Suffix and WINS, etc., can be found under the **Advanced** tab. When you are finished configuring your network adapters, click the **OK** button and exit back to the NAS Setup NAvigator. Click the **Forward** button to configure your drives.

Note: The NAS 200 comes preconfigured, and you may be able to skip these steps if the factory default settings meet your business needs.

5.1.2 Configuring the storage

1. You will now define RAID Arrays and Logical Drives. You will also write disk signatures, create partitions, and verify disk health. There are separate NAS Setup Navigator screens for each one and can guide you through if needed. We will show all the steps for configuring Pooled Storage together. Click the **Forward** button to advance to the next screen and then click the ServeRAID Manager link to bring up the ServeRAID Manager window.
2. After clicking the ServeRAID Manager link, the manager window opens as shown in Figure 5-7, and lists the array controller properties.

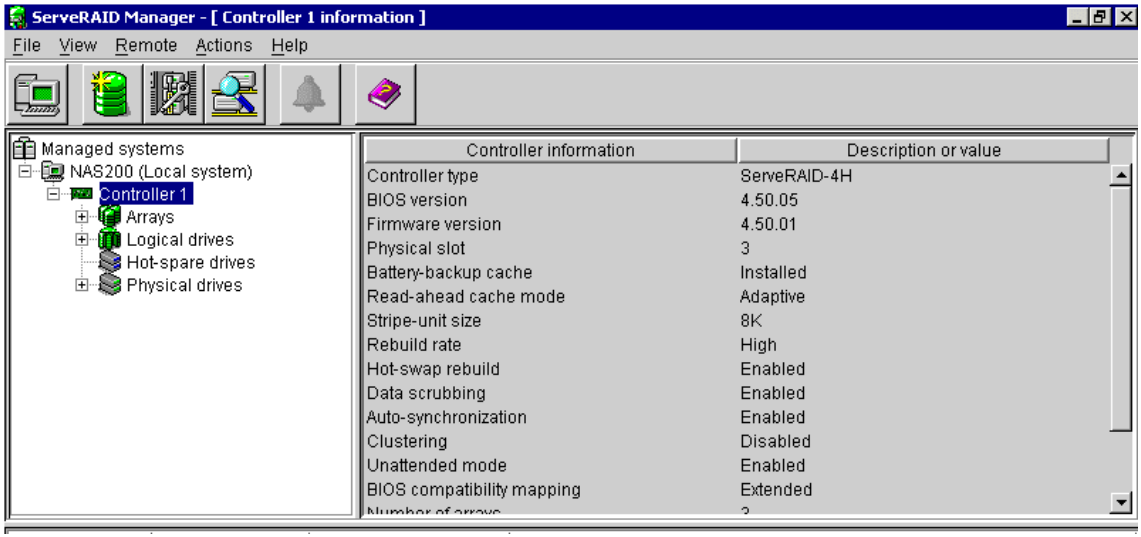

Figure 5-7 Array controller properties

5.1.3 Creating arrays and logical drives

Next we need to create drive arrays and logical drives on the system. The following steps show how we did this:

1. In the main tree, right-click the ServeRAID controller that you want to configure.
2. Click **Create Arrays**.
3. Click the **Custom configuration** button.
4. Click **Next**, and the Create arrays window opens.
5. Right-click the drive or SCSI channel icons in the Main Tree to select the drives that you want to add to your arrays, delete from your arrays, or define as hot-spare drives; then select a choice from the pop-up list. If you want to create a spanned array, click the **Span Arrays** box.
6. After you have selected the ready drives for your arrays and define your hot-spare drives, click **Next**. If you are not creating spanned arrays, you can select the RAID level here.

Figure 5-8 shows our sample array configuration.

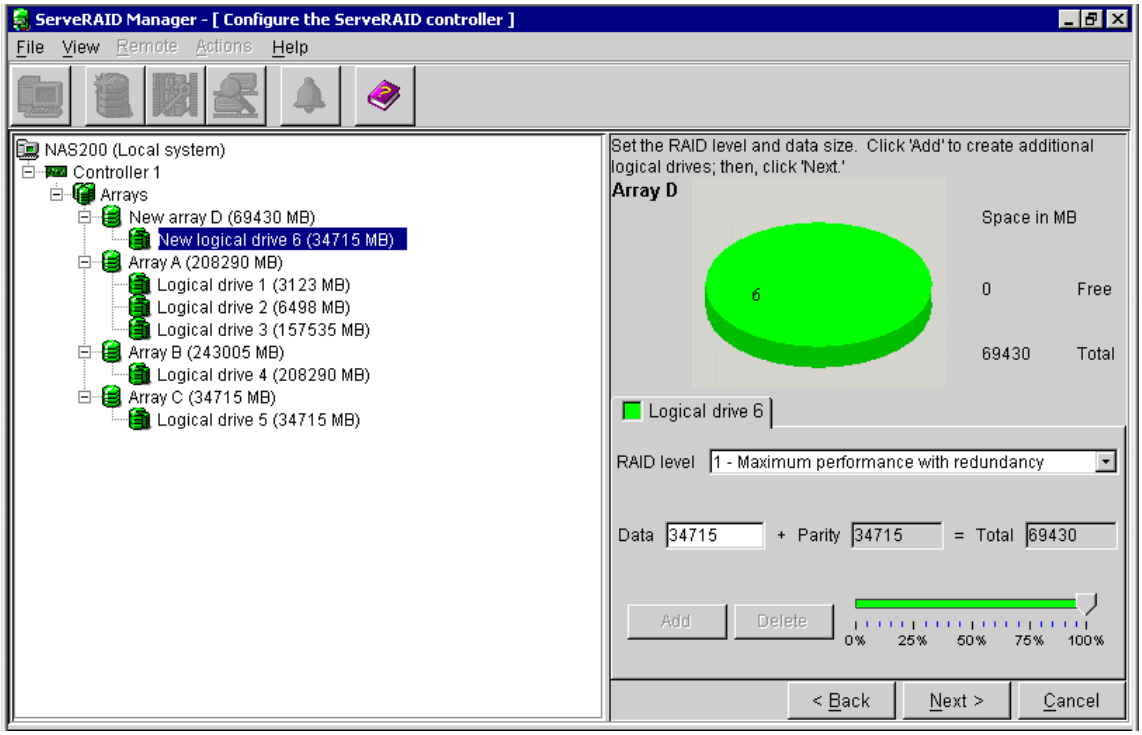

Figure 5-8 Creating arrays

7. To finish the procedure, click **Apply**.
8. Answer **Yes** to the question, “Do you want to apply the new configuration?”

9. Right-click the new array to synchronize it, as shown in Figure 5-9.
10. The synchronization time depends on the RAID level and number of drives.

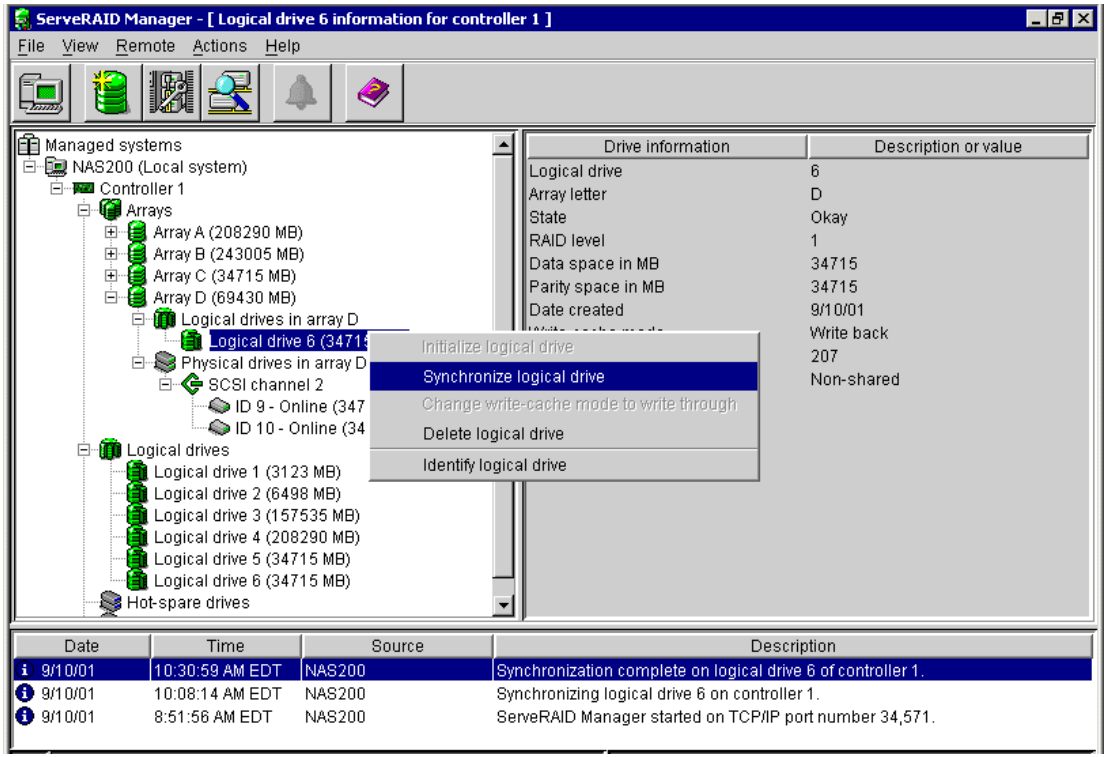

Figure 5-9 Synchronizing arrays

11. When the array is synchronized, close the window to get back to the NAS Setup Navigator. Click the **Forward** button to continue.
12. Click the **Disk Management link** to open the disk management window. Choose a new drive, then right-click to write the signature. The new disk is available now to the operating system.
13. Choose the new drive and right-click **Create partition**.
14. .Click **Next**.

15. You can choose to create either a primary partition or an extended partition, as shown in Figure 5-10.

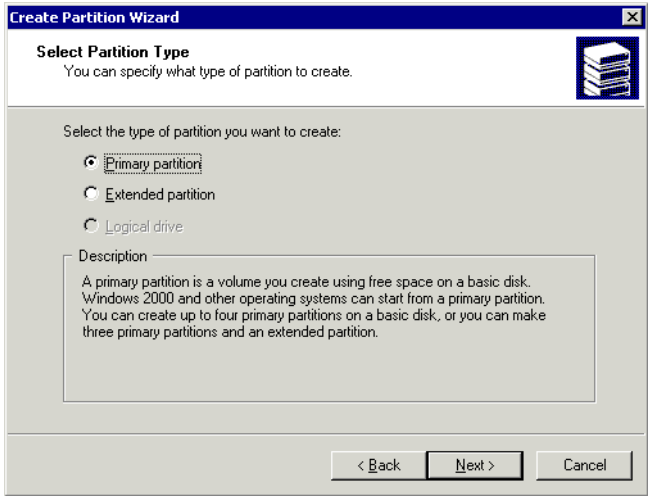

Figure 5-10 Choosing a primary partition

16. You will define the size of the partition here, as shown in Figure 5-11.

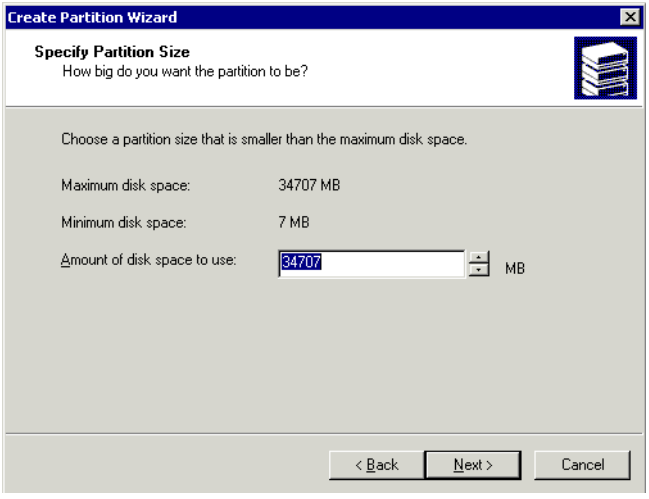

Figure 5-11 Changing the disk space

17. Click **Next**.

18. Choose the drive letter, as shown in Figure 5-12.

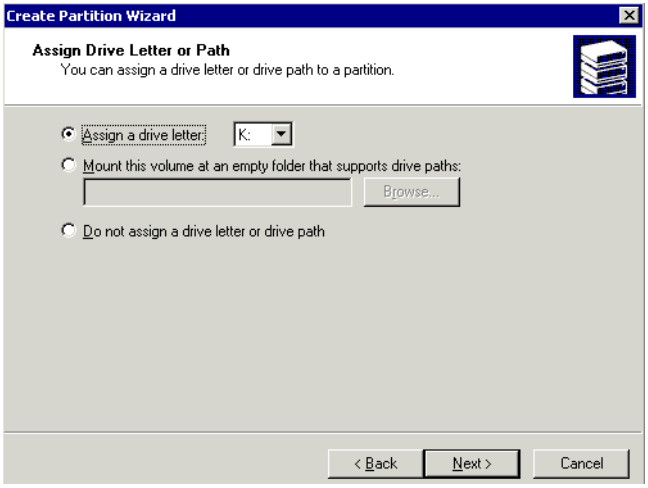

Figure 5-12 Assigning the drive letter

19. After that, you can choose between the NTFS or FAT32 file system. We recommend that you use NTFS. This is shown in Figure 5-13.

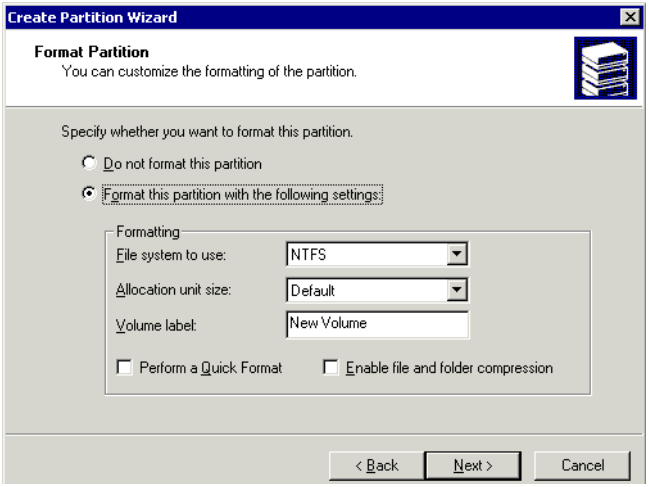

Figure 5-13 Choosing a partition type

20. Click **Perform a Quick Format** and then **Next**, and then click **Finish**. Your storage is now ready to be shared to users.

21. If you need to configure Services for UNIX, share Pooled Storage, etc., refer to Chapter 10, “Cross platform storage” on page 353.

5.2 Managing disks using ServeRAID Manager

You can use the ServeRAID Manager program to configure your ServeRAID controllers, view the ServeRAID configuration and associated devices, create arrays and logical drives, delete an array, dynamically increase the logical-drive size, change RAID levels, and do many more functions regarding your disk management.

Following are the steps for using the ServeRAID Manager:

1. Start a Windows Terminal Services connection and connect to your NAS system.
2. Log on to a NAS 200 with an administrative account.

Note: NAS 300 with the OS mirroring feature also has a ServeRAID and two disks configured in RAID 1 for the operating system that can be managed with the ServeRAID Manager tool.

3. Click **IBM NAS ADMIN.MMC** from the desktop (Figure 5-14).
4. Click the **+** sign beside NAS Management.
5. Click the **+** sign beside Storage.
6. Click the **+** sign beside ServeRAID Manager.
7. Click **Server Raid Manager**.

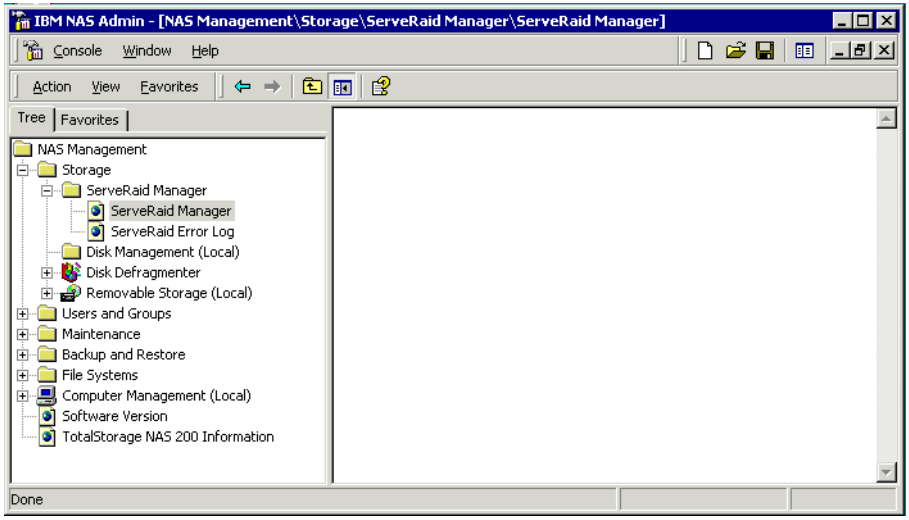

Figure 5-14 Using a IBM NAS ADMIN

After following these steps, we are able to run the ServeRAID Manager, as shown in Figure 5-15, and list the array controller properties.

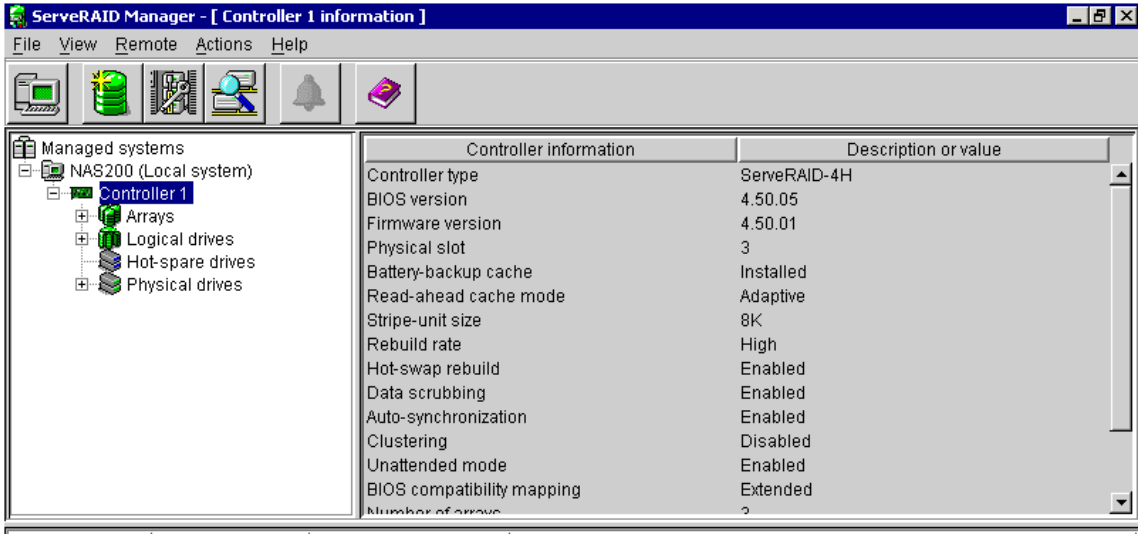

Figure 5-15 Array controller properties

5.2.1 Creating arrays

Next we need to create drive arrays and logical drives on the system. The following steps show how we did this:

1. In the main tree, right-click the ServeRAID controller that you want to configure.
2. Click **Create Arrays**.
3. Click the **Custom configuration** button.
4. Click **Next**, and the *Create arrays* window opens.
5. Right-click the drive or SCSI channel icons in the Main Tree to select the drives that you want to add to your arrays, delete from your arrays, or define as hot-spare drives; then select a choice from the pop-up list. If you want to create a spanned array, click the **Span Arrays** box.
6. After you have selected the ready drives for your arrays and define your hot-spare drives, click **Next**. If you are not creating spanned arrays, you can select the RAID level here.

   Figure 5-16 shows our sample array configuration.

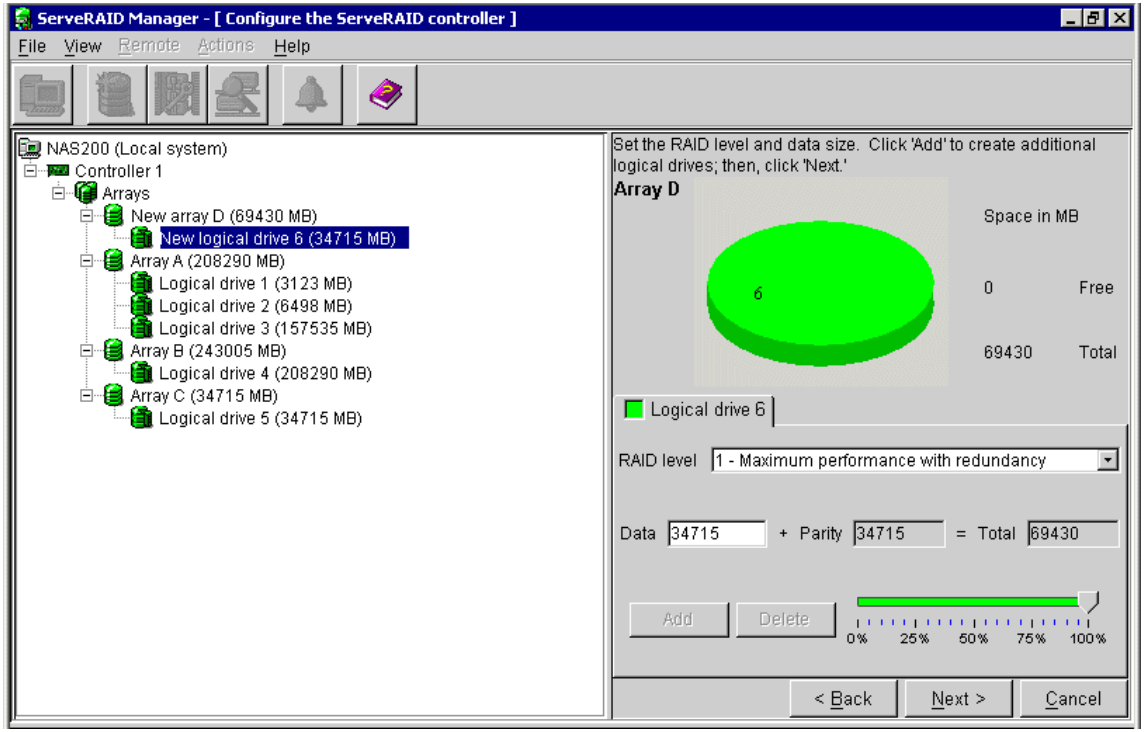

Figure 5-16 Creating arrays

7. To finish the procedure, click **Apply**.
8. Answer **Yes** to the question, “Do you want to apply the new configuration?”
9. Right-click the new array to synchronize it, as shown in Figure 5-17.
10. The synchronization time depends on the RAID level and number of drives.

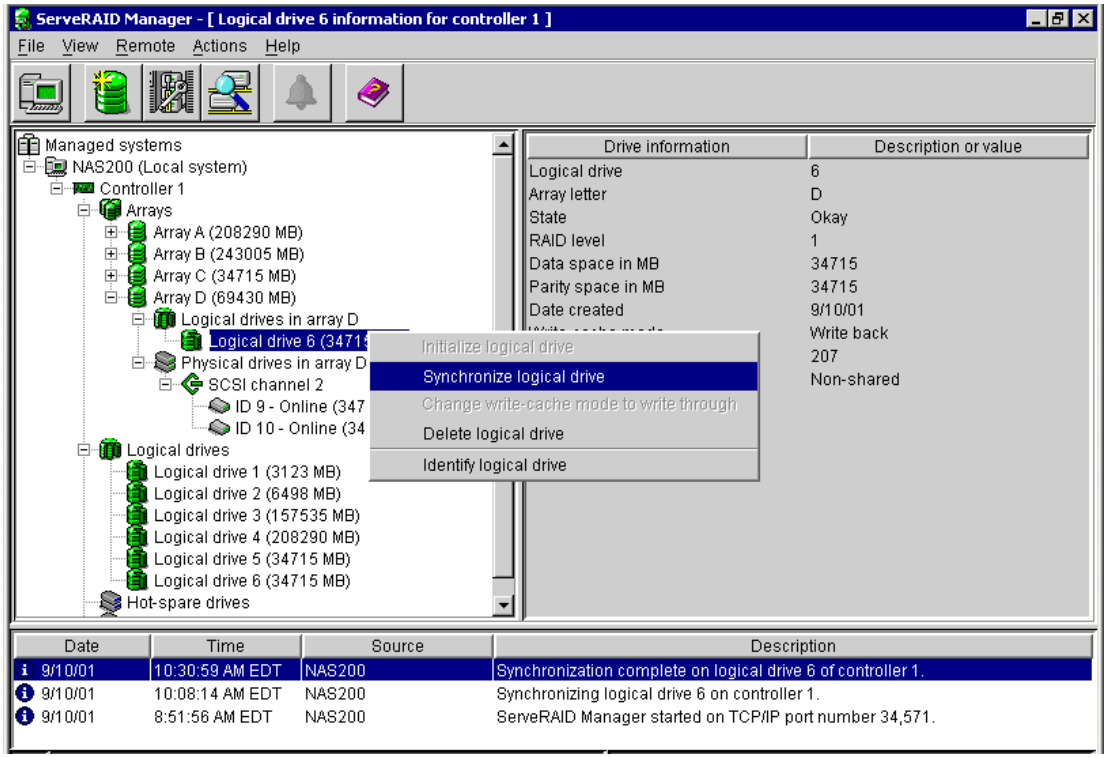

Figure 5-17 Synchronizing arrays

11. When the array is synchronized, open **IBM NAS ADMIN.MMC** from the desktop.
12. Click **Disk Management (Local)** and choose a new drive, then right-click to write the signature. The new disk is available now.

5.2.2 Creating and logical drives

After finishing creating the array, we are now ready to create logical drives.

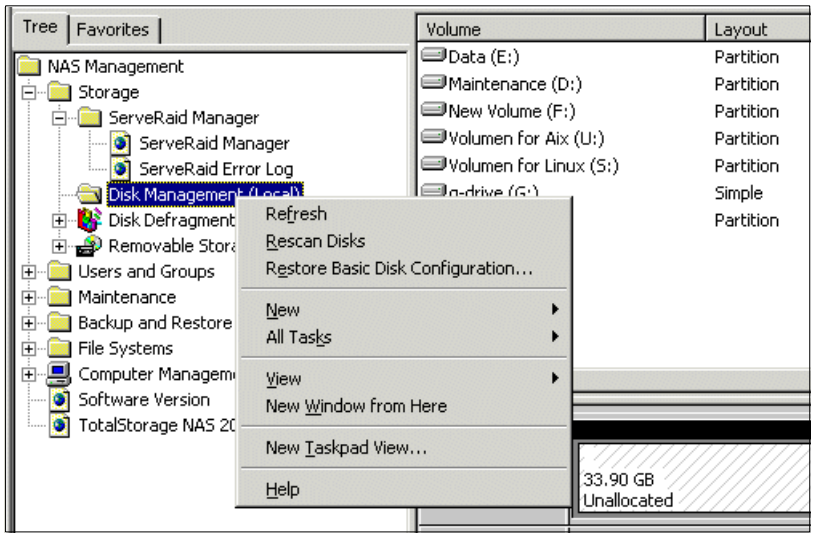

Figure 5-18 Creating a new volume

Now the array is ready to be used by the operating system.

1. Open **IBM NAS ADMIN.MMC** from the desktop.
2. Click **Disk Management (Local)** and choose the new drive (see Figure 5-18).
3. Right-click **Create partition**.
4. Click **Next**.

5. You can choose to create either a primary partition or an extended partition, as shown in Figure 5-19.

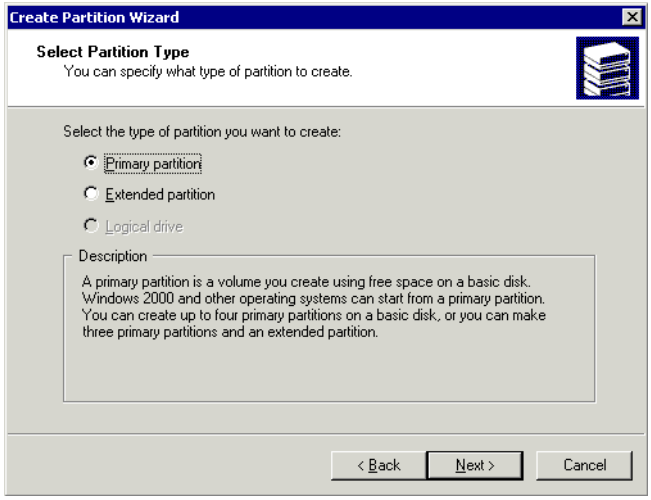

Figure 5-19 Choosing a primary partition

6. You will define the size of the partition here, as shown in Figure 5-20.

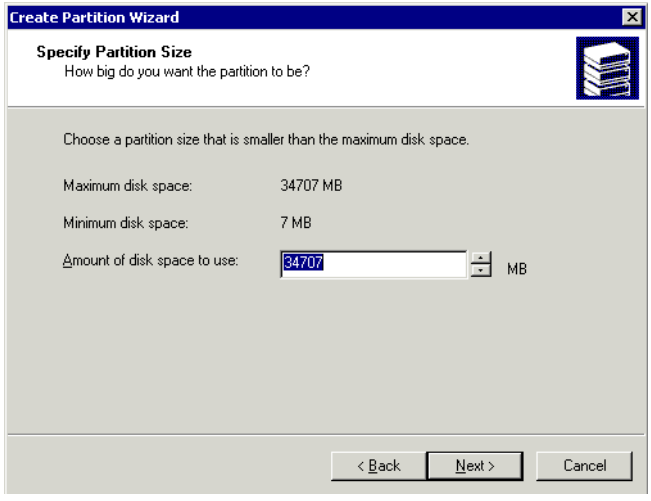

Figure 5-20 Changing the disk space

7. Click **Next**.
8. Choose the drive letter, as shown in Figure 5-21.

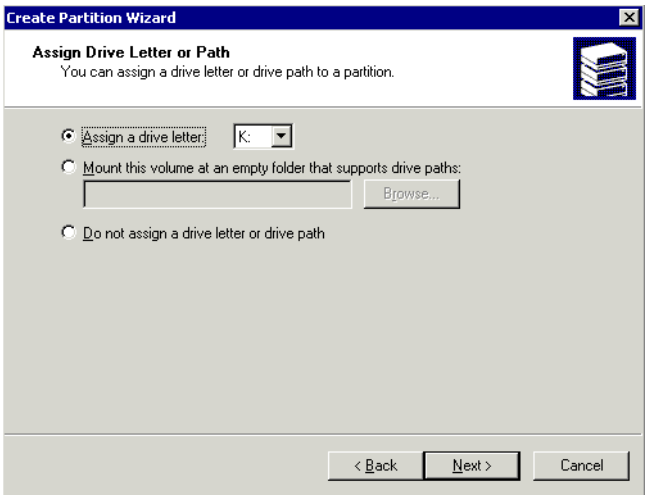

Figure 5-21 Assigning the drive letter

9. After that, you can choose between the NTFS or FAT32 file system. We recommend that you use NTFS. This is shown in Figure 5-22.

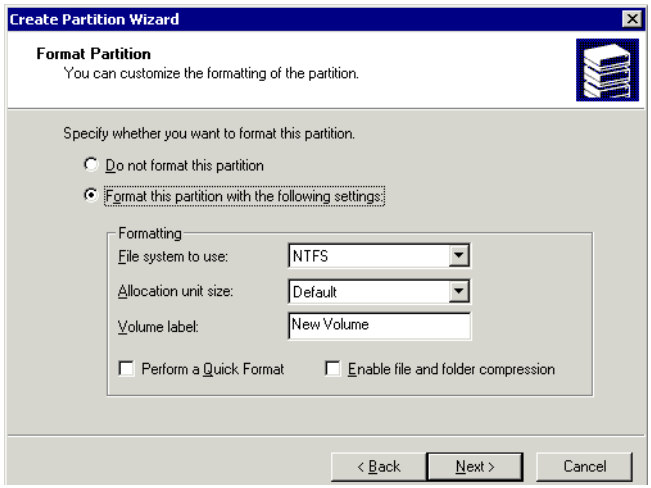

Figure 5-22 Choosing a partition type

10. Click **Perform a Quick Format** and then **Finish**.

Now you just need to wait for a few minutes while that new drive is formatting. If you performed this operation using Terminal Services, you must log off and start a new session with Terminal Services to recognize the new drive.

6

Implementing the IBM TotalStorage NAS 300

In this chapter we explain how to configure the IBM TotalStorage NAS 300 from the very beginning.

We cover these topics:

- Using the Navigator to set up the NAS 300
- Defining arrays and logical drives on the NAS 300
- Setting up Microsoft Cluster Server (MSCS)

6.1 Using the Navigator to set up the NAS 300

In this section we demonstrate how to use the NAS Setup Navigator to create a Microsoft Cluster and install Active Directory on the IBM TotalStorage NAS300.

6.1.1 Basic configuration

1. Connect to one of the NAS 300 systems and logon using an administrative account (default is **Administrator** and **password**).
2. Open the NAS Setup Navigator. The icon is located on the desktop of the IBM TotalStorage NAS system. The Getting Started screen appears first and explains the Navigator. To advance, click the **Forward** button, which brings you to the Information and Setup Options screen, as shown in Figure 6-1.

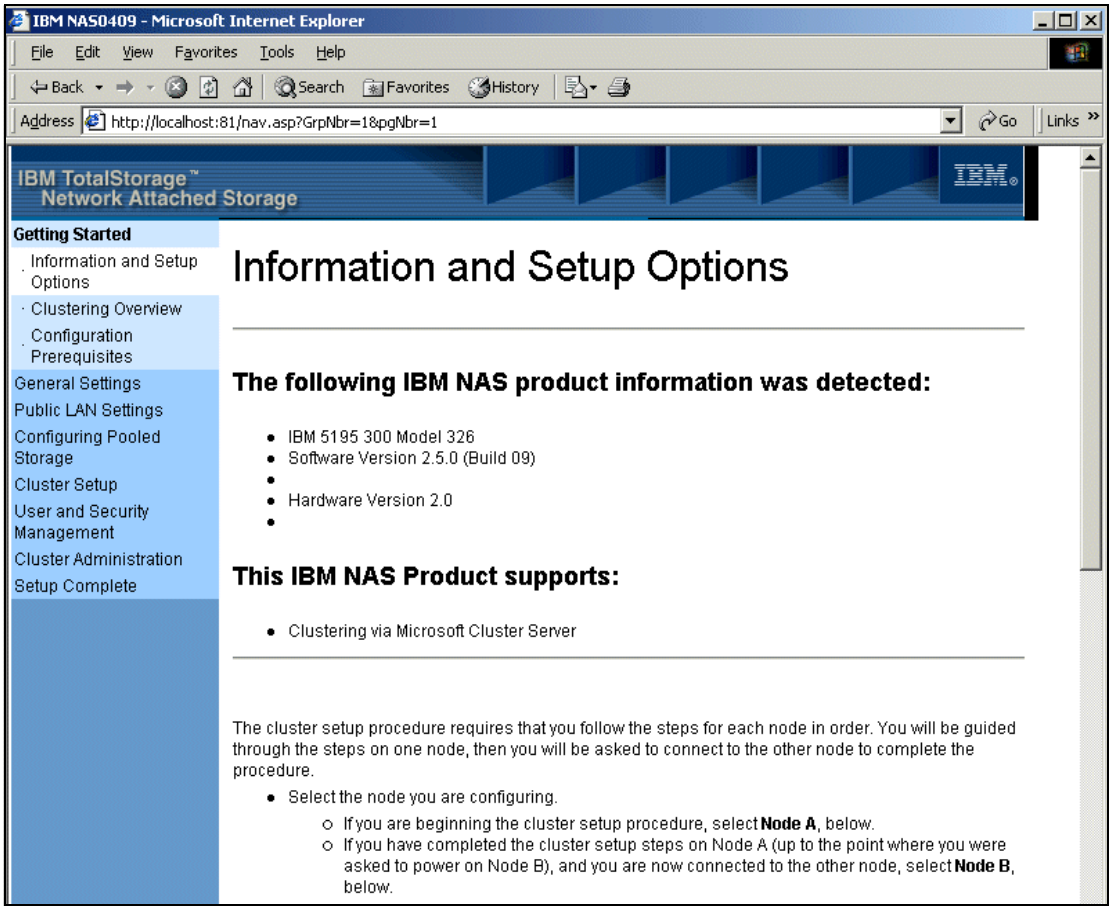

Figure 6-1 NAS Setup Navigator — Information and Setup Options screen

Note: New with NAS V2.5 is the addition of the domainlet feature on the IBM TotalStorage NAS 300. This feature allows an administrator to create a domain on the clustered system, if no NT4 Domain or Windows 2000 Active Directory (AD) currently exists in their infrastructure. The Active Directory Wizard will guide you through the process of installing Active Directory and creating your own domain.

We will call the first system node A and the second system node B.

3. Select the appropriate node to work with (node A) and decide if the domainlet feature is required (if an NT4 domain or AD currently exist and you can join it, then you do not need to create an Active Directory). After you have selected the appropriate radio buttons, click **Apply**. This refreshes the screen and brings you back to the top of the Information and Setup Options screen.
4. Click **Forward**. This brings up the Clustering Overview page (Figure 6-1).

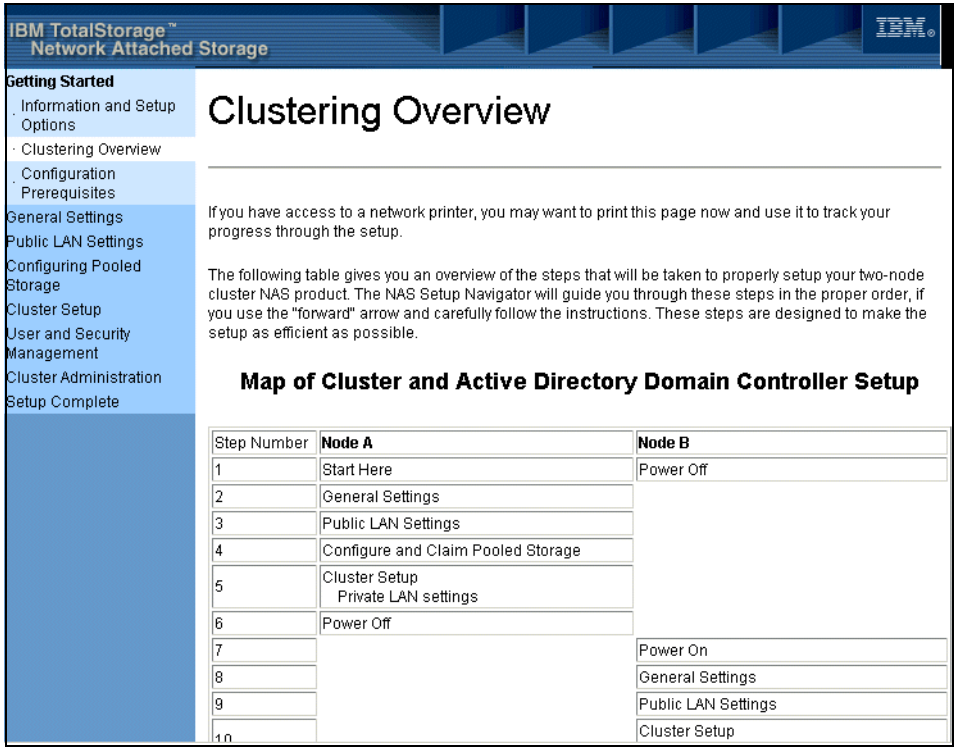

Table 6-1 NAS Setup Navigator — Clustering Overview screen

This screen gives you a “map” for setting up a IBM TotalStorage NAS appliance cluster. In this section we will use the navigator and follow the steps on the map. Before you get started, you should have the necessary information to complete the NAS installation. Gather the information and print it out to have a record of the configuration. This would be helpful, should you ever need to reload a system in the cluster. The information that is needed appears on the next screen, Configuration Parameters. See Figure 6-2.

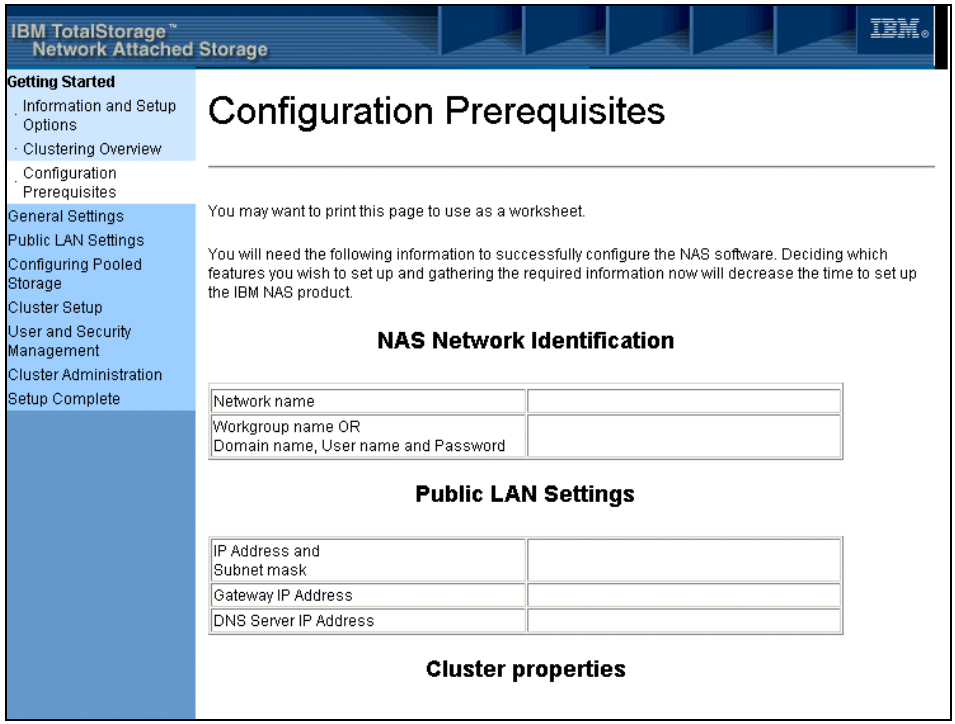

Figure 6-2 NAS Setup Navigator — Configuration Prerequisites screen

As you can see, gathering this information will help the NAS setup and cluster installation go smoothly. The next series of steps in the NAS Setup Navigator are general settings that you will need to change to align with your current IT policies and procedures.

5. Follow the steps and screens and configure node A. When you are done with each screen, click the **Forward** button on the NAS Setup Navigator to advance to the next step. Configure node A with the appropriate information for: System Language, Administrator Password, and Date and Time. Each NAS Setup Navigator screen will have a link to bring up those configuration windows listed above.

6. Configure Network Identification on your NAS 300 by clicking the Network Identification link. This will bring up the system properties window. Click the **Network Identification** tab and then **Properties**. We have chosen to label our server as NAS300A, as you can see in Figure 6-3.

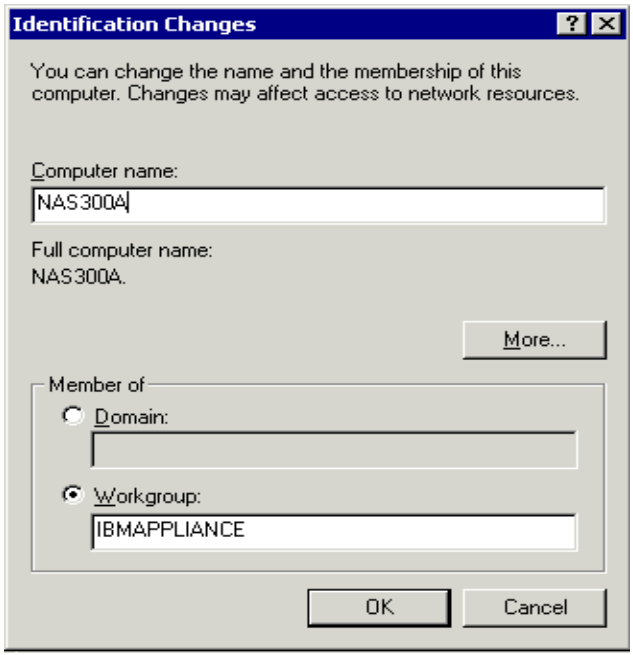

Figure 6-3 Computer Identification NAS 300A

7. Now you have the option to join the NAS 300A to your environment. If you have a domain and want to use it for your cluster go ahead and join, otherwise leave it as workgroup and we will use the domainlet feature later on. If you want the NAS 300 to be part of a domain, then check **domain** and enter the name of the domain you would like to join. A pop-up window will ask for the username and password of a domain administrator or equivalent. We are going to use the domainlet feature, so will leave it as workgroup. Either way, when you are finished, click the **OK** button. A reboot is required.
8. You will now continue with the NAS Setup Navigator by configuring Public LAN settings. Configure your IBM TotalStorage NAS appliance with the appropriate information for Public LAN Settings by clicking the **Network and Dial-up Connections** link. This will bring up the Network and Dial-up Connections window, as shown in Figure 6-4.

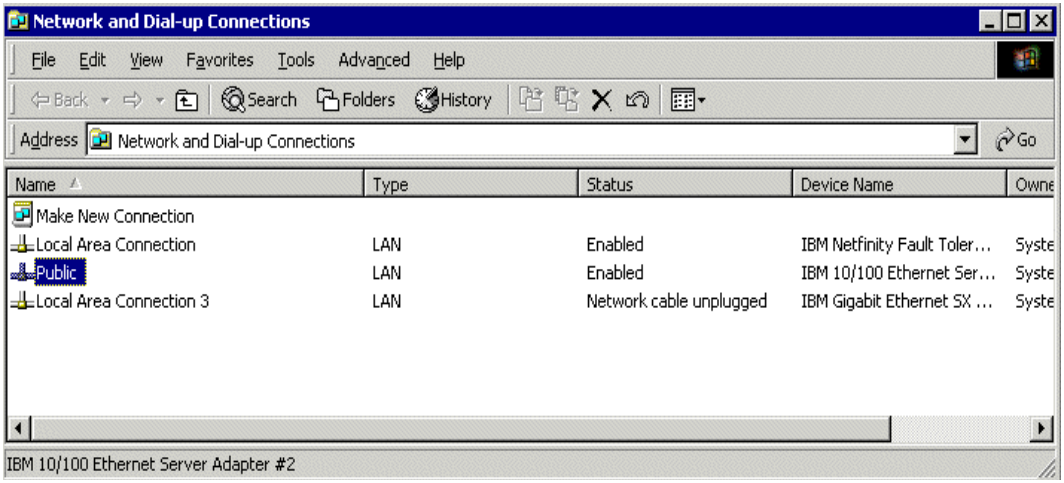

Figure 6-4 Network and Dial-up Connections

9. Select the adapter you are going to use for connectivity to your network. For our example, we will use the IBM 10/100 Adapter. Right-click it, and choose **Rename**. Rename this LAN connection to Public. Right-click again and choose **Properties**.
10. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 6-5).

Note: We will rename both the adapters to Public and Private. This is not required, but it is common practice in the Microsoft Cluster Server (MSCS) setup to know which LAN connection is for the internal cluster and which is for external network connection.

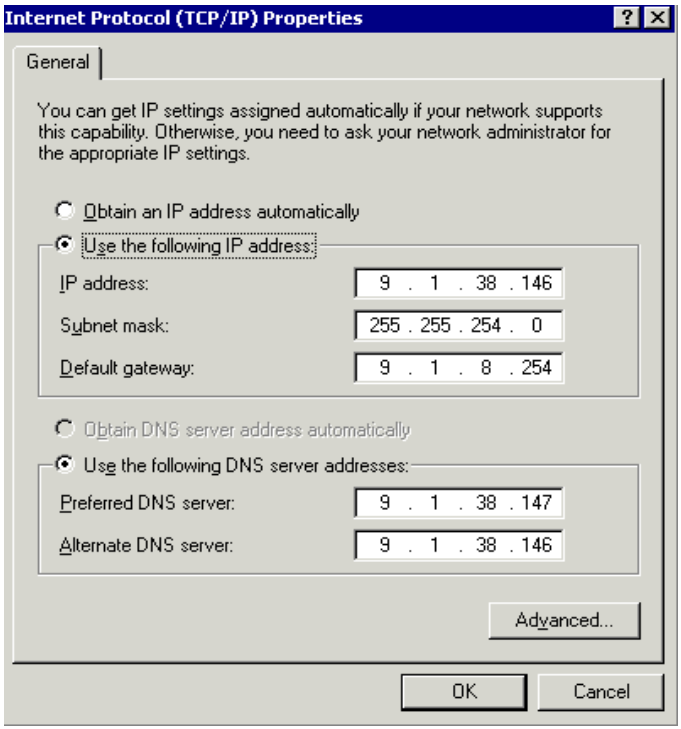

Figure 6-5 Public IP Address for NAS300A

11. Our example shows an IP address of 9.1.38.146. Configure the IP address and settings for your network. You should get the following information from your network administrator:
    - IP address
    - Subnet Mask
    - Default Gateway
    - DNS Server
      (for Domainlet, use one of the nodes for preferred, the other for alternate)
    - DNS Suffix

The DNS Suffix and WINS, etc., can be found under the **advanced** tab. When you are finished configuring your network adapters, click the **OK** button and exit back to the NAS Setup NAvigator. Click the **Forward** button.

6.1.2 Configuring the storage

1. You are now going configure the Storage on the FAStT200. You can accept the storage as it comes from IBM. If so, then follow the directions on the NAS Setup Navigator screens. If you want different settings, then click the **Forward** button on the NAS Setup Navigator and click the IBM FAStT Storage Manager Client link.
2. This will start the IBM FAStT Storage Manager Client utility shown in Figure 6-6.

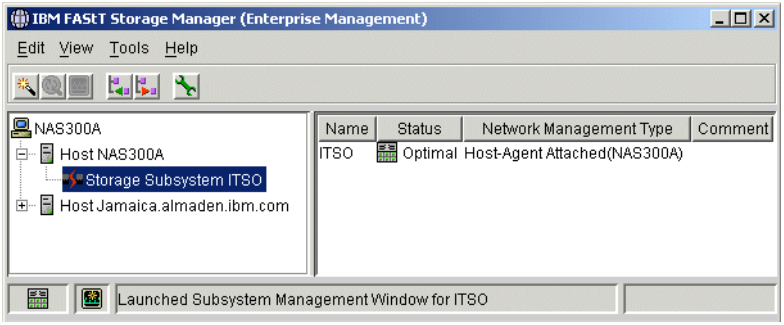

Figure 6-6 IBM FAStT Storage Manager

3. Right-click the Storage Subsystem and select **Manage Device**. (Figure 6-7).

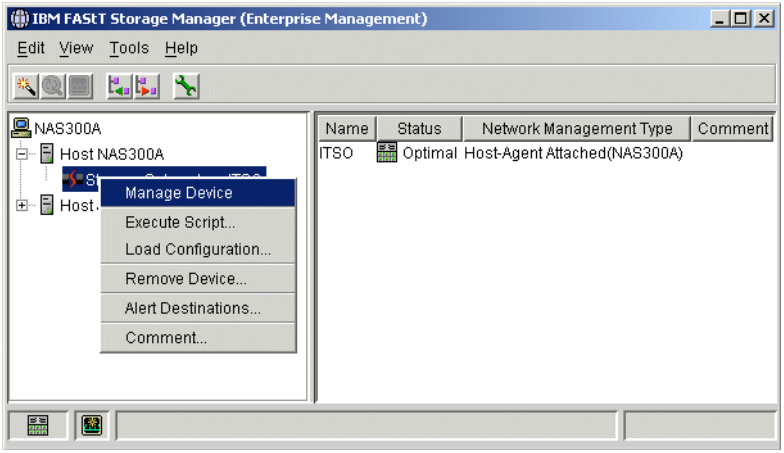

Figure 6-7 Starting Subsystem Management

4. The Logical/Physical View of the Subsystems Management window opens up. It is split into two panes (Figure 6-8). The left pane shows the logical view (arrays, RAID level, logical drives, free capacity, etc.). The right pane shows the physical view (drive, controller enclosure, used disks, hot-spare, etc.).

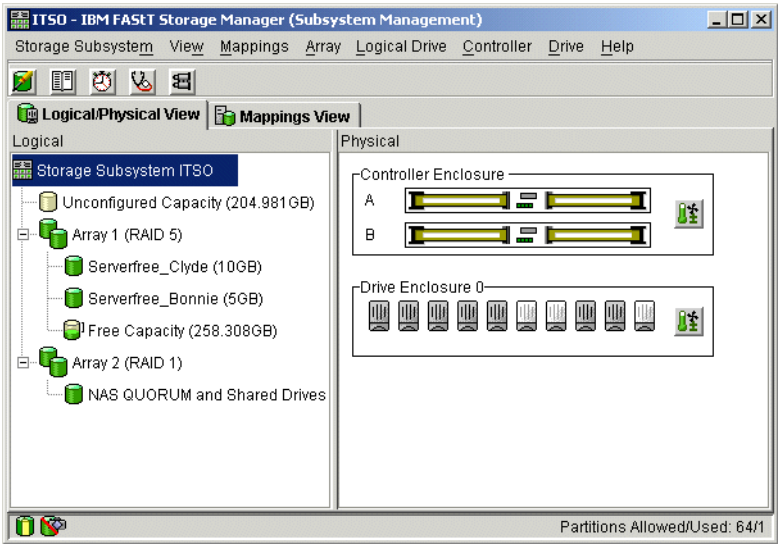

Figure 6-8 Subsystems Management

5. In a larger storage environment with several storage expansion units in a rack, sometimes it is difficult to identify drives belonging to a particular array. Here the Locate function can be useful. To test it, right-click the array in question and select **Locate**, as shown in Figure 6-9. All physical drives belonging to this array will have the indicator lights flashing.

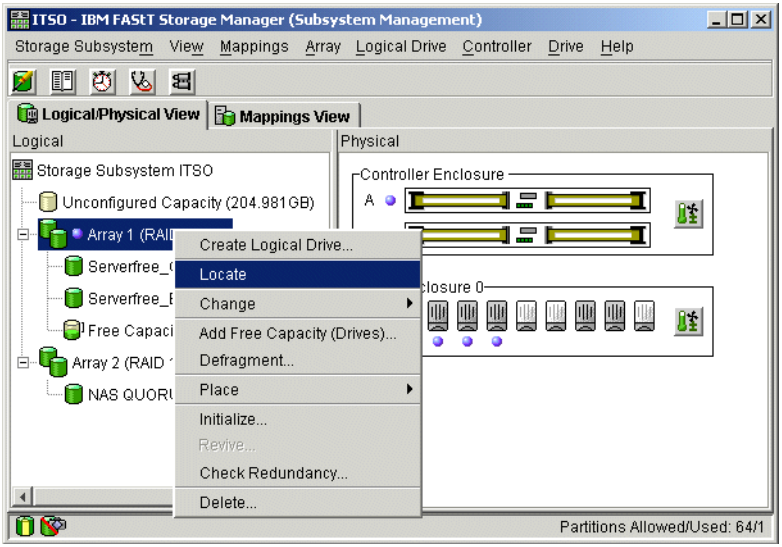

Figure 6-9 Locating drives

6. To stop the indicator lights, simply click **OK** in the window that pops up (Figure 6-10).

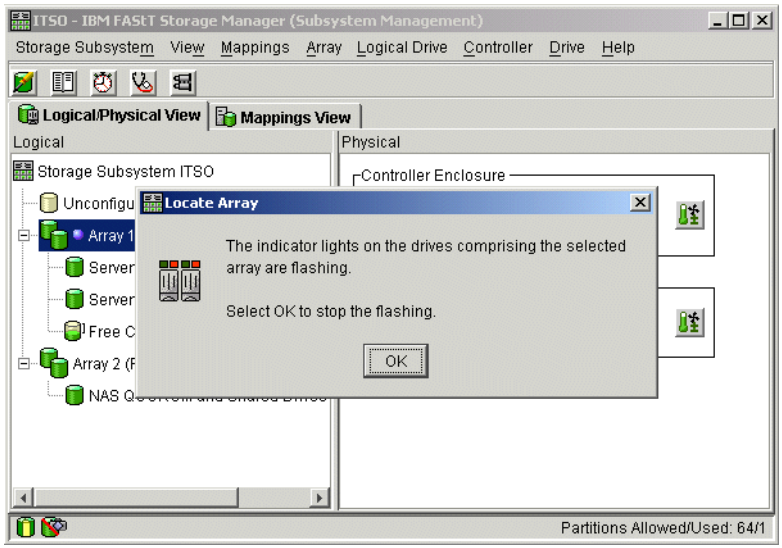

Figure 6-10 Stop the indicator lights

7. To create new storage units, right-click the **Unconfigured Capacity** (available drives will be shown with a purple bullet underneath) and choose **Create Logical Drive**, as shown in Figure 6-11.

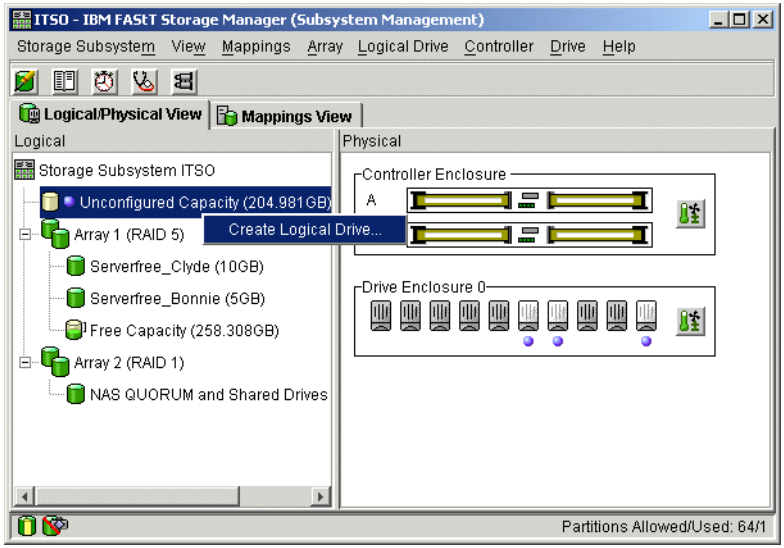

Figure 6-11 Create new storage units

8. If Storage Partitioning has been enabled before, a window similar to the one shown in Figure 6-12 will pop up, informing you of the default host type. If the host system you are planning to attach to this new storage unit is of a different type, you can select it here from the menu.

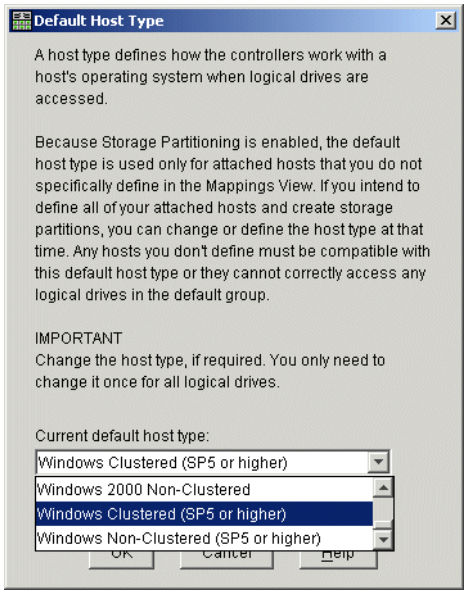

Figure 6-12 Selecting Host Type

9. Now the Create Logical Drive Wizard starts (Figure 6-13). If there are existing arrays, you will be asked whether you want to create a logical drive in the existing array or from the unconfigured capacity (create new array).

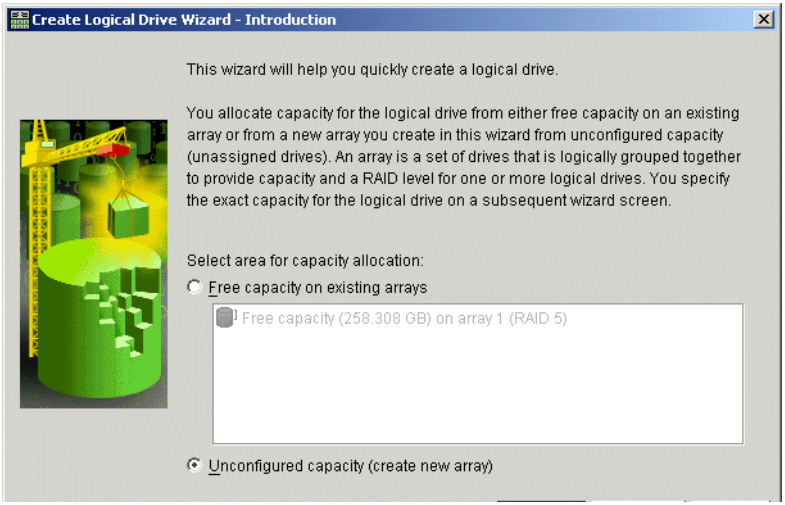

Figure 6-13 Create Logical Drive Wizard

10. On the next screen, under the RAID level selection, choose the RAID level required by clicking the pull-down arrow. Under “Drive selection choices”, select either **Automatic** or **Manual** (Figure 6-14).

11. If you chose Manual, you can define specific drives you want to be members of the array. For multiple drives, press the Ctrl key while selecting the drives. Finally, click **Apply**.

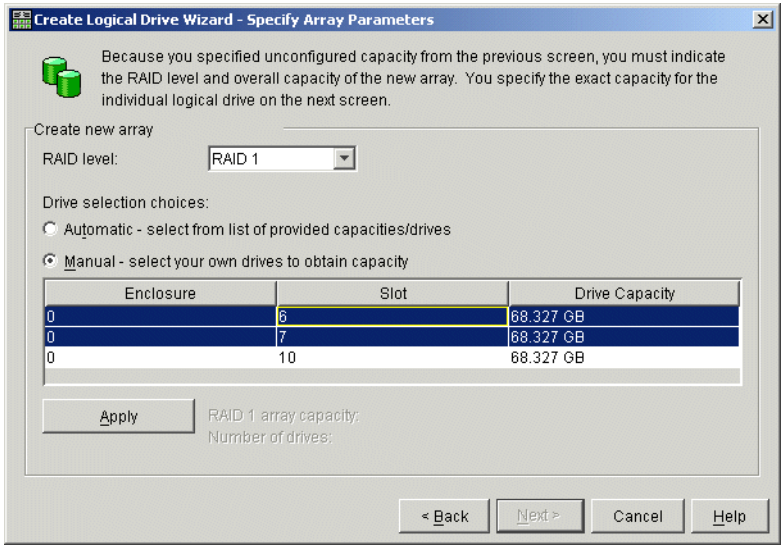

Figure 6-14 Array Parameters

12. Next, define the Logical Drive inside the array you just created.
    For “New logical drive capacity”, type in the size you need and specify it in GB or MB. The default value shown is the maximum capacity of the array.
    The name should be set to something meaningful. If you select **Customize settings** you will be given a chance to define additional settings for this logical drive (Figure 6-15).

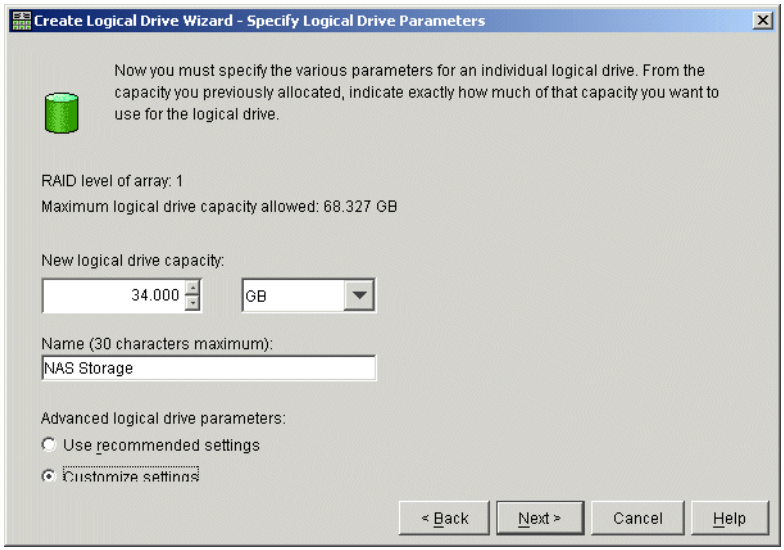

Figure 6-15 Logical Drive Parameters

13. For Advanced settings (Figure 6-16), it is possible to define:

    a. Logical Drive I/O characteristics (optimized for **File system**, **Database**, **Multimedia** or **Custom** — for this option you can customize “Cache read ahead multiplier” and “Segment size”). Since NAS is a file sharing machine, you can leave this setting at default, unless you plan to use it for other type of usage.
    Under “Segment size”, select the optimum value depending on the application that will use the storage on the NAS system. Normally, this could be left at default, which is based on expected usage.

    b. For preferred controller ownership of this logical drive, select either: **Slot A** or **B**.

Tip: To obtain the optimum performance, it is advised that you distribute the ownership of the arrays to both controllers. Controllers are configured for failover by default, so you don’t need to be concerned about a single controller failure.

c. For Logical Drive to LUN mapping, you can choose between **Automatic** or **Map later with Storage Partitioning**.

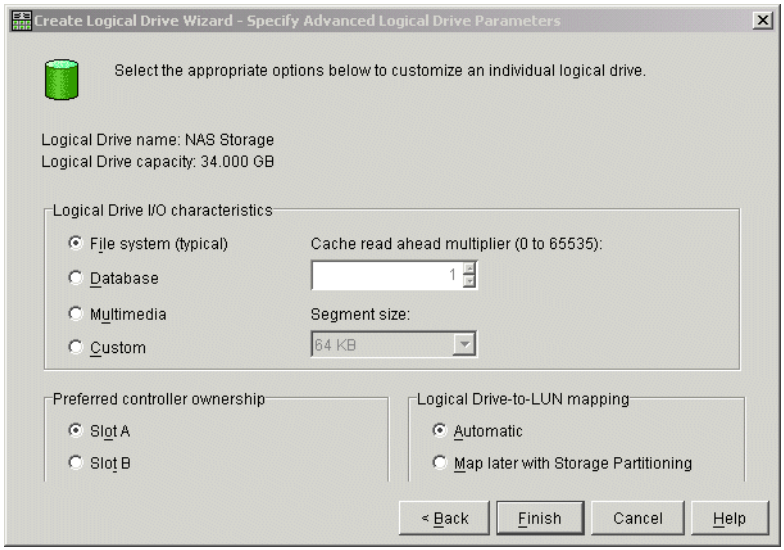

Figure 6-16 Logical Drive Advanced Parameters

14. The logical drive will be created with selected parameters, and you will be asked if you wish to create another logical drive (Figure 6-17).

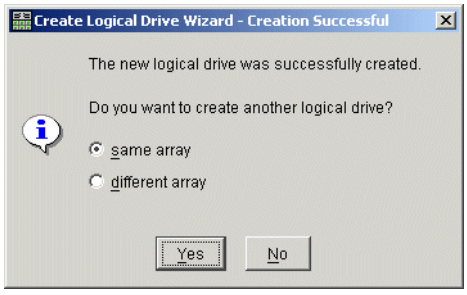

Figure 6-17 Logical Drive created

15. Again you will be notified how to change the Logical Drive to LUN mappings (Figure 6-18).

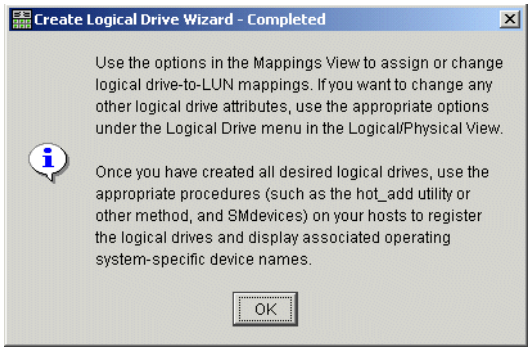

Figure 6-18 Mappings Notice

16. You should now see the array that you just created on the logical view pane, as shown in Figure 6-19. Until the new logical drive has been initialized, it cannot be accessed (shown with a clock symbol in the logical drive icon). Length of the initialization procedure depends on the size of the drive.

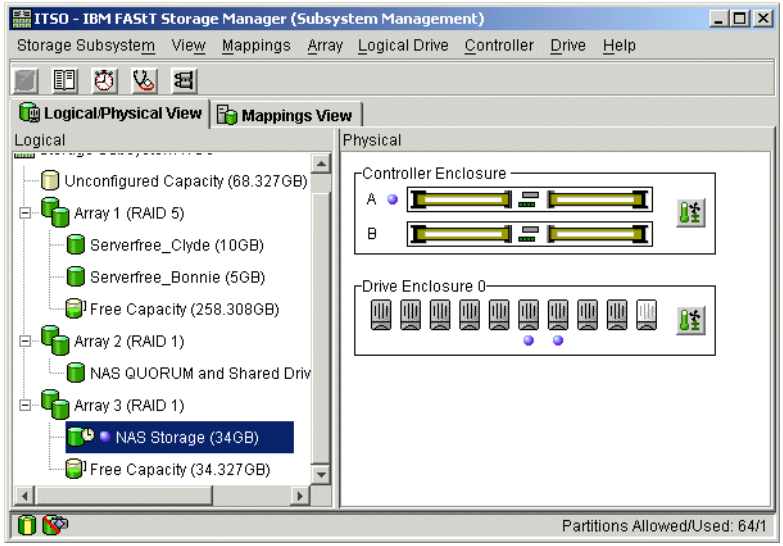

Figure 6-19 Newly created Array and Logical Drive

17. You can check the initialization progress by right-clicking the drive and selecting **Properties**. A window as shown in Figure 6-20 will open, giving you all details about the created drive, with the Initialization Progress bar on the bottom.

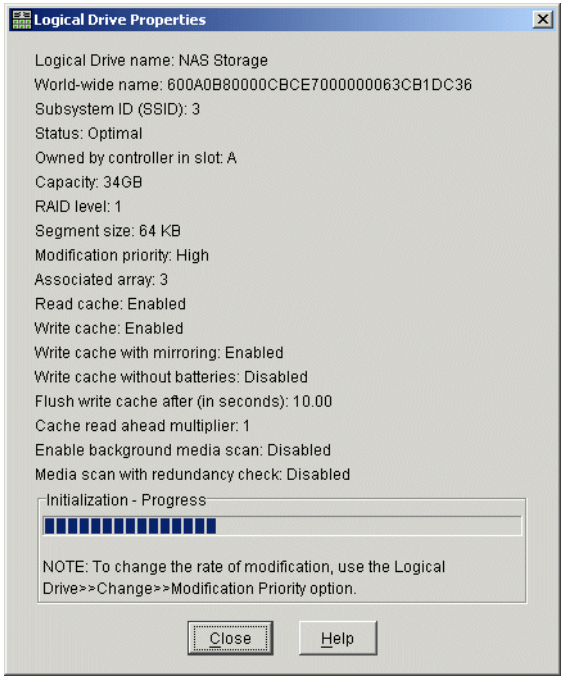

Figure 6-20 Logical Drive Initialization Progress

18. Click **OK** on the Logical Drive Creation Successful window (Figure 6-21).

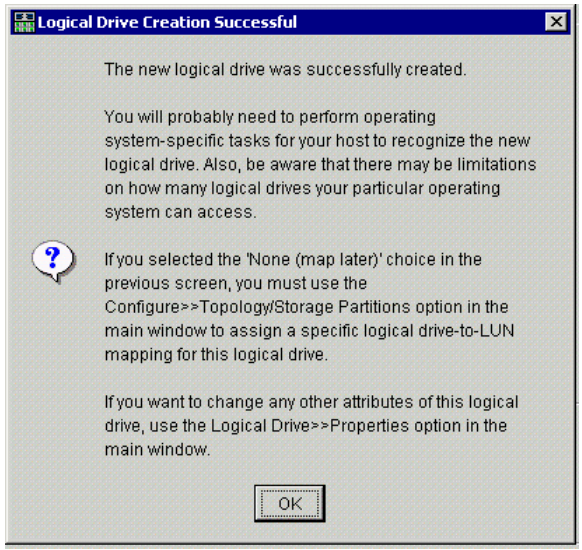

Figure 6-21 Logical drive creation successful

19. Now that you have successfully configured the FAStT200, you must configure the operating system to see those disks. To do this, click the **Disk Management link** to open the disk management window. On this menu, select **Action** and execute **Rescan Disks** as shown in Figure 6-22.

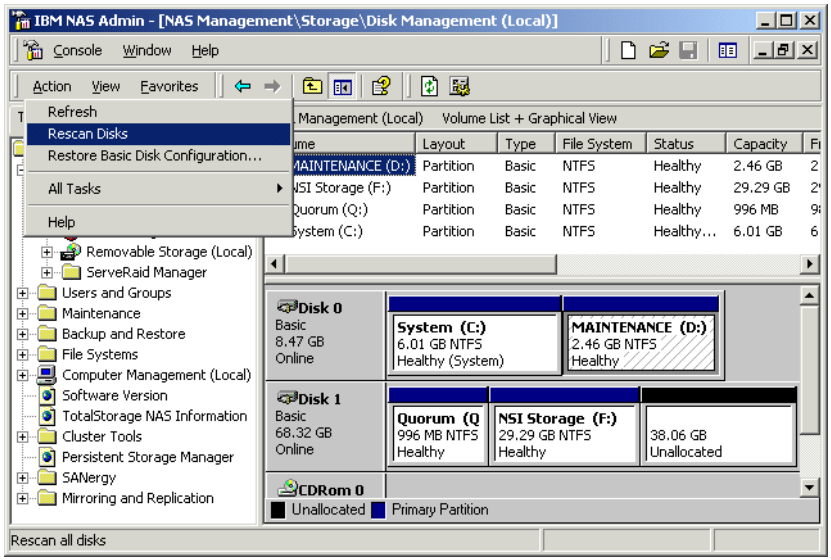

Figure 6-22 Starting Rescan Disks

20. A progress window will open, informing you about the scanning process (Figure 6-23).

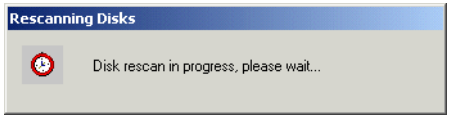

Figure 6-23 Rescanning process

21. Once the rescan is complete, the array that you just created should be shown on the list at the bottom right as unknown and unallocated (Figure 6-24). Before you can use it, the signature has to be written to it first.

Note: Make sure that all disk drives as seen in the Windows 2000 Disk Administrator tool that will be used as clustered devices are of the type “Basic” and not “Dynamic”.

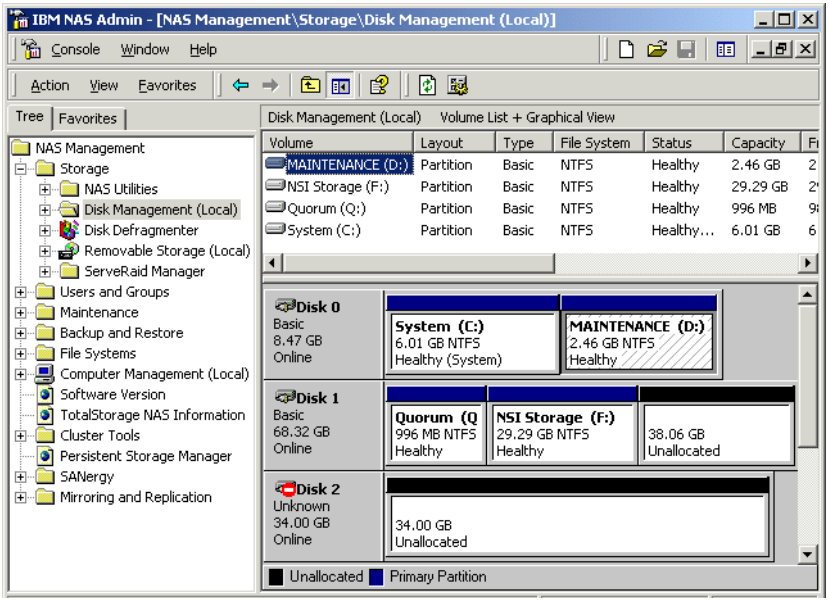

Figure 6-24 Unknown disk found

22. Right-click the unknown disk and select **Write Signature** (Figure 6-25).

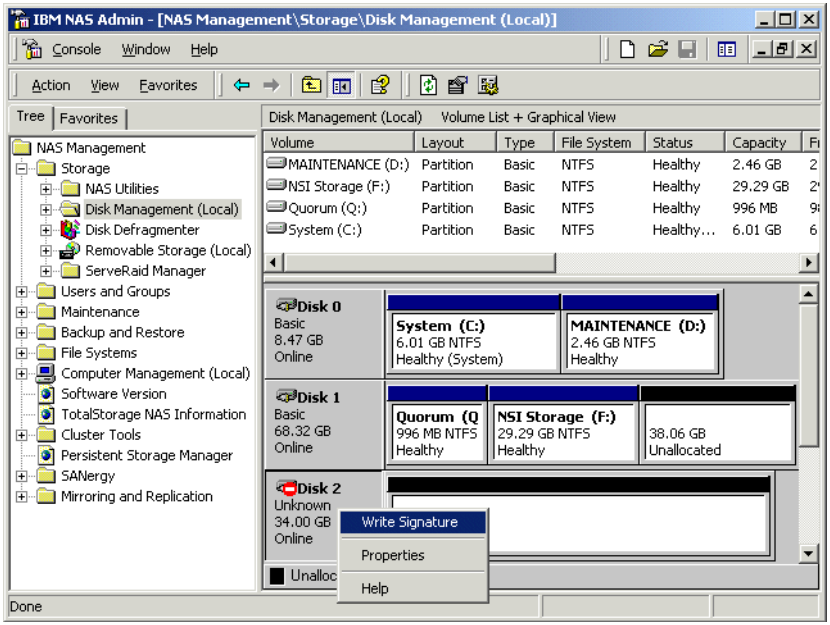

Figure 6-25 Writing a signature to the new disk

23. Click **OK** on the Write Signature window (Figure 6-26).

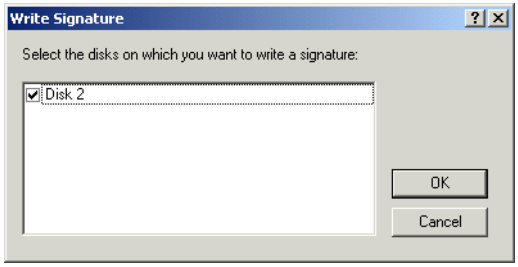

Figure 6-26 Confirming disk signature

24. Now the disk is ready for partitioning. To create a partition inside the new disk, right-click it and select **Create Partition** (Figure 6-27).

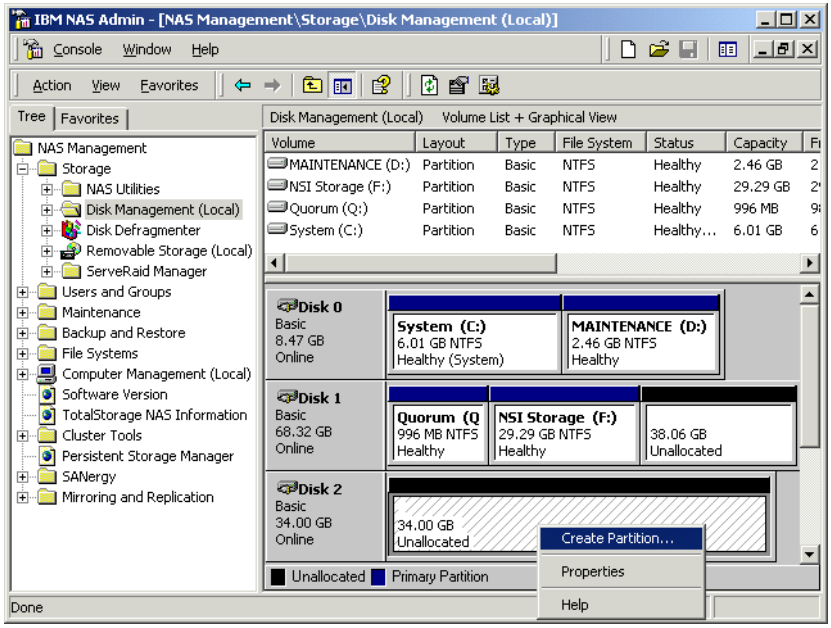

Figure 6-27 Creating a partition

25. The partition wizard will start. Click **Next** (Figure 6-28).

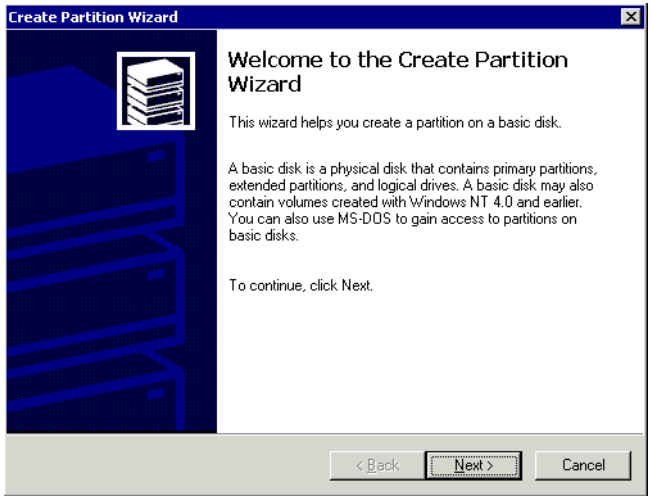

Figure 6-28 Create Partition Wizard

26. On the Partition Type window, select **Primary partition** (Figure 6-29).
27. Then click **Next**.

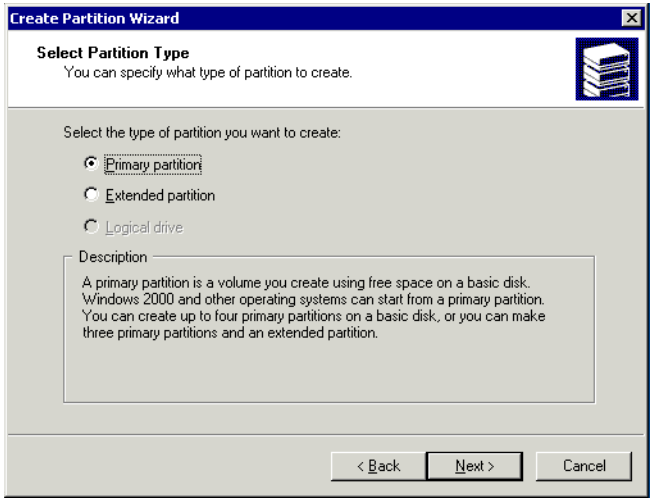

Figure 6-29 Select partition type

28. For “Specify Partition Size” you can set the partition size. We left it at the default, which should be the maximum (Figure 6-30).

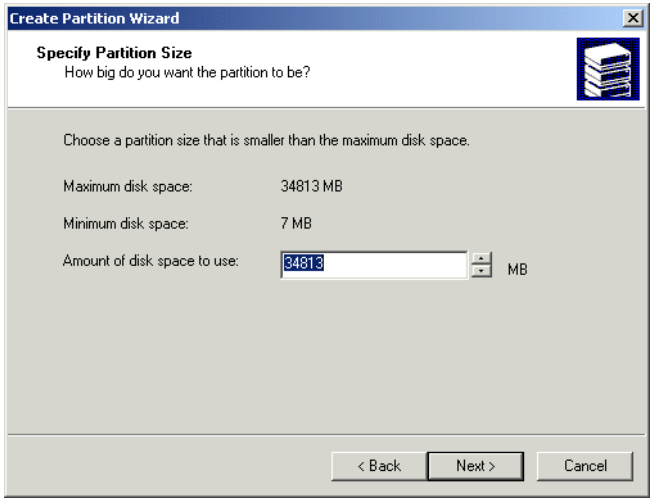

Figure 6-30 Specify partition size

Important: Every drive on a cluster must be single disk single partition. The Microsoft Cluster Server does not support multiple partitions/logical drives on a single disk.

29. On the Assign Drive Letter or Path window, assign a drive letter to this partition (Figure 6-31). Then click **Next**.

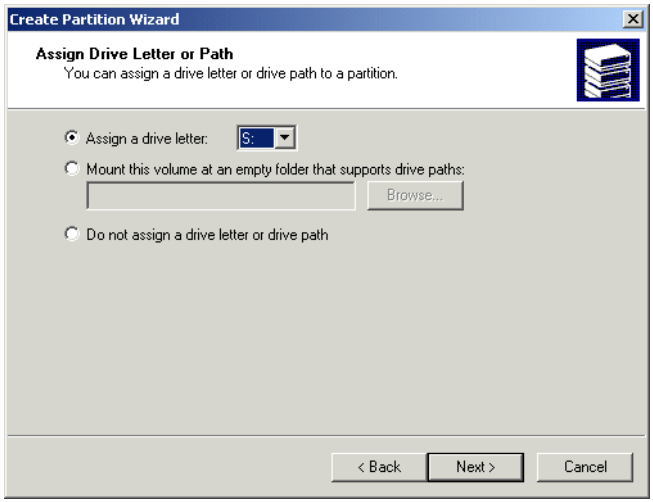

Figure 6-31 Assign drive letter

On the Format Partition window, select **Format this partition with the following settings** (Figure 6-32).

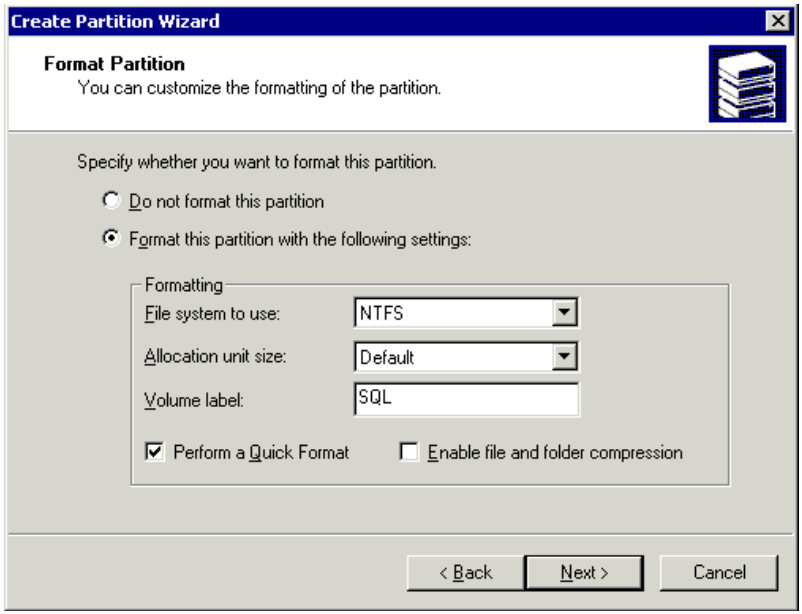

Figure 6-32 Format partition settings

For “File System to use:” make sure that NTFS is selected.

Important: Microsoft Cluster only supports NTFS partitions.

You can specify the allocation unit size or leave it at the default value.
For “Volume label”, type in the name you want to associate with this partition.
Click **Perform a Quick Format** (this will save you some time), then click **Next**.

Important: Do ***not*** select **Enable File and Folder Compression**, since this is not supported by Microsoft Cluster.

30. On the final window, just click **Finish** (Figure 6-33).

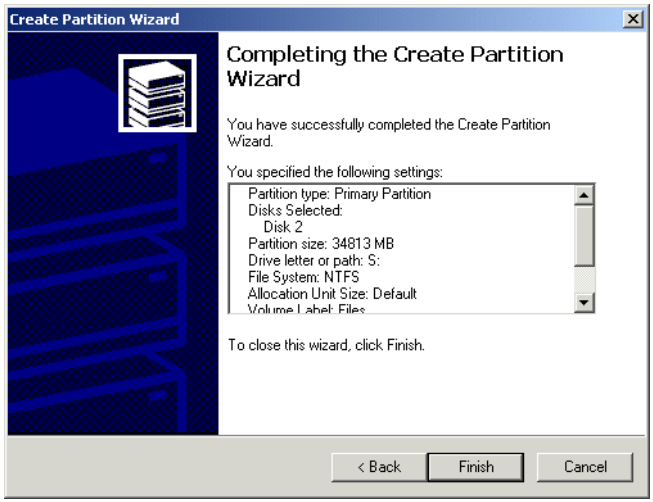

Figure 6-33 Completing the Partition Wizard

31. Your partition is now ready for use (Figure 6-34). Close all the screens to get back to the NAS Setup Navigator.

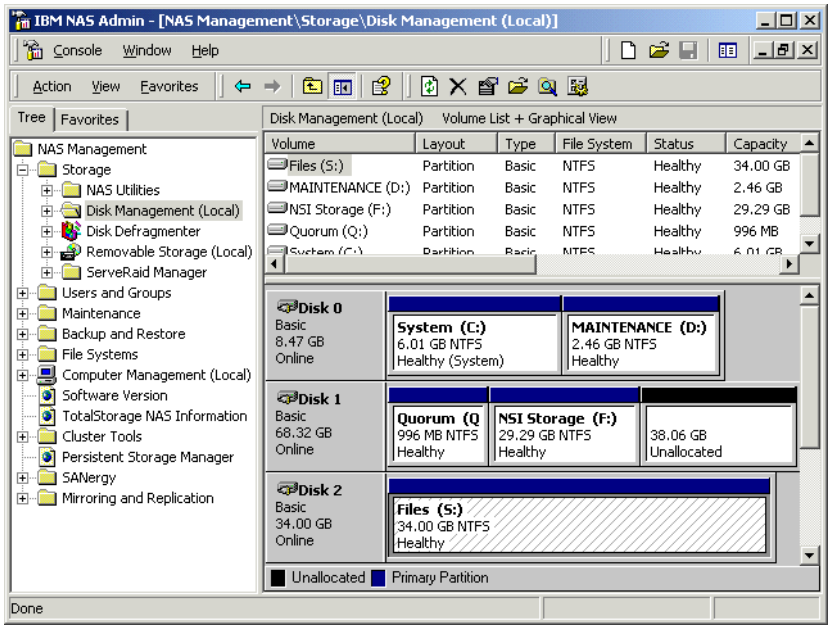

Figure 6-34 Partition ready

6.1.3 Preparing the Cluster nodes

1. We will now configure the network settings needed for the cluster. Click the **Network and Dial-up Connections** link which brings up the corresponding window. Right-click the IBM Netfinity Fault Tolerance PCI adapter and click **Properties**. Rename the adapter to **Private** as in (Figure 6-35). Right-click again and choose **Properties**.

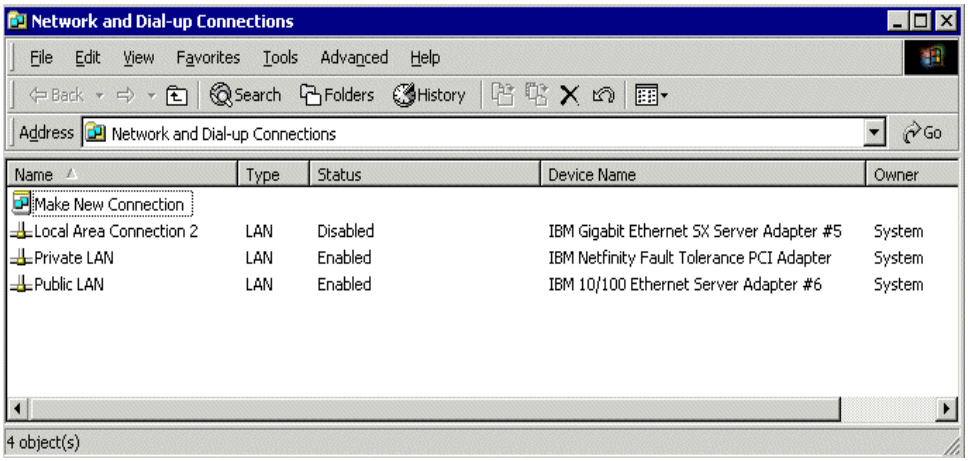

Figure 6-35 Private Network LAN Connection

2. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 6-36). Configure the private IP address for the cluster with a static IP address. We have chosen the values shown in Table 6-2.

Table 6-2 IP configuration — 10.1.1.1

Config Name	Value
IP Address	10.1.1.1
Subnet Mask	255.255.255.0
DHCP	OFF
WINS — NetBIOS over TCP/IP	Disabled

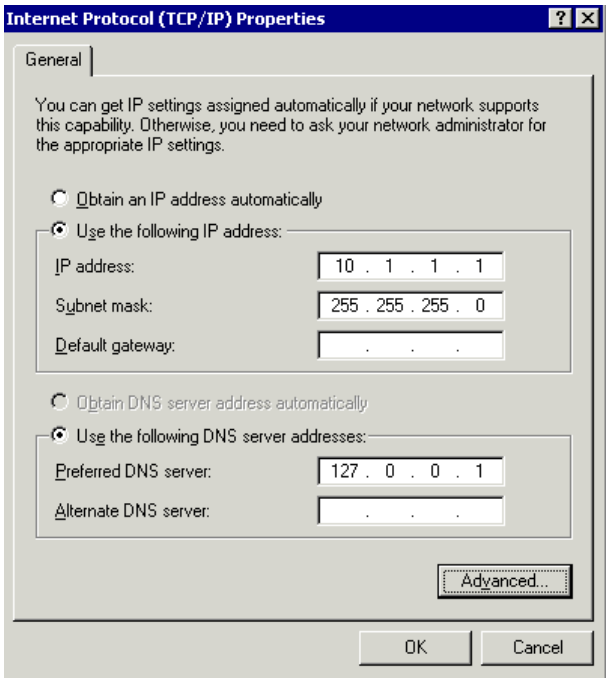

Figure 6-36 Private Address for Cluster Node

3. When you have finished configuring the private LAN, click **OK** to get out of both windows and then **Close**. You will be back at the NAS Setup Navigator screen.

Note: Microsoft Cluster setup requirements dictate that both nodes have seen the resources before they can join/create a cluster. The steps you have performed up until now have allowed the first system to create and see the resources. The second node will actually be the node that creates an Active Directory (if needed) and the cluster.

4. After setting up the private LAN settings, turn off node A and turn on node B.

Attention: Make sure you turn off Node A before configuring Node B.

5. Connect to node B. Most of the steps you performed on Node A you will now perform on Node B.

6. For configuration of node B, you follow the NAS Setup Navigator screens just as you did with node A. After you have selected Node B and Setup active Domain Controller in the Information and Setup Options screen, click **Apply**. Scroll down to the bottom of the screen and click **Forward**. You will see the Cluster Overview and Configuration prerequisite screens again.
7. Go ahead and insert the appropriate configuration parameters for System Language, Administrator Password and Date and Time, advancing through each NAS Setup Navigator screen as you go.
8. Configure Network Identification on your NAS 300 by clicking the **Network Identification** link. This will bring up the system properties window. Click the **Network Identification** tab and then **Properties**. We have chosen to label this server as NAS300B, as you can see in Figure 6-37.

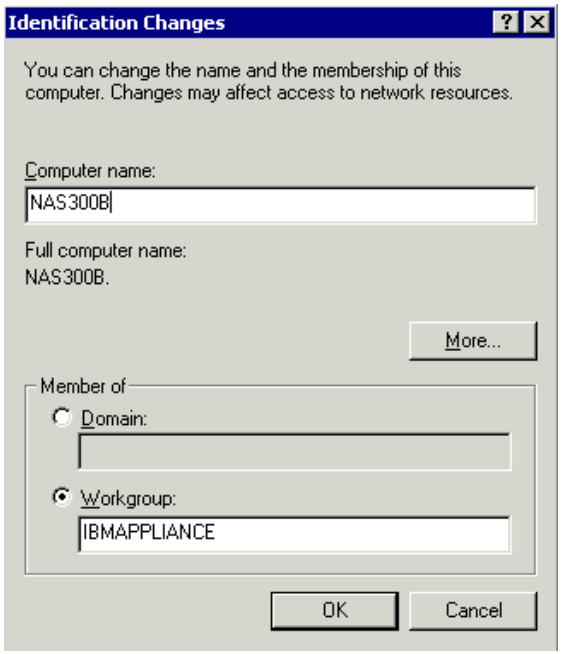

Figure 6-37 Computer Identification NAS 300B

- Now you have the option to join the NAS 300B to your environment. If you had joined a domain with the NAS300A server, then join the same domain. We are going to use the domainlet feature, so we are leaving them in a workgroup. When you are finished, click the **OK** button. A reboot is required.

9. You will now continue with the NAS Setup Navigator by configuring Public LAN settings. Configure your IBM TotalStorage NAS appliance with the appropriate information for Public LAN Settings by clicking the **Network and Dial-up Connections** link. This will bring up the Network and Dial-up Connections window as shown in Figure 6-38.

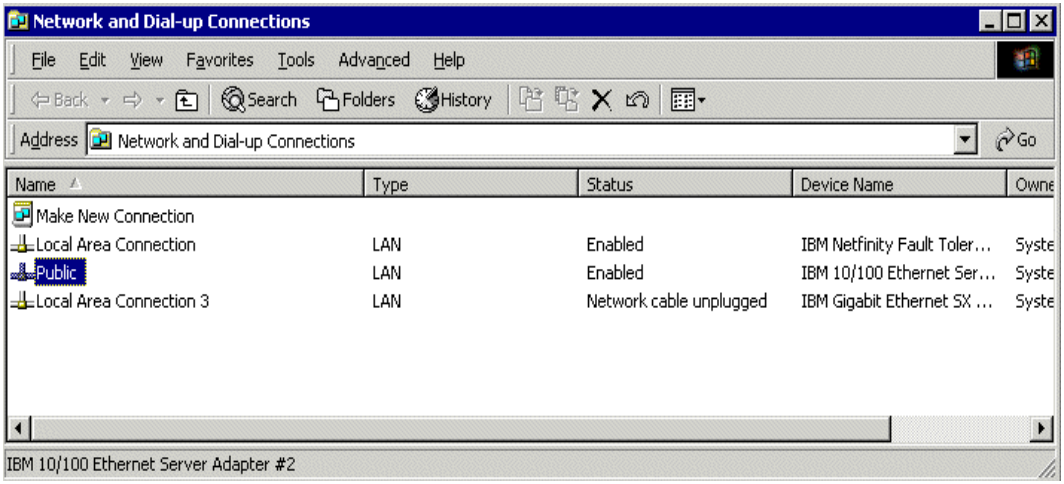

Figure 6-38 Network and Dial-up Connections

10. Select the adapter you are going to use for connectivity to your network. For our example we will use the IBM 10/100 Adapter. Right-click it, and choose **Rename**. Rename this LAN connection to Public. Right-click again and choose **Properties**.

11. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 6-39).

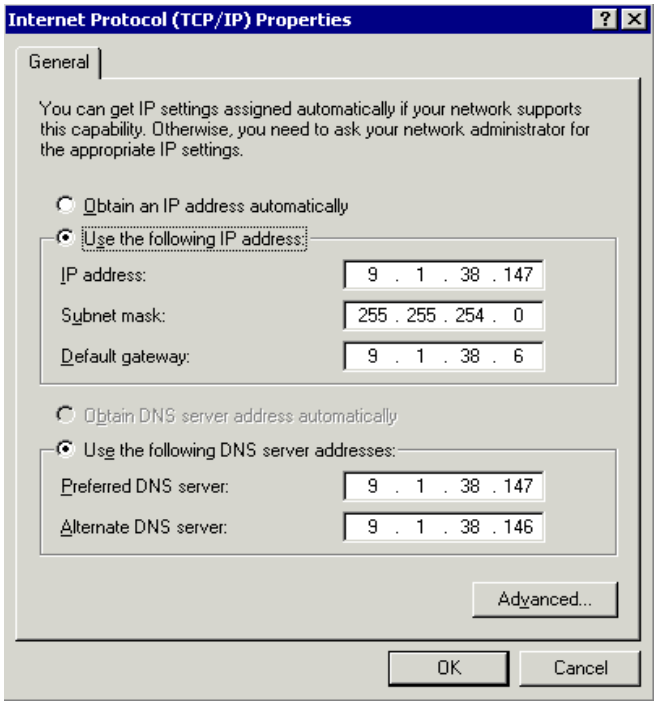

Figure 6-39 Public IP Address for NAS300B

12. Our example shows an IP address of 9.1.38.147. Configure the IP address and settings for your network. You should get the following information from your network administrator:
    - IP address
    - Subnet Mask
    - Default Gateway
    - DNS Server
      (for Domainlet, use one of the nodes for preferred, the other for alternate)
    - DNS Suffix

The DNS Suffix and WINS, etc., can be found under the **Advanced** tab. When you are finished configuring your network adapters, click the **OK** button and exit back to the NAS Setup Navigator. Click the **Forward** button.

Note: You will not need to set up the initial pooled storage or FAStT200 drives, as node B will see those resources which were created on Node A.

13. We will configure the network settings needed for the cluster. Click the **Network and Dial-up Connections** link which brings up the corresponding window. Right-click the IBM Netfinity Fault Tolerance PCI adapter and click **Properties**. Rename the adapter to **Private** as shown in Figure 6-40. Right-click again and choose **Properties**.

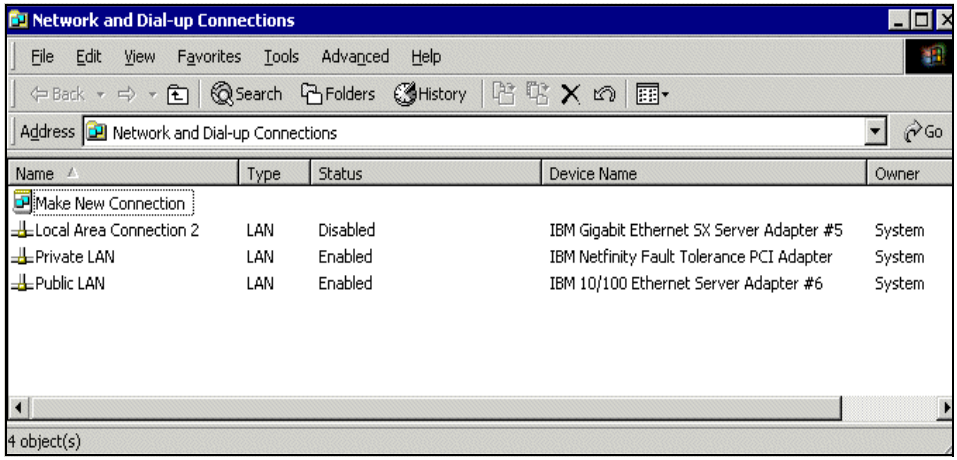

Figure 6-40 Private Network LAN Connection

14. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 6-41). Configure the private IP address for the cluster with a static IP address. We have chosen the values shown in Table 6-3.

Table 6-3 IP configuration — 10.1.1.2

Config Name	Value
IP Address	10.1.1.2
Subnet Mask	255.255.255.0
DHCP	OFF
WINS — NetBIOS over TCP/IP	Disabled

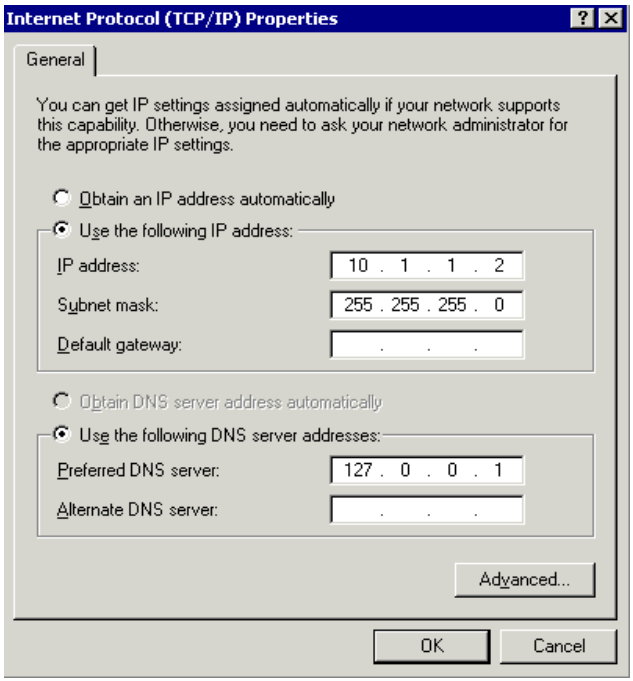

Figure 6-41 Private Address for Cluster Node B

15. When you have finished configuring the private LAN, click **OK** to get out of both windows and then **Close**. You will be back at the NAS Setup Navigator screen.

6.1.4 Configure Active Directory integration

At this point in the installation you may need to configure Microsoft's Active Directory. The NAS Setup Navigator will link you to the Active Directory wizard as shown in Figure 6-42. Pay special attention to the steps and notes included on this page.

Note: There is a requirement for a Domain Name Services DNS or Dynamic DNS. If a DNS server exists, you can have the IBM TotalStorage NAS appliances use that DNS. If no DNS is present and is needed, the Active Directory Wizard will inform you of this and lead you through the steps to create one. If a DNS server is created on one of the Nodes in the cluster, the other node will require it also for failover purposes.

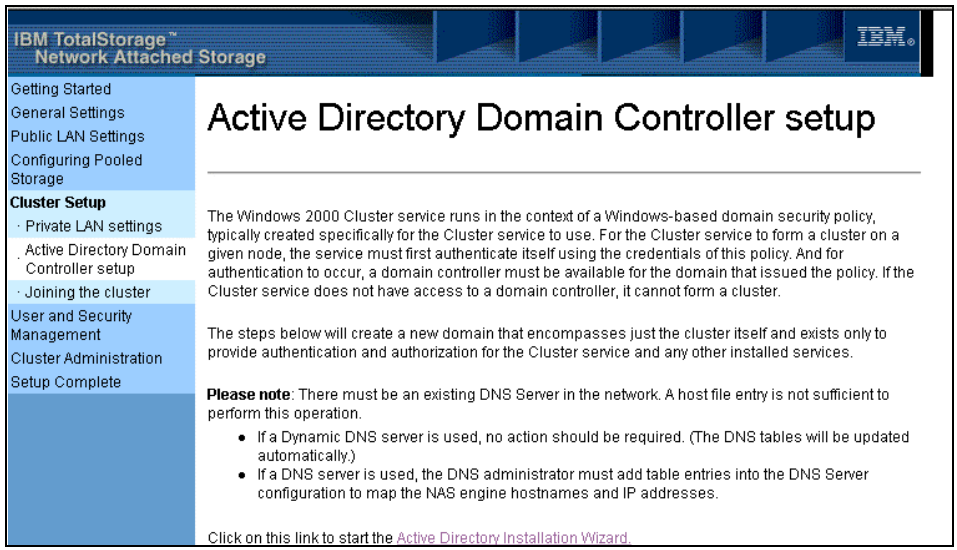

Figure 6-42 NAS Setup Navigator — Active Directory Setup

Note: If you already have a domain or Active Directory to use and do not need to create one on the NAS appliance, skip the next step and continue with “After node B restarts”.

16. Click the **Active Directory Installation Wizard** link to create the Active Directory on node B. (Since we are using the domainlet feature, we will also require a DNS server.) This will bring up the welcome screen. Click **Next** to continue.

17. You will be asked if you want to create the create the domain controller in a new or existing domain. Select **Domain controller for a new domain** as shown in Figure 6-43.

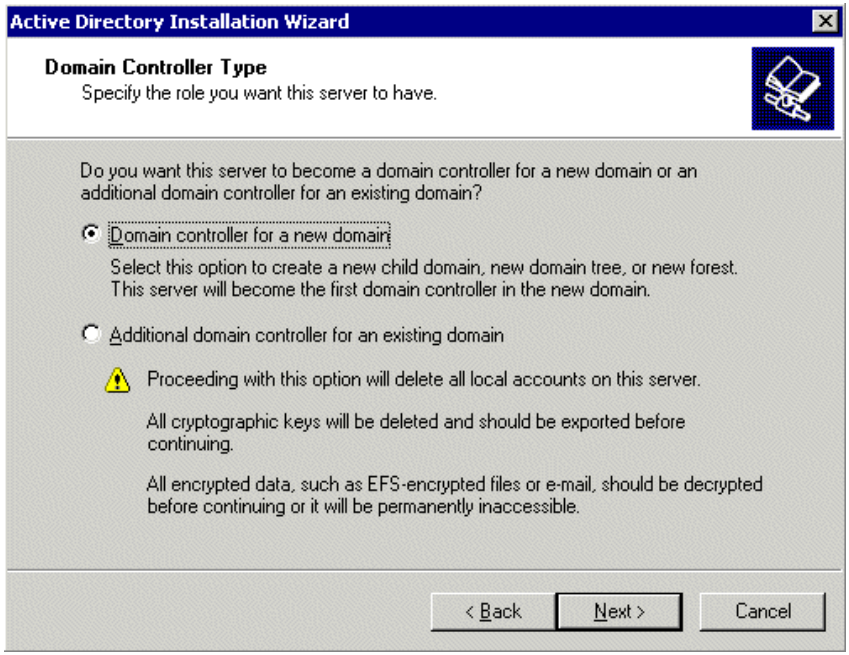

Figure 6-43 Active Directory Wizard — New Domain

18. Click **Next** and the screen shown in Figure 6-44 will appear.

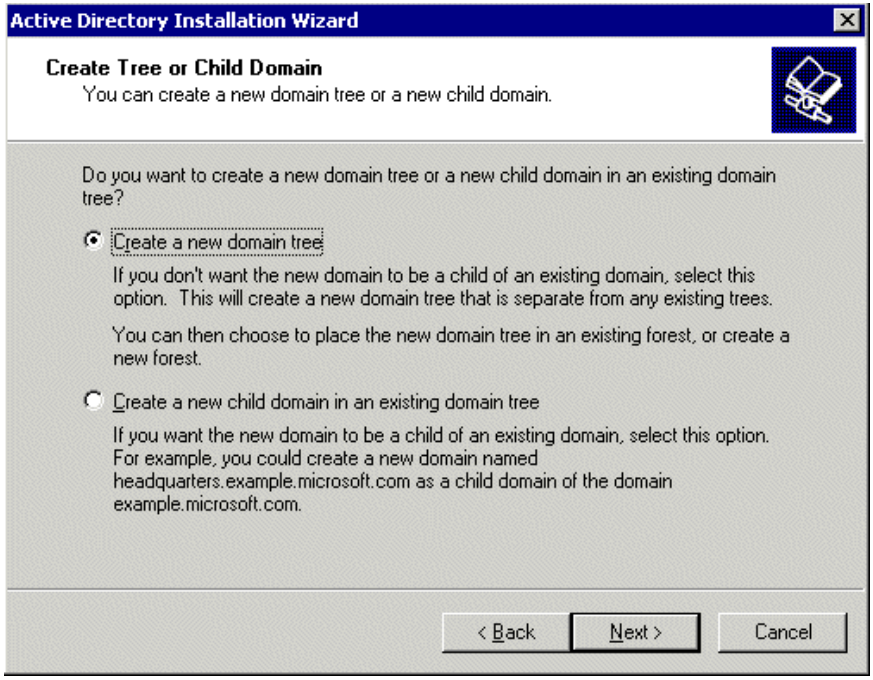

Figure 6-44 Active Directory Wizard — New Domain Tree

19. Next, select **Create a new domain tree** as in Figure 6-44.
20. Click **Next**.

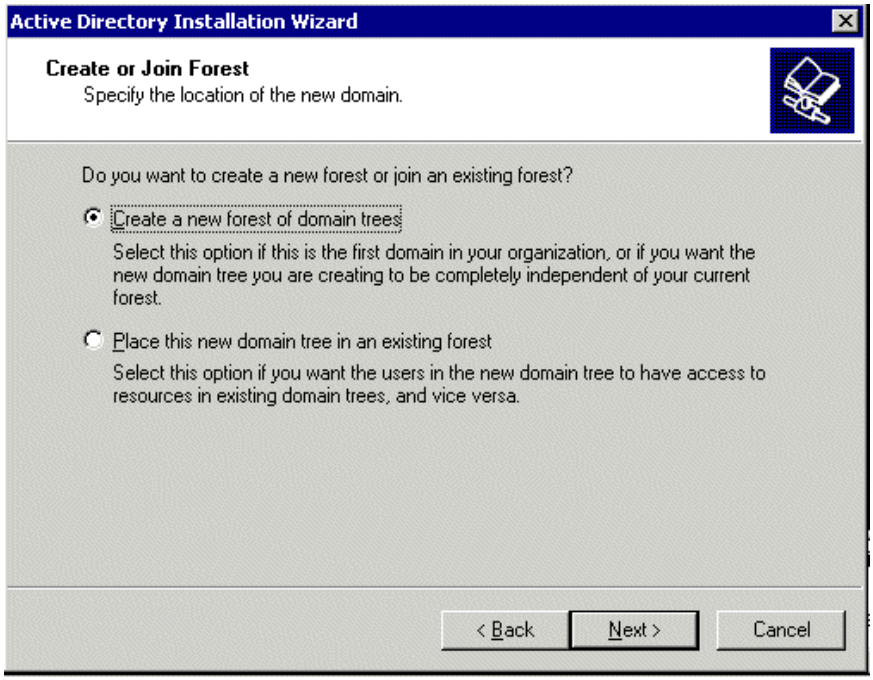

Figure 6-45 Active Directory Wizard — Create or Join Forest

21. Select **Create a new forest of domain trees**, as in Figure 6-45, and then click **Next**.

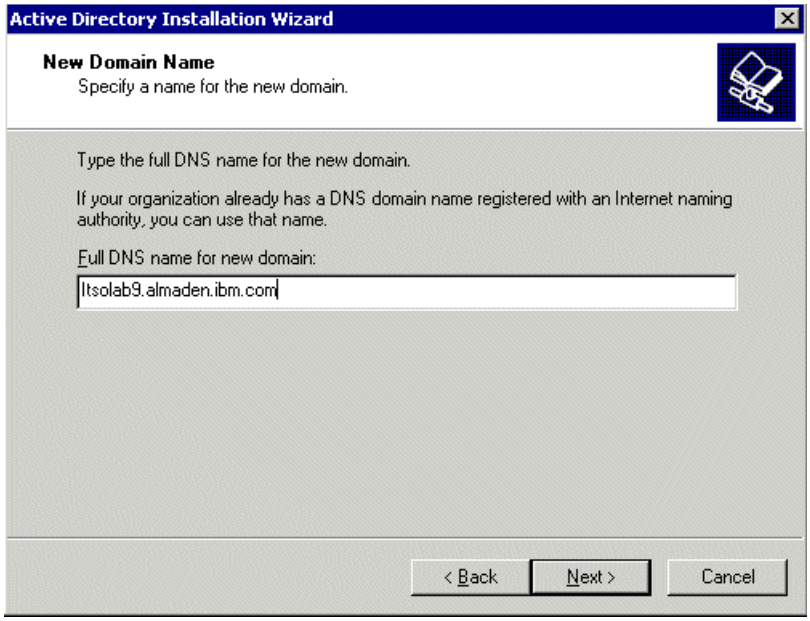

Figure 6-46 Active Directory Wizard — New Domain Name

22. You will now supply a domain name for the new domain as in Figure 6-46. When finished, click **Next**.
23. You will be asked to specify a NetBIOS name; click **Next** to continue.
24. The next two steps will ask you where you would like to store the Database and Log files and also the Shared System Volume. Accept the defaults for each one and click **Next** for each of the screens (Figure 6-47 and Figure 6-48).

Figure 6-47 Active Directory Wizard — No DNS Available

25. After clicking **Next** on the shared system volume screen, a warning will come up saying that the computer could not contact the DNS server. Click **OK** and the following screen will ask you if you would like to install and configure DNS on the computer, as shown in Figure 6-48. Leave it as-is and click **Next**.

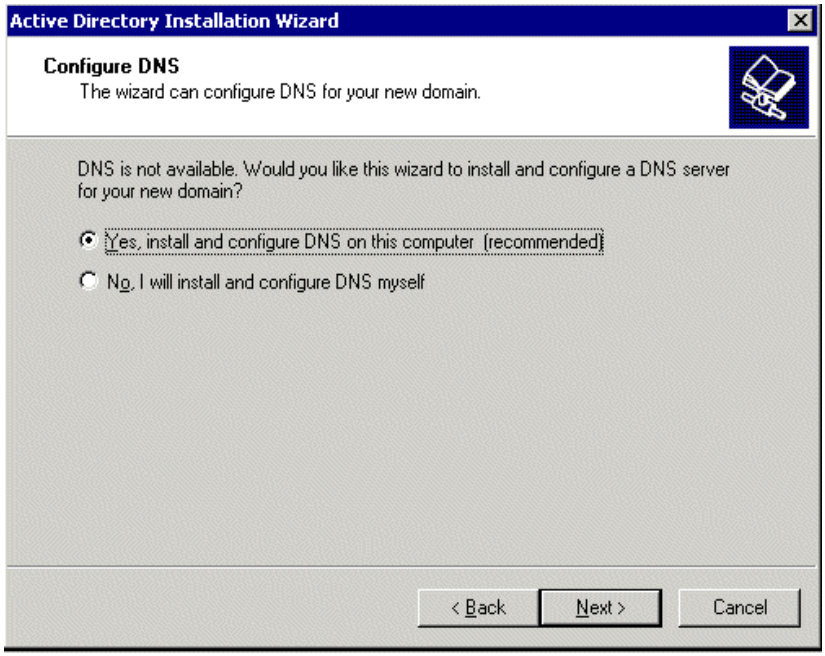

Figure 6-48 Active Directory Wizard — Configure DNS

26. You will now be asked to select the default permissions for users and groups. You can leave it at the default of Permissions compatible with pre-Windows 2000 servers, as shown in Figure 6-49. Click **Next**.

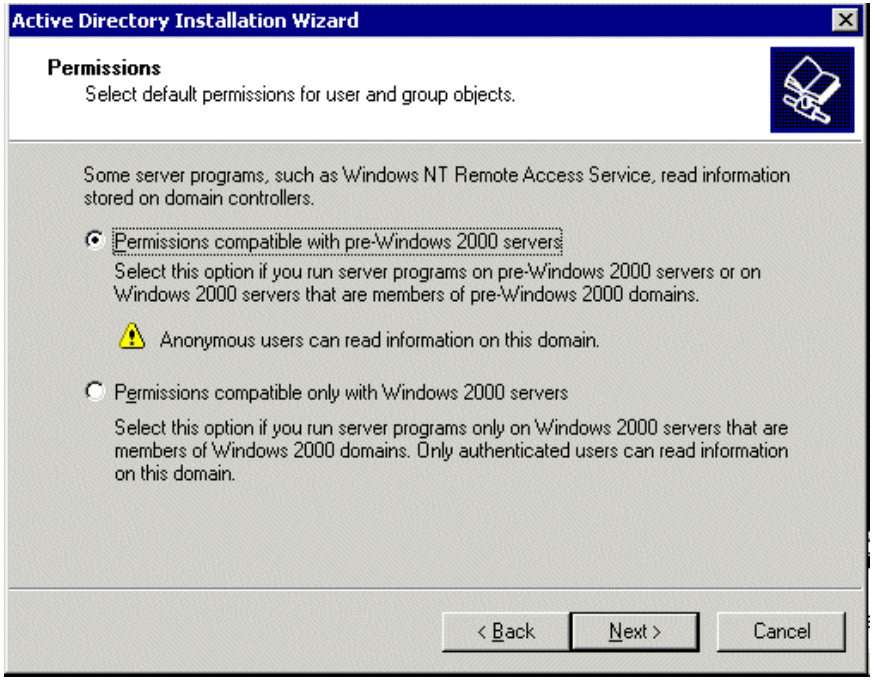

Figure 6-49 Active Directory Wizard — Permissions

27. You will be asked to supply an administrative password to be used when the account is started in Directory Services Restore Mode. Supply a password as shown in Figure 6-50 and click **Next**.

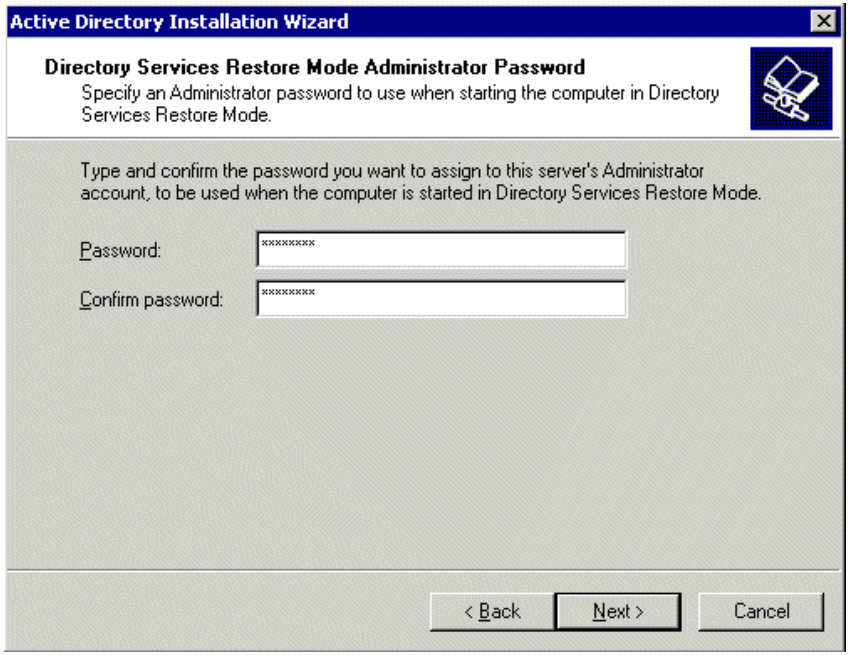

Figure 6-50 Active Directory Wizard — Administrative Password

28. You will be provided with a screen showing a summary of the settings. Click **Next** and the DNS and active directory installation will begin, as shown in Figure 6-51.

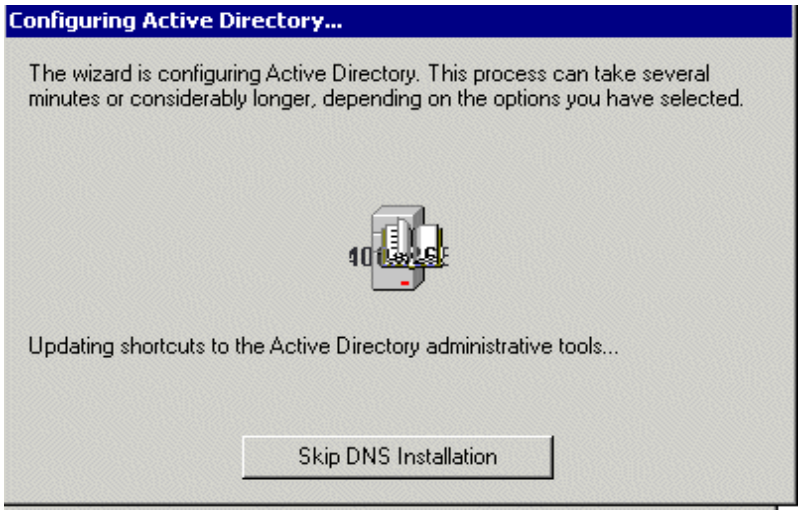

Figure 6-51 Active Directory Wizard — Configuring

29. You will be asked to supply some files to continue the installation, as shown in Figure 6-52. They are located on the drive in c:\I386. Click **OK**, supply the correct path, and then continue with the installation.

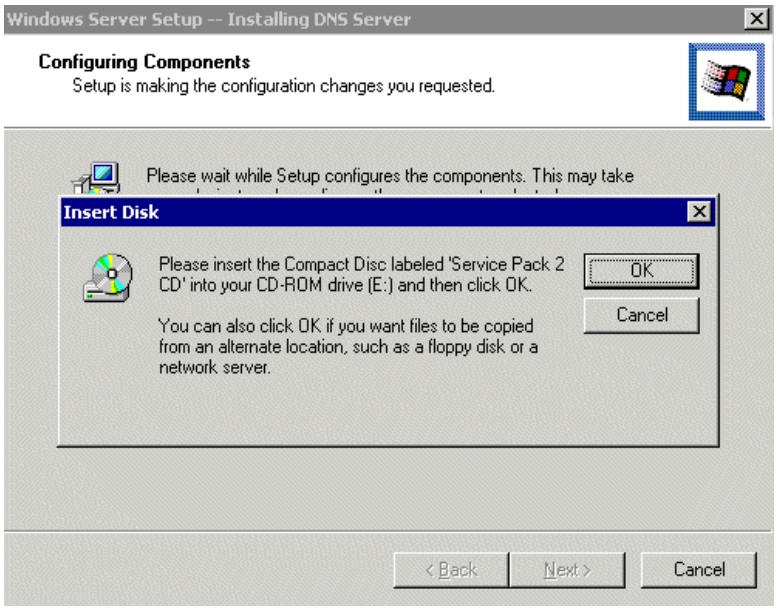

Figure 6-52 Active Directory Wizard — Files Required

30. The installation will continue and you will get a **Completing the Active Directory Installation Wizard** screen. Click **Finish** and then reboot. (Figure 6-53).

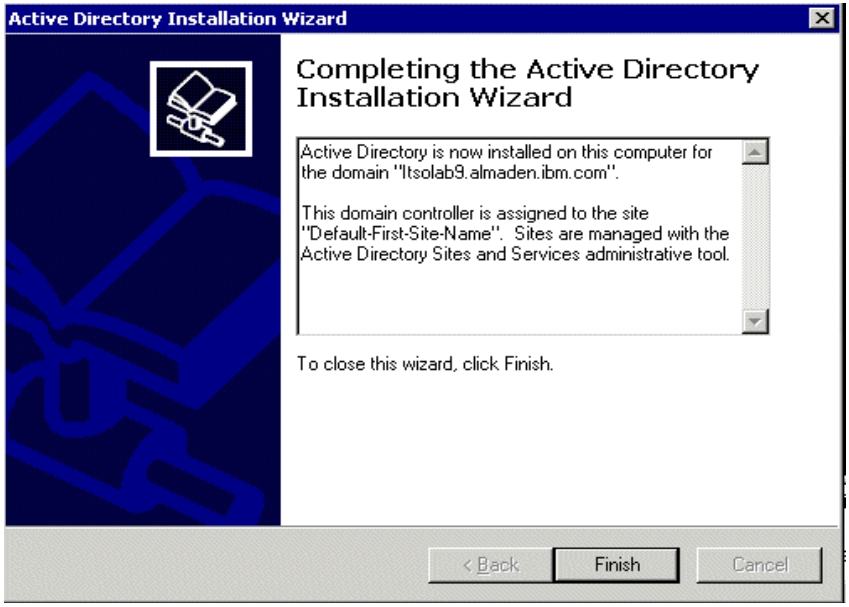

Figure 6-53 Active Directory Wizard — Final Steps

6.1.5 Cluster configuration and setup

1. After Node B restarts, you are ready to create a cluster. In the NAS Setup Navigator screen under Cluster Settings, you will see Creating a Cluster. That NAS Setup Navigator screen will have a Cluster Configuration Wizard link on it. Go ahead and open up the Cluster Configuration Wizard, and select **Continue** (Figure 6-54).

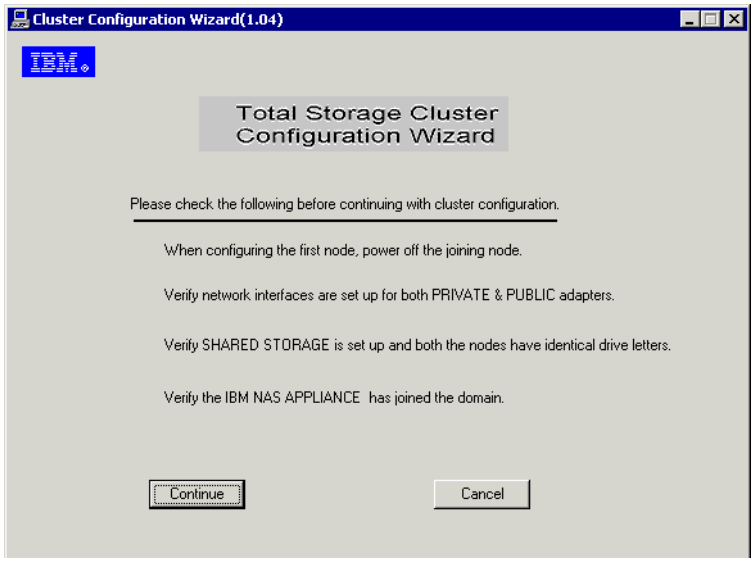

Figure 6-54 Total Storage Cluster Configuration Wizard

2. On the Nodes window, select **First Node** (Figure 6-55).

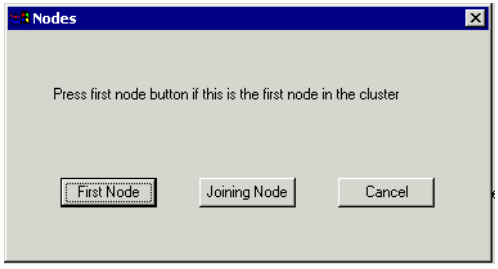

Figure 6-55 Node selection

3. On the Cluster Information window (Figure 6-56), fill in the required information:
   - **Administrator's ID:** The domain administrator account.
   - **Password:** The password for that domain administrator account.
   - **Domain name:** The domain which this machine is a member of.
   - **Cluster name:** The name by which the clients will address the NAS 300.
   - **Cluster IP address:** The IP address bounded to the Cluster Name.
   - **Subnet mask:** The network to which this machine belongs.
   - **Quorum drive:** The drive that will contain the cluster information.

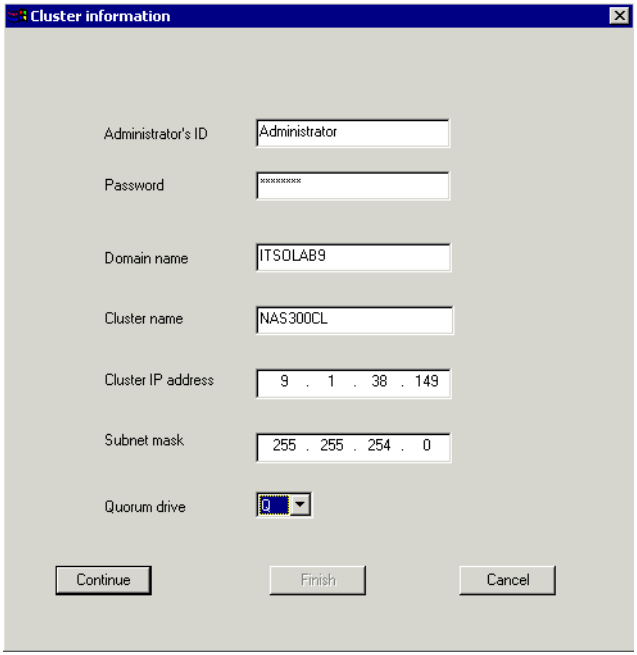

Figure 6-56 Cluster configuration

4. On the confirmation window, select **Yes**.
5. Once the configuration and copying of the necessary files are completed, the *Cluster Information* window will re-appear.
6. Click **Finish**.

The cluster administration utility will automatically start, showing the first node with its groups and the resources, with the corresponding information you supplied earlier (Figure 6-57).

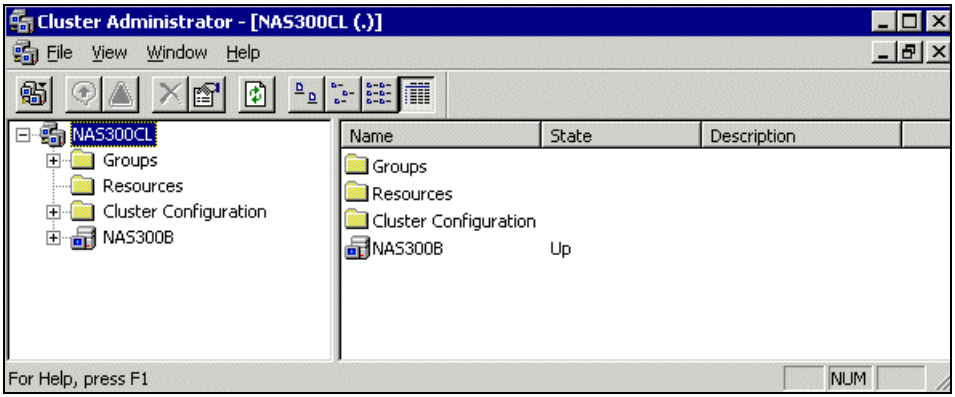

Figure 6-57 Cluster Administration with first node

7. After Node B restarts, power on the first node (node A) and connect to it.

Note: Both nodes should be powered on at this time.

At this point in the installation process you have already configured one machine (Node B) with Microsoft's Active Directory, a DNS server, and a Microsoft Cluster.

Node A NAS Setup Navigator should continue where you had stopped. Follow the screens to continue with Active Directory setup on node A.

If no DNS server existed and you were required to create one on Node B back in step 8, you will need to also create one now on node A.

8. If required, create a DNS server on node A. Do this by **Start** —> **Settings** —> **Control Panel** —> **Add/Remove Programs** —> **Add/Remove Windows Components** —> **Networking Services**. Click the box to check Networking Services and then click the **Details** button. The Subcomponents of Networking Services box comes up. Click the box next to Domain Name System, then click **OK.**
9. Reboot Node A when required.

10. Logon to Node A. The navigator will step you through joining the cluster you had created on node B. This can be seen in Figure 6-58. The same Cluster Configuration Wizard will come up. This time the computer will be the **Joining Node**. Supply the name or IP address of the first node in the cluster.
11. When Node A has successfully joined the cluster, you should use the cluster administrator to verify the resources.

Once the configuration and copying of the necessary files are completed, Cluster Administration should automatically start. The second node should now appear as part of the cluster.

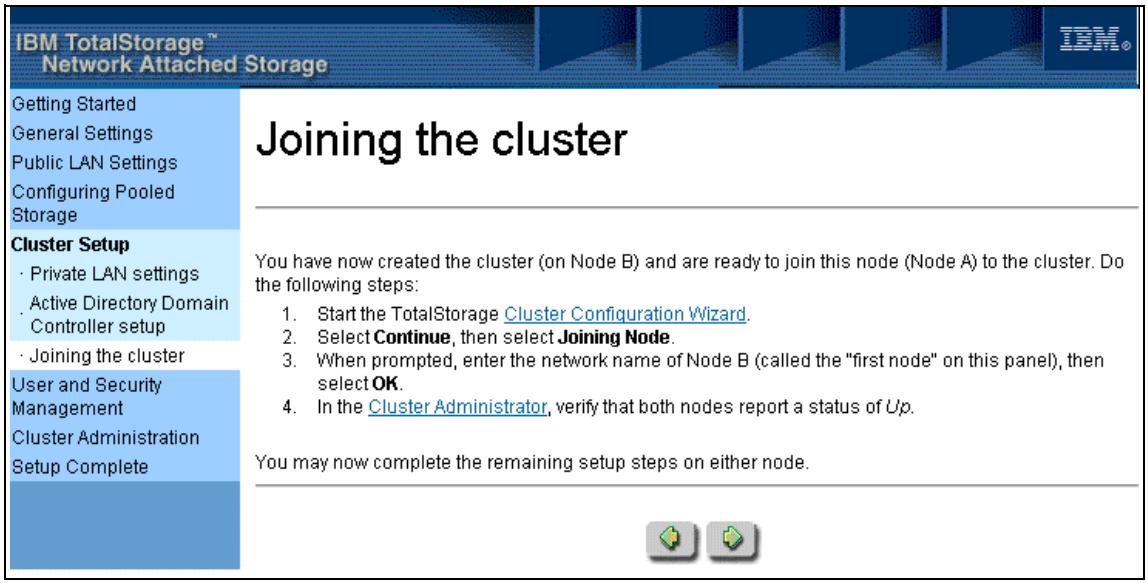

Figure 6-58 NAS Setup Navigator — Joining the cluster

12. For NFS and Windows file sharing, creating clustered shares and user administration, refer to Chapter 10, “Cross platform storage” on page 353.

6.2 Defining arrays and logical drives on the NAS 300

Just like the NAS 200, the NAS 300 system uses a disk utility to manage its storage. Storage in NAS 300 is fibre attached, and the utility to manage it is called the IBM FAStT Storage Manager. As already mentioned, it is possible to manage it either via a local console or remotely via browser or through Windows Terminal Services Client. Here are the steps to perform tasks required on your storage.

Creating drive arrays

Whichever method of connection you choose, it will give you access to the IBM NAS Admin utility. Inside you will see several management tools, among them the IBM FAStT Storage Manager:

1. Double-click the **IBM NAS Admin** utility and click the **+** sign beside Storage and NAS Utilities (Figure 6-59).

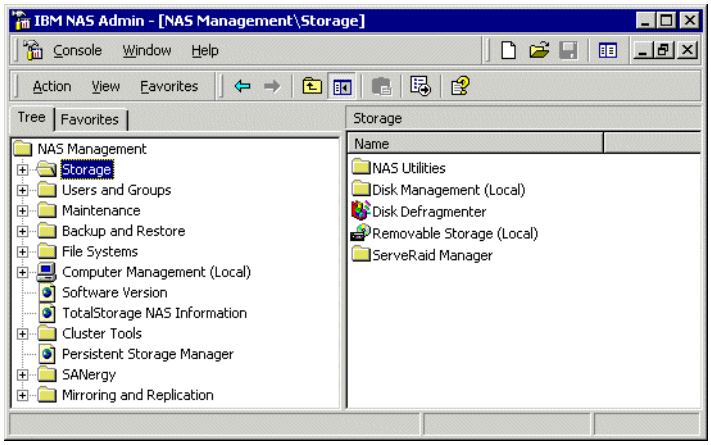

Figure 6-59 IBM NAS Admin Utility

2. Select **IBM FAStT Storage Manager Client**.

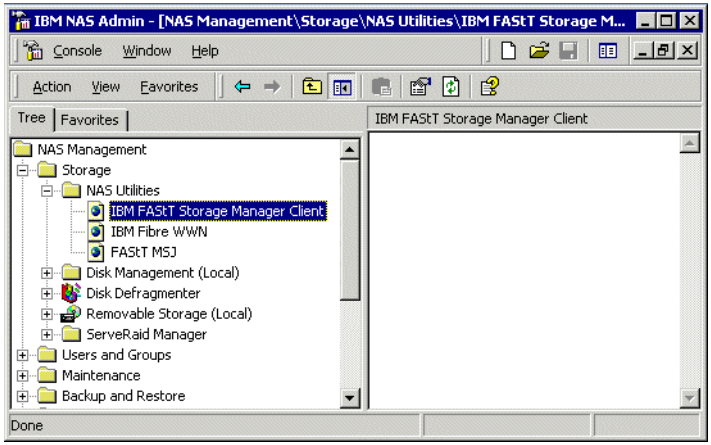

Figure 6-60 Starting IBM FAStT Storage Manager Client

3. This will start the IBM FAStT Storage Manager Client utility shown in Figure 6-61.

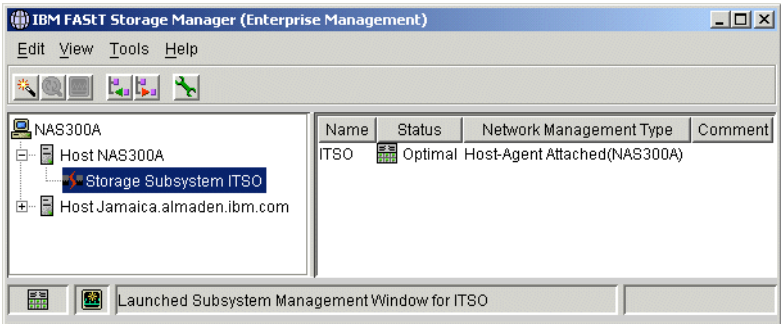

Figure 6-61 IBM FAStT Storage Manager

4. Right-click the Storage Subsystem and select **Manage**. (Figure 6-62).

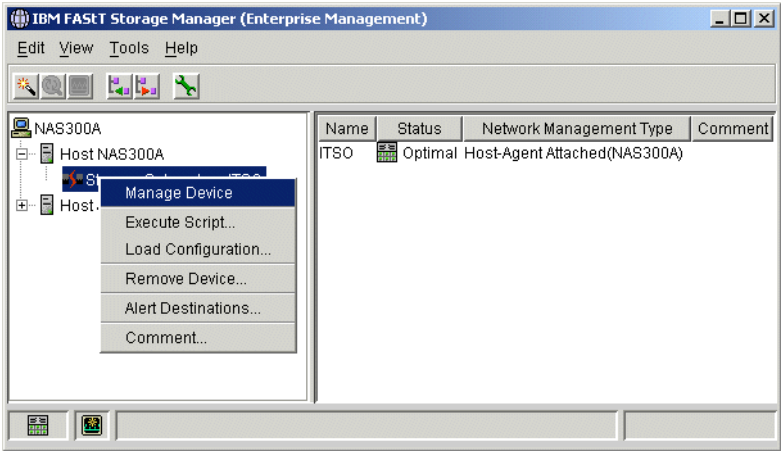

Figure 6-62 Starting Subsystem Management

5. The Logical/Physical View of the Subsystems Management window opens up. It is split into two panes (Figure 6-63). In the left pane you see the logical view (arrays, RAID level, logical drives, free capacity, unconfigured capacity). The right pane shows the physical view (drive, controller enclosure, used disks, hot-spare, and so on).

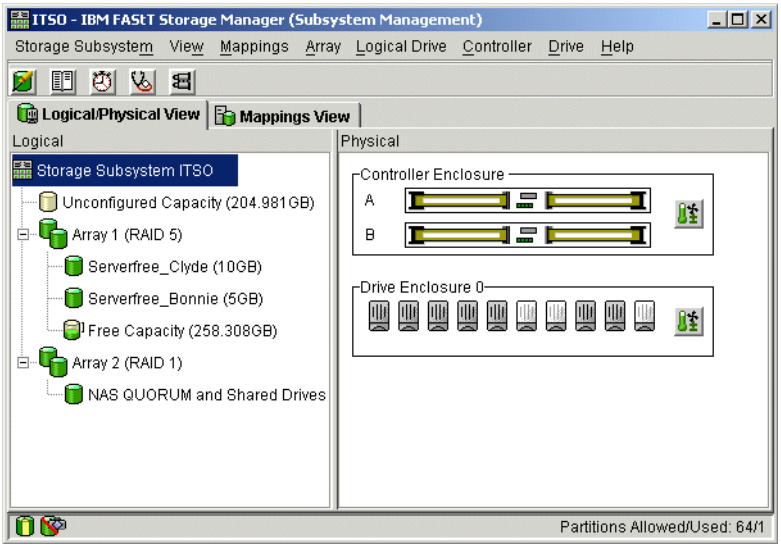

Figure 6-63 Subsystems Management

6. In a larger storage environment with several storage expansion units in a rack sometimes it is difficult do identify drives belonging to particular array. Here the Locate function can be useful. To test it, right-click the array in question and select **Locate**, as shown in Figure 6-64. All physical drives belonging to this array will have the indicator lights flashing.

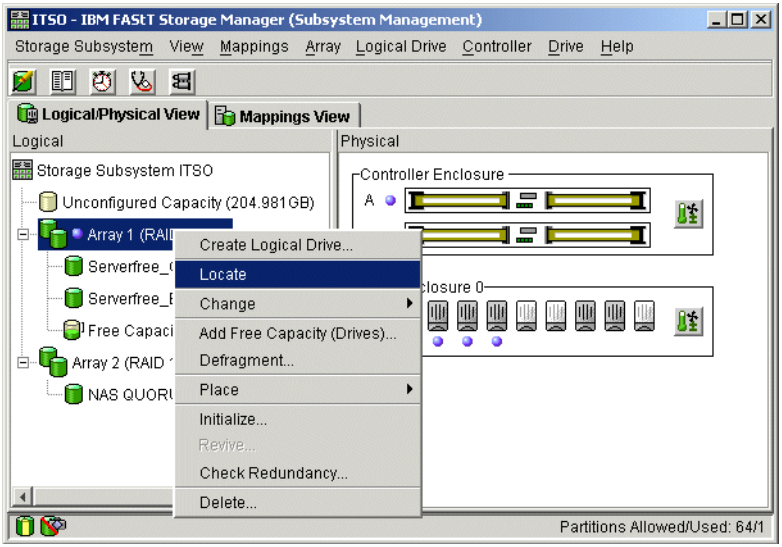

Figure 6-64 Locating Drives

7. To stop the indicator lights, simply click **OK** in the window that pops up (Figure 6-65).

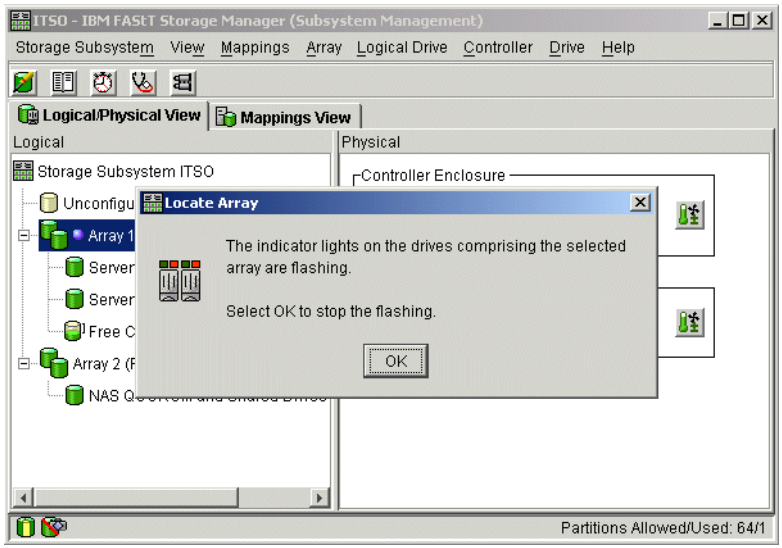

Figure 6-65 Stop the indicator lights

8. To create new storage units, right-click the **Unconfigured Capacity** (available drives will be shown with a purple bullet underneath) and choose **Create Logical Drive**, as shown in Figure 6-66.

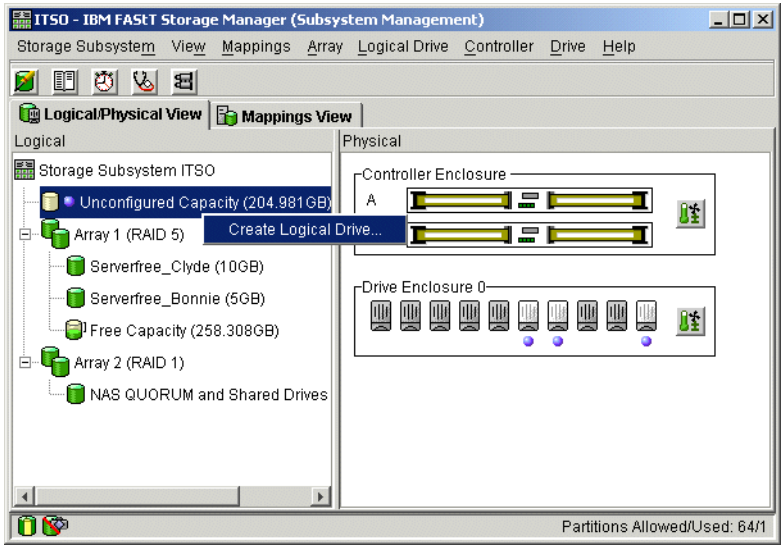

Figure 6-66 Create new storage units

9. If Storage Partitioning has been enabled before, a window similar to the one shown in Figure 6-67 will pop up informing you of the default host type. If the host system you are planning to attach to this new storage unit is of different type, you can select it here from the menu.

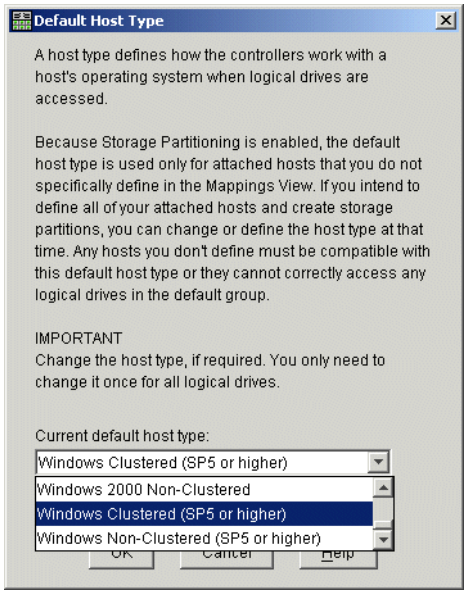

Figure 6-67 Selecting Host Type

10. Now the Create Logical Drive Wizard starts (Figure 6-67). If there are existing arrays, you will be asked whether you want to create logical drive in the existing array or from the unconfigured capacity (create new array).

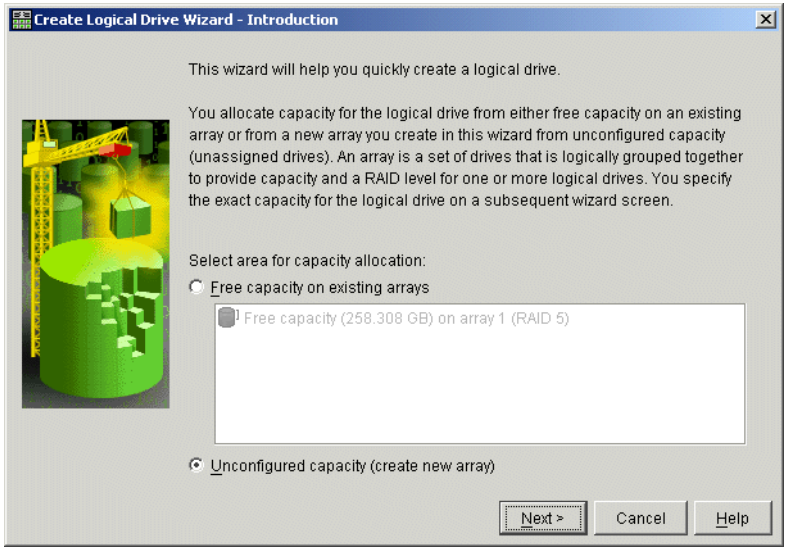

Figure 6-68 Create Logical Drive Wizard

11. On the next screen, under the RAID level selection, choose the RAID level required by clicking the pull-down arrow. Under “Drive selection choices”, select either **Automatic** or **Manual** (Figure 6-69).

12. If you chose Manual, you can define specific drives you want to be members of the array. For multiple drives, press the Ctrl key while selecting the drives. Finally, click **Apply**.

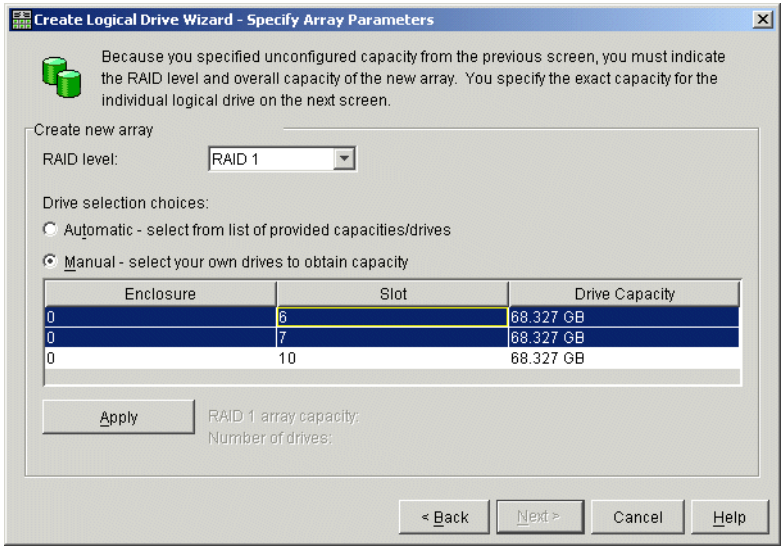

Figure 6-69 Array parameters

13. Next, define the Logical Drive inside the array you just created. For “New logical drive capacity”, type in the size you need and specify it in GB or MB. The default value shown is the maximum capacity of the array.

    The name should be set to something meaningful. If you select **Customize settings** you will be given a chance to define additional settings for this logical drive (Figure 6-70).

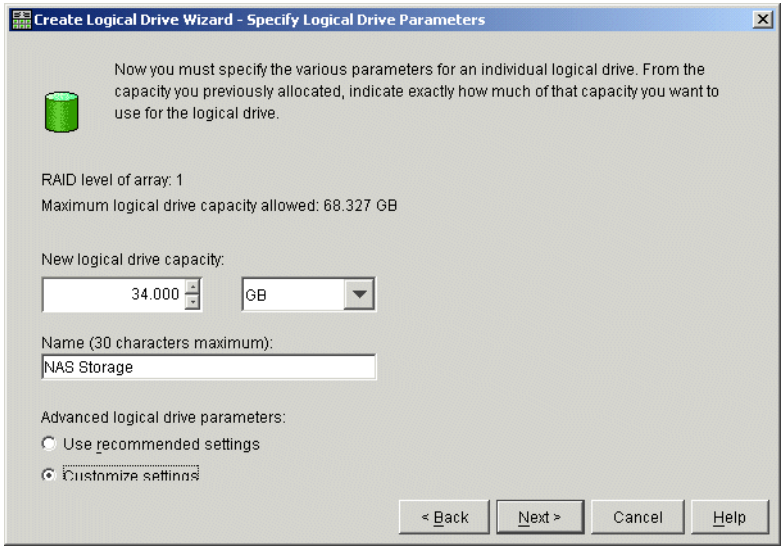

Figure 6-70 Logical Drive Parameters

14. For Advanced settings (Figure 6-71), it is possible to define:

    a. Logical Drive I/O characteristics (optimized for **File system**, **Database**, **Multimedia** or **Custom** — for this option you can customize “Cache read ahead multiplier” and “Segment size”). Since NAS is a file sharing machine, you can leave this setting at default, unless you plan to use it for other type of usage.

       Under “Segment size”, select the optimum value depending on the application that will use the storage on the NAS system. Normally, this could be left at default, which is based on expected usage.

    b. For preferred controller ownership of this logical drive, select either: **Slot A** or **B**.

Tip: To obtain the optimum performance, it is advised that you distribute the ownership of the arrays to both controllers. Controllers are configured for failover by default, so you don’t need to be concerned about a single controller failure.

    c. For Logical Drive to LUN mapping you can choose between **Automatic** or **Map later with Storage Partitioning**.

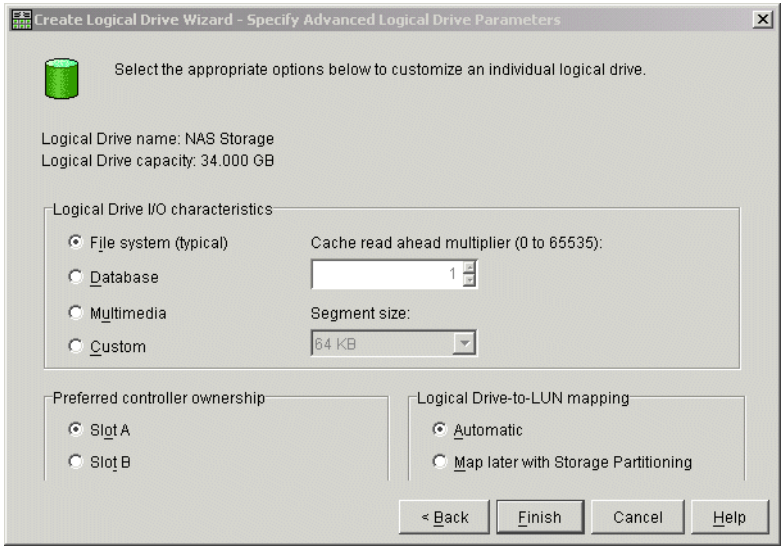

Figure 6-71 Logical Drive Advanced Parameters

15. The logical drive will be created with selected parameters and you will be asked if you wish to create another logical drive (Figure 6-72).

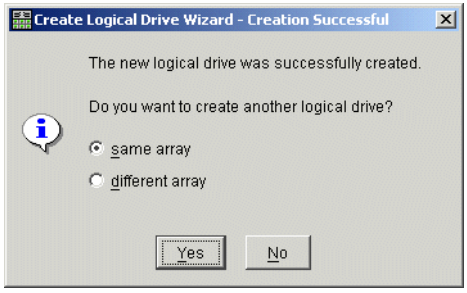

Figure 6-72 Logical Drive created

16. Again you will be notified how to change the Logical Drive to LUN mappings (Figure 6-73).

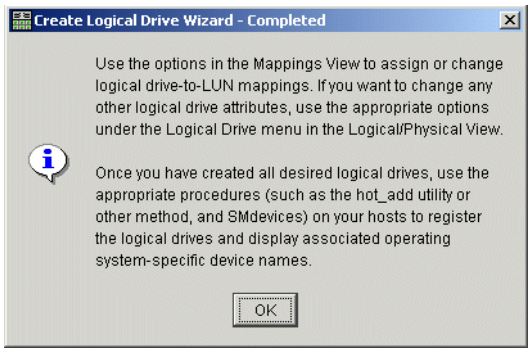

Figure 6-73 Mappings Notice

17. You should now see the array that you just created on the logical view pane, as shown in Figure 6-74. Until the new logical drive has been initialized, it cannot be accessed (shown with a clock symbol in the logical drive icon). Length of the initialization procedure depends on the size of the drive.

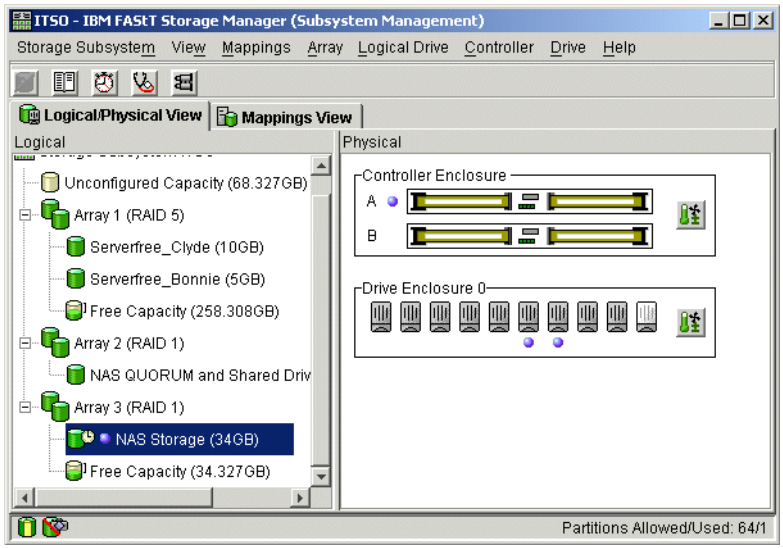

Figure 6-74 Newly created Array and Logical Drive

18. You can check the initialization progress by right-clicking the drive and selecting **Properties**. A window as shown in Figure 6-75 will open, giving you all details about the created drive, with the Initialization Progress bar on the bottom.

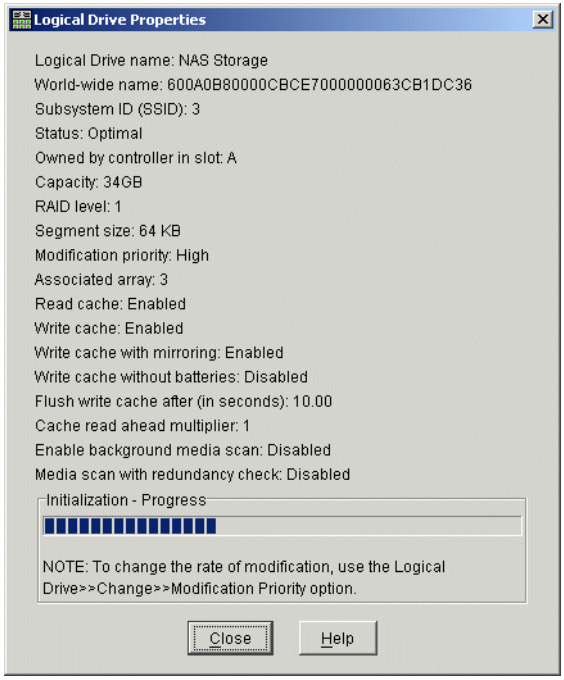

Figure 6-75 Logical Drive Initialization Progress

19. Click **OK** on the Logical Drive Creation Successful window (Figure 6-76).

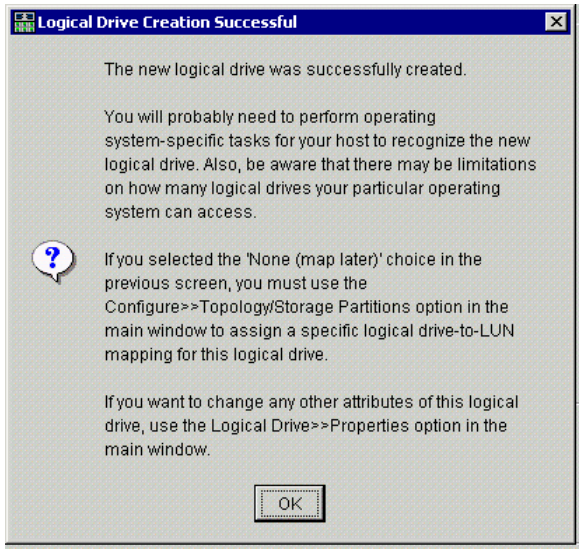

Figure 6-76 Logical drive creation successful

Defining newly created drives

Now that the logical drives are completely set up and ready, they need to be configured with the Windows 2000 disk management tool. Follow these steps to accomplish the task:

1. Open the IBM NAS Admin by clicking its icon on the desktop.
2. Click the plus (+) sign beside the Storage folder.
3. Select **Disk Management (Local)**. The disk management tool will open as shown in Figure 6-77.

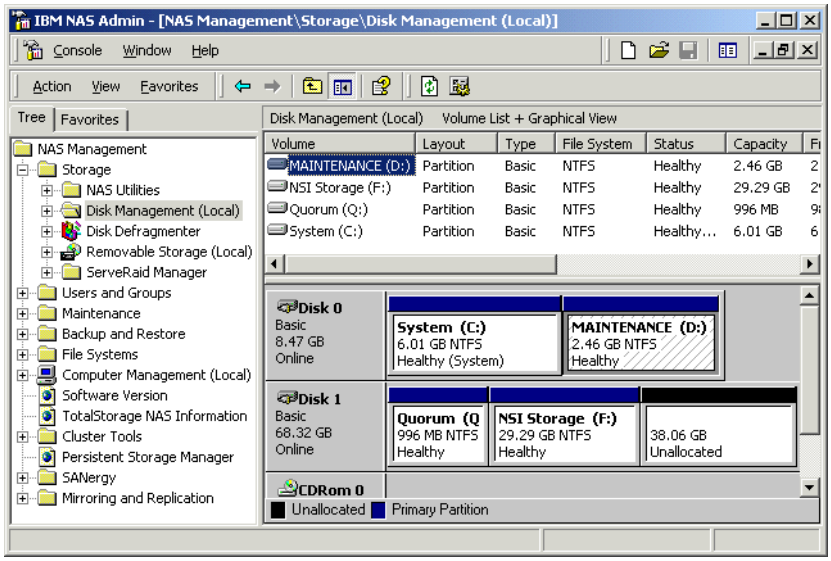

Figure 6-77 Disk management tool

4. On this menu, select **Action** and execute **Rescan Disks** as shown in Figure 6-78.

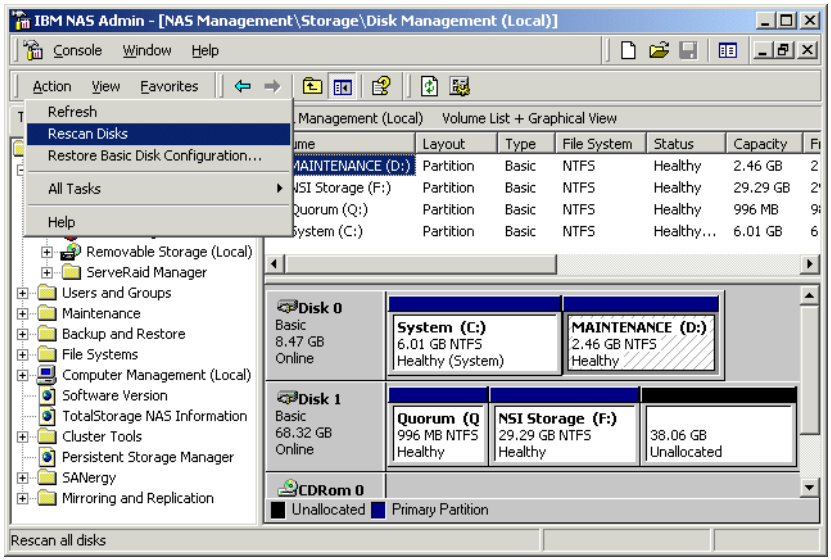

Figure 6-78 Staring Rescan Disks

5. A progress window will open, informing you about the scanning process (Figure 6-79).

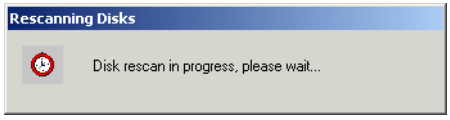

Figure 6-79 Rescanning process

6. Once the rescan is complete, the array that you just created should be shown on the list at the bottom right as unknown and unallocated (Figure 6-80). Before you can use it, the signature has to be written to it first.

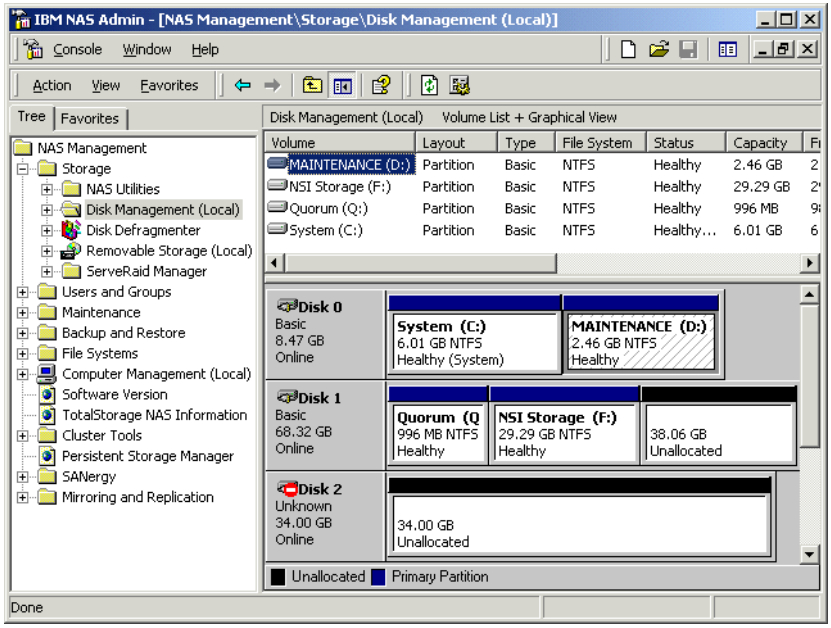

Figure 6-80 Unknown disk found

7. Right-click the unknown disk and select **Write Signature** (Figure 6-81).

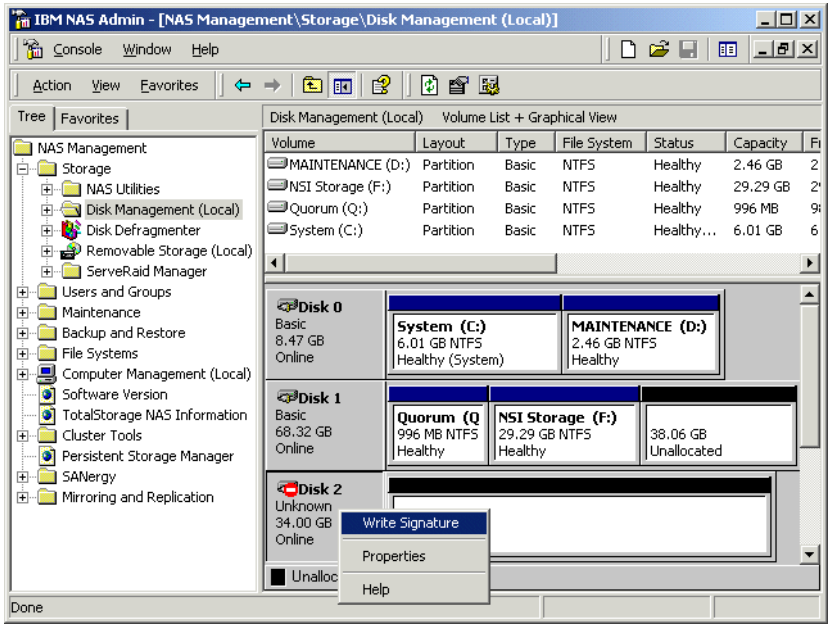

Figure 6-81 Writing a signature to the new disk

8. Click **OK** on the Write Signature window (Figure 6-82).

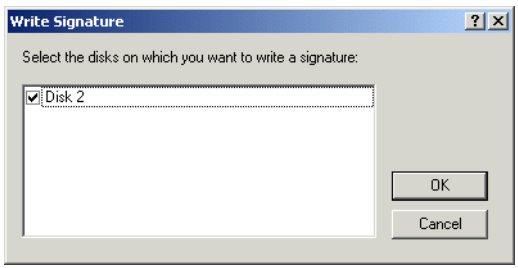

Figure 6-82 Confirming disk signature

9. Now the disk is ready for partitioning. To create a partition inside the new disk, right-click it and select **Create Partition** (Figure 6-83).

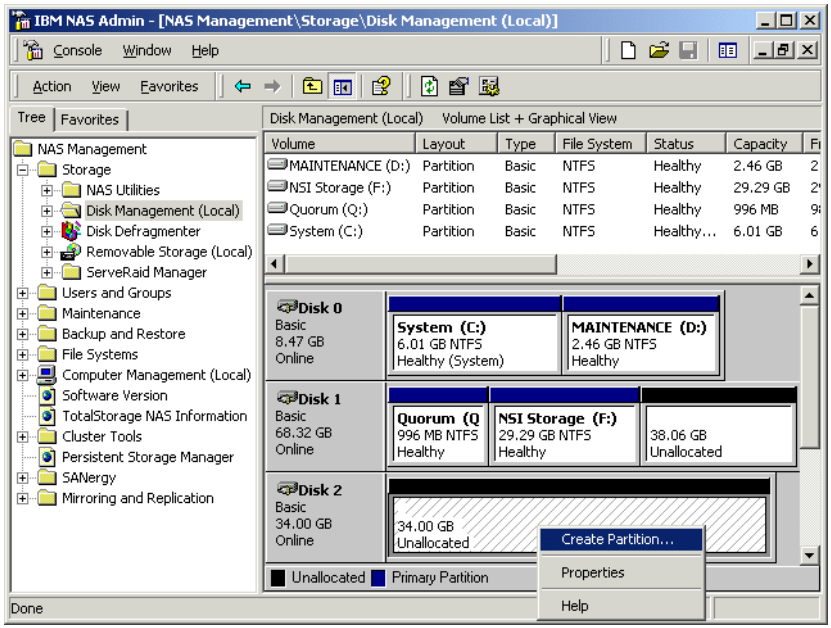

Figure 6-83 Creating a partition

10. The partition wizard will start. Click **Next** (Figure 6-84).

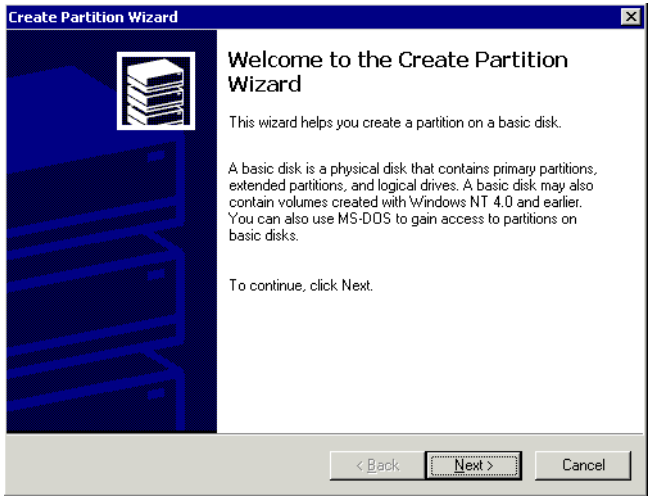

Figure 6-84 Create Partition Wizard

11. On the Partition Type window, select **Primary partition** (Figure 6-85)
12. Then click **Next**.

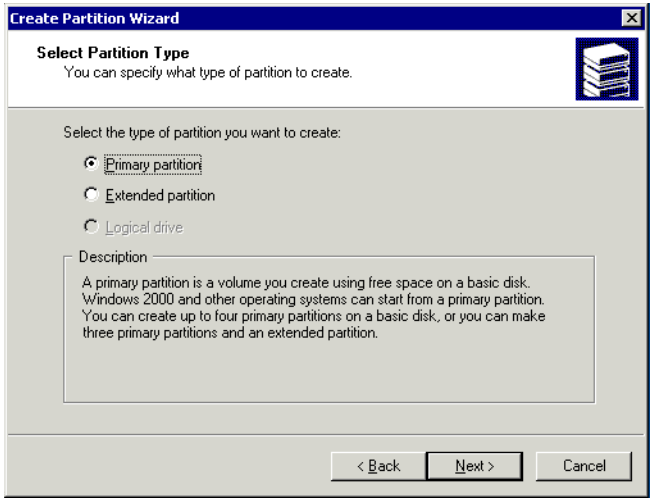

Figure 6-85 Select partition type

13. For “Specify Partition Size” you can set the partition size. We left it at the default, which should be the maximum (Figure 6-86).

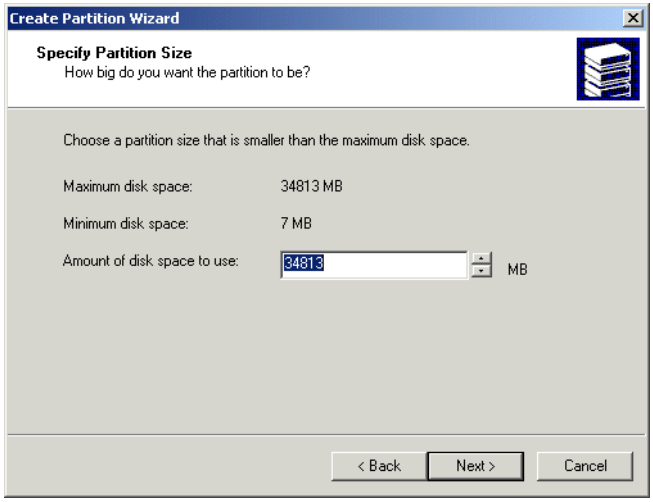

Figure 6-86 Specify partition size

Important: Every drive on a cluster must be single disk single partition. The Microsoft Cluster Server does not support multiple partitions/logical drives on a single disk.

14. On the Assign Drive Letter or Path window, assign a drive letter to this partition (Figure 6-87). Then click **Next**.

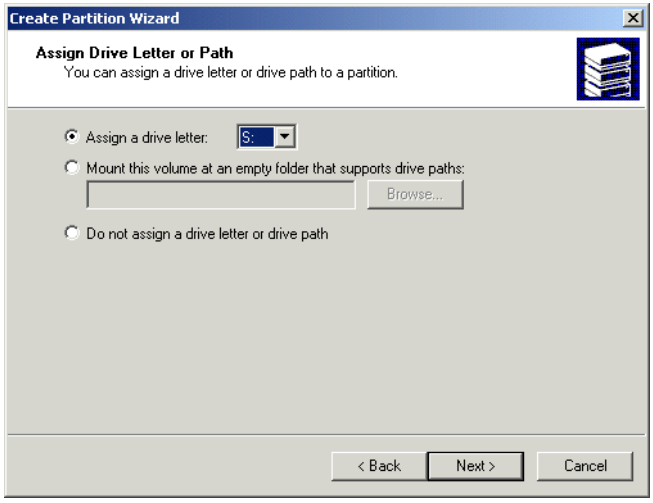

Figure 6-87 Assign drive letter

15. On the Format Partition window, select **Format this partition with the following settings** (Figure 6-88).

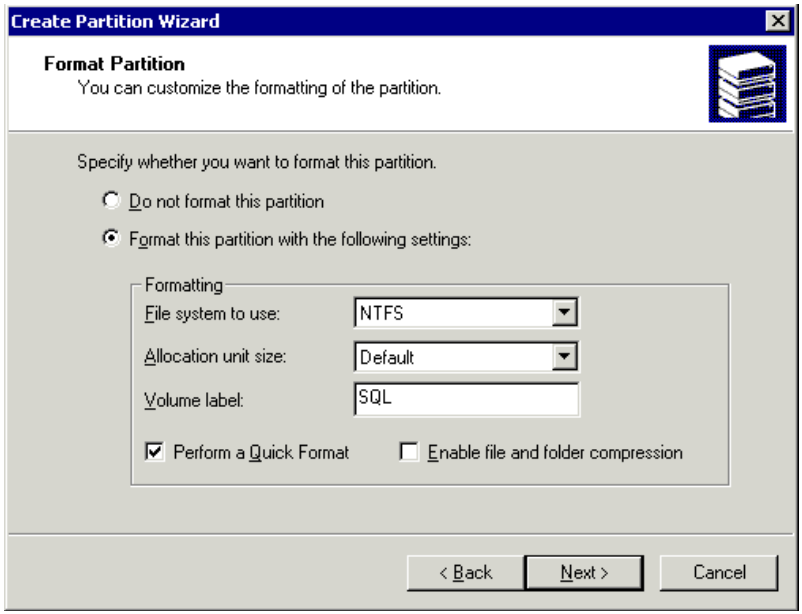

Figure 6-88 Format partition settings

For “File System to use:” make sure that NTFS is selected.

Important: Microsoft Cluster only supports NTFS partitions.

You can specify the allocation unit size or leave it at the default value.
For “Volume label”, type in the name you want to associate with this partition.
Click **Perform a Quick Format** (this will save you some time), then click **Next**.

Important: Do ***not*** select **Enable File and Folder Compression**, since this is not supported by Microsoft Cluster.

16. On the final window, just click **Finish** (Figure 6-89).

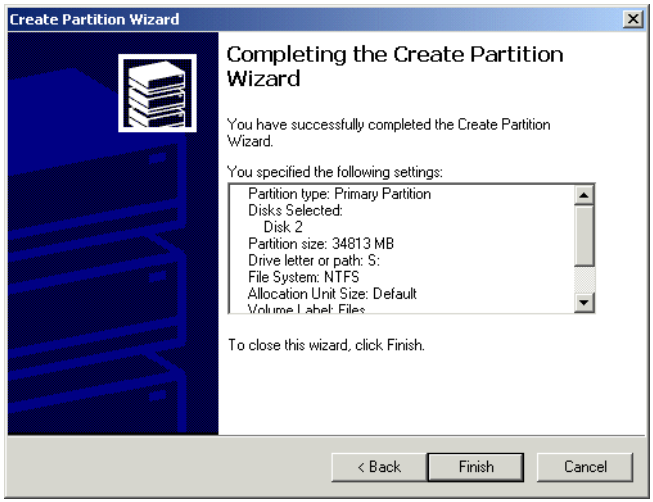

Figure 6-89 Completing the Partition Wizard

17. Your partition is now ready for use (Figure 6-90).

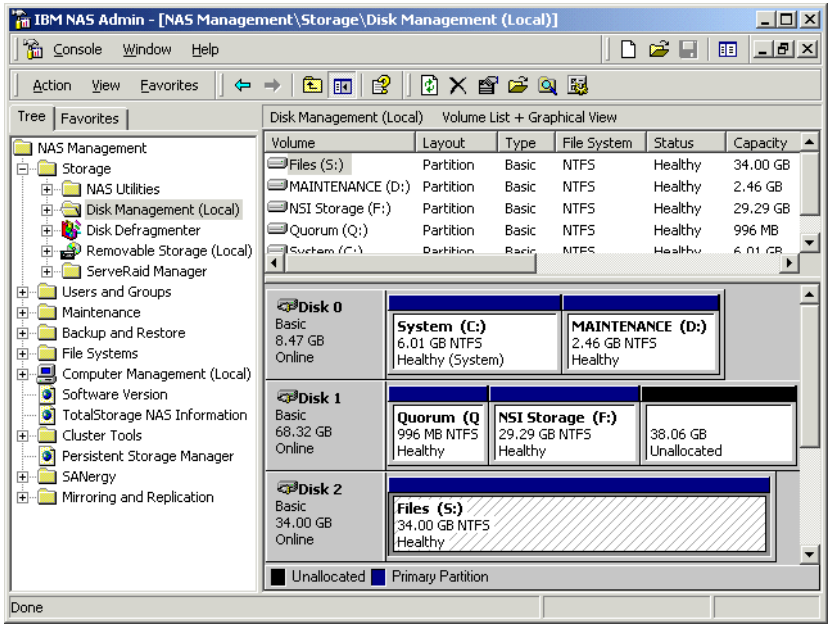

Figure 6-90 Partition ready

6.3 Setting up Microsoft Cluster Server (MSCS)

The IBM TotalStorage NAS 300 enables you to cluster to maintain high availability. Clustering the NAS 300 is similar to clustering in a client/server environment.

This section guides you through the configuration and setup of a NAS 300 cluster using the Microsoft Cluster Server (MSCS) software, which is included with the NAS 300. Our environment is depicted in Figure 6-91.

Our goal in this section is to clarify a few steps. By following our procedures carefully, you should be able to implement the sometimes dreadfully complicated cluster service, relatively easily and painlessly. Pleasant clustering!

Prerequisites and checklist

Make sure that the following information is ready for the cluster:

- **Virtual Name:** This is the hostname the clients will use to address the NAS 300. This hostname can bind and failover to any of the engines/nodes on the appliance.
- **Virtual IP address:** This is the address bound to the Virtual name.
- **Drive letter of the Quorum drive:** The Quorum drive is where the cluster information is kept. It is recommended to use drive G (which is the default on the NAS 300).

Note: The MSCS needs a Windows Domain or Active Directory controller. You must integrate your clustered system in such a security context. If your environment does not provide this, use the *domainlet* function to create a small Active Directory structure.

Figure 6-91 shows our environment and how we set up the sample cluster.

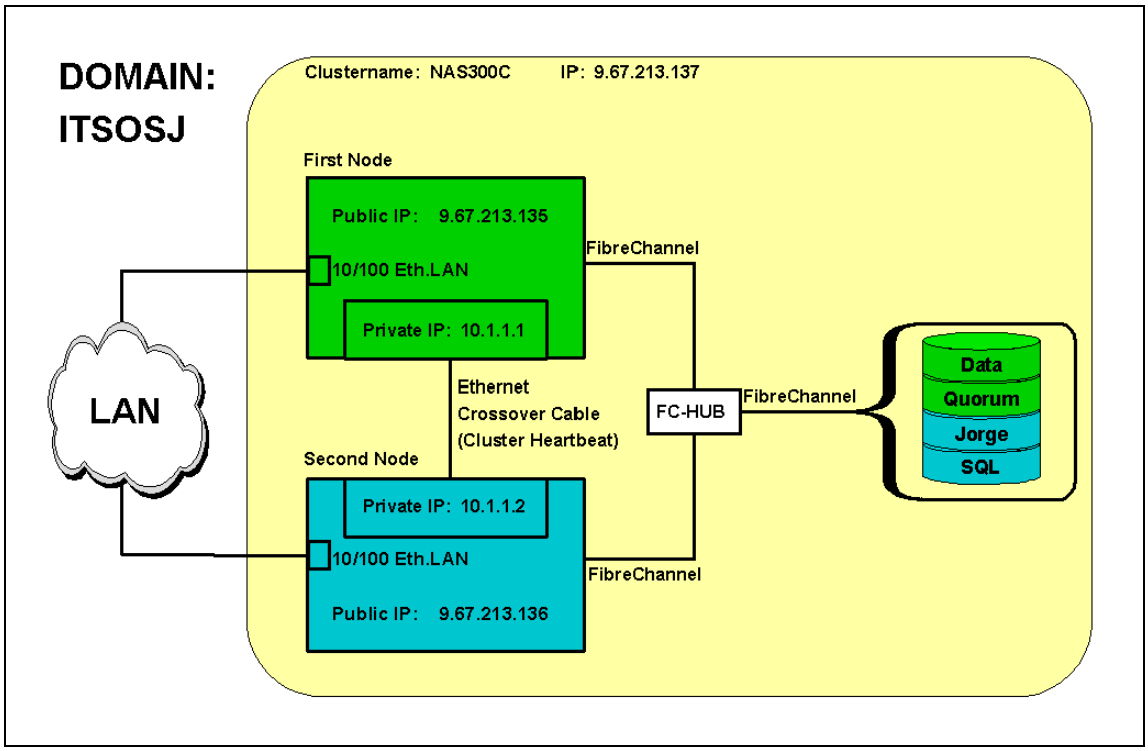

Figure 6-91 Cluster setup diagram

Before the actual setup of the cluster, you need to perform the following on each NAS engine/node:

- Assign sticky drive letters on the cluster drives.
- Configure network adapter settings.

Assigning sticky drive letters

You need to assign **sticky** drive letters to the Quorum and data drives, so that each drive will be recognized exactly the same on each node. In other words, the drive mappings on both nodes need to be set up exactly the same way. Also, you need to make sure that they are configured as *Basic*. These are requirements of the Microsoft Cluster Server. To achieve this, do the following on each node (one at a time):

1. Open a Terminal Services Client session to the first engine (either using Internet Explorer or Terminal Services Client Program as discussed in 3.3.2, “Getting access with Internet Explorer” on page 53).

2. Once logged on, open the IBM NAS Admin utility and select **Storage**.
3. Choose **Disk Management (Local)**.
4. Right-click the drive and select **Change Drive Letter and Path** (Figure 6-92).

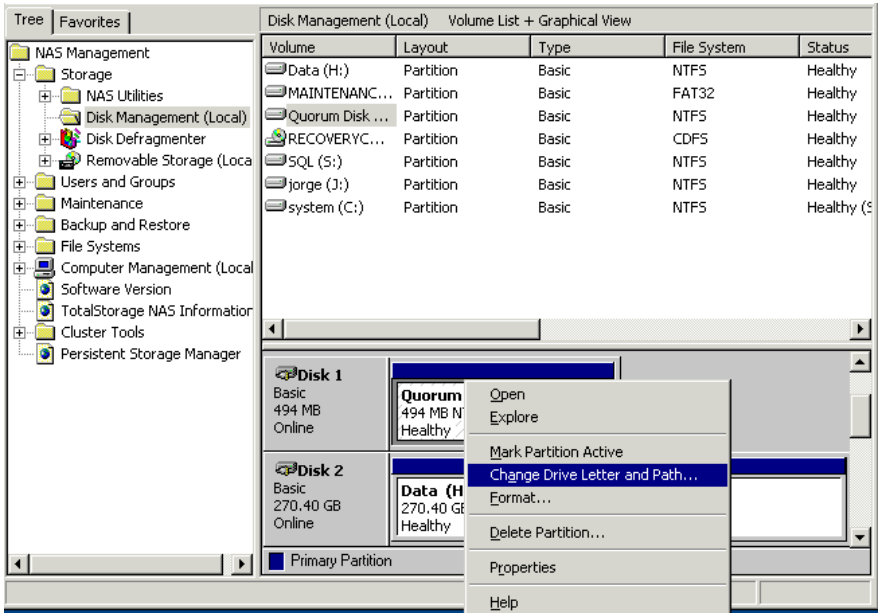

Figure 6-92 Change drive letter

5. On the Change Drive Letter window, select **Edit** (Figure 6-93).

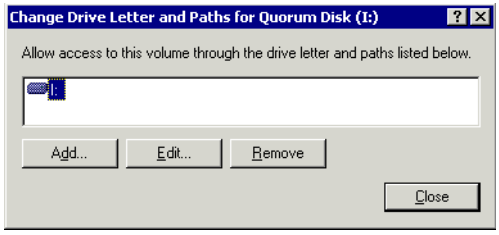

Figure 6-93 Change drive letter

6. On the Edit Drive Letter or Path, click the pull-down arrow and select the drive letter you want (again, it is recommended to use G: for the Quorum). See Figure 6-94 for details.
7. Then click **OK**.

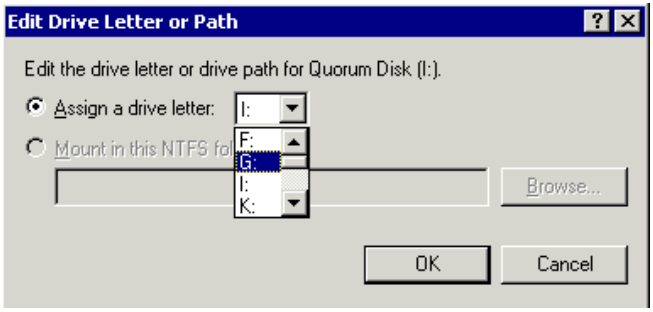

Figure 6-94 Select drive G: here for the Quorum

8. When asked to confirm, click **Yes**.

The Quorum drive should now have the letter you just assigned. By doing this manual assignment, Windows will not change the drive letter even if you reboot the engine (otherwise, it could happen that Windows changes the drive letter because it is assigned dynamically). Thus the **sticky** drive letter is fixed, as shown in Figure 6-95.

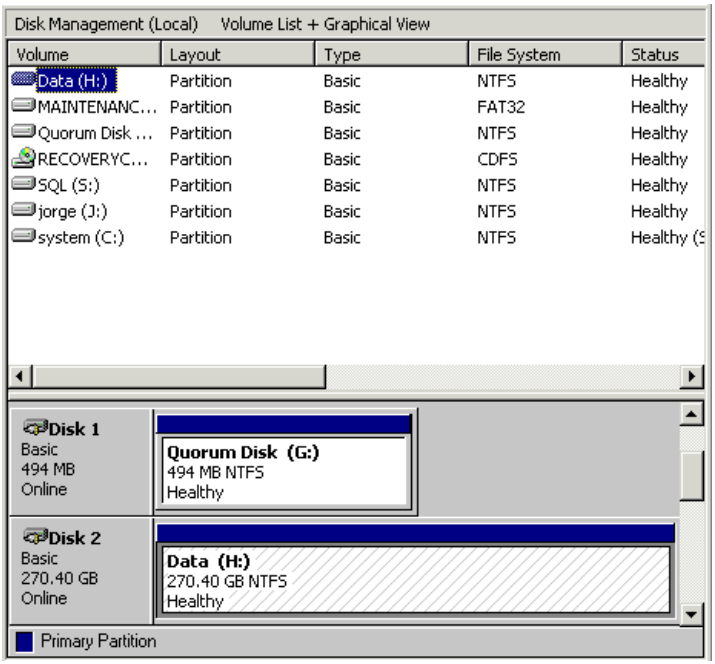

Figure 6-95 Sticky drive letter G: assigned to Quorum

Perform these steps for the remaining drives that will be clustered.

Note: It is recommended to reserve drive letter F: for Persistent Storage Manager images, and use E: for the CD-ROM drive. For more information on PSM, refer to Chapter 11, “Backup considerations” on page 431.

Now perform the same steps for the second node.

After you are done, and the disks have **sticky** drive letters assigned, they are recognized as *Basic*, and they are exactly the same on both nodes. It is time to configure the network on both machines. We now describe the steps to perform.

Cluster network configuration

Now that we have considered all prerequisites for the cluster setup, we can start with the first node.

First node preparation

These are the steps for preparing the first node:

1. Get connected to your NAS system via the Web or Terminal Server interface.
2. Once logged on, right-click **My Network Places** (Figure 6-96).
3. Go to **Properties**.

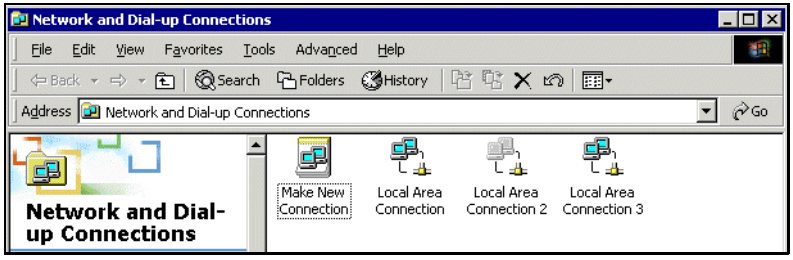

Figure 6-96 Network configuration window

4. On the Network and Dial-up Connections, right-click **Local Area Connection**.
5. Then select **Rename** and type in `private`. This creates the private connection for the cluster.

Note: We recommend that you rename the heartbeat and production networks on the nodes, but this is not a cluster requirement. However, it helps for easy identification to use a more descriptive name.

6. Right-click **Private**.
7. Choose **Properties**. Make sure that *IBM 10/100 NetFinity Fault Tolerant Adapter* is displayed under *Connect using*. Our example is shown in Figure 6-97.

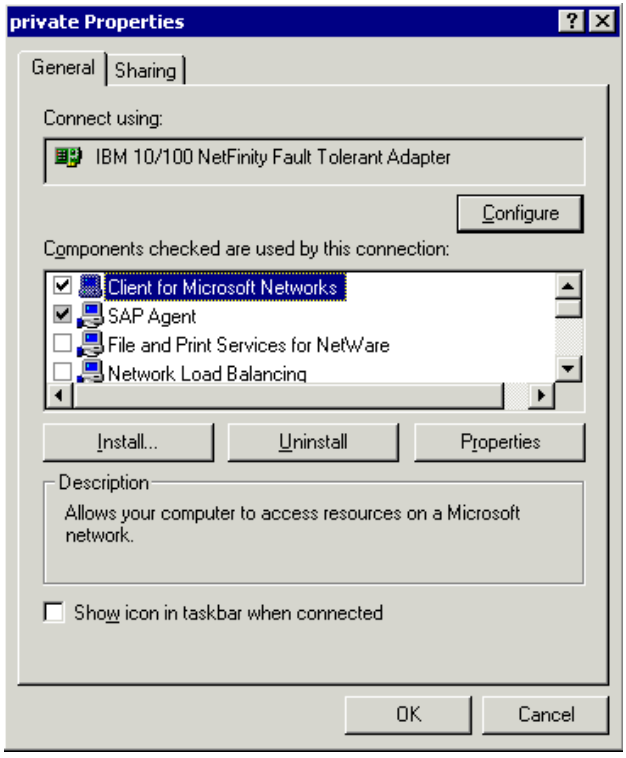

Figure 6-97 Private network properties

8. Highlight **Internet Protocol (TCP/IP)**, then click **Properties**. Now you will see the TCP/IP protocol configuration screen (Figure 6-98).

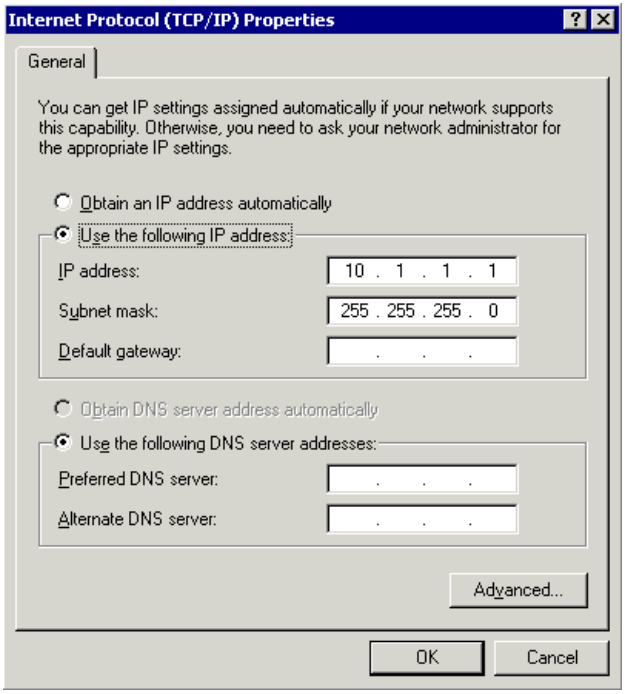

Figure 6-98 IP properties for on-board adapter

Because this network will be used only for cluster communication (also called heartbeat), it needs to have a different IP network. By default, the first engine's on-board adapter has an IP address of 10.1.1.1 and subnet mask of 255.255.255.0. We highly recommend that you use the default, but you can redefine the addresses, as long as you use a different network from the public/production network.

1. Click **Advance** and you will see a screen similar to Figure 6-99.

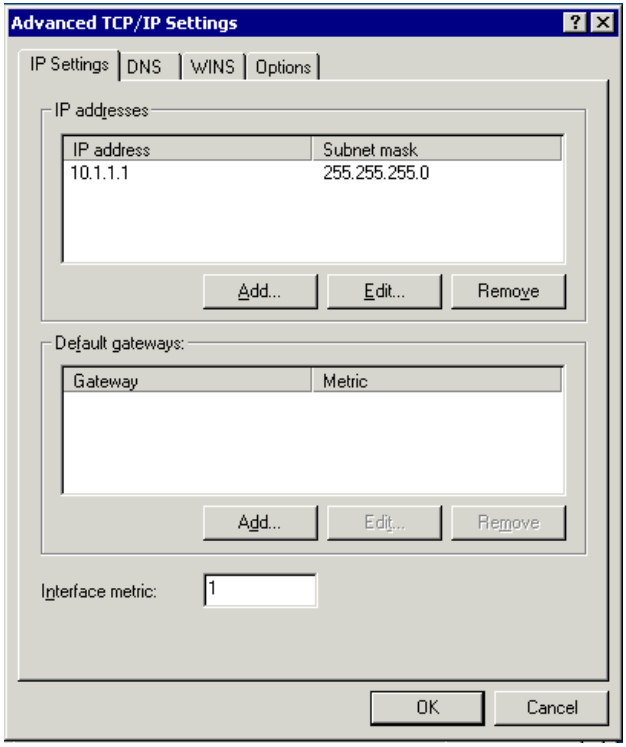

Figure 6-99 Advanced TCP/IP settings

2. Select the **WINS** tab.
3. At the bottom, click **Disable NetBIOS over TCP/IP** (Figure 6-100).
4. Then click **OK**.

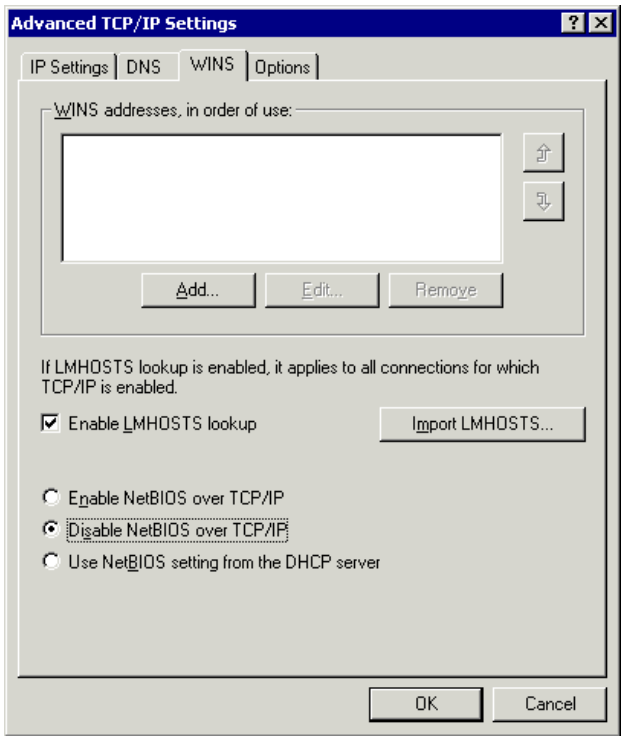

Figure 6-100 WINS TCP/IP settings

5. When prompted with *This connection has an empty primary WINS address. Do you want to continue?* (Figure 6-101), choose **Yes**.

   In case you are using WINS for name resolution, you need to specify the WINS server name.

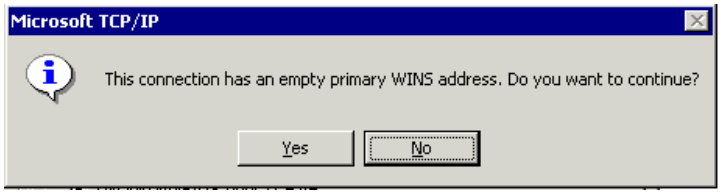

Figure 6-101 WINS warning

6. On the Internet Protocol (TCP/IP) Properties window, click **OK** (Figure 6-98).
7. On the Private Properties window (Figure 6-97), click **Configure** to set up the network card. You will see the window in Figure 6-102.

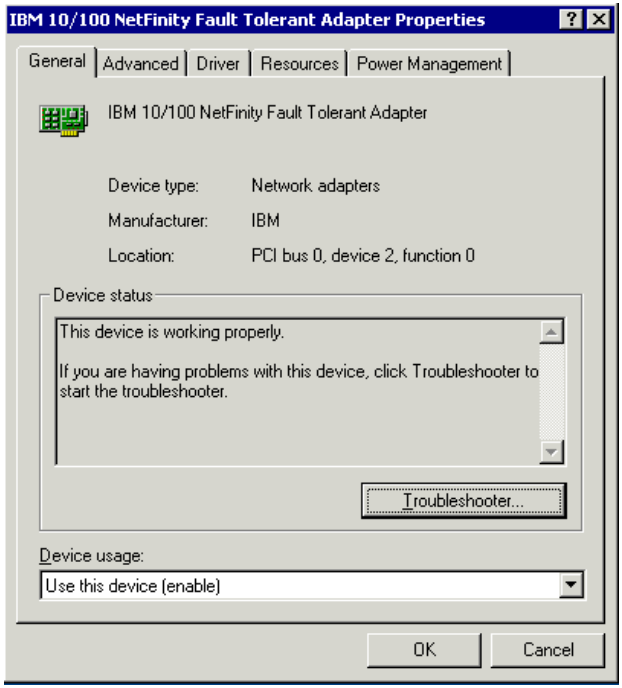

Figure 6-102 Private network adapter properties

8. Click the **Advanced** tab (Figure 6-102) and you will get the window shown in Figure 6-103.

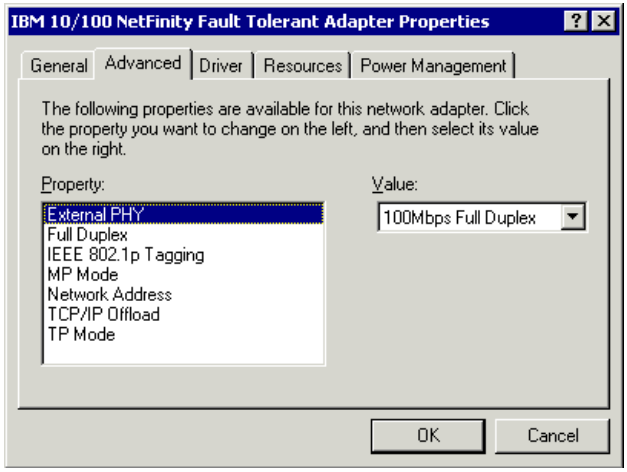

Figure 6-103 On-board adapter advanced properties

9. Here, make sure that the following settings are correct. They should be set up as shown in Table 6-4.

Table 6-4 Advanced settings for the private network

Property	Value
External PHY	100Mbps Full Duplex
Full Duplex	UTP — Full Duplex
TP Mode	On (Default)

10. Once you have check those settings, click **OK** *and again* **OK**.

11. Highlight **Private**, then on the menu, select **Advanced** *and again* **Advanced Settings** (Figure 6-104).

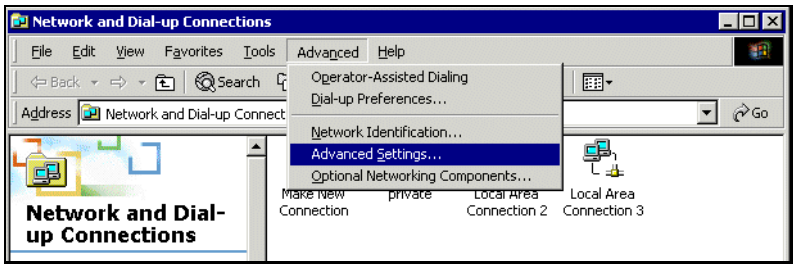

Figure 6-104 Configuring on-board adapter advanced settings

12. Select **Adapters and Bindings** tab. Under **Connections** make sure that *private* is first in the list (which should be the default). If not, re-sort the list to make it first by clicking the up and down arrows at the right (Figure 6-105).

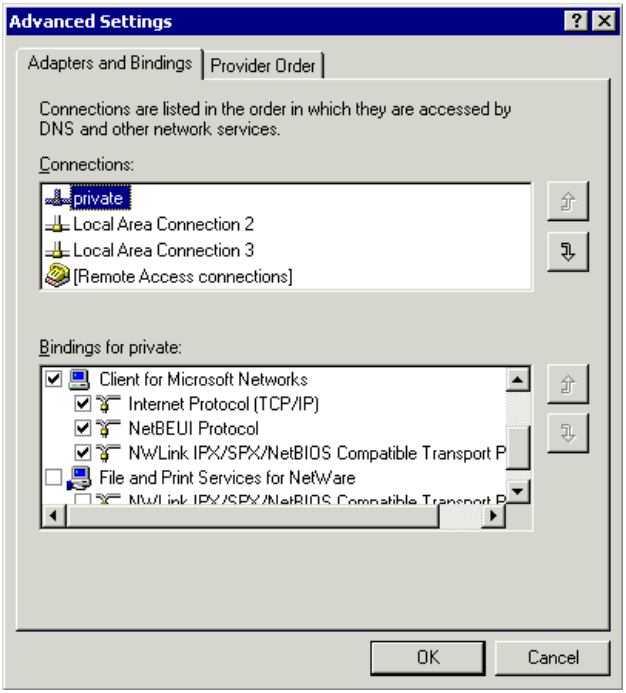

Figure 6-105 Network advanced settings

13. Click **OK**.
14. When prompted to reboot (this is necessary due to the changing of the bindings), select **Yes**.
15. Once the machine has restarted, connect again to the NAS system. On the Network and Dial-up Connections, right-click **Local Area Connection 3** (which is the IBM 10/100 Ethernet Server Adapter).
16. Then select **Rename** and type in `public`.

Note: If you plan to use a Gigabit network adapter for the public network of NAS 300, then perform the last steps for the Gigabit card also.

17. Configure the following IP settings on your *public* network adapter (IBM 10/100 Ethernet Server Adapter). You should get the information from your network administrator:
    - IP address
    - Subnet Mask
    - Default Gateway
    - Preferred DNS Server
    - DNS Suffix

After performing all the steps, the first node is now ready for cluster setup.

Second node preparation

These are the steps for preparing the second node:

1. Once again get connected to your NAS system and follow the steps described in Section , “First node preparation” on page 181.
2. Right-click **Local Area Connection**.
3. Choose **Properties**. Make sure that IBM 10/100 NetFinity Fault Tolerant Adapter is displayed under “Connect using”. Refer to Figure 6-97 on page 182.
4. Highlight **Internet Protocol** (TCP/IP), then click **Properties**.
5. The difference here is in selecting the right IP address. Because this network will be used for cluster communication (heartbeat) with the first node, it needs to be on the same IP network as the private adapter on the first node. By default, the second node’s on-board adapter has an IP address of 10.1.1.2 and subnet mask of 255.255.255.0. It is highly recommended that you use the default, but you can redefine the addresses accordingly.
6. Repeat all the other steps that you did on the first node also on the joining node.

Steps in cluster setup

Now that the two nodes have been prepared, the Microsoft Cluster Server is ready to be set up. Next we describe the steps to follow.

Note: The Microsoft Cluster Server (MSCS) setup on the NAS system is different compared to the original Microsoft implementation. IBM provides a NAS optimized setup routine.

First node setup

1. Get connected to your NAS system and log on (either locally or to the domain) and open the IBM NAS Admin utility (Figure 6-106).

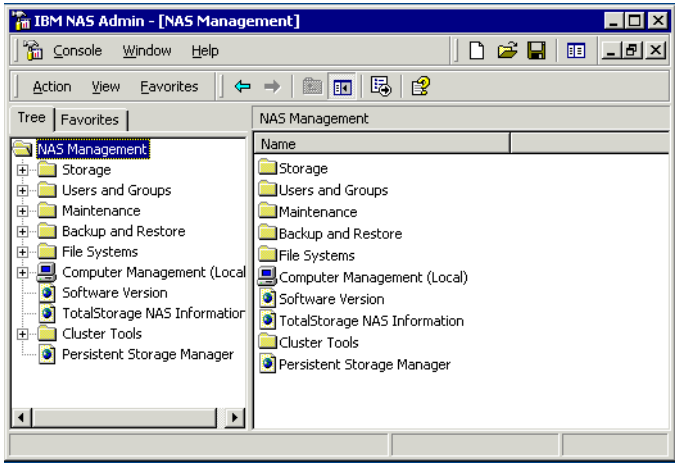

Figure 6-106 NAS Admin utility

2. Click **Cluster Tools**.
3. Select **Cluster Administration**.
4. On the *Total Storage Cluster Configuration Wizard*, select **Continue** (Figure 6-107).

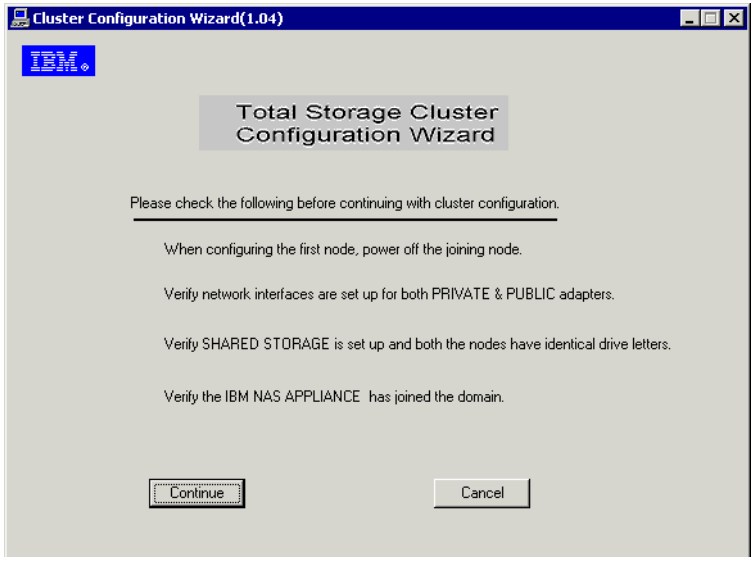

Figure 6-107 Total Storage Cluster Configuration Wizard

5. On the *Nodes* window, select **First Node** (Figure 6-108).

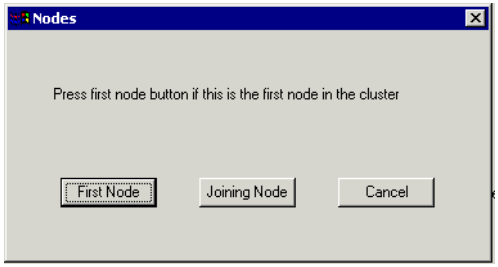

Figure 6-108 Node selection

6. On the *Cluster Information* window (Figure 6-109), fill in the required information.
   - **Administrator's ID:** The domain administrator account.
   - **Password**: The password for that domain administrator account.
   - **Domain name:** The domain which this machine is a member of
   - **Cluster name:** The name by which the clients will address the NAS 300.
   - **Cluster IP address:** The IP address bounded to the Cluster Name.
   - **Subnet mask:** The network to which this machine belongs.
   - **Quorum drive:** The drive that will contain the cluster information.

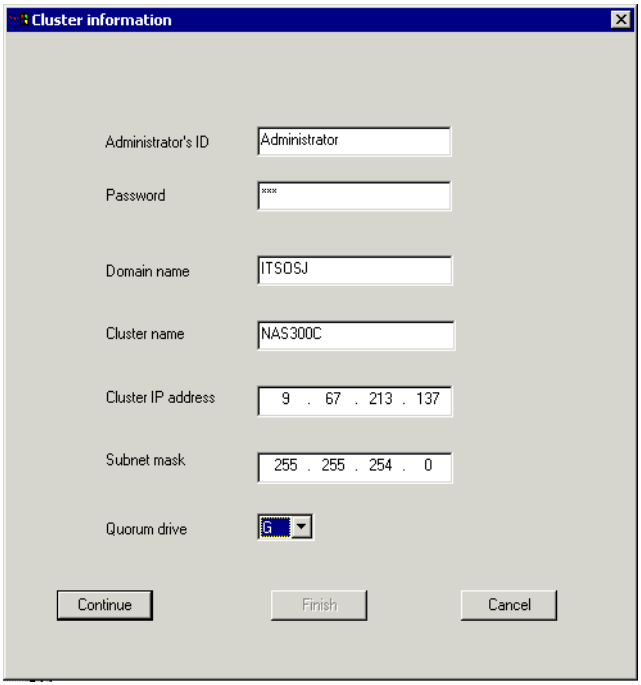

Figure 6-109 Cluster configuration

7. On the confirmation window, select **Yes**.
8. Once the configuration and copying of the necessary files are completed, the *Cluster Information* window will re-appear (Figure 6-110).
9. Click **Finish**.

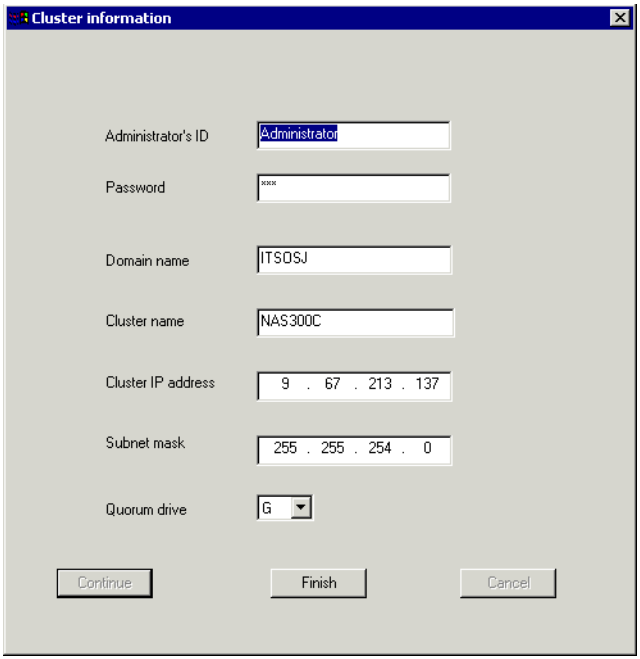

Figure 6-110 Cluster information

The cluster administration utility will automatically start, showing the first node with its groups and the resources, with the corresponding information you supplied earlier (Figure 6-111).

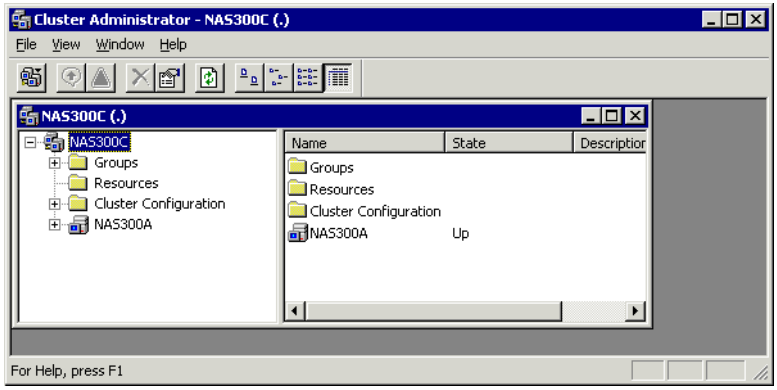

Figure 6-111 Cluster Administration with first node

Joining/second node

These are the steps for joining nodes:

1. Get connected to your NAS system and log on (either locally or to the domain) and open the IBM NAS Admin utility (Figure 6-106)

Important: If you are going to logon locally, make sure that the local account you will use is also an Admin account on the domain (with the same password), since this account will be authenticated on the domain when you join the second node to the cluster. Better yet, logon to the domain directly using the same Admin account you supplied for the first node in Figure 6-109 on page 192.

2. Click **Cluster Tools**.
3. Choose Cluster Administration.
4. On the *Total Storage Cluster Configuration Wizard* (refer to Figure 6-107 on page 190), select **Continue**.
5. On the *Nodes* window (refer to Figure 6-108 on page 191), select **Joining Node.**
6. On the First Node Information window, type the hostname or IP address of the first node (Figure 6-112).

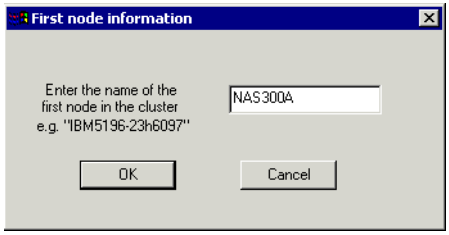

Figure 6-112 Second cluster node joining setup

Once the configuration and copying of the necessary files are completed, Cluster Administration should automatically start. The second node should now appear as part of the cluster, as shown in Figure 6-113.

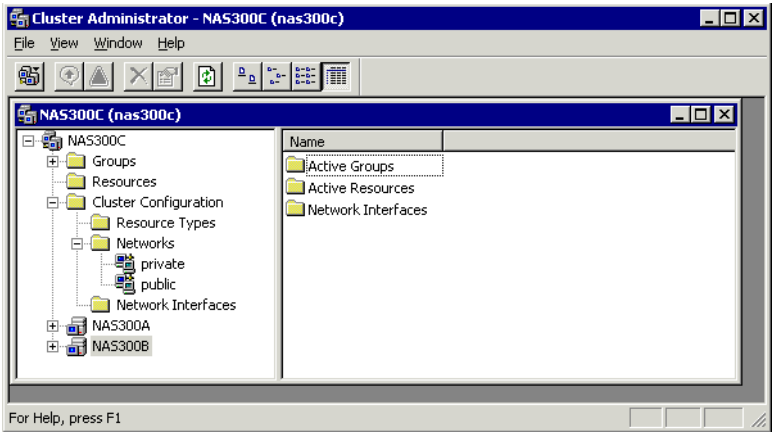

Figure 6-113 Cluster Administrator with two member nodes

Congratulations! Now you have successfully set up the cluster for the NAS 300. Next it is time to test the functionality of the cluster.

Testing the cluster

Now that the cluster is successfully set up, it should be tested for failover function. A simple test to verify this is to **move** the resources from one node to another. This is equivalent to shutting down (normal or abnormal) the node that owns the resources. If the ownership moves, the cluster is functioning properly.

Moving groups of resources from first node to second node

Following are the steps for moving groups of resources from the first node to the second node:

1. Start Cluster Administrator by opening the IBM NAS Admin Utility, select **Cluster Tools**, and start the **Cluster Administration**.
2. When prompted with *Open Connection to Cluster*, type the cluster name (Figure 6-114).

Figure 6-114 Cluster connection

3. Once the cluster window is opened, select **Groups**. It is shown in Figure 6-115 that the current owner of the groups or resources is Node A (which is NAS300A in our example).

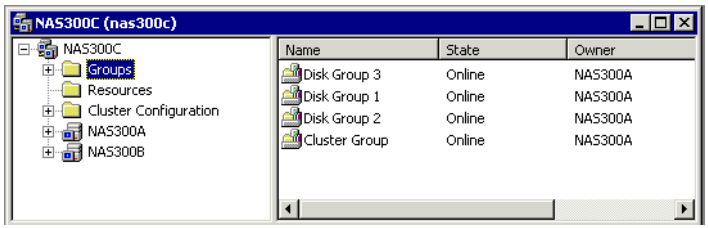

Figure 6-115 Cluster groups owned by NAS300A

4. On the right pane, right-click **Cluster Group** and select **Move Group** (Figure 6-116).

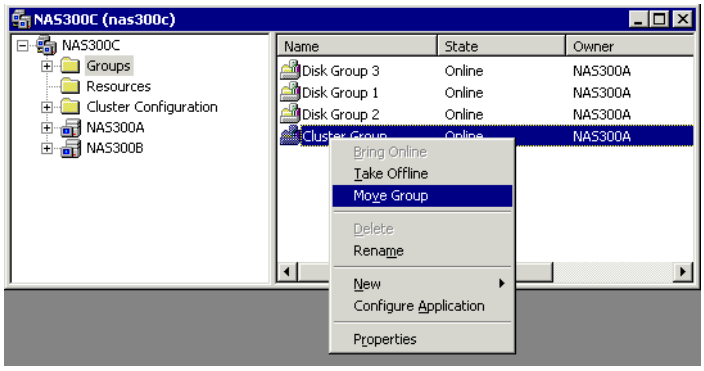

Figure 6-116 Moving a cluster group

5. The group's state should change from Online on Node A to Online on Node B as shown in Figure 6-117.

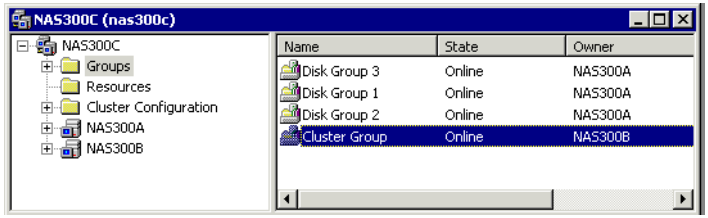

Figure 6-117 Cluster group moved

6. Perform all previous steps on the other groups. If all can be moved, the cluster setup is successful (Figure 6-118).

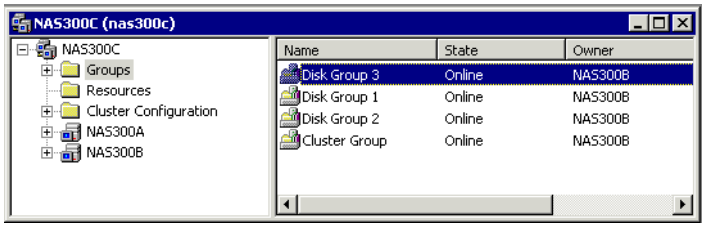

Figure 6-118 Cluster group with everything moved

Now you can try shutting down Node B, and the groups should move back to Node A.

If you want more information about cluster setup and administration on IBM NAS systems, refer to Chapter 4 in the redbook, *Implementing the IBM TotalStorage NAS 300G, High Speed Cross Platform Storage and Tivoli SANergy!*, SG24-6278.

http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246278.html?Open

7

User and security management

In this chapter we describe the integration of the NAS appliances into a secure environment. We have skimmed over this subject pretty lightly so far, but in keeping with its role as an appliance, the NAS 100, 200, and 300 are designed to plug right into your existing user and security management system.

We cover these topics:

- Active Directory, NT 4 Domains, and Workgroups
- UNIX NIS integration

7.1 Active Directory, NT 4 Domains, and Workgroups

The NAS 200 and 300 will integrate with all of the Microsoft Operating System versions that you have in your current network environment. It will work with any existing user and security management for those systems, including:

- Windows Workgroup Computing (NAS 200 only)
- Windows NT 4 Domains
- Windows 2000 Active Directory (mixed and native mode)

A full description of user and security management for Windows is beyond the scope of this book, so we simply give you a quick overview. A more detailed look at Active Directory is provided in Chapter 13, “Microsoft Active Directory and IBM TotalStorage NAS” on page 467. Also, refer to the literature regarding Microsoft Operating Systems. Some examples are listed in “Related publications” on page 637.

To change the security environment for your NAS 200 and 300, from the desktop (or from a Windows Terminal Service session of your remote administrative console), right-click **My Network Places** and choose **Properties**. In the resulting window, select the **Advanced** menu and choose **Network Identification**.

Now you have the option to join the NAS 200 and 300 to your environment, as shown in Figure 7-1.

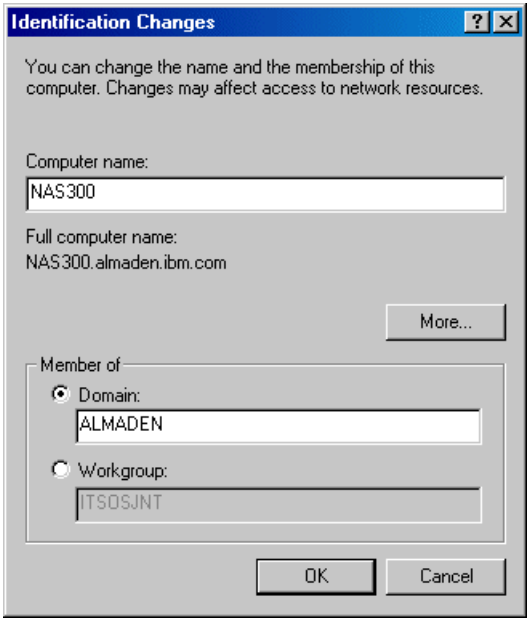

Figure 7-1 Setting up network identification for Active Directory

If you want the NAS 200 to be a storage device for your workgroup, check **Workgroup** and type in the name of your workgroup. Be aware that all security management within such a workgroup is local. This means you have to administer all user accounts on the NAS 200. Also be aware that the default behavior when sharing a network drive with Windows is to grant all users full access to the data.

Important: The NAS 300 needs to be member of a Windows Domain or Active Directory because the Microsoft Cluster requires a Windows domain controller. Remember that Release V2.5 allows you to create a DNS server and Domain on the NAS 300 servers themselves.

When you want to join an existing Windows NT 4.0 Domain, check the **Domain** box and type in the name of your domain. Press **OK** to confirm your choice and you will soon be a happy member of an existing Windows NT 4.0 Domain.

Tip: When joining an existing NT 4 Domain across subnets or via routed paths, define the PDC as the primary WINS server, even if the WINS service is not running on the PDC. This way, the joining client will find the PDC easily.

Joining an Active Directory tree is almost identical to joining an NT 4 Domain. Check **Domain** and type in the name of your Active Directory tree.

Important: When joining an Active Directory, it is essential that your TCP/IP configuration and DNS name resolution be working properly. Make sure both machines can ping each other using the IP address and the fully qualified domain name before joining the domain. For example, type:

```
ping NAS200.nas.redbook.com
```

7.2 UNIX NIS integration

The UNIX Network Information System (NIS) services work just like using the yellow pages. While the feature set of the NAS 200 and 300 includes support for NIS, the security standard of NIS is not very high. Therefore, we do not recommend the use of NIS with this product.

The NAS 200 and 300 come preconfigured with Microsoft Services for UNIX 2.2. Within these services you have a NIS migration wizard. This tool allows you to migrate a NIS. The tool takes your NIS source files and migrates them into Active Directory.

The Server for NIS feature allows a Windows Domain controller to be an NIS master server or an NIS subordinate (slave) by integrating NIS into Active Directory. When using the NIS server as a slave, the NIS master server must be a Windows 2000 Server.

For detailed information, check the following Web site:

http://www.microsoft.com/WINDOWS2000/sfu/default.asp

Password synchronization

Another tool that is included within the Microsoft Services for UNIX 2.2 is a password synchronization tool (2-way). It allows you to synchronize password changes between Windows NT or Windows 2000 and UNIX. Pre-compiled single sign-on daemons are available for:

- IBM AIX 4.3+
- Linux (Red Hat 5.2, 6.0, 7.x)
- Sun Solaris 2.6+
- HP-UX 10.3+
- Compaq Tru64 UNIX

Note: Even if your UNIX version is not on the list — it may still work. Microsoft provides the source code for the password synchronization tool

Part 3

Advanced IBM TotalStorage NAS topics

In Part 3 of this book we explore more advanced NAS configuration topics:

- First we take a look at the advanced configuration of the NAS appliance systems (Chapter 8, “Advanced NAS configuration” on page 207).
- After we have finished with the configuration process, we then show how to manage the NAS appliance using IBM systems management tools (Chapter 9, “Systems management for NAS” on page 285).
- Next we explain how to set up the NAS appliance as a cross platform storage solution (Chapter 10, “Cross platform storage” on page 353).
- Then we discuss some backup and restore considerations, and we integrate the Persistent Storage Manager (PSM) into this picture to show how it can be used with Tivoli Storage Manager (Chapter 11, “Backup considerations” on page 431).
- Next we show how to rebuild the NAS appliance from scratch in case of a disaster (Chapter 12, “NAS recovery procedures” on page 453).
- Finally, we show how to integrate IBM TotalStorage NAS and Microsoft Active Directory (Chapter 13, “Microsoft Active Directory and IBM TotalStorage NAS” on page 467).

8

Advanced NAS configuration

In this chapter we provide a detailed description of tools and procedures for advanced configuration of your IBM TotalStorage NAS 200 and 300.

We cover these topics:

- Quota management and WQuinn StorageCeNTral
- File screening
- Dynamic storage expansion on IBM NAS
- Double-Take
- Persistent Storage Manager (PSM)
- Additional tools
- Uninterrupted Power Supply support

8.1 Quota management and WQuinn StorageCeNTral

Previous releases of IBM TotalStorage NAS have offered some basic quota functionalities. Release 2.5 includes several advanced quota functionalities from a limited version of the WQuinn StorageCeNTral suite.

New functions included are:

- Real-time quotas on directories only.
- The Storage Reports set of tools, including duplicate files, file type summary, files not being backed up, large files, and most-commonly-used file reports.
- A screening tool to block any types such as MP3, graphic files, or executables from being stored on server appliances.

Note: The WQuinn StorageCeNTral suite includes more than those included in the Secure Attention Key (SAK). The suite comprises three products: QuotaAdvisor, DiskAdvisor, and FileScreen.

These functionalities are accessible on the Disks screen, as shown in Figure 8-1.

To access this screen, open your Web browser and use the following example, modified as necessary for your environment:

```
http: //computername or ip address:8099
```

When prompted for a username and password, use the administrative account (for example, administrator, password). The NAS main screen will appear; select **Disks** (Figure 8-1).

Figure 8-1 Disks — main screen

8.1.1 Disk quotas

The following steps allow you to set up disk quotas on the NAS devices:

1. Click **Disk Quota** in the Disks Main screen in Figure 8-1.
2. You get the Disk Quota screen (Figure 8-2) where you can select a volume to define disk quotas:
   - On a per-volume basis (click **Quota**).
   - On a per-user basis (click **Quota Entries**).

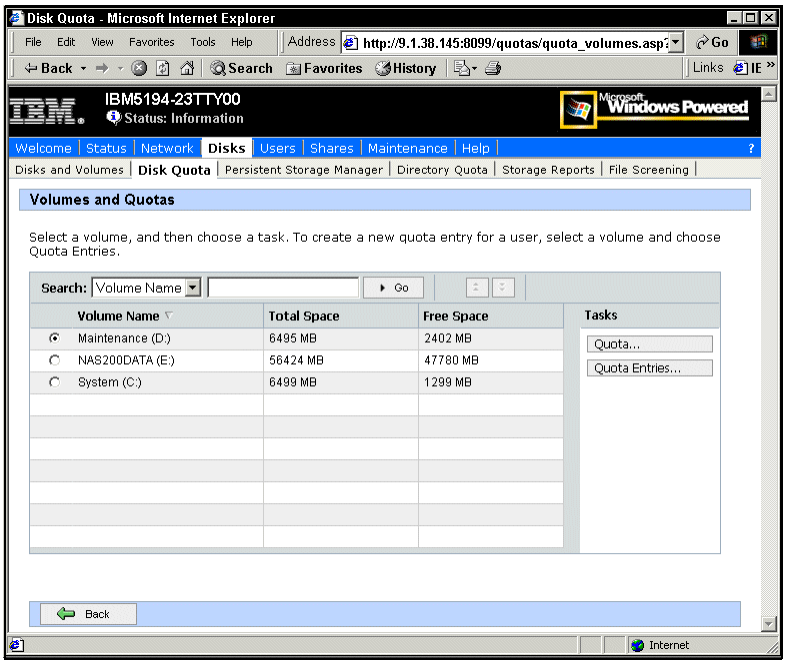

Figure 8-2 Disk Quota screen

3. Click **Quota** to set all quota parameters for a volume, as shown in Figure 8-3.

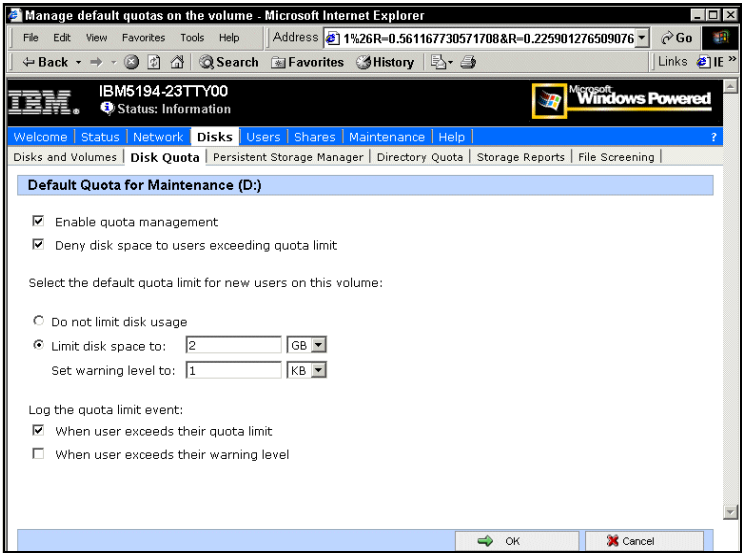

Figure 8-3 Disk Quota settings screen

4. When finished, click **OK** to get back to the Disk Quota screen.
5. Click **Quota Entries** to get the Quota Entries for the selected volume screen (Figure 8-4).

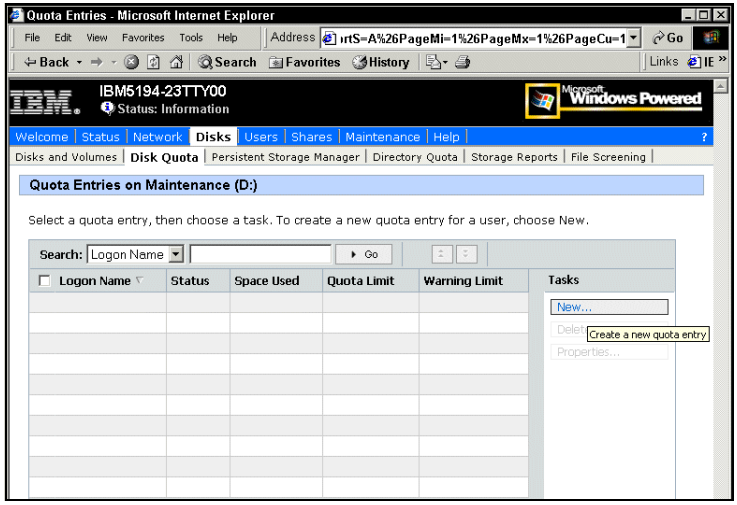

Figure 8-4 Quota Entries screen

6. Click **New** to get the New Quota Entry screen (Figure 8-5) to create a quota for a user in the selected volume. Configure the settings and click **OK**.

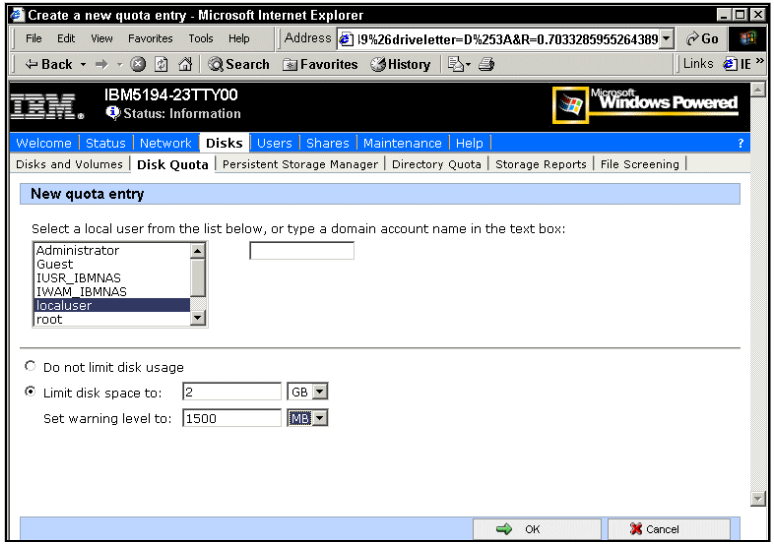

Figure 8-5 New Quota Entry screen

7. Now you get back to the Quota Entries screen that shows the new user quota (Figure 8-6).

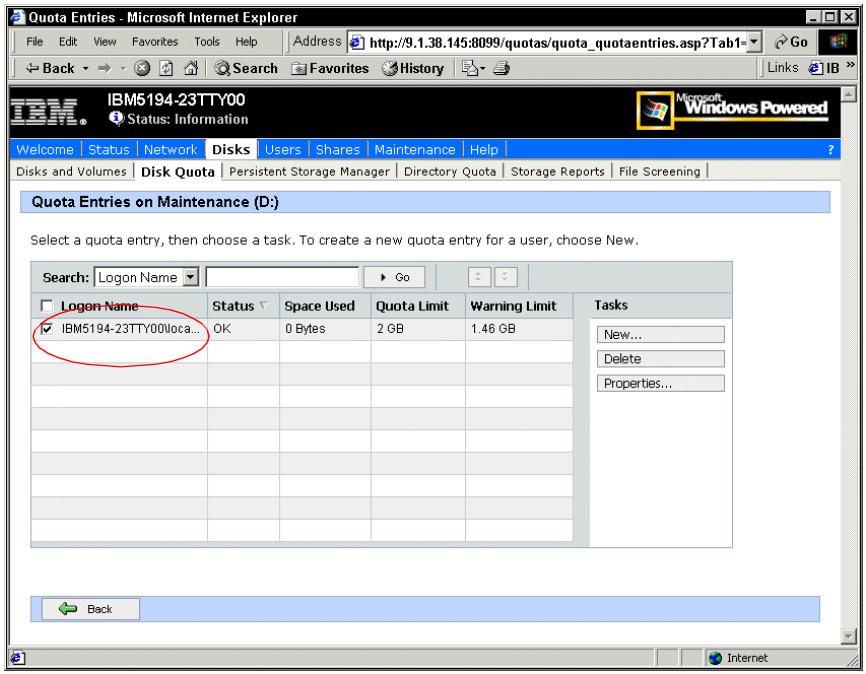

Figure 8-6 Quota Entries screen

Click **Back —> Back** to get back to Disk Main screen (Figure 8-1 on page 209).

8.1.2 Directory quotas

Directory quotas allow you to add, delete, monitor, and change space limits for selected directories on the server appliance. Directory quotas provide disk space monitoring and control in real time, and support active and passive limits with two real-time alarms.

The directory quota feature includes the following components:

- Active and passive space limits on directories
- Best practice storage resource management policies
- Two levels of alarm threshold: Severe and Warning
- Auto discovery of drives
- Customized messages
- Alarms sent to the user and/or to the event log
- Storage Reports that can be sent to an intranet Web site
- Custom script

To configure directory quotas:

1. In the Disk screen, click **Directory Quota** (Figure 8-1).
2. In the Directory Quota screen, click **New**.
3. In the Initial wizard screen, click **Next** to get the Create New Directory Quota screen.
4. Set the directory name for which you want to assign a quota, and then choose the policy setting for this directory (Figure 8-7).

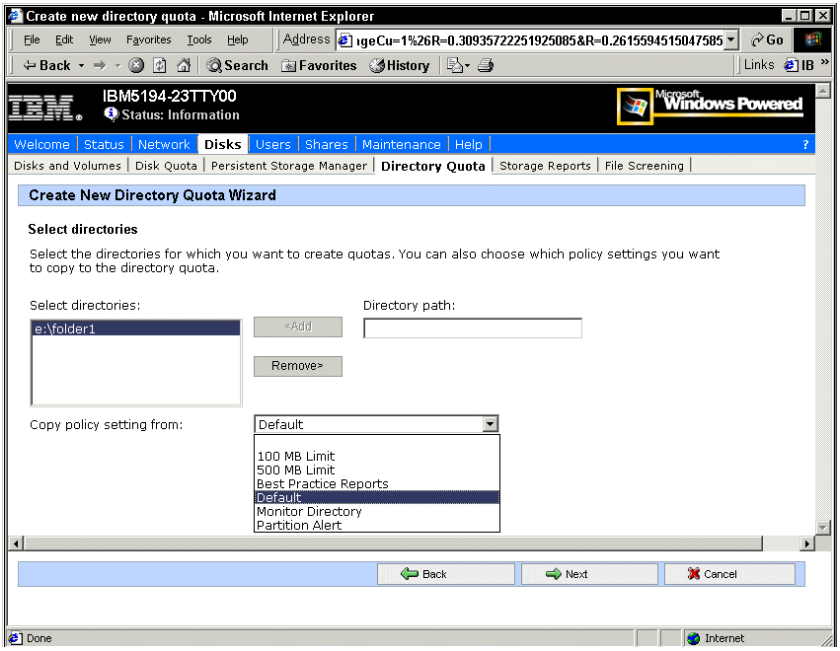

Figure 8-7 Directory quota settings

5. Click **Next**.
6. Click **Finish** to exit the wizard.

7. You will get back to the Directory Quota screen that now shows the recently created directory quota. See Figure 8-8.

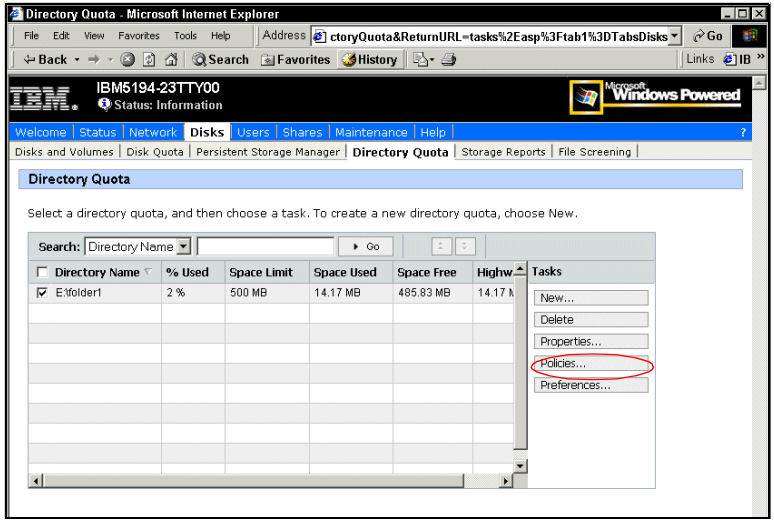

Figure 8-8 Directory Quota screen

In this screen you can define advanced settings for your directory quota management. You have three main options:

- Properties
- Policies
- Preferences

Policies

By clicking **Policies** in the Directory Quota screen (Figure 8-8), you will get to the Directory Quota Policies screen (Figure 8-9). In this screen you can create, delete, or modify a storage policy.

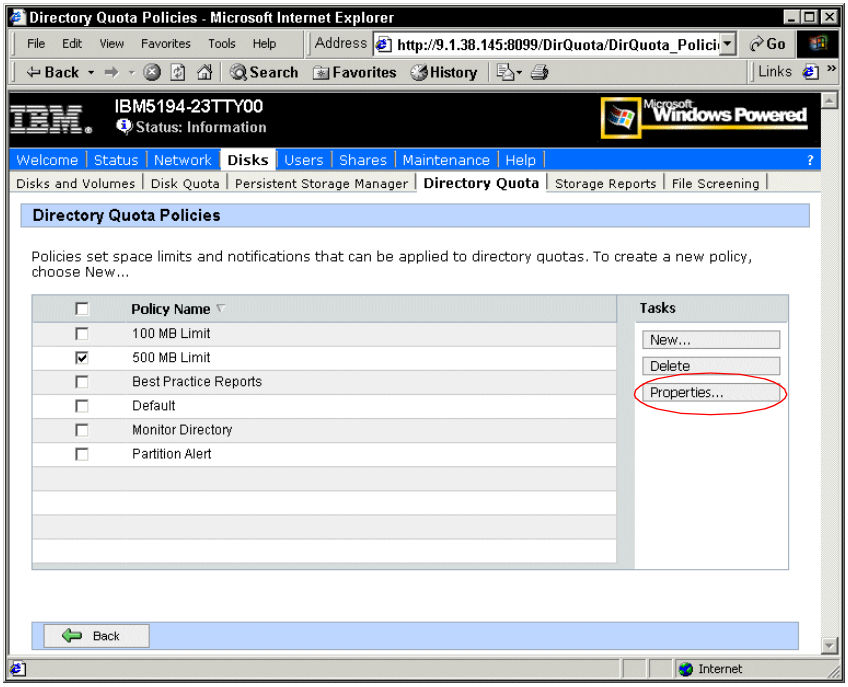

Figure 8-9 Directory Quota Properties screen

There are some predefined policies that can also be modified by selecting and clicking **Properties** in the Directory Quota Policies screen (Figure 8-9).

Now you get the Directory Quota Policy properties for the selected policy. In this screen, you have three tabs that allow you to configure several options:

- Disk Space Limit
- Severe Alarm
- Warning Alarm

You can access these tabs by clicking in them:

- **Disk Space Limit tab:** As shown in Figure 8-10, you can set the disk space limit in this folder. Also, in the case of reaching the quota limit for the directory, you can decide whether only alarms should be activated, or whether the usage of additional space in the folder is forbidden.

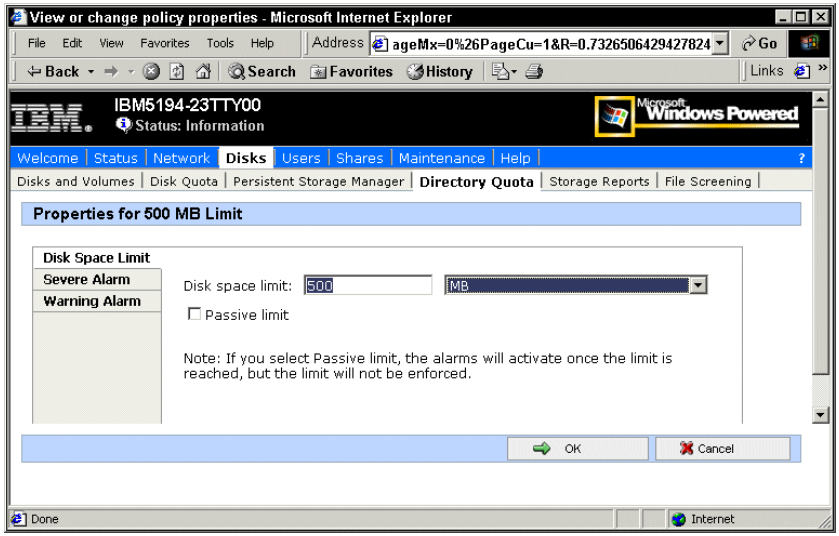

Figure 8-10 Directory Quota Policy Properties: Disk Space Limit

- **Severe Alarm tab:** As shown in Figure 8-11, you can configure up to two alarms with different level of severity: warning and severe. The severe alarm has a default threshold value of 90% and has the following configurable settings:

  Notification message This is the message that will be sent if the severe alarm triggers. You can customize this message and decide if it will be sent to the user, logged it in the event log or both.

  Action to perform After the alarm is activated, you can choose which action will occur: run a program, run a Storage Report (you can find information on this in 8.2.1, “Storage Reports” on page 221) or both.

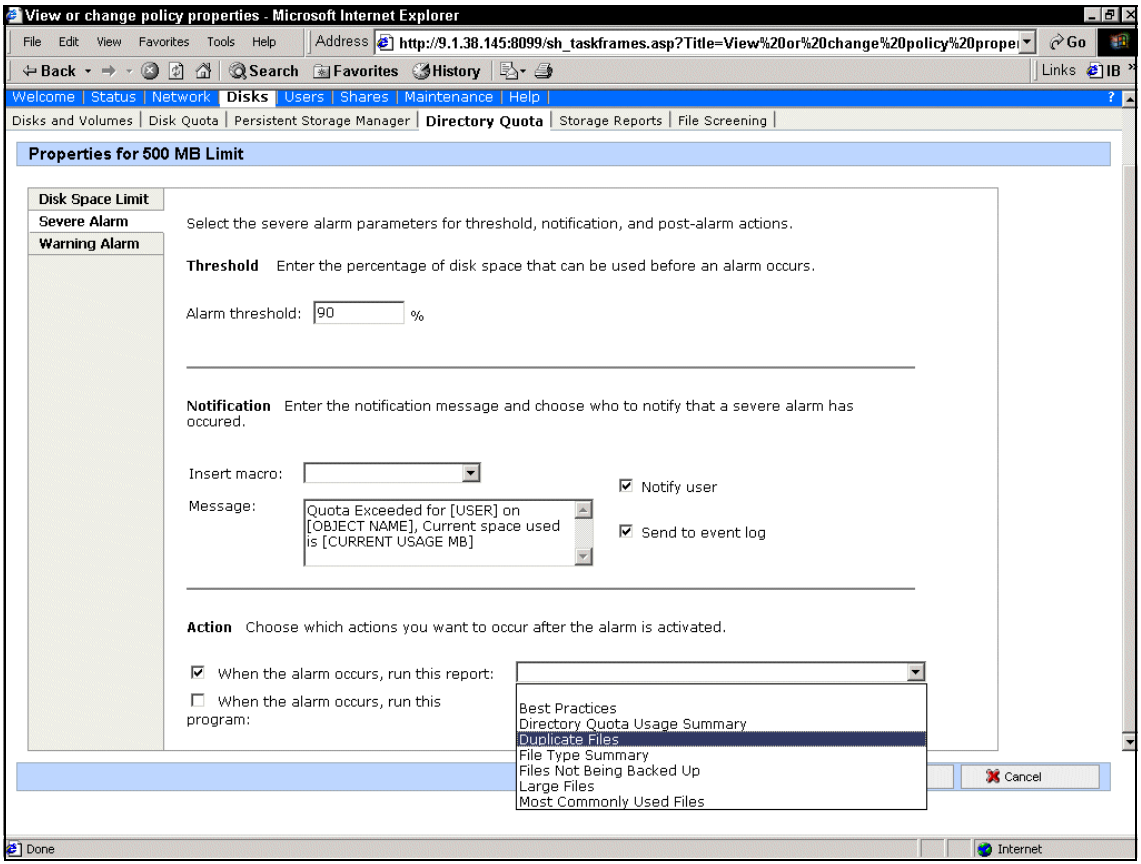

Figure 8-11 Directory Quota Policy Properties: Severe Alarm

- **Warning Alarm tab:** The warning alarm has the same configuration settings as the severe alarm (shown previously in Figure 8-11). However, the threshold default value for the warning alarm is 70% of the capacity.

Properties

You get the Directory Quota preferences screen for a folder (Figure 8-12) by selecting a share and clicking **Properties** in the Directory Quota screen (Figure 8-8).

In this screen you can configure all settings seen in the policies section for the folder selected; or alternatively, you can copy the settings from one of the already created policies (Figure 8-12).

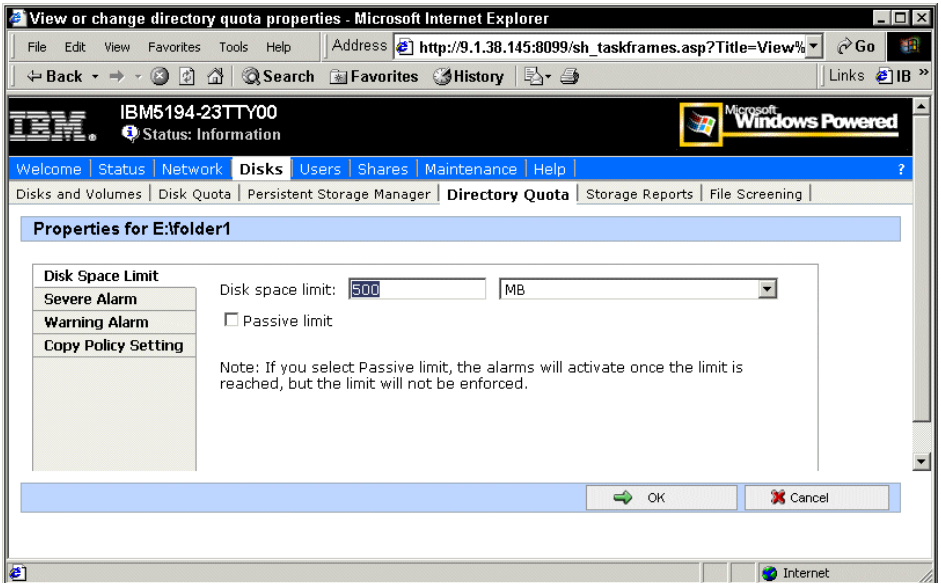

Figure 8-12 Directory Properties screen

Preferences

You get the Set Directory Quota preferences screen (Figure 8-13) by clicking **Preferences** in the Directory Quota screen (Figure 8-8). In this screen, by checking the box labeled “Auto detect disks”, you can automatically detect new mounted drives and apply one of the storage management policies, as explained in the previous section.

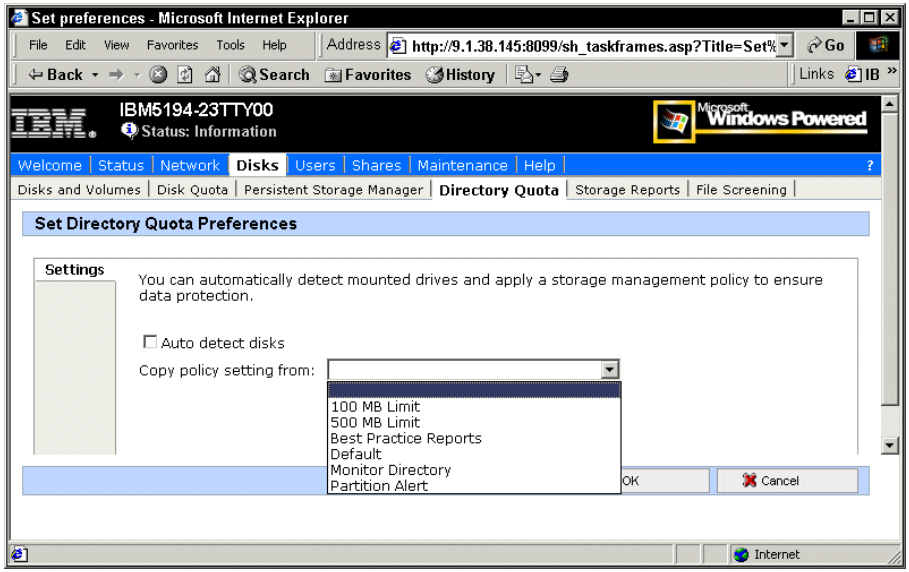

Figure 8-13 Set Directory Quota Preferences screen

Tip: Use caution when placing a directory quota on the system partition. If the system does not have sufficient quota to write files, it may fail. Also, if the system partition does not have enough space to write temporary files during system boot, the system may not restart. You can avoid this by using caution when placing quotas on the system partition or system directories.

8.2 File screening

File screening prevents unwanted files from being stored on the NAS appliance. These files increase backup times and consume valuable disk space, and they can also put your company at risk of violating piracy and discrimination laws. File screening includes the following features:

- Active and passive file screening on directories
- Best practice file screening policies
- Notification alarm when file screening policy is violated
- Audit database containing screened files
- Customized alarm messages
- Alarm messages to the event log
- Alarm messages to a user
- Storage Reports when alarm is activated and sent to intranet Web site
- Custom script when alarm is activated

To get the File Screening screen (Figure 8-14), click **File Screening** in the Disk main screen (Figure 8-1).

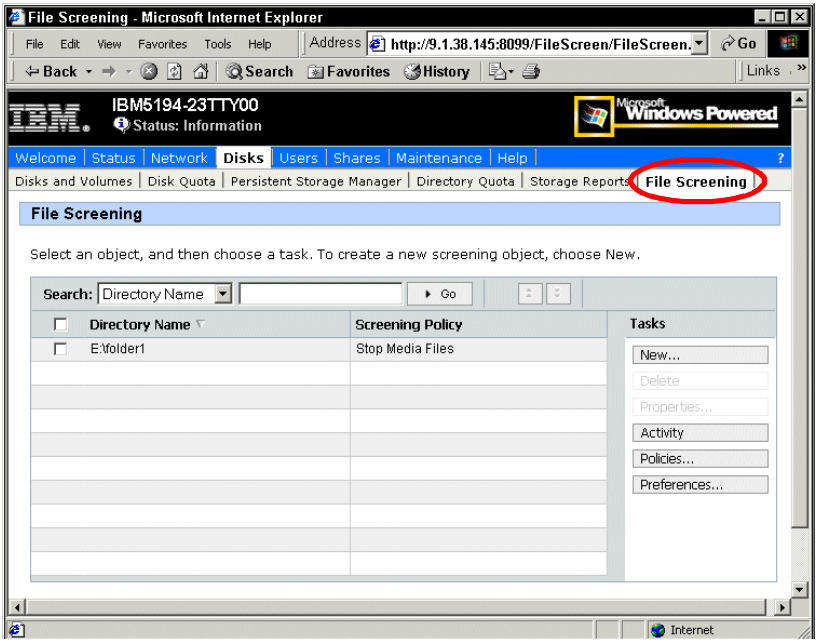

Figure 8-14 File Screening screen

In this dialog box, you can:

- Create or delete a screening object (a folder to screen).
- Modify the properties of this screening object.
- View the file screening activity on the selected folder.
- Create or modify policies for screening files. There are predefined screening policies that can also be modified (Figure 8-15).
- Set the parameters related to the audit database.

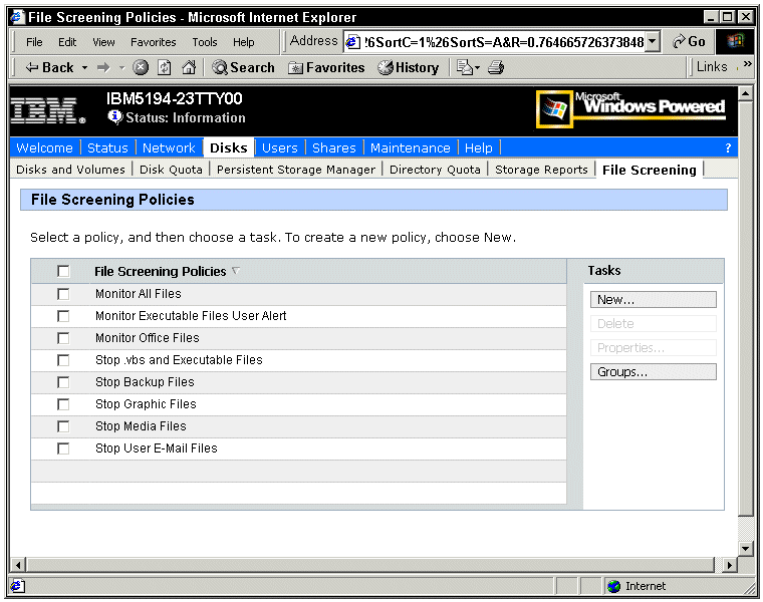

Figure 8-15 File Screening Policies window

8.2.1 Storage Reports

The Storage Reports tool addresses disk usage, wasted space, file ownership, security, and administration. Reports can run interactively, scheduled on a regular basis, or run as part of a storage resource management policy when disk space utilization reaches a critical level as seen in the preceding sections.

To get the Storage Reports screen (Figure 8-16), click **Storage Reports** on the Disk main screen (Figure 8-1 on page 209).

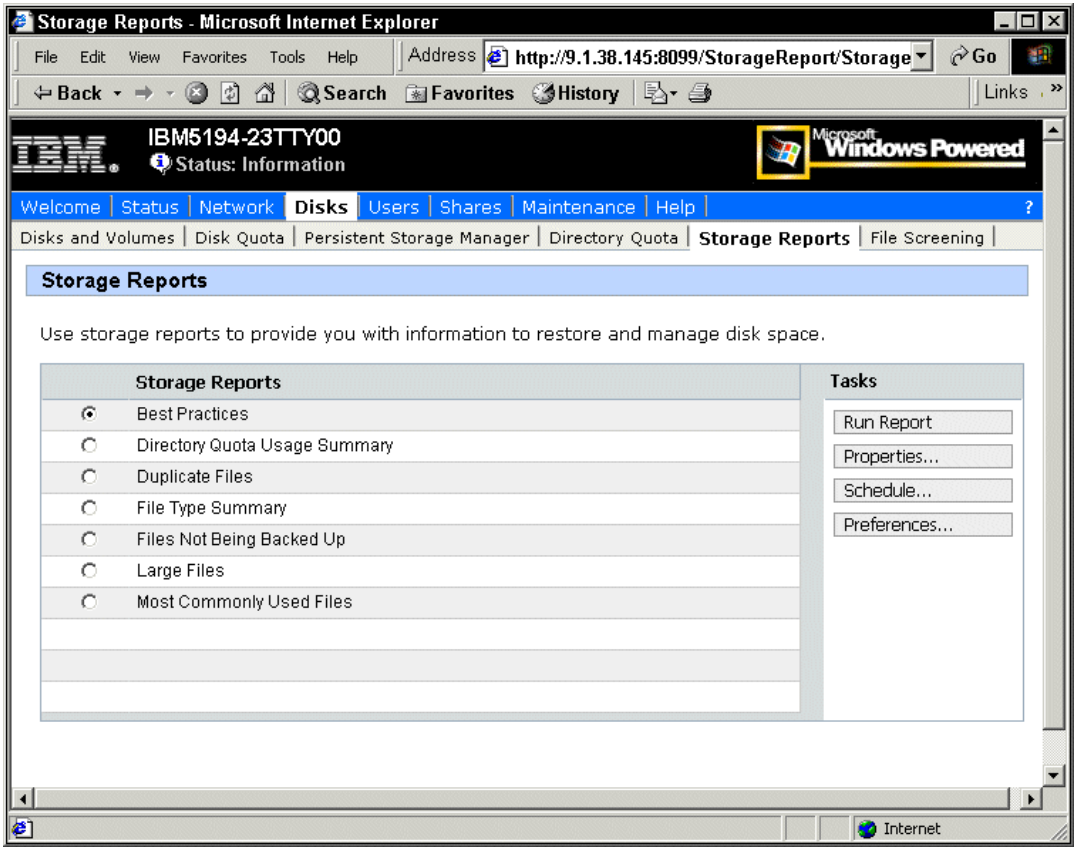

Figure 8-16 Storage Reports main screen

In this screen, you can:

- Run or schedule one of the predefined Storage Reports.
- Define e-mail addresses or Web site as a target for the Storage Report.
- Change the properties for the predefined Storage Reports.

For an example of a Storage Report, see Figure 8-17.

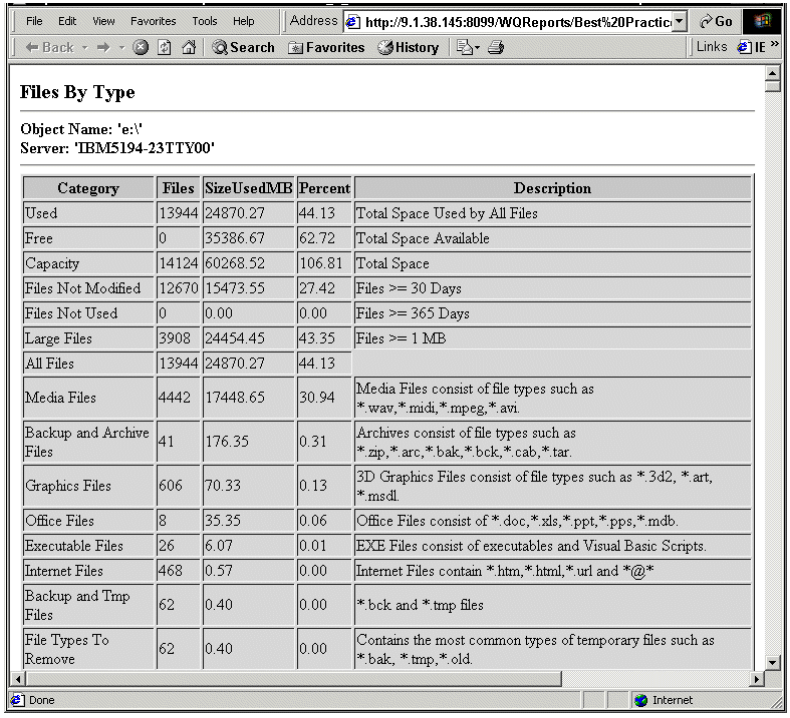

Files By Type

Object Name: 'e:\'
Server: 'IBM5194-23TTY00'

Category	Files	SizeUsedMB	Percent	Description
Used	13944	24870.27	44.13	Total Space Used by All Files
Free	0	35386.67	62.72	Total Space Available
Capacity	14124	60268.52	106.81	Total Space
Files Not Modified	12670	15473.55	27.42	Files >= 30 Days
Files Not Used	0	0.00	0.00	Files >= 365 Days
Large Files	3908	24454.45	43.35	Files >= 1 MB
All Files	13944	24870.27	44.13	
Media Files	4442	17448.65	30.94	Media Files consist of file types such as *.wav,*.midi,*.mpeg,*.avi.
Backup and Archive Files	41	176.35	0.31	Archives consist of file types such as *.zip,*.arc,*.bak,*.bck,*.cab,*.tar.
Graphics Files	606	70.33	0.13	3D Graphics Files consist of file types such as *.3d2, *.art, *.msdl.
Office Files	8	35.35	0.06	Office Files consist of *.doc,*.xls,*.ppt,*.pps,*.mdb.
Executable Files	26	6.07	0.01	EXE Files consist of executables and Visual Basic Scripts.
Internet Files	468	0.57	0.00	Internet Files contain *.htm,*.html,*.url and *@*
Backup and Tmp Files	62	0.40	0.00	*.bck and *.tmp files
File Types To Remove	62	0.40	0.00	Contains the most common types of temporary files such as *.bak, *.tmp,*.old.

Figure 8-17 Example of a Storage Report

8.3 Dynamic storage expansion on IBM NAS

As the need for storage capacity grows, customers want to be able to expand their storage. IBM provides tools to help the customer accomplish this task easily. In this section we show how to increase the disk capacity on the IBM TotalStorage NAS 200 and 300.

8.3.1 Tools for dynamic storage expansion

Several tools can used for expanding a LUN in a NAS device:

- **FAStT Storage Manager:** This utility is intended to manage the FAStT Storage Subsystem in the NAS 300. With this tool, you can add new disks to an array adding free space, and, add this free space to a LUN in the same array. The use of this tool is explained in 6.2, “Defining arrays and logical drives on the NAS 300” on page 156.
- **ServeRAID Manager:** This utility, explained in detail in 5.2, “Managing disks using ServeRAID Manager” on page 105; is used to manage the ServeRAID adapter in the NAS 200. You can expand an array and a LUN when using ServeRAID adapters in a NAS 200.
- **The diskpart utility:** This command line utility is part of the Windows 2000 Resource Kit and preinstalled on the IBM TotalStorage NAS devices. It allows you to manage disks in a more advanced way than using the Disk Management tool. It can also be used with scripting.

We take advantage of the expanding volume functionality of this tool with a Windows 2000 basic disk.

Note: Help for the **`diskpart`** utility is available in the IBM NAS Admin tool by selecting: **NAS Management —> Storage —> Diskpart —> Diskpart NAS Specific Commands**.

For a detailed description of the **`diskpart`** utility, read the article, “A Description of the Diskpart Command-Line Utility” (Q300415), in the Microsoft Knowledge Base. The article can be found at:

`http://support.microsoft.com/default.aspx?scid=kb;en-us;Q300415`

8.3.2 Increasing volume size for a NAS 200 disk

In this section we explain how to add disk space in an IBM TotalStorage NAS 200.

Expand the array and expand the LUN

To check the initial environment, we can open the IBM NAS Admin tool and select **Storage —> Disk Management Local**.

As you can see in our example in Figure 8-18, we have three volumes: System, Maintenance, and the data disk (Disk H: NAS200DATA), with 50.87GB.

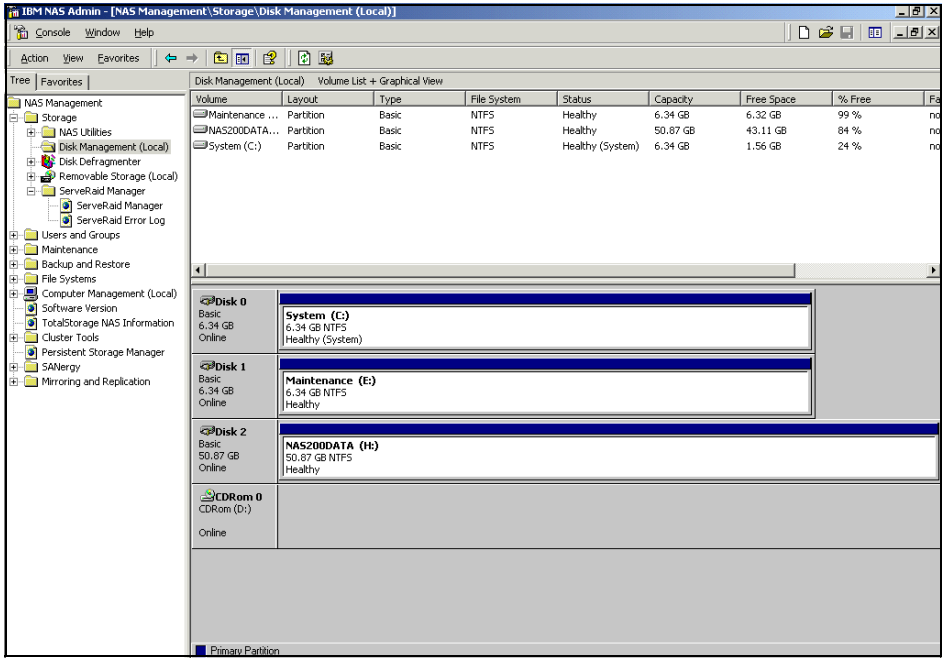

Figure 8-18 Initial partition sizes

Open the **ServeRAID Manager tool, IBM NAS Admin tool —> Storage —> ServeRAID Manager —> ServeRAID Manager**.

In Figure 8-19, you can see that we have three disks in a RAID 5 array with three logical drives or LUNs.

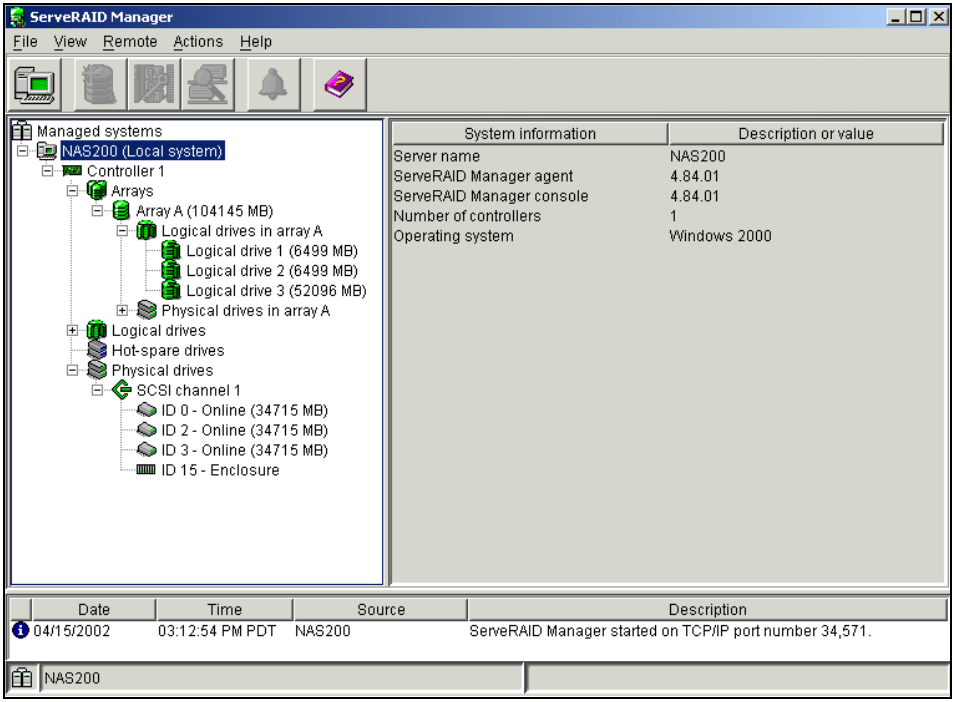

Figure 8-19 Initial ServeRAID Manager information

These are the steps that we will perform:

1. Add the disk physically to the server.
2. In the ServeRAID Manager, select the controller you are using by clicking it. (**Controller 1** in our example).
3. Select **Actions -> Scan for new or removed ready drives** from the menu.
4. When finished, click **Done**.
5. Now, as you can see, the new physical disk appears in Figure 8-20.

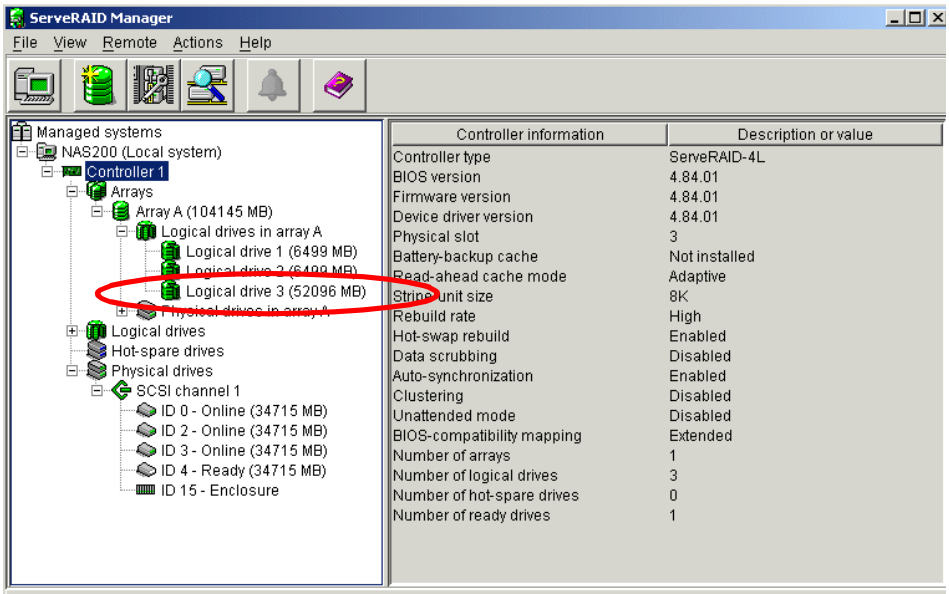

Figure 8-20 ServeRAID Manager — New disk

6. Select the array you want to expand by right-clicking the array and selecting: **Logical-drive migration —> Increase Logical Drive size**.

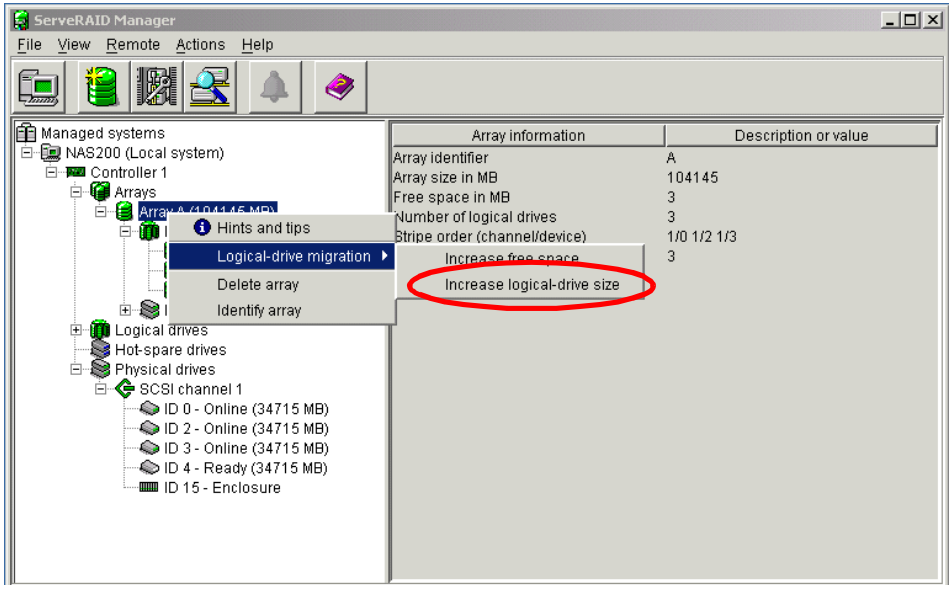

Figure 8-21 ServeRAID Manager — increase logical drive size

7. You get a dialog box (Figure 8-22), to modify the Array. You can choose the disk to add. We only have one free disk ready. Select this disk and add it to the Array A by pressing the **Add** button.

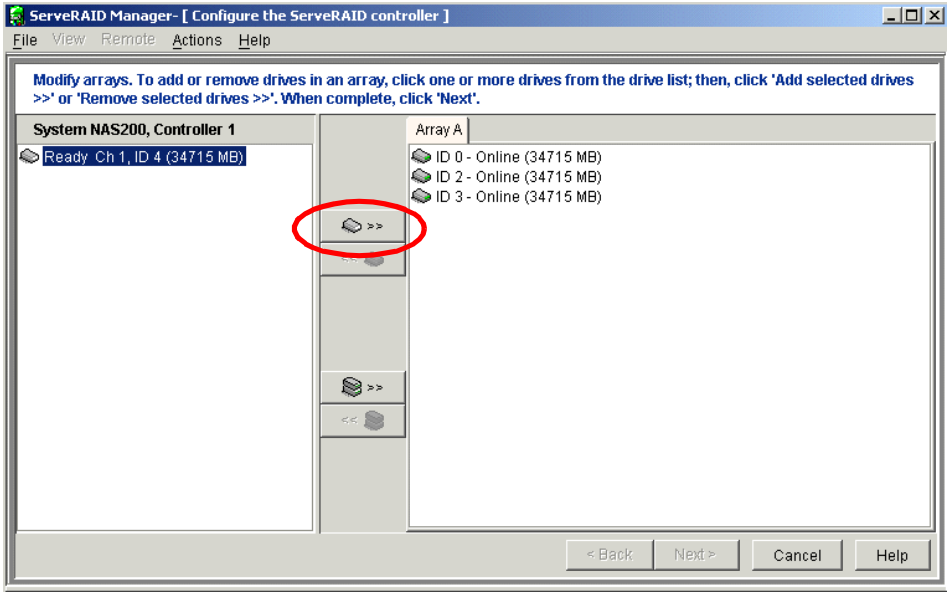

Figure 8-22 Array modification dialog box

8. Then, click **Next -> Apply**.
9. Click **Yes** in the confirmation Dialog Box.
10. The Logical Drive Migration (LDM) process starts for all LUNs in the array. We can see it in the progress bar. In Figure 8-23, notice that there is a new LUN. This LUN is a temporary one used during the process and has a different color (grey instead of green).

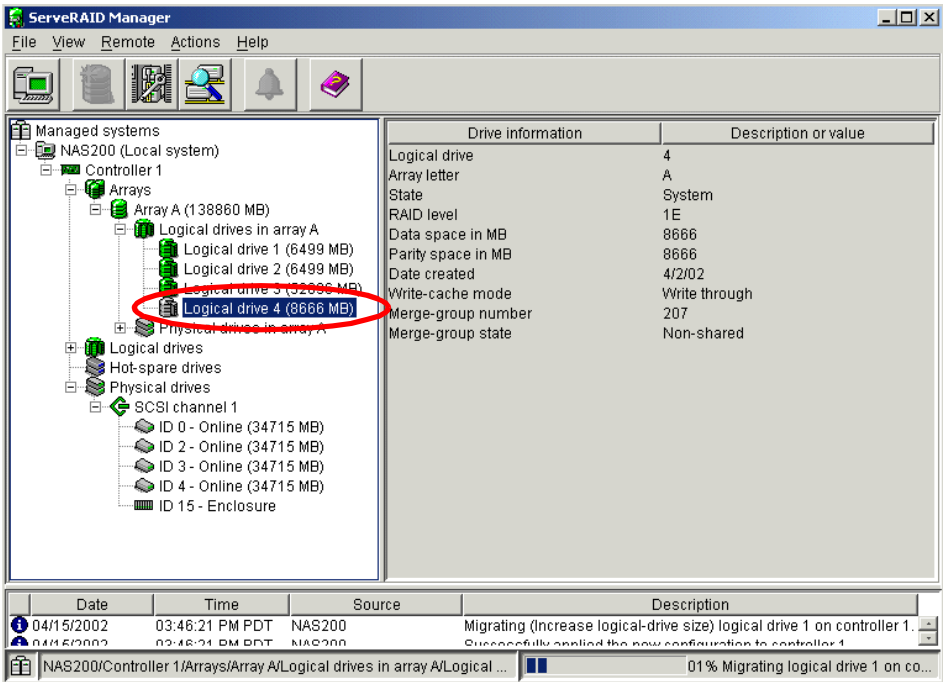

Figure 8-23 Logical Drive Migration in progress

11. This operation will take some time, but all data remains accessible during the process.

12. When the process ends, the temporary LUN disappears and the size of the other LUNs has been increased. See Figure 8-24.

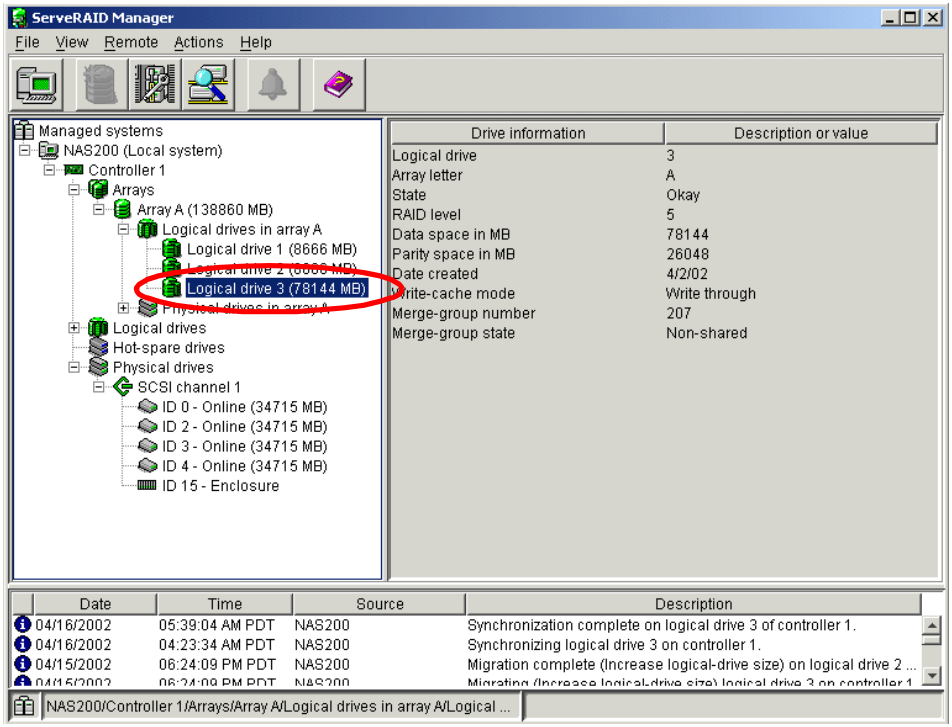

Figure 8-24 Logical Drive Migration process finished

13. We can now close this tool (**File —> Exit**) and check the Disk Management tool: **IBM NAS Admin tool —> Storage —> Disk Management Local**.

14. Now, notice that free space is added at the end of each partition during the ServeRAID logical Drive Migration process, as shown in Figure 8-25.

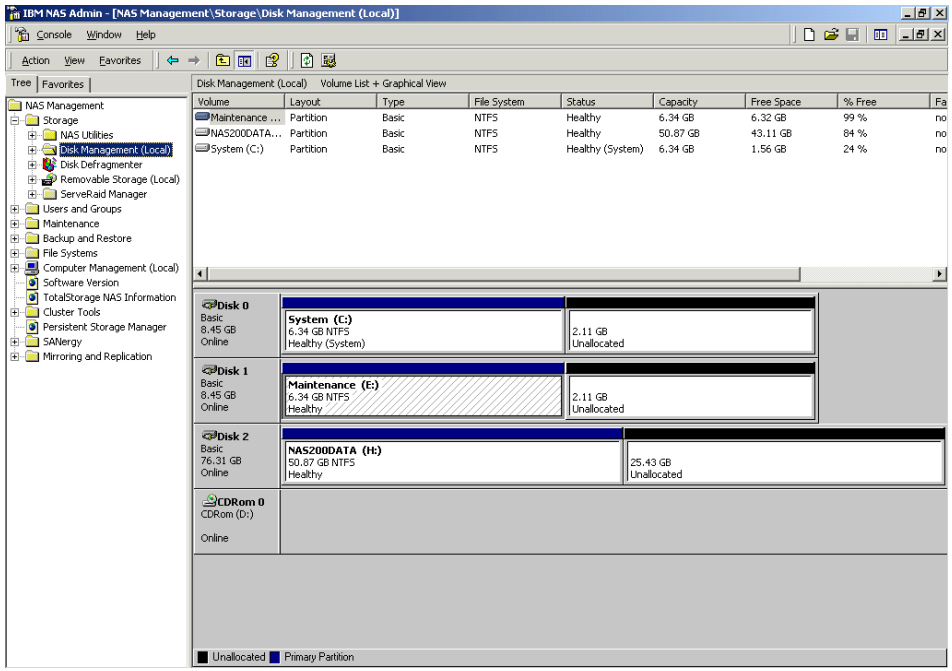

Figure 8-25 Free space added

15. If the volume doesn't have a label, this is a good time to give it one. Select the partition, right-click, choose **Properties**, and write the label in the label box.

Expanding the volume

In this section, we shown how to expand the volume with the recently created free space. We can find disks of two types: basic or dynamic.

Basic disks

This is the standard disk type. Disks created in an IBM NAS are basic disks. We use the **`diskpart`** utility, described in “Expanding the volume” on page 243, to expand these disks.

> **Tip:** This utility is preloaded in the IBM TotalStorage NAS server, and can also be downloaded from the Microsoft Windows 2000 Resource Kit Web site:
>
> `http://www.microsoft.com/windows2000/techinfo/reskit/tools/`

Here are the steps to follow:

1. Open a command line window and issue the **`diskpart`** command. Or, in the IBM NAS Admin Tool, select **`Storage —> Diskpart —> Diskpart`**.

2. This opens an interactive session with the **diskpart** utility. Now you can list the volumes in the computer by issuing the **List Volume** command.
3. Check the volume number your volume has by comparing the label. The volume we want to extend is the volume 3 (label is NAS200DATA). Therefore, we need to select it with the command **select volume 3**.
4. Then, issue the command **extend** to extend the selected volume.
5. We rescan the disks with the **rescan** command, and again we list all volumes to see that the capacity of volume 3 has been increased.
6. Now we can exit the utility with the **exit** command.

Important: Always perform a backup of your disks before you use the **diskpart** utility.

7. The entire process of expanding the volume is shown in Figure 8-26.

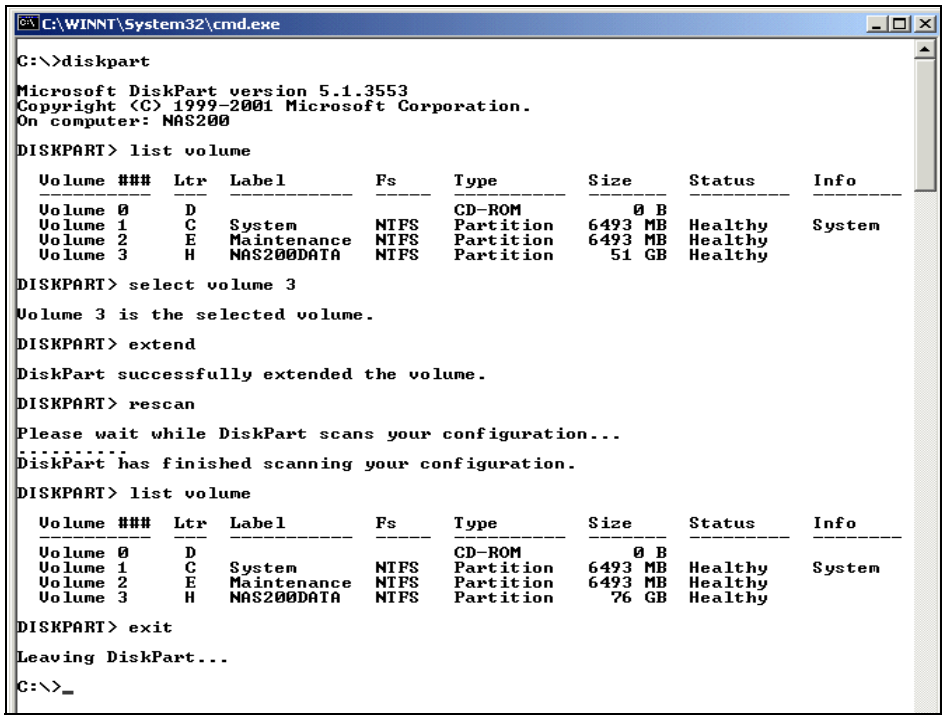

Figure 8-26 Diskpart process

8. Open the Disk Management tool again to see that the new storage is recognized for the operating system, as shown in Figure 8-27.

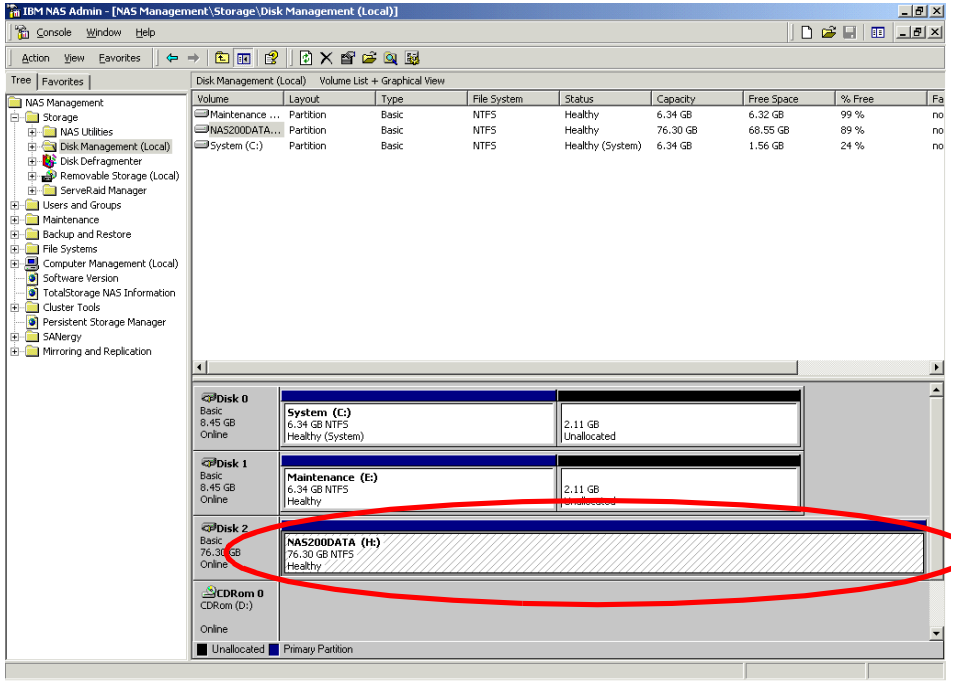

Figure 8-27 Final size

9. We can repeat the process for all the volumes.

Tip: This procedure can also be executed with a script file (**`script_file_name`**) by issuing the command:

```
diskpart /s script_file_name
```

Now you have successfully expanded your storage in an IBM TotalStorage NAS 200.

Dynamic disks

If the disk was created as a dynamic disk, we can use the standard Windows Disk Management console to expand it with the new free space created:

1. Open the Disk Management tool, IBM NAS Admin tool, and select **Storage —> Disk Management Local**.
2. As an example, in Figure 8-28, you can see that there is new free space at the end of a dynamic disk (Disk4) labeled migrate.

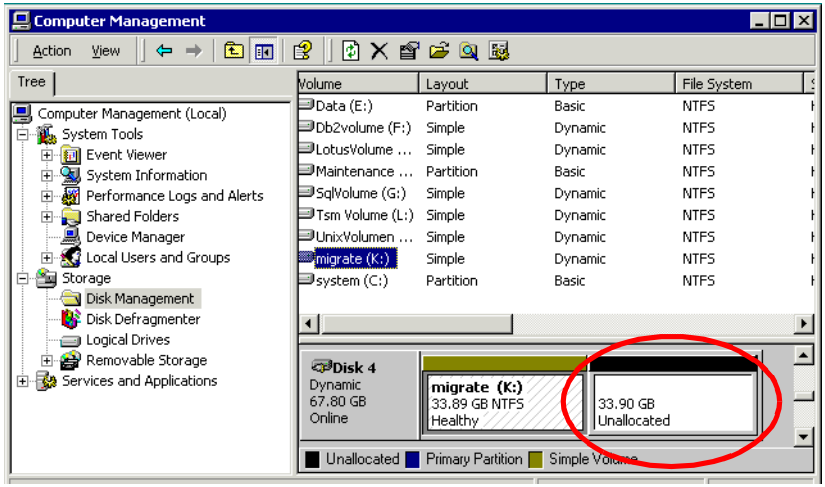

Figure 8-28 Free space on a volume

3. Right-click **Disk 4 —> Extend Volume** to extend the volume (Figure 8-29).

Note: In order to extend this disk, it has to be created as a Dynamic Disk.

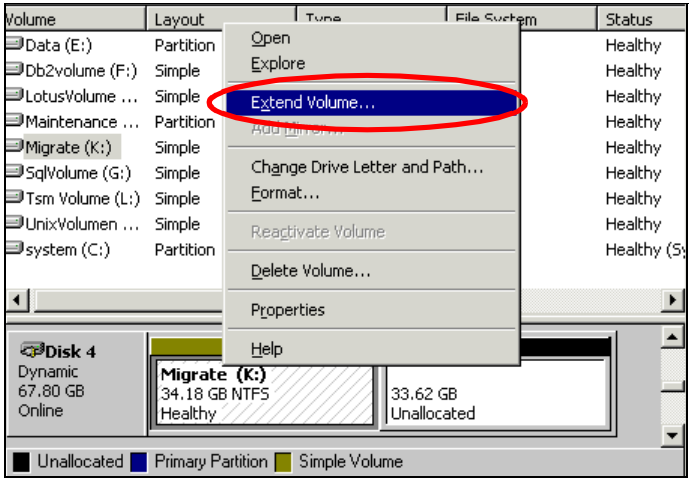

Figure 8-29 Extending a volume

4. Select the disk space that you want to extend, as shown in Figure 8-30, and click **Next**.

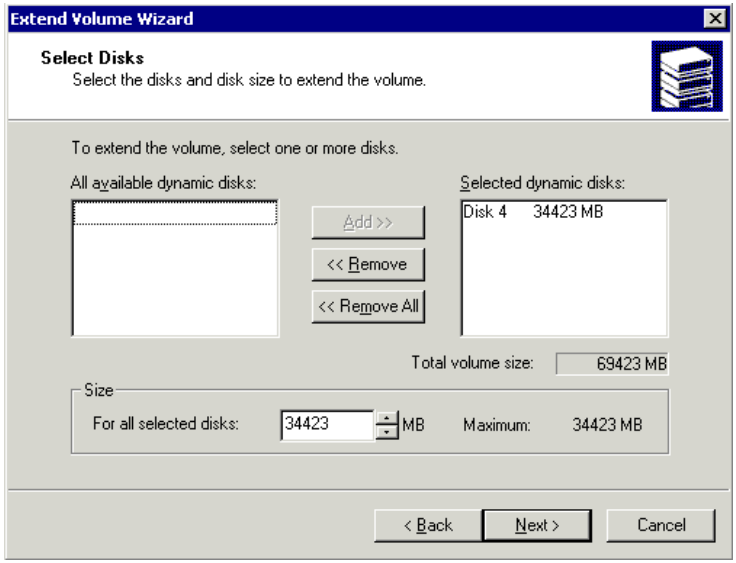

Figure 8-30 Selecting disk available

5. Now you can see the results in Figure 8-31. The disk volume size is 67.80 GB and is available in drive K.

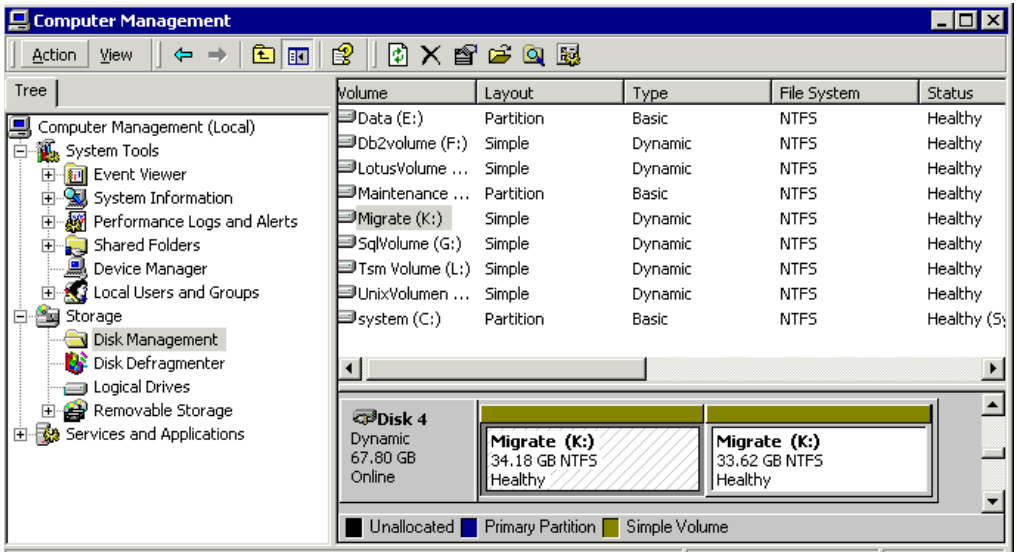

Figure 8-31 View of new volume

6. The next display, Figure 8-32, shows the disk properties for the new drive.

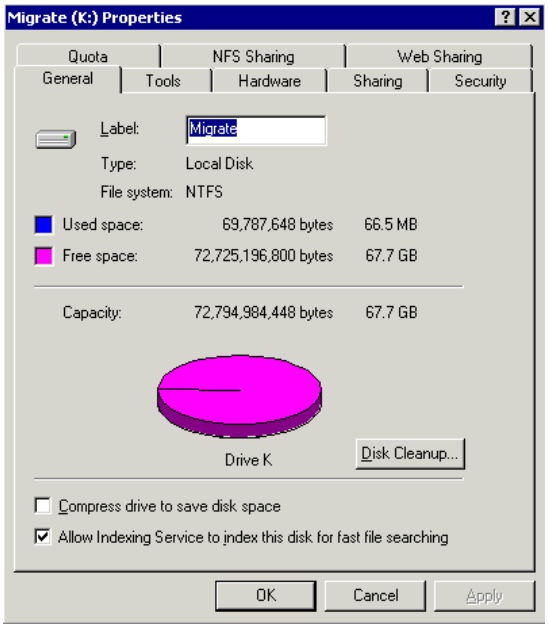

Figure 8-32 View of disk properties

You are now finished with this procedure.

8.3.3 Increasing volume size for a NAS 300 shared disk

In this section we explain how to add disk space to a currently shared disk in an IBM TotalStorage NAS 300. This procedure is also valid for a NAS 300G connected to a FAStT Storage device using Storage Manager.

8.3.3.1 Adding space to your LUN

In this section we show how to add space to an existing LUN using the FAStT Storage Manager Client. The first step to perform is to physically add the new disk or disks. After that, you can follow these steps.

1. From the IBM NAS Admin tool, select **Storage —> NAS Utilities —> SM8**.
2. Choose the Storage Subsystem to manage by double-clicking it.
3. Select the LUN you want to manage. In the example in Figure 8-33, we choose the LUN labeled NAS Data on Array 3.

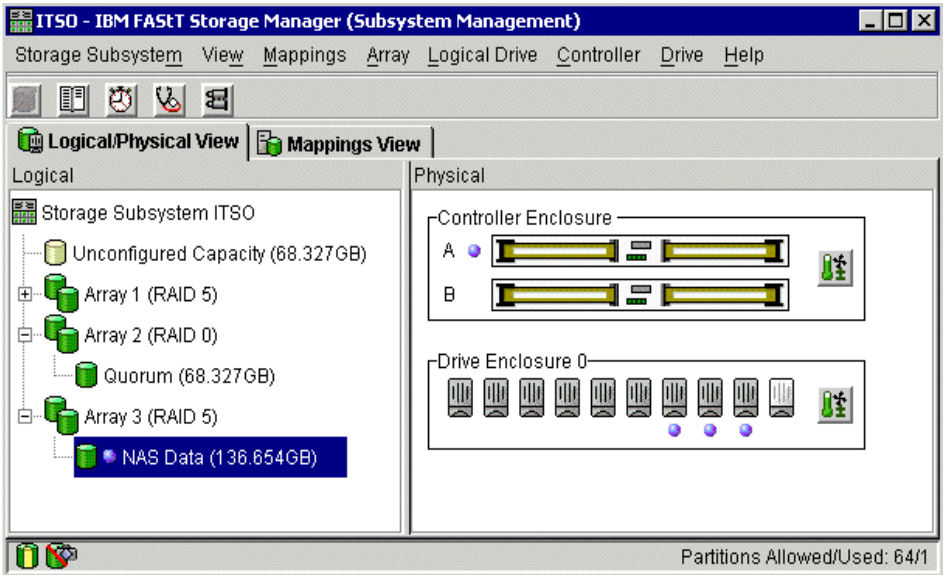

Figure 8-33 Choosing array to expand

4. Right-click the selected LUN and choose **Increase Capacity** (Figure 8-34).

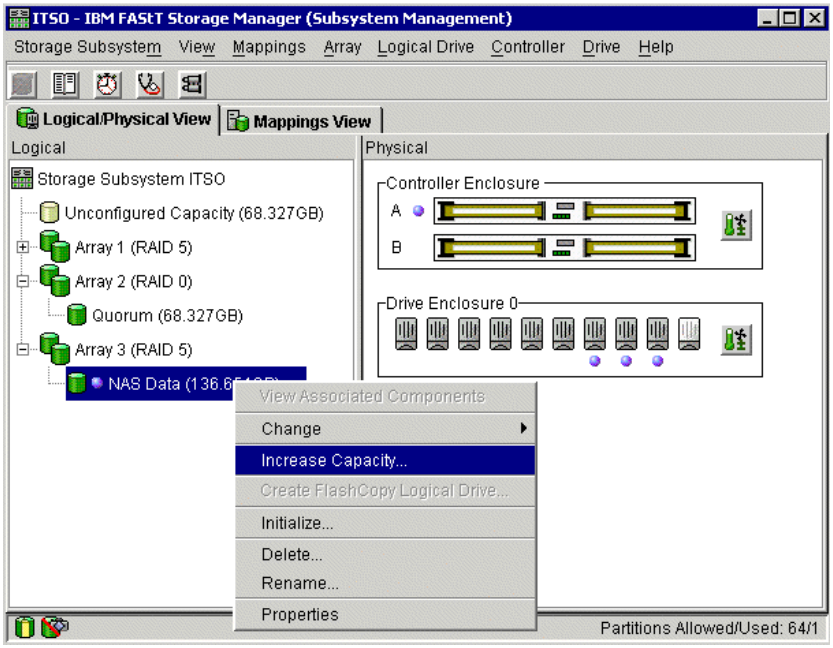

Figure 8-34 Increasing capacity in the LUN

5. You get the Increase Logical Drive Capacity dialog box shown in Figure 8-35.

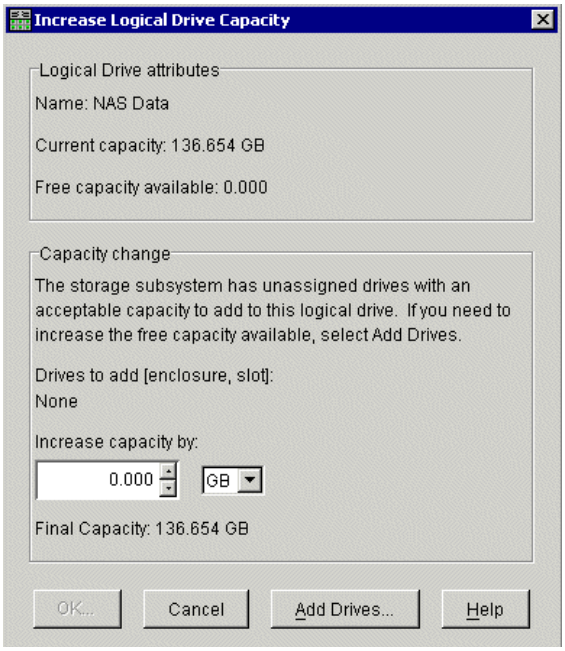

Figure 8-35 Increasing Logical Drive Capacity dialog box

6. To have free space, you need to add drives, so click **Add Drives**.

7. You get a dialog box that allows you to choose what drives to add. Select the newly added drives and click **Add** (Figure 8-36).

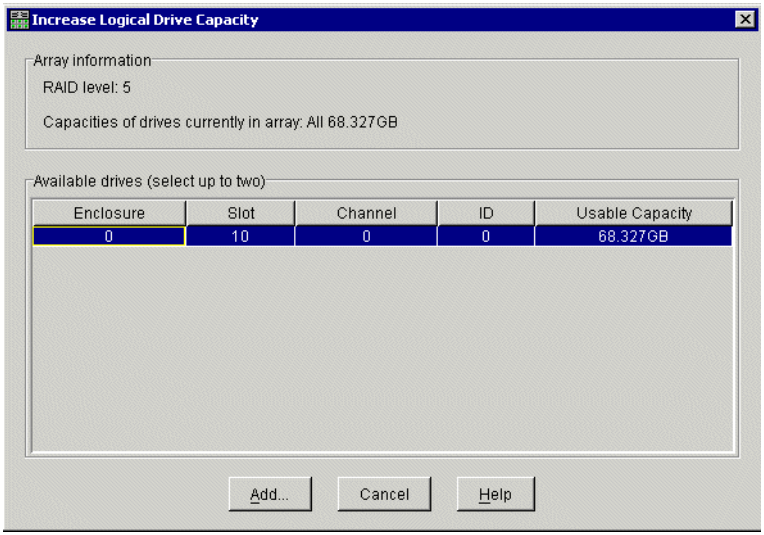

Figure 8-36 Choosing drives to add

8. At this point, you see a screen with information about which operating systems support this operation (not depicted here). This screen will be shown again during the process. Just click **OK** to continue.

9. You get the Increase Logical Drive Capacity dialog box, but now you have free space to add to the LUN, as shown in Figure 8-37.

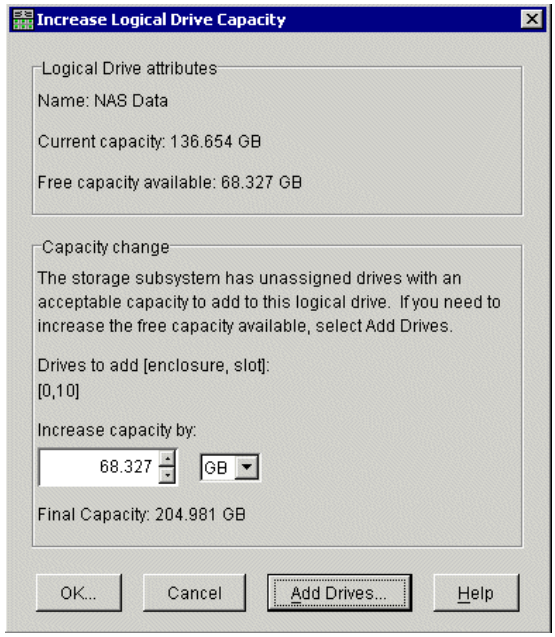

Figure 8-37 Increasing Logical Drive Capacity

10. Click **OK** to add the free space.
11. In the confirmation dialog box, type **yes**, and click **OK** to continue.
12. The process is now adding the free space. You can see in the Storage Manager Client console that there is a clock symbol near the previously selected LUN icon in Figure 8-38. This means that the operation is in progress.

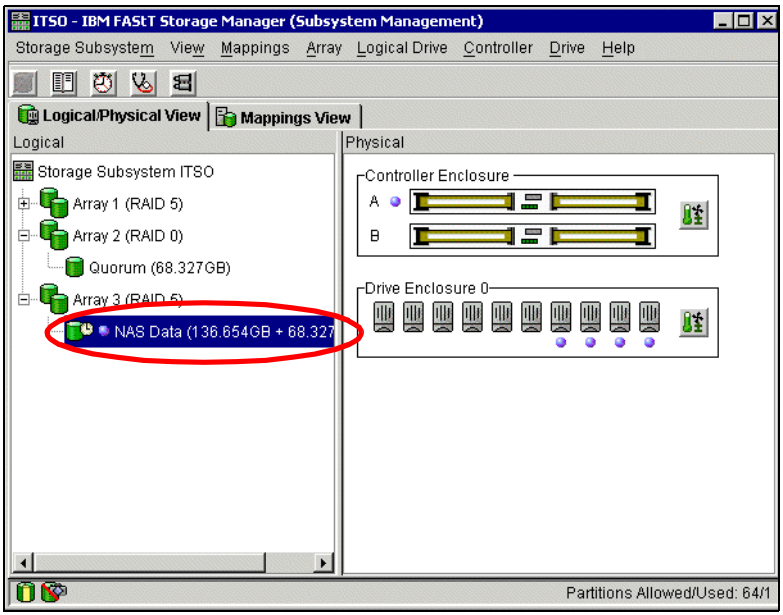

Figure 8-38 Increasing Logical Drive Capacity in process

13. This process may take some time. Once the capacity has been increased, the clock icon disappears, and the free space is added at the end of the LUN, as shown in Figure 8-39.

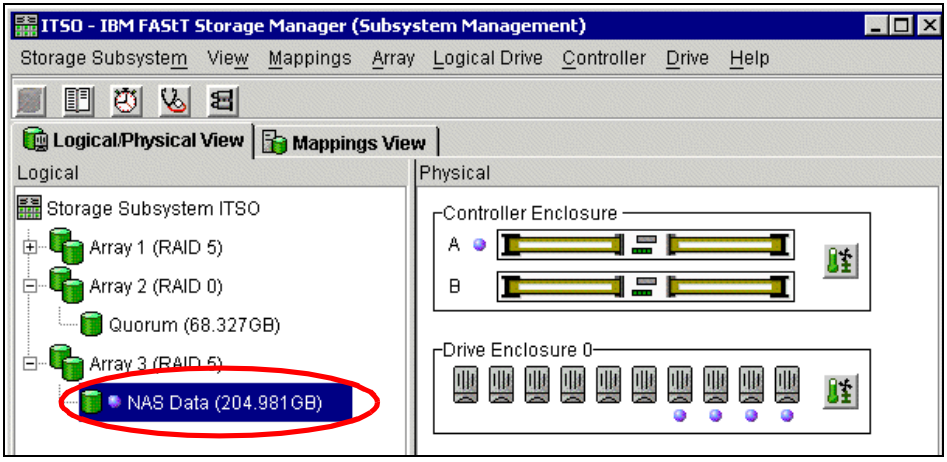

Figure 8-39 Increasing Logical Drive Capacity finished

Note: While the operation is in progress, the data in the array is accessible for users.

14. Now, you have a LUN with empty space immediately following the actual volume. To check it, just open the IBM NAS Admin tool and select **Storage —> Disk Management (local)**.

15. In our example in Figure 8-40, you can see that Disk 2 has 63.36 GB of unallocated space.

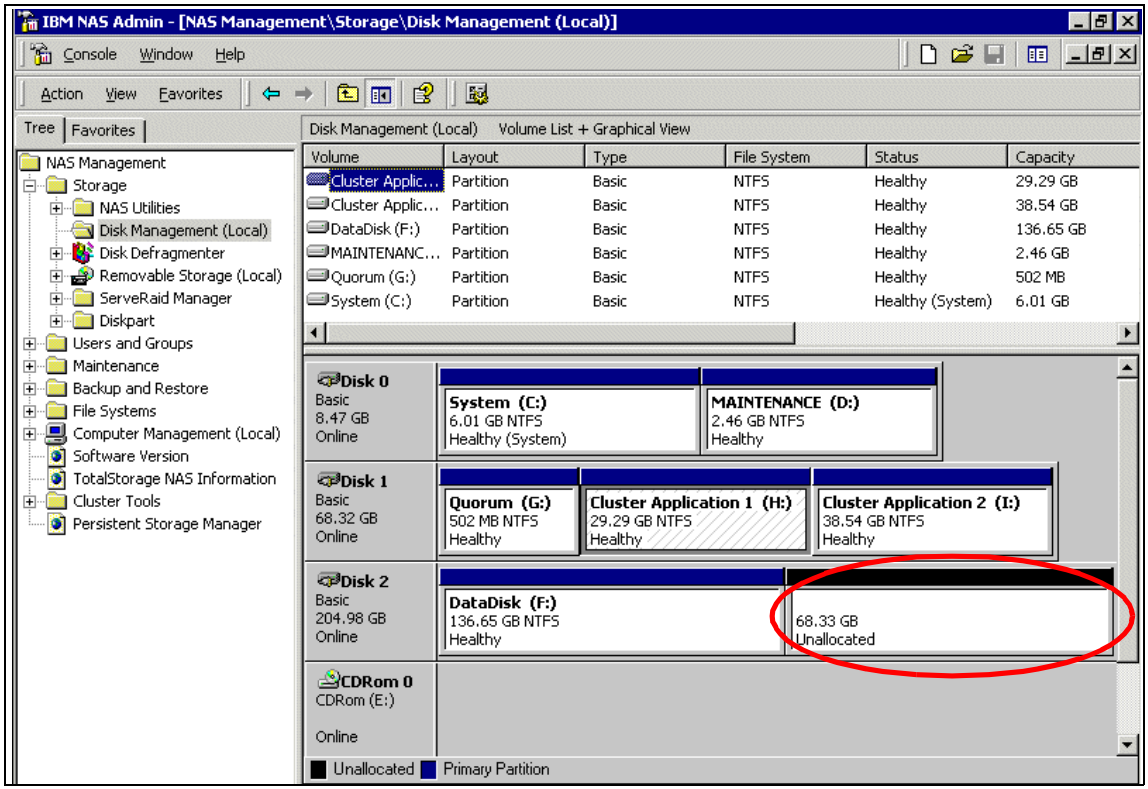

Figure 8-40 The free space is added at the end of the LUN

16. If the volume doesn't have a label, now is a good time to assign a label to the volume in order to identify it later. Right-click the volume, select **Properties**, and write the label (DataDisk in our example).

Expanding the volume

IBM TotalStorage NAS 300 takes advantage of the Microsoft Cluster Server Technology in order to provide high availability. The disks used for shared storage must be the *basic* type, since *dynamic* disks are not allowed in a Microsoft Clustered environment, therefore we will use the **`diskpart`** utility to expand Windows basic disks.

> **Note:** This utility is preloaded in the IBM TotalStorage NAS server. Additionally, it can be downloaded from the Microsoft Windows 2000 Resource Kit Web site:
>
> `http://www.microsoft.com/windows2000/techinfo/reskit/tools/`

To expand a volume, you first need to add free space at the end of the volume you want to expand. This free space is now directly behind the existing volume that is to be extended.

Check what computer is the owner of the volume you want to expand:

1. Open the IBM NAS Admin Utility and select **Cluster Tools —> Cluster Administration**.
2. Select the group where the disk is located, and see what computer is the owner of the disks. In our example, in Figure 8-41, Disk F: is owned by NAS300A.

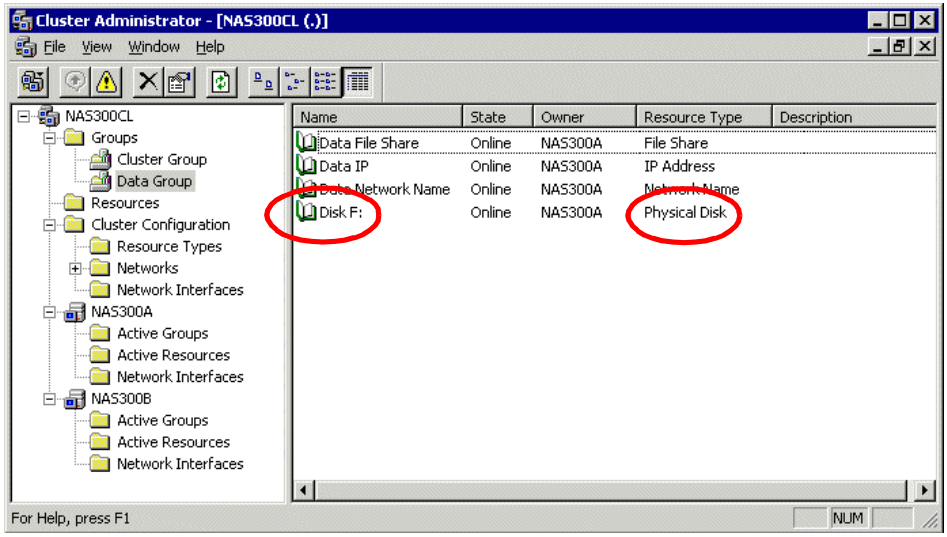

Figure 8-41 See which node is the owner

3. Shut down the other server (NAS300B). You can move all cluster resources to the other server prior to this shutdown by right-clicking the group in the Cluster Administrator and choosing **Move Group**.

Important: Stop all I/O to the disk while performing this procedure by setting offline all the resources in the cluster group that contains the disks. One way to do this is by bringing offline the cluster group that contains the disk in the cluster administration utility. And then bring online only the physical disk. This should close all open handles to the disk, as shown in Figure 8-42.

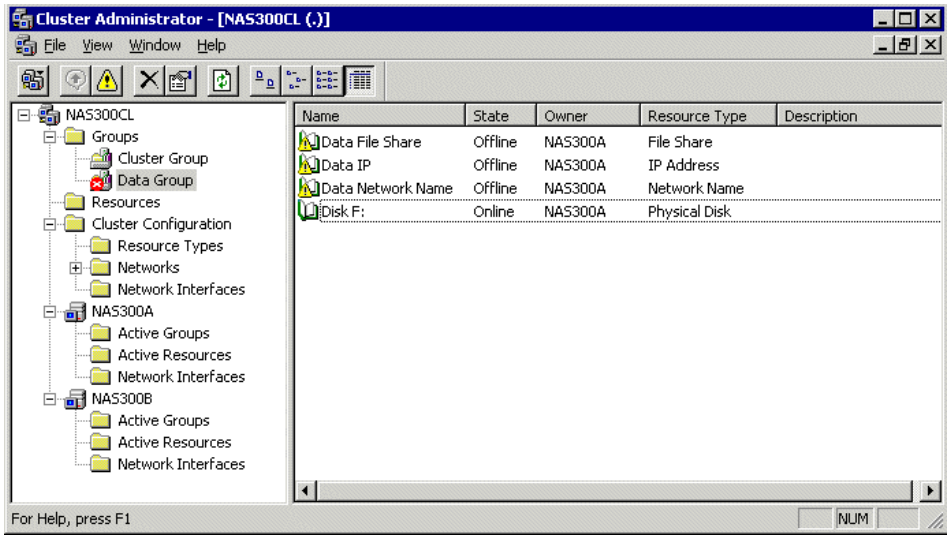

Figure 8-42 Bring offline the resources in the groups

4. Open a command prompt window and issue the **`diskpart`** command. Or in the IBM NAS Admin Tool, select **Storage —> Diskpart —> Diskpart**.
5. All the steps for this process are shown in Figure 8-43.

Attention: We recommend that you always perform a backup of your data before using the **`diskpart`** utility.

6. This opens an interactive session with the **`diskpart`** utility. Now you can list the volumes in the computer by issuing the **`List Volume`** command. Check the number for your volume by comparing the label. In our example, the volume number for “DataDisk” is 6.

7. The volume we want to extend is the volume 6 (label is DataDisk). Therefore, we need to select it with the command **`select volume 6`**.
8. Then we issue the command **`extend`** to extend the selected volume.
9. We rescan the disks with the **`rescan`** command and list again all volumes with the capacity changed.
10. Now we can exit the utility with the **`exit`** command. See Figure 8-43.

Tip: This procedure can also be executed with a script file (**`script_file_name`**) by issuing the command:

```
diskpart /s script_file_name
```

```
c:\winnt\DISKPART.EXE

Microsoft DiskPart version 5.1.3553
Copyright (C) 1999-2001 Microsoft Corporation.
On computer: NAS300A

DISKPART> list volume

  Volume ###  Ltr  Label        Fs     Type        Size     Status     Info
  ----------  ---  -----------  -----  ----------  -------  ---------  --------
  Volume 0     E                       CD-ROM          0 B
  Volume 1     C   System       NTFS   Partition   6158 MB  Healthy    System
  Volume 2     D   MAINTENANCE  NTFS   Partition   2518 MB  Healthy
  Volume 3     G   Quorum       NTFS   Partition    502 MB  Healthy
  Volume 4     H   Cluster App  NTFS   Partition     29 GB  Healthy
  Volume 5     I   Cluster App  NTFS   Partition     39 GB  Healthy
  Volume 6     F   DataDisk     NTFS   Partition    137 GB  Healthy

DISKPART> select volume 6

Volume 6 is the selected volume.

DISKPART> extend

DiskPart successfully extended the volume.

DISKPART> rescan

Please wait while DiskPart scans your configuration...
..........
DiskPart has finished scanning your configuration.

DISKPART> list volume

  Volume ###  Ltr  Label        Fs     Type        Size     Status     Info
  ----------  ---  -----------  -----  ----------  -------  ---------  --------
  Volume 0     E                       CD-ROM          0 B
  Volume 1     C   System       NTFS   Partition   6158 MB  Healthy    System
  Volume 2     D   MAINTENANCE  NTFS   Partition   2518 MB  Healthy
  Volume 3     G   Quorum       NTFS   Partition    502 MB  Healthy
  Volume 4     H   Cluster App  NTFS   Partition     29 GB  Healthy
  Volume 5     I   Cluster App  NTFS   Partition     39 GB  Healthy
  Volume 6     F   DataDisk     NTFS   Partition    205 GB  Healthy

DISKPART> exit_
```

Figure 8-43 Using the diskpart utility

11. Bring online all resources in the cluster group containing the disk by right-clicking the group and selecting **Bring Online**.

12. After this procedure, boot the server that is down and move the group that own the disk to ensure proper operation. When the server is up, perform these steps:

    a. Open the IBM NAS Admin Tool and select **Cluster Tools -—> Cluster Administration**.

    b. Select the group containing the disk. Right-click and choose **Move Group**.

13. Once the group has been moved check that the volume size has been increased. To check it just open the IBM NAS Admin tool and select **Storage —> Disk Management (local)**. You can see in Figure 8-44 that the volume size has been increased.

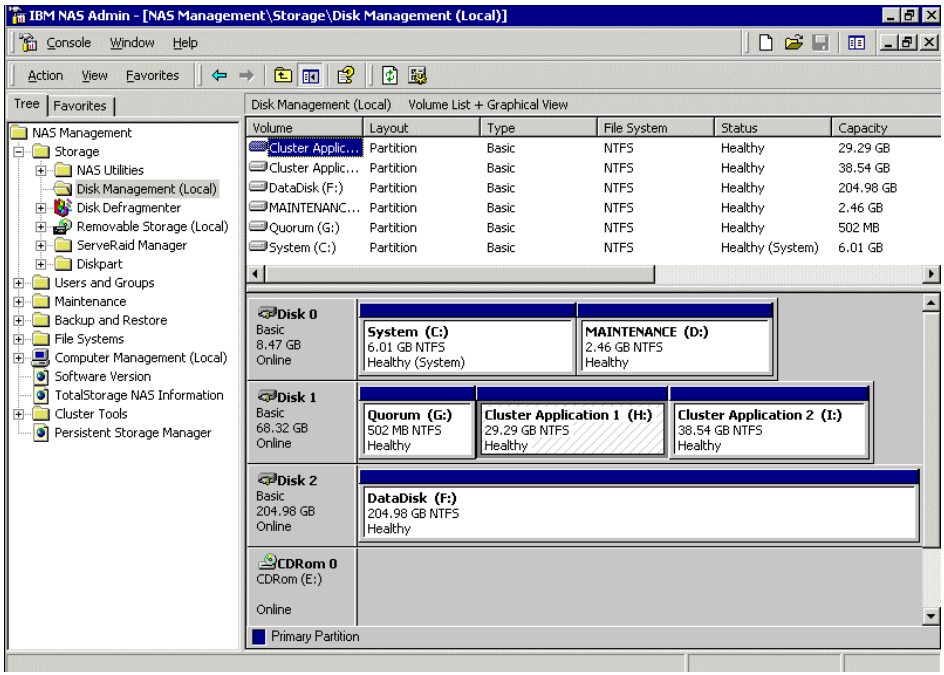

Figure 8-44 The space is finally added

8.4 Double-Take

In this section we describe Double-Take and its features. We show how to install, set up, and use Double-Take on NAS systems. Double-Take can be configured in many different ways to meet your data protection and availability plans. The following section covers the replication configuration based on two servers.

This product is not included in the current NAS release. If you are interested in these features, simply contact NSI. You can find more information at:

http://www.nsisoftware.com/

8.4.1 Description of Double-Take

Double-Take is a real-time software product that can be used for increased data availability through failover and replication of the data. By minimizing downtime and data loss, Double-Take is one more tool in your data protection arsenal. It accomplishes this with only a small amount of impact on existing communications and network resources.

Double-Take allows the user to select specific volumes and directories to replicate, in real-time, from a source server to a target server. The target server(s) can be local or remote, giving the user the ability to have off-site backups in real-time. It has the ability to move only changes in data allowing minimal impact on resources.source. There are many ways to implement Double-Take in a data protection environment, including these:

- Local data replication services
- Offsite data replication services

8.4.2 Double-Take features

Double-Take has many features which can help you to create an improved data protection environment. These features include: high availability and disaster recovery management features, resource utilization features, multi-platform features and improved backup benefits.

- High availability and disaster recovery management features:
  - Continuous replication, lowering your environment's exposure to data loss by real-time replication of file changes to target system.
  - Failover capability, which may reduce system downtime enabling businesses to continue to function.
  - Multiple configuration options, which include a one-to-one, many-to-one, one-to-many, and others across any IP based LAN or WAN.
  - Ability to configure a many-to-one disaster recovery plan, which may allow for reduced total cost of ownership.
  - Options that allow reporting and verification of data protection systems.
- Resource utilization features:
  - Selective file replication, allowing you to choose the specific directories and drives that require replication saving on valuable bandwidth.

  - File update technology, which only transmits file updates, not whole disk blocks or whole files. This may lower the bandwidth requirement and allow for more efficient use of resources.
  - Asynchronous mode of operation, allowing processing to continue even in a congested link environment.
- Multi-platform features:
  - Multi-platform consistency provides identical functionality on Windows 2000/NT and Solaris platforms. This allows for simplification of training and support. Single management of both platforms may allow for lower Total Cost of Ownership. Double-Take can be used on the following operating systems:
    - Windows 2000 and Windows NT 4.0
    - Solaris 2.51, 2.6, 7, and 8
    - Improved backup benefits
  - Double-Take allows for mirroring of open files, eliminating the need to schedule system downtime for complete backups.

8.4.3 Double-Take operations

Double-Take performs four basic types of operations: mirroring, replication, failure monitoring/failover, and restoration.

- **Mirroring:** The process of creating an identical copy of data from the source system to the target system. File attributes and permissions are retained. After initial mirroring, only subsequent changes to files are updated and transferred.
- **Replication:** This is the real-time transmission of file changes from source to target. Double-Take's replication process operates at the file system level and can independently track file changes from the file's related application.
- **Failure monitoring/failover:** This setup allows for the a target system to stand in for a failed source system. This results in user and application requests being redirected to the target machine. Failover can be automatic or manual, depending on the configuration.
- **Restoration:** Restoration is an easy way to replicate data back from target system to source system. This can be accomplished with specific data or the whole replication set. Double-Take provides a method for this with minimal configuration.

8.4.4 Double-Take Installation

Install Double-Take according to the manufacturer's recommendations.

8.4.5 Setting up Double-Take

Setting up Double-Take is fairly straightforward:

1. After you have installed Double-Take, start the application with the Double-Take Management Console. You will see a list of all the machines that have the Double-Take software installed. This is shown in Figure 8-45.

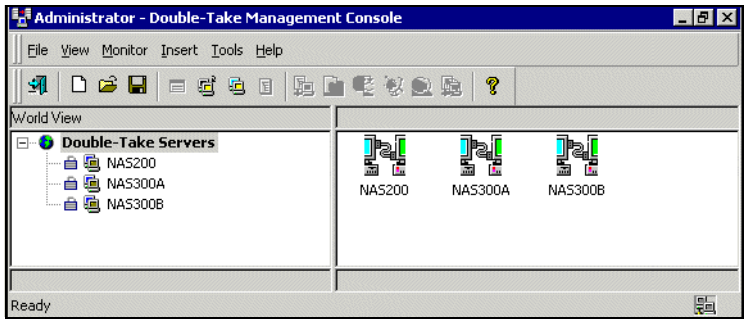

Figure 8-45 Double-Take — Server List

2. Valid activation codes are now required to enable the Double-Take software. This is done by logging into each machine, by either right-clicking each server listed and choosing logon, or by double-clicking the appropriate server. Supply a username and password as shown in Figure 8-46.

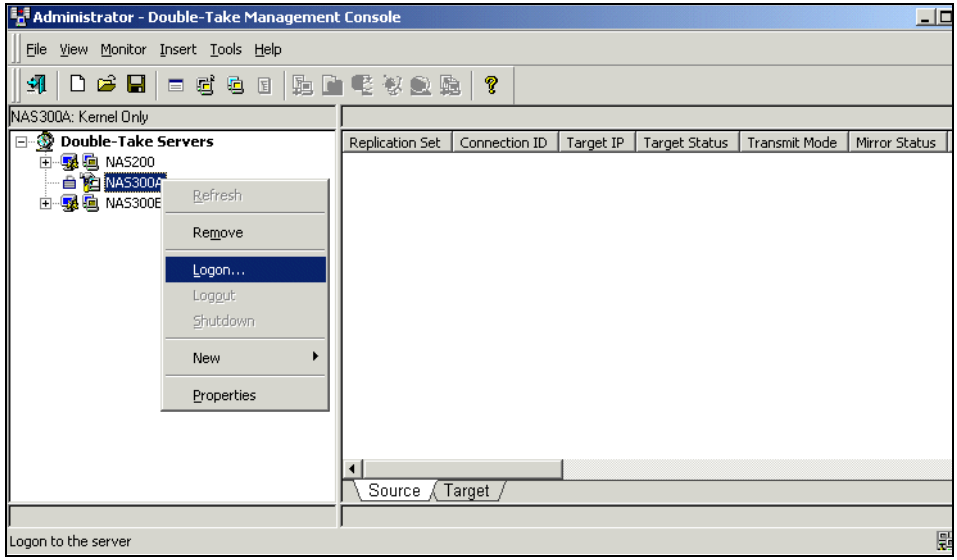

Figure 8-46 Double-Take — Logon to Server

3. Once you have logged on, you will get a message stating that you need to supply a valid activation code to enable Double-Take. This is shown in Figure 8-47.

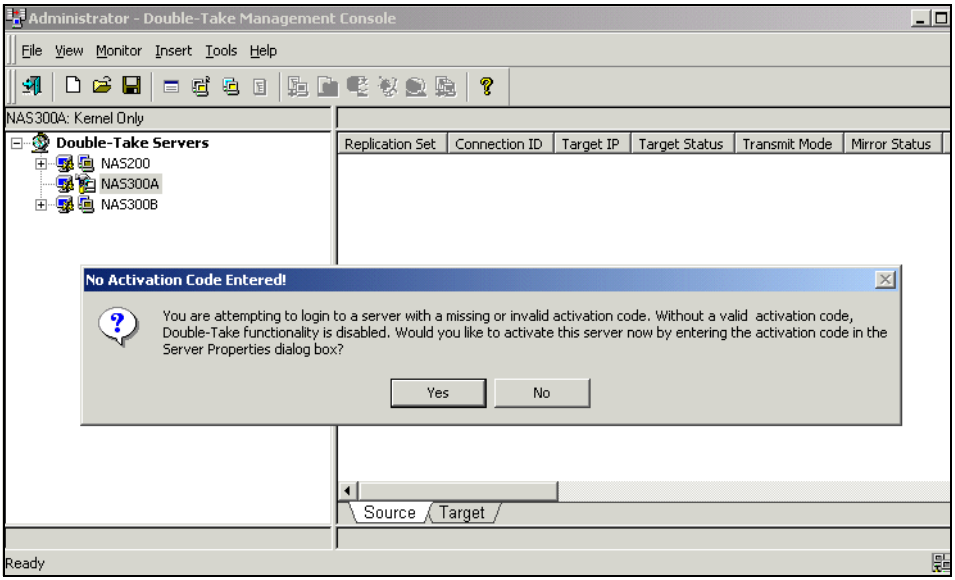

Figure 8-47 Double-Take — No Activation Code Entered

4. At this time you will need to decide on the role of the system. Is the NAS appliance going to serve as a source, or a target, or both? If you choose not to select a box, or only one of the boxes, you can always go back in later and change the settings. Select the appropriate check boxes and type in the activation code as shown in Figure 8-48. Click the **OK** button to continue.

Tip: Deselecting a check box may cause one of your replication sets to stop working if that server is part of a current replication set.

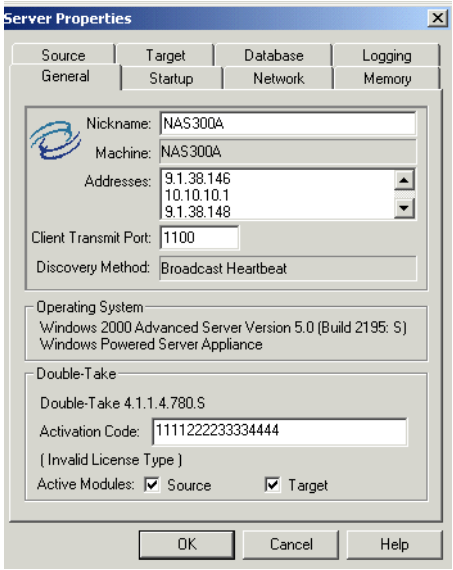

Figure 8-48 Double-Take — Activation Code Server Properties

5. Once you have entered a valid activation code for each system, if you double-click any server in the Double-Take Servers list, it will show all the available drives under each corresponding server. Now you are ready to start using Double-Take.
6. If, however, you had entered a valid activation code but did not check either the Source or Target boxes, your server will show up in the list but no drives will show up if you double-click that server. If this is the case, simply right-click the appropriate server and chose **Properties** from the menu. Make the necessary changes and click the **OK** button to continue.

8.4.6 Creating a mirror

Double-Take performs four basic types of operations: mirroring, replication, failure monitoring/failover, and restoration. We are going to demonstrate how to create a mirror between two servers:

1. The first step in mirroring is to decide which machine is going to be the source and which is going to be the target. Once you have decided this, you can bring up the connection wizard by clicking the **Tools** menu and selecting **Connection Wizard**. This can be seen in Figure 8-49.

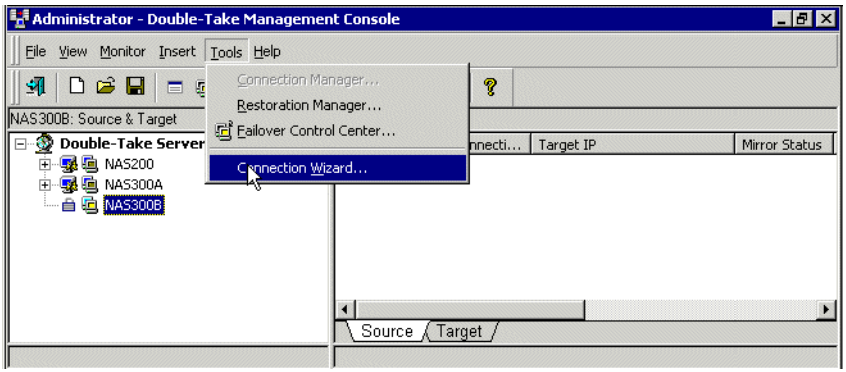

Figure 8-49 Double-Take — Connection Wizard

2. You may also bring up the wizard by dragging the source machine onto the target machine. Either way will bring up the Double-Take Connection Wizard welcome screen, as shown in Figure 8-50.

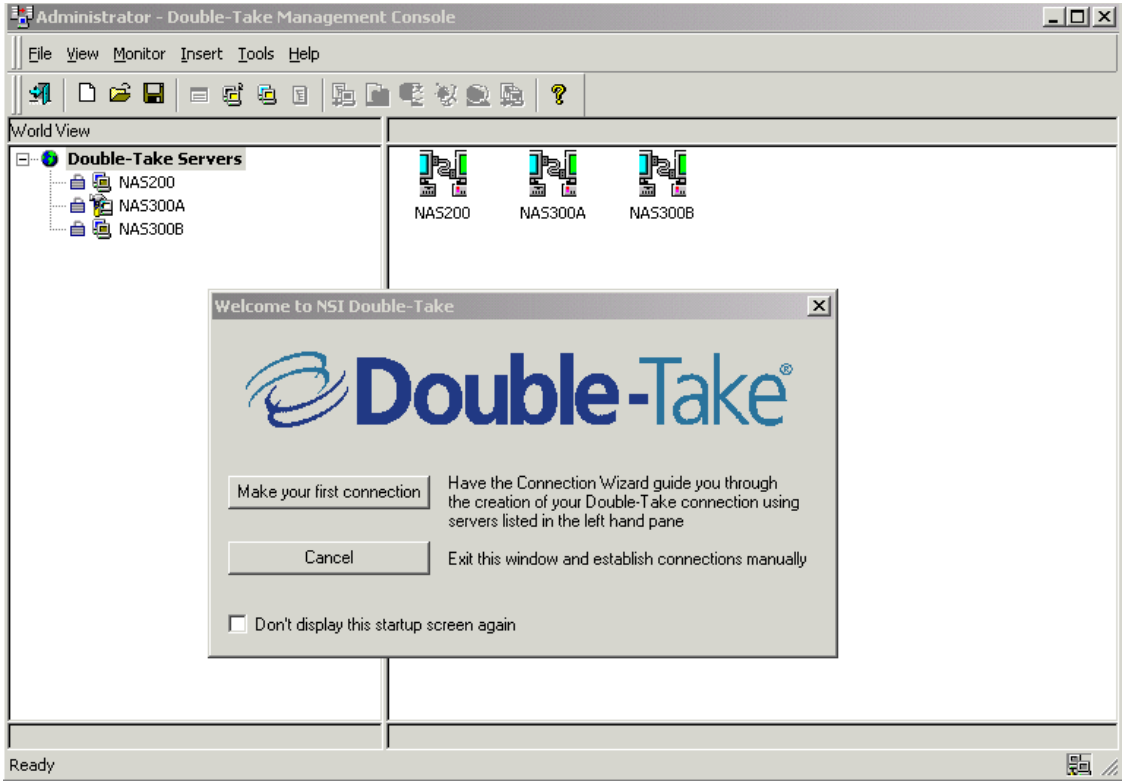

Figure 8-50 Double-Take — Connection Wizard

3. The connection wizard will take you through the process of setting up a connection with a source and target system. Click the **Make your First Connection** button and follow the wizard through the process of creating a connection between two NAS appliances.
4. The wizard will inform you that it will take you through the basic steps of setting up a connection. It will ask you what system is going to be the one with the information that needs to be protected, which will be the source device.
5. All the systems that have Double-Take software loaded, that have a valid activation code installed, and that you had checked to be a source, will show up here. Use the drop-down box to select the appropriate NAS appliance and click **Next** as shown in Figure 8-51.

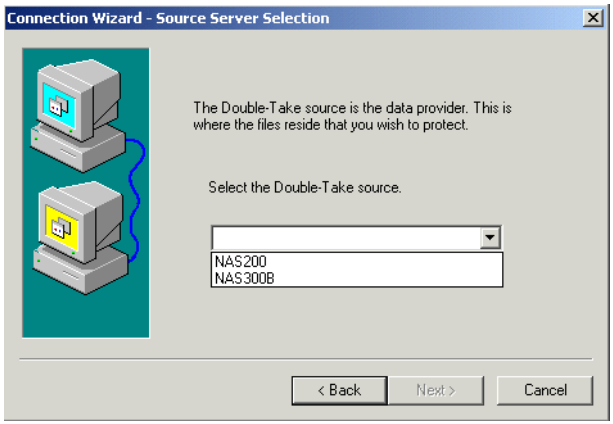

Figure 8-51 Double-Take — choosing source system

6. The wizard will then ask for the target machine, which will be receiving the data. Again, the only ones that will show up in the drop-down box will be the ones with the code installed, have a valid activation code, and you had selected to be a target system. Use the drop-down box to select the NAS appliance that will be used as the target system and click **Next** as shown in Figure 8-52.

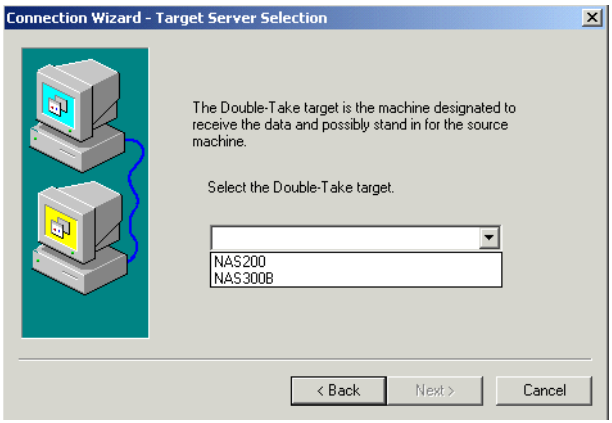

Figure 8-52 Double-Take — choosing target system

7. The wizard will take you through creating a replication set. This is a list of all the files that need to be protected. You can either create a replication set or use an existing one. As you can see in Figure 8-53, we have decided to create a new replication set and call it “Test Mirror”. When you have finished, click **Next**.

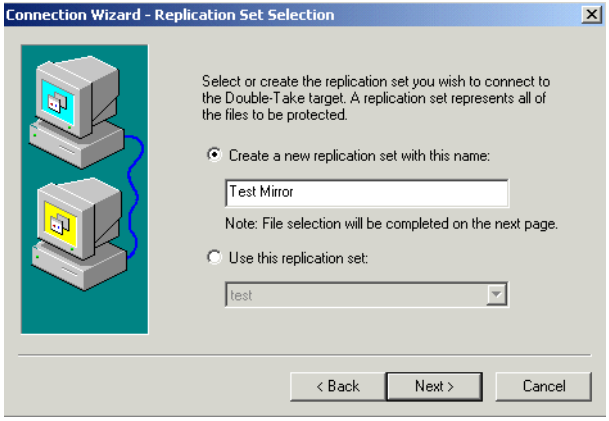

Figure 8-53 Double-Take — creating replication set

8. The wizard will ask you to select the volumes and folders that you want to protect. You are beginning to define the replication set, that is, selecting the volume(s) and directories that will be included in the mirroring process. This is shown in Figure 8-54. When you are finished, click **Next** to continue.

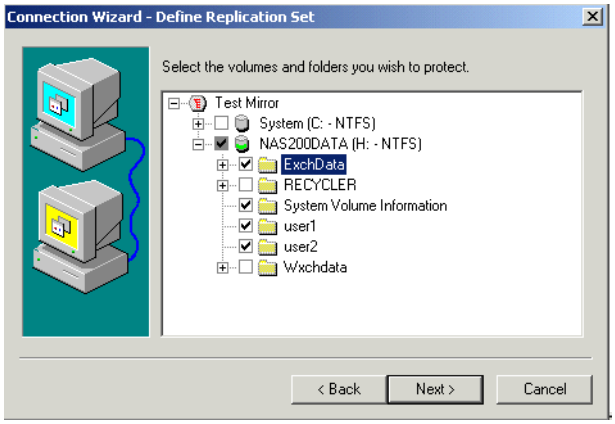

Figure 8-54 Double-Take — defining replication set

9. You have defined the specific NAS appliance and data that will be the source; now you will supply the wizard with a location on the target machine. The target system had already been defined at the beginning of the wizard. You now need to select the appropriate path. You have the option of sending all the data to a single path or the same path. This is shown in Figure 8-55.

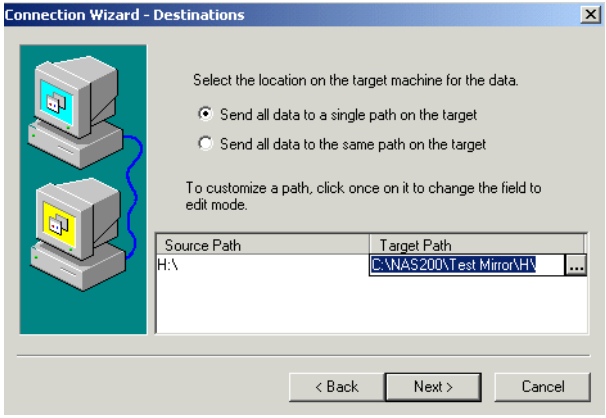

Figure 8-55 Double-Take — destination of data

10. The Target Path allows you to browse to a location. By default, the pathing will be the name of the server, the name of the replication set, and source drive. Double-click the **Target Path** and a browse window will appear to the right of the path as shown in Figure 8-56.

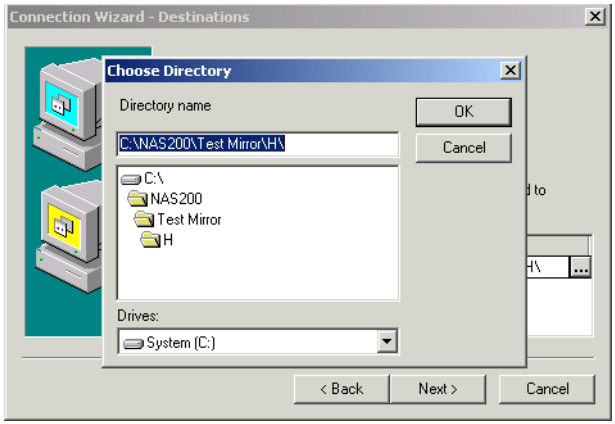

Figure 8-56 Double-Take — Destination browse window

11. You can browse to a different destination. Once you have selected the destination directory name, click **OK** and then **Next** as seen in Figure 8-57.

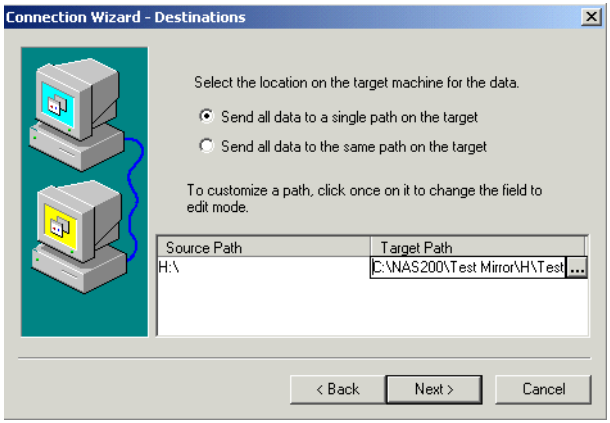

Figure 8-57 Double-Take — Destination complete

12. The wizard will bring up a final screen showing a summary of the connection information. You can continue if satisfied with the settings, otherwise click the **Back** buttons to change the appropriate setting. An **Advanced Options** button will allow you to adjust mirror, schedule, and failover options. Click **Finish** to complete the setup. See Figure 8-58.

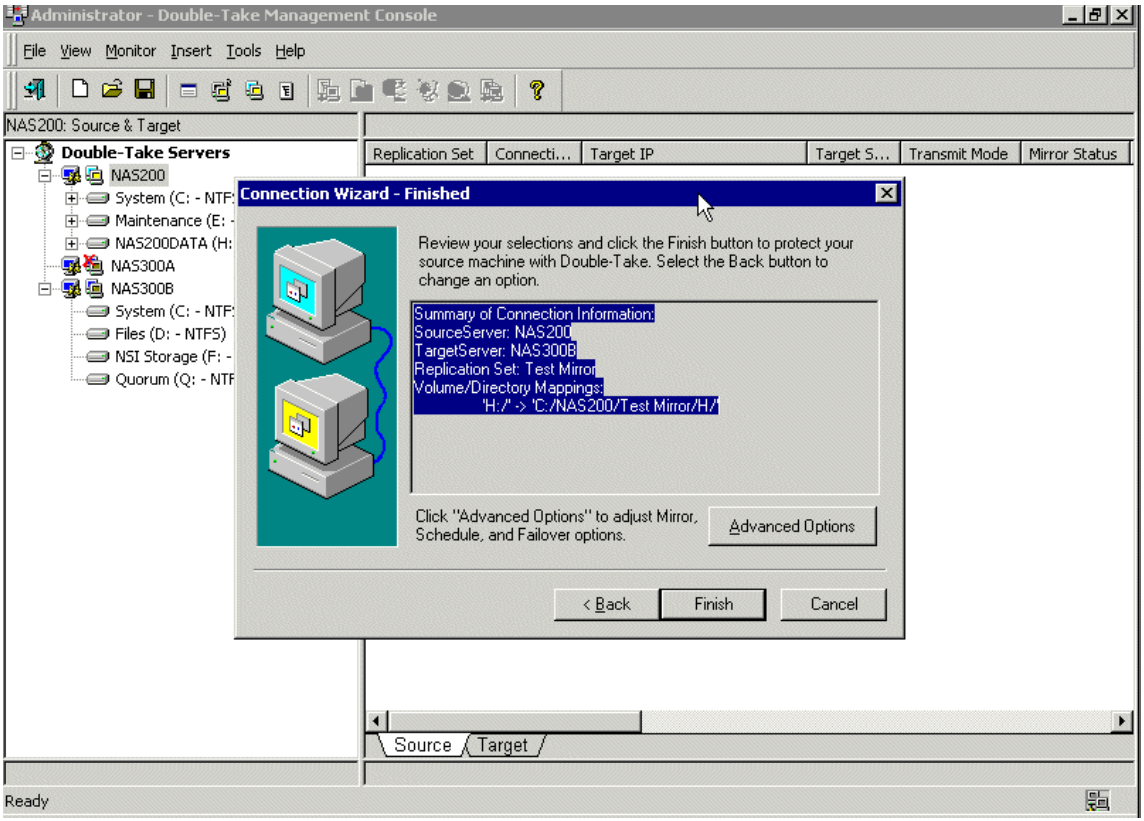

Figure 8-58 Double-Take — Connection Wizard — Finished

13. You can now check to see if the mirror is operational. One way is to double-click the source server. There will be an icon and the name of your replication set. Open up the replication set and observe the settings. You can also verify by going back to the main window and clicking the source. It will show the replication sets and other pertinent information like IP address, target and mirror status, etc. These checks are shown in Figure 8-59.

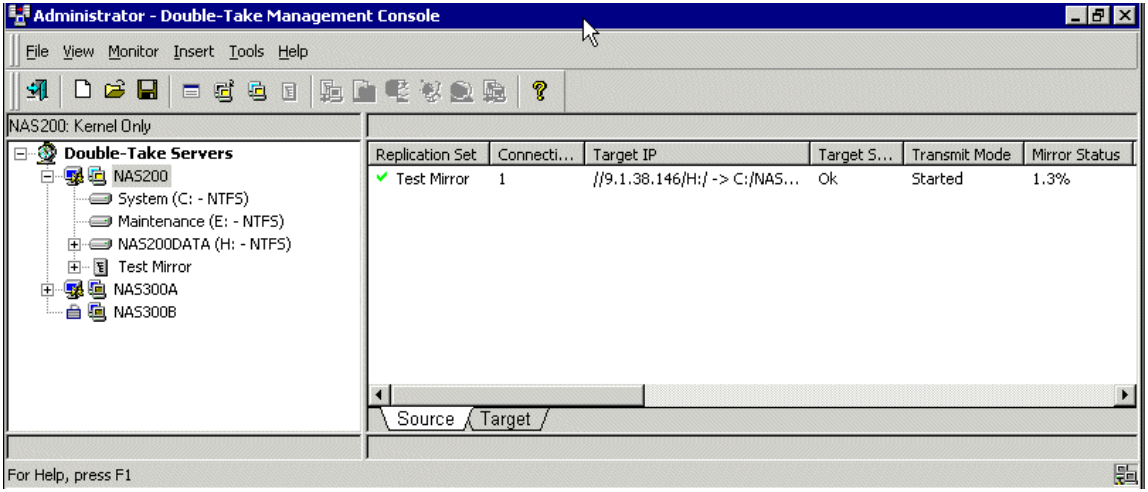

Figure 8-59 Double-Take — verifying replication set

You have just completed a mirror that will take your source data and replicate it to the target system in real-time. The IBM NAS appliances and Double-Take together are just more ways that IBM ensures your company's data availability.

8.5 Persistent Storage Manager (PSM)

Persistent Storage Manager (PSM) is a utility that creates point-in-time images of the file system on the NAS unit.

This is done by using a copy-on-write technique that uses, for each volume, an area of pre-allocated storage (the PSM cache file) that keeps only those data blocks which have been written since the time you made a persistent image of the volume.

This image then can be used to restore accidentally deleted files, corrupted data, or can even be used to back up data to another location (disk or tape). Also, this image is sometimes called Persistent True Image (PTI), because it is:

- **Persistent:** Images survive accidental or intentional reboots, file corruption, and system crashes, and are highly resistant to virus attacks
- **True Image:** Data is managed at the block level of the device, even though files and folders are presented to the users

PSM is preloaded in the IBM TotalStorage NAS devices.

PSM's most important function is its capability to create a True Image (other terms are point-in-time image, snapshot) of open files, thus eliminating the necessity to shut down applications. It does this using its component called Open Transaction Manager (OTM). With point-in-time images, you can run your backup while the system I/O continues.

8.5.1 How PSM works

PSM runs and operates below the file system as a storage filter class driver, intercepting all write actions to the NAS volumes.

When the command to create a True Image is executed, PSM begins monitoring the file system looking for a quiescent period. The quiescent period provides sufficient time for completion of writes and for the various software buffers to flush.

The premise is that, by the end of the quiescent period, a volume will be created which is in a "stable" state. This means that the volume is at rest, the caches has been flushed, and the data is consistent.

If the volume is captured in a "stable" state, then all its contents (files and folders) can be returned to a "usable condition" for user access. If quiescence is not achieved within the allocated period, the True Image will not be created

Following the quiescent period, PSM creates the persistent True Image, a virtual point-in-time representation of the volume. This True Image is presented in folders and files structure in exactly the same manner as they are presented on the source volume (see Figure 8-74 on page 274). Metaphorically, the True Image contains the data that was overwritten on the live volume from the time of the previous True Image creation to the creation of the current True Image.

In reality, at the time of True Image creation, PSM sets up junction points, to the Diff Data maintained in the PSM area. The actual creation of the True Image requires minimal resources and time.

Writing to a PSM NAS volume

When a write I/O is sent to the production volume after a True Image has been created, PSM intercepts and pauses the request, reads the data (or the blocks) that is to be overwritten, and saves the data in a Diff directory within the PSM-specific cache file. After the original data has been copied to the cache, PSM releases the new data and it is written to the live (production) volume. This process is called “copy-on-write”. It is shown in Figure 8-60.

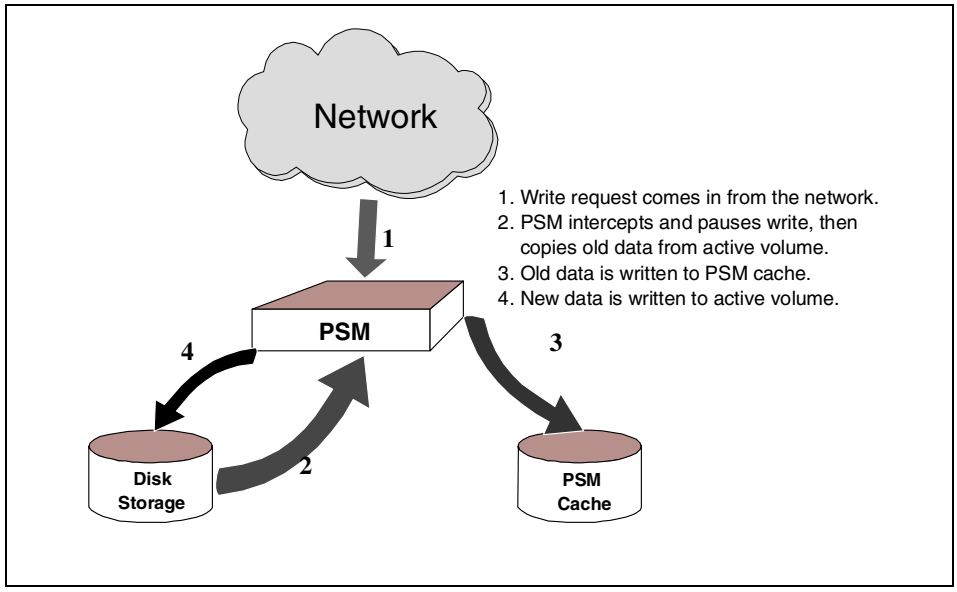

Figure 8-60 PSM’s copy-on-write process

Reading a True Image

True Images may consist of data on both the active volume and PSM cache. If this is the case, and there is a requirement to read the True Image (for backup purpose or data retrieval), PSM determines whether data has changed (data is now on the PSM cache) or is still unchanged on the active volume. Then PSM retrieves the data accordingly (whether on PSM cache or active volume) and presents it to the user. See Figure 8-61 for the process flow.

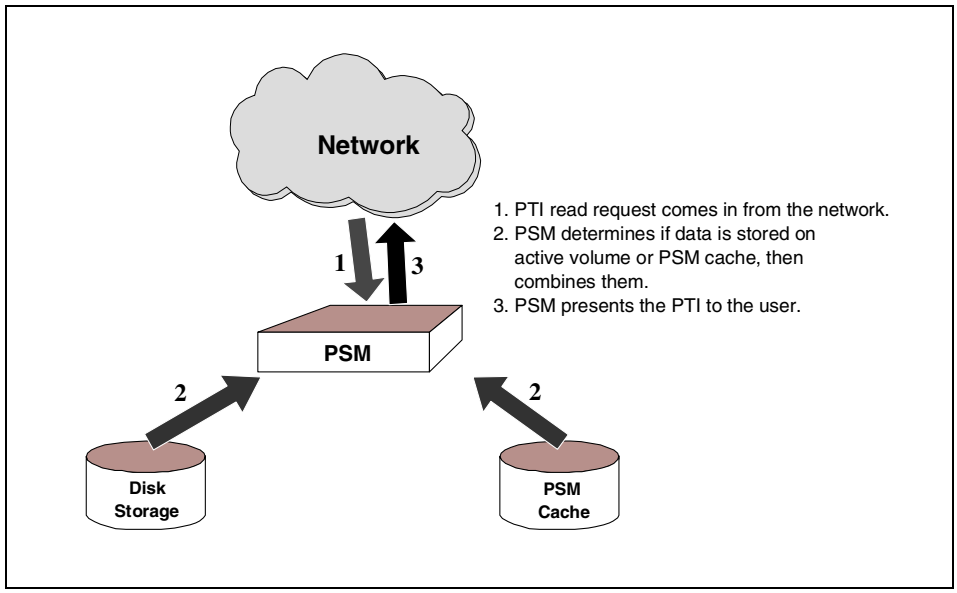

Figure 8-61 Process flow of reading a True Image

Processes such as backup or restore having access through a persistent image, have a lower priority than normal read and write operations. Therefore, if the NAS device is experiencing heavy client utilization, and at the same time a backup program is launched to access the True Image, the latter will have lesser priority for minimal performance impact.

While creating the PSM images happens very quickly, it may take a few minutes before that image is visible and available to the users. Generally, the very first image will take much longer to create than subsequent images.

Performance impact of PSM

The performance considerations for PSM can be subdivided into write performance and read performance.

Write performance
PSM creates minimal additional I/O overhead which is limited to writes. The copy-on-write process adds one read (the write is paused to read the old data from the live volume) and one write (old data is copied to PSM cache) to each write system request.

Read performance
Reads are merely affected, since typically 90% of all I/O activities are reads directly from the live volume, which causes no interaction with PSM. However, when access to True Images is required (backup, prototypes, compatibility testing), this causes interaction with PSM for data retrieval from live volume, PSM cache, or both. This, as discussed earlier, has a lower priority.

Note: PSM is designed for the main purpose of quick data retrieval, as well as creating readily available images for other functions (such as backup and development testing) even with open files (no need for application shutdown). Although it can be used for backup purposes, backup performance is not an issue PSM was designed to address.

8.5.2 Creating images with PSM

To access the PSM functions, you need to get connected to the NAS system via Web interface as seen in the next steps:

1. Open your Internet Explorer (you may also connect by using Terminal Services or locally) and use the following example, modified for your environment:

   ```
   http: //computername or ip address:8099
   ```

2. When prompted for a username and password, use the administrative account (for example, administrator, password).

3. The NAS main screen will appear, as shown in Figure 8-62.

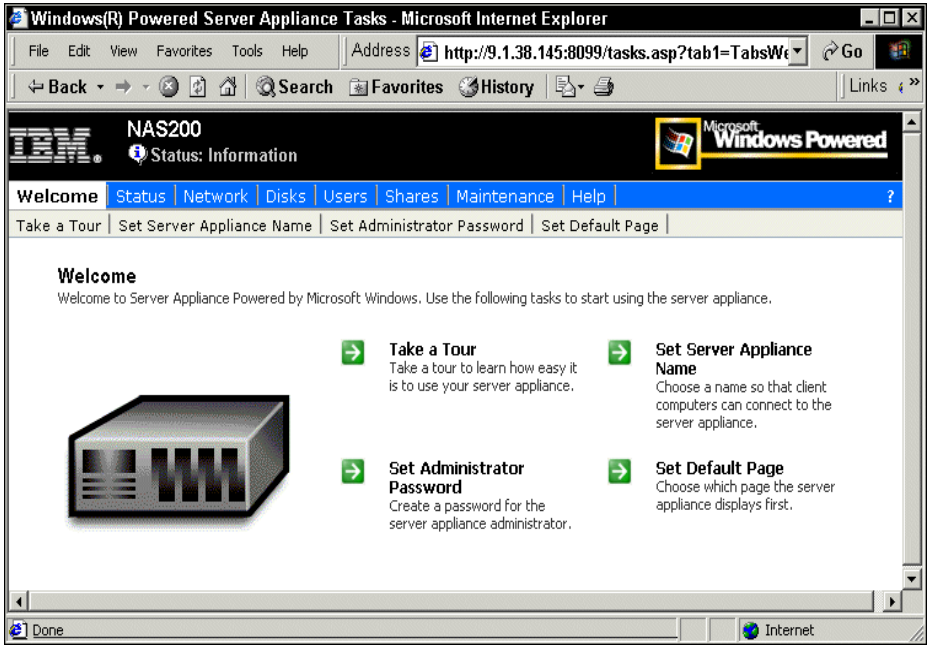

Figure 8-62 Microsoft Windows 2000 for NAS main screen

4. From the main screen, select **Disks** (Figure 8-63).

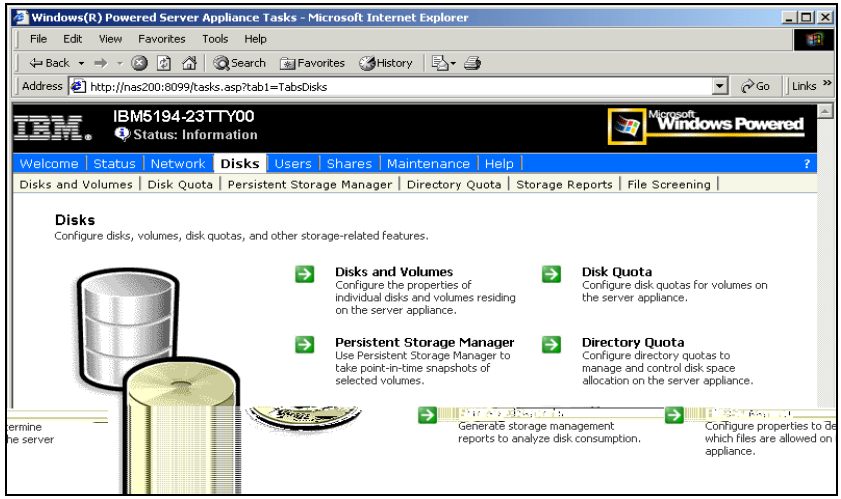

Figure 8-63 Disks screen

5. Select **Persistent Storage Manager** and you will see Figure 8-64.

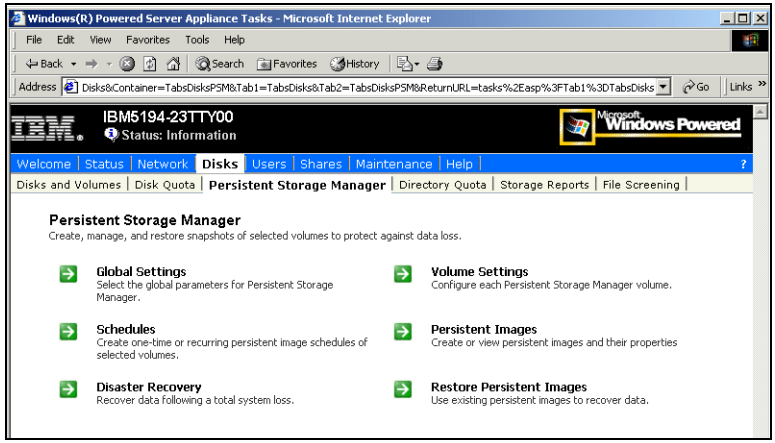

Figure 8-64 PSM main screen

Now you are in the PSM main screen ready to configure PSM.

8.5.3 Configuring PSM

Before you create images, you need to configure PSM first. The following steps are intended to guide you through the PSM configuration:

1. Configure the Global settings. From the PSM main screen, click **Global Settings** (see Figure 8-65).

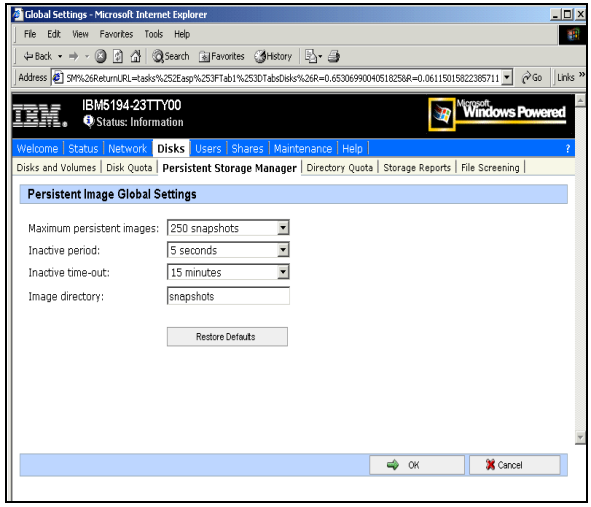

Figure 8-65 PSM Global Settings screen

Here, you can set these PSM Global Settings:

– **Maximum persistent images:**

  This corresponds to the maximum number of active images that you can create per volume.

  The default value is *250*.

– **Inactive period:**

  This is the idle time (on the volume) PSM will wait before creating a persistent image.

  The default value is *5 seconds*.

– **Inactive time-out:**

  This is the time that PSM will wait for inactivity. If the Inactive period (for example, 5 seconds) does not occur within the specified Inactive time-out (for example, 15 minutes), PSM will not create a persistent image.

  The default value is *15 minutes*.

– **Persistent image directory name:**

  This is the name of the directory that will contain the image of your volume. The default here is “snapshot”, but you can change it to any name you want.

  The default directory is *snapshot*.

When finished, click **OK** to get back to the PSM main screen.

2. Configuring the Volume settings. Click **Volume Settings:**

   This is where you can configure the specific volume attributes (see Figure 8-66).

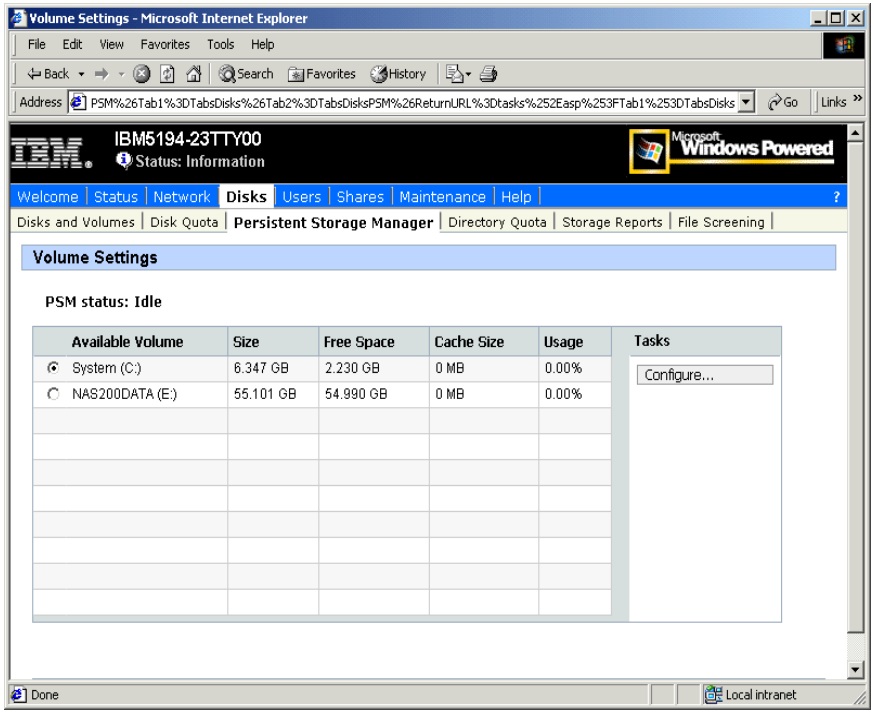

Figure 8-66 PSM Volume Settings screen

You can select a volume and configure the specific attributes by clicking **Configure** (See Figure 8-67).

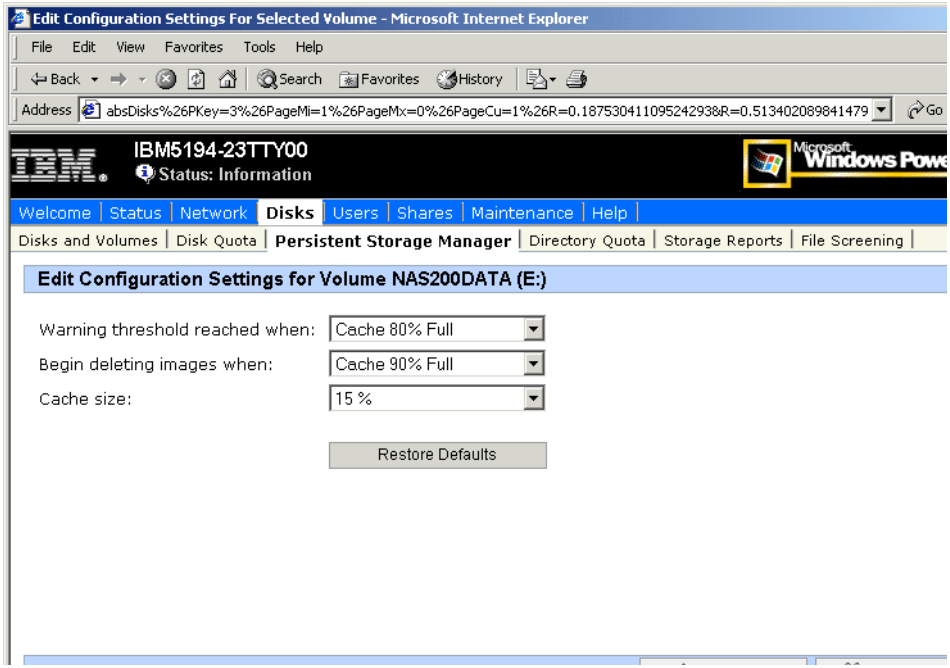

Figure 8-67 PSM attributes of a volume

– **Warning threshold:**

  This is the percentage of the cache size before warnings are sent. This is done to inform the NAS administrator that it is time to save the images before unwanted deletion of the first persistent images occurs. The logs for this option are saved in the Windows Event Log, so you can check for it using either Internet Explorer or a Terminal Services Client.

  The default value is *Cache 80% Full.*

– **Begin deleting images:**

  This is the percentage of cache size that, if reached, will begin deleting images on first in first out basis

  The default value is *Cache 90% Full.*

– **Cache size:**

  This is the size of the PSM cache allocated from the PSM volume location. Is expressed in a percentage of the volume size. Make sure that you have enough space in you volume to hold the cache file.

  The default value is *15%.*

Click **OK** to get back to the Volume settings screen, and click **Back** to come back to the PSM main screen.

8.5.4 Creating a PSM image

You have two options for creating a PSM image:

- Immediate
- Scheduled

Creating an immediate PSM image

We start by showing the steps to create an immediate PSM image:

1. At the PSM main screen (see Figure 8-64 on page 264), click **Persistent Images**. You will see the screen in Figure 8-68 that contains the already created images.

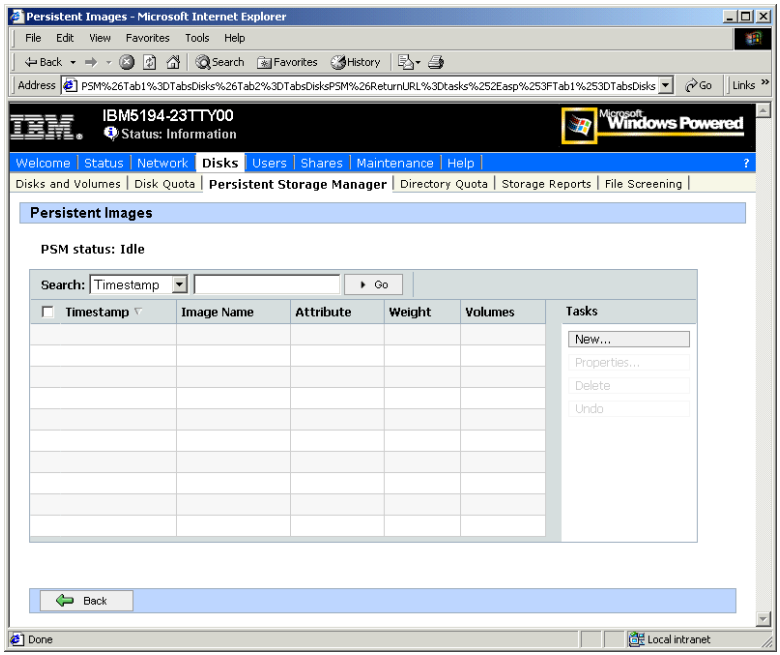

Figure 8-68 PSM Already created Images screen

2. Click **New** to see the Create Image screen (see Figure 8-69).

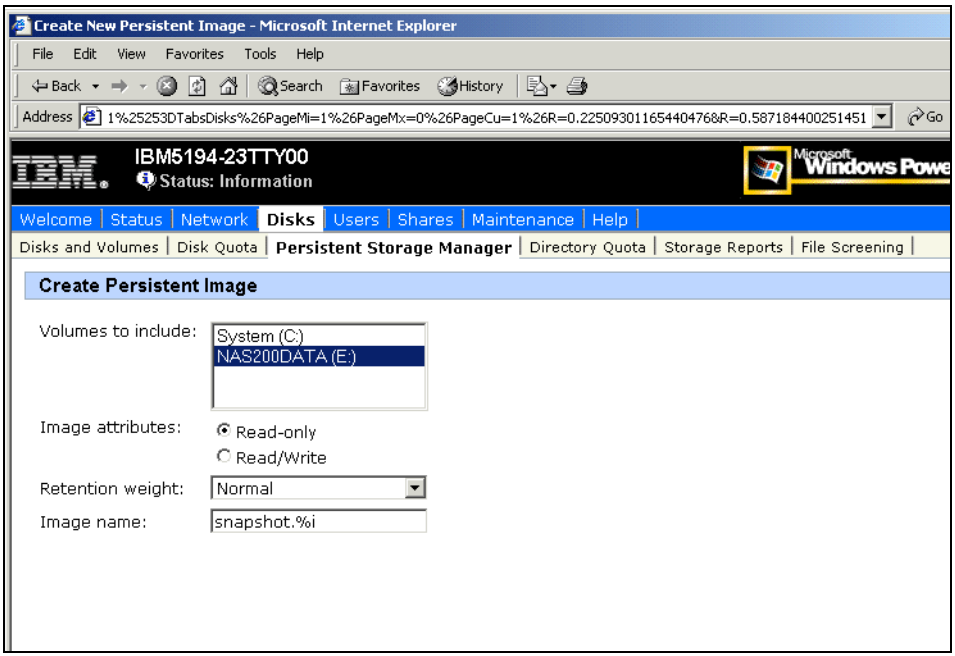

Figure 8-69 Create Image screen

3. In the *Volumes to Include:*, select the drives that you want to create an image of. For multiple volumes, press the Ctrl or Shift key while selecting the drives you want to have PSM images created of.
4. You can choose if the image has read-only attributes or if it read/write.
5. You can also give the image a relative retention weight and a name. The retention weight is important when PSM needs to delete some persistent images of a volume because the cache file for the volume has reached a certain threshold.
6. Click **OK**. You will be taken to the Persistent Image List window showing the volumes.

Note: After PSM images are created, you might have to wait for a few seconds or minutes in order for PSM to update its write-back queues and caches. In particular, the very first image will generally take much longer than subsequent images. Hence, if the system is heavily utilized, this update may take a while. After this, you should be able to access the images. One other thing to keep in mind is that — by design — PSM will run at a lower priority than regular I/O.

7. After a while, click the **Refresh** icon on your Internet Explorer. You should now see the new images you created on the list (Figure 8-70). Using this screen, you can change the properties, undo changes made to the image, or delete the image.

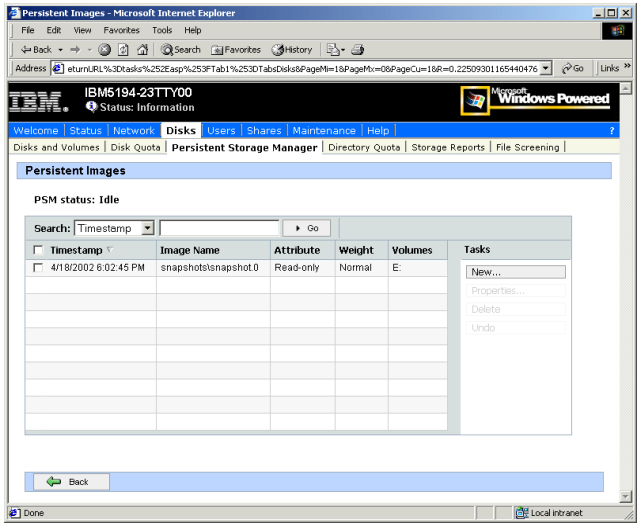

Figure 8-70 Persistent Image List screen

8. To check if the images contain exactly the same data as the volumes you selected, logon to the NAS using Terminal Services Client from the Maintenance menu on your Web Browser (or do it locally).
9. After that, open a Windows Explorer window.

   As shown in Figure 8-71, a *snapshot* directory has been created on each volume (E:) that was selected during the image creation. The mounted volumes in turn contain the directories (and files) that were in each volume at the time you created the images.

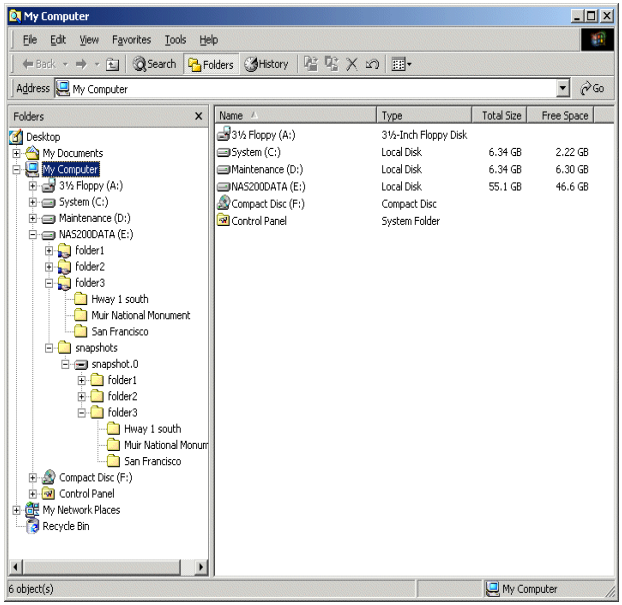

Figure 8-71 Screen showing the image created

Creating scheduled PSM images

An automated version of creating a PSM image is also available. You can schedule your job tasks so that those actions can take place during the night and after business hours. The following steps demonstrate how to create a scheduled PSM image:

1. From the PSM main screen, select **Schedules** (Figure 8-64 on page 264).
2. On the Persistent Image Schedules screen, click **New**.
3. On the Create Persistent Image Schedules screen (Figure 8-72), select the entries for the following by clicking the pull-down arrow.
   - **Start at:**
   - **Repeat Every:**
   - **Begin:**
   - **Volumes to Include:**
   - **Image Attributes:**
   - **Retention Weight:**
   - **Number of images to save:**
   - **Image name:**

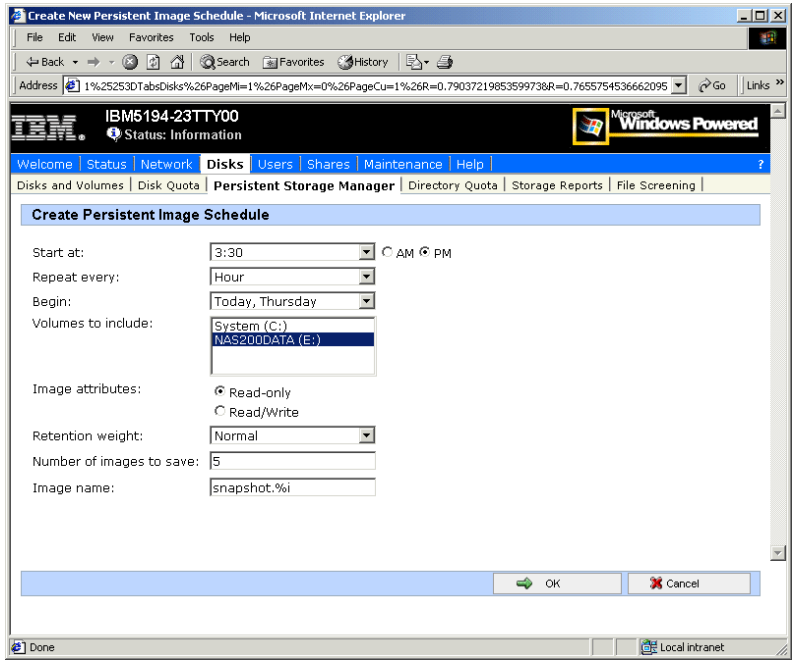

Figure 8-72 Screen for creating a new scheduled persistent image

4. Click **OK**.
5. You reach the Scheduled Persistent Images window showing the volumes, time and date, and repetition you selected earlier (Figure 8-73). In this screen you can also change the properties of the scheduled Image or delete it.

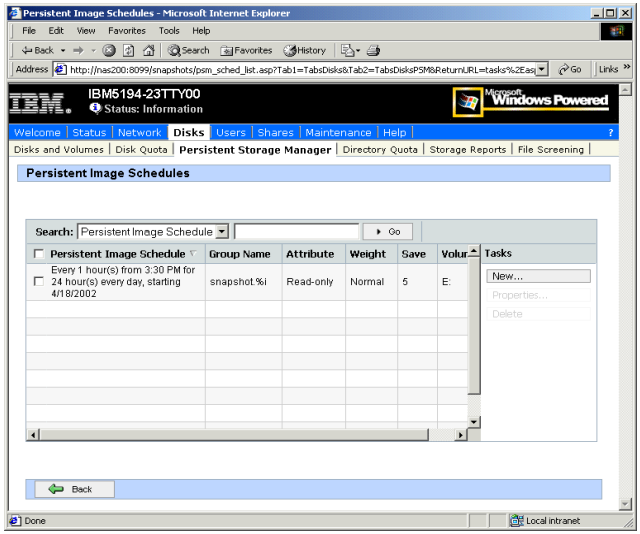

Figure 8-73 Screen showing scheduled persistent images

8.5.5 Restoring a Persistent Image

If you need to use the information stored in an image, you have two choices:

- File system access
- Restore the complete image

File System Access

You can access the files stored in the Persistent Images just as any other file in your system. See Figure 8-74.

1. Open the Windows Explorer.
2. Go to the Persistent Images directory (*snapshot* in our example).
3. Choose the image you want to use (*snapshot.1* in our example).
4. Choose and work with the file as usual, browsing with the Windows Explorer and finding the file.
5. You can now do the following:
   a. Drag-and-drop it on to the volume to replace the actual file.
   b. Edit the file directly, if the image was created as read/write.
   c. Edit the file after copying it to the volume, if the image was created as read-only.

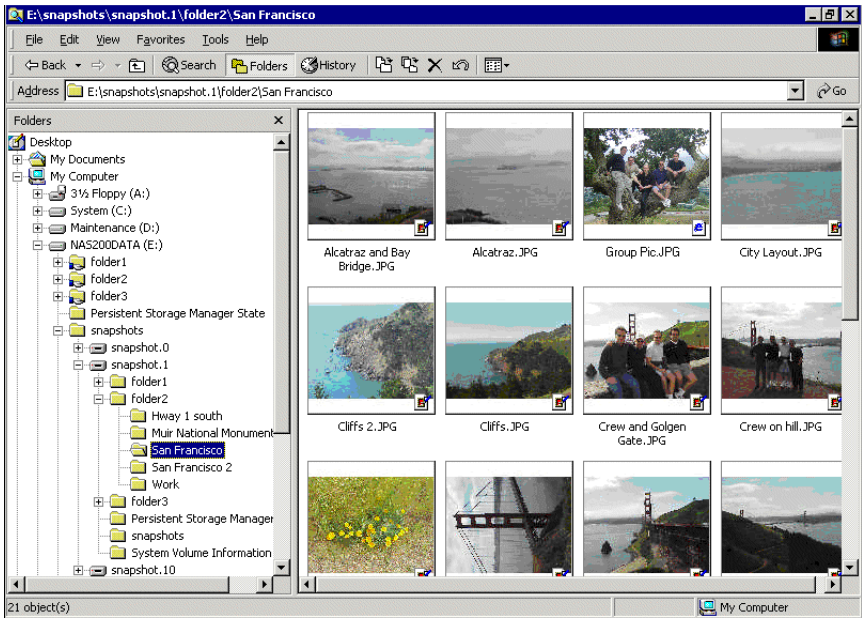

Figure 8-74 Using files in a Persistent Image

Restore the complete image

You restore the complete volume image by clicking Restore Persistent Images in the PSM main screen.

1. In the Persistent Images to Restore Screen, you can choose the Image to be restored (see Figure 8-75):
   a. Select the Image you want to restore.
   b. Click **Details** to see more information.
   c. Click **Restore** to restore the image.

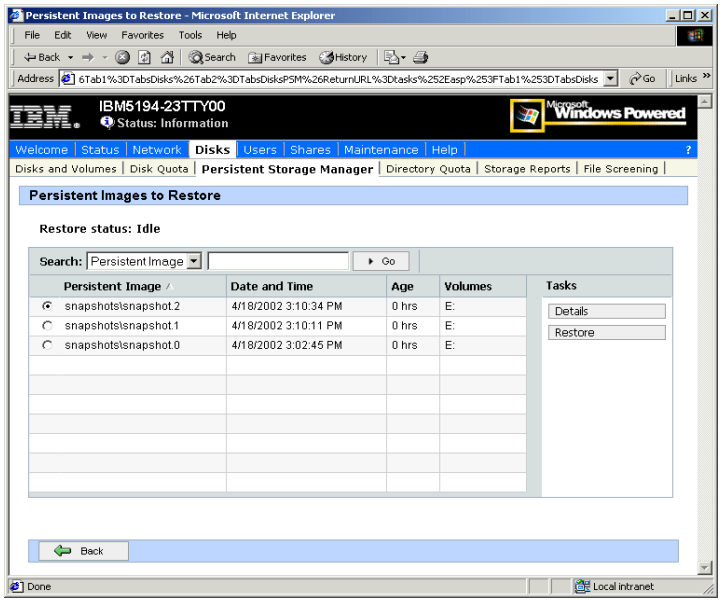

Figure 8-75 Choose the Persistent Image to restore

2. Click **OK** in the confirmation screen.

Now you have successfully restored the Persistent Image.

8.5.6 Disaster Recovery with PSM

In the event that you need to recover the operating system volume from the Recovery CD all systems settings have to be recreated. PSM provides a procedure for backing up the system partition in a network share or local disk. This procedure also creates a boot diskette to boot the machine, and recovers the system partition from the image (unattended):

1. To start the process, click **Disaster Recovery** in the PSM main screen. You are now in the Disaster Recovery screen (see Figure 8-76). This is an informative screen with the current status of the disaster recovery tool. You can start an immediate backup, create a boot diskette, or configure the properties.

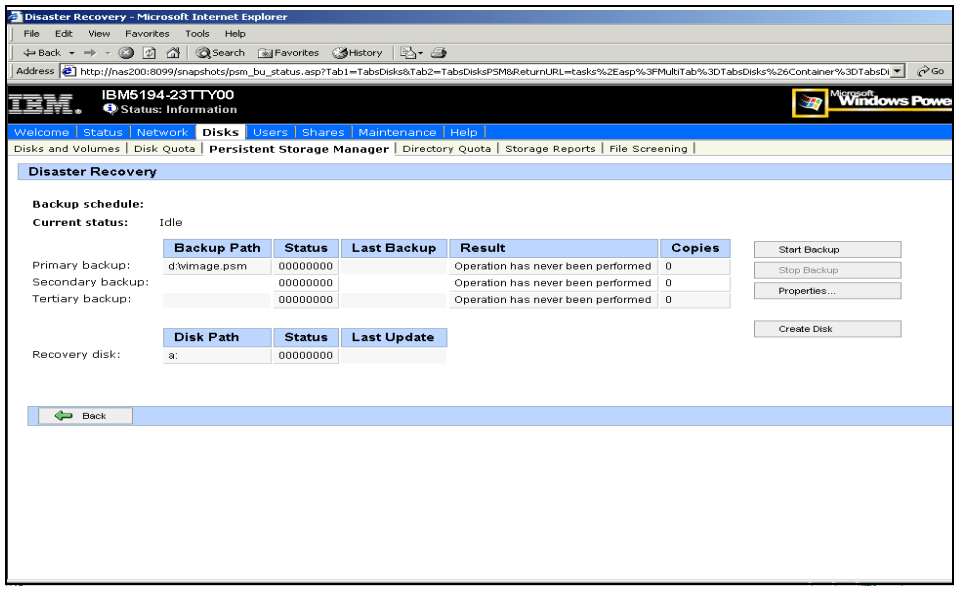

Figure 8-76 PSM Disaster Recovery screen

2. Click **Properties** to configure the settings. Now you are in the PSM Disaster Recovery Properties screen, as shown in Figure 8-77.

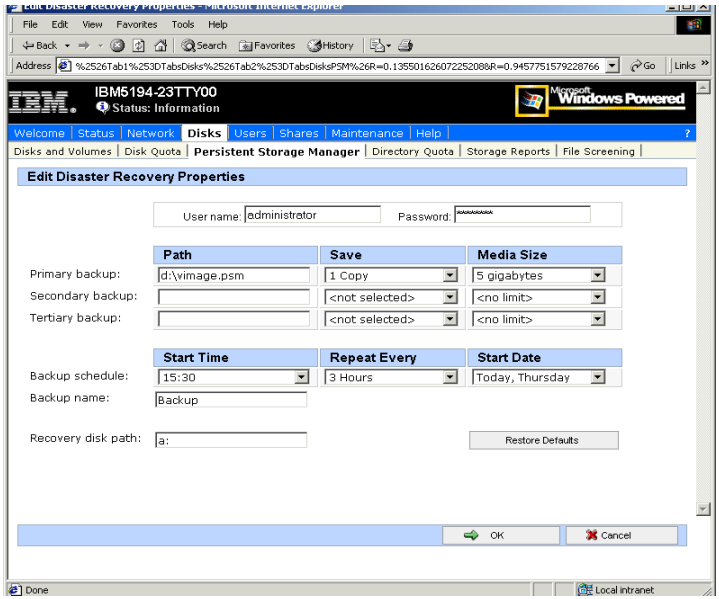

Figure 8-77 PSM Disaster Recovery Properties screen

3. Provide the parameters needed, such as these:
   - Location of backup (up to three sites, including network shares)
   - Number of copies for each site
   - Size limit for each site
   - Settings for scheduling the command
   - Backup name
   - Username and password that attaches to the network shares during a system backup or disaster recovery
4. Now you can click **OK** and come back to the PSM Disaster Recovery screen.
   a. Click **Start** to start the backup.
   b. Click **OK** again in the confirmation screen.
5. You can monitor the progress of the image creation in the PSM Disaster Recovery screen, as shown in Figure 8-78.

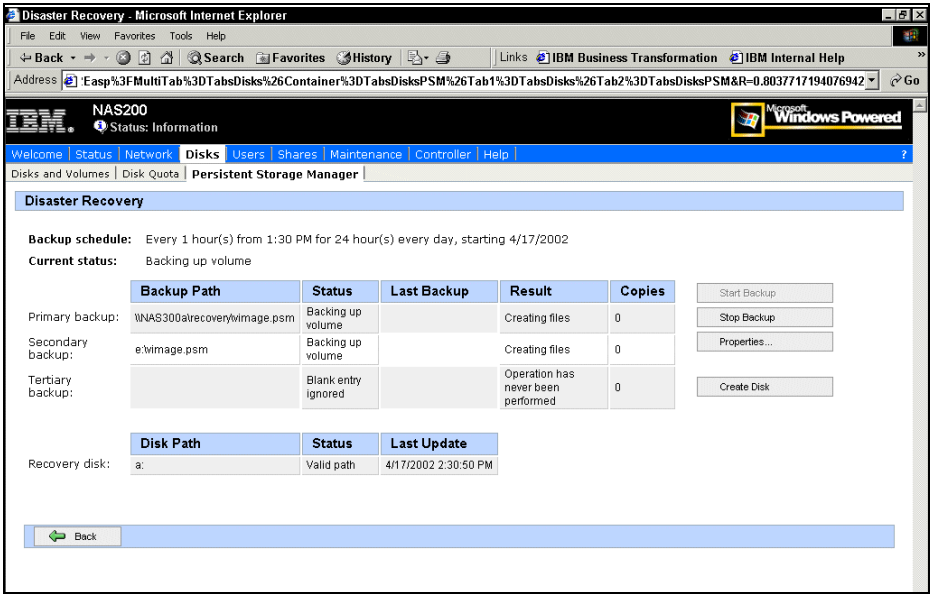

Figure 8-78 Backing up Disaster Recovery Image

6. When the image is copied, you can see the results in the PSM Disaster Recovery screen in Figure 8-79.

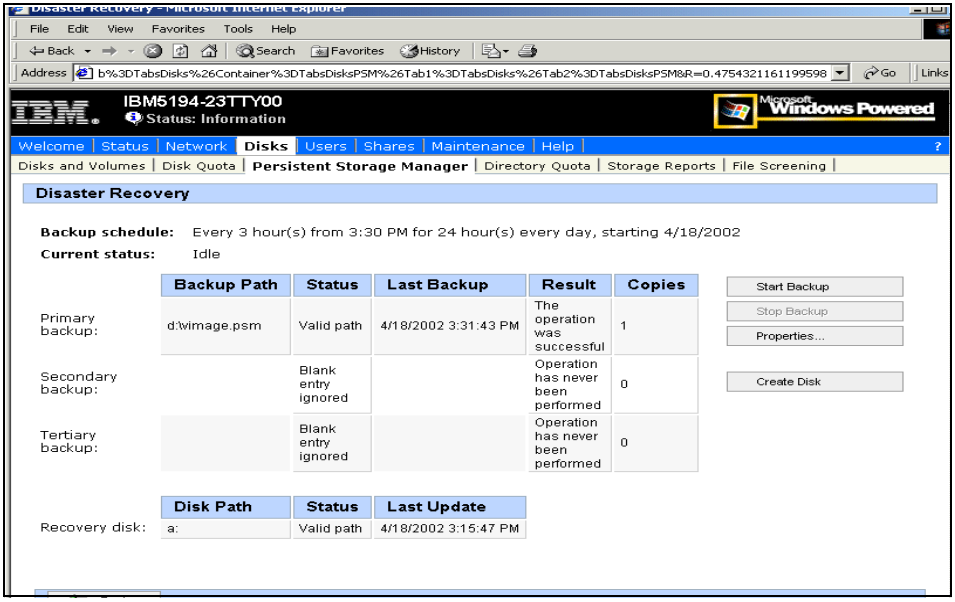

Figure 8-79 PSM Disaster Recovery Image created

7. You should now create the boot disk:
   a. Insert a formatted floppy disk in the diskette drive of the NAS device.
   b. Click **Create Disk** in the PSM Disaster Recovery screen.
   c. Click **OK** in the confirmation screen. This can take some time.
   d. When finished click **Back** to get back to the PSM Disaster Recovery screen.
   e. To make the disk bootable, run the **`fixboot.bat`** file on the floppy disk.
8. The process to recover the system volume from a disaster recovery image is straightforward. Just boot the NAS appliance with the boot diskette inserted. The recovery process starts automatically and will try the first given path that contains a valid image and load it. After that you only have to reboot the machine.

Restriction: Restoration of a PSM backup image over the network is not supported for the Gigabit Ethernet Adapter. If you have only Gigabit Ethernet adapters installed, it is recommended that you perform PSM backup of each node to its maintenance partition (D: drive), which would allow you to recover if the system volume is corrupt and/or unbootable.

Attention: Be aware that the partition D, the location were you store the image, has to use the FAT filesystem. Otherwise, if you boot with a diskette that does not include a NTFS driver, you will not be able to read the image.

More information about the FAT file system and NTFS can be found at:

http://www.sysinternals.com/ntw2k/freeware/ntfswin98.shtml

8.5.7 Special considerations for PSM in a clustered environment

This section reviews some special considerations to take in account when PSM is used in a clustered environment like in the IBM TotalStorage NAS 300.

The use of PSM in clustered volumes is supported, and there are no big differences with PSM standard volumes, as explained in 8.5.2, “Creating images with PSM” on page 262. However, the following points are worth noting:

- Creating and accessing a snapshot is the same procedure as for unclustered systems.
- When failing over a disk from one node (or after a failure), the other node mounts the volume and all the snapshots in the volume. After that the snapshots are accessible via the other node.
- After a failover, there is a delay between the moment the volume is accessible until the snapshots are accessible that vary depending on the number and size of the snapshots.
- In the event of a failover, PSM has to load the images before the volume can be brought online on the node. The default time-out for pending state in a cluster is 180 seconds. It is recommended to change this value to a greater one to avoid that the disk is set offline, and not accessible for either node.

  This is done by opening the following: **IBM NAS Admin Tool —> Cluster Tool —> Cluster Administrator**. Select a group with clustered disks, and in the right panel you can see all the resources for this group. For each resource listed, do the following:

  a. Right-click the resource name and select **Properties**.

  b. Select the **Advanced** tab.

  c. Change the **Pending timeou**t value to 1200 (seconds).

  d. Click **Apply**, then click **OK**.

  e. Repeat for all groups with clustered disks.

- Scheduling a Persistent Image is always done on one of the nodes. If the disk is moved to the other node, the image is not created, and an error message is generated as shown in Figure 8-80.

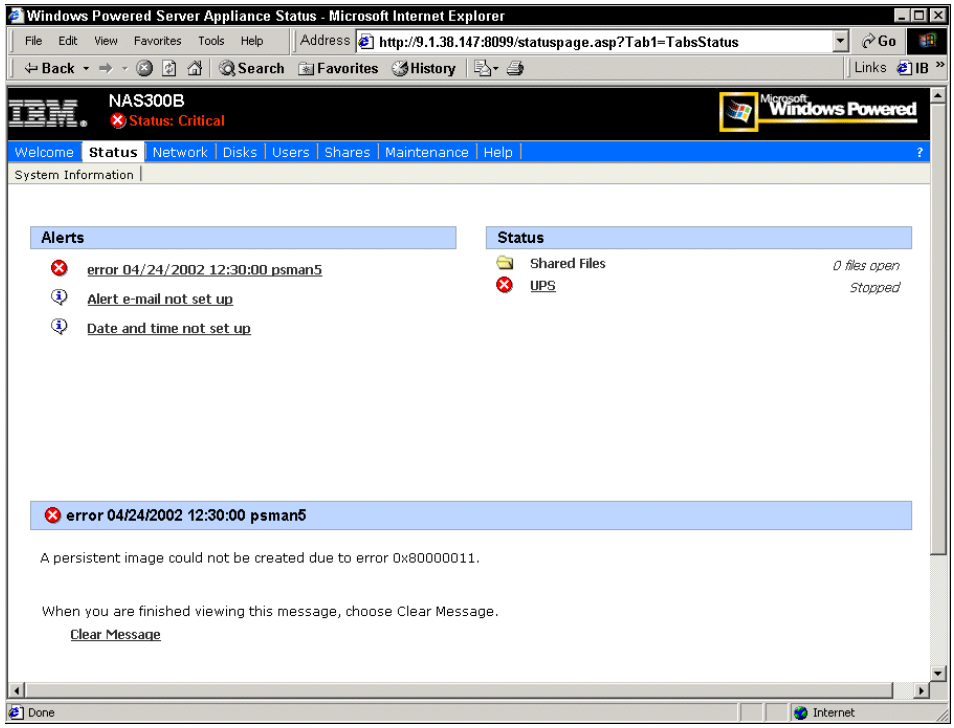

Figure 8-80 PSM error when failover a disk

Tip: If you want to assure that the image is created no matter which node owns the disk, configure the scheduled images on both nodes. Now you are sure that the image is generated on one node and that the error message is generated on the other one.

8.6 Additional tools

IBM provides you with a range of tools to help manage and maintain your IBM TotalStorage NAS device. Some tools (like the IAACU, NSI Double-Take, and CDP Persistent Storage Manager) have already been described in this chapter. In the following sections, we provide information on some additional tools.

8.6.1 IBM World Wide Name utility

The World Wide Name utility displays to the user the World Wide Name (WWN) of the IBM Host Bus Adapters (HBA) resident within the appliance. This utility is accessible via an MMC Snap-in located within the IBM NAS Admin Tool. The WWN — which is used in a similar way to MAC addresses in an Ethernet network — is a unique address provided by the manufacturer of the Fibre Channel adapter. Whenever you configure SAN devices (for example, zoning or LUN masking), the WWN is used to identify individual HBAs and storage devices. As the IBM TotalStorage NAS 300 is shipped with integrated SAN devices, the WWN utility can be used to support administrative staff in managing and troubleshooting the system. See Figure 8-81 for an example of the utility.

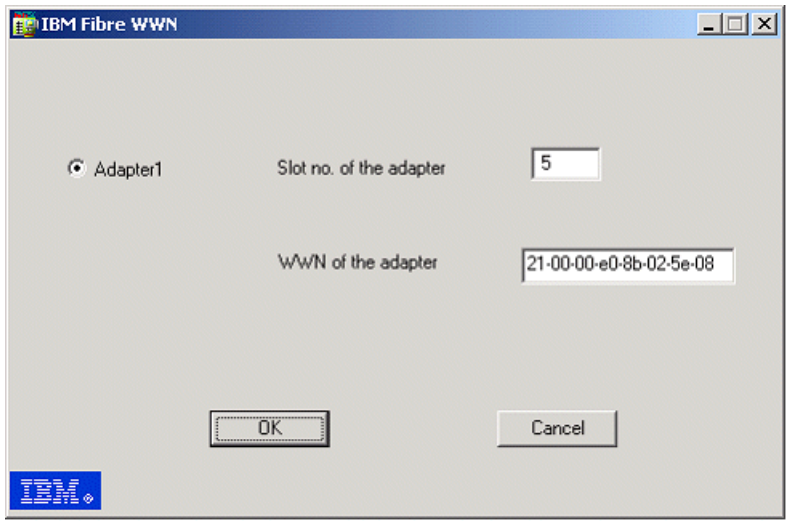

Figure 8-81 IBM WWN utility

8.6.2 IBM FAStT Management Suite Java

The IBM FAStT Management Suite Java (MSJ) Diagnostic and Configuration Utility is a network capable application that connects to remote systems. You can use FAStT MSJ to manage and diagnose one or more IBM FAStT Host Bus Adapters that are installed on a local or remote computer. Also, FAStT MSJ can be used in conjunction with problem determination procedures (PDPs) on new and existing installations to assist in diagnosing Fibre Channel (FC) subsystem configurations (in adapters and devices that are attached to the FC fabric or loop). This application is accessible via an MMC Snap-in located within the IBM NAS Admin tool. See Figure 8-82 for an example of this tool.

Figure 8-82 IBM FastT Management Suite Java — HBA View

8.6.3 IBMSNAP

IBMSNAP is used to create a snapshot of the specified drive letter (on a separate unused drive letter), call the specified batch file which runs a backup utility to perform a backup of the snapshot drive letter, delete the snapshot, then exit. IBMSNAP writes the drive letter to the file IBMSNAP.DAT, which the batch file will call to perform the backup. The tool is described in more detail in 11.2.1, “IBMSNAP utility” on page 433.

8.7 Uninterrupted Power Supply support

Integrated within the appliance is the support for IBM Uninterrupted Power Supplies (UPS). Uninterrupted power supply (UPS) provides emergency backup power for a specific period of time when the local power fails. This power comes from batteries housed within the UPS. High-performance surge suppression protects your server appliance from electrical noise and damaging power surges. During a power failure, the UPS instantly switches your NAS appliance to emergency battery backup power.

Once you have installed a UPS for your server appliance, you can set options for its operation using the UPS task on the Maintenance page of the NAS admin tool. The UPS task enables you to control how the UPS service works on your NAS appliance. The available UPS settings depend on the specific UPS hardware installed on your system.

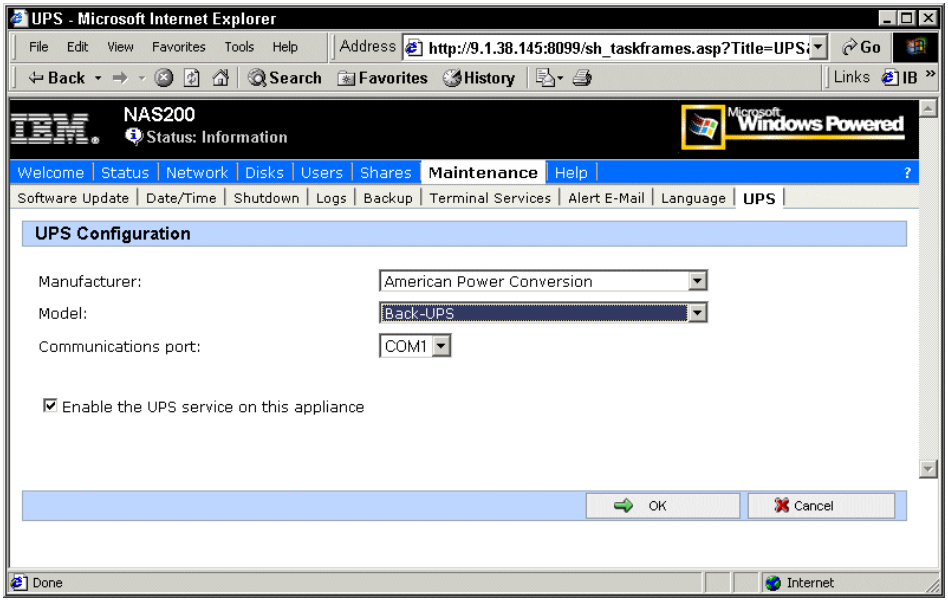

Figure 8-83 UPS configuration screen

To ensure that the server appliance is protected from power failures, test it by simulating a power failure. Do this by disconnecting the main power supply from the UPS device. Your server appliance and peripherals connected to the UPS device should remain operational, messages should be displayed, and events should be logged. Wait until the UPS battery reaches a low level to verify that a graceful shutdown occurs. Restore the main power to the UPS device, and check the event log to verify that all actions were logged and there were no errors.

All detected power fluctuations and power failures are recorded in the event log, along with UPS service start failures and server shutdown initiations. Critical events may change the status of the server appliance.

9

Systems management for NAS

Thus far, we have described some of the systems management tools. In this chapter, we point out additional tools that can be found in the IBM TotalStorage NAS appliances.

We cover these topics:

- Description of IBM Director
- Hardware aspects of systems management
- IBM Director preloaded on NAS devices
- Using IBM Director on a NAS appliance
- How to install IBM NAS Extensions to IBM Director
- Microsoft Multiple Device Manager (MDM)

9.1 Description of IBM Director

In this section, we introduce the IBM Director systems management solution. If you are already familiar with this topic, you might want to skip directly to 9.4, “Using IBM Director on a NAS appliance” on page 289.

IBM Director (the current version is 3.1) is a powerful client/server workgroup manager, built on industry standards and designed for ease-of-use. It can leverage your existing enterprise or workgroup management environments to securely and efficiently access and manage physically dispersed IT assets. With IBM Director, IT administrators can view the hardware configuration of remote IBM and non-IBM systems in detail to monitor the usage and performance of critical components, such as processors, disks, memory and network subsystem. Advanced management capabilities, such as remote control, process management, event management, and inventory are part of this solution.

IBM Director provides a convenient management user interface to execute tasks on specific systems or groups of systems. On the client side it uses a single, powerful agent which provides the client services to manage your systems intelligently. It supports systems management standards from legacy SNMP to DMI and emerging CIM. A wide variety of clients running operating systems from Microsoft, Novell, IBM and the major Linux distributors, and an assortment of protocols including TCP/IP, IPX, SNA, NetBIOS, SLIP and HTTP can be used.

The three major components of IBM Director are the Director Server, Director Console, and Director Agent.

IBM Director Server: This is installed on a Windows NT 4.0 or Windows 2000 Server operating system in the managed environment. As the heart of IBM Director, the Director Server provides application logic and an up-to-the-minute data store of management information.

The IBM Director Server provides basic functionality through the following actions:

- Discovery of remote systems
- Presence checking
- Security and authentication
- Persistent store of inventory information, including operating system, installed applications and hardware
- Management Console support
- Support for each administrative task

IBM Director includes Microsoft JET as its built-in SQL-compliant database, but can be installed with an existing database if needed. Microsoft SQL Server (6.5 or higher), IBM DB2 Universal Database Version 5.2 and Oracle 7.5 are supported. As long as the Director Server is available, the information about the managed system will be available even if the client system in question is not responding.

IBM Director Console: This is a Java based user interface from which administrative tasks are performed. The console provides comprehensive hardware management based on a single click or drag-and-drop operation. All system specific data gathered by the Director Console is stored on the Director Server. The IBM Director Console can communicate with only one IBM Director Server at a time. Multiple Director Consoles, however, can be open at the same time, each communicating with the same or a different Director Server.

IBM Director Agent: This includes all of the code and interfaces needed on the managed systems. It can be used to manage the local system through a Web browser, through Microsoft Management Console or via IBM Director Console. In this case, it communicates with the Director Server, executes the queries and provides requested information. The Director Agent comes preloaded on all IBM TotalStorage NAS appliances.

IBM Director Server Extensions: These offer an additional component that can be used to achieve advanced systems management functionality. They work with the Advanced Systems Management Processor or other systems management functions contained in the IBM TotalStorage NAS and xSeries server hardware. They help track resource utilization and provide recommendations to improve performance; enhance reliability by helping administrators discover, set up, and manage clustered servers from a single GUI; and let administrators configure, monitor, and manage IBM ServeRAID adapters and arrays without taking the RAID system off-line to perform maintenance.

Here we provide a short description of each tool:

- **Management Processor Assistant:** Through the management processor, built into IBM xSeries servers and IBM TotalStorage NAS appliances, you get exceptional control of remote systems even if the targeted system is not powered on. Management Processor events can be upward integrated into supported third-party workgroup/enterprise systems management applications.
- **Capacity Manager:** Tracks resource utilization, identifies multiple levels of existing or potential bottlenecks, and makes recommendations to improve performance. Enhancements for Version 3.1 include the ability to run Capacity Manager on groups of systems.

- **Cluster Systems Management:** Enhances reliability of Microsoft Cluster Server (MSCS) clusters by helping you control and manage clustered servers and alerting administrators to any event in the cluster.
- **ServeRAID Manager:** Configures, monitors, and manages the ServeRAID™ Controller through a graphical display. In IBM Director Version 3.1, this tool is enabled for CIM, allowing its management information to be included in IBM Director's upward integration into higher-level management products.
- **Software Rejuvenation:** Predicts pending operating system failures that could lead to costly downtime and can automatically refresh the software for optimal operation. Enhancements for IBM Director Version 3.1 include a "culprit list" of the applications most likely contributing to the server software degradation, and a Trend Viewer, which graphically depicts the software aging over time.
- **System Availability:** Tracks and provides a variety of graphical views of system downtime or uptime for an individual system or group of systems. On Windows-based operating systems this tool queries information from the Event Log. Enhancements for IBM Director Version 3.1 include the ability to differentiate between planned and unplanned outages, and the persistent store of system availability data.
- **Rack Manager:** Configures and manages a rack by dragging and dropping elements on a realistic graphical depiction of the rack and provides health status information of the rack and its components. Enhancements for IBM Director Version 3.1 include the ability to drag-and-drop components between racks, a capability to drill down for detailed system health data, and incorporation of Rack Manager data into the IBM Director hardware inventory database.

For detailed information about IBM Director and its usage, see the redbook: *Implementing IBM Director Management Solutions*, SG24-6188-00.

9.2 Hardware aspects of systems management

As IBM TotalStorage NAS appliances are based on IBM xSeries hardware, they have the Advanced System Management Processor integrated into the planar.

The processor provides the administrator with extensive remote management capabilities — even when the system has been switched off or when it has failed. The processor is an integrated subsystem solution independent of the hardware and operating system. It complements the server and IBM NAS hardware instrumentation by monitoring, logging events, reporting on many conditions, and providing full remote access independent of server status.

The ASM Processor controls IBM's innovative Light Path Diagnostics. In conjunction with the Light Path and Predictive Failure Analysis (PFA), the processor provides extensive alerting and real-time diagnostics. These indicate when a component such as a hard drive, power supply, or fan is failing. PFA can send notifications about the component, anticipating problems to help keep business up and running. The processor logs and sends alert for PFA events on the CPU, voltage regulating modules (VRMs), and Error Correction Code (ECC) memory, as well as on power supplies and fans.

9.3 IBM Director preloaded on NAS devices

IBM Director Agent and Server Extensions (Version 3.1) are integrated within the NAS appliance preload (Version 2.5). This integration provides the NAS appliance with a client management application for networked computers, centralizing control of NAS appliance and lowering TCO.

Previous releases of the IBM NAS appliances were preloaded with IBM Director Agent Version 2.2. Release 2.5 preloads the latest version of the Director Agent, V3.1. Here are some of the facts on interoperability:

- IBM Director Server V3.1 requires Console V3.1 and vice versa.
- IBM Director Server V3.1 can manage Agents V2.2.1, V2.2, or V3.1.
- Agent V3.1 requires an IBM Director Server V3.1.
- For ServeRAID, the Agent, Console, and Server must be at V3.1.

IBM Director services use some of the latest systems management standards, including Common Information Model (CIM), Web-Based Enterprise Management (WBEM), and the Extensible Markup Language (XML). Because of this, the management of NAS devices can be integrated with several enterprise management environments (Tivoli Enterprise, Tivoli NetView, Computer Associates Unicenter, HP Openview, Microsoft SMS, and Intel LANDesk Management Suite).

9.4 Using IBM Director on a NAS appliance

In this section we describe some real-life usage of IBM Director systems management solution on the IBM NAS appliance. To be able to understand and use them, a general level of knowledge in IBM Director configuration and usage is assumed.

For extensive information about installation, configuration, and usage of IBM Director, see the redbook: *Implementing IBM Director Management Solutions*, SG24-6188-00.

9.4.1 Discovering NAS systems

All IBM TotalStorage NAS appliances come preloaded with IBM Director Agent code. To be able to manage them from the IBM Director Console, they have to be discovered by the IBM Director Server first. The procedure is described here.

1. Provided the IBM Director Server is already installed in your network, start the Console by clicking **Start —> Programs —> Director —> Management Console**.

Login by providing the Director Server name (or its IP address), user ID and password, as shown in Figure 9-1. The user ID and password must be an authorized account on the IBM Director Server.

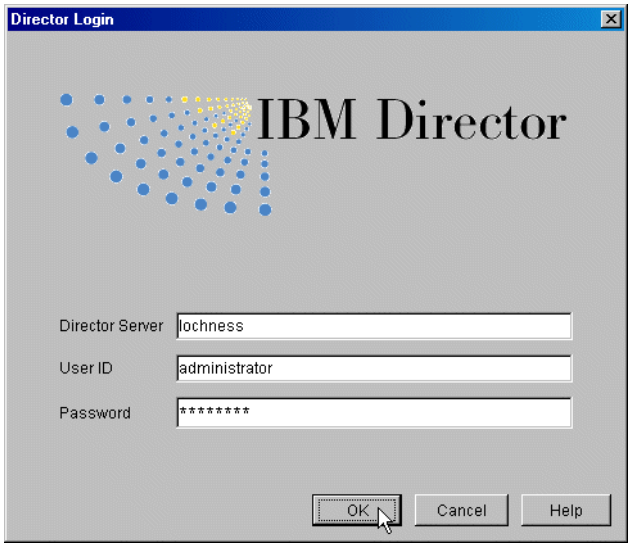

Figure 9-1 IBM Director console login

2. After a successful login, the Console window opens. It is made up of a Menu bar on the top, a Tool bar below it, Groups, Group Contents, and Tasks panes in the middle part, and a Ticker Tape and Status bar in the bottom part of the window. You can start discovering new systems by clicking **Discover All Systems** icon in the Console's Toolbar, as shown in Figure 9-2.

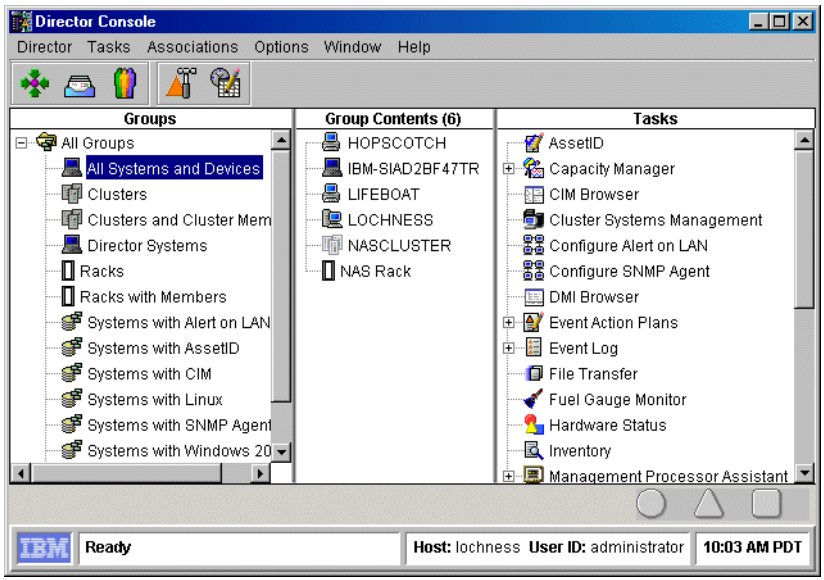

Figure 9-2 IBM Director console — start discovering systems

Tip: If you want to be selective, only native IBM Director systems can be discovered by choosing **Tasks —> Discover Systems —> Director Systems**.

3. Newly discovered systems are represented with solid icons in the Group Contents (middle pane), as shown in Figure 9-3.

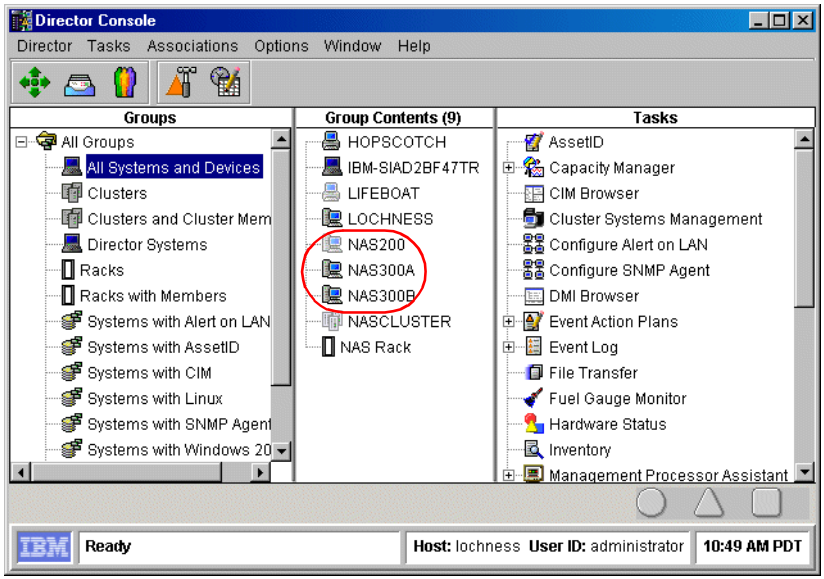

Figure 9-3 IBM Director — discovered systems

9.4.2 Executing tasks

You can start IBM Director tasks and operations in different ways. Some tasks are performed by right-clicking the system's icon and selecting the desired task from the context menu or by dragging and dropping task icons onto the managed systems. Some operations, on the other hand, can only be done by selecting operations from pull-down menus (for example, discovering new systems). This section gives you a very basic guidance to try the techniques most often used to manage IBM TotalStorage NAS devices.

Context-sensitive menu

Discovered systems are stored in the Director database on the Director Server, so even if they are off-line (represented by a greyed-out icon), their management properties are still available. To view system attributes, right-click the system's icon and select **Open** (Figure 9-4).

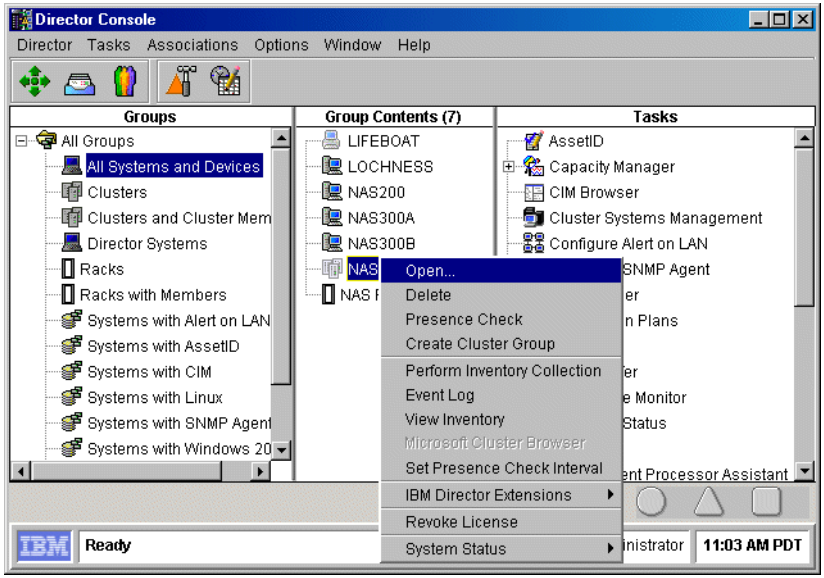

Figure 9-4 IBM Director — opening system attributes

A window opens showing several attributes, among them the System Name, Unique System ID, Operating System Version, System State, and Presence Checking Interval (Figure 9-5).

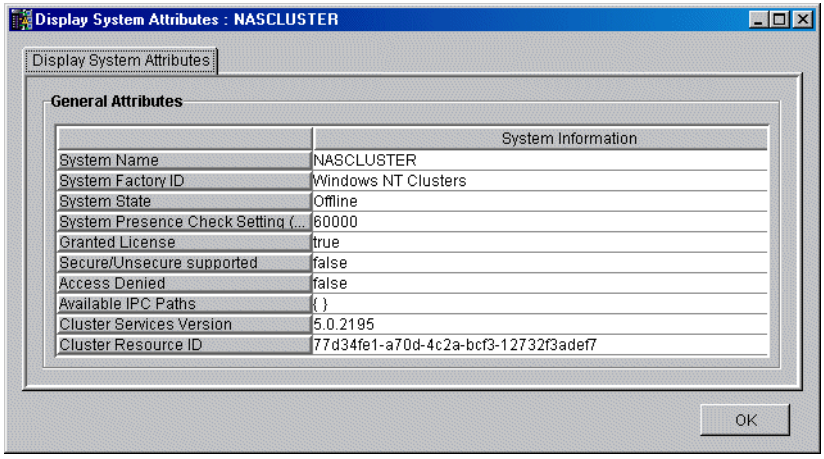

Figure 9-5 IBM Director — System Attributes

Drag-and-drop operations

Several windows displayed in IBM Director consist of two or more panes. In most instances you can drag-and-drop task and target icons between these panes.

However, you cannot perform drag-and-drop operations between two separate IBM Director windows.

To execute a task on a managed system in the Console, drag the managed system icon from the Group Contents pane and drop it onto the task icon in the Tasks pane, as shown in Figure 9-6.

Tip: You can also drag the task icon from the Tasks pane and drop it onto the desired managed system or group of systems icon in the left two panes.

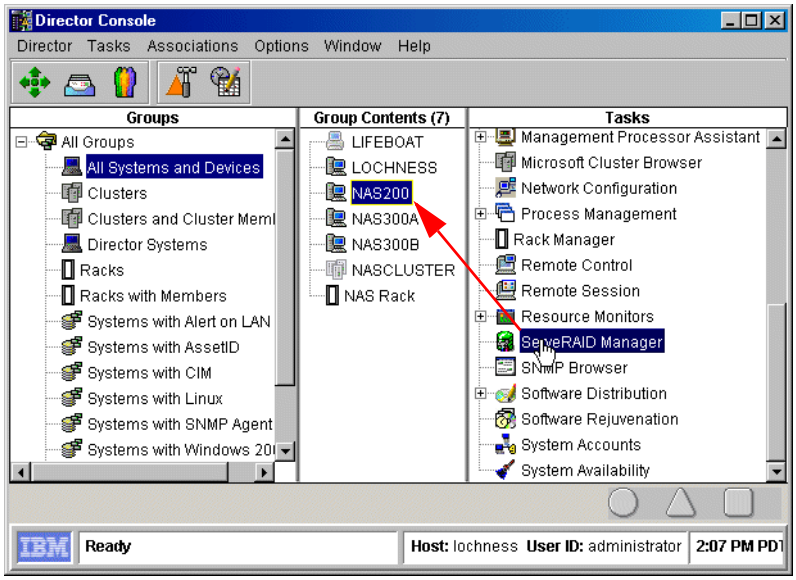

Figure 9-6 Dropping the system onto management task

9.4.3 Grouping systems

Groups consist of logical sets of managed systems. An example of a group might be one that contains only NAS devices. When you first log on to the Console, a number of default groups are created. Included in this default list is the All Systems and Devices group, which contains all discovered devices. After the group is created and populated with desired systems, sharing the same attributes, it is easier to run systems management tasks on the whole group.

9.4.4 Event and action management

Event management in IBM Director is done using the Event Action Plan tool. It identifies and categorizes system events and initiates actions in response to those events. For example, it enables an administrator to determine when a file server's hard disk drive is approaching its full capacity, so that he or she can provide a fix before users are affected. The administrator would use the Resource Monitor task to set a storage threshold on the file server; that configuration would generate an event when the remaining free space on the main data drive drops below, say, 500 MB. Using Event Management, the administrator can then configure an Event Action Plan that triggers a notification (console pop-up or page alert, for example) when the threshold is exceeded.

Event Action Plans consist of two components:

- **Event Filters:** A set of specific criteria which determine if an incoming event will trigger an action
- **Actions:** The tasks that are executed as a result of the event

Event Action Plans can be associated with a system or a group of systems.

Using system threshold to generate an event (alert)

When an IBM Director Server is installed, it already has some thresholds associated with subsystems like CPU, memory, disk, power supply. If there is a pending failure or critical usage of those resources, the administrator will receive notification about it without the need to configure it beforehand. However, in many environments there will be a need for customized monitoring of resources, and for that reason the administration personnel may have to modify existing thresholds or create new ones.

Probably the most important resource on the NAS appliance is disk space used for file storage. Here we describe the procedure to set up a disk threshold:

1. Login to Director Console and drag the Resource Monitors task from the Tasks pane (on the right) onto the selected system, as shown in Figure 9-7.

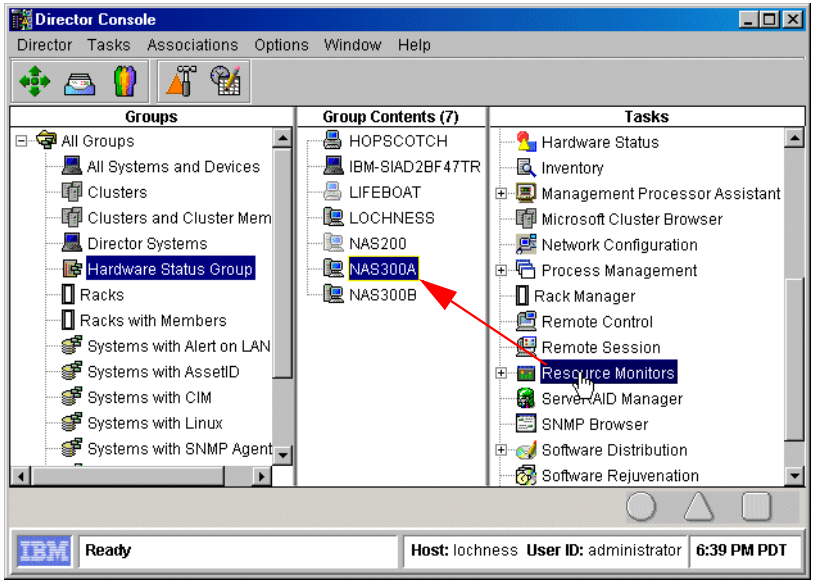

Figure 9-7 IBM Director — start resource monitor

2. The Resource Monitor window for selected system opens. By clicking the + sign in front of **Director Agent** and **Disk Monitors,** select **Drive D: Space Remaining**. Right-click it and select **Add to selected Resource Table** (Figure 9-8).

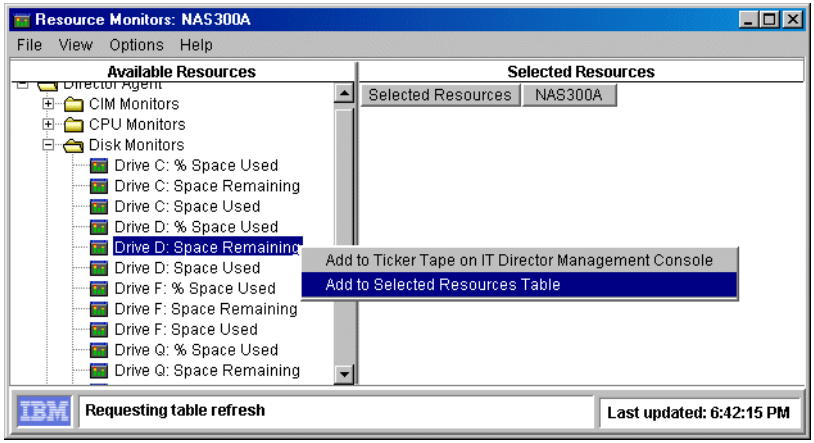

Figure 9-8 IBM Director — add disk resource monitor

3. The Disk Resource Monitor will be added to the Selected Resources pane, and the Director will start collecting data. In a few seconds, current resource data — remaining disk space in our case — will be shown under the device name (Figure 9-9).

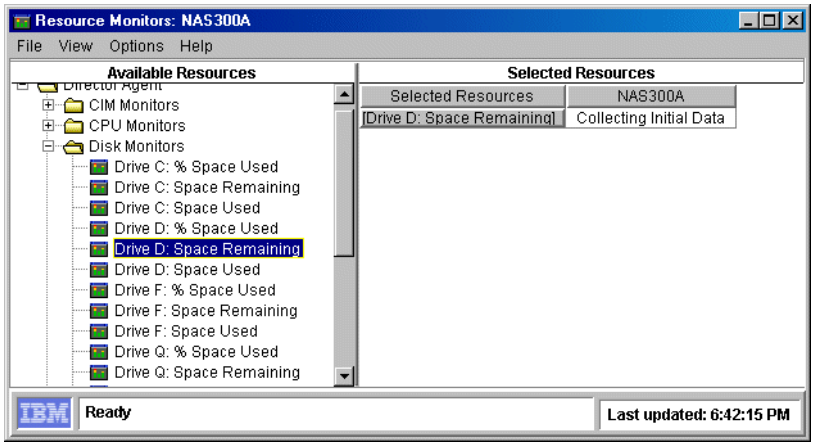

Figure 9-9 IBM Director — collecting disk data

4. Right-click the data value field and select **Individual Threshold**, as shown in Figure 9-10.

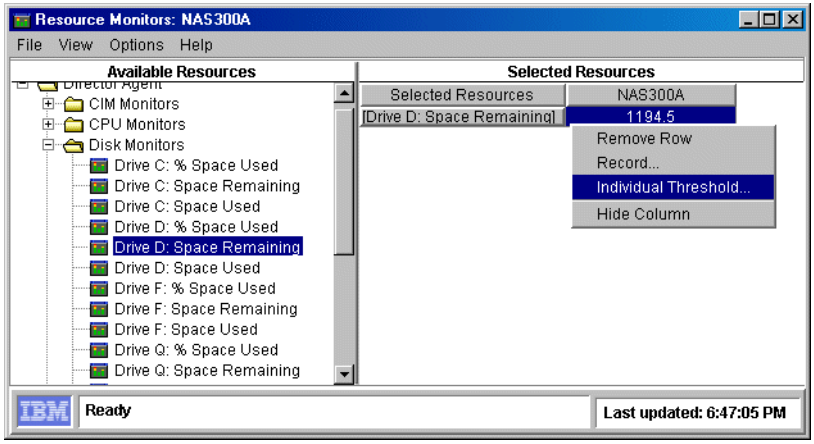

Figure 9-10 IBM Director — individual threshold

5. In the Threshold Settings window, enter the values for the remaining disk space, which will trigger a Warning and an Error condition. Click **OK** to close the window (Figure 9-11).

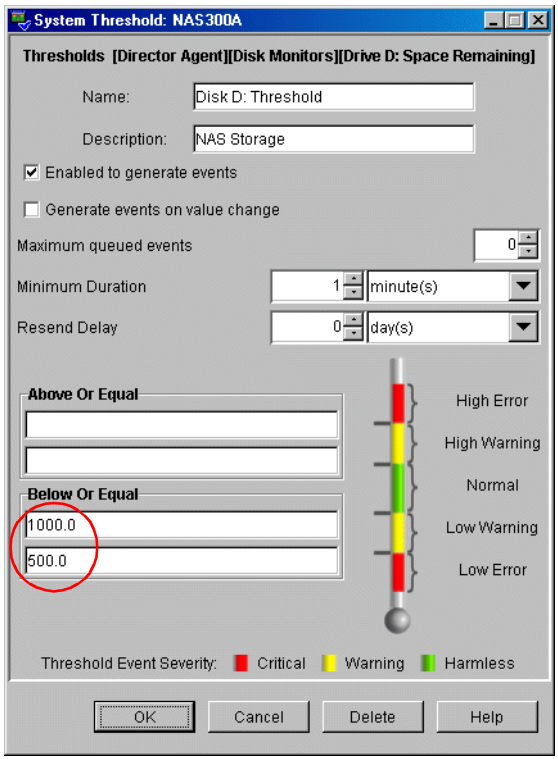

Figure 9-11 Threshold Settings

6. Now save the Resource Monitor by selecting **File —> Save As**, as shown in Figure 9-12.

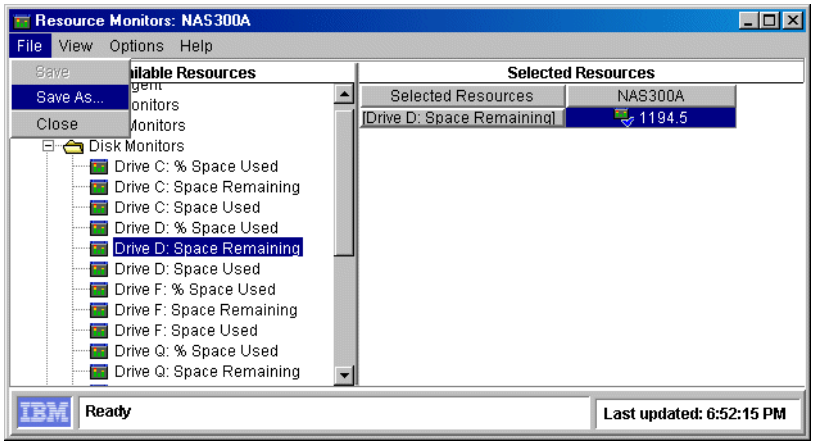

Figure 9-12 Saving Disk Space Resource Monitor

7. Provide a meaningful name for the Resource Monitor you just created (Figure 9-13).

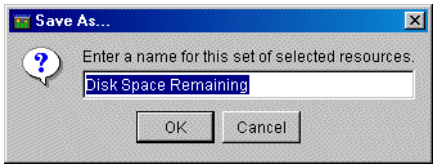

Figure 9-13 Resource Monitor Name

8. The last step needed is to activate the Resource Monitor we just created. In the main Console window, drag the monitor from the Tasks pane and drop it onto the selected system (Figure 9-14).

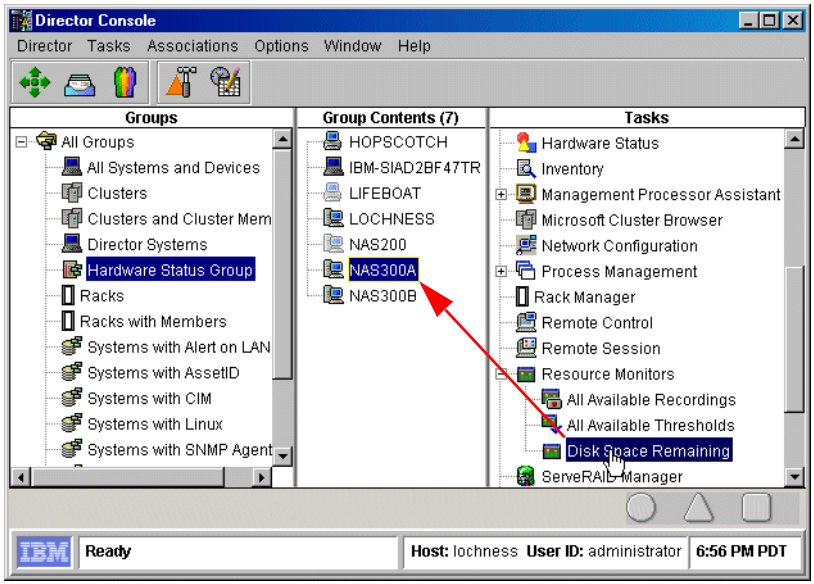

Figure 9-14 Activating Monitor

9. The Activated Monitor will be represented with a blue check mark in the icon, as shown in Figure 9-15.

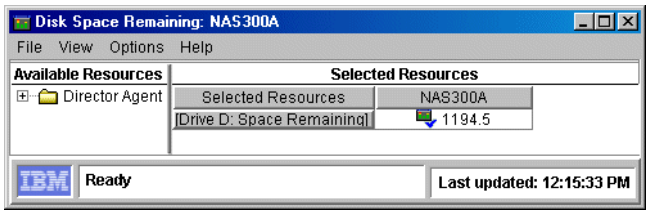

Figure 9-15 Monitor Activated

Viewing the Event Log

If you want to verify the Threshold Monitor we just created, drag the Event Log from the Tasks pane onto your system (Figure 9-16).

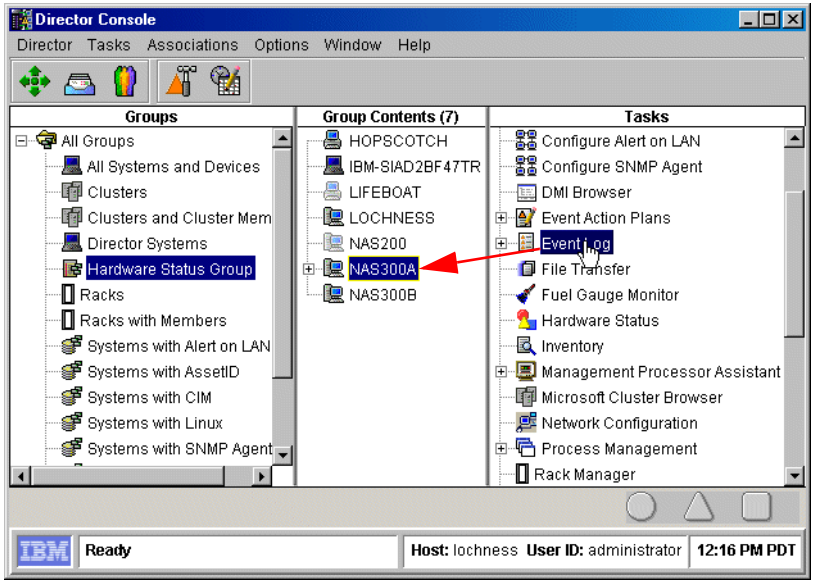

Figure 9-16 IBM Director — starting Event Log

The Event Log window opens, as shown in Figure 9-17.

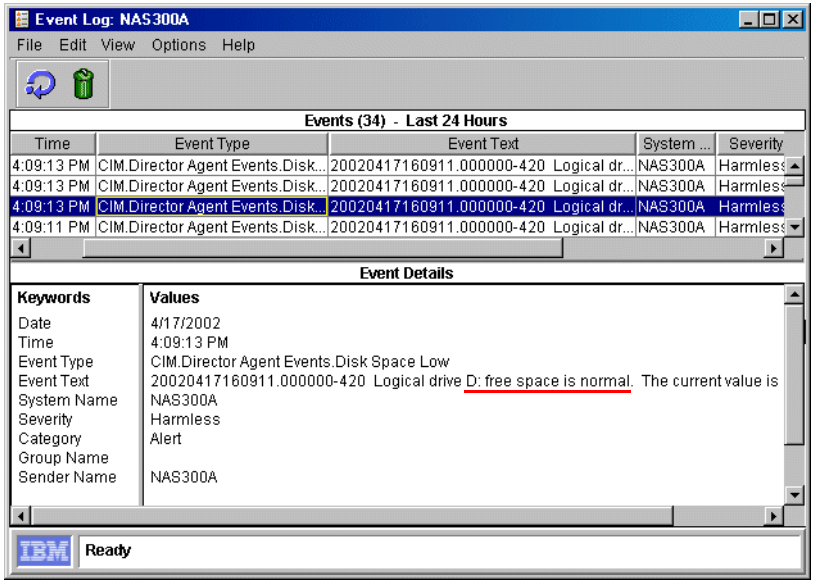

Figure 9-17 Disk Space OK in Event Log

Now copy some files to the disk to fill it up until it reaches the predefined threshold. A Warning or an Error event should show in the Event Log now, as shown in Figure 9-18.

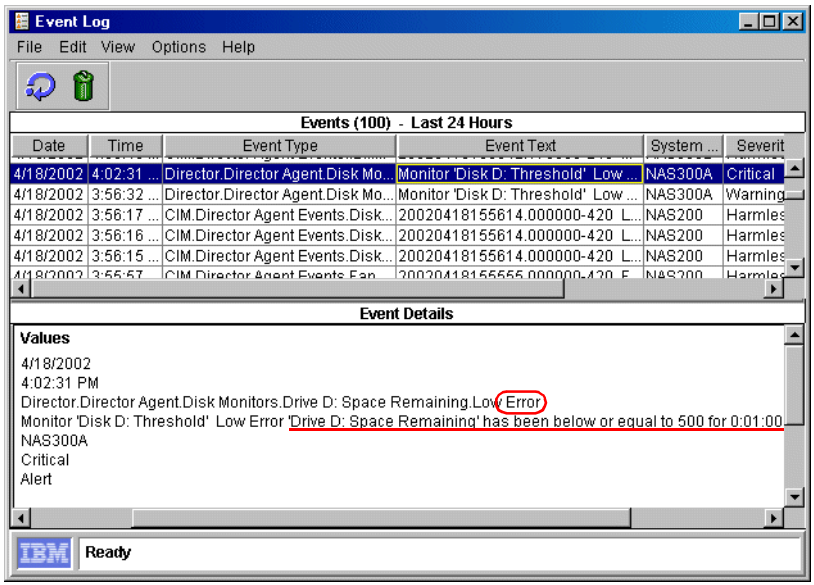

Figure 9-18 Disk Space Critical Error

Defining Event Filters

Now, as the Threshold Monitor is created and resource is being monitored, we need to define the Event Filter, which will receive all events and, based on the criteria defined, decide whether the event should trigger an action. Here is how you create a custom filter:

1. Click the Hammer icon in the Toolbar. This will open the Event Action Plan Builder window, right-click **Threshold Event Filter**, and select **New**, as shown in Figure 9-19.

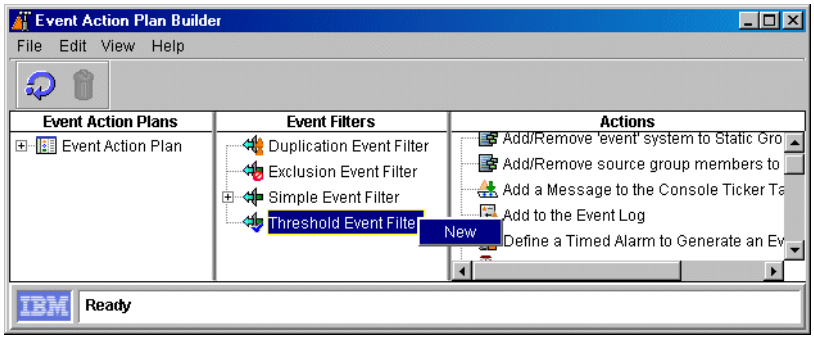

Figure 9-19 Event Action Plan Builder

2. In the next window, remove the checkmark next to **Any** and expand the tree by clicking the + sign next to **Director —> Director Agent —> Drive D: Space Remaining (**Figure 9-20). Select the desired event type(s) and then save the Event Filter by clicking the diskette icon and providing a meaningful name for the Filter.

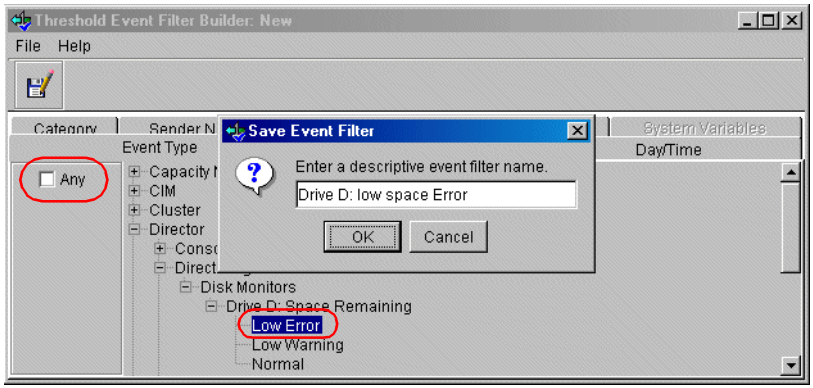

Figure 9-20 Selecting Event Type

Defining Event Action

Now we need to define the action that will be triggered upon reaching the preset threshold. IBM Director offers a wide variety of actions that can be chosen, depending on the situation and the severity of the alert. Here are some of them:

- Add a message to the console Ticker tape
- Add to the Event Log
- Log to textual Log file
- Post to a news group (NNTP)
- Send an alphanumeric or numeric page
- Send an event message to the Console user
- Send an Internet (SNMP) E-mail
- Send an SNMP trap
- Start a program on a system with the event or on the server
- Start a task on a system with the event or on the server

With the exception of Add Event to Event Log, each type of event has to be customized by right-clicking it and selecting the **Customize** option.

In our example we will add the “Remaining disk space” message to the Ticker tape area of the main console. Here are the steps to achieve this:

1. In the Actions pane of the Event Action Plan Builder window right-click **Add a Message to the Console Ticker Tape** and select **Customize** (Figure 9-21).

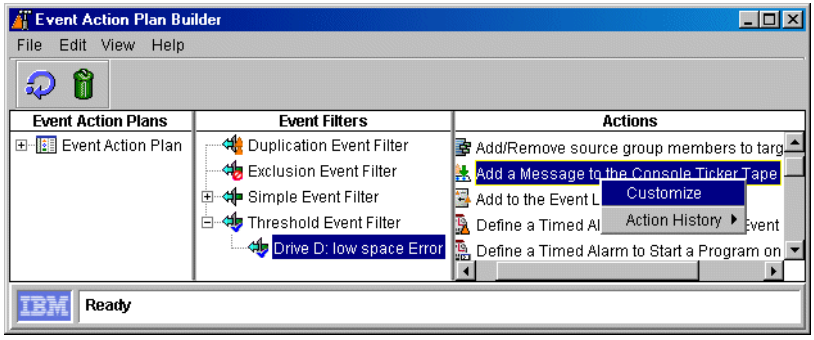

Figure 9-21 Customize Action

2. In the Message field enter the desired text, combined with variables for system (&SYSTEM), event text (&TEXT), time (&TIME) and date (&DATE) of the event, as shown in Figure 9-22.

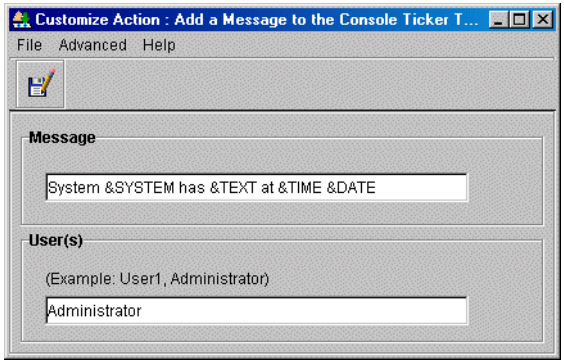

Figure 9-22 Define the Message

Tip: If you want the message to be displayed on all management consoles, you can enter asterisk (*) character instead of specific usernames.

3. Click the **Diskette** icon and save the Action by entering a descriptive name and clicking OK (Figure 9-23).

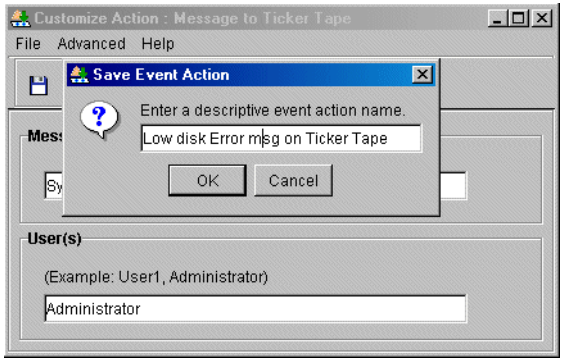

Figure 9-23 Save Event Action

Creating an Event Action Plan

Now that we have created the Event Filter and Event Action, we only need to combine them into an Event Action Plan. Following are the steps for doing this:

1. Right-click the **Event Action Plan** icon in the “Event Action Plan Builder” window and select **New**, as shown in Figure 9-24.

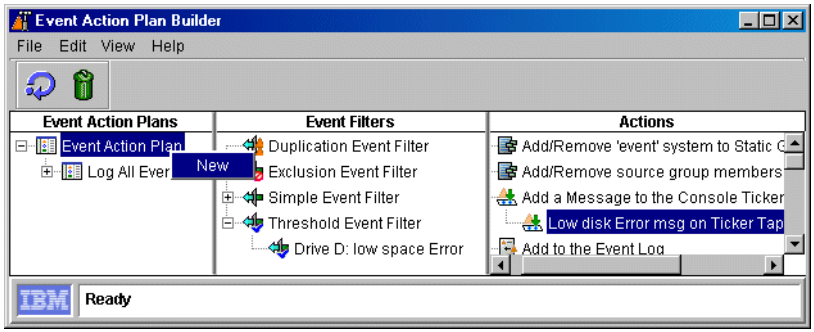

Figure 9-24 Create Event Action Plan

2. Name the new Action Plan and confirm the save operation by clicking the **OK** button (Figure 9-25).

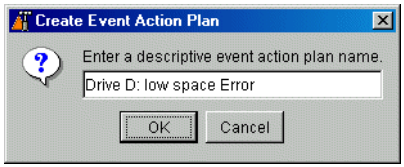

Figure 9-25 Save Event Action Plan

3. Add the customized Filter to the new Action Plan by dragging it from the middle pane to the left pane (Figure 9-26).

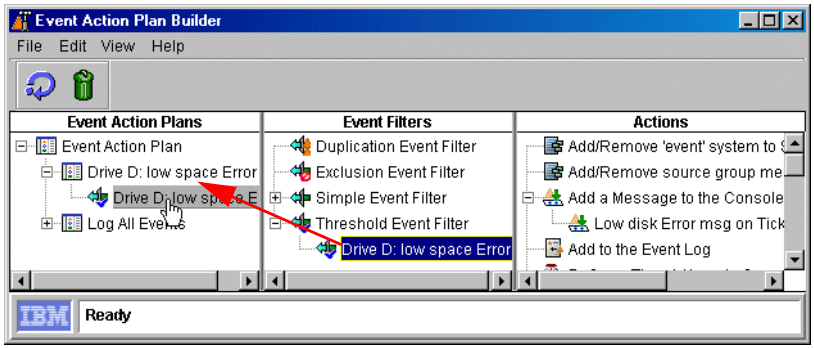

Figure 9-26 Add Filter to Action Plan

4. Next, drag-and-drop the customized Action to the new Action Plan, as shown in Figure 9-27.

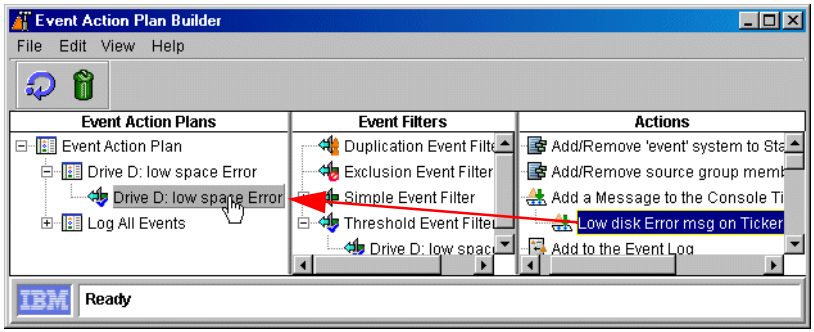

Figure 9-27 Add Action to Action Plan

5. The Action Plan is set up, as shown in Figure 9-28, so close the Builder window and return to the main Console window.

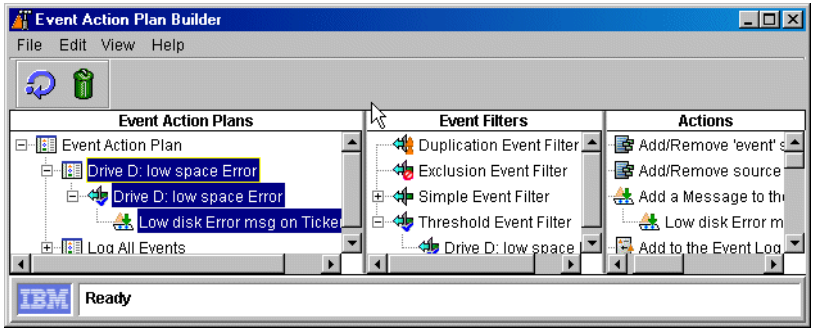

Figure 9-28 Event Action Plan created

6. The last step is to activate the new Action Plan. Drag it from the Tasks pane onto the selected system in the Group Contents pane, as shown in Figure 9-29.

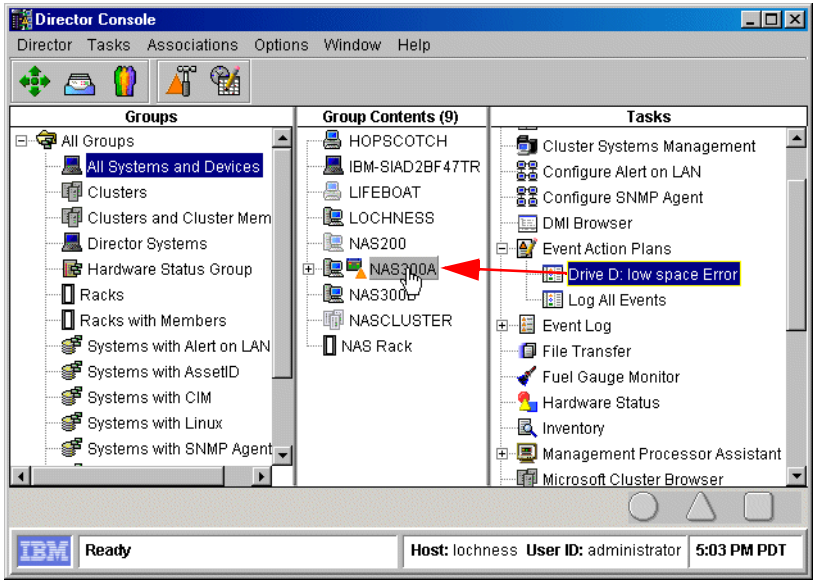

Figure 9-29 Activate the Event Action Plan

7. When the remaining disk space falls below the defined threshold, the event is processed and you will notice the message in the bottom part of the main Console — the Ticker Tape area, scrolling from right to left, as shown in Figure 9-30.

Tip: Even though the whole procedure might appear lengthy, the real value of Event Action Plans lies in their “reusability”. Once different Thresholds, Event Filters and Actions are created, they can be simply combined into several Action Plans and easily applied to different systems or groups of systems.

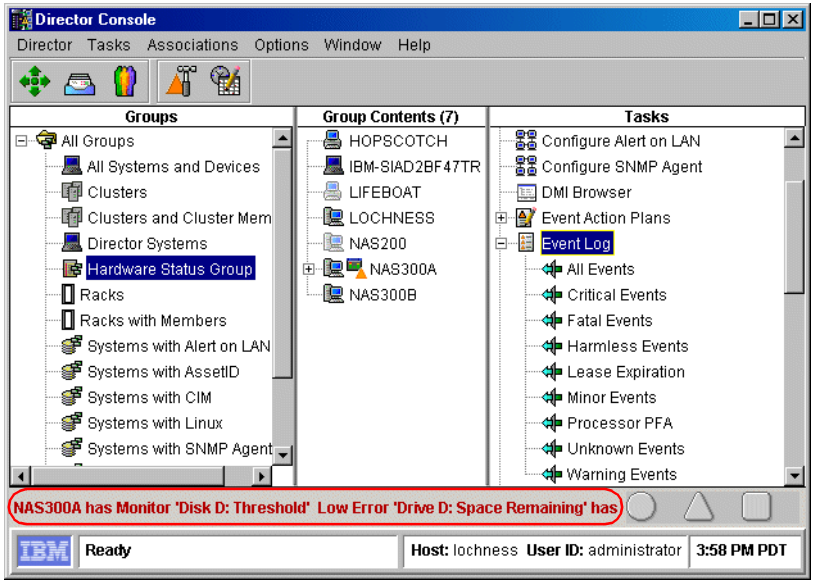

Figure 9-30 Ticker Tape Message

Instead of adding the message to the Ticker Tape, any other action could be selected. For example, an E-mail could be sent to an administrator's E-mail or text messaging account. This makes the Event Action task a universal notification tool, able to deliver the message regardless of the administrator's location.

Another example might include running a batch file using the **`diskpart`** utility, expanding the disk partition and thus giving needed disk space for users connected to the IBM TotalStorage NAS appliance.

9.4.5 Rack Manager

Rack Manager is a a graphical representation of systems and devices, physically placed inside the rack. It can be used for status monitoring and management of racks and their components. See Figure 9-31.

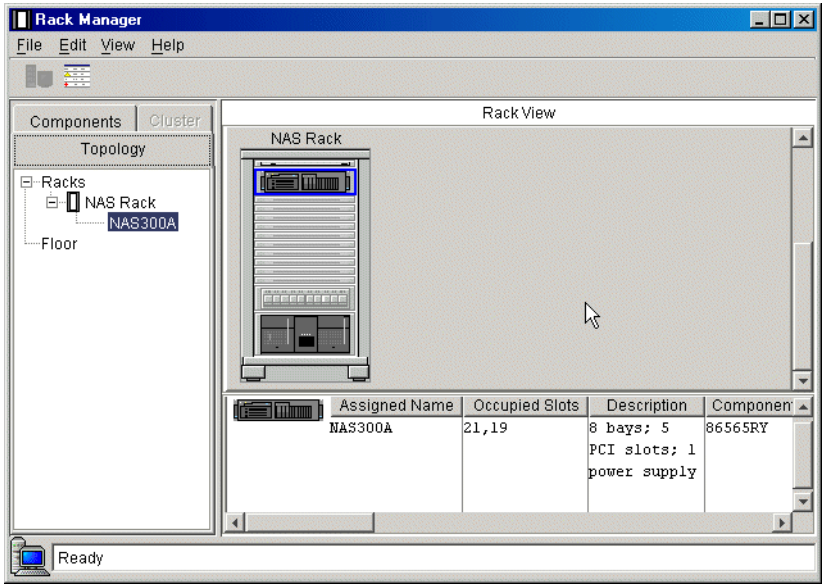

Figure 9-31 Rack Manager

For example, it is possible to visualize all components of a rack and then start management tasks from the same window by right-clicking the component and selecting the needed task. See Figure 9-32.

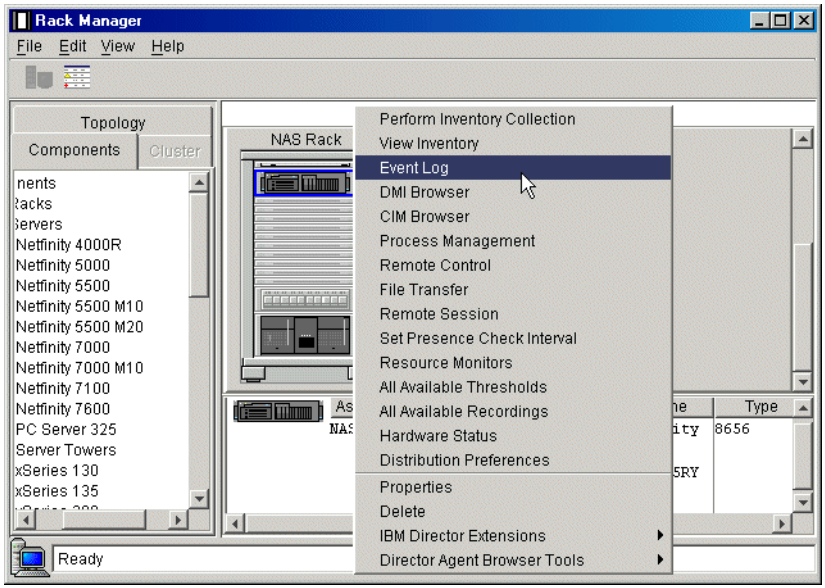

Figure 9-32 Starting Event Log from Rack Manager

9.4.6 System Availability

The System Availability tool can be used to track the availability of your IBM TotalStorage NAS appliance or a group of systems. The information gathered can be viewed as reports or graphical representations.

Note: This tool uses information from the operating system log (Event Log in Windows). If you clear the logs, all System Availability data will be lost as well.

The System Availability tool is able to distinguish between planned and unplanned outages. It will classify outages based on duration to help you track down reasons for system problems.

To start it, simply drag the System Availability task from the right pane onto the selected system or group of systems, as shown in Figure 9-33.

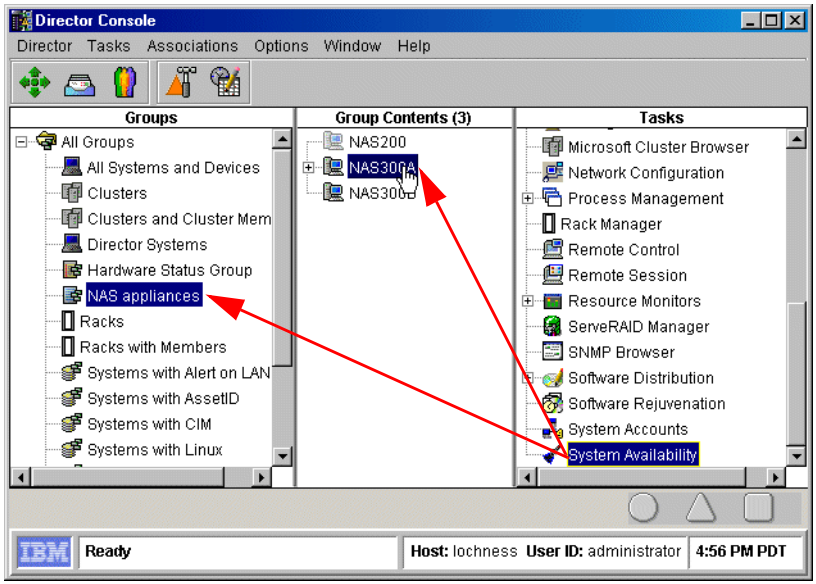

Figure 9-33 Starting System Availability tool

To select the time period for which the availability data will be displayed, select **File —> Set Time**. In the window that opens, you can choose between All time, a week, a month, 3 months, a year, or a customized period (Figure 9-34).

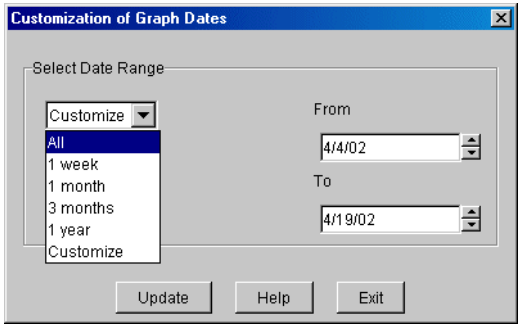

Figure 9-34 Setting the time period

Different types of availability information can be chosen from the **View** menu. The following views are available:

- Distribution of system outages (Figure 9-35):

  Displays a pie chart of how long the systems were down before being restarted, represented as a percentage.

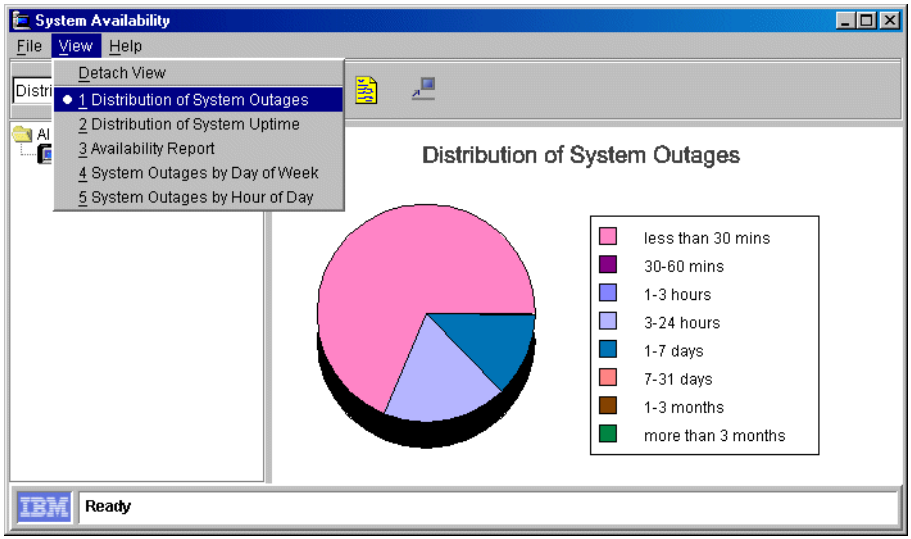

Figure 9-35 Distribution of System Outages

- Distribution of system uptime (Figure 9-36):

  Displays a pie chart of how long the systems were up before being restarted.

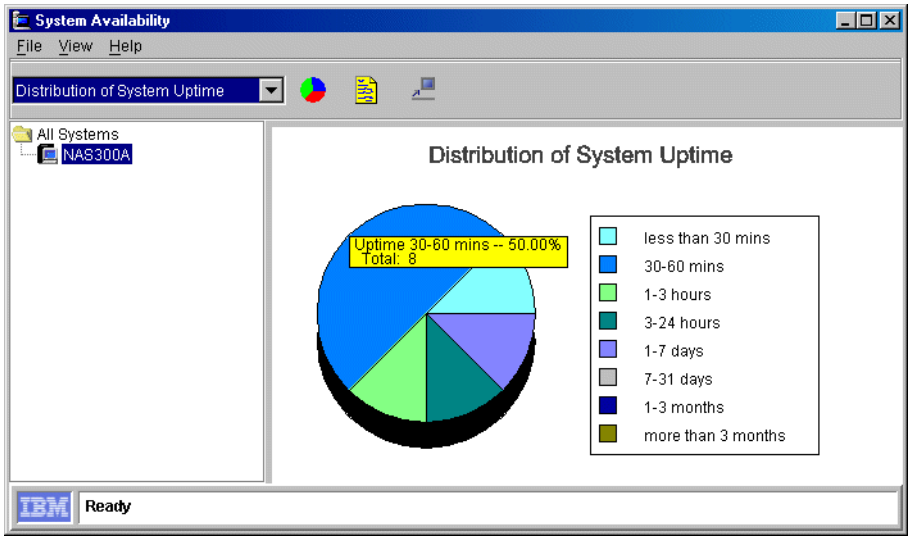

Figure 9-36 Distribution of System Uptime

Tip: If you position the cursor on a particular segment of the chart, detailed statistics of that segment (duration and number of outages, uptimes, daily statistics) will be shown in a pop-up box.

- System outages by day of week (Figure 9-37):

  Displays the day of the week that outages occurred for the system or group as a percentage of downtime.

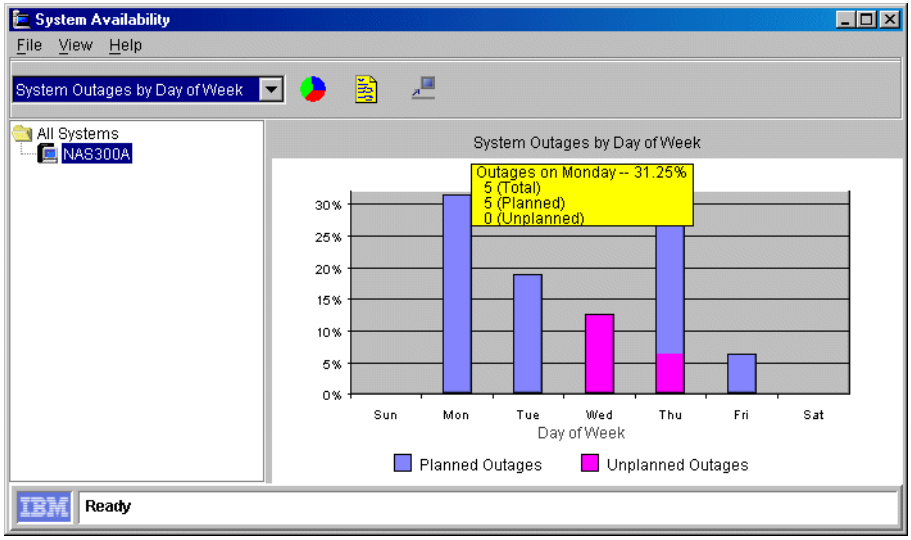

Figure 9-37 System Outages by Day of Week

- System outages by hour of day (Figure 9-38):

  Displays the time of the day that outages occurred for the system or group as a percentage of downtime.

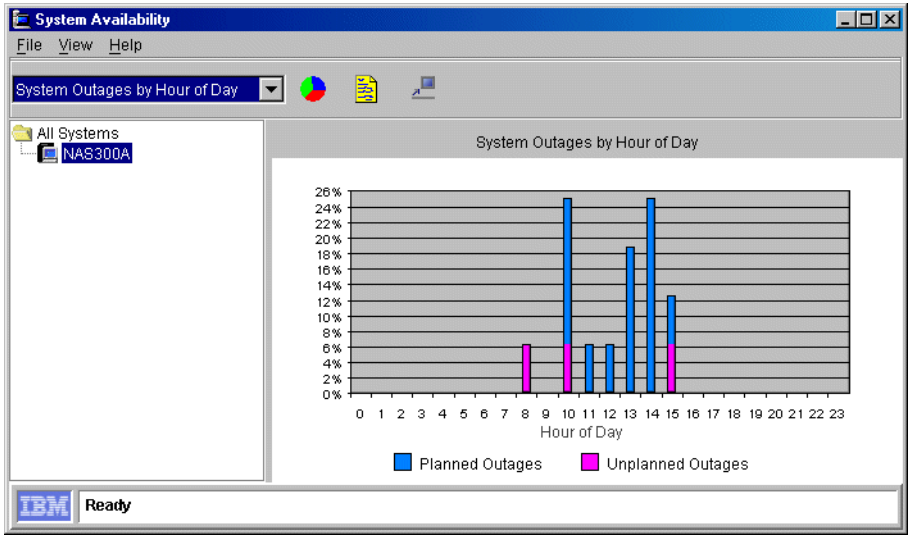

Figure 9-38 System Outages by Hour of Day

- Report of system availability (Figure 9-39):

  Displays a detailed report of the system's availability including the list of all uptimes and downtimes. Additional statistical parameters are displayed, including total restarts, number of systems restarted (if the task was started against a group) and mean time between unplanned outages.

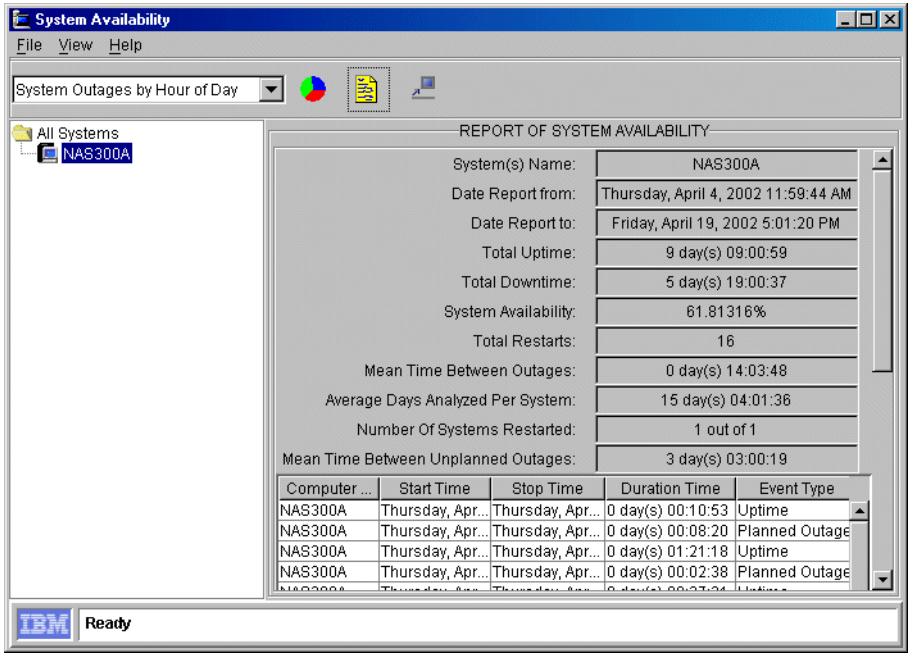

Figure 9-39 Report of System Availability

9.4.7 Capacity Manager

Capacity Manager is an efficient system management tool integrated into IBM Director to help you measure the potential bottlenecks of various subsystems. You can use this tool to forecast performance degradation of IBM xSeries servers and TotalStorage NAS appliances. You can plan for an appropriate action to overcome the bottleneck well in advance, so as to prevent overall performance degradation.

The key concept to understand about Capacity Manager is that the data is always being gathered. Unlike Windows 2000's System Monitor, you do not have to start the logging of data. With Capacity Manager, you simply specify what data you want retrieved from the systems in your network and it is gathered and displayed graphically for you. Up to one month's worth of data is automatically saved by every system.

Collected data can also be exported into a spreadsheet for further analysis. These reports show at a glance potential bottlenecks within the selected systems. Analysis and ability to predict bottlenecks is critical when planning for future upgrades. Capacity Manager gives you the ability to plan the allocation of hardware upgrades for the systems that really need them before a capacity bottleneck occurs.

Capacity Manager is part of the IBM Director Server Extensions, so it comes preloaded on the IBM TotalStorage NAS devices.

When you start the IBM Director Console, the Capacity Manager icons will appear in the task pane of IBM Director as shown in Figure 9-40.

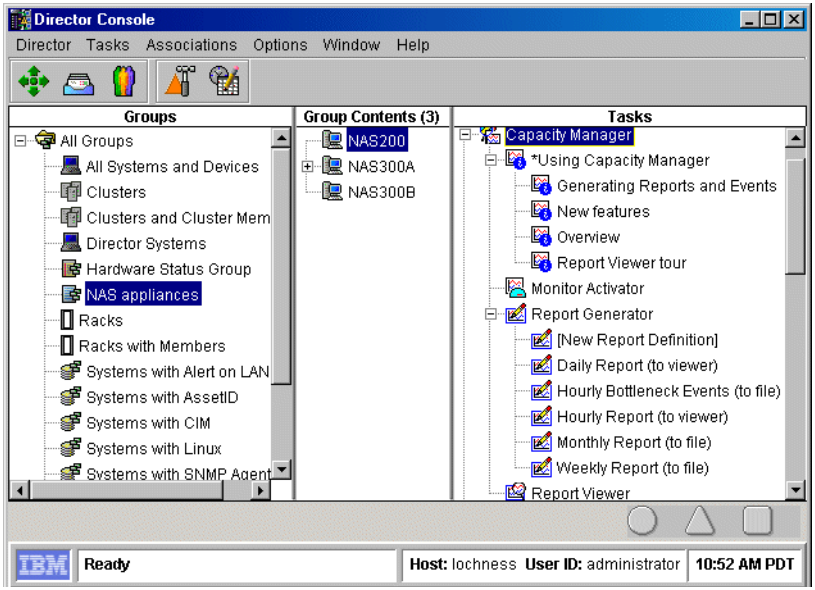

Figure 9-40 Capacity Manager tasks

You can perform three functions from this menu:

1. Double-click any of the four Using Capacity Manager tasks (as shown in Figure 17-2) to learn about Capacity Manager:
   - New features
   - Overview
   - Report Viewer tour
   - Steps to create a report

   If you are new to IBM Director and/or Capacity Manager, we suggest you review each of these help topics.

2. Change what data is recorded on specific clients using Monitor Activator (see “Monitor Activator” on page 420).
3. Generate a report either directly to the viewer or to a report file using Report Generator (see “Report Generator” on page 421).
4. View a report that has already been generated using Report Viewer (see “Report Viewer” on page 432).

Monitor Activator

The Monitor Activator function is where you specify what data is to be gathered on specific clients or groups of clients. Simply drag the Monitor Activator icon onto a group or a single client to activate it. The window shown in Figure 9-41 appears.

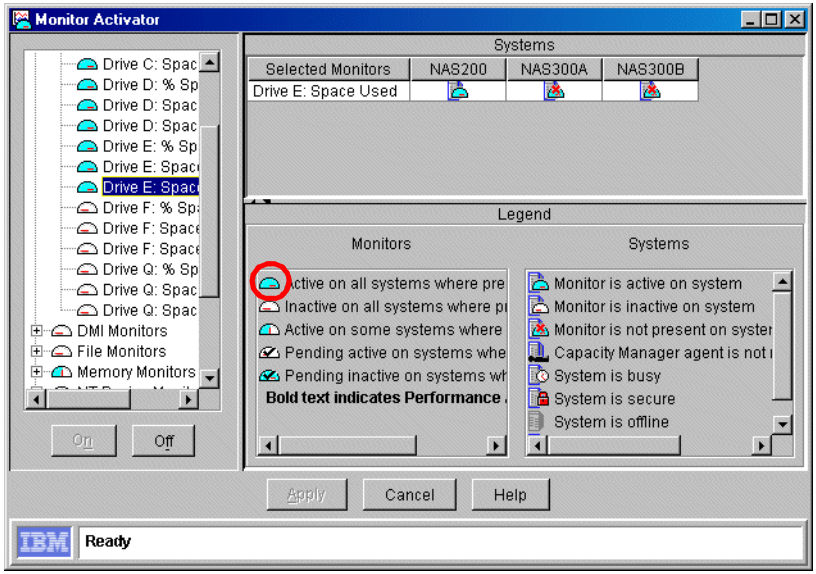

Figure 9-41 Monitor Activator window

Some of the counters (CPU utilization, process count, disk workload and usage, memory usage, network interface) are enabled (active) by default. They are represented with corresponding icons for an active monitor (see the Legend pane in Figure 9-41).

Other monitors have to be enabled by clicking the corresponding monitor in the left pane, and selecting the **On** button. When you are done enabling the monitors, you need to apply the changes by clicking the **Apply** button.

Report Generator

With this task, you gather data from specific systems and either display it on the screen (using the Report Viewer) or save it to a report file. Here are the predefined report definitions (as shown in Figure 9-40 on page 315):

- Daily Report (to viewer)
- Hourly Bottleneck Events (to file)
- Hourly Report (to viewer)
- Monthly Report (to file)
- Weekly Report (to file)

You can also create a new report by double-clicking **New Report Definition** in the IBM Director management console.

You can do the following with predefined reports:

- Execute it, just as you can with report definitions you've created, by dragging and dropping one onto a client or group of clients.
- Edit it by double-clicking the entry in the management console.
- Delete it by right-clicking the entry and clicking **Delete**.
- Changing the output definition from viewer to file or from file to viewer by right-clicking and clicking **To viewer or file**.

To generate a report, simply drag the report definition onto a client or a group of clients as shown in Figure 9-42.

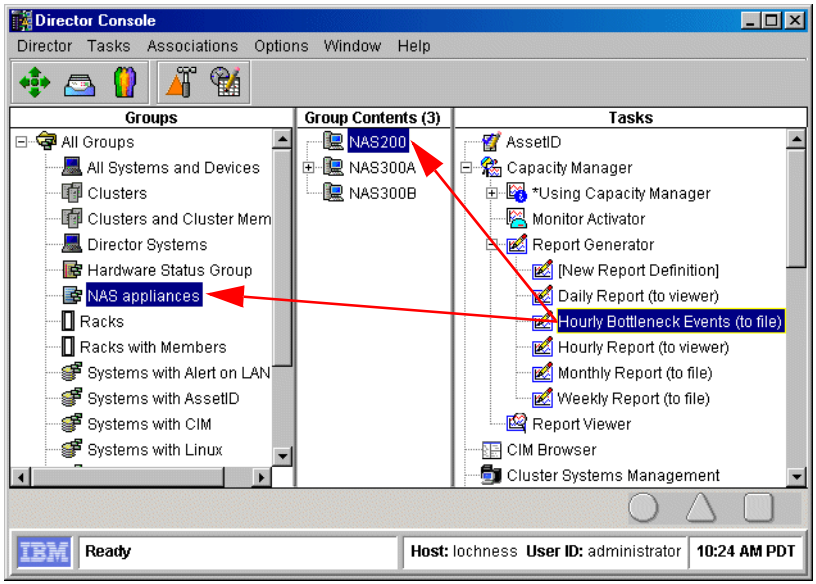

Figure 9-42 Generating a report

If you choose to output the report to the Report Viewer, you will see a progress window showing the status of data being gathered from each system. The Report Viewer then loads and displays the results.

If your report definition specified to output to a file, then you will see the following dialog box asking if you want to execute it immediately or schedule the task to be executed at a later time.

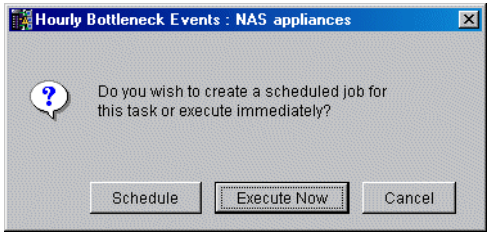

Figure 9-43 Output to file dialog box

If you select **Execute Now**, the progress window (Figure 9-43) will appear, informing you about the status of the reporting process.

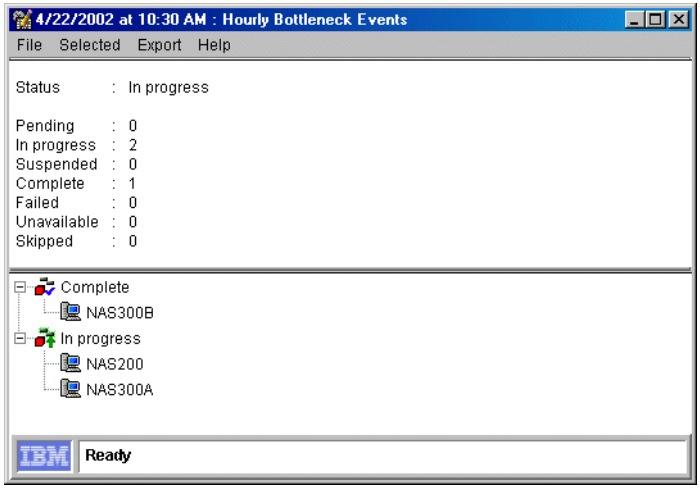

Figure 9-44 Report to file — progress window

Regardless of your choice, the report is saved to a file on the IBM Director Server (not the console). The name of the report is the name of the report definition, plus the time and date the report was created.

Report Viewer

The Report Viewer is used to examine reports you have requested to be gathered immediately or to examine reports you have saved to a file. The viewer starts automatically if the report definition you used specifies the output to go to the viewer.

To view a report that was saved to a file, double-click the **Report Viewer** icon from the task pane in the management console. You will then be prompted to select a report file (.CMR or .TXT) from the IBM Director server's C:\PROGRAM FILES\DIRECTOR\REPORTS directory.

A typical Report Viewer window is shown in Figure 9-45.

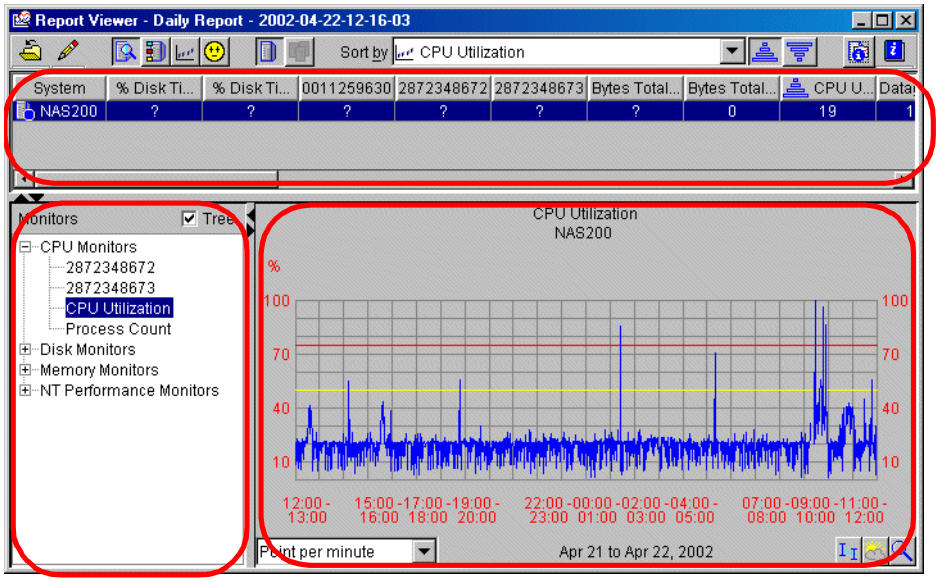

Figure 9-45 Report Viewer

As you can see, this window is made up of three panes:

- The System pane, on top, shows the systems chosen for the report.
- The Monitor pane, on the lower left-hand side, lists the monitors you have chosen in the Report Generator. Here you can select only one monitor at a time to be displayed in the Graph pane.
- The Graph pane, on the lower right-hand side, displays the chosen monitor in graphical format.

To display data on the graph, select a monitor from the Monitor pane, then one or more systems from the System pane (select more than one system with the Shift or Ctrl key). Figure 9-45 shows the CPU Utilization monitor selected for the NAS200 system. To make the graph larger, select the edge of the pane with your mouse and drag the panel up or left.

If you want to analyze a particular part of the graph, you can zoom into it by clicking the **Zoom** button and then clicking into a particular part of the graph, as shown in Figure 9-46.

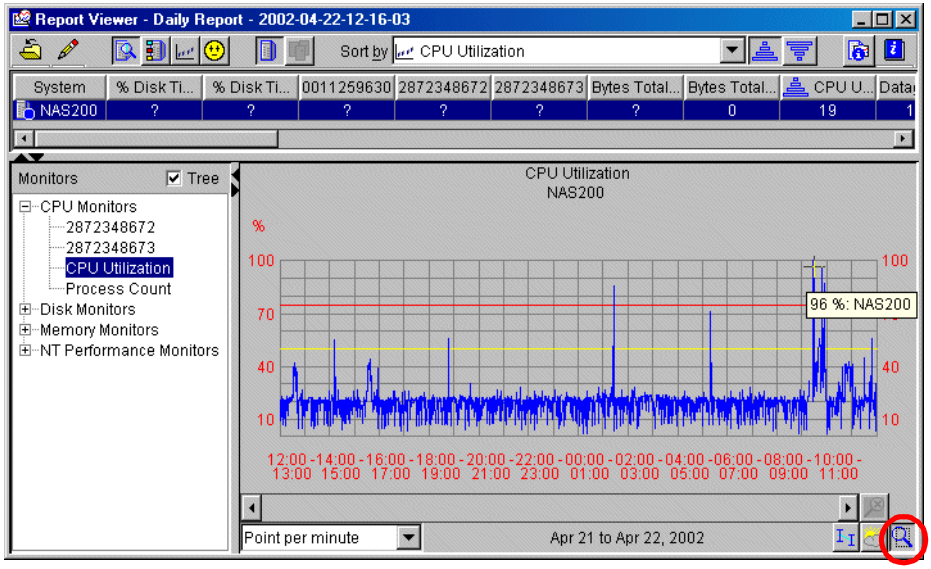

Figure 9-46 Zooming into Graph pane

Forecast

The forecast function is available by clicking the **Forecast** button while viewing the Capacity Manager report, as shown in Figure 9-47. The function allows you to see Capacity Manager's prediction of the future performance of your selected systems. The forecast is for whatever monitor you currently have selected. To see a forecast for another monitor, click its name in the monitor box.

For the forecast to be valid, Capacity Manager needs a minimum of 21 days of previously collected data where the system monitors have been running at least 50% of the time.

The forecast line is a dashed line with an arrow at the end. The forecast interval is a multiple of your data collection period. The default prediction period is set to the same length as the data collection period. For example, if you have a month of collected data, the forecast will be for a month into the future.

The confidence interval is represented by the dotted lines above and below the forecast line. The vertical bar at the beginning of the forecast data depicts the range. The gap between the actual collected data and the beginning of the predicted data serves as a separator between these two data sets.

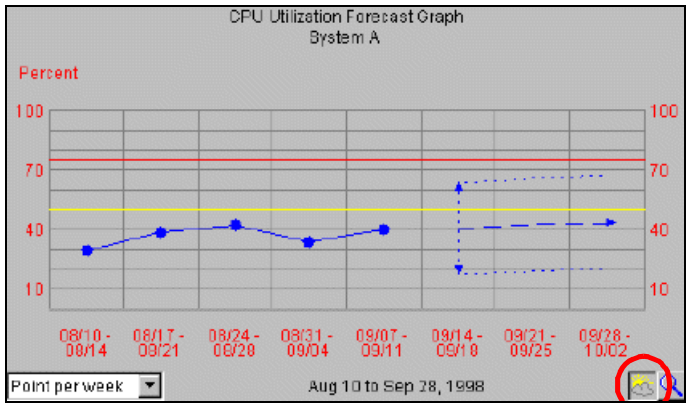

Figure 9-47 Forecast Graph

Capacity Manager will display one of two warnings if your forecast is not valid. Invalid forecasts should not be used to make decisions about your systems:

- “Data collection period too short for a valid forecast.” To generate a valid forecast, you need at least 21 days of data.
- “System 'X' does not have enough data for forecasting”, or “Multiple systems do not have enough data for forecasting.”

One of these two messages will appear when you have a sufficiently long period for data collection, but one or more monitors were not on for at least 50% of the time during the data collection period.

Performance Analysis feature

Performance Analysis is a new artificial intelligence feature that probes for bottlenecks in system’s hardware performance, diagnoses the problem, and suggests ways to improve performance. The performance analysis algorithm is based on the experiences of experts. The algorithm can find many but not all system problems. A minimum of a month’s worth of data is needed to make accurate predictions.

Note: Performance analysis is only available on Windows NT and Windows 2000 systems.

The algorithm monitors four system functions:

- Memory
- Disk subsystem
- CPU
- Network

The report produced by the performance analysis function consists of two main sections:

- **Recommendations:** A summary of the actions recommended is shown in the top part of the Report window.
- **Details:** All analysis results are graphically represented in the bottom part of the Report window.

To see the results of the performance analysis on your data, click the **Performance Analysis** button on the tool bar or use the **ALT+N** keyboard combination. A window similar to Figure 9-48 appears.

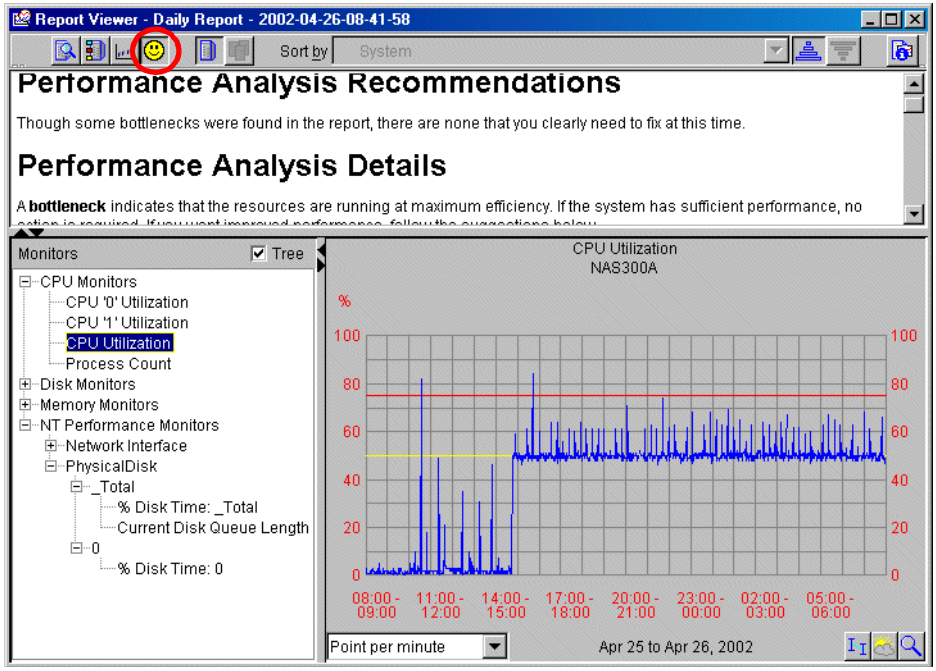

Figure 9-48 Performance Analysis Report

The report presents the bottleneck information first as a summary of the recommendations, then in a more detailed format, as shown in Figure 9-48. It also has links to the supporting graphic data. Keep in mind that bottleneck detection and analysis are complicated. If a monitor seems to be missing in one bottleneck, it may be because it is contributing to another one.

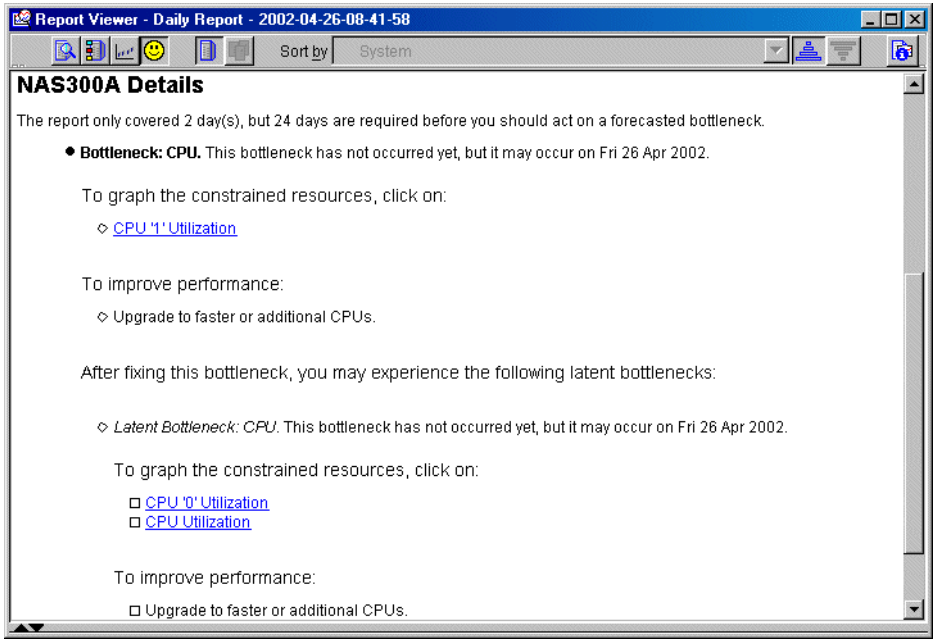

Figure 9-49 NAS300 Performance Analysis Report

The performance analysis report is available online and as an HTML file.

For more detailed explanation of Capacity Manager, see the redbook, *Tuning Netfinity Servers for Performance*, SG24-5287.

9.4.8 Cluster Systems Management (CSM)

You can use Cluster Manager to monitor and manage the Microsoft Cluster Services based functionality of your IBM TotalStorage NAS 300 and NAS 300G systems.

To start the Cluster Systems Management (CSM) task, just drag it from the Tasks pane and drop it into the chosen cluster system. The CSM window opens, as shown in Figure 9-50.

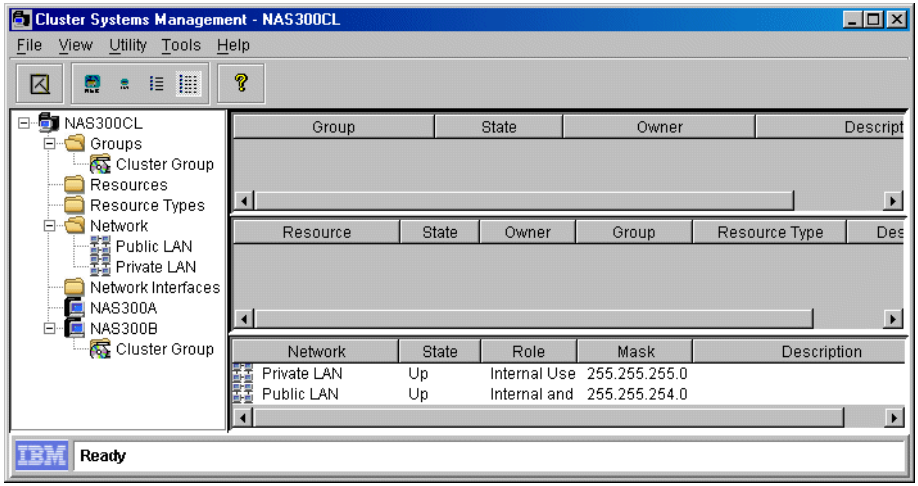

Figure 9-50 Cluster Systems Management window

The user interface displays cluster components, including nodes, groups and resources. You can work on a single cluster, or cluster components onto a node, using the drag-and-drop method.

A cluster name represents the top component in the Cluster Systems Management tree structure and owns all of the entities, such as resources, resource groups, nodes, networks, and network interfaces. You can create, delete or move cluster entities, as well as change their properties.

It also includes the Cluster Configuration Expert Wizard to create and define new resource groups. In the MSCS environment, defining high-availability resource groups is the most important task that must be completed. MSCS must know which resources make up the resource groups and know their dependencies on each other. You can use the Cluster Expert Wizard to support file-share, print-spooler, and Internet Information Server (IIS) resource groups. You can create and define new resource groups in existing resource groups.

Creating a new file share

Before creating a new file-share resource, make sure a physical disk drive is available:

1. Click **Utility —> Cluster Expert Wizard —> File Share**. The Expert Wizard window will open, as shown in Figure 9-51.

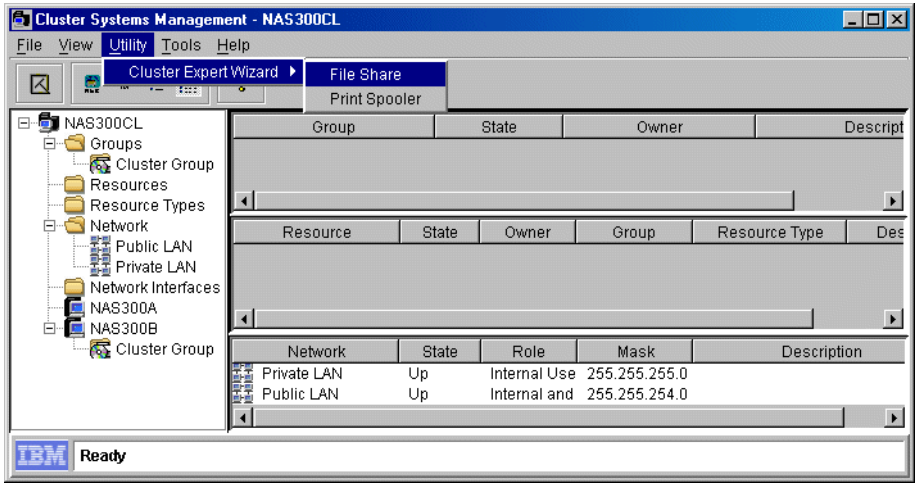

Figure 9-51 Starting the Cluster Expert Wizard

2. If the IP Address Range has not been set up before, you will be able to enter it in the next window (Figure 9-52).

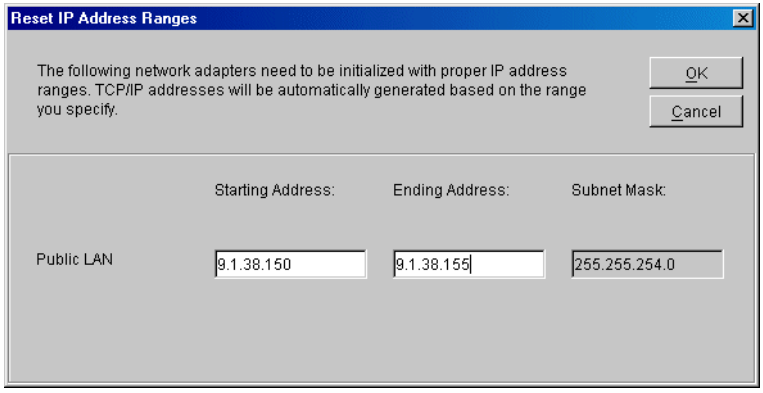

Figure 9-52 IP Address range

3. Click the **Create a new group radio** button, as shown in Figure 9-53.

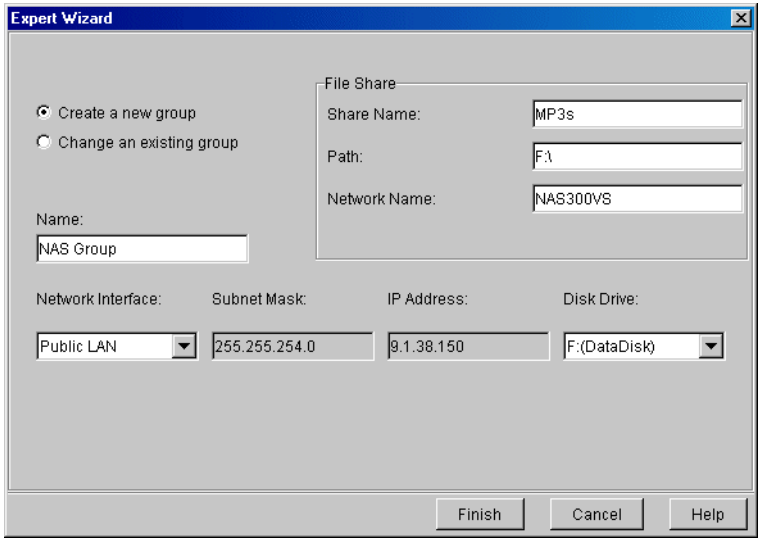

Figure 9-53 Cluster Expert Wizard

4. Type the group name in the **Name** field.
5. Type the Share Name, Path, and Network Name in the File Share fields.
6. Select the Network Interface from the drop-down list box.
7. Select the physical disk drive letter from the drop-down list box.
8. Click **Finish**.

The group name is displayed in the main CSM window, as shown in Figure 9-54.

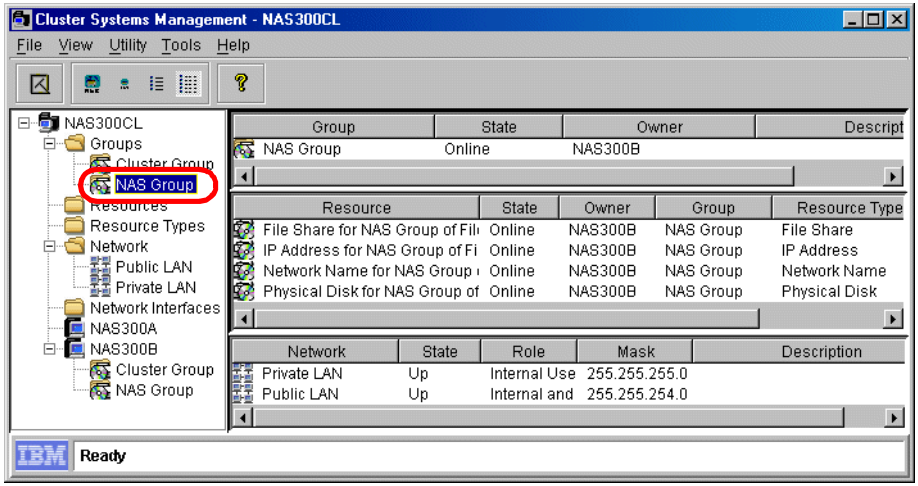

Figure 9-54 New Group in the CSM window

9.4.9 ServeRAID Manager

IBM Director Server Extensions provides both monitoring and remote control functions for ServeRAID. The ServeRAID Manager task lets you view and manage ServeRAID configuration. The ServeRAID Manager agent also automatically sends any RAID alerts to the IBM Director Server for processing using event action plans. No configuration is required to send these alerts.

The ServeRAID Manager task lets you perform functions similar to the CD-ROM-based stand-alone ServeRAID Manager, with an almost identical interface. However, some stand-alone ServeRAID Manager functions are not available with the ServeRAID Manager task in IBM Director:

- Configuring SNMP trap alerting
- Configuring Event Notification to ServeRAID Manager
- Configuring ServeRAID Manager Security.

To start the ServeRAID Manager task, just drag its icon from the Task pane onto the selected system, as shown in Figure 9-55.

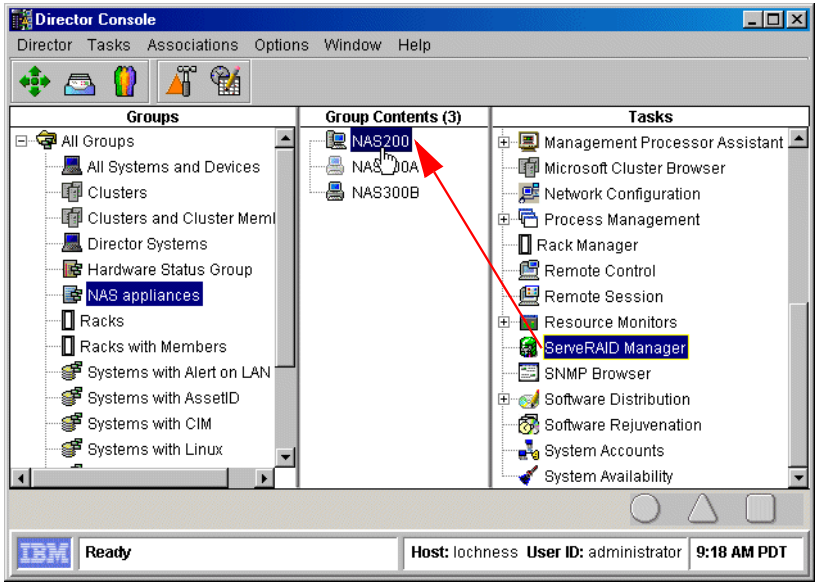

Figure 9-55 Starting ServeRAID Manager task

The ServeRAID Manager window opens. By clicking the + signs in front of the system name, controller, and array, you will be able to get details about the controller, arrays, and logical drives, as shown in Figure 9-56.

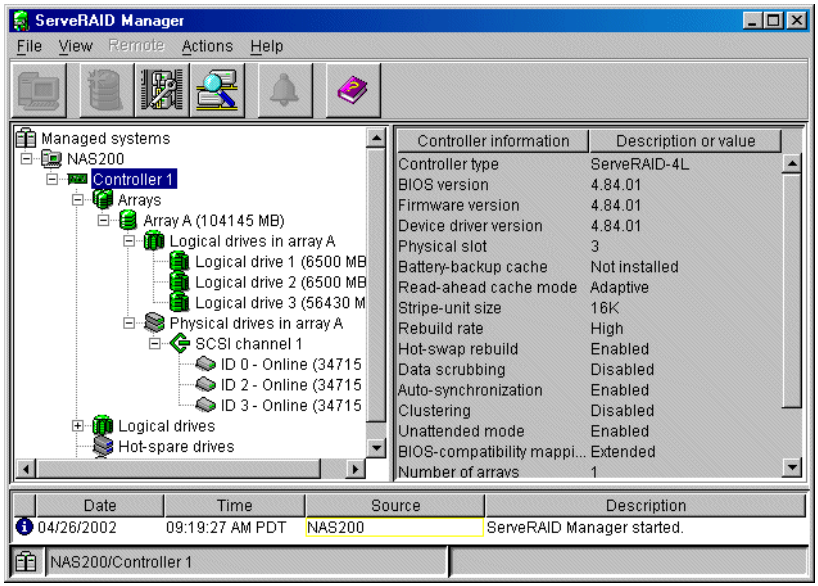

Figure 9-56 ServeRAID Manager window

Any action that happens to the RAID subsystem will be notified in the Event log, shown in the bottom part of the window. If you want to find out some details about a particular event, right-click it and select **Configuration Event Detail** (Figure 9-57).

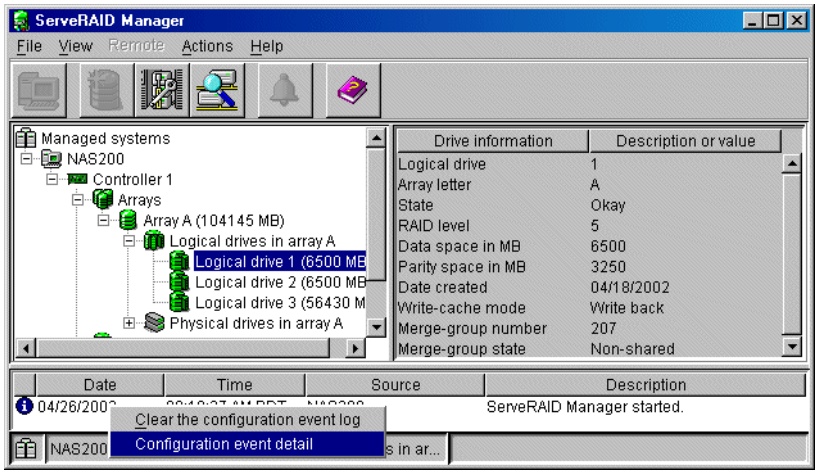

Figure 9-57 ServeRAID Event

The Event detail window will open and you will be able to see the date, time, and description of this event (Figure 9-58).

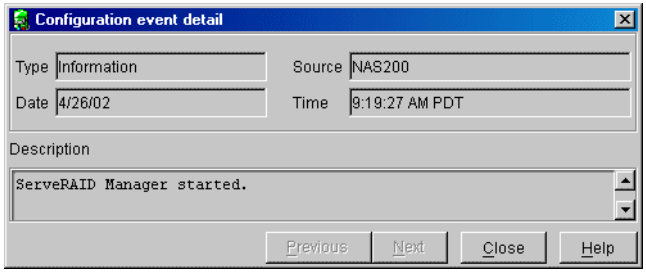

Figure 9-58 Event details

More detailed use of the ServeRAID Manager, specifically creation of the arrays and logical drives, is described in 9.4.9, “ServeRAID Manager” on page 327.

9.4.10 Usage tips for the NAS 100

Some of the IBM Director functions are not supported, but with a little bit of tweaking of your configuration, you can make them work. This goes beyond the supported scope of the NAS 100 system, and these tips will not be supported officially.

A typical example is the Remote Control function. Sometimes it would be useful to see what is happening on the appliance at that exact moment. As the NAS 100 is designed as a completely headless appliance, it doesn't have a video card built-in, and no PS/2 keyboard or mouse port are integrated. This however, is a prerequisite for Remote Control to function. However, during our testing, we found a little trick that can be used to make it work:

1. Enable USB support in BIOS (see "Enabling USB support" on page 592).
2. Connect a USB keyboard to the bottom USB port on the back of the NAS 100 appliance.
3. Install the correct keyboard driver (if the generic one is not supported with your type of the keyboard). A mouse does not need to be attached.
4. If needed, restart the NAS 100 device.
5. Connect to the appliance with Terminal Services Client and run the TWGIPCCF.EXE application. In the Network Driver Configuration window, remove the **Require User Authorization for Screen Access** checkbox.

Attention: This action opens a security hole into your system. We recommend that you use this only in emergency situations, and return to the normal settings immediately afterwards.

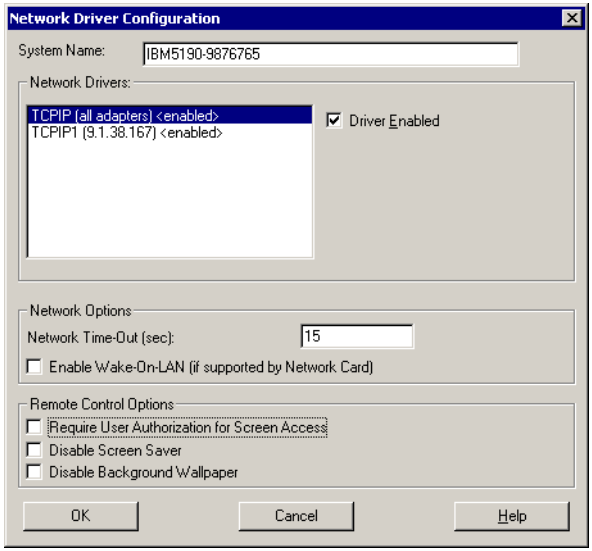

Figure 9-59 IBM Director — Network Driver Configuration

6. Now the Remote Control function should work as on any other controlled system. On the IBM Director Console, drag the Remote Control task from the right pane onto the NAS 100 appliance, as shown in Figure 9-60.

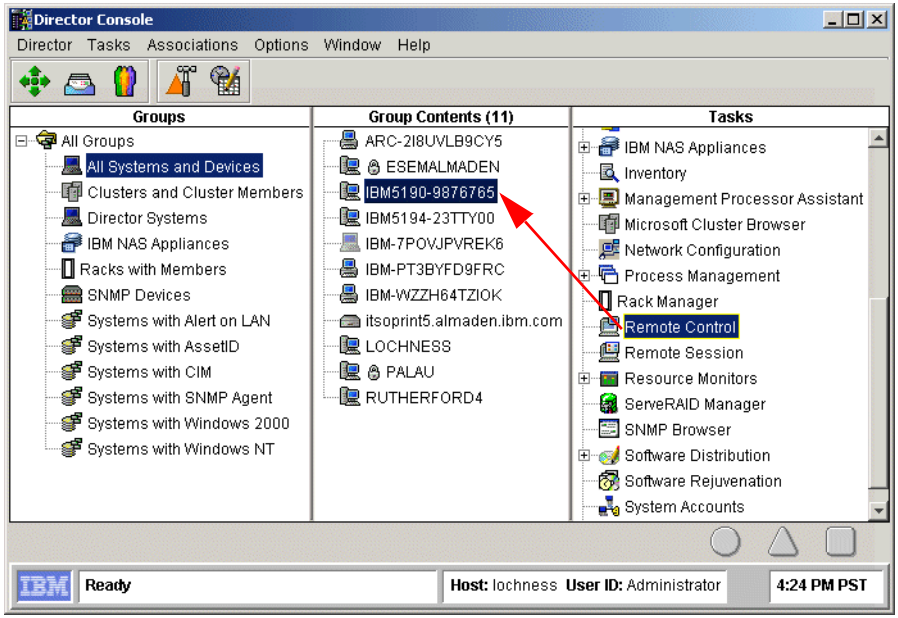

Figure 9-60 Starting Remote Control

A new window opens with the NAS 100 desktop presented (Figure 9-61).

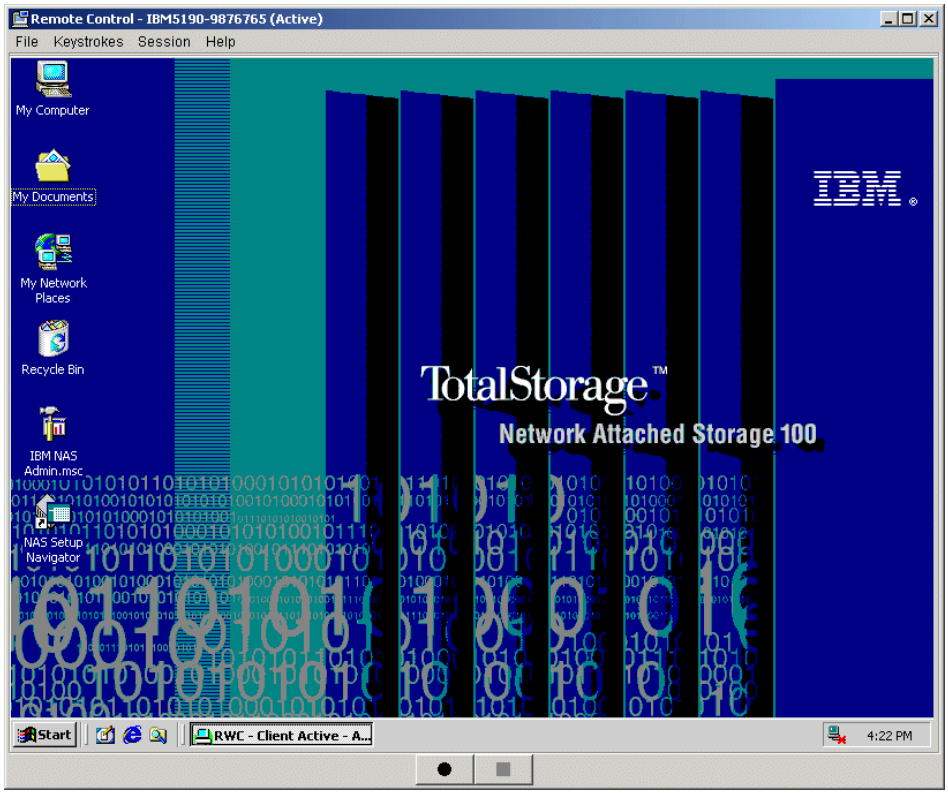

Figure 9-61 Remote Control of the NAS 100 appliance

9.5 How to install IBM NAS Extensions to IBM Director

The IBM NAS Extensions to IBM Director provide capabilities to IBM Director that are specific to the IBM NAS Appliances. This enables detection and grouping of NAS appliances from Director.

In this section we show you how to install these extensions to IBM Director. You have to install the extensions on the machine with the management console.

The files needed are available on the Supplemental CD1 that is shipped with the NAS 100 appliance. The subdirectory name is **/IBM NAS Extensions For IBMDirector**.

1. Insert the CD-ROM into the machine's CD-ROM drive. There is no need to stop the IBM Director services, because the installation program will do this by itself.

2. Start the **setup.exe** program. The install shield wizard will start (Figure 9-62).

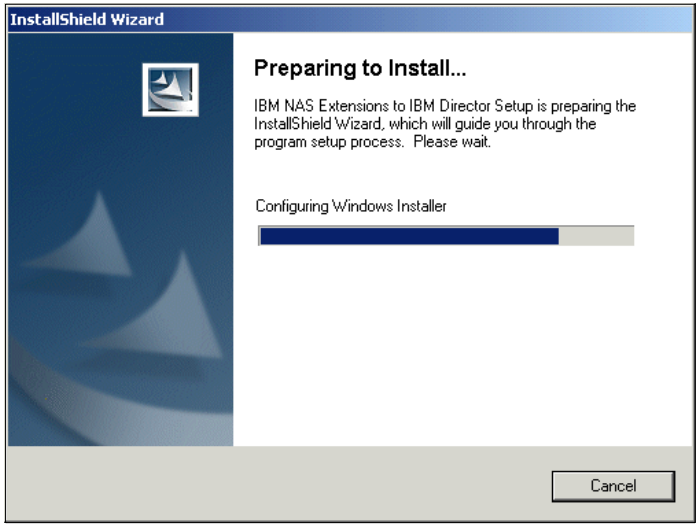

Figure 9-62 Install shield wizard preparing to install extensions

3. The Welcome screen will appear. Click **Next** to continue (Figure 9-63).

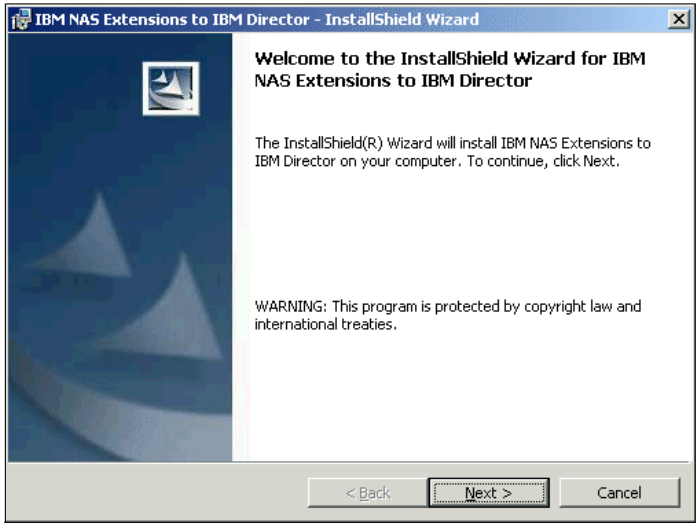

Figure 9-63 Welcome screen

4. Choose a destination folder or just click **Next** to continue (Figure 9-64).

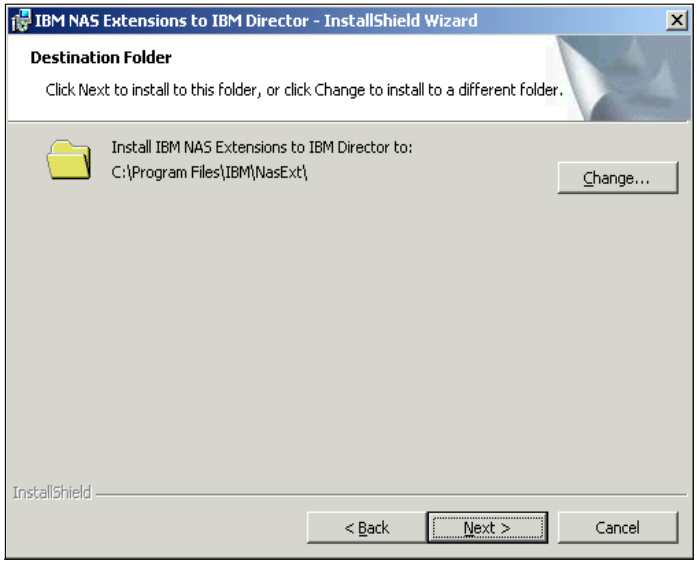

Figure 9-64 Chose destination folder

5. The installation wizard is now ready for the installation process and asks for any changes. Click **Install** to start the installation (Figure 9-65).

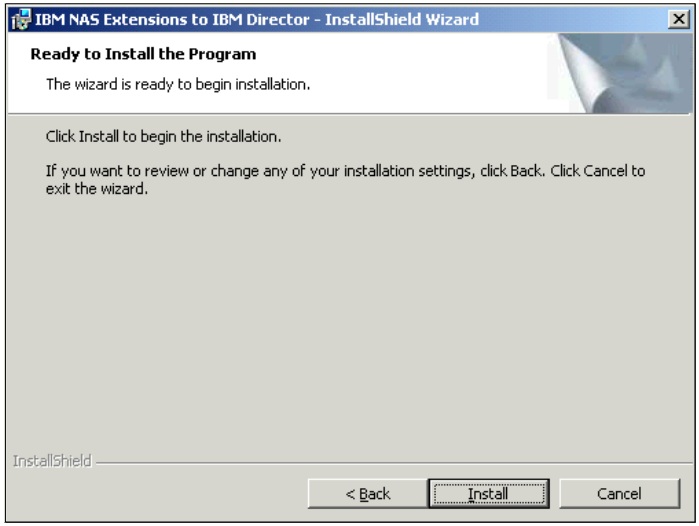

Figure 9-65 Ready to install the extensions

6. The installation wizard will install the extensions now. First it is generating scripts (Figure 9-66).

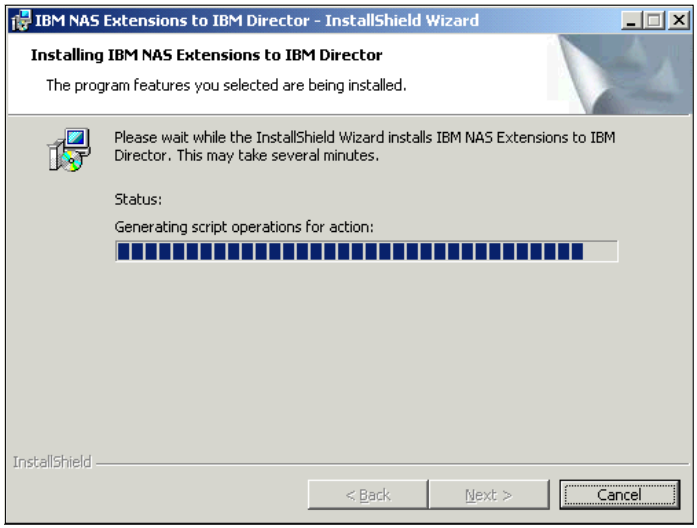

Figure 9-66 Generating scripts

7. A command line interface will open and the wizard starts its scripts (Figure 9-67, Figure 9-68).

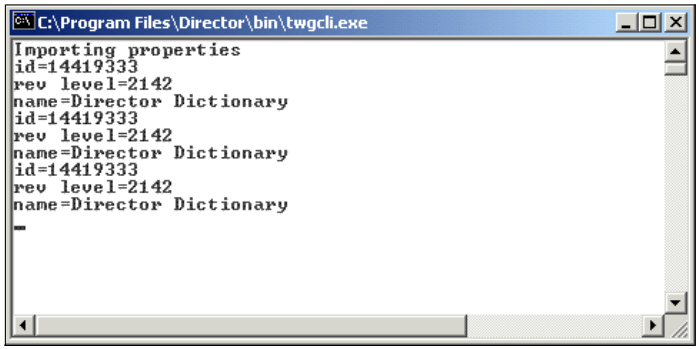

Figure 9-67 Wizard is executing its scripts

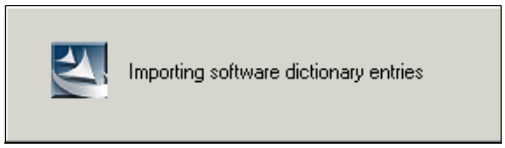

Figure 9-68 Importing software dictionary entries

8. The wizard will restart the IBMDirector services now (Figure 9-69, Figure 9-70).

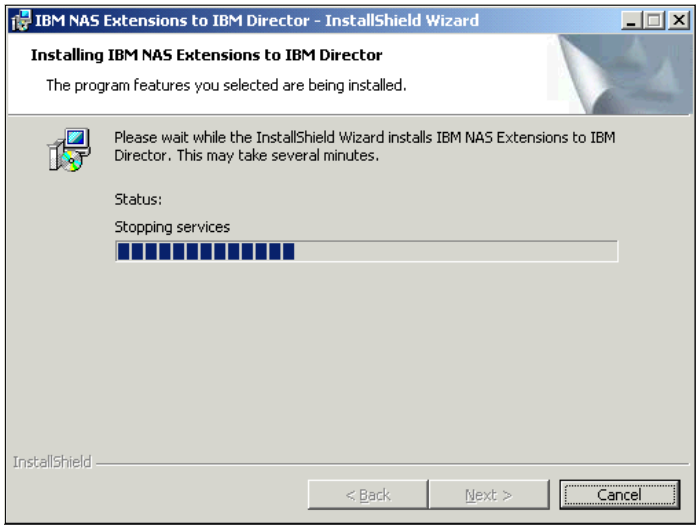

Figure 9-69 Wizard is stopping the services

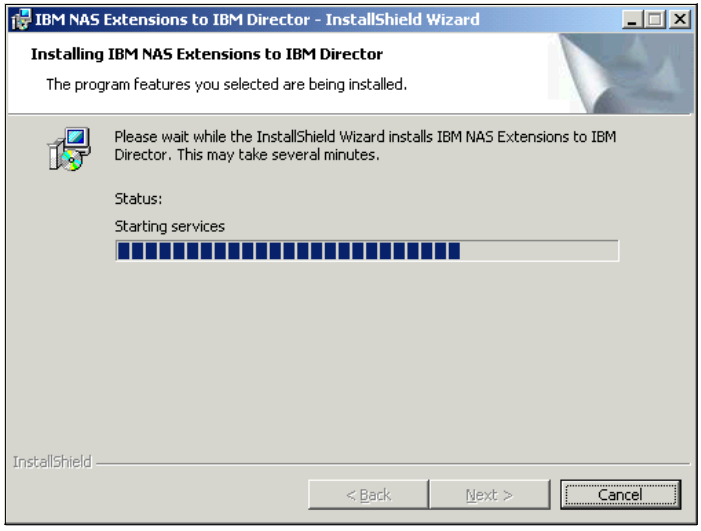

Figure 9-70 Wizard is starting the services

9. Now the backup files will be removed (Figure 9-71).

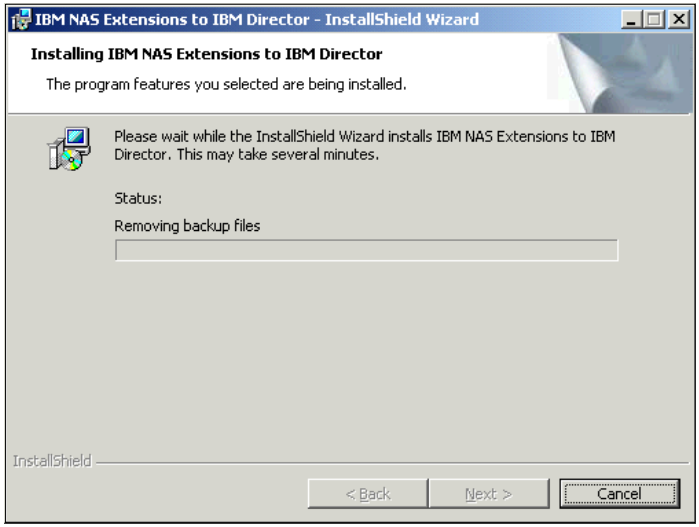

Figure 9-71 Removing backup files

10. The installation is now complete (Figure 9-72).

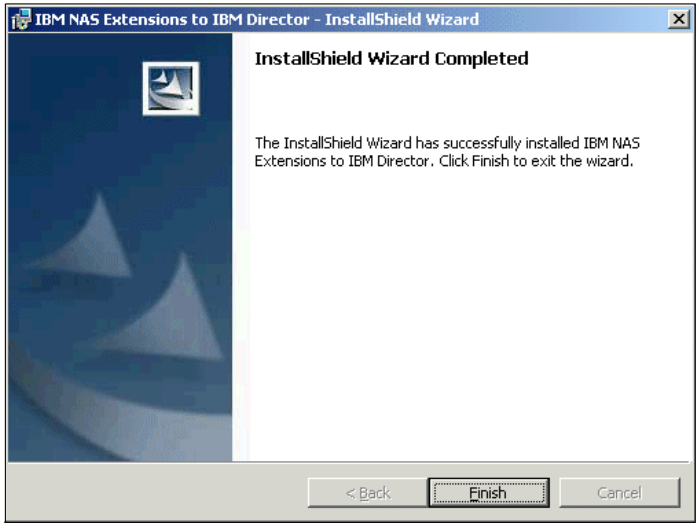

Figure 9-72 Installation completed

Now start the IBM Director to see the new features (Figure 9-73). In the Groups window you can see a new group called **IBM NAS Appliances**. Whenever IBM Director recognizes a NAS appliance in the network, it will put it into this group.

Note: This feature is not supported by the current BIOS version of the NAS appliances but will be implemented soon. The IBM Director will show the NAS appliances in the **All Systems and Devices** group.

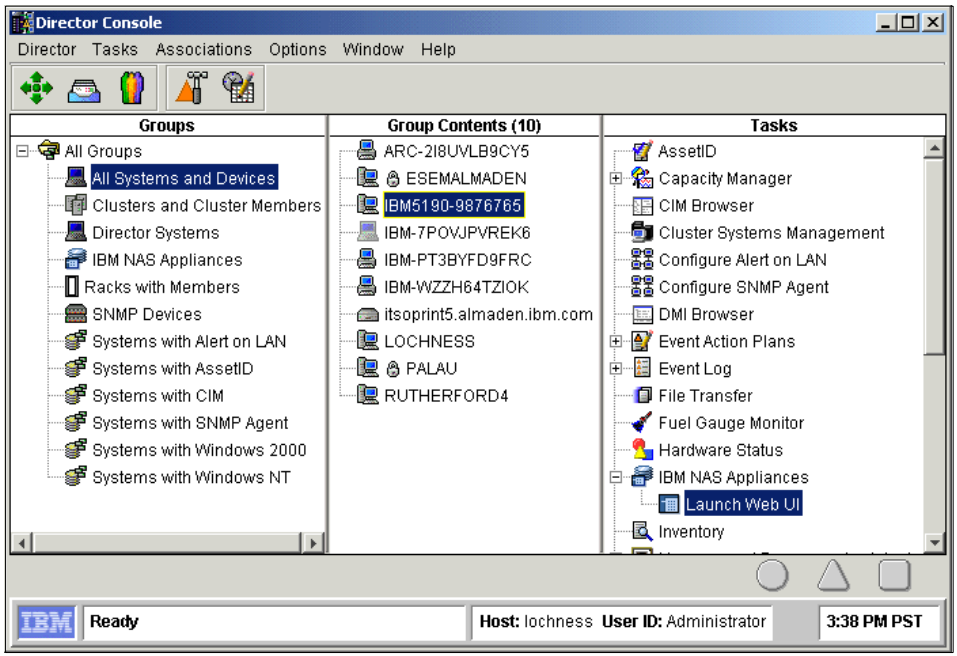

Figure 9-73 Director console with new tasks for NAS appliances

You can see anther new feature in the Tasks window. This has been implemented to support the Web UI from the NAS appliances. Open the directory IBM NAS Appliances in the Tasks window and drag the Launch Web UI to the NAS 100 icon. The NAS Web UI will open in a new screen (Figure 9-74).

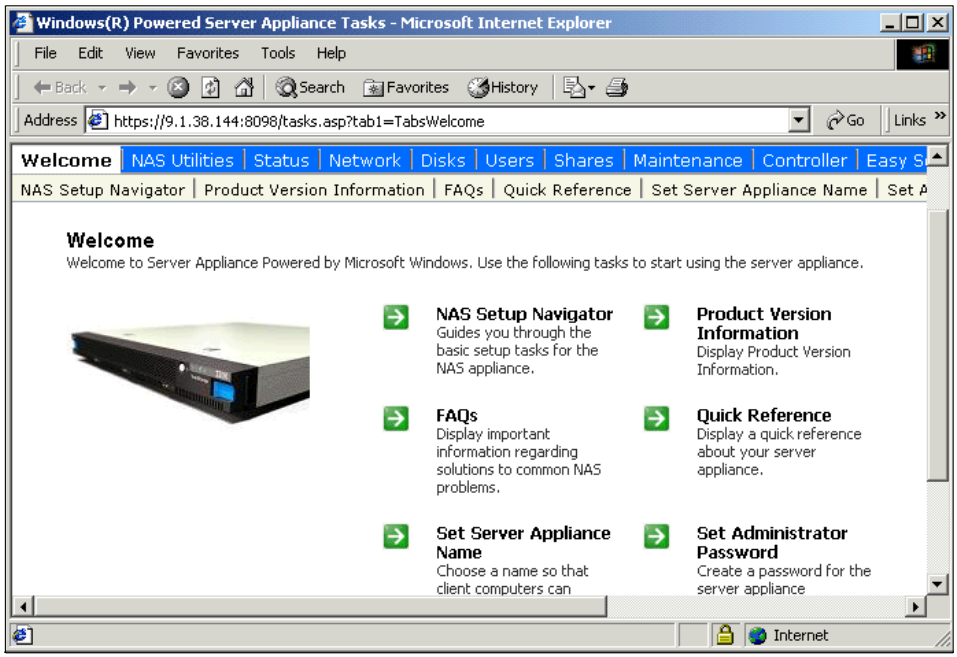

Figure 9-74 NAS Web UI

9.6 Microsoft Multiple Device Manager (MDM)

Newly announced in NAS 100 is support for Microsoft's new Multiple Device Manager (MDM), which offers the ability to remotely manage devices.

Microsoft has just started shipping the Microsoft Device Manager, and it is not a part of the standard SAK code.

MDM provides centralized management of the system with system discovery capabilities, and can manage multiple server appliances from a central location.

Other features listed include auto-discovery of new devices, ability to track and audit operational histories, grouping capabilities for reports, and alerts based on aggregation of sets of devices.

MDM provides new abilities for administrators to manage multiple appliances and to have greater control over access to the devices. It offers the capability to manage multiple server appliances from a central location.

MDM features include:

- Auto-discovery of new devices added to the network
- The ability to track and audit operational histories of the devices
- Grouping capabilities for management purposes
- Updated security
- SAK Alert aggregation for sets of devices

9.6.1 NAS 100 and MDM

The NAS 100 is remotely managed from the Headquarters LAN by the Utility Administrator.

This customer segment often has no IT skill at all in the remote LAN locations, and sometimes these locations are manless operations.

Using the 10/100/1000 Copper Gigabit Ethernet adapters, the customer can centrally manage many multiples of NAS 100s at remote locations.

Some of the tools that can be used include a Web browser, IBM Director, Microsoft Terminal Services, IBM Advanced Appliance Configuration Utility, and Microsoft Multiple Device Manager. SNMP and standard MIBs are also supported. In the following figures we show some of the features and their uses:

- Software distribution (Figure 9-75):
  - Uses MS software install tools
  - Distributes upgrades simultaneously

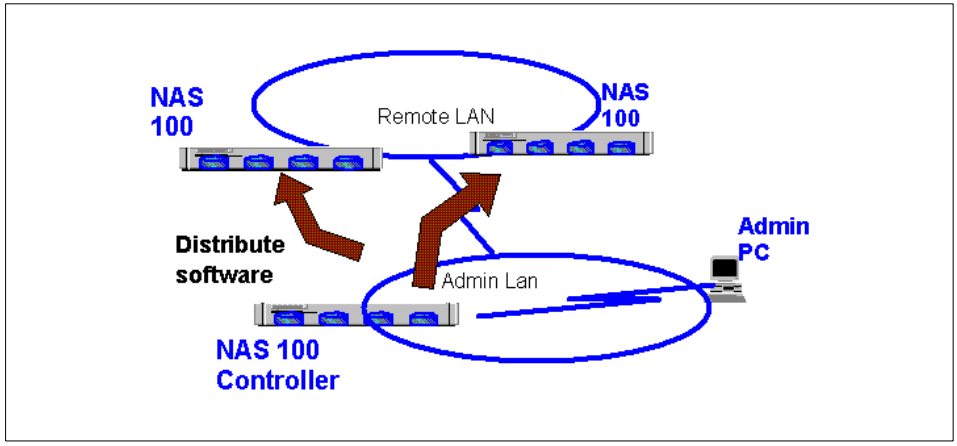

Figure 9-75 Distribute software

- Change configuration (Figure 9-76):
  - Create local users and groups
  - Join domains or workgroups
  - Change passwords
  - Create shares

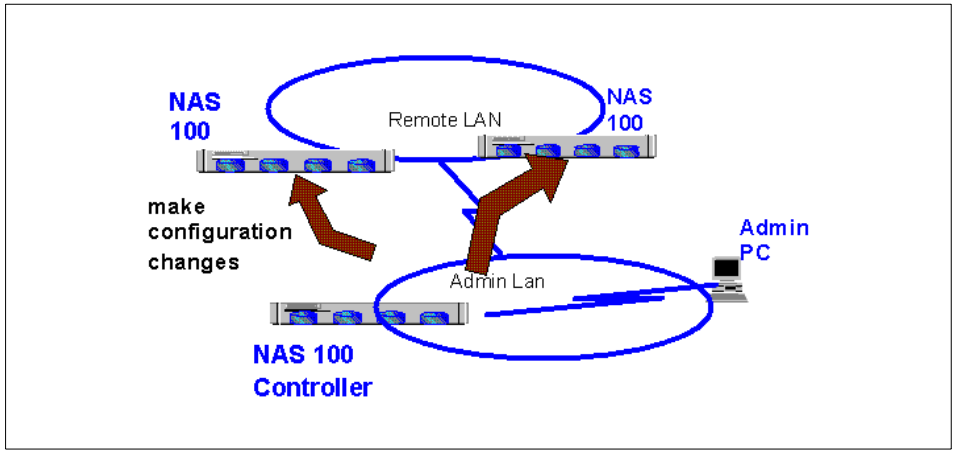

Figure 9-76 Change configuration

- Run and schedule jobs (Figure 9-77):
  - Disk tools
  - Track job history

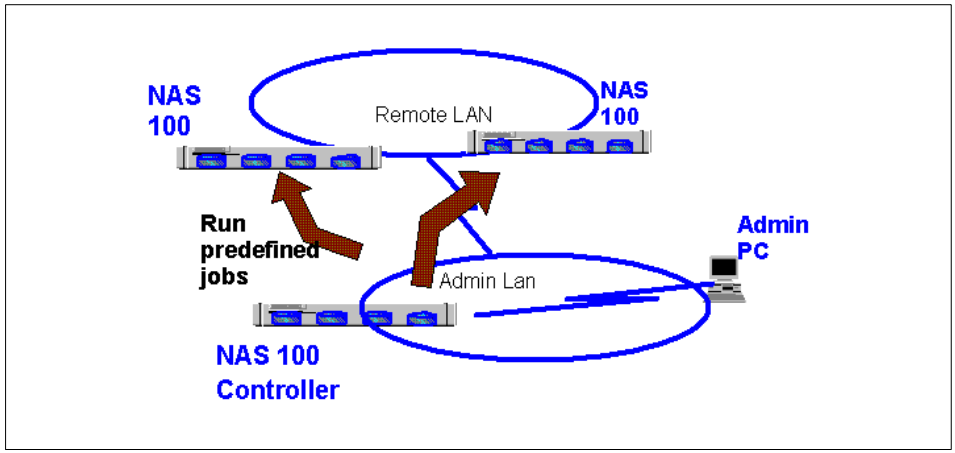

Figure 9-77 Run and schedule jobs

- Alerts (Figure 9-78):
  - Receive e-mail alerts
  - Receive error log entries

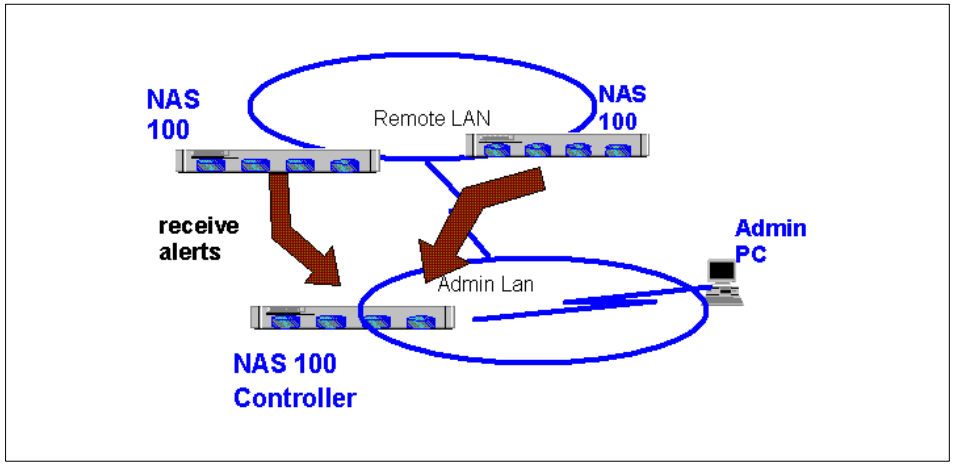

Figure 9-78 Receive alerts

- Here are some general characteristics:
  - Auto-discovery of other controlled devices is possible.
  - Controller must be booted first.
  - Appliances send out a special network packet which can be identified by the controller.
  - Multicast packet sent to 239.255.255.250 port 1900.
  - Other devices which are not in the multi-cast domain can be added by explicitly specifying the server name.

Warning: Do not leave a device in an “able to be controlled” state without an active controller in sessions with it. This would be a security issue.

9.6.2 Controller installation on NAS 100 appliance

You can access the Microsoft MDM settings through the NAS Web UI.

Check the box **Install Controller** to install MDM on the appliance. After that the NAS 100 has to reboot.

You do also have the possibility to control NAS 100 from another appliance:

1. To give the box the ability to be managed from another box, choose **Be Controlled** from the Welcome screen of NAS Web UI. The **Be Controlled** tab will be shown on the screen (Figure 9-79).

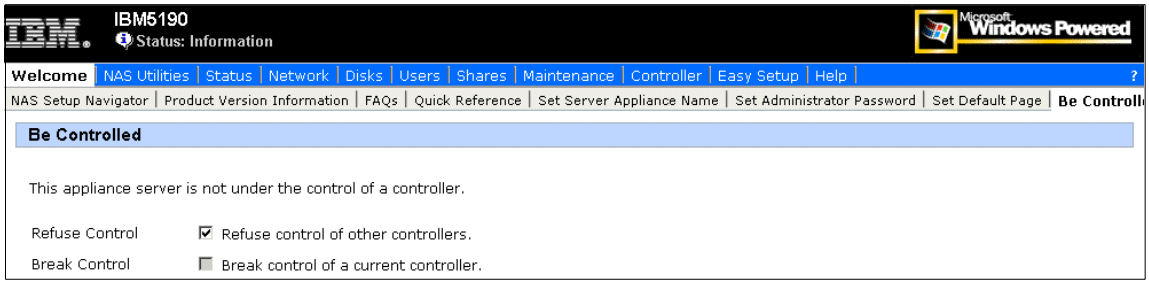

Figure 9-79 Be controlled tab

2. Uncheck the **Refuse Control** option and reboot.
3. Now the appliance is able to be controlled from another NAS appliance.

9.6.3 MDM functions

The MDM Welcome Screen gives you a survey of the MDM tools (Figure 9-80).

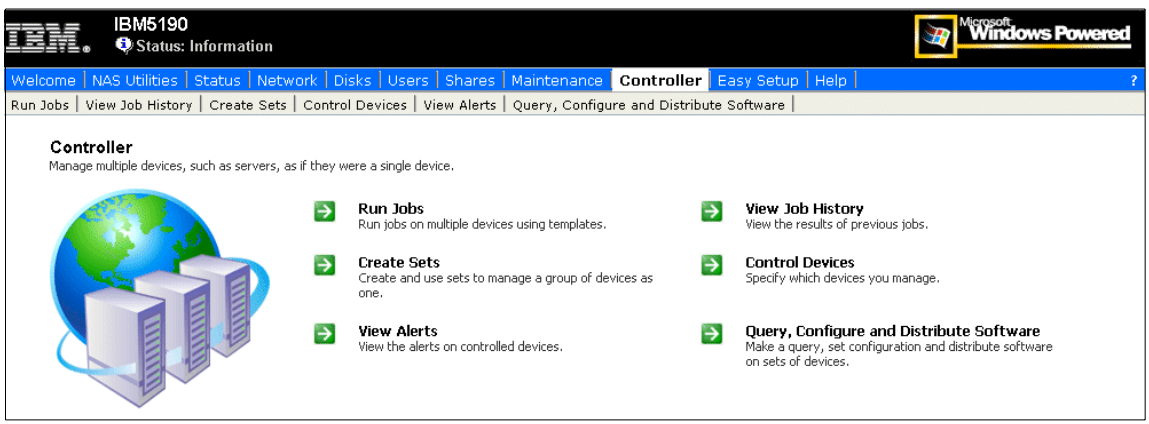

Figure 9-80 MDM Welcome Screen

Run Jobs

In the **Run Jobs** tab, you can configure new jobs for remote appliances or sets (Figure 9-81).

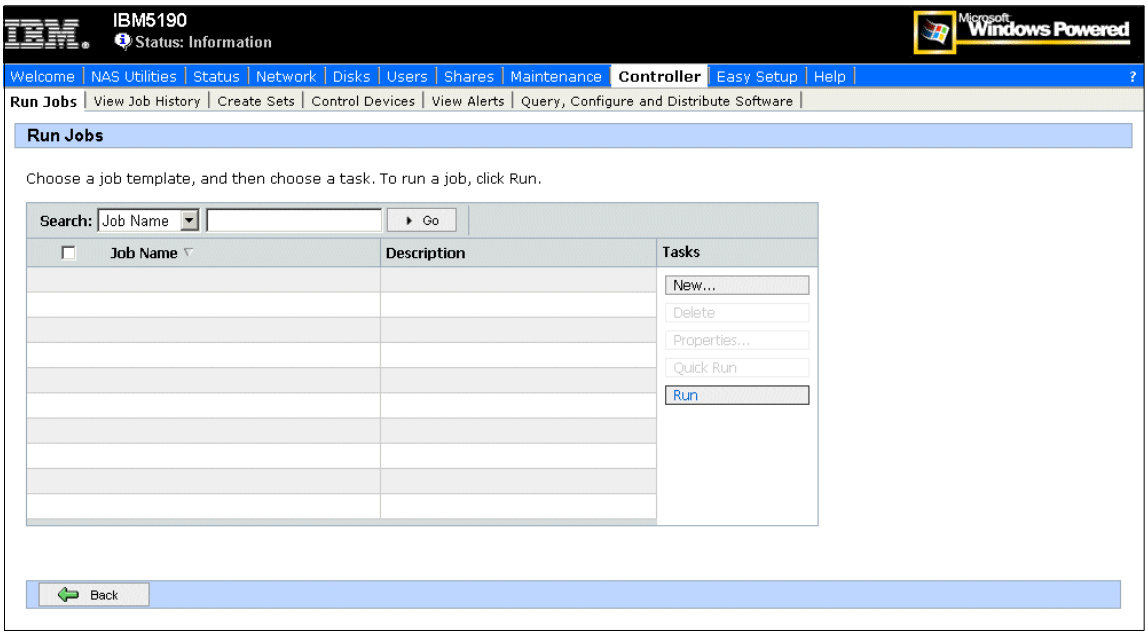

Figure 9-81 Run Jobs tab

Click the **New...** button and you will get the **Job Template Wizard** (Figure 9-82). You can create new jobs and schedule them easily.

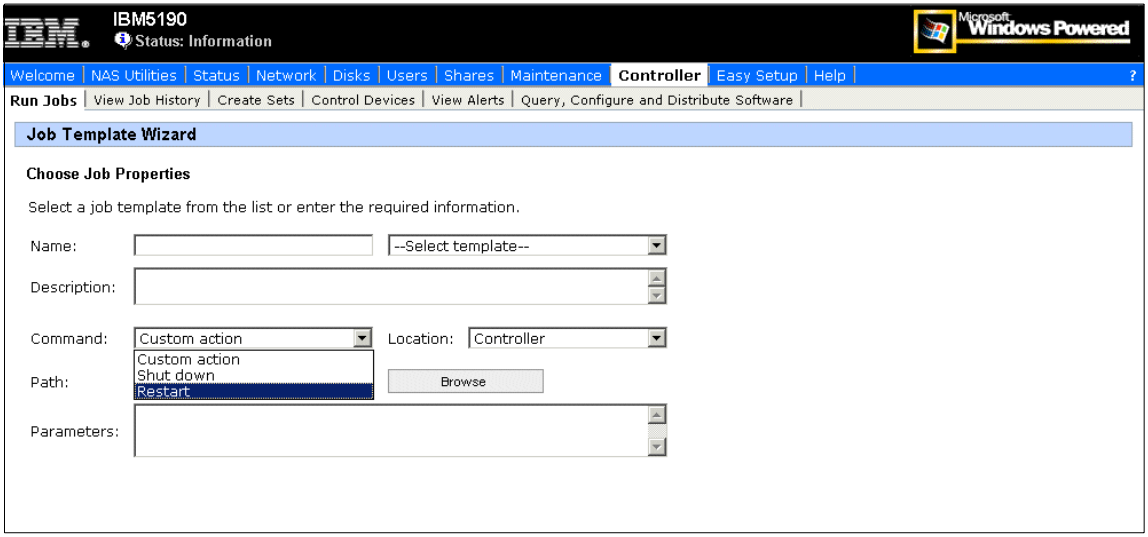

Figure 9-82 Job Template Wizard

Create Sets

In the **Create Sets** tab, you are able to configure new sets of NAS appliances to run or configure all at one time (Figure 9-83). We will create a set called IBMAppliance.

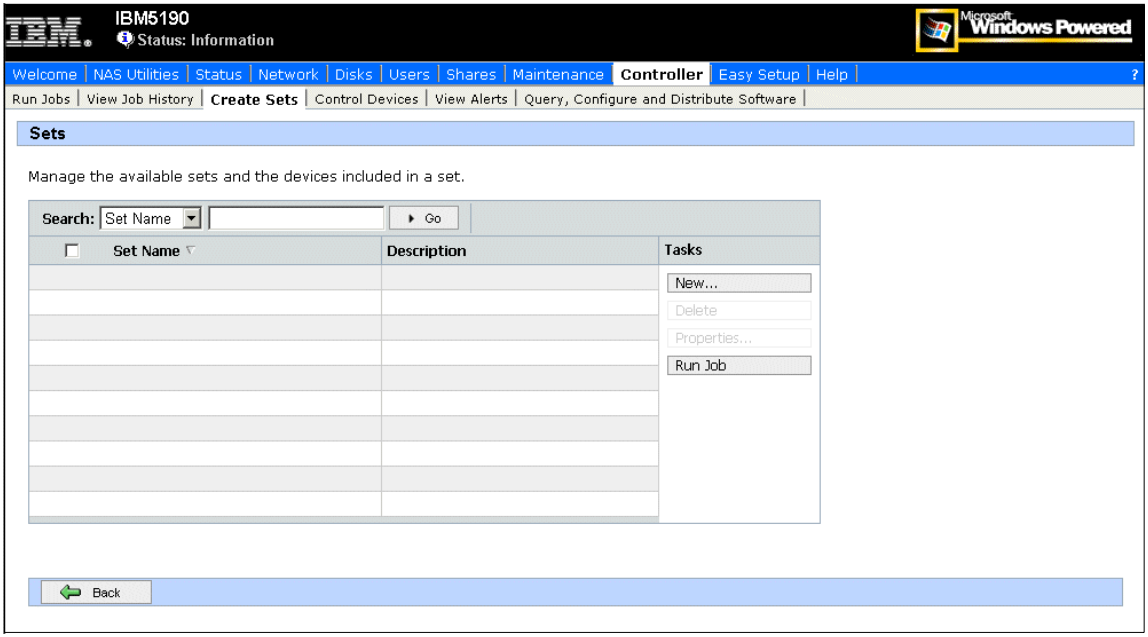

Figure 9-83 Create Sets tab

1. Click the **New...** button and the Create Set screen will open (Figure 9-84).

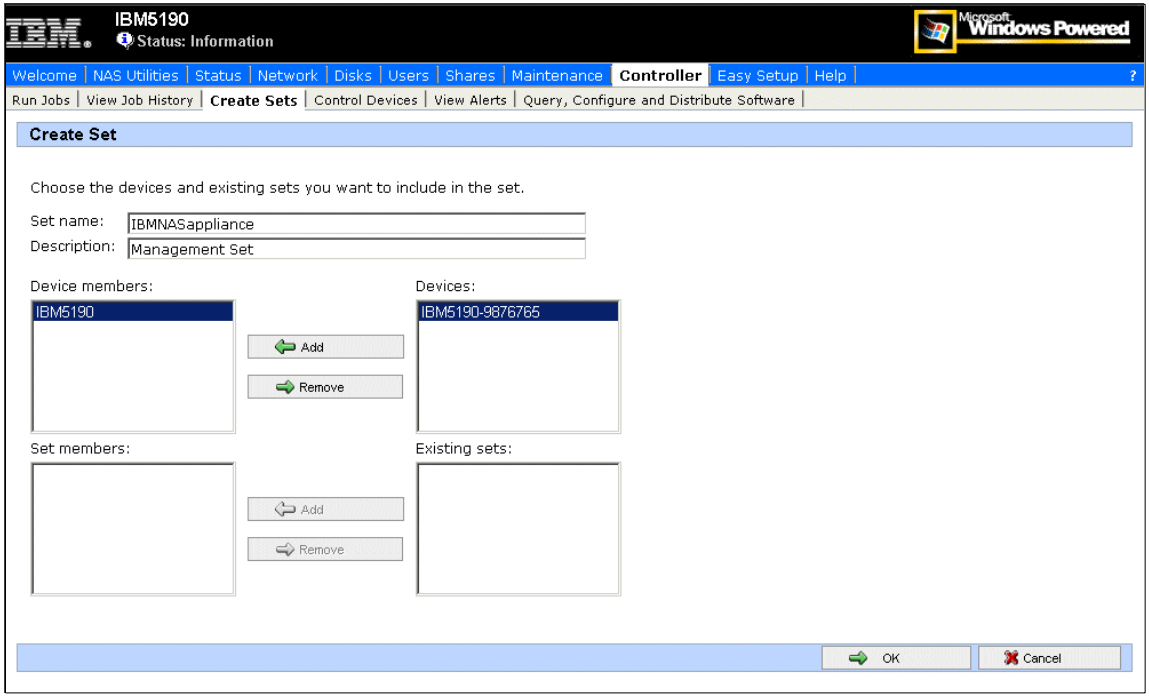

Figure 9-84 Create Set screen

2. Insert a Name and Description for the set and add NAS appliances to your set by clicking the appliance and clicking the **Add** button.
3. Create the set by clicking the **OK** button (Figure 9-85).

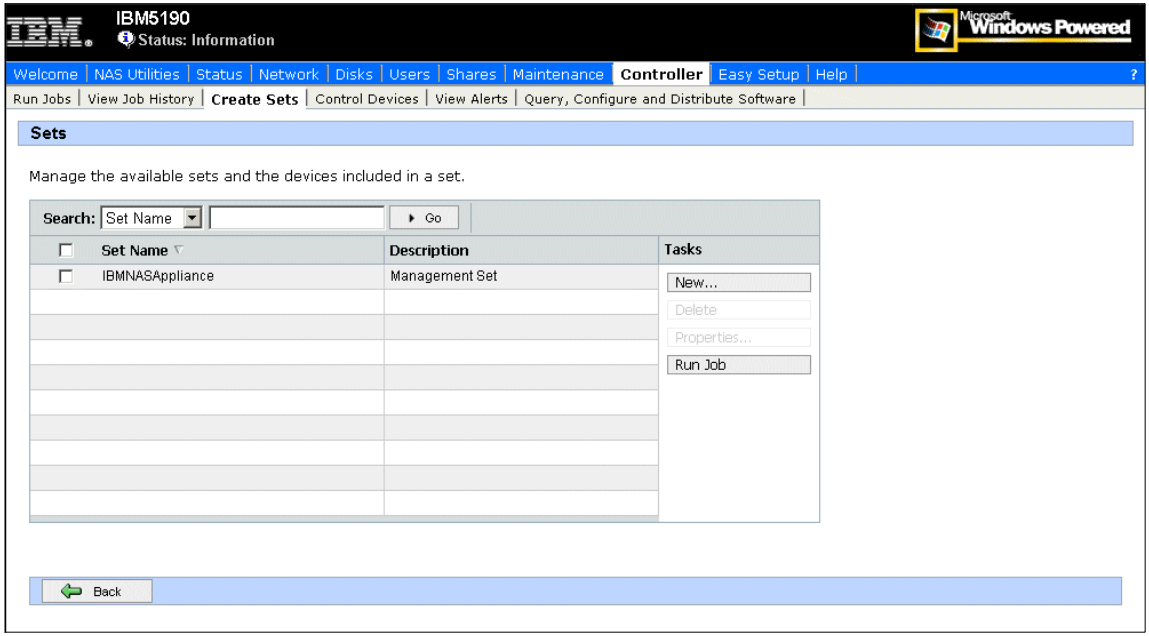

Figure 9-85 Set screen with the created IBMAppliance set

Control Devices

You can see the devices that can be managed by your controller in the **Control Devices** tab. New discovered devices can be identified by reading the description (Figure 9-86).

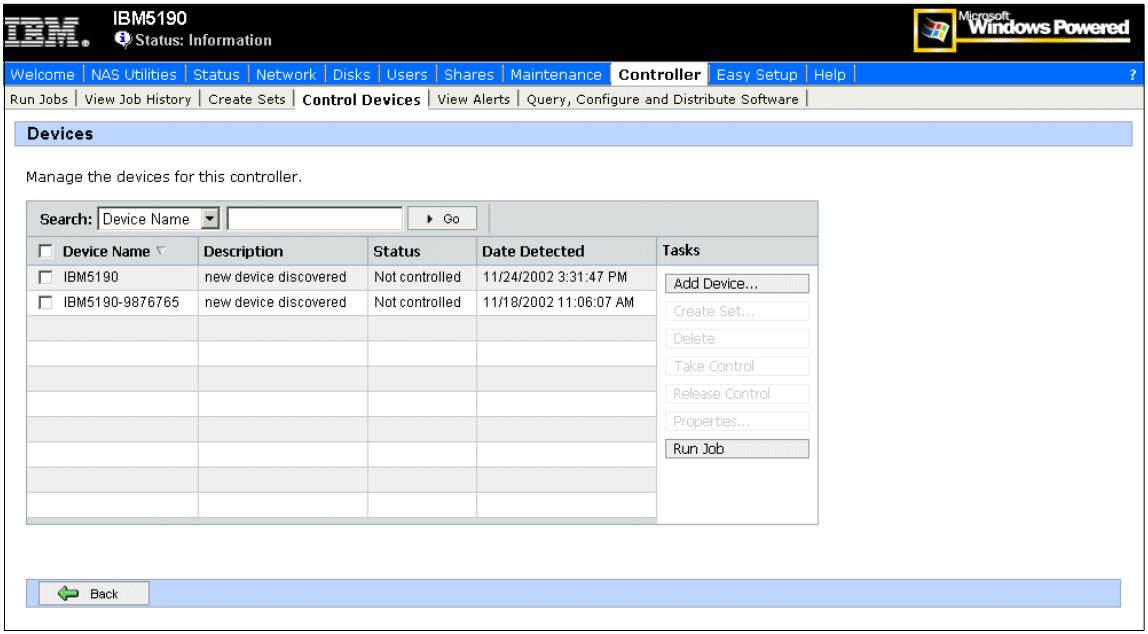

Figure 9-86 Device tab

Query, configure, and distribute software

You have the capability to query for appliances, set configuration, and distribute software on sets of devices (Figure 9-87).

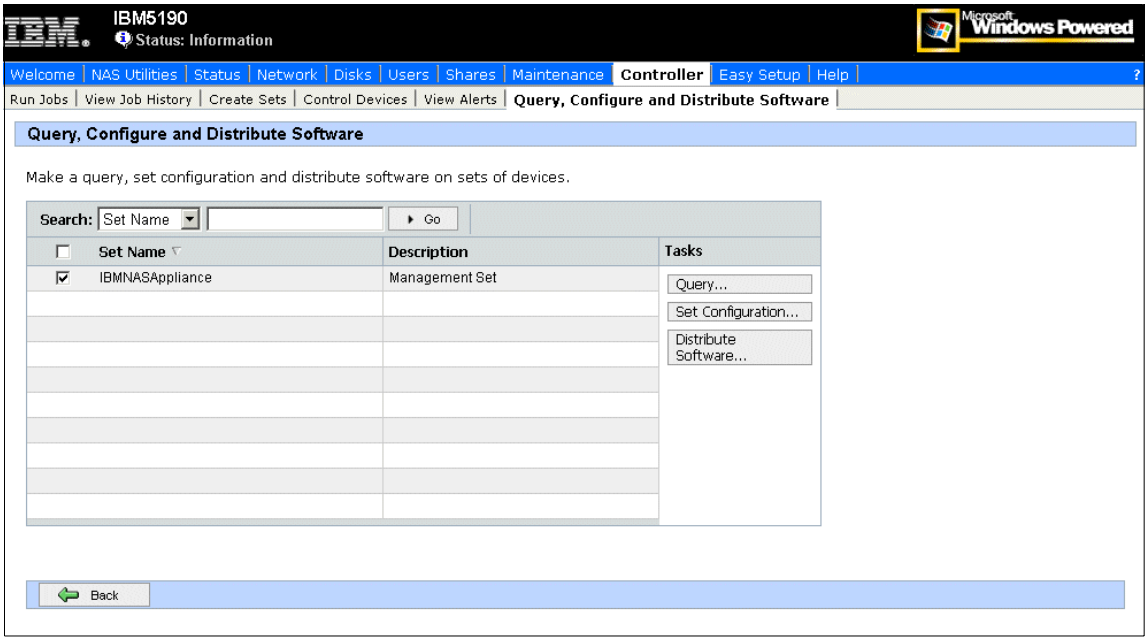

Figure 9-87 Query, Configure, and Distribute Software tab

If you check the box beside the newly created set, it is possible to make configuration updates and have them distributed to the specified NAS Appliance set (Figure 9-88).

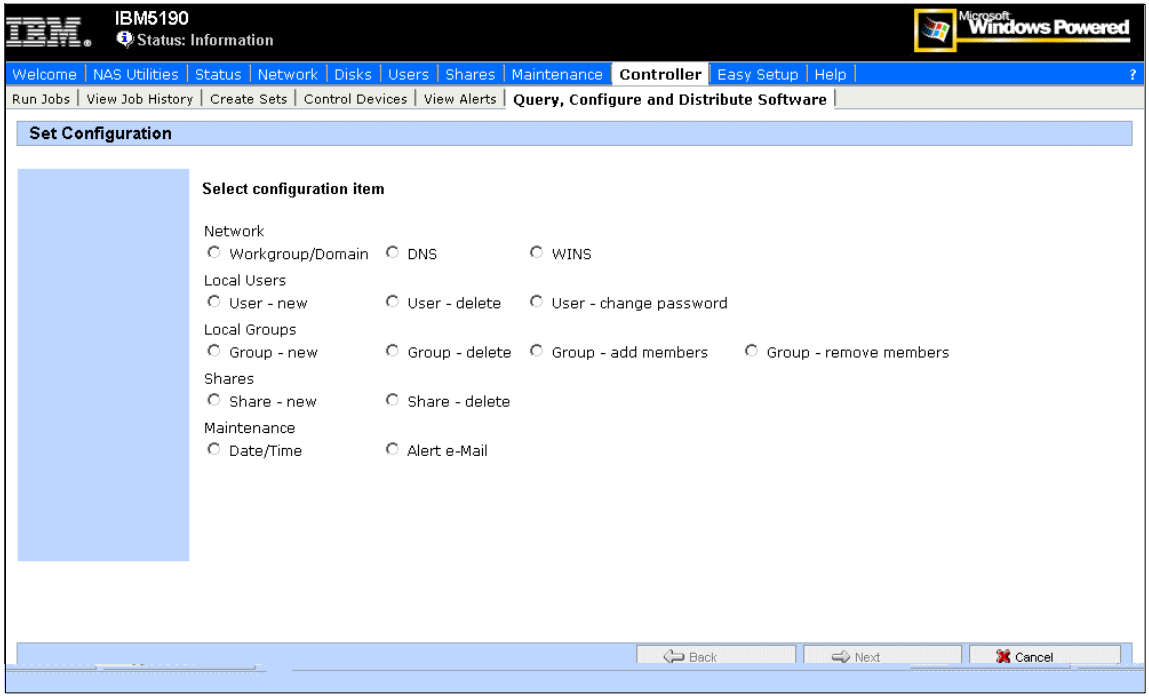

Figure 9-88 Configuration options for the set

But if you choose **Distribute Software...** you can easily distribute Software without an additional tool (Figure 9-89).

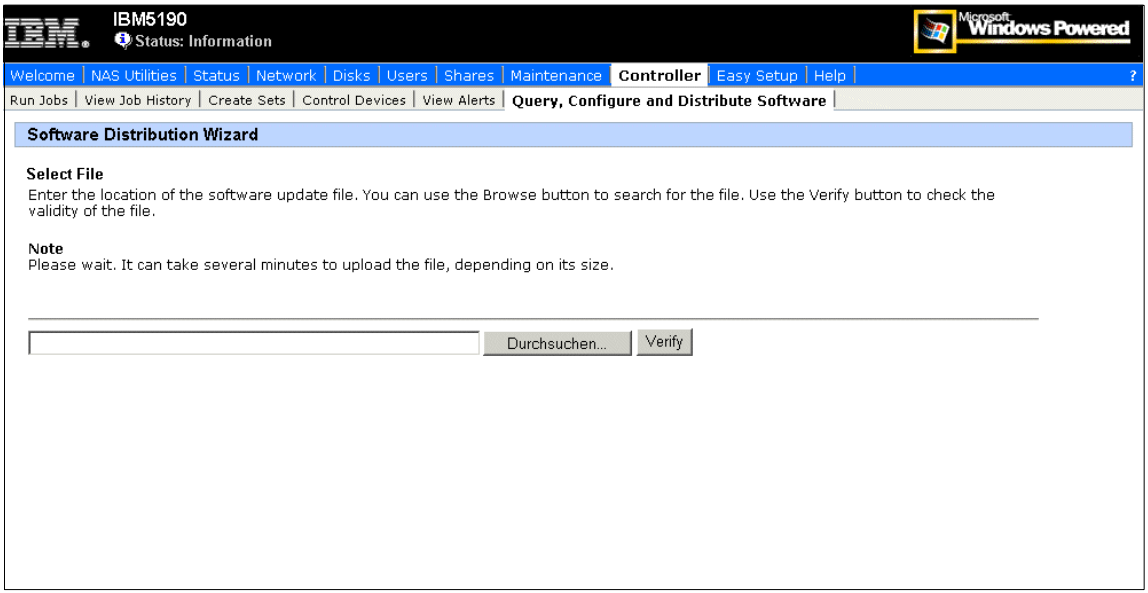

Figure 9-89 Software Distribution Wizard

For more Information about the new Microsoft Multiple Device Manager (MDM), refer to:

http://www.microsoft.com/windows/Embedded/sak/evaluation/whatsnew/default.asp#mdm

10

Cross platform storage

We now have plenty of storage space that is owned by our NAS system, and we want to give the other machines in our network access to it. In this chapter we provide all of the information necessary to use NAS appliances from Windows and UNIX clients. Your operating environment influences the type of file sharing to be configured on the NAS appliance:

- If your operating environment consists entirely of Windows clients, consult the material in 10.1, “File sharing for Windows clients” on page 354 and 10.2, “Accessing the shares from our Windows clients” on page 370.
- If your operating environment consists entirely of UNIX clients, the material in 10.3, “File sharing for UNIX clients” on page 372 and 10.4, “How to configure Services for UNIX (SFU)” on page 385 will be helpful. You should also review 10.4.5, “Accessing the shares from our UNIX clients” on page 421.
- If your operating environment consists of a mix of UNIX and Windows clients, in addition to 10.4, “How to configure Services for UNIX (SFU)” on page 385, you should carefully read the material in 10.4.4, “Mapping the Gateway for NFS share from a Windows client” on page 419, which describes how to access Network File System (NFS) shares from Windows clients.

Finally, in “Accessing the shares with the Samba client” on page 423, we show how to use the Samba client to connect to the NAS appliance in Linux and AIX hosts.

10.1 File sharing for Windows clients

This section describes how to create shares for Windows clients. To ensure access for Windows users, we recommend having the NAS appliance join a Windows NT Domain or Active directory. However, if you do this, you must also create a separate account in NAS appliances.

10.1.1 CIFS file sharing in a non-clustered environment

The process of creating a share in a Windows server is an easy one. These are the steps you should follow to share a folder:

1. Login on the NAS 100 with the appropriate rights to create a share. The *administrator* account can always be used for doing this.
2. Create a folder that you want to share.
3. Right-click the folder and click **Sharing** to get the Sharing dialog box.
4. When you are creating a share for a complete disk, the Sharing dialog box shows that there is an administrative share name in place. The administrative share name ends with a dollar sign ($), as shown in Figure 10-1.

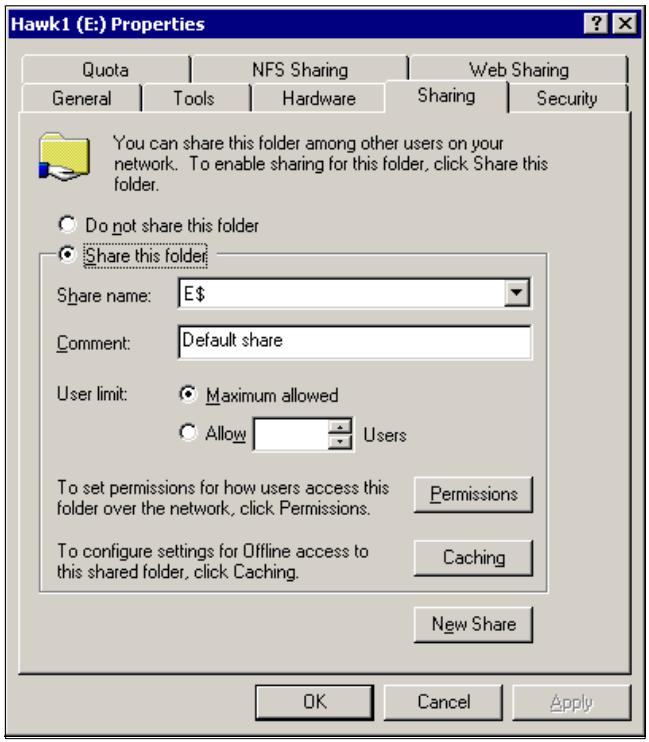

Figure 10-1 Administrative share

Important: To ensure access for Windows users other than administrators, you should always define an explicit share. Although on Windows NT and 2000 there is already a hidden administrative share present by default, you should not use it as a share for users. Instead, you need to create designated user shares.

5. In order to set up a share that would work well for us later on, we selected the option **Do not share this folder** and clicked **Apply** (Figure 10-2).

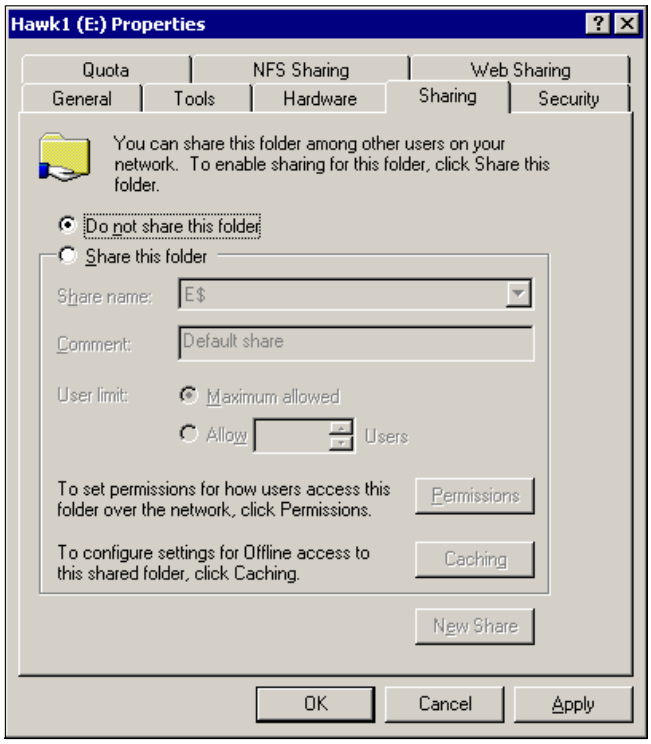

Figure 10-2 Getting rid of the administrative share

6. Next we selected the option **Share this folder** and supplied a share name (Figure 10-3).

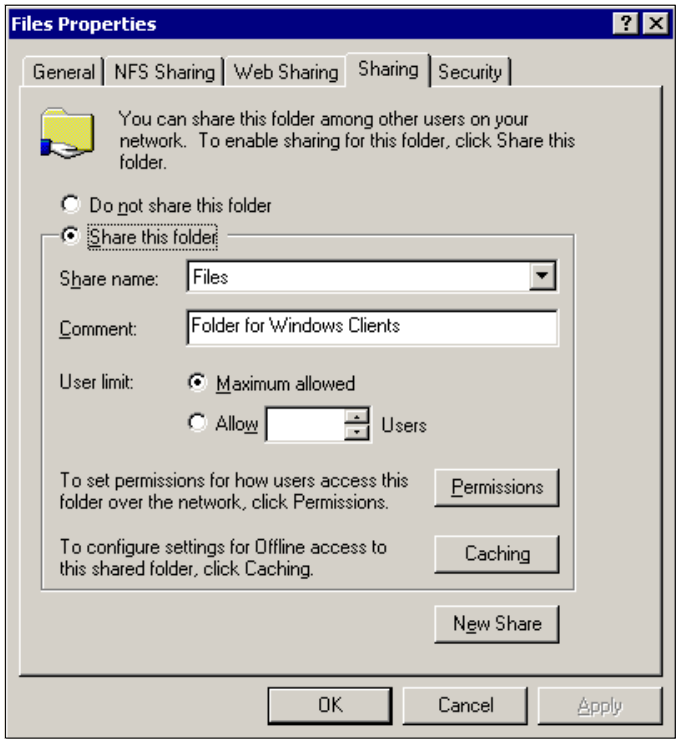

Figure 10-3 Sharing a folder

7. Before accepting this share, we clicked the **Permissions** button so security for the share could be adjusted to meet our needs (Figure 10-4).

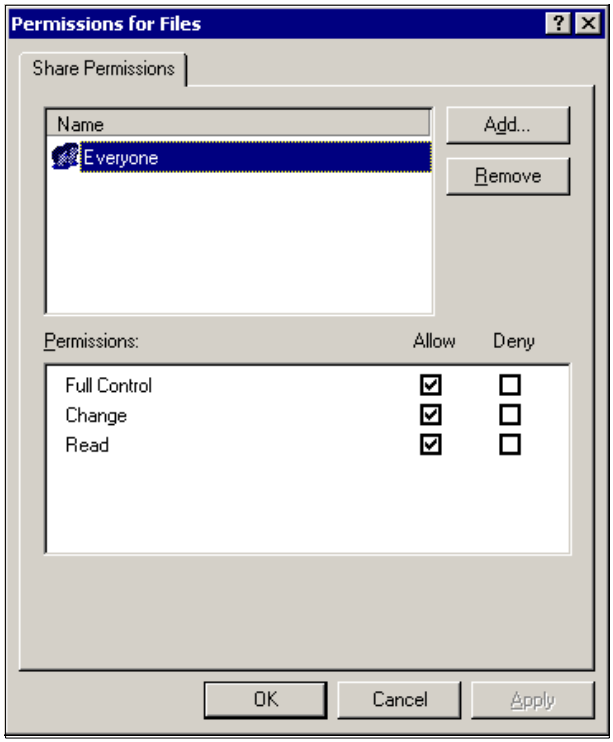

Figure 10-4 Permissions for the Windows share

8. You can also modify the NTFS permissions. Click the **Security** tab (refer back to Figure 10-3 on page 357) and then click **Add**.

Important: Permissions on Windows shares are a combination from network share and local NTFS file system share permissions. Those permissions work cumulatively.

9. Choose all users and groups that you want to grant access to this folder (Figure 10-5).

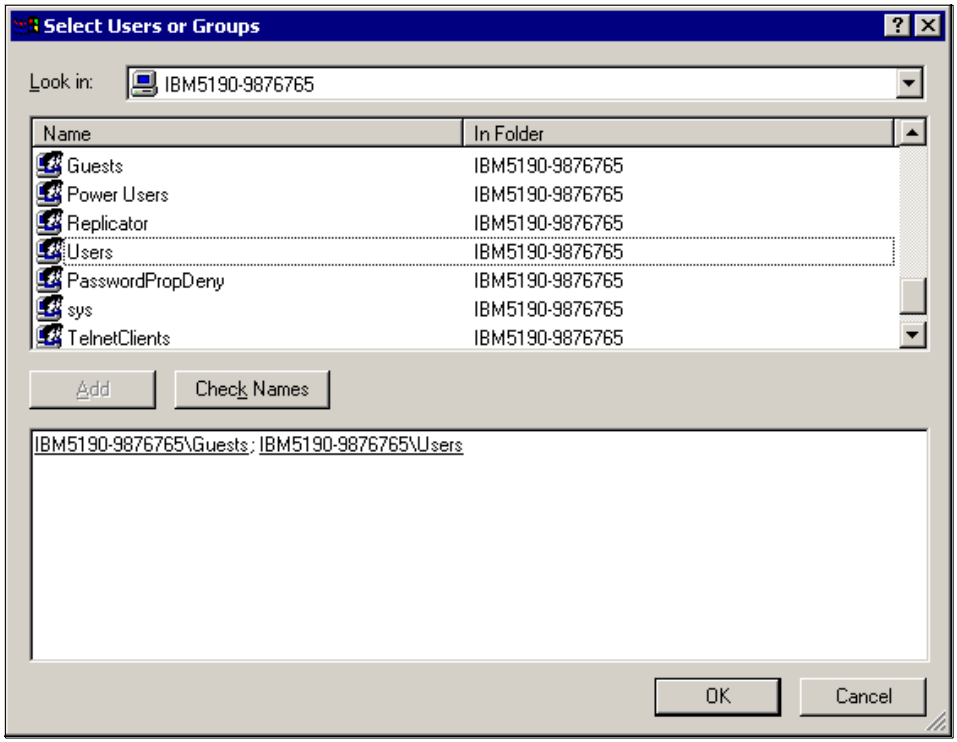

Figure 10-5 Select users

10. Configure the correct permissions for those users and groups, and finish by clicking **Apply** (Figure 10-6).

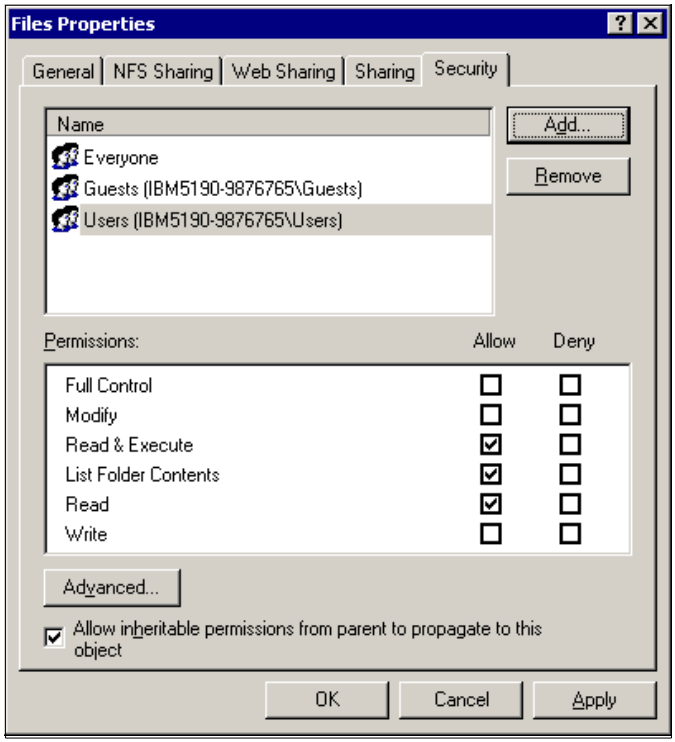

Figure 10-6 Modifying permissions

Once the shares were ready, they appeared, as “handed out”, in the Windows Explorer.

To check your current Windows shares from a DOS prompt, you can use the **`net share`** command to display all your current shares.

10.1.2 CIFS file sharing in a clustered environment

This section will describe the steps necessary to create a CIFS share on a clustered system as the IBM TotalStorage NAS 300.

Important: You must define file shares for Windows users in Cluster Administrator, to allow the file shares to fail over between nodes if a node goes down. Any file shares defined via Windows Explorer will not fail over as they are only valid for the node on which the share was created.

The creation of clustered file shares involves dependencies on other resources in the group as a physical disk, a static IP address, and a network name. The various dependencies allow resources that are defined to the same disk group to start in the correct sequence. The dependencies also assure necessary access for the given resource.

The steps on this page guide you through creating the resources and dependencies.

Note: To follow this procedure you can be direct attached to the NAS device or to manage it through a Terminal Services connection.

Creating the IP address resource

These are the steps for creating the IP address resource:

1. Open the Cluster Administrator (Figure 10-7), in the IBM NAS Admin tool and select **Cluster Tools —> Cluster Administration**.

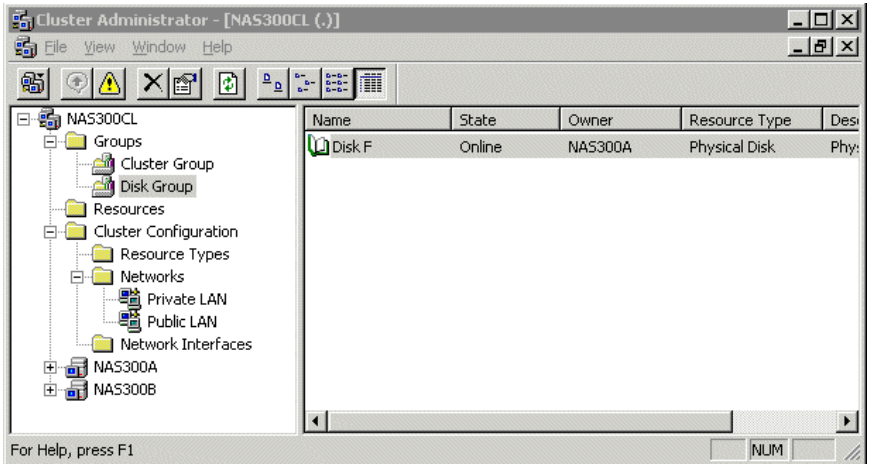

Figure 10-7 Cluster Administrator

2. From Cluster Administrator (Figure 10-7), select **File —>New—> Resource** (Figure 10-8).

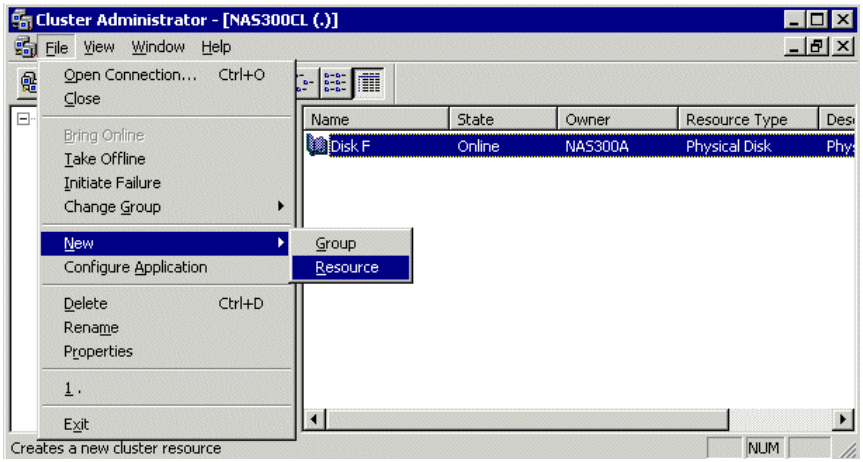

Figure 10-8 Creating New Resource

3. Enter the IP Address resource name, select **IP Address** as the resource type, select the disk group the resource should reside in as the Group, then click **Next** (Figure 10-9).

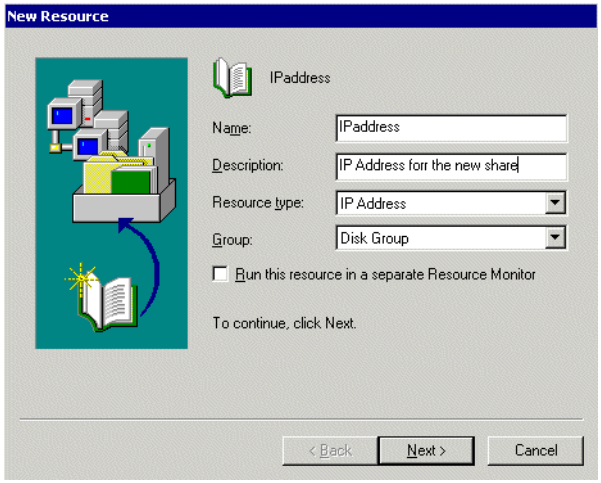

Figure 10-9 Defining an IP Address Resource

4. Add both nodes as possible owners using the **Add** button, then click **Next** (Figure 10-10).

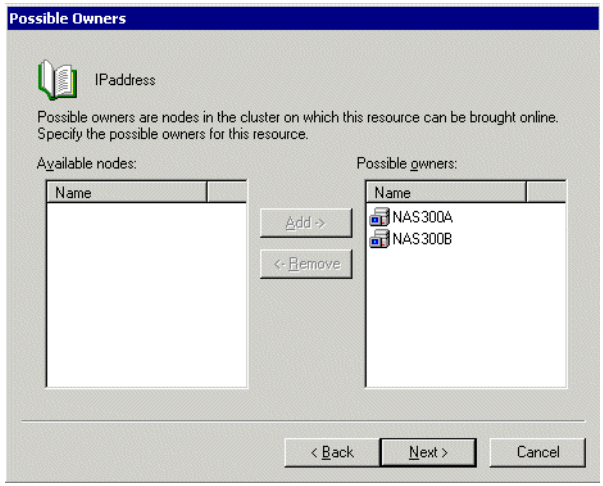

Figure 10-10 Choosing possible owners

5. Leave the Resource Dependencies window empty and click **Next** (Figure 10-11).

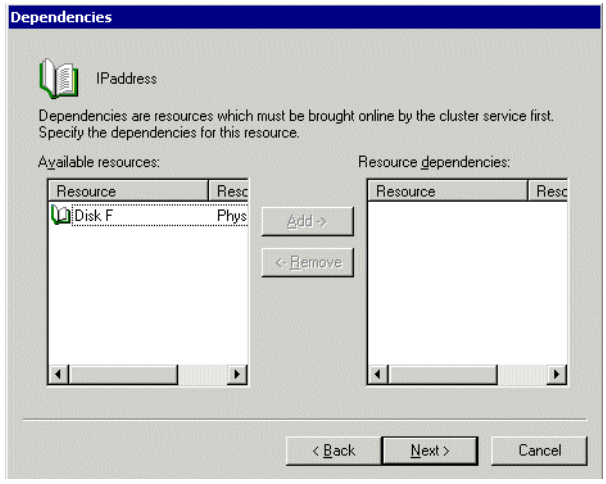

Figure 10-11 Setting resource dependencies

6. Enter your IP address and subnet mask and click **Finish** (Figure 10-12).

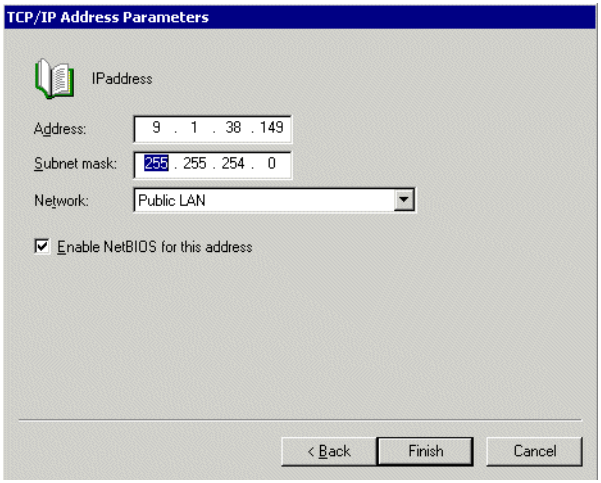

Figure 10-12 Setting the IP Address parameters

7. Right-click the resource and select **Bring Online** to make it available (Figure 10-13).

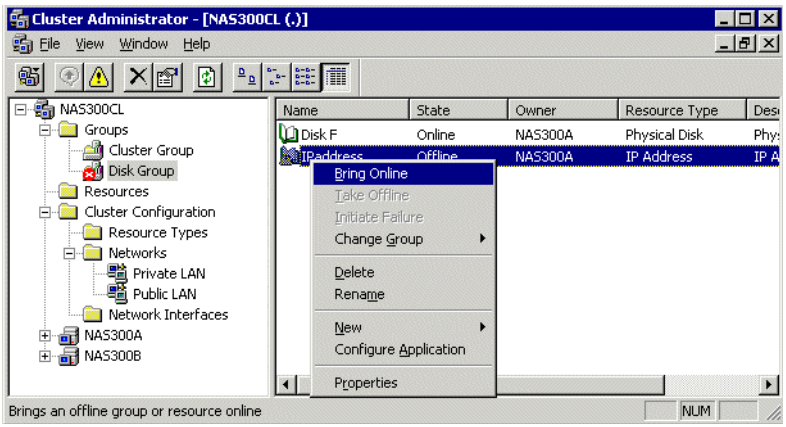

Figure 10-13 Bringing the resource online

Creating the network name resource

These are the steps for creating the network name resource:

1. From Cluster Administrator, select **File —>New—> Resource**.
2. Enter the network resource name, select **Network Name** as the resource type, select the disk group the resource should reside in as the Group (this should be the same disk group to which you added the IP Address resource), and then click **Next** (Figure 10-14).

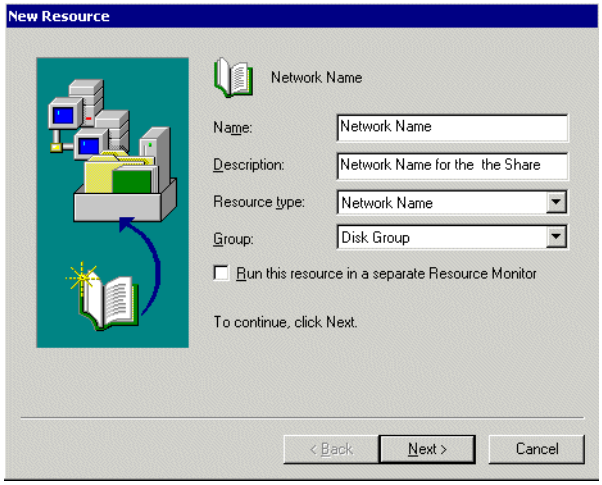

Figure 10-14 Creating Network Name resource

3. Add both nodes as possible owners using the **Add** button, then click **Next**.
4. In the Dependencies dialog box, select the **IP resource name** in the Available resources window and click the **Add** button (Figure 10-15).

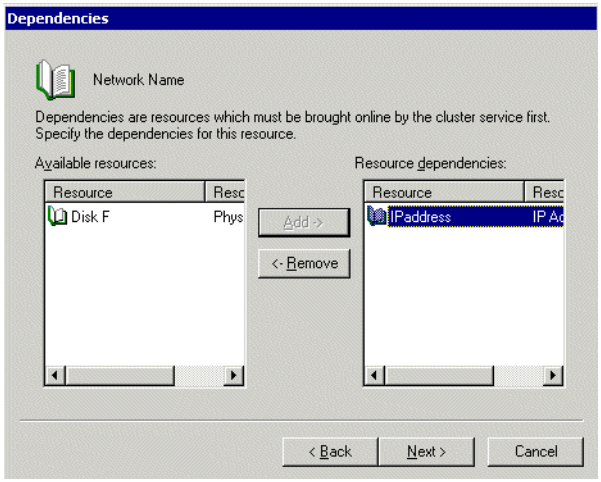

Figure 10-15 Setting Dependencies for the Network name

5. Enter the network name and click **Finish** as shown in Figure 10-16. (The network name is the name your network clients will use when connecting to the disk resource.)

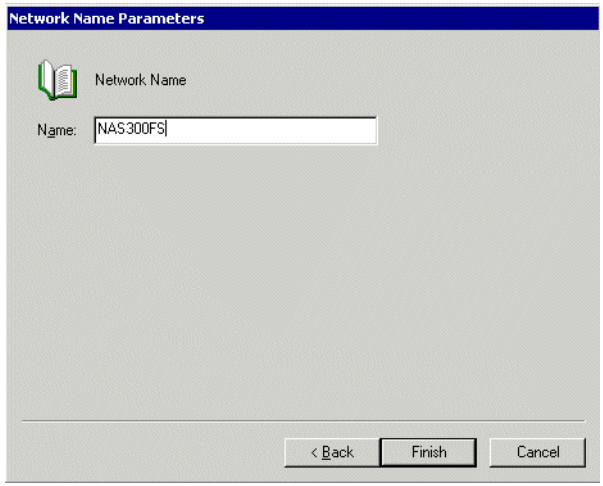

Figure 10-16 Setting the Network name parameters

6. Right-click the resource and select **Bring Online** to make it available.

Creating the file share resource (CIFS)

These are the steps for creating the file share resource:

1. From Cluster Administrator, select **File -> New -> Resource**.
2. Enter the file share resource name, select **File Share** as the resource type, select the disk group the resource should reside in as the Group (this should be the same disk group to which you added the IP Address and Network Name resources), and then click **Next** (Figure 10-17).

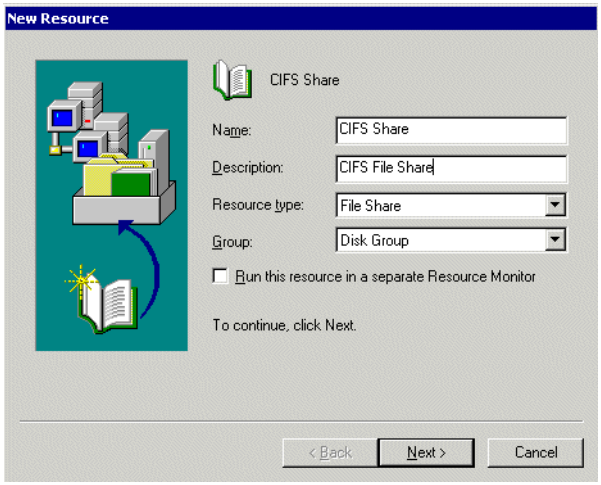

Figure 10-17 Creating a CIFS file share

3. Add both nodes as possible owners using the **Add** button, then click **Next**.
4. In the Dependencies dialog box, highlight the network name and physical disk and click the **Add** button to add these resources as dependencies and then click **Next** (Figure 10-18).

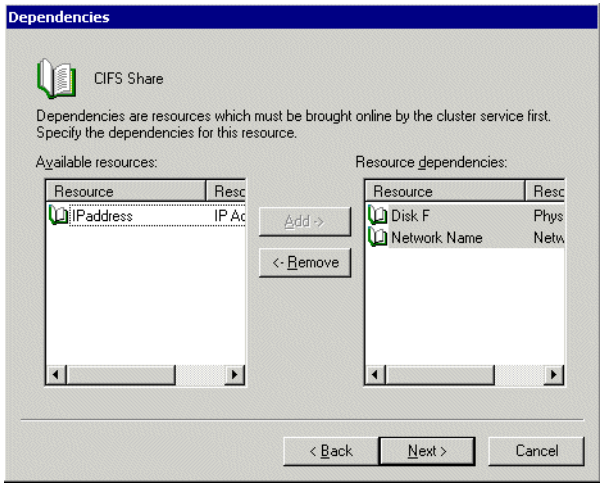

Figure 10-18 Setting the dependencies for CIFS file share

5. In the Share name field, enter the file share name (this is the actual name that clients will use to access the share).

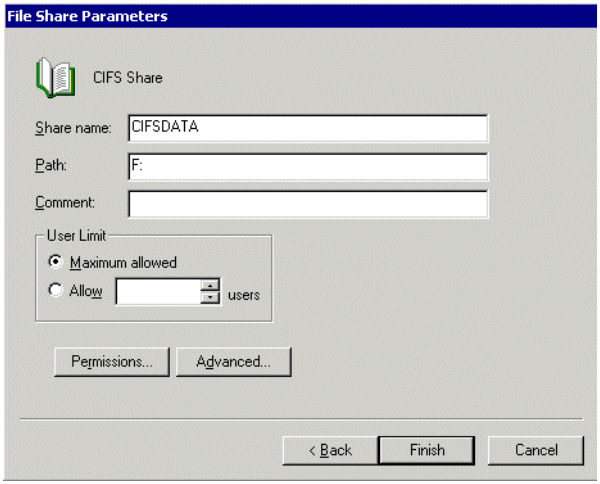

Figure 10-19 Setting the CIFS share parameters

6. In the Path field, enter the fully qualified path name of a directory on the disk that is being shared; this can just be the root directory if you want to share the entire disk under this file share (for example, “F:\”).
7. Click the Permissions button to set up access permissions. Change the permissions of the file share, then click **OK**.
8. From the File Share Parameters window, click **Advanced**.

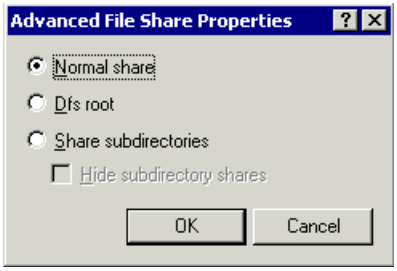

Figure 10-20 Selecting Normal share

9. Select the Normal share radio button, then click **OK** (Figure 10-20).
10. Click **Finish** from the File Share Parameters window.
11. .Right-click the resource and select Bring Online to make it available (Figure 10-21).

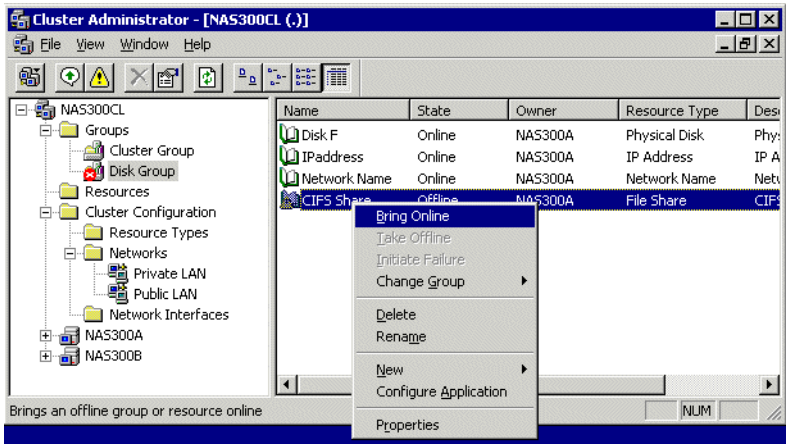

Figure 10-21 Bringing online the CIFS share

You have now configured your share in a clustered environment.

10.2 Accessing the shares from our Windows clients

From Windows, accessing the share was extremely straightforward. We just went into the Network Neighborhood (or My Network Places, as Windows 2000 prefers to call it), drilled down to the NAS 200, supplied a user name and password, right-clicked the shared directory, and chose **Map Network Drive** (Figure 10-22).

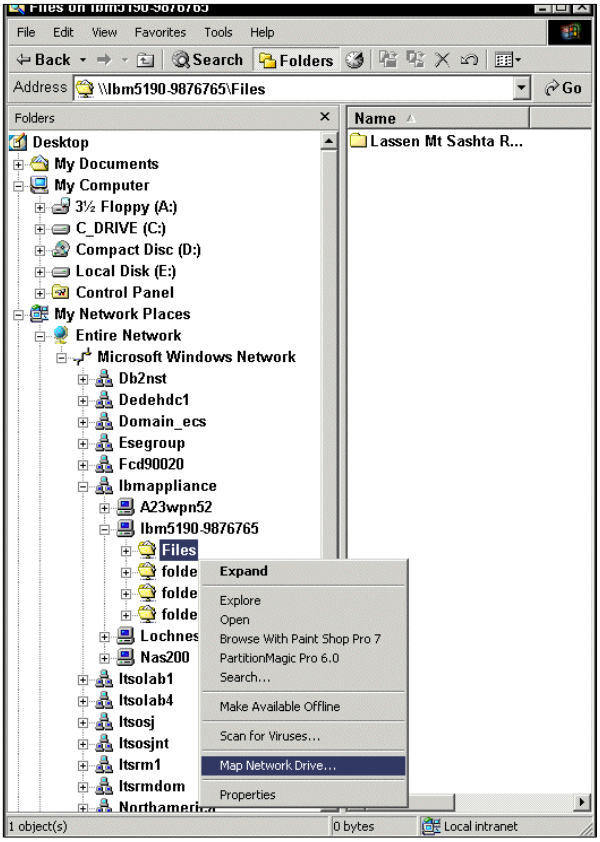

Figure 10-22 Map Network Drive

We were presented with a window requesting a drive, a folder, and were asked whether or not we wanted to reconnect at login (Figure 10-23).

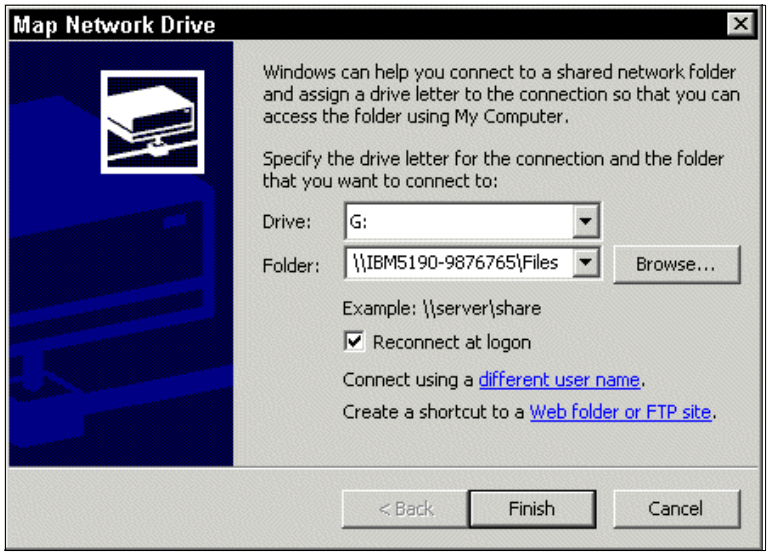

Figure 10-23 Windows mapping information

Once we supplied that information, we were able to see the shared disk, read from it, and write to it. You can check it using Windows explorer or executing the **net use** command from a DOS command prompt.

Another way of getting connected is by using the **net use** command. If you prefer the command line style, see Example 10-1.

Example 10-1 The net use command

```
C:\>net use /?
The syntax of this command is:
NET USE [devicename | *] [\\computername\sharename[\volume] [password | *]]
        [/USER:[domainname\]username]
        [/USER:[dotted domain name\]username]
        [/USER:[username@dotted domain name]
        [[/DELETE] | [/PERSISTENT:{YES | NO}]]
NET USE {devicename | *} [password | *] /HOME
NET USE [/PERSISTENT:{YES | NO}]

NET USE EXAMPLE: net use x: \\ibm5190-9876765\users password
/user:administrator
```

10.3 File sharing for UNIX clients

In this section we describe how to create shares for Windows clients in non-clustered and clustered environments.

Note: The Server for NFS is needed for NFS sharing and is already installed in the IBM TotalStorage NAS system.

10.3.1 NFS file sharing for non-clustered environments

Enabling access for UNIX systems requires just one more step. From the same dialog, we click the **NFS Sharing** tab and set it up as well. This is shown in Figure 10-24.

Important: This procedure will result in a breakdown of group/user properties which are critical to system administrators in the UNIX world. You will experience some group/user ID problems while following these steps. We will show a different approach in 10.4, “How to configure Services for UNIX (SFU)” on page 385 that might better fit your needs.

There are five tabs at the top of this property folder. There are two tabs that require changes for UNIX sharing: NFS Sharing and Security.

1. Click **NFS Sharing**. Then click the radio button for **Share this folder**.

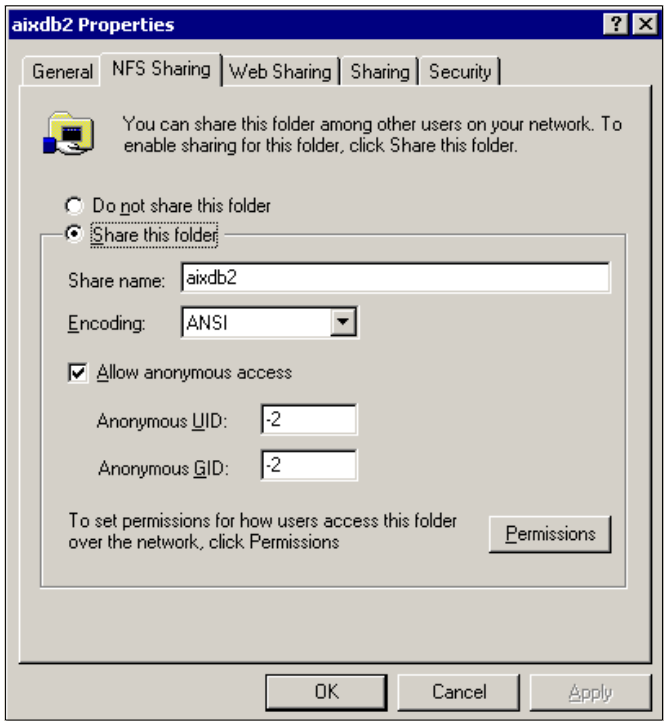

Figure 10-24 NFS sharing

2. Once again, we assigned a name to the share. We chose to use the same name as we used for the Windows clients. This conveniently allows the shared directory to be mapped/mounted in the same way from both UNIX and Windows clients. In Services for UNIX 2.3 you have some extra functions in addition to those you have in SFU 2.2. You can choose international encoding formats to get connections to NFS-clients with different character sets (Japanese, Korean, Taiwanese, etc.). See Figure 10-25.

Figure 10-25 SFU 2.3 international character sets

Note on anonymous access: See Figure 10-26. It is strongly recommended that you do not disable anonymous access. If a client presents a UID that is not recognized, Server for NFS can still grant that client a very limited form of access as a special *nobody* user. This is known as *anonymous* access, and you can enable or disable it on a per-share basis. This anonymous user will have very limited access to resources on the NAS: it has only the permissions that are granted to the *Everybody* group in Windows, which corresponds to the *other* (or *world*) bits in a POSIX permissions mode.

Allowing anonymous access is not a security risk, so disabling it might provide a false sense of security. (The real security risk is to grant everyone access to resources that should be protected.) And disabling anonymous access has one severe consequence: It is so unexpected by NFS clients that they may not be able to connect as NFS Version 3 clients at all, and may instead downgrade the connection to use the NFS V2 protocol.

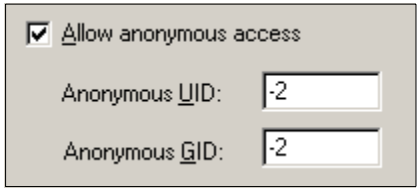

Figure 10-26 Anonymous access for NFS sharing

3. Since access permissions in Windows and UNIX are significantly different, however, we checked the **Permissions** dialog for NFS Sharing (Figure 10-27).

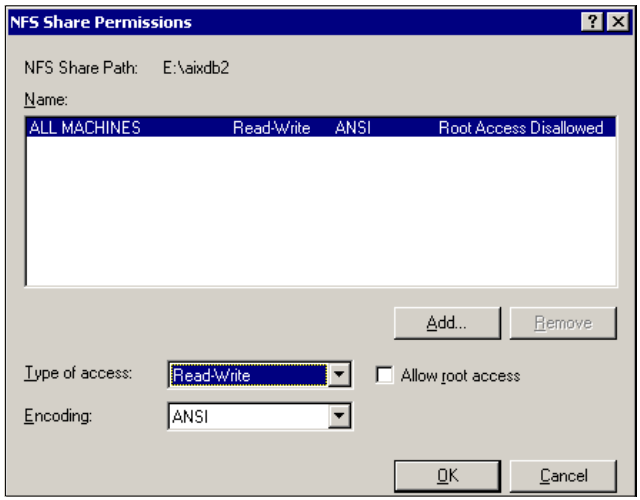

Figure 10-27 NFS share permission

4. Click **Add** to add the hosts you will allow to mount this share with Root permissions. The screen display is shown in Figure 10-28.

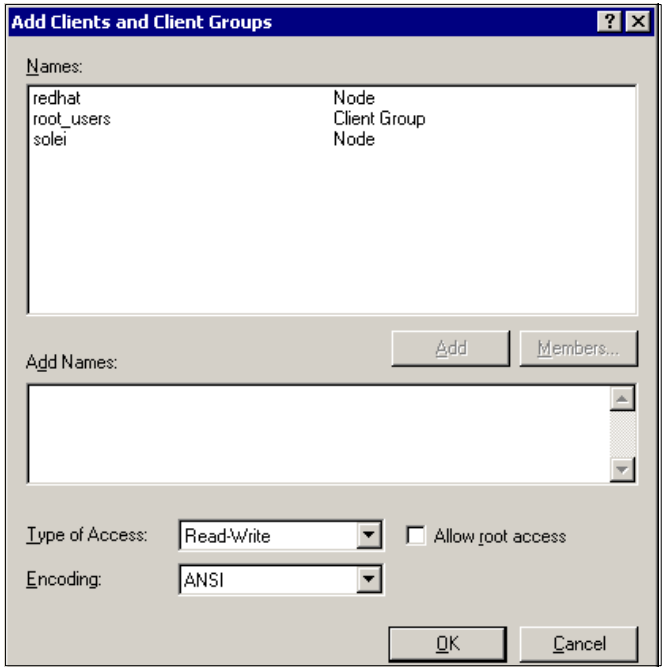

Figure 10-28 Add host to NFS share

5. In this case we added one Linux host as a client running RedHat 8.0. The hostname is *redhat*.
6. Select that hostname and click **Add**. This action will display the *redhat* client in a smaller rectangle box below (Figure 10-29). Then click the down arrow, and select **Read-Write**. Check the **Allow root access** box to give this host root access to the share.

   Most UNIX systems can mount a file system as *root*. The *root* use has the highest administrative rights on the UNIX system and is equivalent to the administrator account in Windows NT/2000 systems.

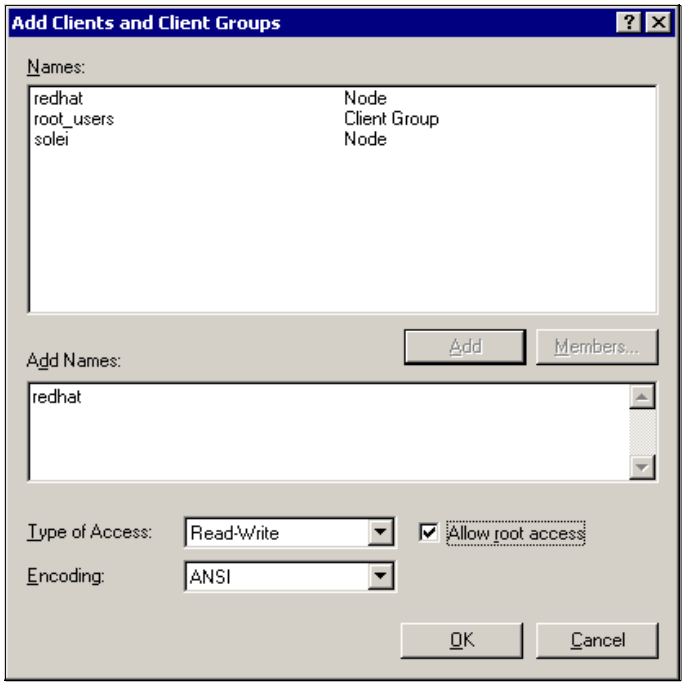

Figure 10-29 Set Root access for UNIX host

7. Then click **OK**, which brings you back to previous screen where you click **OK** again. Then, on the main NFS share window, click **Apply**.

Now you have completed the folder sharing configuration for UNIX.

To check your current NFS shares from a DOS prompt, you can use the **`nfsshare`** command to display all your current shares. This tool comes with Microsoft Services for UNIX.

Tip: To make your administration easier, we recommend using the same share name for Windows and UNIX shares.

10.3.2 NFS file sharing for clustered environments

In this section we describe the steps necessary to create a NFS share on a clustered system as the IBM TotalStorage NAS 300.

You must define file shares for Windows users in Cluster Administrator, to allow the file shares to fail over between nodes if a node goes down. Any file shares defined via Windows Explorer will not fail over, as they are only valid for the node on which the share was created.

The creation of clustered file shares involves dependencies on other resources in the group as a physical disk, a static IP address, and a network name. The various dependencies allow resources that are defined to the same disk group to start in the correct sequence. The dependencies also assure necessary access for the given resource. The steps on this page guide you through creating the resources and dependencies.

Note: To follow this procedure you can be direct attached to the NAS device or to manage it through a Terminal Services connection.

Creating the IP address resource

These are the steps for creating the IP address resource:

1. Open the Cluster Administrator (Figure 10-30), in the IBM NAS Admin tool, select **Cluster Tools —> Cluster Administration**.

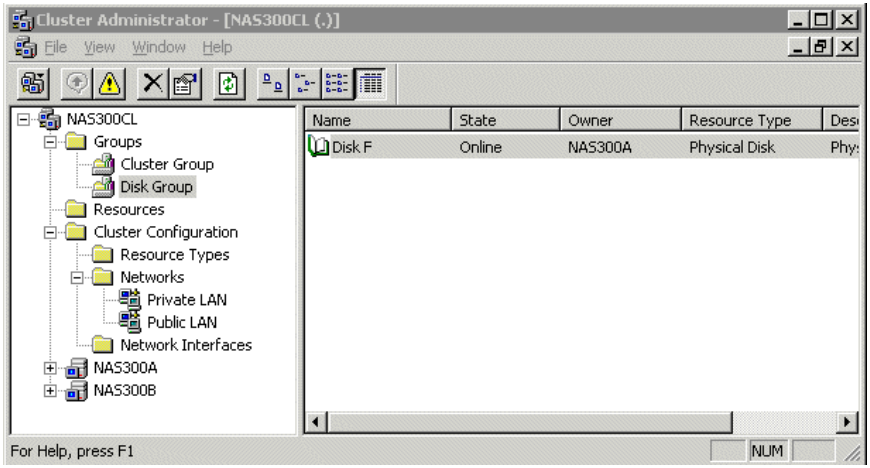

Figure 10-30 Cluster Administrator

2. From Cluster Administrator, select **File —>New—> Resource** (Figure 10-31).

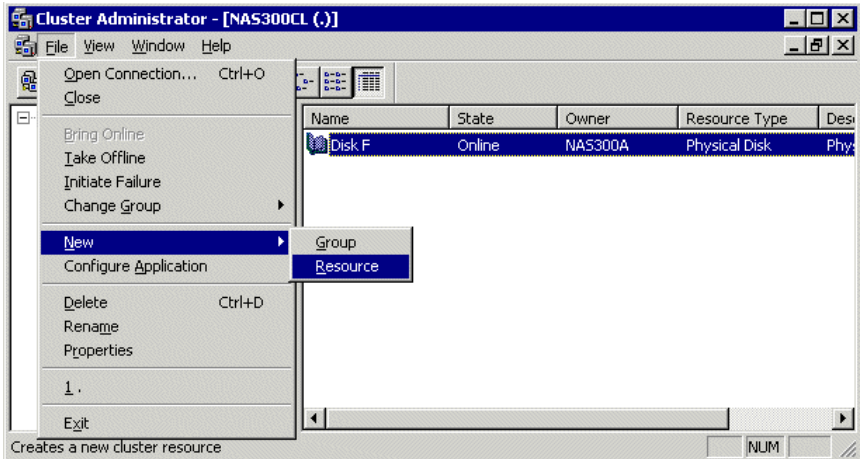

Figure 10-31 Creating New Resource

3. Enter the IP Address resource name, select IP Address as the resource type, select the disk group the resource should reside in as the Group, then click **Next** (Figure 10-32).

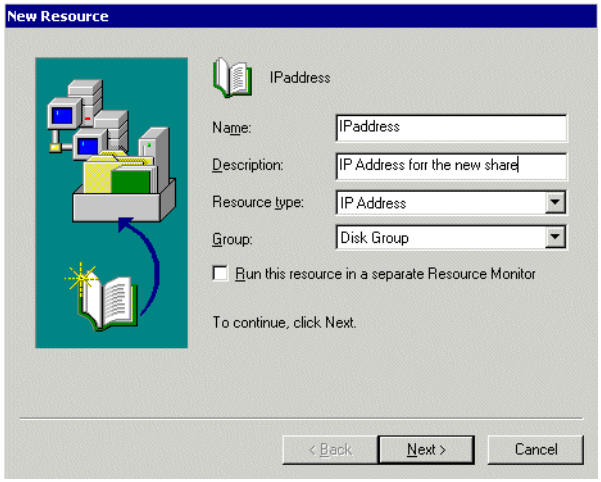

Figure 10-32 Defining an IP Address Resource

4. Add both nodes as possible owners using the **Add** button, then click **Next** (Figure 10-33).

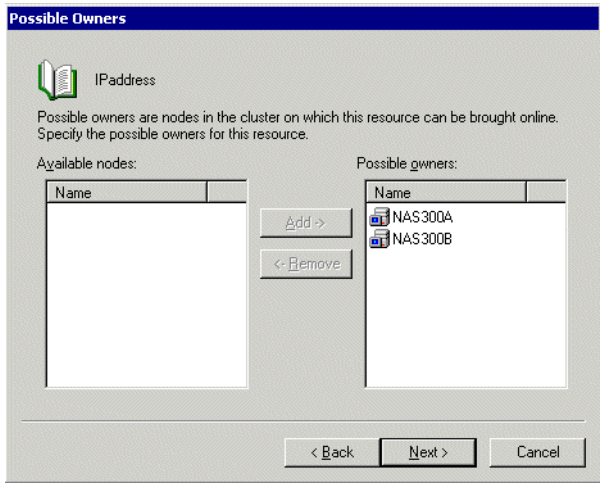

Figure 10-33 Choosing possible owners

5. Leave the Resource Dependencies window empty and click **Next** (Figure 10-34).

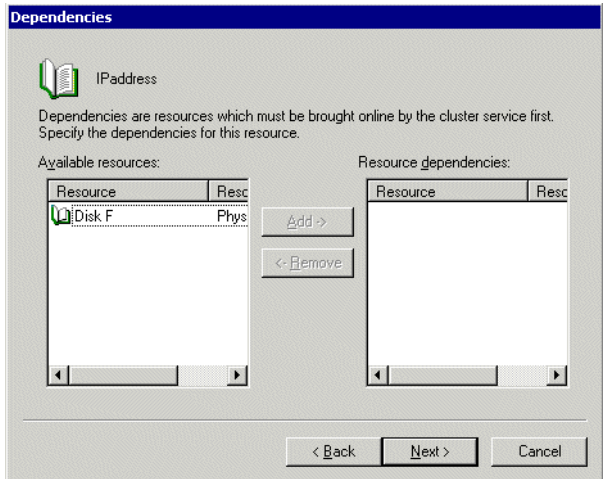

Figure 10-34 Setting resource dependencies

6. Enter your IP address and subnet mask and click **Finish** (Figure 10-35).

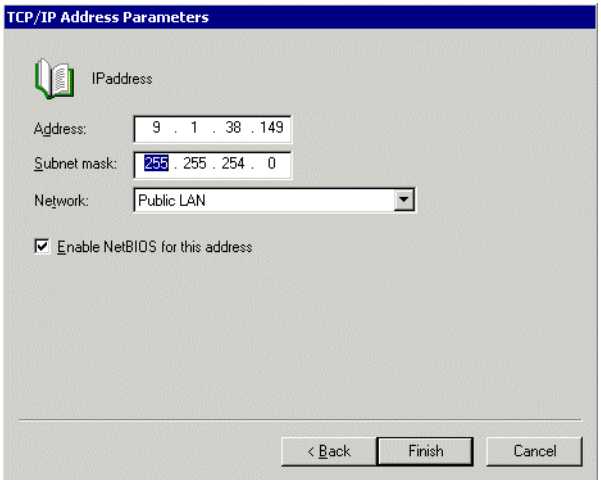

Figure 10-35 Setting the IP Address parameters

7. Right-click the resource and select **Bring Online** to make it available as shown in Figure 10-36.

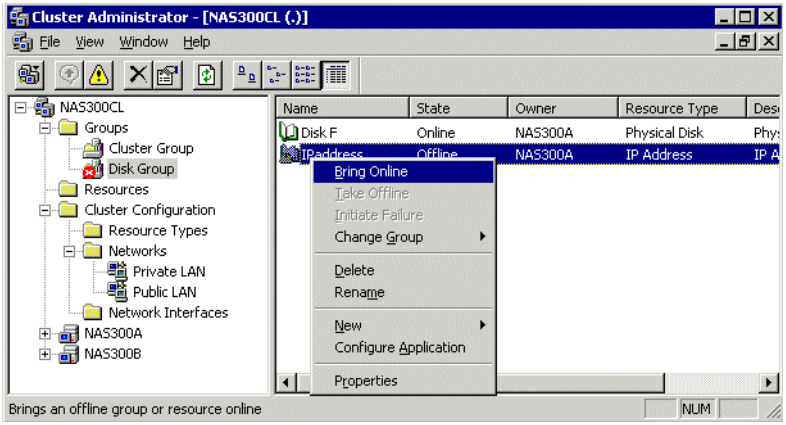

Figure 10-36 Bringing the resource online

Creating the network name resource

These are the steps for creating the network name resource:

1. From Cluster Administrator, select **File —>New—> Resource**.
2. Enter the network resource name, select Network Name as the resource type, select the disk group the resource should reside in as the Group (this should be the same disk group to which you added the IP Address resource), and then click **Next** (Figure 10-37).

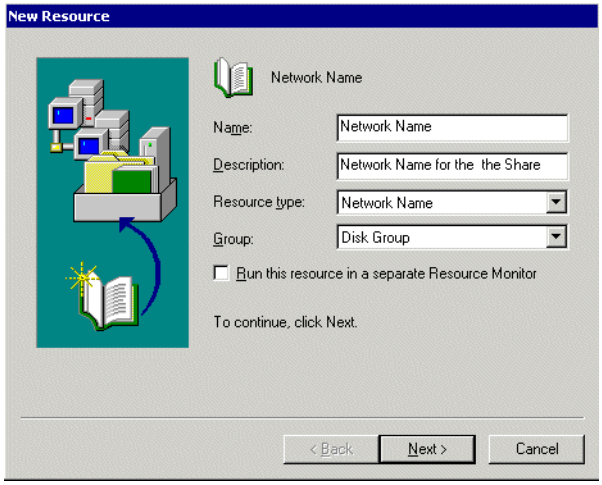

Figure 10-37 Defining Network Name Resource

3. Add both nodes as possible owners using the **Add** button, then click **Next**.
4. In the Dependencies dialog box, select the **IP resource name** in the Available resources window and click the **Add** button (Figure 10-38).

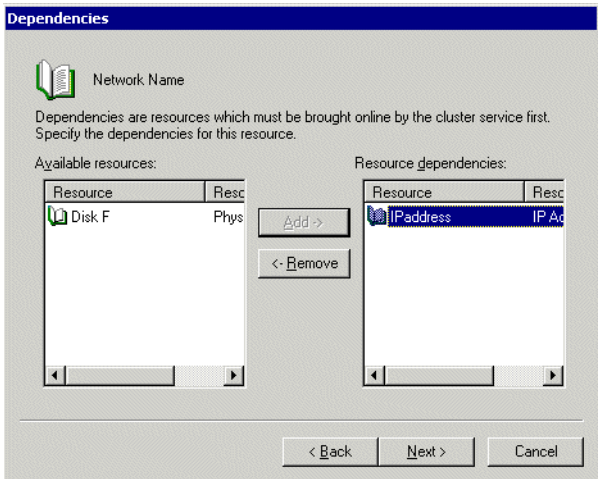

Figure 10-38 Setting Dependencies for the Network name

5. Enter the network name and click **Finish** as shown in Figure 10-39. (The network name is the name your network clients will use when connecting to the disk resource.)

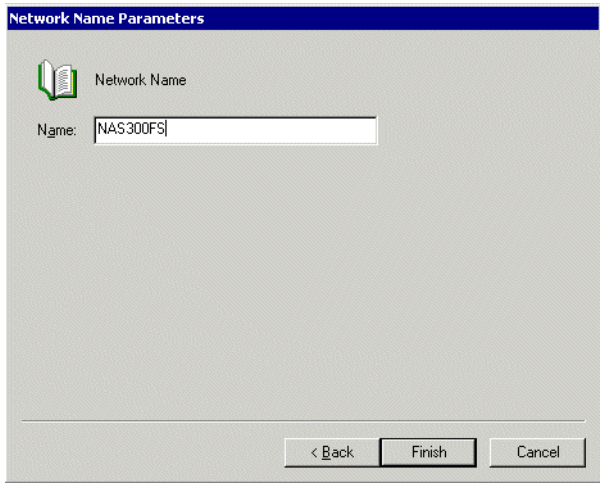

Figure 10-39 Setting the Network name parameters

Right-click the resource and select **Bring Online** to make it available.

Creating the NFS file share resource (CIFS)

These are the steps for creating the file share resource:

1. From the cluster Administrator, select **File —> New —> Resource**.

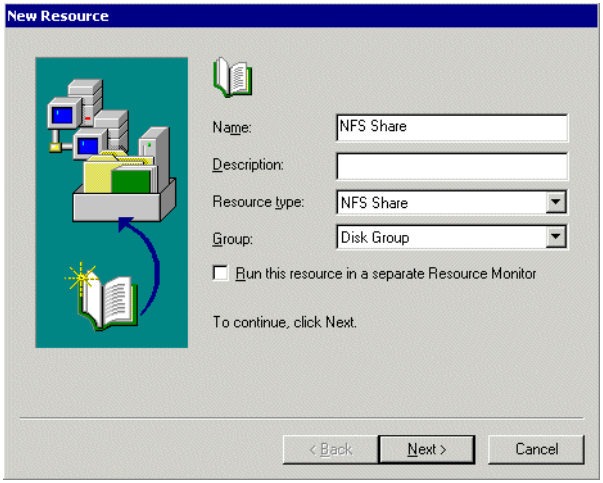

Figure 10-40 Creating the NFS share

2. Enter the file share resource name and select **NFS Share** as a resource type. Select the disk group (this should be the same group to which you created the IP Address and Network Name resource), and then click **Next** (Figure 10-40).
3. Add both nodes are possible owners using the **Add** button. Click **Next** (Figure 10-41).

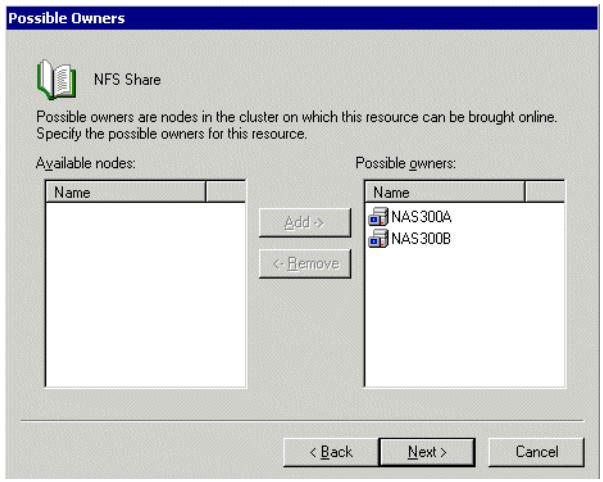

Figure 10-41 Choosing possible owners

4. Add the resource dependencies for the physical disk and network name that the file share will use and click **Next** (Figure 10-42).

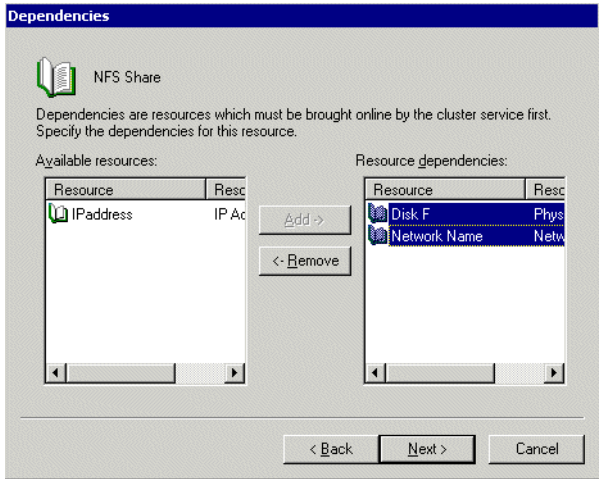

Figure 10-42 Setting Dependencies

5. In the File share field, enter the share name and the path to the disk in this group, either drive or sub-directory as shown in Figure 10-43. You can then also set permissions.
6. Click **Finish** to create the resource.

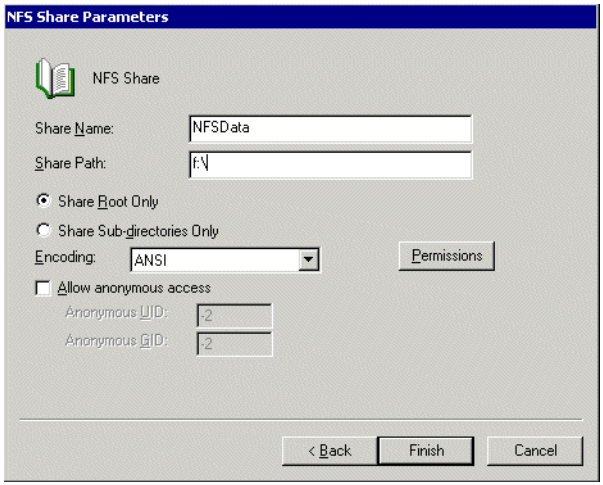

Figure 10-43 Setting the NFS file share parameters

7. Right-click the resource and select **Bring online** (Figure 10-44).

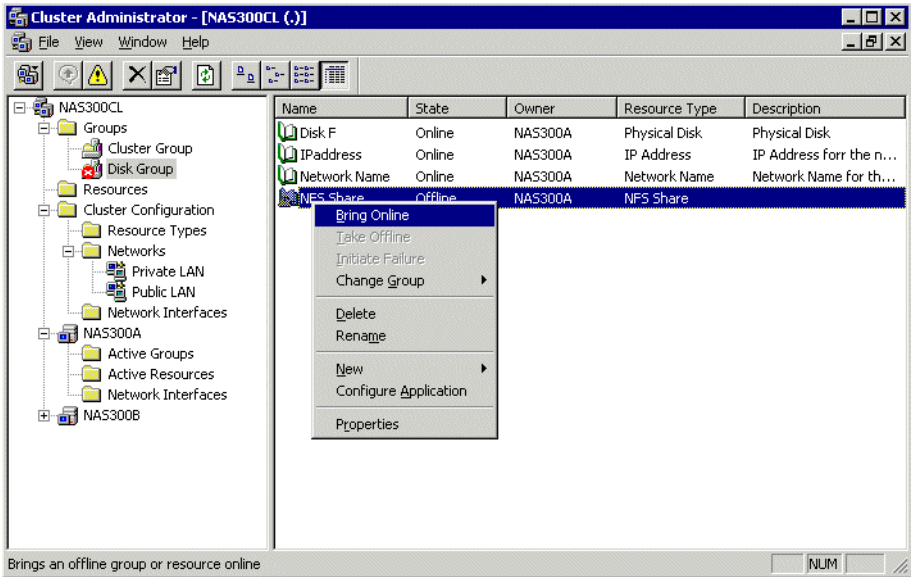

Figure 10-44 Bringing online the NFS share resource

10.4 How to configure Services for UNIX (SFU)

In a mixed environment, where UNIX and Windows clients will be accessing the storage of the NAS appliance, special care must be taken to ensure a consistent view of the data for all users who are granted access rights to it. If the storage is configured correctly, as described in this chapter, access to the data should be transparent to the user, whether they are accessing it from a Windows or UNIX workstation.

The Network File System (NFS), described in 10.3, “File sharing for UNIX clients” on page 372, is the common file system for access in an environment where there are both UNIX and Windows clients. NFS is the “binding” agent that makes inter operability between these two, very diverse, operating systems possible. All accesses in this environment are through the NFS Server, which is running in the NAS appliance.

Also, 10.3, “File sharing for UNIX clients” on page 372, describes the steps that are taken to export an NFS share from the NAS appliance. Then, 10.4.5, “Accessing the shares from our UNIX clients” on page 421 describes how to mount this NFS share from the UNIX client using the command line *mount* utility.

In this section we explain all of the steps that are required to access the NFS exported share from the Windows client. This includes, most importantly, the introduction of a new component — the *Gateway for NFS* share.

The NFS share is not native to the Windows operating system and, therefore, cannot be accessed directly from a Windows client on the network. Special NFS client software could be added to the Windows workstation allowing its user to access the share in much the same way as the UNIX client. However, there is a simpler way to accomplish this using another feature of the NAS appliance.

What is an NFS Gateway?

An existing NFS share can be used to create a new Windows drive which can then be mapped directly from a Windows workstation on the network. This new drive, called a *Gateway for NFS* share, is accessible from the network like any other Windows share. The significant difference is that file access requests for this share are passed through to the Services for UNIX components, NFS Server and User Name Mapping. In this way, a consistent view of the data is preserved for all users, and access to the data is transparent to the user whether they are accessing it from a Windows or UNIX workstation.

A sample access to the Gateway for NFS is shown in Figure 10-45.

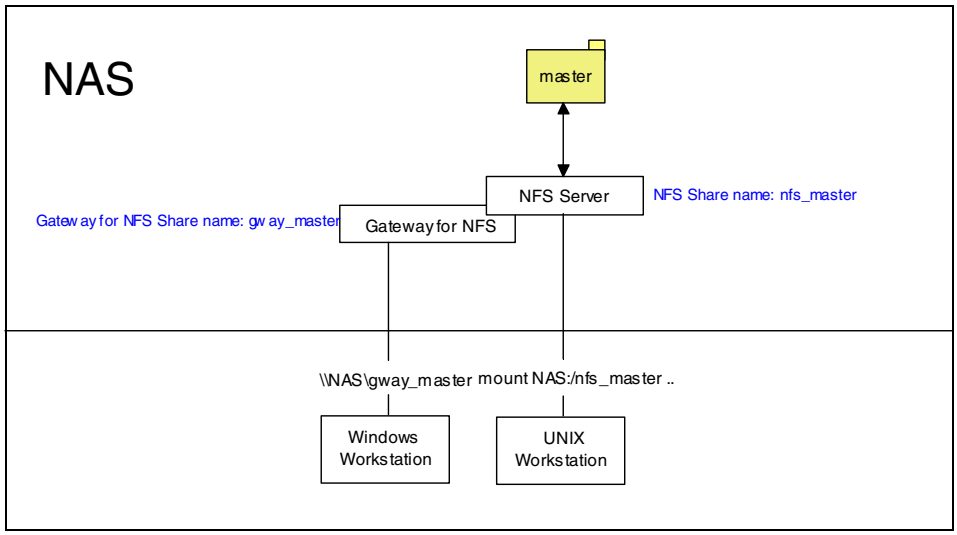

Figure 10-45 Gateway for NFS access

We strongly recommend the use of NAS Code Release 2.5 (SFU 2073.1) for heterogeneous file sharing.

10.4.1 Configuring a cross platform share in a Windows 2000 Domain

In this section we give a brief introduction on how to set up Services For UNIX (SFU) for heterogeneous file sharing between UNIX and Windows. Our intention is to cover the basic installation in the most common customer environments with a Windows 2000 Domain.

Configuring the Windows 2000 Domain controller

Before you can use SFU in your network environment, it is required to install the SFU “Server for NFS Authentication” on your Domain controller.

1. Insert the Supplemental CD 2/2 shipped with the NAS100 box and select the folder SFU_xxxx (Figure 10-46).

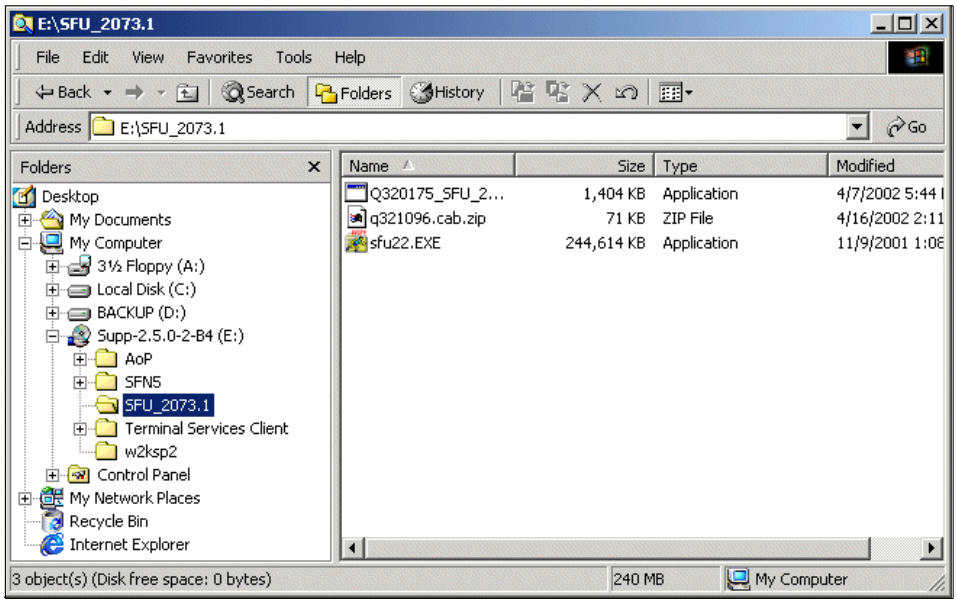

Figure 10-46 SFU 2073.1 on the Supplementary cd 2/2

2. Unzip the SFU22.exe by double (Figure 10-47).

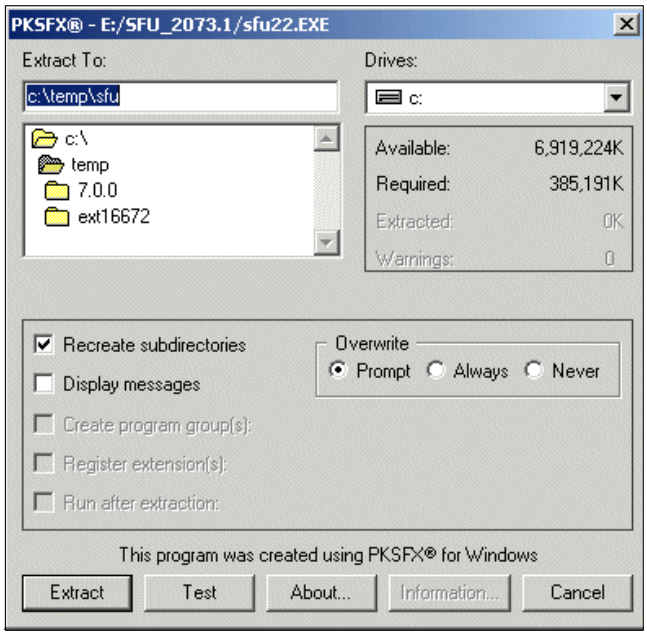

Figure 10-47 Unzip screen for SFU22.exe

3. Choose a different temporary location or click **Extract**. Confirm to create the temporary directory with **Yes** (Figure 10-48).

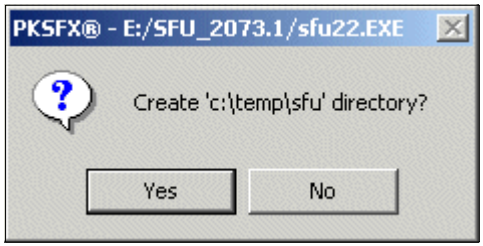

Figure 10-48 Create temporary directory

4. The program will start to inflate the zipped files (Figure 10-49).

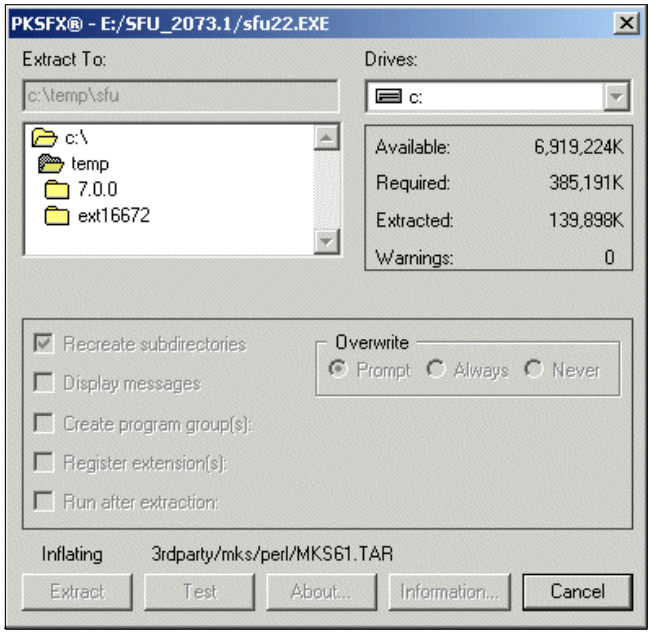

Figure 10-49 Inflating the installation files

5. Switch to the */temp/sfu* directory and start the installation by double clicking the **OEMsetup.msi** file (Figure 10-50).

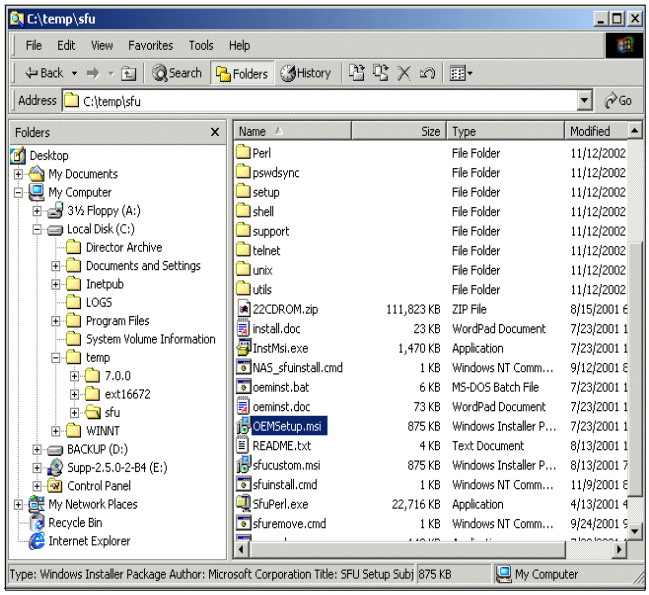

Figure 10-50 Start installation with OEMSetup.msi

6. The welcome screen will appear. Continue the process by clicking **Next** (Figure 10-51).

Figure 10-51 SFU 2.2 welcome screen

7. Accept End-user license agreement (Figure 10-52).

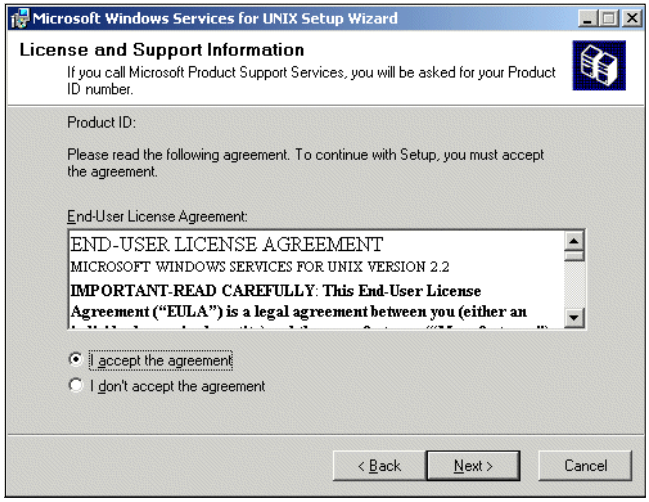

Figure 10-52 End-User license agreement

8. After inserting the customer's information (name and company name), the Installation Options screen will appear (Figure 10-53). Select **Customized installation** and click **Next**.

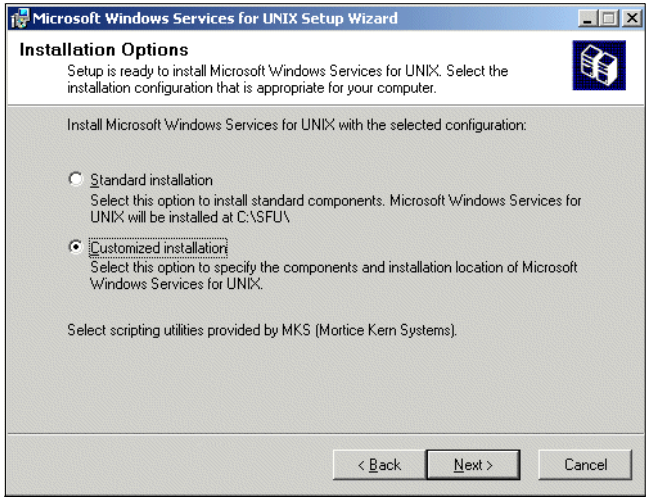

Figure 10-53 Installation Options screen

9. Unselect everything except “Server for NFS Authentication” (Figure 10-54).

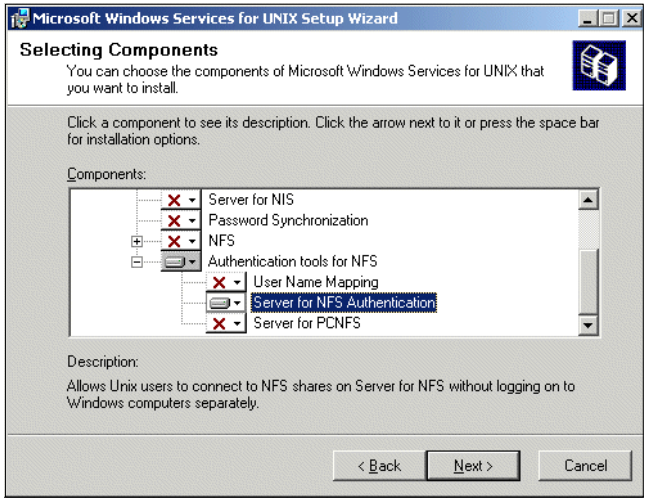

Figure 10-54 Select “Server for NFS Authentication”

10. Choose a different location for the program files or click **Next** (Figure 10-55).

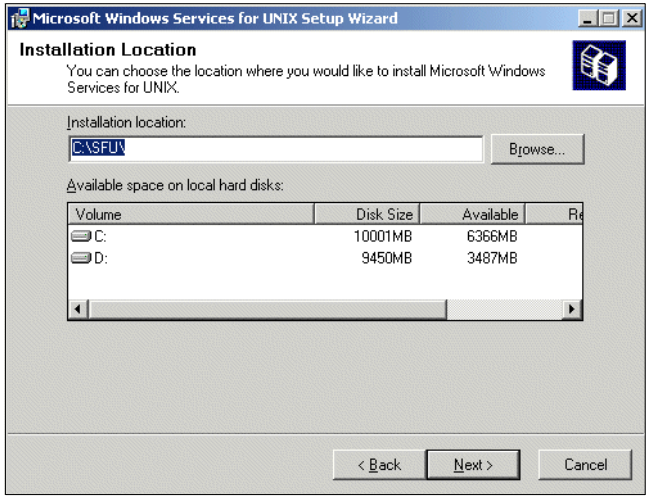

Figure 10-55 Location for SFU

11. Now the installation program will install the Service on the Domain controller (Figure 10-56).

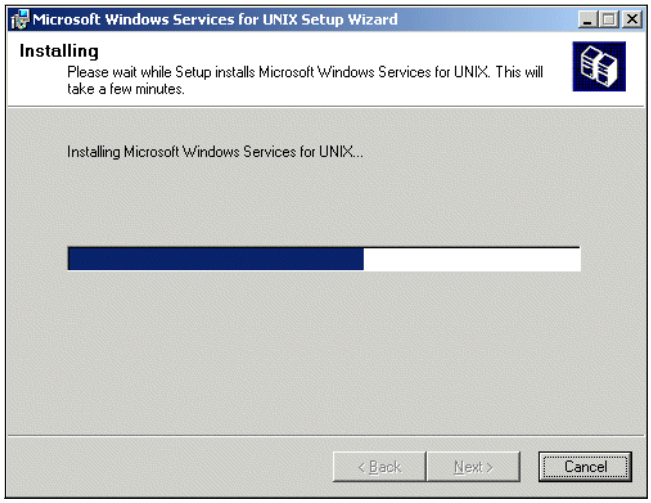

Figure 10-56 Installation of SFU

12. Click **Finish** after the installation process (Figure 10-57).

Figure 10-57 Completing Setup screen

13. Create additional Windows and Groups for heterogeneous file sharing in the Domain controllers Active Directory (only if needed).

Note: The primary group must be set to a group that is mapped in SFU for every user that will be used for heterogeneous file sharing, otherwise “nobody” (-> `ls -al` command) will be displayed.

14. Execute **Start->Programs->Administrative Tools->Active Director Users and Computers**.
15. Select the User, select the tab **Member of** and set the **primary group** to one that will be mapped in SFU (Figure 10-58).

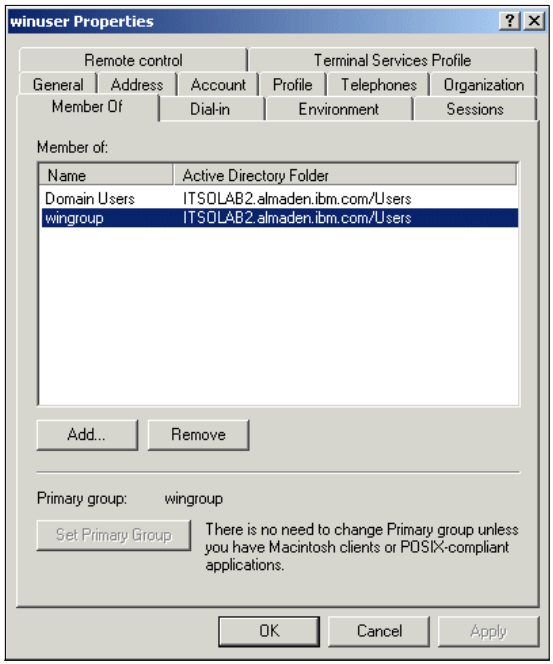

Figure 10-58 Primary group from winuser set to wingroup for SFU

Configuring the AIX Client

We assume your AIX Client is up and running.

Create Users and Groups for heterogeneous file sharing if needed.

Configuring the Windows Client

We assume your Windows Client is up and running.

Connect the Windows Client to the Windows Domain Controller (logon to the Domain with a valid Domain User Account).

Configuring the NAS 100

The next step is to set up Microsoft Services For UNIX (SFU) on the NAS 100. This is an additional software that runs on the Windows Powered OS on the NAS server. There are various steps of configuration that need to be done. Those steps are documented step-by-step in this section.

Before you configure the users and groups on Windows 2000 server and Services for UNIX, you need to identify all users on the UNIX side which will be accessing this shared resource. The root user must be defined, so the root user can mount a shared NAS resource as a file system on UNIX server. The user and group information can be found in two systems files. These are */etc/passwd* and */etc/group* on most UNIX systems, unless a Network Information Services (NIS) service has been set up. We did not use a NIS setup, so we used the */etc/passwd* and */etc/group* files for the information.

We assume that you are an experienced Windows administrator, so we are covering the process of how to create users and groups on a Windows 2000 server. We will need this information for mapping as a part of the configuration.

Apply Fixes if needed

The NAS 100 box has SFU 2.2 installed with the HotFixes QFE 320174 and QFE 321096. Any NAS box that has installed the SFU QFE Q320175 installed may encounter a problem after 8/1/02 due to the inclusion of a beta expiration check into that QFE. This check will render the component unusable.

http://www-1.ibm.com/support/docview.wss?uid+ssg1S1001387

Get the Password and Group files from UNIX

FTP /etc/passwd and /etc/group from the AIX Client to the NAS 100 (NAS 100 path: c:\winnt\system32\drivers\etc). FTP the passwd and group files in binary mode from the AIX workstation to the NAS 100 (Figure 10-59):

1. Open the command prompt.
2. Change to the \winnt\system32\drivers\etc\ folder on your c: drive.
3. To transfer the files, type **`ftp xxx.xxx.xxx.xxx`** and press Enter.
4. Login with a valid UserID and password.
5. Type **`bin`** for binary mode and press Enter.
6. Type **`get /etc/passwd`** to transfer the first file.
7. Type **`get /etc/group`** to transfer the second file.
8. Close the FTP session by entering **`bye`** and pressing Enter.

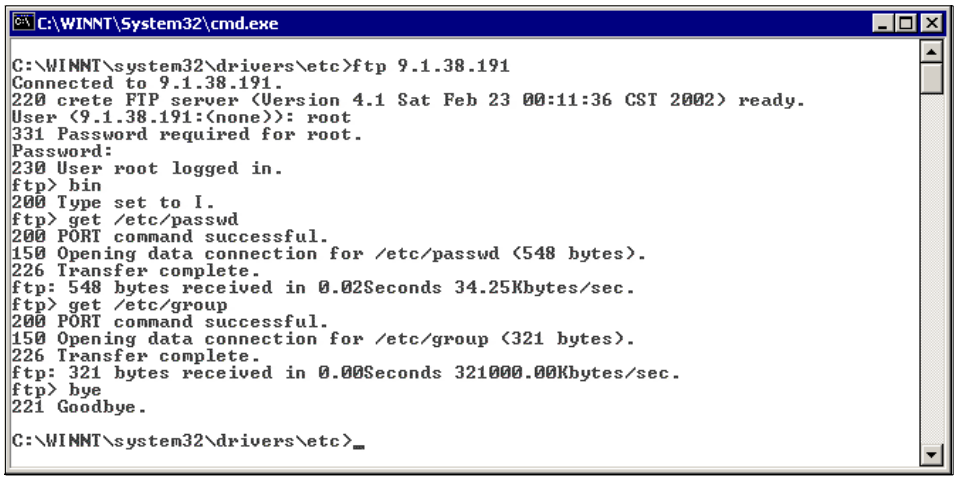

Figure 10-59 FTP /etc/passwd and /etc/group to the NAS 100

The /etc/group file contains information that is not needed by the PCNFS Server. The Server would not accept the group information from the file. For that reason, it is important to modify the group file for further use.

9. Open the Windows Explorer and click to move to the directory /winnt/system32/drivers/etc (Figure 10-60).

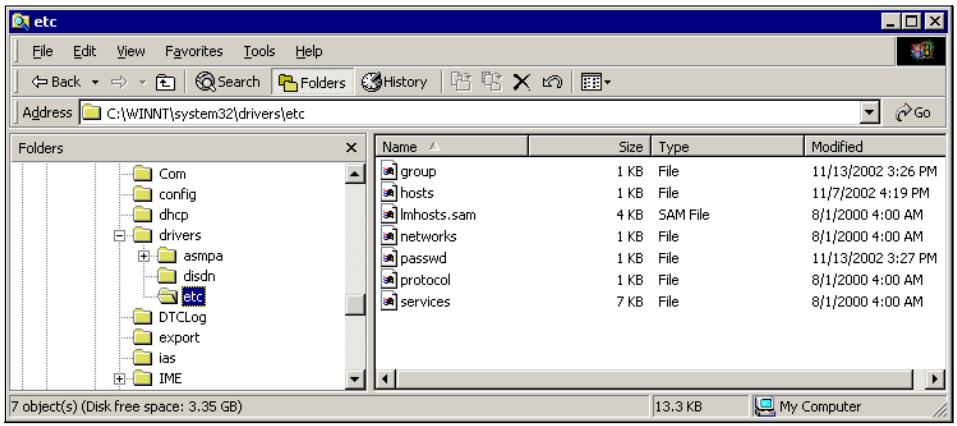

Figure 10-60 Directory for /etc/group and /etc/passwd

10. Right-click the *group* file (Figure 10-61).

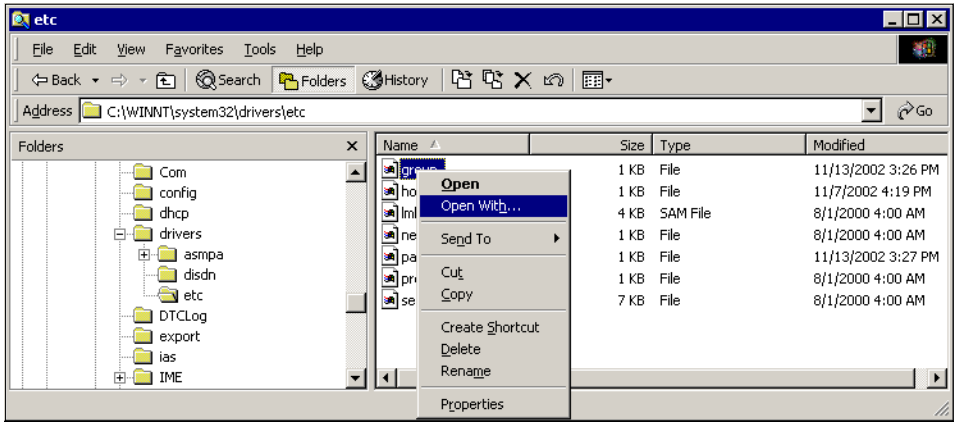

Figure 10-61 Open /group with...

11. Choose **Notepad** to open the file (Figure 10-62).

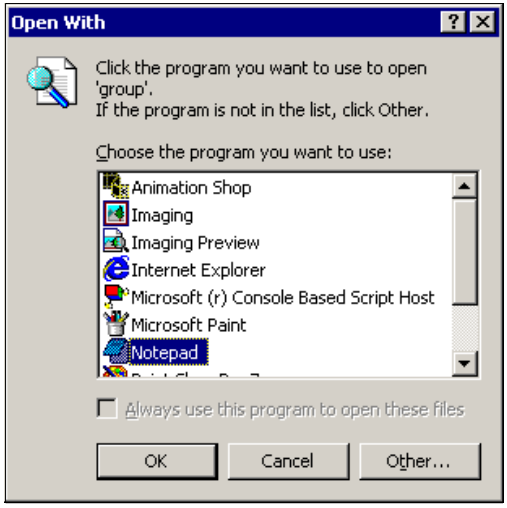

Figure 10-62 Choose Notepad

12. Click the **Edit** tab and choose **Replace...** (or Ctrl-H) (Figure 10-63).

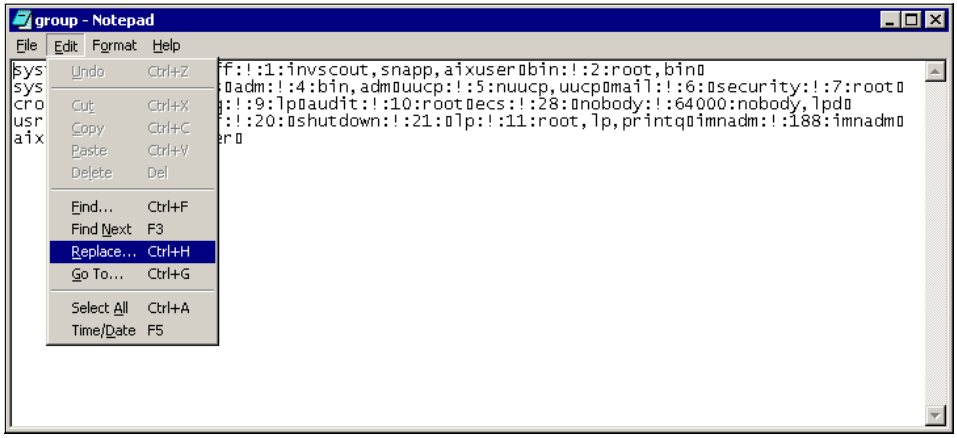

Figure 10-63 Edit Replace tab

13. Enter an exclamation mark (**!**) in **Find what** and click the **Replace All** button (Figure 10-64).

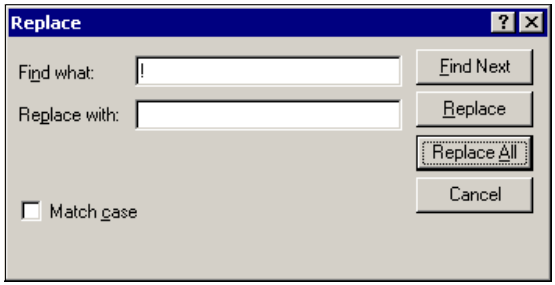

Figure 10-64 Replace exclamation mark with nothing

14. Notepad will erase all exclamation marks. Save and close the file after the procedure. The file can now be used with the PCNFS Server.

Configuring SFU on the NAS 100

This is the procedure to follow:

1. Click the IBM NAS **admin.msc** icon on your desktop.
2. Then click the **+** next to the file system and select **Services for UNIX**. You will see the Services for UNIX screen, including all details and release levels. Here, we have *Services for UNIX V2.2* in our environment (Figure 10-65).

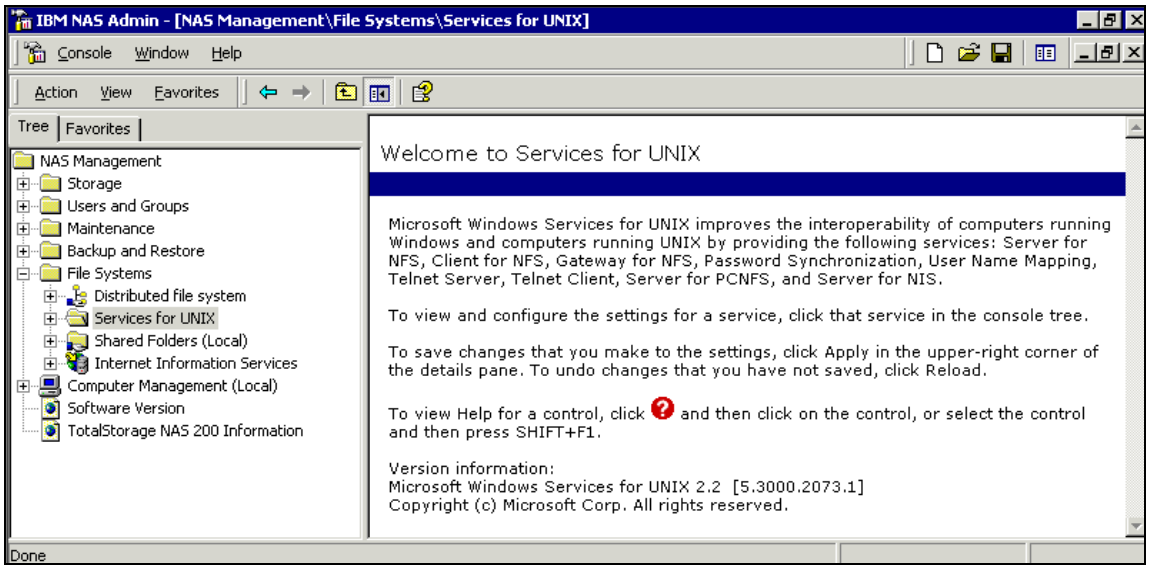

Figure 10-65 Services for UNIX main screen.

3. Click **Server for NFS**.
4. Click the **User Mapping** tab and you will see a display like the one in Figure 10-66.

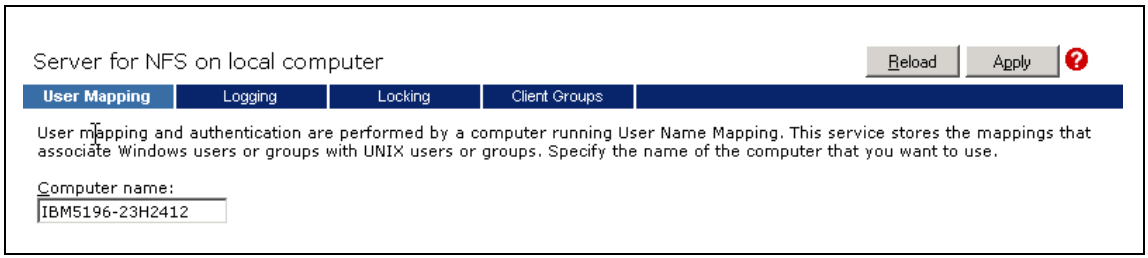

Figure 10-66 Server for NFS user mapping

Specify the computer name of the server that is running the *User Name Mapping* Service. If *User Name Mapping* will be configured to run on this NAS appliance, specify localhost or the machine name.

Note: Make sure to apply every change by clicking the **Apply** button in the upper right corner of the window.

5. Select **User Name Mapping** from the left navigation bar (Figure 10-67).

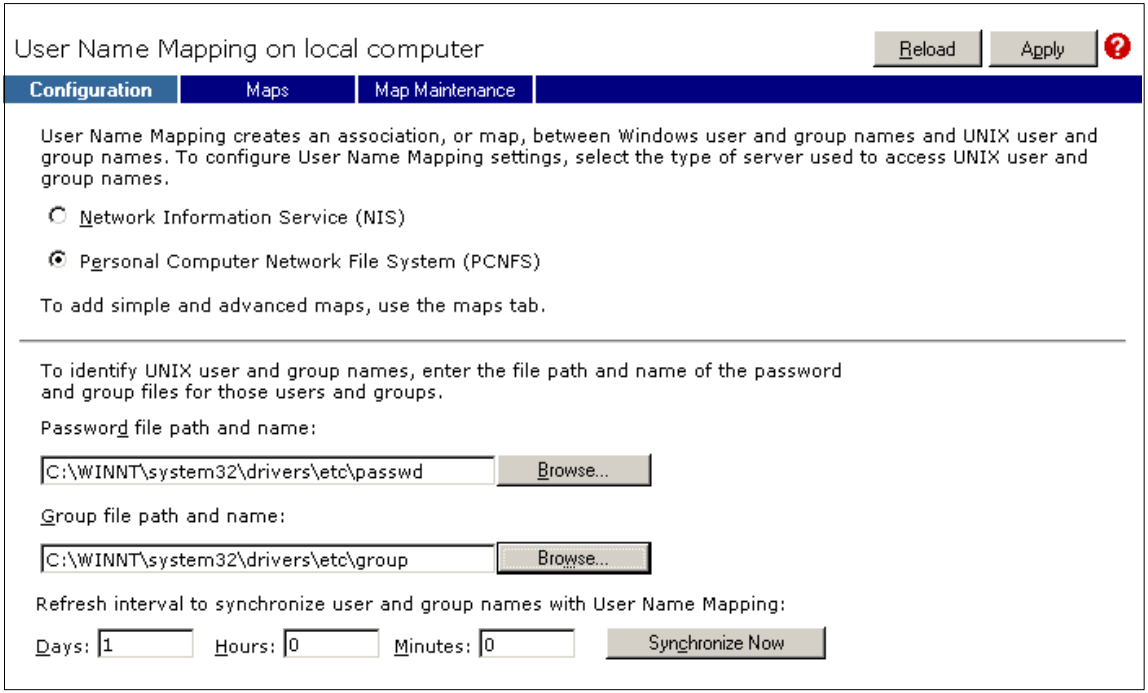

Figure 10-67 User Name Mapping

6. Select Personal Computer Network File System (PCNFS) since we are not using a NIS Server in our environment.
7. Point the Password and Group file path to the “passwd” and “group” file FTP location on the NAS 100, as it is shown.
8. Select **Maps** in the right window.

9. Select Simple Maps in order to map users with the same user name in the Windows and UNIX environment. Make sure you have selected your Window Domain (Figure 10-68).

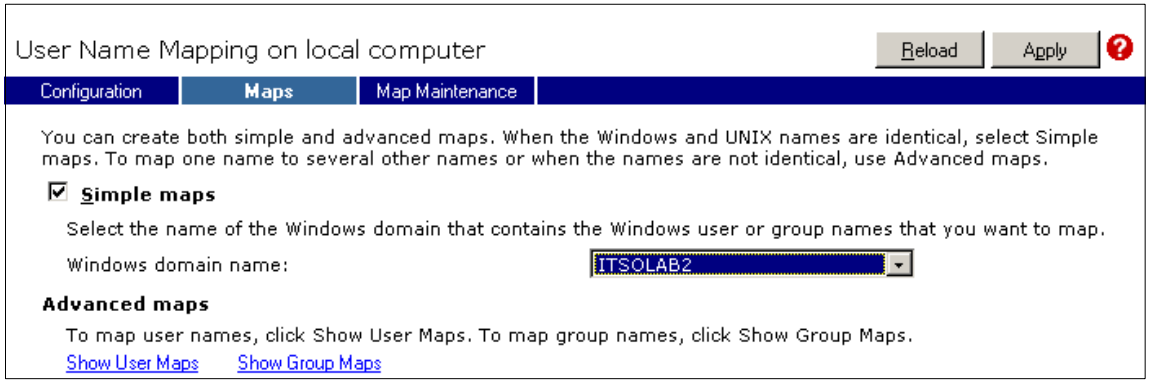

Figure 10-68 User Name Mapping / Maps

10. Click the **Show User Maps** tab.
11. Select the **List Windows Users** and **List UNIX Users buttons**.
12. Select the Windows user **root** from the Windows users list and the UNIX user **root** from the UNIX users list.
13. Click the **Add** button to add a mapping for these two users.
14. You can create mappings between other UNIX and Windows users at this time. This is shown in Figure 10-69.

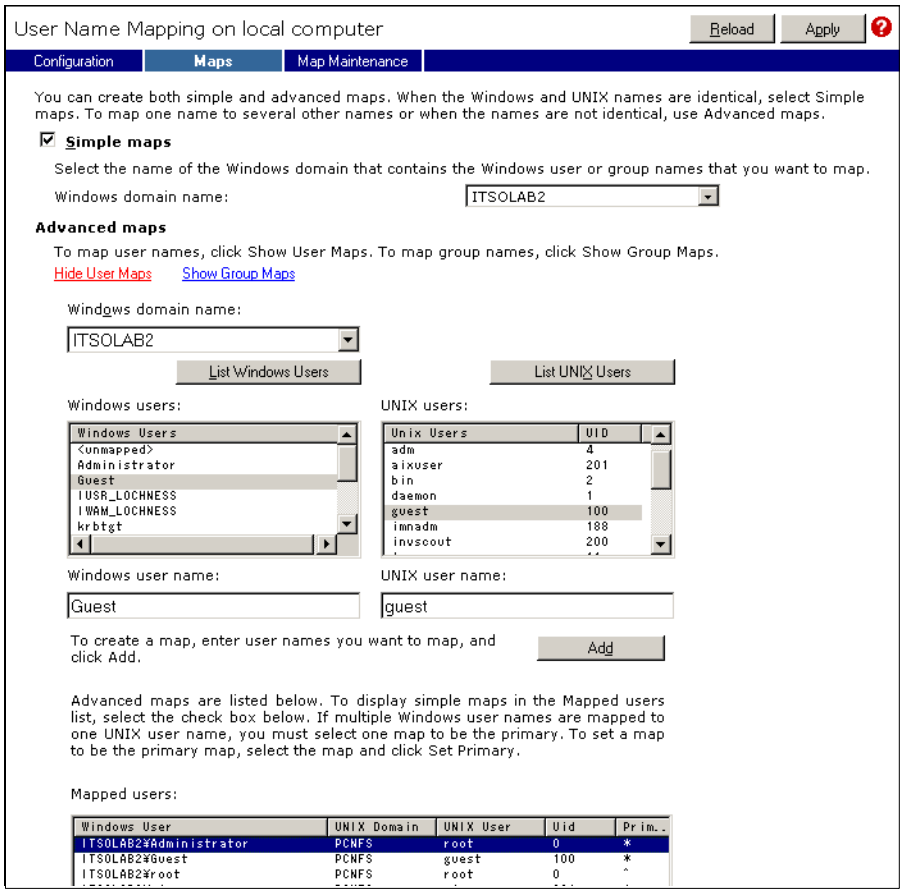

Figure 10-69 User name mapping on NAS 100

15. Click the **Show Group Maps** link.
16. Select the **List Windows Groups** and **List UNIX Groups** buttons.
17. Select the Windows group **Domain Admins** from the **Windows groups list** and the UNIX group **system** from the **UNIX groups list**.
18. Click the **Add** button to add a mapping for these two groups.
19. Map all needed groups from the Windows system to according groups from the UNIX system (Figure 10-70).

Important: Be aware of the Windows domain box, which says \\ITSOLAB2. That is where users and groups will be mapped to, and that is where authentication will be done. If you do not select the right domain there, then you will have file access problems on the UNIX client side.

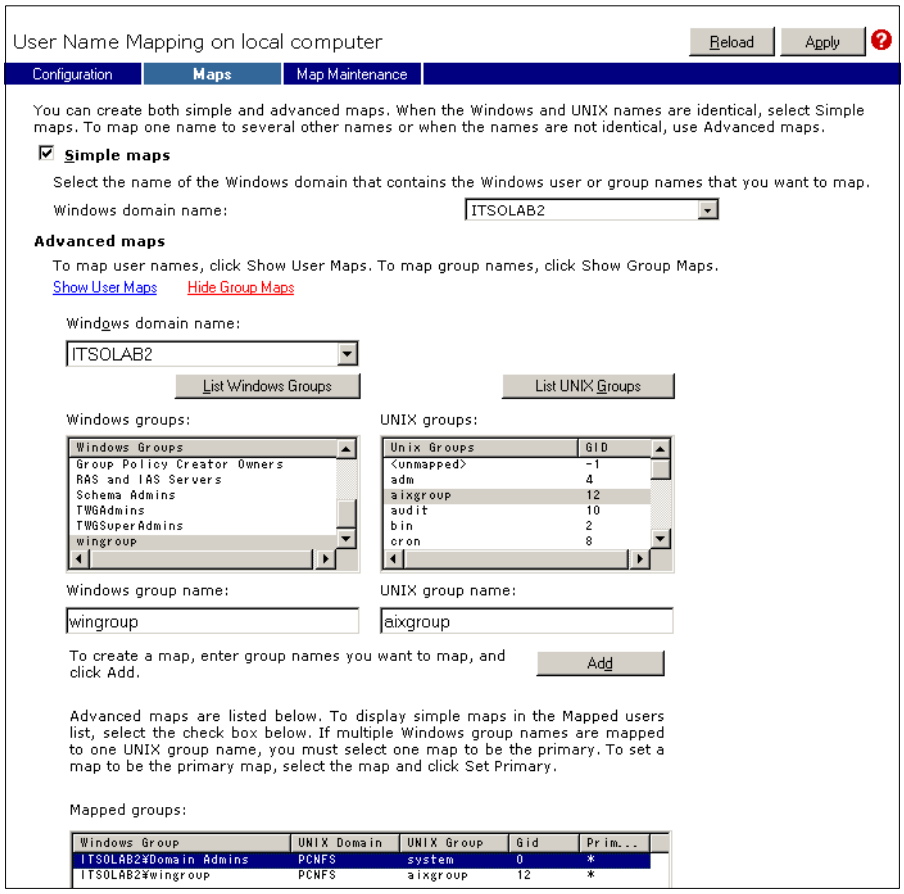

Figure 10-70 Group mapping on NAS 100

The Microsoft Services for UNIX are now configured within the Windows Domain.

Note: Make sure everything is applied before you close the SFU configuration. Step back to the **Configuration** tab and click the **Synchronize Now** button to refresh user and group names.

20. Open a DOS prompt and verify the User and Group mapping by the DOS command **mapadmin list -all** (Figure 10-71).

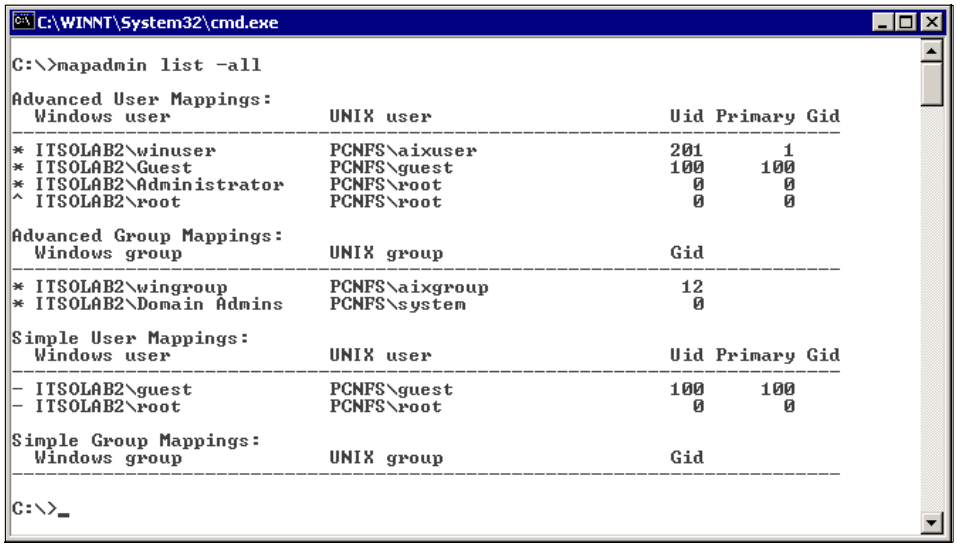

Figure 10-71 Verifying Maps by DOS command mapadmin

Congratulations! Now you have successfully completed your Services for UNIX setup. You should be able to mount shared folders into the file system to your UNIX clients.

For further steps, refer to “Configuring the shared storage” on page 412.

10.4.2 Configuring cross platform share without a Domain Controller

Log on via Terminal Services Client or NAS 100 Web GUI interface->Maintenance->Terminal Services to a NAS 100 with administrative access. Open the **IBM NAS Admin** by clicking its icon on the desktop.

Creating new Windows User accounts:

These are the steps for creating new Windows User accounts:

1. Click the **+** sign beside NAS Management.
2. Click the **+** sign beside Users and Groups.
3. Click the **+** sign beside Local Users and Groups (Local).
4. Click **Users**.
5. Create a new user called **root**.

6. Add the new user **root** to the **Administrators** group.
7. Create a new group called **support**.
8. Create a new user called **charlie**.
9. Add the new user **charlie** to the **support** group.

Add any other new Windows users that you will want to map to corresponding UNIX users at this time.

Configuring the Services for UNIX components

These are the steps for configuring the Services for UNIX components:

1. Click the **+** sign beside NAS Management.
2. Click the **+** sign beside Maintenance.
3. Click **Services**.

Scroll down the list of services and verify that the following services are Started:

- Gateway for NFS
- Server for NFS
- Server for PCNFS
- User Name Mapping

Configuring Server for NFS

These are the steps for configuring Server for NFS:

1. Click the **+** sign beside NAS Management.
2. Click the **+** sign beside File Systems.
3. Click the **+** sign beside Services for UNIX.
4. Click **Server for NFS**.
5. Click **User Mapping** tab and you will see a display like the one in Figure 10-72.

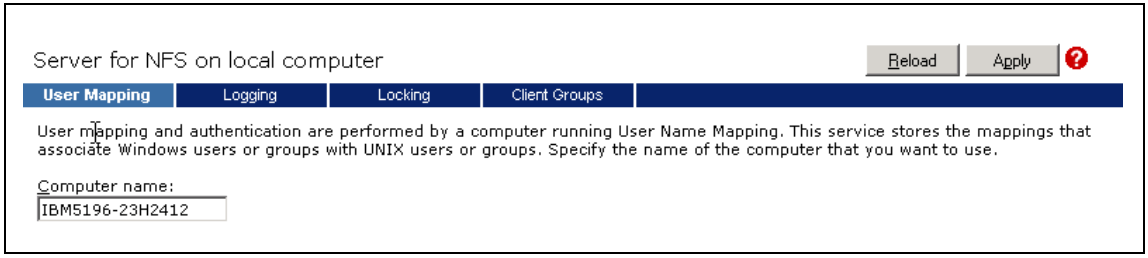

Figure 10-72 Server for NFS user mapping

Specify the computer name of the server that is running the *User Name Mapping* Service. If *User Name Mapping* will be configured to run on this NAS appliance, specify localhost or the machine name.

6. Click the **Client Groups** tab (refer to Figure 10-73).

   Root access to the NFS share is granted on a per-machine basis. It is very common that some machines will need root access to the NFS share. Here we will create a special client group (or netgroup) and add machine names to the group:

7. Create a client group called **root_users** by supplying a name in the **Group name**: field and selecting the **New** Button. The new group should appear in the **Current groups** list. Select the **root_users** group from the **Current groups** list and select the **Advanced** link.

8. Add the IP address, or hostname, and click **Add Clients** for each machine that you want to add to this group.

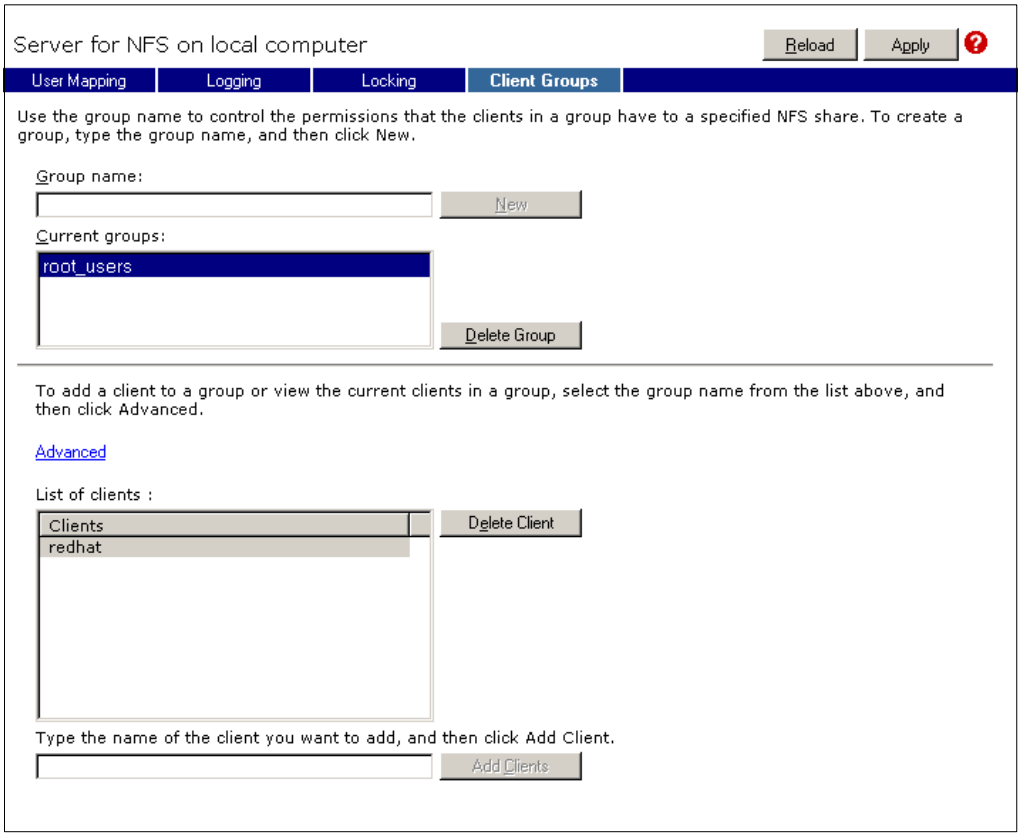

Figure 10-73 Server for NFS client groups

9. You can create other client groups at this time if needed. This completes the configuration of **Server for NFS**. Click the **Apply** button in the upper right corner to save your changes.

Configuring the Gateway for NFS

Next we describe how to configure the Gateway for NFS.

1. Click **Gateway for NFS**.
2. Select the **Authentication** tab and specify the computer name running the **User Name Mapping Service**. In this case we will specify the NAS appliance host name (localhost would also work here!). This is shown in Figure 10-74.

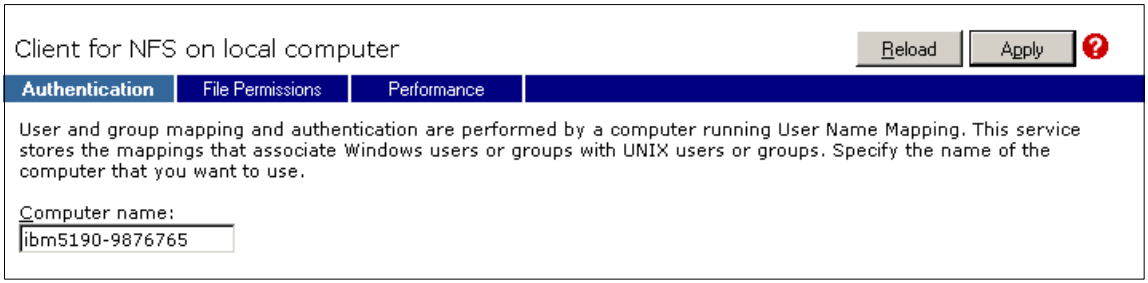

Figure 10-74 Gateway for NFS authentication

This completes the configuration of Gateway for NFS. Click the **Apply** button in the upper right corner to save your changes.

Configuring the Server for PCNFS

Next we describe how to configure the *Server for PCNFS*.

3. Click Server for PCNFS, and the window shown in Figure 10-75 will appear when you click the **Groups** tab.

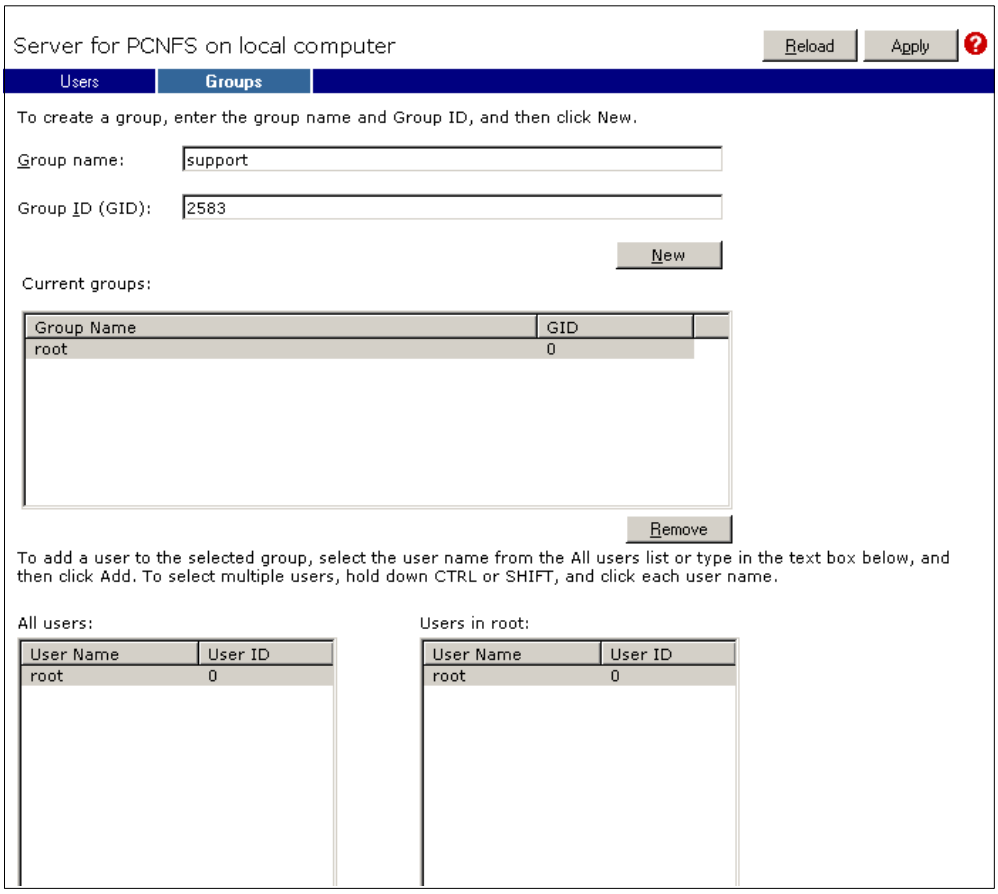

Figure 10-75 Server for PCNFS new group

4. Create a new group called **root** with GID (group ID) of 0. Create another new group called **support** with GID of 2583. Add other UNIX groups at this time being careful to specify the GID exactly as it is specified in the */etc/group* file in your UNIX clients or NIS database.
5. Create a new user **root** with a UID (user ID) of 0. The primary group for this new user is root. Create a new user **charlie** with GID of 505 and primary group **support**. Add other UNIX users at this time being careful to specify the UID exactly as it is specified in the */etc/passwd* file in your UNIX clients or NIS database.
6. Click the **Groups** tab again.

7. Add the new user **root** as a member of group **root.** Add the new user **charlie** as a member of group support. This is shown in Figure 10-76.

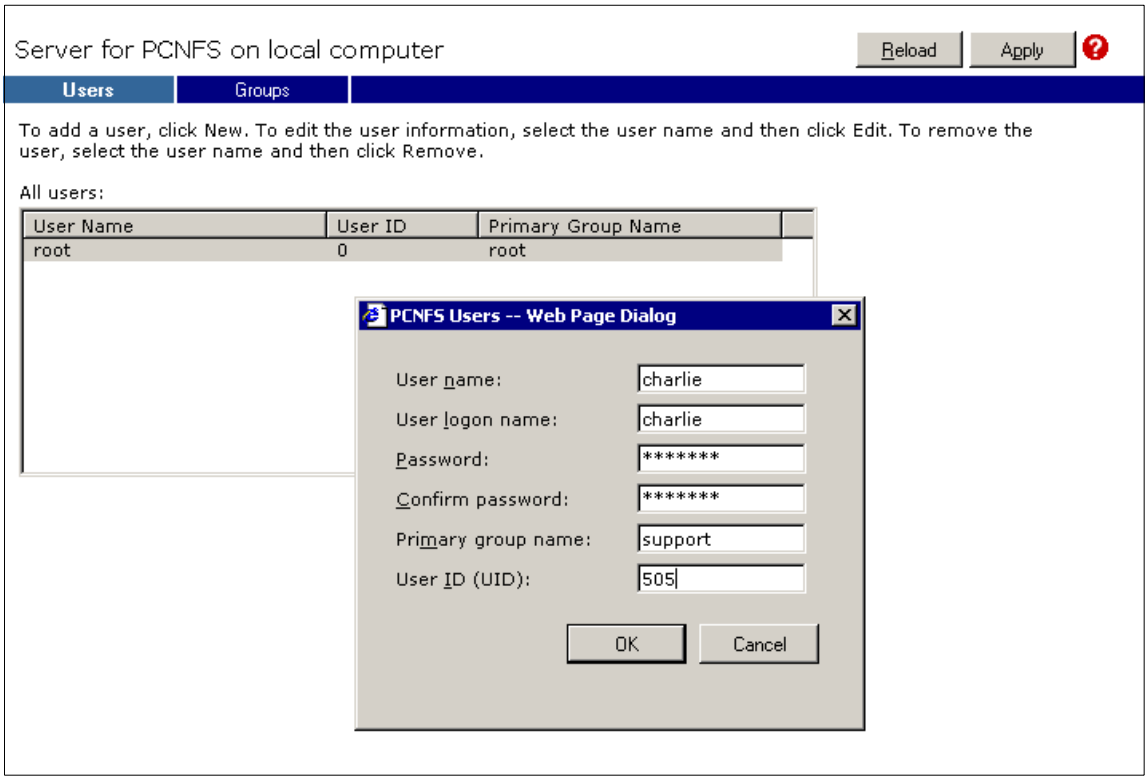

Figure 10-76 Server for PCNFS new user dialog

8. Add other UNIX users to their respective groups as is specified in the */etc/passwd* file in your UNIX clients or NIS database.

This completes the configuration of *Server for PCNFS*. Click the **Apply** button in the upper right corner to save your changes.

Configuring the User Name Mapping

Next we describe how to configure the *User Name Mapping*.

1. Click **User Name Mapping**.
2. Select the **Configuration** tab.

   Make sure that the **Personal Computer Network Files System (PCNFS)** option is selected. Verify that the group and password file paths are correct as shown in Figure 10-77.

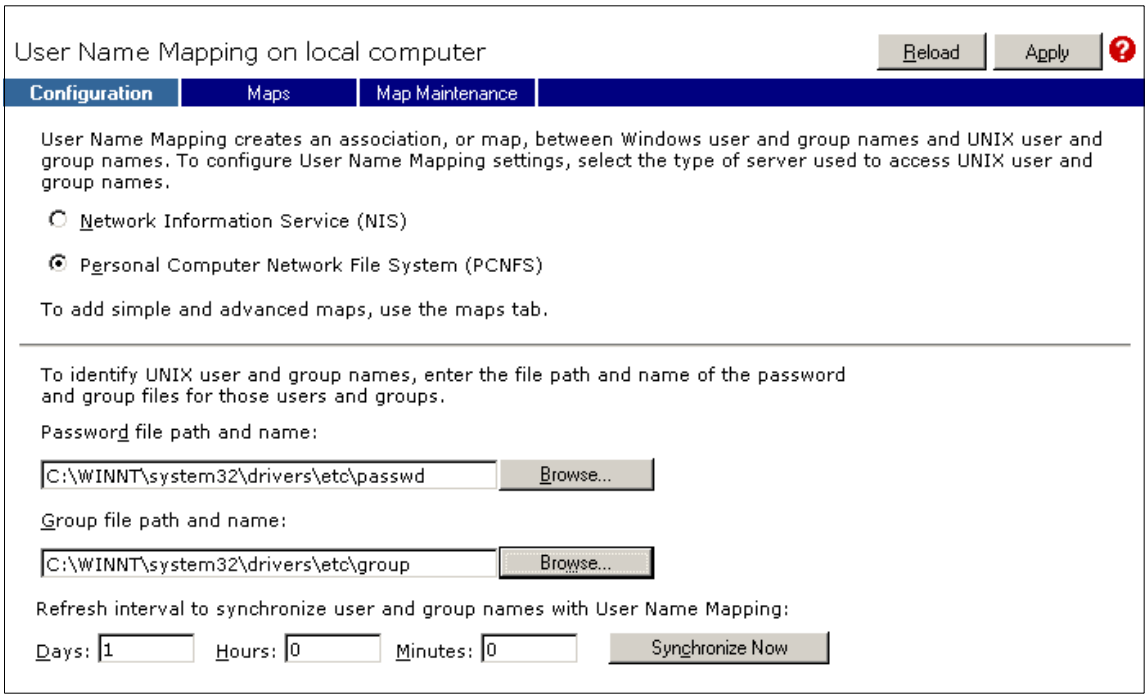

Figure 10-77 User Name Mapping

3. Click the **Maps** tab.
4. Click the **Show User Maps** tab.
5. Select the **List Windows Users** and **List UNIX Users buttons**.
6. Select the Windows user **root** from the Windows users list and the UNIX user **root** from the UNIX users list.
7. Click the **Add** button to add a mapping for these two users.
8. Repeat the process for the user **charlie**. You can create mappings between other UNIX and Windows users at this time. This is shown in Figure 10-78.

Important: Be aware of the Windows domain box, which says \\ibm5190-9876765. That is where users and groups will be mapped to, and that is where authentication will be done. If you do not select the right domain there, then you will have file access problems on the UNIX client side.

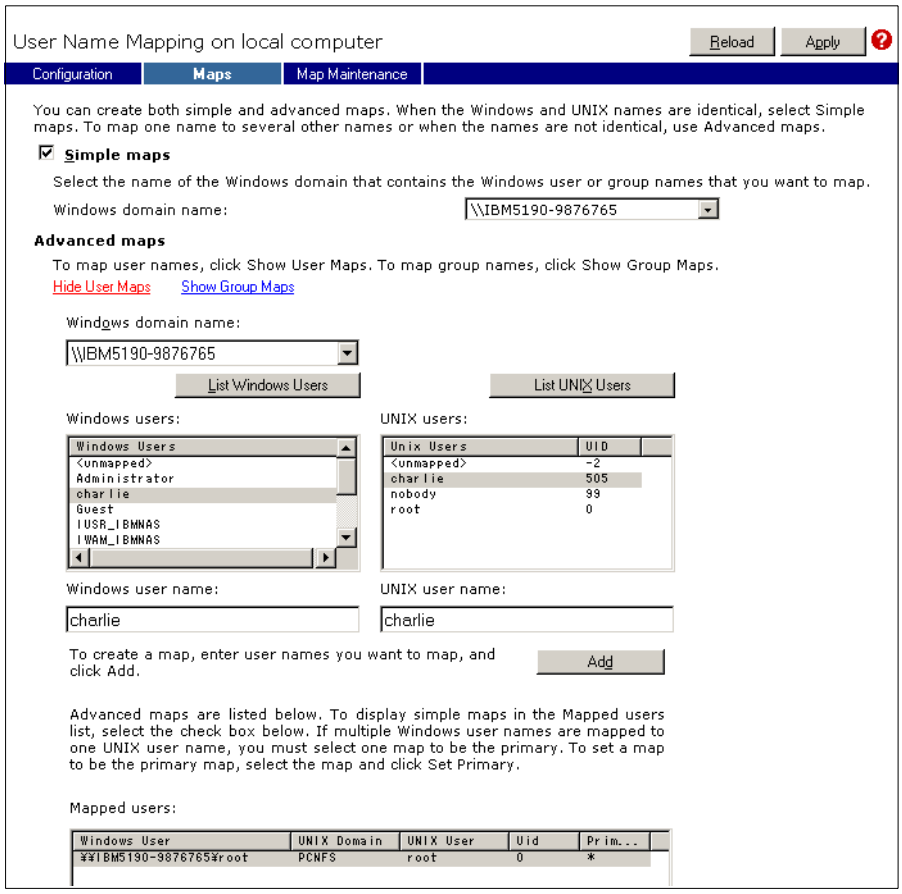

Figure 10-78 User Name Mapping user configuration

9. Click the **Show Group Maps** link.
10. Select the **List Windows Groups** and **List UNIX Groups** buttons.
11. Select the Windows group **Administrators** from the **Windows groups list** and the UNIX group **root** from the **UNIX groups list**.
12. Click the **Add** button to add a mapping for these two groups.
13. Repeat the process for the group **support**. The result is shown in Figure 10-79.

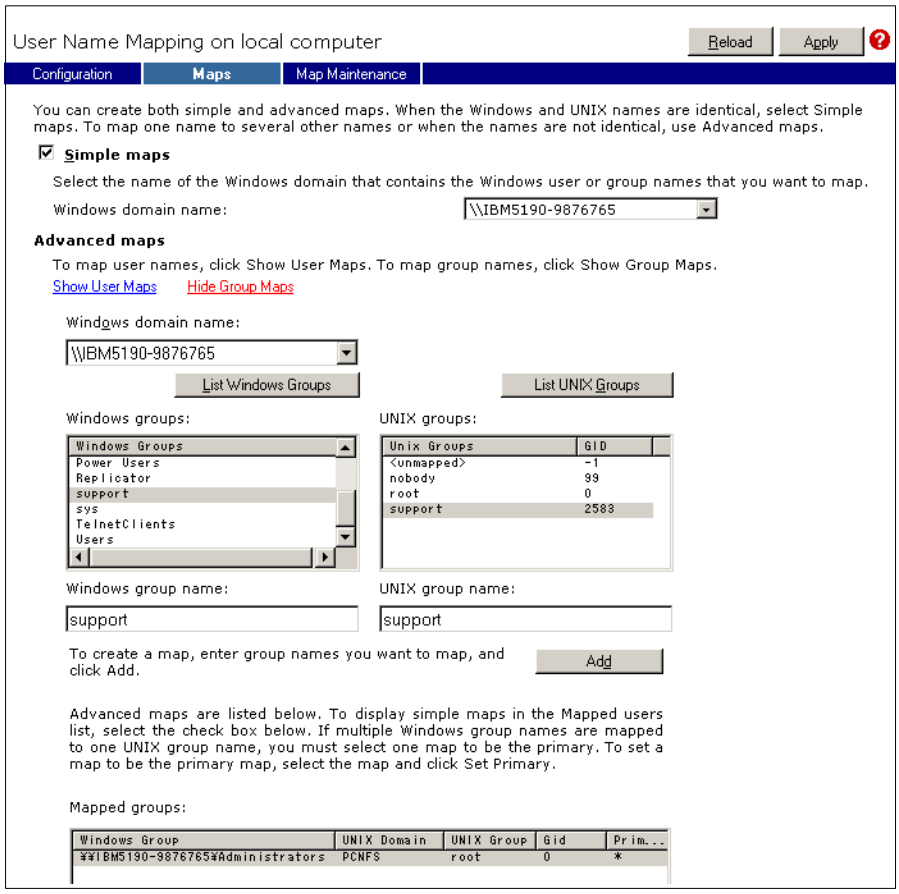

Figure 10-79 User Name Mapping group configuration

14. Create mappings between other UNIX and Windows groups at this time.

This completes the configuration of *User Name Mapping*. Click the **Apply** button in the upper right corner to save your changes.

Congratulations! Now you have successfully completed your Services for UNIX setup. You should be able to mount shared folders into the file system to your UNIX clients.

10.4.3 Configuring the shared storage

From Windows explore, create a new folder and name it **interop**. Right-click the folder to edit the *Properties* sheets.

1. Select the **NFS Sharing** tab.
2. Select the **Share this folder** radio button.
3. Name the share **nfs_interop** (Figure 10-80).

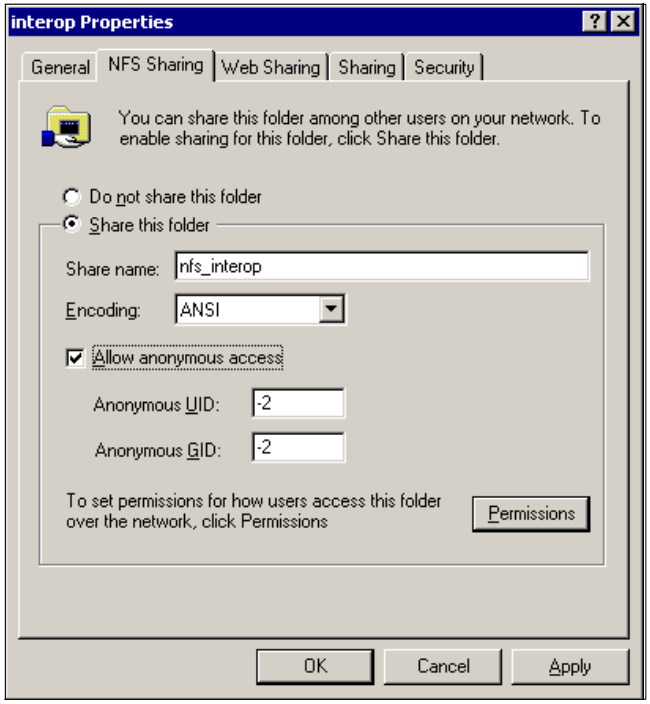

Figure 10-80 NFS sharing tab

4. Check the box for **Allow anonymous access**.
5. Click the **Permissions** button, and the **NFS Share Permissions** dialog will appear.
6. Click the **Add** button to give access to a new machine (Figure 10-81).

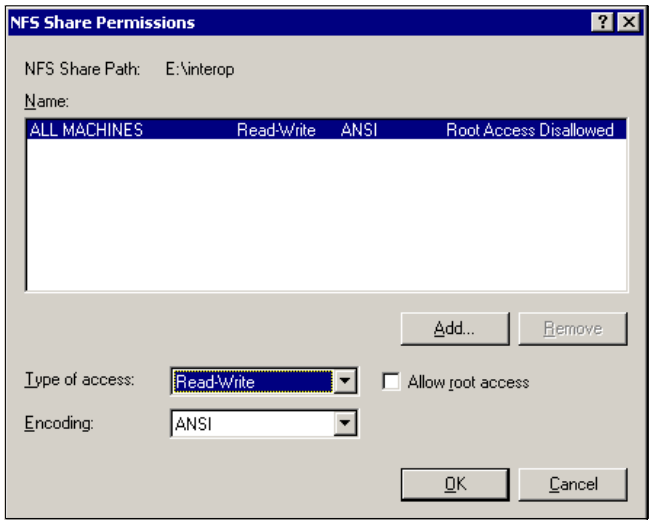

Figure 10-81 NFS share permissions

7. Click the **Add** button.
8. Select the **root_users** client group and then click the **Add** button.
9. Change the **Type of Access** to **Read-Write.**
10. Check the **Allow root access** box and click the **OK** button.
11. Machines in this group will have **root** access to the share (Figure 10-82).

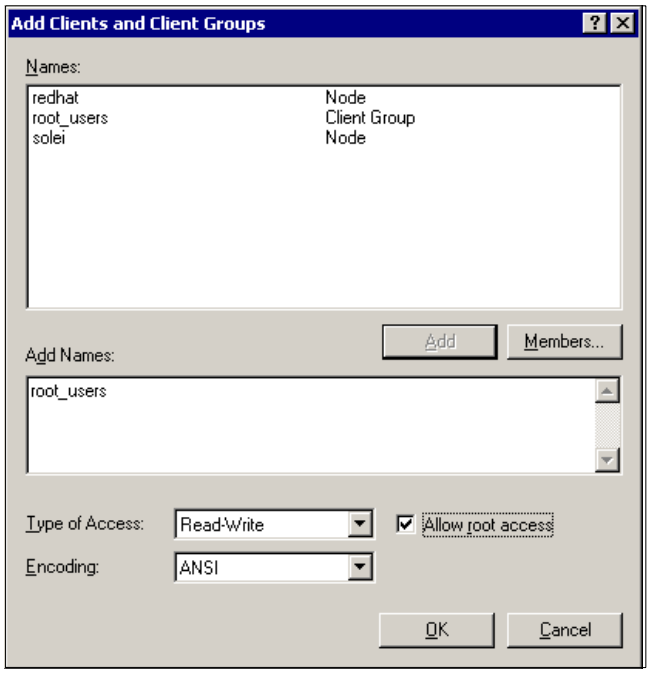

Figure 10-82 NFS share add clients and groups

12. Add the following users/groups to the permissions list:
    - Administrator user, Full Access
    - Administrators group, Full Access
    - User *root*, Full Access
    - User *charlie*, Read and Execute, List Folder Contents, Read

    This is shown in Figure 10-83.

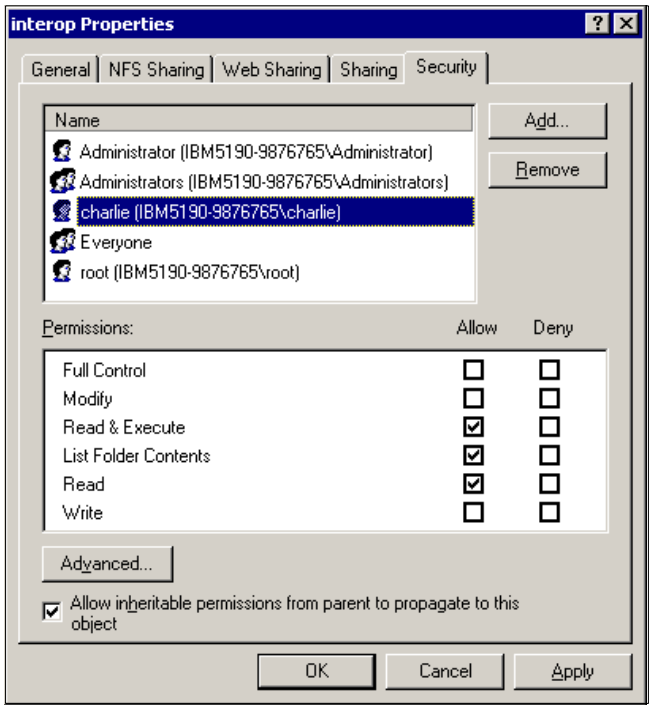

Figure 10-83 Security permissions

13. Add other users and groups as appropriate at this time.
14. unselect the “Allow inheritable permission...” to set the permission for the Windows group everyone which is UNIX nobody (Figure 10-84).

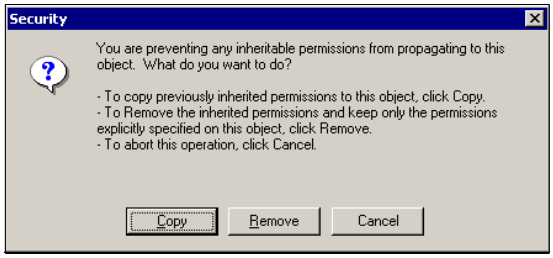

Figure 10-84 Unselect “Allow inheritable permission...”

Select Copy to choose the permissions individually.

15. Click the **Advanced** button.
16. Check the box labelled **Reset permissions on all child objects and enable propagation of inheritable permissions**.

17. Make sure the box labelled **Allow inheritable permissions from parent to propagate** to this object is checked (Figure 10-85).

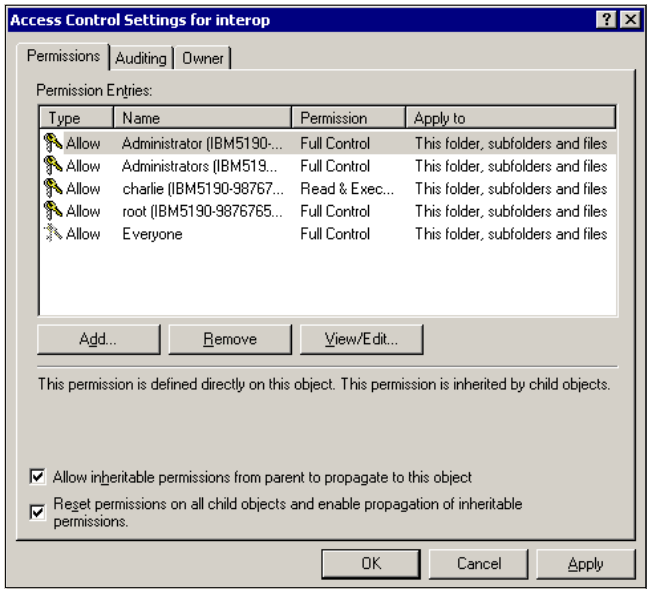

Figure 10-85 Advanced security permissions

18. Click the **Owner** tab.
19. The **Access Control Settings** dialog will appear.
20. In the **Change owner to:** list, select the **Administrator** user.
21. Place a check mark in the box labelled **replace owner on subcontainers and objects**.
22. Click the **Apply** button, then click **OK** (Figure 10-86).

Note: You will have to add every group and also the users that are in the groups to reflect the access rights. Another possibility is to add groups only like you would do in Windows, but then the access rights for the specific “owner” of the file will not be displayed, they are just blank. A workaround is to set the registry key Mapping “Implicit Permission = 1”.

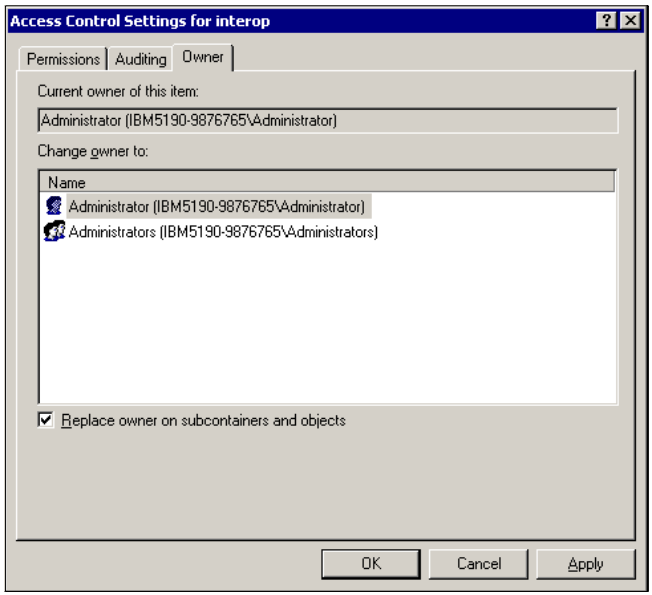

Figure 10-86 Access control settings

23. Exit out of the properties sheet for the new folder.

At this point, the NFS share has been created and is now available for mounting by the UNIX clients. Following the example taken thus far, the **`mount`** command may look something like this:

```
mount ibm5190-9876765:/nfs_interop /nfs_mount,
```

In this command, **nfs_mount** is the local mount point on the UNIX client.

In the next step, we will create a *Gateway for NFS* share that the Windows clients can use.

Gateway for NFS configuration

From the NAS Desktop:

1. Click the **Start** button.
2. Click **Programs**.
3. Click **Windows Services for UNIX**.
4. Click **Gateway for NFS Configuration**.

The Gateway for NFS Configuration screen will appear (Figure 10-87).

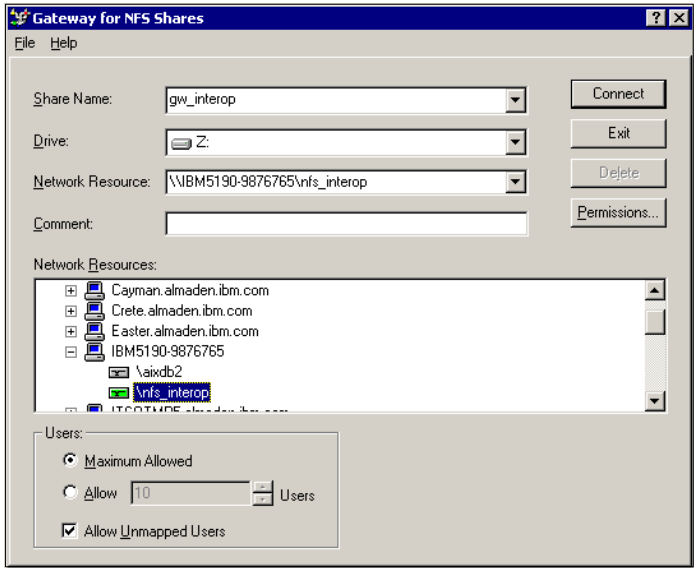

Figure 10-87 Gateway for NFS shares

5. Click the **+** sign beside Default LAN.
6. Click the **+** sign beside the NAS Appliance (for example, IBM5190-9876765).
7. Click the **nfs_interop** drive. The **Network Resource:** field will show the NFS share that was created above.
8. Select the drive letter that you would like to assign on the NAS appliance from the **Drive:** list (for example, Z:) and specify the Gateway for NFS share name in the **Share name:** field.
9. Select **Connect** to create the share.

At this point, the NAS appliance has exported a new share (gw_interop) that can be mapped from any Windows client on the network. In addition, a new drive (Z:) has been created on the NAS appliance. When a mapped user (User Name Mapping) with access privileges to the share is logged on at the NAS desktop, they can add or edit files and directories from this drive.

In the next section we illustrate how to map the gateway share from a Windows client on the network.

10.4.4 Mapping the Gateway for NFS share from a Windows client

This example is for a Windows 2000 Professional client on the same network with the NAS appliance.

1. Right-click **My Network Places**.
2. Click **Map Network Drive**, and the **Map Network Drive** dialog will appear.
3. In the **Folder:** field, specify the machine\resource that you want to map. The machine name is abbreviated somewhat in order to show that we have specified the share called **gw_interop** as was specified when we created the Gateway for NFS share previously (Figure 10-88).

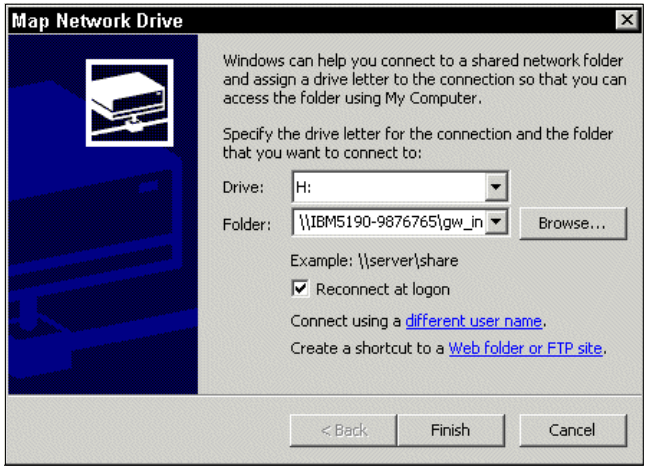

Figure 10-88 Client mapping

4. Click **Connect using a different user name**. We will log on as the user **charlie,** since that user has been given access to the share.
5. Enter the user ID and password and click **OK** (Figure 10-89).

Figure 10-89 Client mapping with a different user

6. Then click **Finish** to create the mapped drive.

When the drive is created, you can access the drive from the Windows Explorer to create new files and folders.

Logon to your UNIX client as user *charlie* and mount the **nfs_interop** share that is exported by the NAS appliance. Verify that the files created from the Windows client (Figure 10-90) have the **charlie:support** ownership (Figure 10-91). Create new files from the UNIX client and verify that the files can be seen from the Windows client.

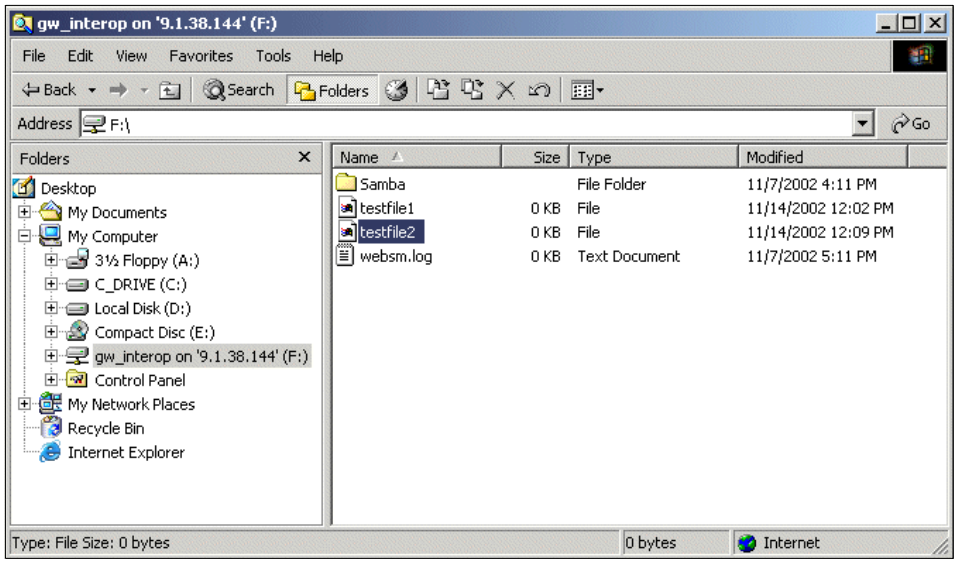

Figure 10-90 Created testfile2 from Charlie on the shared gw_interop folder

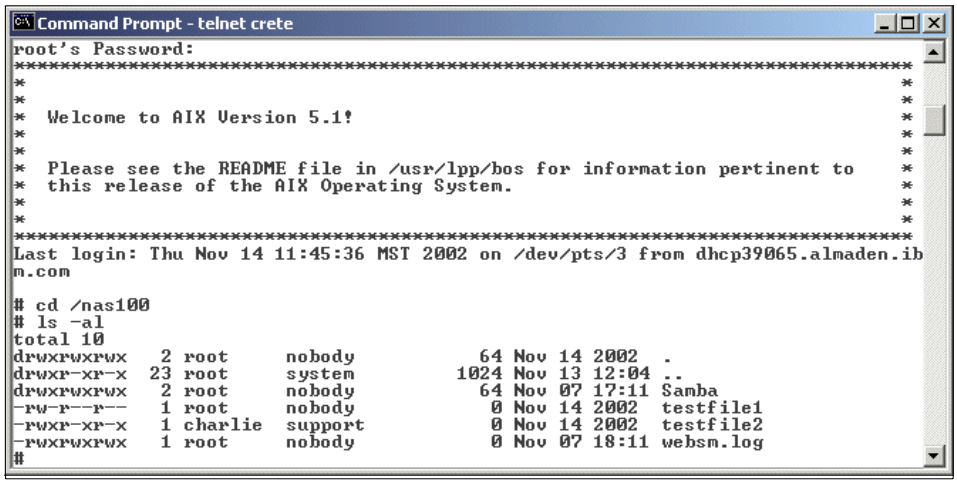

Figure 10-91 testfile2 created from charlie:support in Windows on NFS-share

How to enable character translation (Q289627)

We discuss the problems that can occur due to different character mapping on Windows and UNIX operating systems. In some cases you may encounter a problem while sharing files between both environments.

There are several restrictions used in Windows and UNIX on valid characters for naming a file. To solve any problem that can relay on an invalid use of characters and character translation, refer to the Microsoft Web site:

`http://support.microsoft.com/default.aspx?scid=KB;en-us;289627&`

10.4.5 Accessing the shares from our UNIX clients

We will now explain how to get access to the shares created on the NAS system from the UNIX world.

Accessing the shares from our Linux/Solaris/HP-UX clients

Connecting to the shared disks from our RedHat Linux 8.0 client, *dhcp39070*, was just as easy. First, we modified the `/etc/fstab` file to include a listing for the shared disk.

Note: Under Solaris, the /etc/vfstab file is updated, rather than the /etc/fstab.

This is shown on the last line in Figure 10-92.

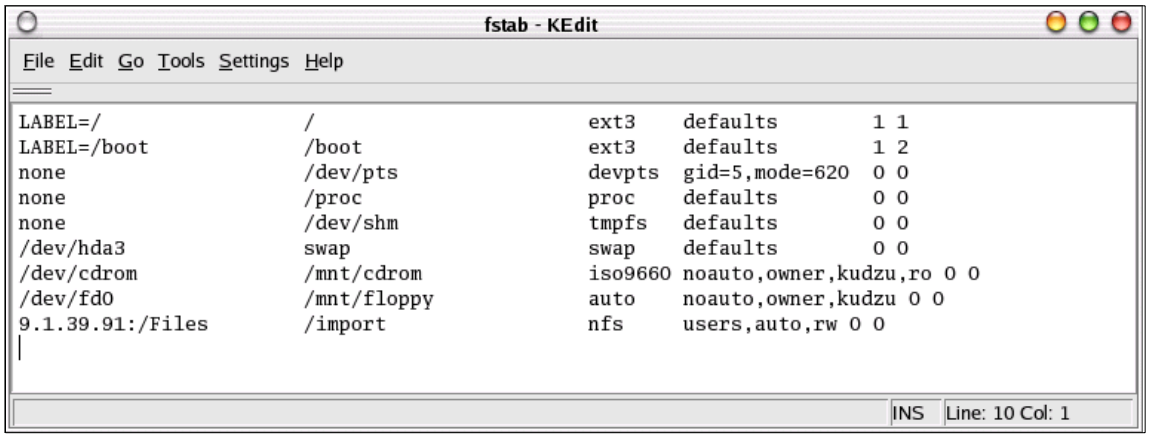

Figure 10-92 Adding the NAS's shared disk to the Linux fstab file

Once that was done, we created a directory named */import* and mounted the shared directory */Files* to it normally using the **mount** command, as shown in Figure 10-93. We then changed directory to */import* and were immediately able to see all of the data on the shared disk.

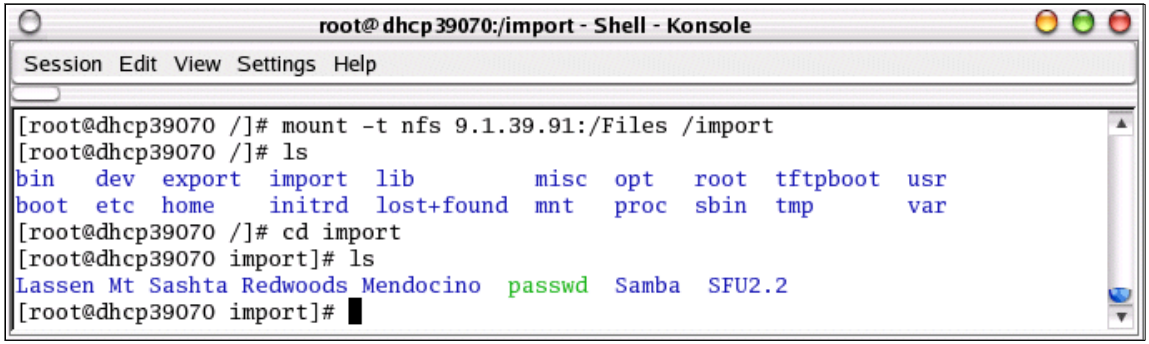

Figure 10-93 Mounting the NAS's shared directory from a Linux client

We opened some of the files and created new ones. This worked beautifully and we were able to see the changes from the Windows clients. As a further test, we created a text file from the Linux client, saved it, and left it open. We then tried accessing the file from one of the Windows clients. While we were able to open the file normally, we were pleased to note that Windows recognized that the file was still in use on the Linux system and did not let us overwrite it.

With a minimum of effort, we were able to use the NAS 100 to safely share storage among heterogeneous LAN/WAN clients.

Accessing the shares from our AIX clients

Now we will show you the few commands we used to get AIX ready to use the NAS 200 and 300's shared disks.

First we needed to update /etc/filesystems. This was accomplished by using the **crfs** command, as shown in Example 10-2.

Example 10-2 Using the crfs command

```
# crfs -v nfs -m /nas100 -n ibm5190-9876765 -d nfs_interop -A yes
# cat /etc/filesystems
/nas100:
        dev             = nfs_interop
        vfs             = nfs
        nodename        = ibm5190-9876765
        mount           = true
        account         = false
```

Finally, we mounted the share, as shown in Example 10-3.

Example 10-3 Mounting the share

```
# mount -v nfs ibm5190-9876765:nfs_interop /nas100
# mount
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
/dev/hd4         /                jfs    Nov 11 16:50 rw,log=/dev/hd8
/dev/hd2         /usr             jfs    Nov 11 16:50 rw,log=/dev/hd8
/dev/hd9var      /var             jfs    Nov 11 16:50 rw,log=/dev/hd8
/dev/hd3         /tmp             jfs    Nov 11 16:50 rw,log=/dev/hd8
/dev/hd1         /home            jfs    Nov 11 16:51 rw,log=/dev/hd8
/dev/lv00        /usr/welcome_arcade jfs    Nov 11 16:51 rw,log=/dev/hd8
/dev/lv01        /usr/welcome     jfs    Nov11 16:51 rw,log=/dev/hd8
ibm5190-9876765 nfs_interop /nas100 nfs3 Nov 11 17:17
```

10.5 Accessing the shares with the Samba client

Samba is an implementation of the Server Message Block (SMB) protocol, a subset of Common Internet File System (CIFS), that can be run on almost every variant of UNIX in existence. Samba is an open source project just like Linux. The entire code is written in C so it is easily ported to all flavors of UNIX.

In a Windows network, with IBM TotalStorage NAS in place, if you have a Linux host, it can connect to the IBM NAS machine without having additional protocols but CIFS in your network. This will simplify your network setup and maintenance.

Samba has two components:

Samba server This allows UNIX systems to move into a Windows “Network Neighborhood” in the same way a Windows NT/2000 file server does. With Samba, UNIX servers are acting like any other Windows server, offering their resources to the SMB/CIFS clients.

Samba client This allows Linux or Unix hosts to access any shared directory or printer on Windows NT/2000 servers or Samba servers, by allowing these machines to access Windows files with CIFS/CIFS protocol.

10.5.1 Setting up the Samba client on a RedHat Linux 8.0

The Samba client will allow you to take advantage of the Samba File System (SMBFS). With SMBFS you can mount any share from a Windows NT/2000 server or Samba server into your directory structure.

This is available on Linux only, and it gives you two choices about how to access Windows NT/2000 files:

- Mounting a NAS Share into the Linux file system
- Using the smbclient program

The first step is to check if Samba is already installed in your system. To do this, query the RPM database:

```
rpm -q samba
```

This command will return either the version number of the installed package or a message indicating that the package is not installed. If Samba is not installed, mount the Red Hat 8.0 CD-ROM (Disc 2) and install the following packages:

```
mount /mnt/cdrom
# rpm -ivh /mnt/cdrom/RedHat/RPMS/samba-common-2.2.5-10.i386.rpm
# rpm -ivh /mnt/cdrom/RedHat/RPMS/samba-client-2.2.5-10.i386.rpm
```

10.5.2 Mounting a NAS Share into the Linux file system

We want to have access from our RedHat 8.0 machine to a file share configured in the IBM TotalStorage NAS 100 named *ibm5190-9876765*. The share we want to access is named *Files*.

One-time mount

If we try to **mount** the share as a user without sufficient rights, we are allowed to do it. But, when accessing the files, permission is denied.

```
#mount -t smbfs -o username=user1,password=pwd1 //ibm5190-9876765/Files
/nas100
#ls /nas100
```

We get the following message:

```
#ls: /nas100: Permission denied
```

We have to **umount** the share and **mount** it again as a user with sufficient rights

```
#mount -t smbfs -o username=user2,password=pwd2 //ibm5190-9876765/Files
/nas100
#ls /nas100
#file1 file2
```

Now, we are able to access and manage the files in *//ibm5190-9876765/Files* as they were stored in the local Linux host disk.

Permanent mount

If you don't want to map the share every time you need it, there is a procedure to automatically map it each time the client machine boots.

We included the following line at the end of the */etc/fstab* file:

```
//ibm5190-9876765/Files    /nas100    smbfs    username=user2,password=pwd2
   0   2
```

There is a better solution if you don't want to show your password in the /etc/fstab file. A credentials file can be created in your private home directory, so nobody (but root) can read it. Reference this file in */etc/fstab* in the following way:

```
//ibm5190-9876765/Files    /nas100    smbfs
   credentials=/root/credentials_file    0   2
```

Create a file called *credentials_file* in the */root* directory with the following content:

```
username = user2
password = pwd2
¿ domain = XXXXXx ?
```

Note: The user must be a valid local account on the NAS host or a valid account in the Windows NT/2000 domain.

10.5.3 Using the smbclient program

The **smbclient** is an FTP-like client to access SMB/CIFS resources on servers. This client can open a connection to an SMB/CIFS server in the same way as the FTP client does. Using this command we don't need to **mount** a share, just establish a connection and talk with the NAS machine.

The command line is as follows:

```
smbclient {servicename} [password] [options]
```

All the options are listed in the **man** pages. The most common options are:

servicename	The **servicename** is the name of the service you want to use on the server. Our case //nas200/fileshare.
password	The password required to access the specific share.
-R <name resolve order>	This option is used by the programs in the Samba suite to determine what naming services and in what order to resolve host names to IP addresses.
-N	If specified, this parameter suppresses the normal password prompt from the client to the user.
-n NetBIOS name	By default, the client will use the local machine's hostname (in uppercase) as its NetBIOS name. This parameter allows you to override the host name and use whatever NetBIOS name you wish.
-p port	This number is the TCP port number that will be used when making connections to the server. The standard (well-known) TCP port number for an SMB/CIFS server is 139, which is the default.
-l logfilename	If specified, **logfilename** specifies a base filename into which operational data from the running client will be logged.
-h	Print the usage message for the client.
-I IP-address	IP address is the address of the server to connect to.
-U username[%pass]	Sets the SMB username or username and password.
-A filename	This option allows you to specify a credentials file.

We can open a connection to *fileshare* in *nas200* without parameters just by issuing the following command:

```
#smbclient //ibm5190-9876765/Files pwd2 -U user2
```

Once the client is running and connected, you will see a prompt like this:

```
smb:\>
```

In this prompt, the backslash ("\") shows us the current working directory and will change if the current working directory is changed. The prompt indicates that the client is ready and waiting to fulfill a user command. Commands like **cd**, **ls**, **del**, **help**, **mkdir**, **dir**... can be used. You can find a complete list of commands in the **man** pages of you Linux host.

Commands like **get**, **put**, **mget**, **mput** allow us to send or retrieve files in the IBM NAS TotalStorage server.

Use the **exit** command to close the connection.

10.5.4 Samba client configuration on AIX

Sometimes, you want to access your Samba server or shares from a Windows machine using AIX. You can use the **smbclient** program to do this. The **smbclient** program is a client that can communicate with a SMB/CIFS server. If you have installed Samba using the default path, you will find it in */usr/local/samba/bin*. It is a good idea to include this path in your user profile. If you want to do this, you have only to add the following line in your profile:

```
PATH=$PATH:/usr/local/samba/bin
```

This client has an interface very similar to the **ftp** program. You can use **smbclient** to get files from the server to the local machine, put files from the local machine to the server, retrieve directory information from the server, and so on.

Now that you have set up your profile, you can access your files in the Samba server. You can use some of the options in the command line shown in Example 10-4.

Example 10-4 Samba client command line interface

```
added interface ip=9.3.187.230 bcast=9.3.187.255 nmask=255.255.255.0
Usage: smbclient service <password> [options]
Version 2.0.6
        -s smb.conf           pathname to smb.conf file
        -O socket_options     socket options to use
        -R name resolve order use these name resolution services only
        -M host               send a winpopup message to the host
        -i scope              use this NetBIOS scope
        -N                    don't ask for a password
        -n netbios name.      Use this name as my netbios name
        -d debuglevel         set the debuglevel
        -P                    connect to service as a printer
```

```
  -p port               connect to the specified port
  -l log basename.      Basename for log/debug files
  -h                    Print this help message.
  -I dest IP            use this IP to connect to
  -E                    write messages to stderr instead of stdout
  -U username           set the network username
  -L host               get a list of shares available on a host
  -t terminal code      terminal i/o code {sjis|euc|jis7|jis8|junet|hex}
  -m max protocol       set the max protocol level
  -W workgroup          set the workgroup name
  -T<c|x>IXFqgbNan      command line tar
 -D directory          start from directory
  -c command string     execute semicolon separated commands
  -b xmit/send buffer   changes the transmit/send buffer (default: 65520
```

If you want to connect to the server without specifying any other parameter, you can use the following command:

```
smbclient //<Netbios Server Name>/<Service> -U <Username>
```

Note: You can use \\ instead of each / if you wish. You have to use two backslashes for each slash that you want to substitute. The first backslash acts as a character escape for the second one.

You can also use some options to modify the way that you are going to connect to the server. Here are some options:

-N This option is used to suppress the normal password prompt from the client to the user. This option is very useful when you want to access a server that does not require a password to be accessed.

-p This option is used to specify the TCP/IP port that you will use when making connections. The standard TCP/IP port number for a SMB/CIFS server is 139; so, if you do not use this option, your client will try to connect to the server using the 139 port.

-I This option is used to specify the IP address of the Samba server to which you are trying to connect. This is very useful if your client is having problems using the NetBIOS name resolution.

-O This option is used when you want to specify the socket option. Here is a list of the valid options:

- SO_KEEPALIVE
- SO_REUSEADDR
- SO_BROADCAST
- TCP_NODELAY
- IPTOS_LOWDELAY

- IPTOS_THROUGHPUT
- SO_SNDBUF
- SO_SNDLOWAT
- SO_RCVLOWAT

The last four options take an integer argument.

If you are successful in connecting to the server, you will be prompted for a password. If you enter a valid password, you will see the **smbclient** prompt as shown in Example 10-5.

Example 10-5 Samba client connection

```
# smbclient //lva111a/test -U root
added interface ip=9.3.187.230 bcast=9.3.187.255 nmask=255.255.255.0
Password:
Domain=[DOMAIN01] OS=[Unix] Server=[Samba 2.0.6]
smb: \>
```

If you have problems connecting to the server, you can use the **-R** option before the **-U** option to specify which name resolution services to use when looking up the NetBIOS name. The options are:

lmhosts: This option will use the Samba lmhosts file. You can find this file in the same directory as the smb.conf file. If you have installed your Samba server using the default path, you will find this on /usr/local/samba/lib.

host: This option uses the /etc/hosts file to resolve the names. This method of name resolution depends on the operating system that you are using.

wins: Use the WINS server set up in the smb.conf file. If you do not have one specified, this method will be ignored

bcast: This option does a broadcast on the interfaces listed in the interfaces parameter in the smb.conf file. This is not a good option to choose because it depends on the target host being on a locally-connected subnet.

Now that you are accessing the Samba server, you can execute the **smbclient** commands. The following is a list of some **smbclient** commands that you can use to work with your files:

cd: Changes the current working directory to the specified directory. This operation will fail if the specified directory does not exists or if you do not have access.

dir:	List the files in the current working directory. You can also use **ls** to list files.
mkdir:	Create a new directory on the server. You can use also **md**.
rmdir:	Remove a directory from the server. You can use also **rd**.
lcd:	Change the local machine directory to the one specified. If the specified directory does not exist or if you do not have access to this directory, the operation will fail.
get:	Copy the specified file from the current working directory on the server to the client. You can also use the **mget** command to copy multiple files that match a mask that you specify.
put:	Copy the specified file from the current working directory on the local machine to the remote server. You can also use the **mput** command to copy multiple files that match a mask that you specify.
Del.:	Delete all files in the current working directory that match the mask that you specify.You can also use the **rm** command.
help:	Display a brief description of the command, if you have specified one. If not, it will display a list of all available commands. You can use **?** instead of using the **help** command.
lowercase:	Toggle the option to get the files from the Samba server only in lowercase.
prompt:	Toggle the option for filename prompts during the operation of the **mget** and **mput** commands.
recurse:	Toggle the directory recursion for the **mget** and **mput** commands. When the toggle is on, this option will process all the directories in the source directory and will recourse into any that match the mask specified to the command.
setmode:	This option works like the **attrib** command in DOS. If you want to change the permission of a certain file to read only, you can, for example, use **setmode example.txt +r**.
exit:	This terminates the connection with the server and exits from the smbclient. You can also use **quit**.

10.5.5 Sources and additional information

You can find more information on the official Samba project Web site at:

http://www.samba.org

For good how-to documents, see the Linux documentation project home page:

http://www.linuxdoc.org/

11

Backup considerations

In our computing world today, data is considered the most important competitive differentiating factor. Temporary inaccessibility or the complete loss of data has a huge financial impact, and can drive companies out of business. The inability to manage data can have a negative impact on a company's profitability and limit its ability to grow. Storing, protecting, and managing data growth has become one of the major challenges of today's businesses.

The NAS 200 and 300 are designed to plug into your current data protection scheme — unlike many NAS appliances, they do not depend on special vendor-provided software, but work with out-of-the-box backup software. If you use Tivoli Storage Manager (TSM) as your enterprise backup solution, you will be pleased to know that the NAS 200 and 300 are shipped with a TSM client already installed on it. Also, if you do not already have a data protection scheme in place, the NAS 200 and 300 comes bundled with its own complete backup solution.

In this chapter we describe various configurations for backing up and restoring the NAS 200 and 300, either by using the native backup solution, or by integrating the NAS 200 and 300 into an existing TSM environment.

Tip: For more information on backup and recovery solutions for IBM TotalStorage NAS appliances, see the IBM Redbook, *IBM TotalStorage NAS Backup and Recovery Solutions*, SG24-6831, which is available at:

`http://www.redbooks.ibm.com/`

11.1 NAS 200 and 300 and their native backup solution

The NAS 200 and 300 come with a rich set of utilities for data management. One of the key advantages of using the NAS 200 and 300 is their ability to capture point-in-time image copies without the need for a long downtime window, by means of the Persistent Storage Manager (PSM) software. In the following sections we describe PSM and show how it can be used in conjunction with NTBackup to help increase productivity in backup and recovery of your mission critical data.

11.1.1 NAS 200 and 300 backup

The NAS 200 and 300 use two types of backup: *point-in-time image* copies and *archival backup.*

Point in time backup

Point-in-time images can provide a near-instant virtual copy of an entire storage volume. These point-in-time copies are referred to as *persistent images* and are managed by the Persistent Storage Manager (PSM) software.

These instant virtual copies have the following characteristics:

- Normal reads and writes to the disk continue as usual, as if the copy had not been made.
- Virtual copies are created very quickly and with little performance impact, as the entire volume is not truly copied at that time.
- Virtual copies appear exactly as the original volume when the virtual copy was made.
- Virtual copies typically take up only a fraction of the space of the original volume.

These virtual copies are created very quickly and are relatively small in size. As a result, functions that would otherwise have been too slow, or too costly, are now made possible. Use of these persistent images may allow individual users to restore their own files without any system administrator's intervention. With the pre-loaded code, the NAS administrator can schedule the PSM to automatically perform an instant virtual copy at regular intervals.

The administrator can also grant end users access to their specific virtual copies. If a particular user accidentally deletes or corrupts a file, he or she can just drag-and-drop the virtual copy of that file to their storage without any administrator involvement. If you would like to know more about this topic, refer to the redbook, *IP Storage Networking: IBM NAS and iSCSI Solutions,* SG24-6240-00.

Archival backup

Archival backup is used to make full, incremental, or differential backup copies, which are typically stored to tape. A common problem with these backups is that some files, which were open at the time the backup ran, often fail to get backed up. PSM on the NAS 200 and 300 is not hindered by open files, so it can successfully make backup copies in a 24x7 operation. Special precautions must be made to ensure data consistency when continuous write operations occur. See 11.2.1, “IBMSNAP utility” on page 433.

11.2 Using PSM with backup software solutions

For systems that need to be available 24 hours a day and 7 days a week, backup is a challenge. Traditionally, to ensure data integrity, it is necessary to shut down the application and close any open files before performing a backup. However, if a significant amount of data is involved, backup may take a while, especially when the transfer is directed to tape. This amount of downtime may not be acceptable for most systems, for various reasons, such as delayed transactions or lost profit.

By using PSM, there is no need to shut down the application, because of PSM’s capability to capture snapshots, even with open files. In combination with a backup solution, systems that have NAS boxes can now operate 24x7.

For this topic, we show how to take advantage of PSM image using the following backup solutions:

- NT Backup
- Tivoli Storage Manager (TSM)

11.2.1 IBMSNAP utility

To provide flexibility for your existing backup solution, persistent image functionality is implemented in such a way that any backup application that supports command line backups will be supported. This includes Tivoli Storage Manager (TSM), Veritas Backup Exec, and many others.

The IBM NAS systems come with a command line utility called *IBMSNAP.EXE* that allows you to use PSM technology with your existing backup software solution. Usage of this utility requires knowledge of Windows batch file processing plus the backup software command line utility. *IBMSNAP.EXE* is a command line utility that creates a PSM image, launches the backup batch file, and then sets the archive bits accordingly on the drive being backed up.

The *IBMSNAP.EXE* file is located in the *C:\ibm\NASBackup* directory.

Usage of *IBMSNAP.EXE*:

```
ibmsnap /l:{drive} /files:{backup_script_file} /exit
```

In this command:

{drive} = volume you want to have a PSM image of

{backup_script_file} = script or batch file that PSM runs after creating image

/exit = exit and close PSM

11.2.2 Using IBMSNAP with NTBackup

Note: We performed the following examples on a NAS200 with PSM Version 1.0. We did not test PSM on the NAS 300, which has PSM Version 2.0.

To use IBMSNAP with NTBackup, you will need to create a batch file on the NAS box that calls NTBackup with the necessary parameters:

1. Get connected to your NAS system.
2. From a command prompt, type **notepad *yourbatchfile*** (for example, ntback-h.bat). See Figure 11-1.

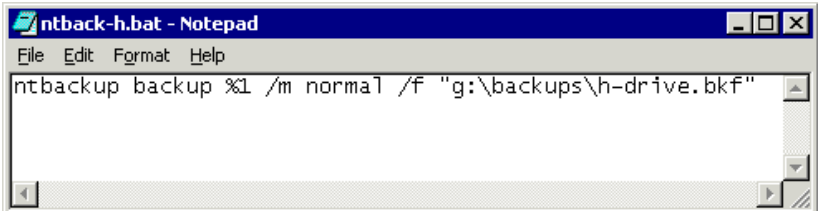

Figure 11-1 Sample batch file calling NTBackup

In this command:

%1 = persistent image virtual drive letter to be supplied by PSM automatically; the drive letter that will be used by PSM is the next available drive on your system

/m = mode (in the example, we selected normal)

/f = the backup file (and its location)

3. Run IBMSNAP (from C:\ibm\NASBackup directory) with the necessary parameters. For example:

```
ibmsnap /l:h /files:c:\winnt\system32\ntback-h.bat /exit
```

In this command:

h = drive to be backed up

ntback-h.bat = batch file to execute NTBackup

4. You should see a screen (on the NAS Terminal Services Client session) similar to the one shown in Figure 11-2.

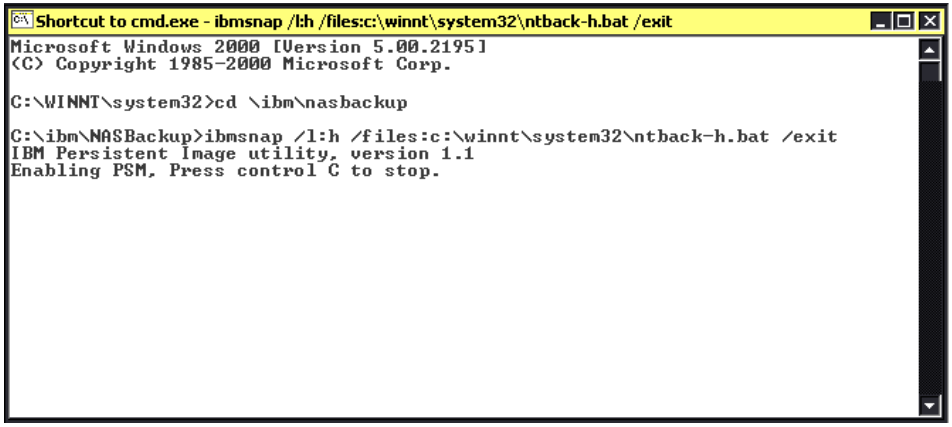

Figure 11-2 Screen showing IBMSNAP running and PSM creating an image

After that, NTBackup should start (Figure 11-3).

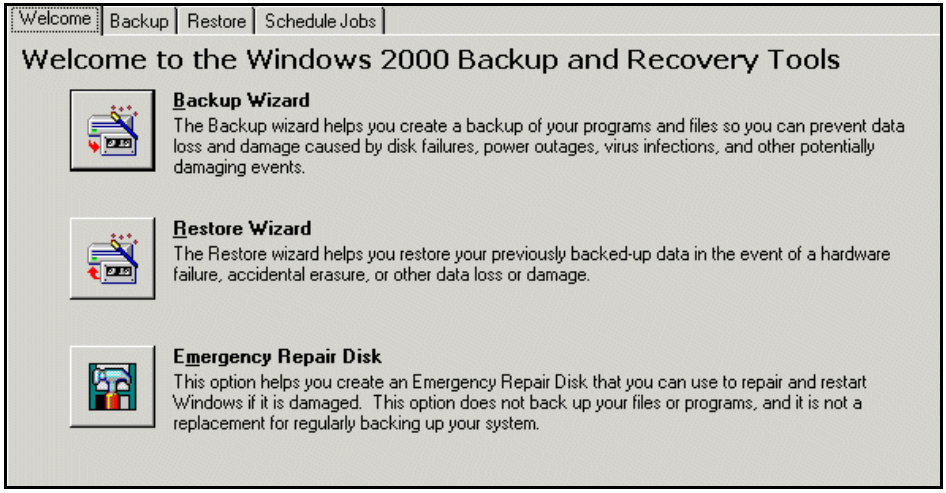

Figure 11-3 NTBackup started automatically by IBMSNAP

5. Next you should see that the backup is being done, as in Figure 11-4. This screen shows an on-going backup of removable disk F, which contains a PSM image of drive H.

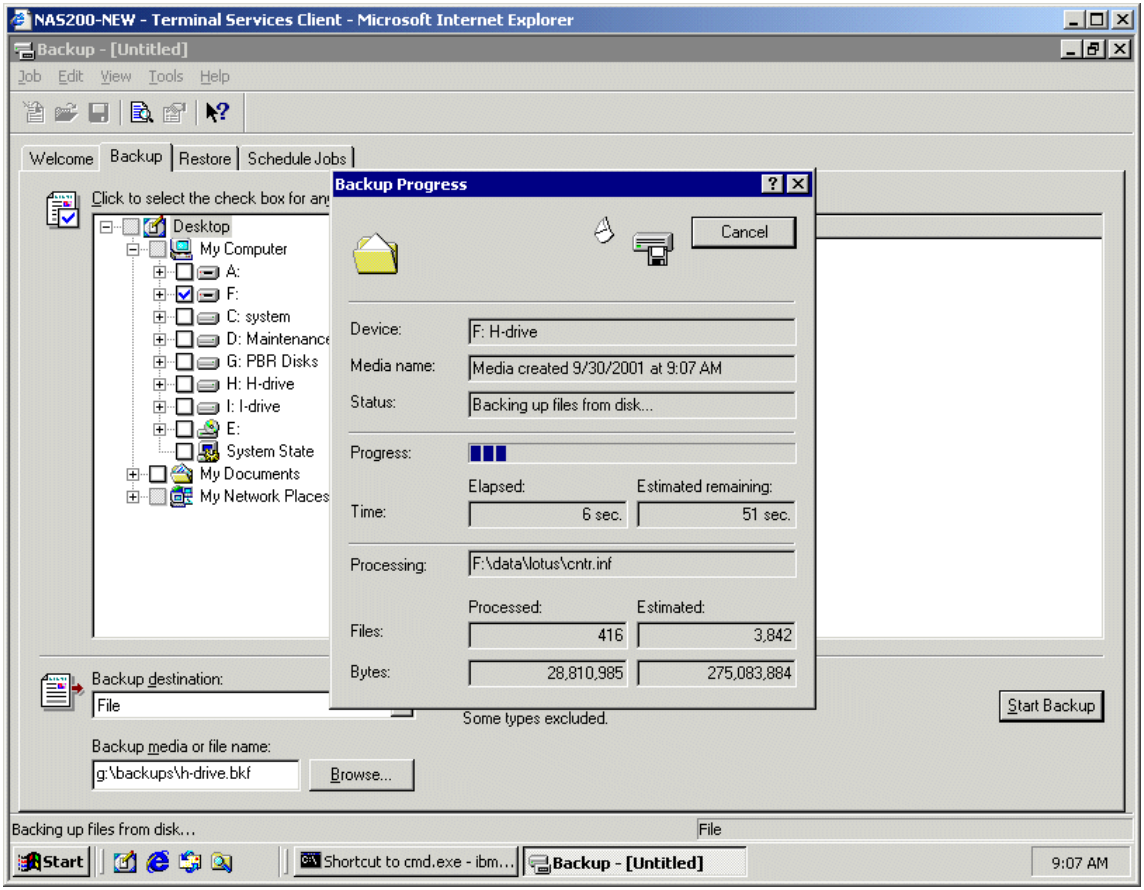

Figure 11-4 On-going backup of removable disk F with PSM image of drive H

6. Once the backup is completed, you will be returned to the command prompt. It should display a screen similar to the one in Figure 11-5, showing that IBMSNAP and NTBackup have completed successfully.

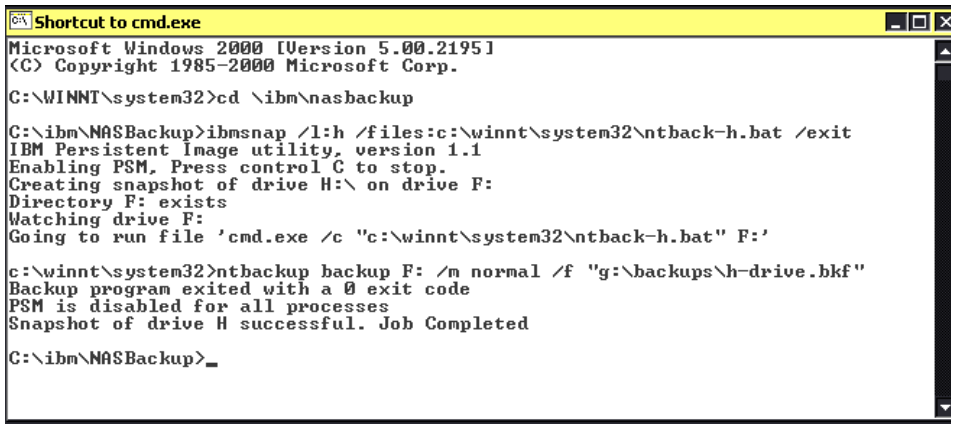

Figure 11-5 Successful completion of IBMSNAP and NTBackup

7. You can check if the backup file has been created (Figure 11-6).

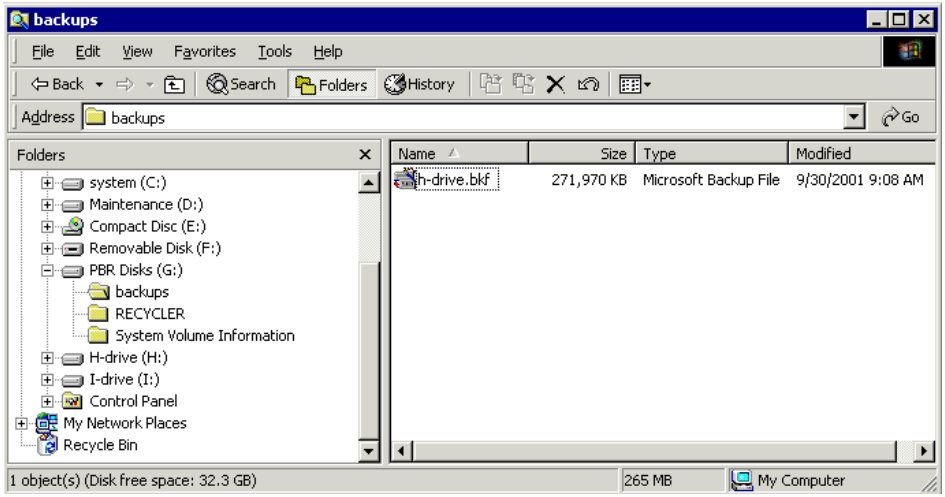

Figure 11-6 Screen showing the backup file created

8. Notice also that there is the “Removable Disk (F:)” on the NAS as shown in Figure 11-6. This is the temporary location to which the persistent image was stored by PSM for the NTBackup purpose. Drive F: was used because it is the next available drive.

   While the backup takes place, you should be able to access the removable disk. But once it has completed and IBMSNAP has finished, this removable disk is no longer accessible. To release the removable disk, you should logoff and logon to your system.

11.2.3 Creating a scheduled NT backup with IBMSNAP

If you want to schedule the creation of a persistent image and then back it up, you need to create two batch files, one that calls IBMSNAP at the scheduled time, and another one that calls NTBackup. Then you need to use the Windows AT command to create an entry on the Task Scheduler to call the IBMSNAP batch file at your specified time.

Here is an example:

1. Create a batch file that calls IBMSNAP (Figure 11-7).

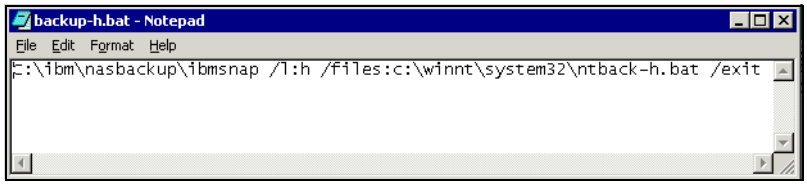

Figure 11-7 Sample batch file that calls IBMSNAP

2. Create the batch file called by IBMSNAP, which in turn calls NTBackup (Figure 11-8).

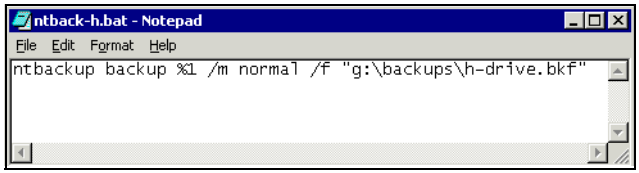

Figure 11-8 Sample batch file that calls NTBackup

3. On the NAS command prompt, type “at *time* “ibmsnap_batch_file”, for example:

```
at 11:00pm "c:\winnt\system32\backup-h.bat"
```

4. You should see that the job is added to Task Scheduler with a specific job ID. To confirm, at the command prompt, type `at` (Figure 11-9).

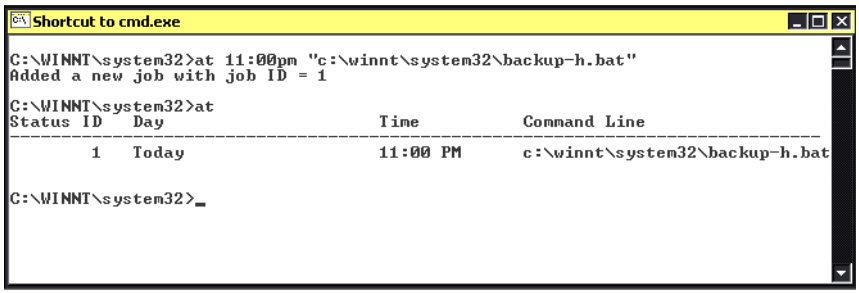

Figure 11-9 Screen showing the scheduled job for IBMSNAP

For more information on NTBackup, refer to the original Microsoft documentation.

Important: Please be aware that if you back up files directly from a PSM persistent image, the entire path name of each backed up file is preserved. As a result, when you restore such a file, it will attempt to restore to the persistent image and not to the original volume. *Restore Using NTBackup* should *only* be used in situations where standard backup (that is, not open file) is deemed sufficient, and you only want to back up a few selected files (as opposed to an entire volume). For all other backups, using NT Backup, the NAS Backup Assistant should be used.

11.2.4 Using IBMSNAP with TSM

The NAS200 and 300 are both preloaded with Tivoli Storage Manager Client Version 3.7. This means that you only need to have a TSM server on your network to have an enterprise storage management solution.

The steps involved in combining IBMSNAP with TSM are outlined below. It is assumed that the NAS machine is new and TSM has not been configured yet.

1. Create a node entry for the NAS on the TSM server.
2. Configure TSM Client on the NAS.

Note: If you already have a pre-configured option file (dsm.opt), just copy it to the TSM *baclient* directory on the NAS system.

3. Create a batch file that calls IBMSNAP.
4. Create a batch file called by IBMSNAP and in turn calls TSM.
5. Execute an automated TSM backup of the persistent image by running the batch file that calls IBMSNAP.

Creating a Node entry for the NAS on the TSM Server

These are the steps for creating a Node entry for the NAS on the TSM Server:

1. Open the Web Admin on your TSM Server.
2. Logon using the administrative account.
3. On the main screen, click **Operation View** (Figure 11-10).

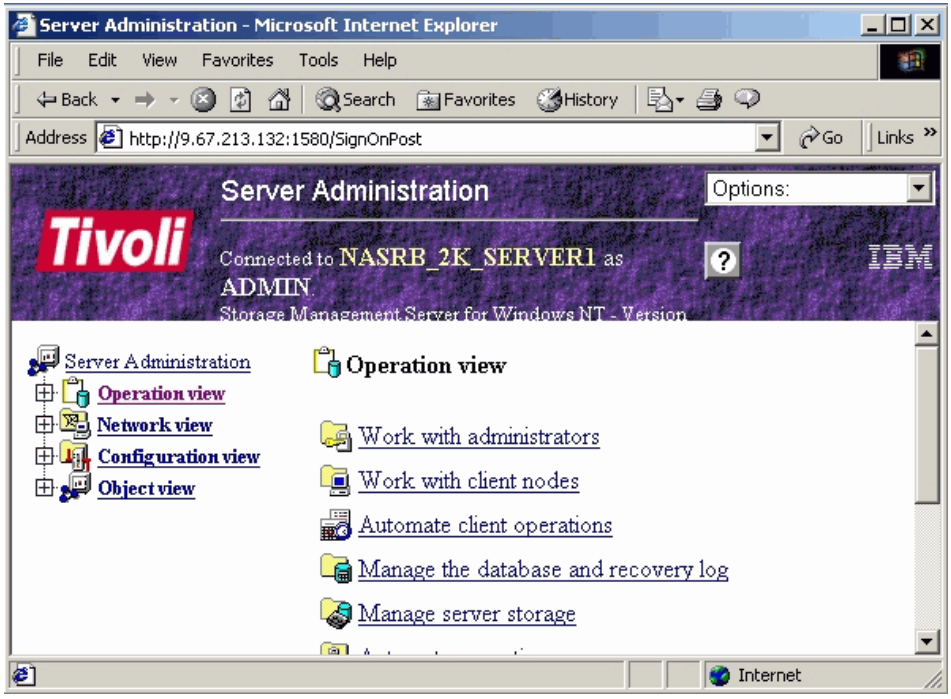

Figure 11-10 TSM Server Operation View

4. Under **Operation View**, select **Work with client nodes** (Figure 11-11).

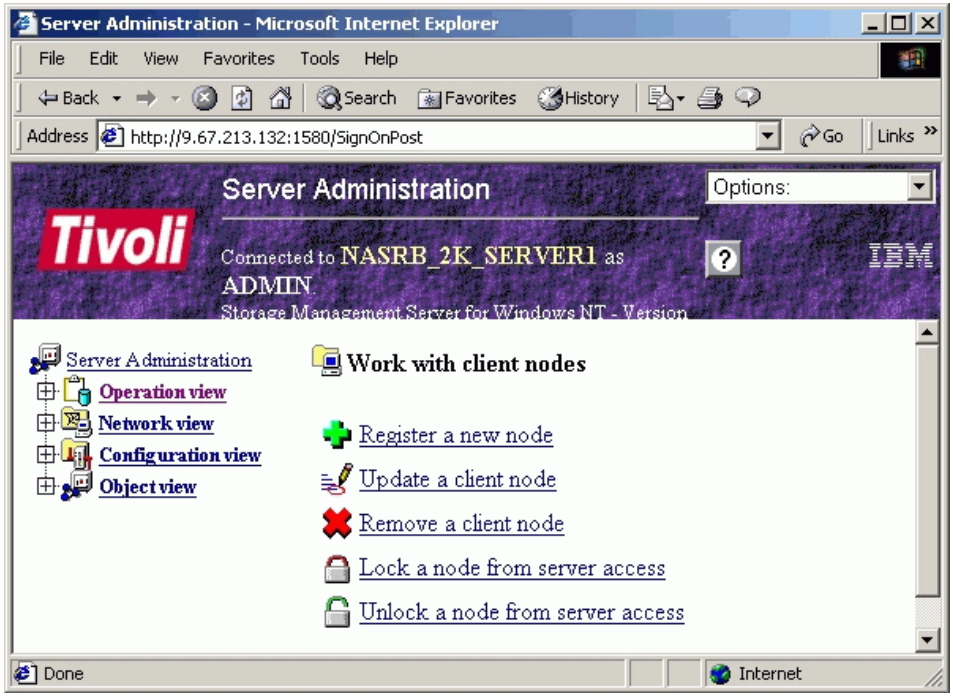

Figure 11-11 TSM Work with client nodes window

5. Fill in the necessary information for the new node as shown in Figure 11-12, and click **Finish**.

Note: For detailed information on TSM and each option under the window, Register a New Node, refer to the redbook, *Getting Started with Tivoli Storage Manager: Implementation Guide*, SG24-5416.

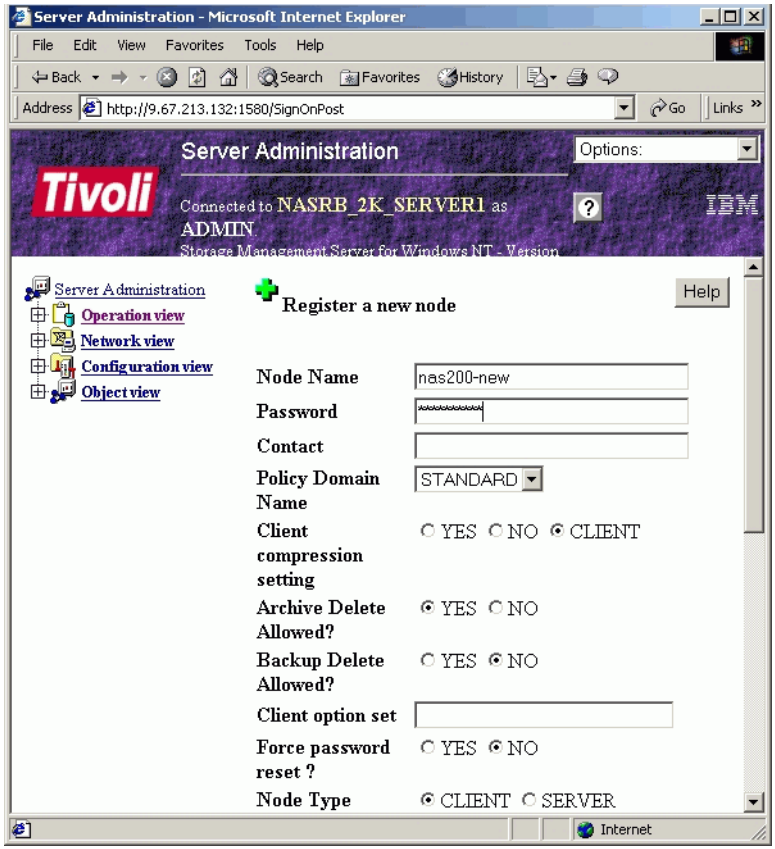

Figure 11-12 TSM server administration for the new node

Configuring a TSM Client on the NAS

These are the steps for configuring a TSM Client on the NAS:

1. Open your TSM Backup Client GUI.
2. On the TSM Client Configuration Wizard, select the options you want to configure by checking the appropriate boxes. Then click **Next** (Figure 11-13).

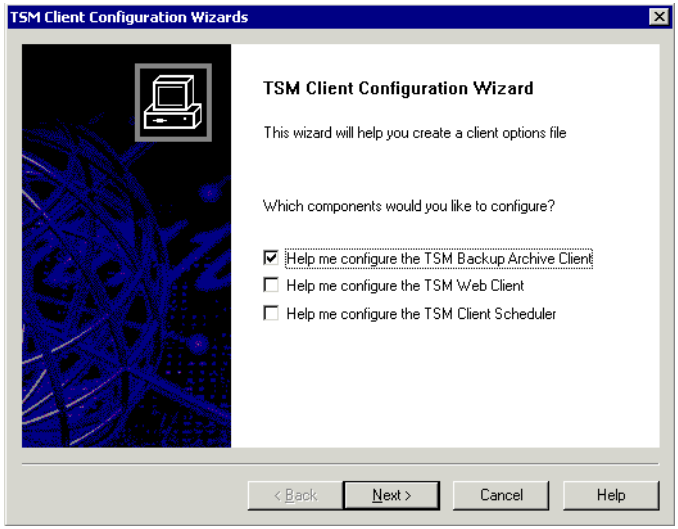

Figure 11-13 TSM Client Configuration Wizard

3. On the Option File Task window, select **Create a new options file,** and then click **Next** (Figure 11-14).

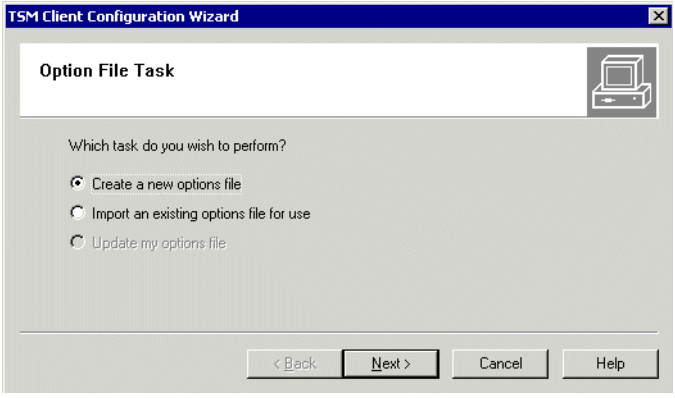

Figure 11-14 Option File Task window

4. On the TSM Authentication screen, input the TSM Node name for the NAS which was defined earlier on the TSM Server (refer to“Creating a Node entry for the NAS on the TSM Server” on page 440). See Figure 11-15.

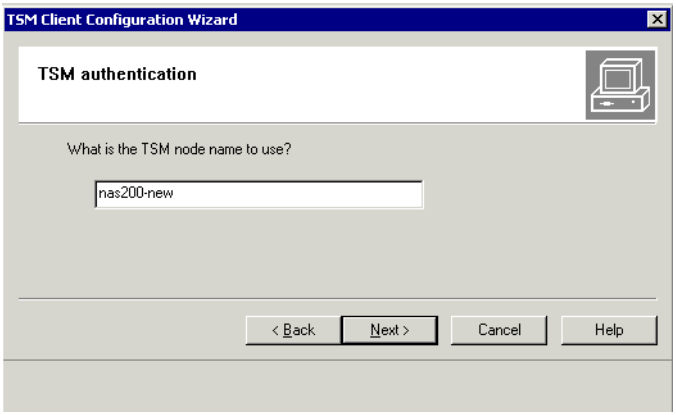

Figure 11-15 TSM client node name definition

5. On the TSM Client/Server Communications, select the appropriate protocol for your network. In this case, we select TCP/IP. Then click **Next** (Figure 11-16).

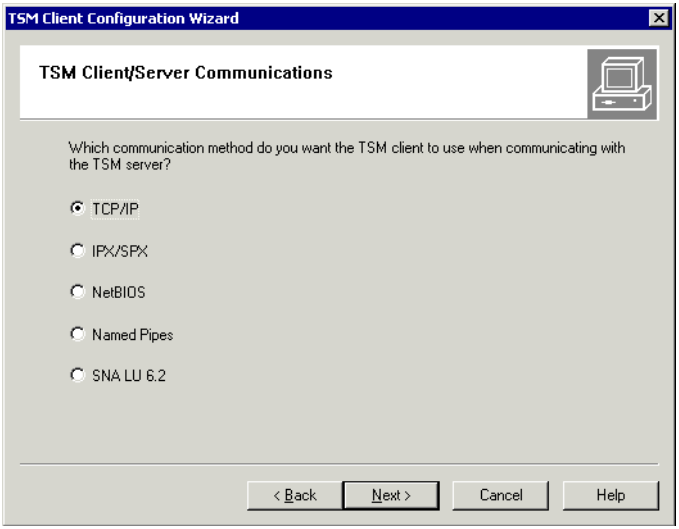

Figure 11-16 TSM client protocol selection

6. On the TCP/IP Parameters screen, type in the IP address of your TSM server. If you have not changed the default TCP/IP port on the TSM server (which is 1500), just accept the default. Otherwise, change it to match that on the TSM server. Then click **Next** (Figure 11-17).

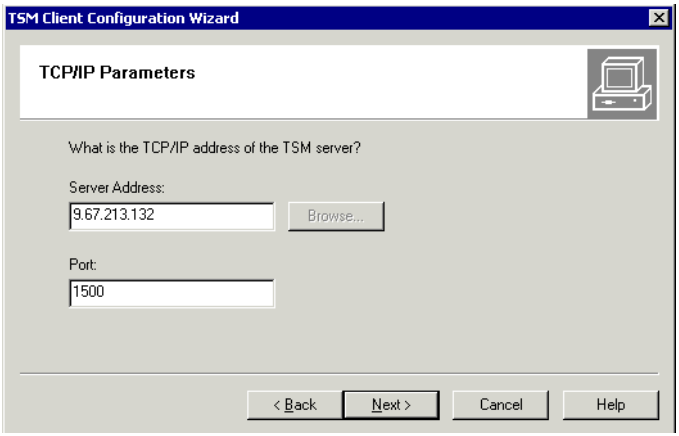

Figure 11-17 TSM client TCP/IP parameters window

7. On the *Domain and include/exclude lists,* click **Edit** to modify the default values. Otherwise, click **Next** (Figure 11-18).

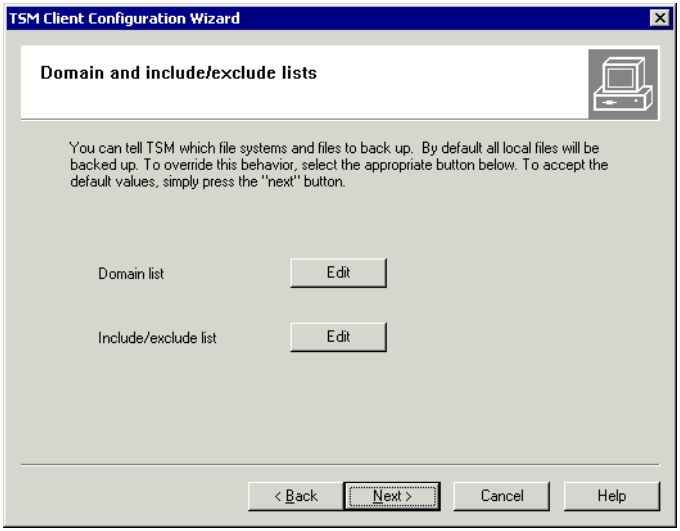

Figure 11-18 TSM client domain include/exclude lists

8. On the last screen of the configuration wizard, click **Finish** (Figure 11-19).

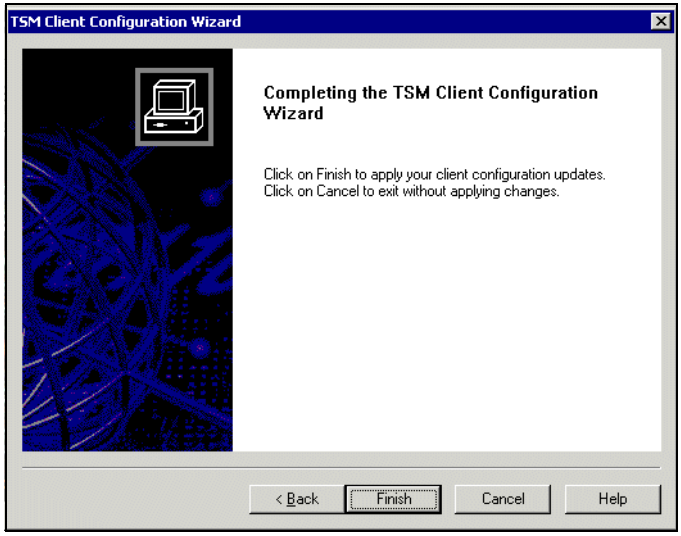

Figure 11-19 TSM client final window

Creating the batch file that calls IBMSNAP

The procedure for running the IBMSNAP is independent of the backup solution (whether NTBackup, TSM, or any other that supports the command line utility). However, for simplification of scheduling the TSM backup of the PSM image, we will create a batch file that executes the IBMSNAP command.

Here is the procedure:

1. On a command prompt, type Notepad *batch_file*". For example:

```
notepad tsmbackup-i.bat
```

2. Enter the **`IBMSNAP`** command with appropriate parameters. Refer to 11.2.1, "IBMSNAP utility" on page 433 for the available parameters and their usage. Save and close Notepad (Figure 11-20).

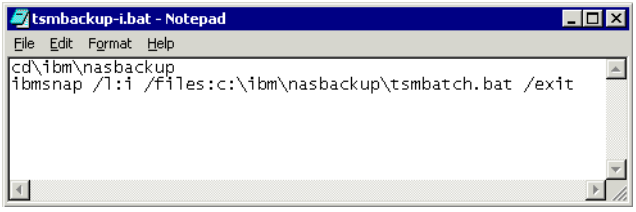

Figure 11-20 Sample IBMSNAP batch file

Creating a batch file that calls TSM

TSM has the **dsmc.exe** utility that enables the command line backup execution. However, to run properly, it has to be executed within the *\baclient* subdirectory. Therefore, the change directory entry has to be added on the batch file.

Unlike NTBackup, TSM uses volume labels to keep track of the drives on the machine. It is therefore important to have unique volume labels on the NAS. This reveals a problem, since IBMSNAP uses the volume label of the source drive for the volume label of the temporary Removable Disk that it creates. For example, suppose your NAS drives are the following:

Table 11-1 Sample NAS disk configuration

Drive Letter	Volume Label
C:	SYSTEM
D:	DATA
E:	CD-ROM
G:	TEST
H	USERS
I:	SQLDATA

While using IBMSNAP, you create a persistent image of D: whose volume label is *DATA*, a temporary Removable Disk will be created by IBMSNAP whose drive letter will be F: (since it is the first available letter), and will assign the volume label *DATA*. So, as far as TSM is concerned, you have two volumes (drives D: and F:) with the same label. If you back up F: (which the batch files do), the process will fail. It is very important, therefore, that you rename the temporary Removable Disk before the TSM backup runs. This can be done by adding a **label** command on your batch file.

Important: While using TSM and PSM together, you need to rename the Removable Disk generated by PSM before starting the TSM backup.

Here is the procedure for creating a TSM batch file:

1. On your command prompt, type “notepad *batch-file*”. For example:

   ```
   notepad tsmbatch.bat
   ```

2. Add the necessary entries for the batch file. Take note that on Figure 11-21, the “label” command is on the first line, to work around the problem of having duplicate volume labels.

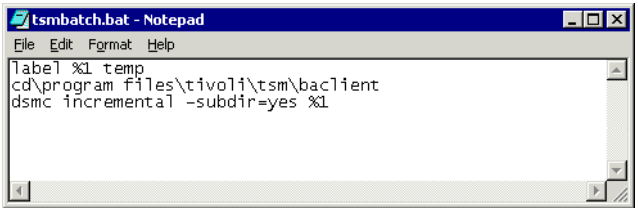

Figure 11-21 Sample batch file calling TSM

3. Remember also that for the %1 variable: IBMSNAP will supply the correct drive letter. Also, you may want to set the “PASSWORDACCESS” parameter on the *dsm.opt* file to “GENERATE”, so that you will not be required to specify the password on the **dsmc** command for security reasons.
4. For more details on the available parameters for “dsmc.exe” and their meaning, refer to the Online Information on your TSM Client Program.
5. Finish with Save and Exit.

Executing automated TSM backup of persistent image

Now that everything is prepared, we can run the automated TSM backup of persistent images. Here are the steps:

1. On a command prompt, execute the batch file that calls IBMSNAP (Figure 11-22). For example:

```
tsmbackup-i.bat
```

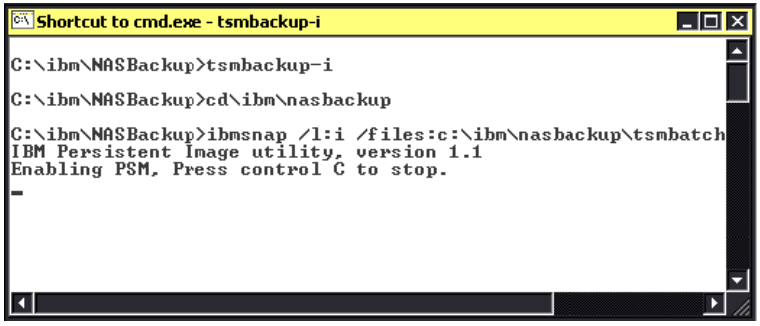

Figure 11-22 Screen right after running IBMSNAP batch file

2. PSM will then run and create a snapshot of the specified volume on the batch file. Also, the screen should show that the Removable Disk volume label has been changed to avoid duplication with the source volume (Figure 11-23).

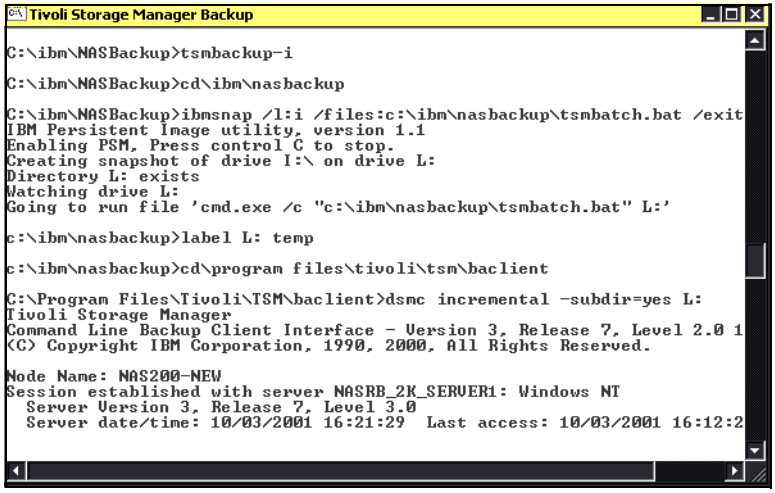

Figure 11-23 Commands of the batch file being executed

3. TSM backup will then proceed (Figure 11-24).

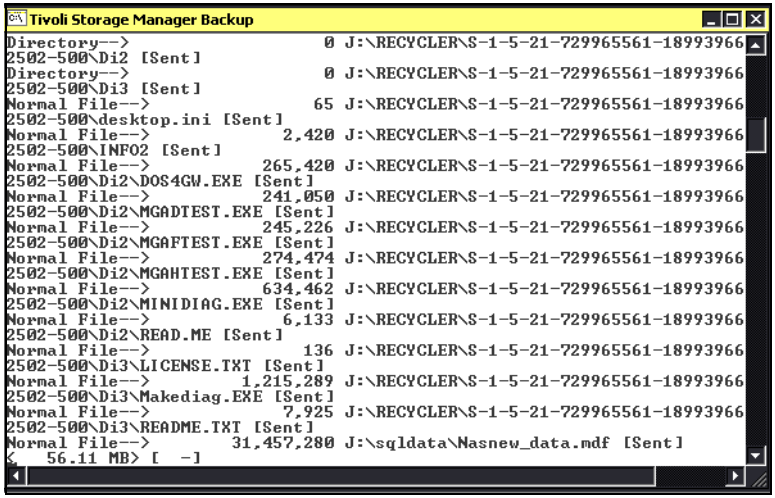

Figure 11-24 Files being backed up by TSM

4. Once the backup is done, the screen should look similar to the one shown in Figure 11-25.

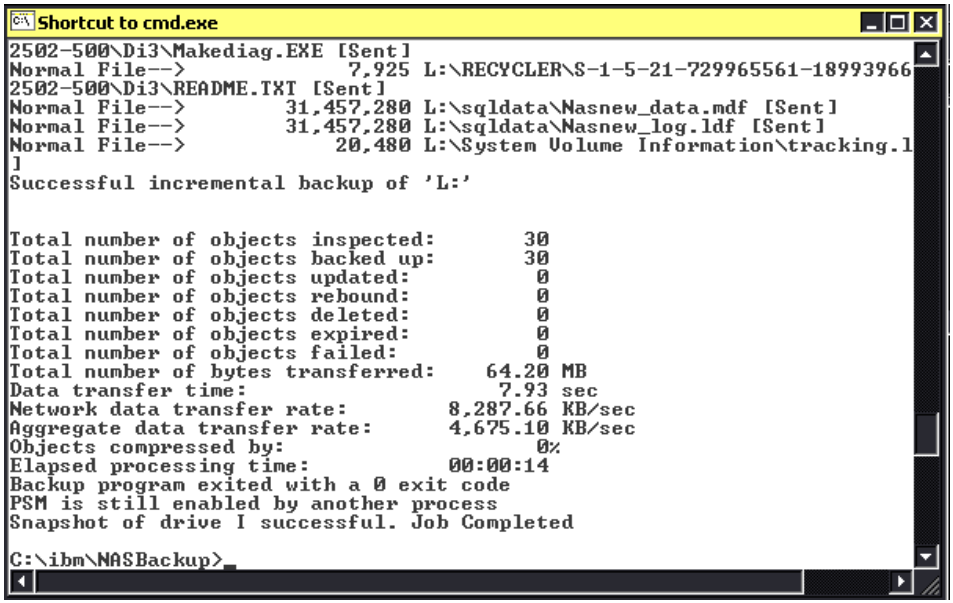

Figure 11-25 Screen showing PSM and TSM processes completing successfully

To verify that the backup has been created on the TSM server, perform the following steps:

1. Open the TSM Server Web Admin window.
2. Click **Object view**, select **Clients**, then select **File Spaces**.

3. Our sample is shown in Figure 11-26.

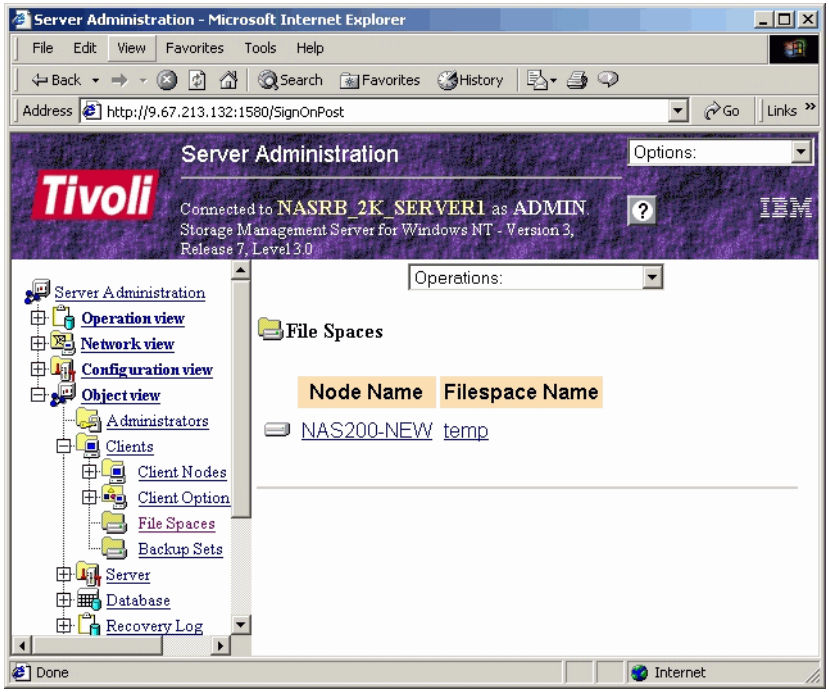

Figure 11-26 TSM Web Admin screen showing the backups available

4. Click the backup you just created to display the details, as shown in Figure 11-27.

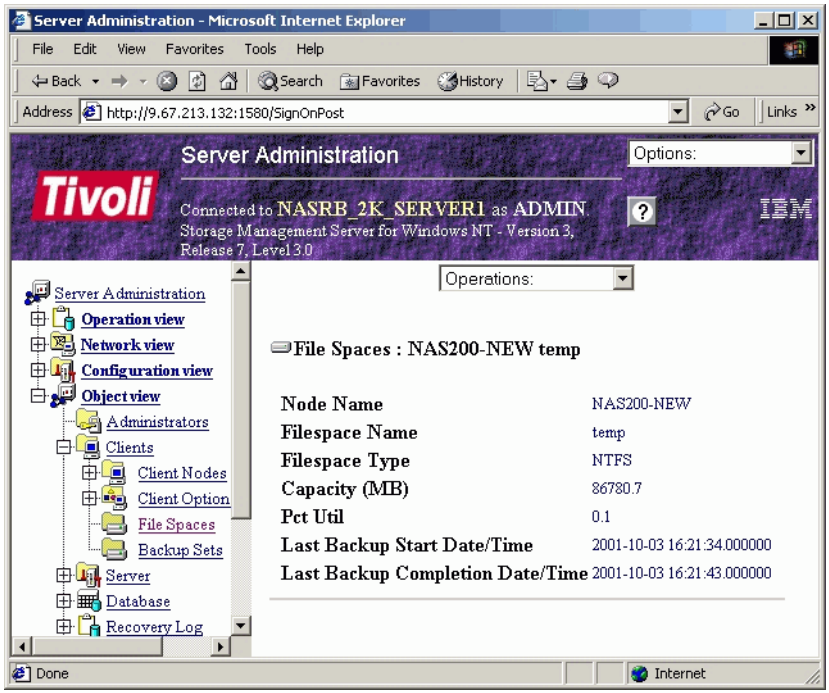

Figure 11-27 Screen showing the backup details

You now have the confirmation that the TSM backup from the NAS client was successful.

11.2.5 Creating a scheduled TSM backup using IBMSNAP

Just as when working with NTBackup, you need to use the **`at`** command utility of Windows 2000 to schedule the start of PSM and TSM. For example, you want to schedule a backup of the I: drive's persistent image at 5pm, you need to execute the following command on a command prompt:

```
at 5pm "c:\ibm\nasbackup\tsmbackup-i.bat"
```

In this command:

tsmbackup-i.bat = batch file that calls IBMSNAP

12

NAS recovery procedures

To demonstrate how easy it is to recover the NAS appliance, and to ensure that our system is in its pristine factory-shipped state, we will begin our work by re-initializing our system.

Re-initializing the unit is very simple, as it comes with CD-ROMs for this purpose. However, the unit also has a protection system to prevent it from being accidentally re-initialized if the first CD-ROM is left in the drive during a reboot. To circumvent this protection, you must use the Recovery Enablement Diskette and follow the procedure described in this chapter.

We discuss recovery techniques for each of the NAS models:

- Recovering the NAS 100
- Recovering the NAS 200
- Recovering the NAS 300

12.1 Recovering the NAS 100

We have created four scenarios to show how to recover your system or data after a failure of a hard drive or partition. We include detailed instructions to help you to handle such a situation.

12.1.1 NAS 100 boot behavior in case of an HDD failure

Regarding failover of the NAS 100, if the watchdog timer is enabled (this is the default, but it can be disabled in the BIOS), and then if there is a failure of the first hard drive, the system will retry two times (a total of three attempts) and if all fail, then it will try the second hard drive. A total of three attempts will be made with the second hard drive, and then the third hard drive will be used. After three more attempts, then the fourth hard drive will be used. Again, after three failures, and since there are no more hard drives to be used, it will not boot at all. Normally one of the hard drives will work.

There is a software component (a Windows service) that resets the watchdog timer. If this service does not start within five minutes of when the watchdog timer is turned on (right before entering Windows), the watchdog timer fires, the attempt is considered failed, and the system reboots and either retries with the same drive or switches to the next drive, as described above.

Note: A hardware failure of a drive will cause it to be skipped in this process. For example, if the first two drives are removed, the NAS 100 will immediately attempt to use the third hard drive (the third and fourth hard drives contain the backup OS).

We assume that you have a backup of your system state and data for a complete recovery of the NAS 100. It will be important in some scenarios to have a backup.

12.1.2 Recovery scenarios

These are the scenarios we considered:

1. Loss of one drive (primary or backup) occurs:
   - The primary and backup OS are in a RAID 1 (mirror) configuration.
   - The data configuration is RAID 5 (striping with parity).

2. NAS 100 does not boot from primary OS. There is a loss of primary OS (both primary drives):
   - The configuration data is lost:
     - The configuration data must be backed up in advance.
     - The data (RAID 5) can be recovered.
3. Loss of both primary or both backup drives (defective) occurs:
   - The configuration data and the data (RAID 5) are lost:
     - The configuration is maintained on the primary drives only.
     - The RAID 5 configuration for data now has only 2 drives.
4. No boot device is available. There is a loss of all 4 drives, primary, and backup. A full system rebuild is required.

Loss of one drive

The loss of a single primary OS hard drive is not disruptive and is redundant in two ways (Figure 12-1):

1. The OS is mirrored on the second drive with RAID 1 mirror.
2. The data is stored in a RAID 5 partition.

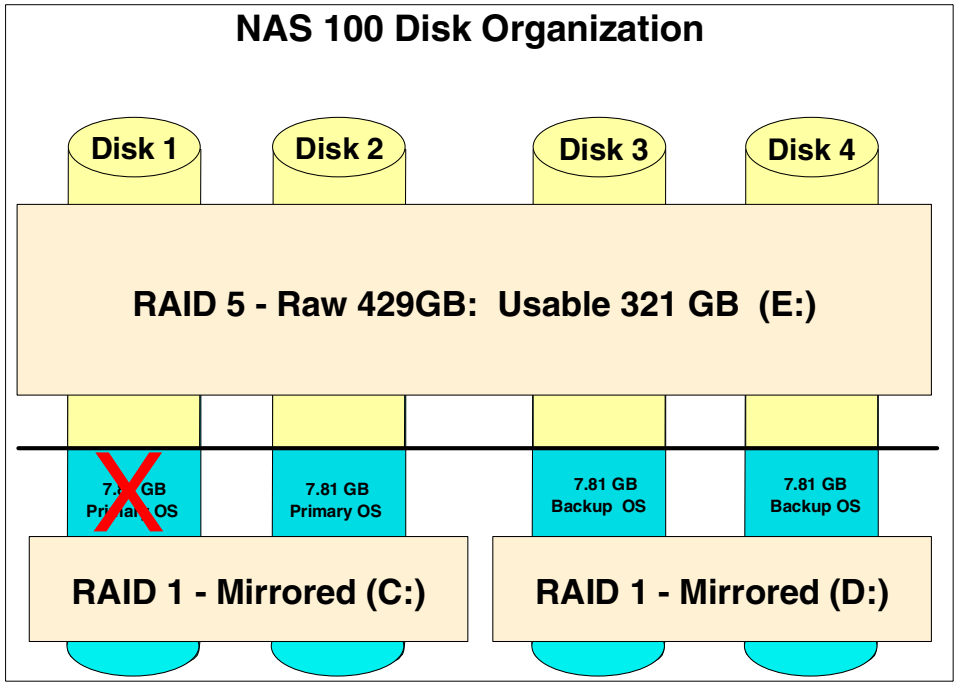

Figure 12-1 Loss of one OS hard drive

Panel lights will reflect the need to rebuild the RAID, as there is no redundancy in this state. The left light blinks green/amber 2 times a second.

1. Replace the defective drive with a new drive.

Note: The drive serial number is stored in the registry and cannot be reused in the same slot

2. Two options exist for rebuilding:
   a. Auto rebuild:
      - The drive can be inserted directly, and RAID is rebuilt.
   b. Manual rebuild:
      - Rebuild can be set for a certain time (for example, midnight).
      - The customer can initiate the rebuild process from the Web GUI by selecting the **Disks -> Admin Initiated** process (Figure 12-2).

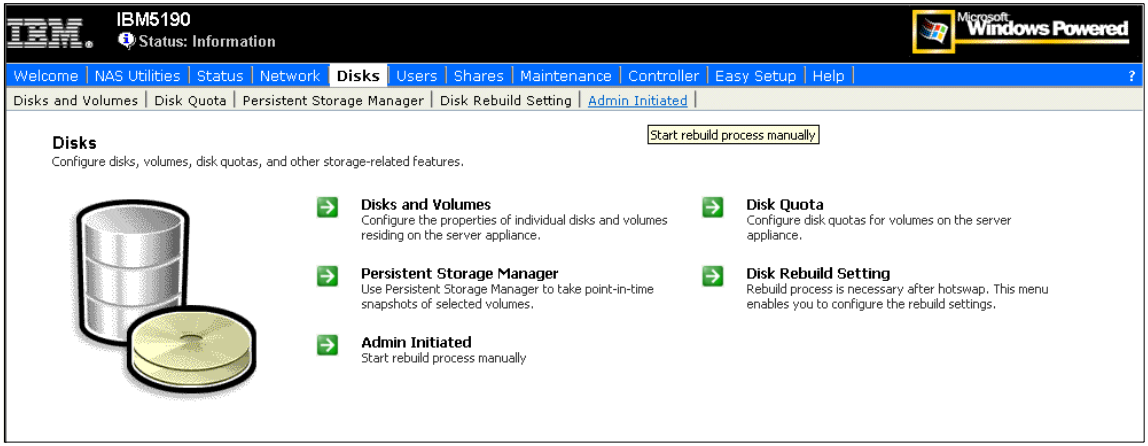

Figure 12-2 Admin initiated RAID rebuild tab

Loss of primary OS (both primary drives)

Recovery scenario 2 is the loss of OS and drives 1 and 2, and prevents the NAS 100 from booting. The drives are still good and the configuration data can be preserved (Figure 12-3):

1. The OS can be recovered from the backup OS on drives 3 and 4.
2. Your data is still available in the RAID 5 partition.
3. The configuration data can be recovered if the system state data has been backed up.

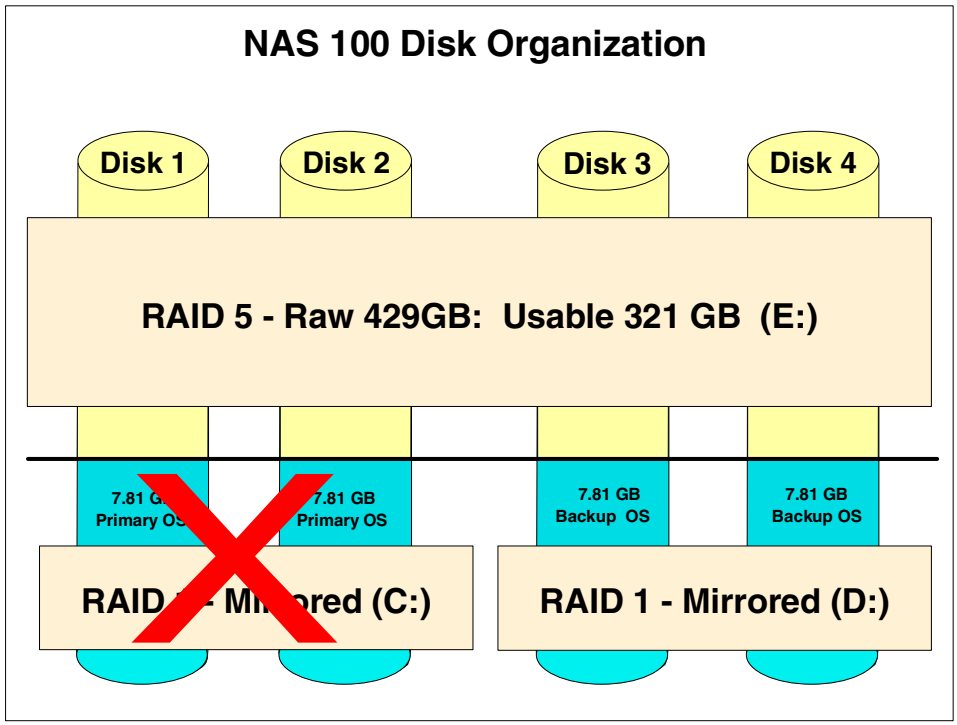

Figure 12-3 Loss of primary OS

The following steps must be performed to recover your NAS 100 appliance:

1. The NAS 100 appliance will boot from the backup OS with the default DHCP/static IP address (192.168.0.1) setting.
2. Access the device by using the device name or the static IP address via terminal services session.
3. Open a DOS window and access the d:\ibm\NAS100 directory.
4. Execute the command **recovery_OS**.
5. The system will rebuild the OS on the first two drives and automatically reboot.
6. Login to the NAS 100 using the default static IP address (192.168.0.1).
7. Access disk management and verify the RAID is correct and other drives are valid.
8. Restore the system state data using the NT Backup Utility.
9. Verify that the original users are in the configuration information.

Loss of both primary or both backup drives (defective)

This recovery scenario is similar to scenario 2, but now there are also two cylinders of the RAID 5 defective (Figure 12-4). This scenario is very unlikely, but we cover it just in case it occurs. In this case, all customer data will be lost and all configuration data will be lost.

1. The OS can be recovered from either from the primary OS or backup OS.
2. The data on the RAID 5 is lost.

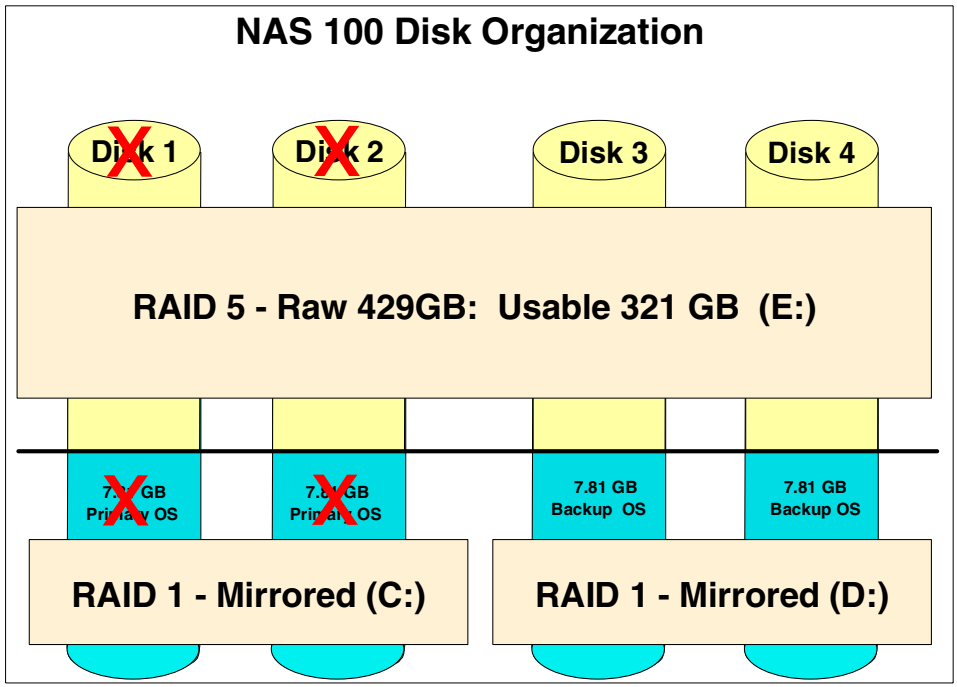

Figure 12-4 Loss of both primary or both backup drives

For the system recovery, proceed with the full system rebuild steps described next.

Full system rebuild

In the case of a complete system breakdown, you have to recover the whole system. You must follow this process in case of a failure of all four disks.

1. On both OS, inclusive, the configuration data is lost.
2. The data on the RAID 5 is lost.

Note: For the full system rebuild, a Recovery Kit is required from service. This kit includes the video adapter, USB keyboard, and a Master-Master drive.

To recover the whole system, follow these steps precisely:

1. Copy the pre-installed BASIC system image onto a blank drive. This procedure requires a video card, display, and USB keyboard.
   a. Remove the cover and install the video adapter in the available PCI slot and connect the display to the adapter.
   b. Connect the keyboard to the bottom USB port on the rear of the box.
   c. Create a recovery hard drive (requires a Master Recovery Hard Drive).
      i. Insert the Master Recovery Hard Drive in bay 1 of the NAS 100 and the drive that will become the Customer's "Recovery Hard Drive" in bay 2.
      ii. Press the button to clear the CMOS as described in "Clearing CMOS data" on page 576 and power the box on, continuing to hold the button until a beep is heard (approximately 10 seconds).
      iii. Press the **F1** key to enter CMOS setup.
         - Under the Advance tab, select **SuperIO Configuration** and disable the **watchdog timer**.
         - Under the Advance tab, select **PCIPnP Configuration** and enable **USB Function** and **Legacy USB Support**.
         - Save changes and exit.
   d. Boot from the **Master Recovery Drive**.
   e. Select **2. Make Another Recovery HDD Template** from the menu.
   f. At the **Enter destination drive number[0,1,2,3]** message, press **1**.
   g. A message will appear indicating that **all data on disk will be lost** and **press any key to continue**. Press Enter.
   h. At the message **Are you sure you want to ZAP Fixed Disk Drive 1(Y/N)?**, press **Y** then **Enter**. A message **Zap of Fixed Disk Drive Complete** and **Press any key to continue** will display. Press **Enter**
   i. DriveImage Pro will complete the creation of the recovery hard drive.
   j. The message **Make another[Y,N]?** will appear. Press **N** and the DOS prompt will appear.
   k. Reboot the NAS 100 (**Ctrl+Alt+Del**).
   l. Select **3. Copy Version Images to Another Recovery HDD (if exist)** from the menu.

m. When the message **Enter destination drive letter[C,D,E,F]?** appears, press **D** and the copy process will begin.

   This process will take approximately 4 minutes.

n. When the message **Copy to another drive[Y,N]?]** is displayed, press **N** and you will return to the DOS prompt.

o. At the DOS prompt enter **\bin\hdd0** and press <enter>.

p. Power down.

2. Create copies of the Primary and Backup OS images. This procedure requires a video card, display, and USB keyboard.

   a. Replace the **Master Recovery Hard Drive** in bay 1 with the Customer's “Recovery Hard Drive” created in the previous step and insert the remaining 3 drives in bays 2,3,4.

   b. Boot from the **Recovery Hard Drive**.

   c. Power on the NAS 100.

   d. Select **1. Start System Recovery Process** from the menu and restoring of the system image will begin (approximately 12 minutes).

   – Drive 1 will have the **basic system image for the primary OS** and drive 3 will have the **basic system image for the backup OS**.

   e. Power down the NAS 100.

   f. Remove the video adapter and keyboard.

3. Configure 4 hard drives on your NAS 100. A VGA card and keyboard are not required for this procedure. The NAS 100 will be used in headless mode (no video or keyboard).

   a. Press the button to clear the CMOS as described in “Clearing CMOS data” on page 576 and power the box on, continuing to hold the Clear CMOS button until a beep is heard (approximately 10 seconds).

      The system will boot on drive 1, the system configuration will start automatically and continue for approximately 1 hour. During this process the following will be performed:

      - Intermittent drive activity LEDs indicate that the build process is underway.
      - Setup for the OS will be done (SID, PID, user name and company, license agreement).
      - The disk will be converted from basic to dynamic.
      - The mirrors will be created.

   b. After approximately 1 hour, watch for all LEDs on all drives to blink green/amber on a 2-second interval indicating build is complete.

c. Log into the NAS 100 using either Terminal services Client or Internet Explorer Browser.

d. From a command prompt, change to “D:\IBM\NAS100” and issue the **`createFD`** command.

4. Create the RAID 5 volume and assign a drive letter and format the partition.
5. Create the RAID 5 drive and assign a drive letter and format the partition.
   a. Right-click **My Computer** and select **Manage**.
   b. Open the **Disk Management** tool.
   c. Right-click the **Unallocated Space** in Disk 0 and select **Create Volume** from the popup menu.
   d. In the Wizard, click **Next**.
   e. Select **RAID-5 volume** from the **Volume Type** menu and **Next**.
   f. Select Disks 1-3 under **All available dynamic disks** and click **Add**.
   g. Click **Next**.
   h. Click **Next**.
   i. Define a Volume label and Format type.
   j. Click **Next**.
   k. Click **Finish** to complete the create volume wizard.
   l. The RAID 5 will begin regenerating then formatting.
   m. This process will complete in approximately 5 hours if you issue a **Quick Format** and 12 hours if you issue a **Full Format**.
   n. After the regenerate/format of the RAIS 5 volume is complete, reboot the NAS 100 appliance.
6. Restore the system state data using the NT Backup Utility.
7. Verify that the original users are in the configuration information.

12.2 Recovering the NAS 200

The IBM NAS 200 can be recovered either using a recovery CD or using the maintenance disk partition. It is assumed you have a keyboard and monitor connected locally.

12.2.1 Using the Recovery CD

This method removes any configuration done on the NAS 200. Before you start the recovery procedure you must delete any persistent images currently present on your system.

To recover the preloaded image on your appliance:

1. Insert the diskette into the diskette drive and CD#1 in the CD-ROM, then restart the appliance. When the Recovery Enablement Diskette has completed loading and modifying your appliance start-up sequence, it will warn you about losing all data on the C drive. If you continue, the CD image process will start. Follow the directions on screen and replace the CDs when asked.
2. Remove the Recovery Enablement Diskette from the diskette drive and the last CD from the CD-ROM drive. When the process is finished, the system will reboot automatically.
3. The original manufacturing preload will be restored. Reinstall all software updates that has been installed on the Model before. Reconfigure shared folders, users, permissions and applications.
4. Apply service packs as needed.

Note: Refer to the backup book for different ways to backup and restore your NAS appliances.

12.2.2 Using the maintenance disk partition

To avoid having to create the users, shared folders, and permissions again, you can use tools such Persistent Storage Manager to create an operating system image. In order for this process to work, you need to meet the following prerequisites:

1. Previous backup of Operating System: See Section 11.2.1, “IBMSNAP utility” on page 433.
2. A diskette format with the following files: BOOT.INI, NTDETECT.COM, NTLDR

Follow these steps to recover the NAS 200 appliances:

1. Change the boot order: Press F1 when the NAS appliance starts and put in this order:
   a. First Startup Device [Diskette Drive 0]
   b. Second Startup Device [CD-ROM]
   c. Third Startup Device [Hard Disk 0]

2. In order to restore the backup, temporarily install Windows 2000 Server into the maintenance partition. The installation processes will show you this partition with 6.4 GB, when the installation is asking you to format the partition. Choose **Leave the current file system intact (No changes)**.
3. Boot on the maintenance partition using a Disk Boot and configure all necessary software (for example, TSM Client) to have the necessary connections to the backup server, if necessary.
4. From the maintenance partition, invoke **Restore To System Partition**.
5. Make sure no errors occurred during the restore processes.
6. Reboot the NAS 200.

12.3 Recovering the NAS 300

The restoration for a single NAS 300 node is the same as described for the NAS 200. We now show how to restore the NAS 300 in a clustered environment.

Reloading a NAS 300 after one node of a cluster fails

This section demonstrates how to reload one node in a clustered environment. We will call this system “new node” for demonstration purposes and are going to use the NAS setup navigator to accomplish the reload. There is an assumption of a local keyboard, mouse, and monitor.

1. Place the enablement diskette and CD#1 into the new node.
2. Reboot the new node.
3. Supply CDs to the system as required. The system should reboot when finished.

Note: Since the new system name will already exist in the cluster and either an NT domain or Active Directory, a few additional steps are required. For best practices we are going to give the new node the same name that it had prior, which is why these steps are necessary.

4. Log on to the working system (not the new node) and go into Cluster Administration. The node that is being reloaded will be represented by an icon with a circle and red slash. You need to evict that node from the cluster. Right-click that node —> **Evict Node**.

Note: The next step will need to be performed on the working system if it contains the Active Directory you are using. If the domain information is located elsewhere, you will need to connect to the appropriate system and remove the name of the new node before recreating.

5. From the same working system, go to **Start** —> **Programs** —> **Administrative Tools** —> **Active Directory Users and Computers**. Open up the computers folder and remove the failed system. This removes the system from Active Directory. (For an NT domain, use **Star** —> **Programs** —> **Administrative Tools** —> **Server Manager**, find the failed system in the list and delete). You are done on this node for now.
6. Log in to the new system with Administrator and password again. The NAS Setup Navigator will start.

Note: Perform step 9 if you have DNS loaded on the cluster. The new node would have had DNS on it.

7. Load DNS on the new server. Do this via **Start** —> **Settings** —> **Control Panel** —> **Add/Remove Programs** —> **Add/Remove Windows Components** —> **Networking Services**. Click the box to check **Networking Services** and then click the **Details** button. The Subcomponents of Networking Services box comes up. Click the box next to Domain Name System, then click **OK**.
8. In NAS Setup Navigator, on the Information and Setup Options screen, select the radio button that applies to the new node and Active Directory (if needed). After you have selected the appropriate radio buttons, click **Apply**. Clicking the **Apply** button will refresh the screen and bring you back to the top of the Information and Setup Options screen.
9. Click the **Forward** button to advance to the next screen.
10. Follow these steps and screens and configure the new node. When you are done with each screen, click the **Forward** button to advance.

Configure the new node with the appropriate information for:

- System Language
- Administrator Password
- Date and Time
- Network Identification
- Public LAN Settings
- Private LAN Settings

Note: At this point in the install of new node, you will either get the Joining a Domain screen or the Active Domain Controller Setup screen. This is determined by the Information and Options screen, where you selected either Joining a Domain or Setup Active Domain Controller.

11. This step requires you to either join a domain or set up Active Directory, depending on what radio buttons were selected in the beginning of the NAS Setup Navigator.
12. Now you will join the cluster. Follow the TotalStorage Cluster Configuration Wizard regarding the joining node.
13. Once you have successfully joined the cluster, verify in Cluster Administration that both nodes are present (the existing node and the new node).

At this time you have reloaded your IBM TotalStorage NAS appliance in a cluster.

13

Microsoft Active Directory and IBM TotalStorage NAS

In this chapter we discuss Microsoft Active Directory and its role in implementing an easy-to-manage and secure IBM TotalStorage NAS infrastructure. It is outside the scope of this book to fully cover the subject of Active Directory. Instead, we provide some basic information about Active Directory and its components, and point out considerations to keep in mind when using Active Directory in conjunction with IBM TotalStorage NAS devices:

- **Introduction to Active Directory (AD):** Benefits; cooperation between AD and Domain Name service (DNS); domain controllers
- **Active Directory architecture:** Objects, sites, domains, organizational units (OUs)
- **Interoperability:** Lightweight Directory Access Protocol (LDAP); Kerberos
- **Active Directory design:** Domain design (sites, trees, forests); OU design
- **Design considerations for IBM TotalStorage NAS:** Physical placement; administration delegation and authentication issues
- **Bringing it all together — best practice:** Leave room for growth

Note: If you are about to plan, design, and/or implement an Active Directory infrastructure, IBM Global Services will be glad to support you in this task. Contact your local IBM Global Services representative or an IBM business partner for more information.

13.1 Introduction to Active Directory (AD)

In this section we introduce Microsoft Active Directory and explain its benefits for today's IT environments. We discuss both Active Directory and the Domain Name Service (DNS) and present the new roles of Domain controllers.

We do not attempt to provide a full explanation of Active Directory that would enable you to plan and implement Active Directory in a global environment. Rather, we offer an overview and then point out considerations that are specific to Windows OS Powered NAS environments.

For more details on Active Directory, see the literature mentioned in 13.7, "Further reading" on page 494, or the massive amount of available literature on that subject. If you are already familiar with Active Directory, you might want to skip the next few sections and proceed directly to 13.5, "Design considerations for IBM TotalStorage NAS" on page 490.

13.1.1 Benefits

Active Directory provides administrators with a central service to organize network resources, to manage servers, clients, users, and applications. Primary features of Active Directory include:

- **Security:** Management of user authentication and access control are fully integrated in Active Directory. Authentication is centralized and access control can be ensured on all properties on every object. Clients can verify the identity of a server before transferring data by using mutual authentication.
- **Simple and flexible administration:** Active Directory allows for fine-grained delegation of administrative tasks, including access control on single attributes of an object. Delegation allows you to specify and distribute administrative tasks in your IT staff without providing administrators with more rights then necessary to do their job. At the same time the use of wizards and new interfaces as well as policy-based administration makes managing today's complex IT environments easier.
- **Scalability:** The use of DNS as a name resolution mechanism and storing directory information in (maybe multiple) domains allows for scalability. The Active Directory can store and access millions of objects.
- **Availability:** Active Directory uses multi master replication to distribute crucial data between multiple domain controllers, thus ensuring a highly available directory service. This covers queries as well as updates in the directory.
- **Open:** By using the Lightweight Directory Access Protocol (LDAP) as a query and update protocol as well as Kerberos as a logon and authentication protocol, the interoperability with non-Windows platforms is ensured (see 13.3, "Interoperability" on page 475).

13.1.2 Active Directory and Domain Name Service (DNS)

A directory is a hierarchical structure that stores information about objects. Objects include resources like servers, shared volumes, shared printers, user and computer accounts. As an example, a directory might store a user account object, including attributes like name, password, e-mail address, work phone and work address.

The Domain Name Service (DNS) — as well as the Active Directory (AD) — is a namespace. In a namespace, a name can be resolved to an address. Looking at the Internet, the process of name resolution requires DNS servers to resolve user-friendly addresses (such as `http://www.ibm.com`) to Internet Protocol (IP) addresses.

DNS consists of a hierarchy of domains, making the entire Internet to one namespace. Although DNS and Active Directory typically share the same name for a domain (for example, `ibm.com`) they manage different namespaces. This way, a name resolution for an object in a directory cannot be achieved by using DNS, as both systems store different data. DNS stores address zones and resource records, while Active Directory stores domains and domain objects.

Each computer in a DNS domain is identified by a fully qualified domain name (FQDN) like `nas200.testlab.itso.ibm.com`; every Windows 2000 domain has a DNS name (for example, `ibm.com`).

In summary, we want to point out that:

- DNS is a name resolution service.
- Active directory is a directory service.

Note: Windows 2000 Active Directory and DNS are tightly integrated. You cannot use Active Directory without DNS. If you have a WINS environment — now you need DNS.

Windows 2000 uses Dynamic DNS (DDNS). If you want to use existing DNS servers together with Active Directory, you have to make sure that these support the Windows 2000 DDNS requirements, including dynamic updates and SRV records. Due to a proprietary secure update mechanism, typically only Windows 2000 DDNS servers support secure DNS updates.

13.1.3 Domain controllers

Transforming a Windows 2000 Server into a domain controller (DC) is achieved by installing the Active Directory code onto it. One domain controller can host exactly one domain. Domain controllers store domain-wide directory data and manage interactions between users and domains, such as logon authentication and directory searches.

In Windows NT Server 4.0, domain controllers were either Primary Domain Controllers (PDCs) or Backup Domain Controllers (BDCs), with PDCs being superior to BDCs. Changes to the domain data had to be committed to the PDC and then mirrored to the BDCs. With the introduction of multi-master replication in Windows 2000, changes to any domain controller are mirrored to all other domain controllers in the same domain.

Although all domain controllers should be equal, there is still one (or more than one) domain controller that is more equal. It is the Global Catalog (GC) server. The Global Catalog server is created automatically on the first domain controller in a forest (see “Forests” on page 473). Each forest must have at least one GC server.

The Global catalog server has two main functions:

- **Logon:** The Global Catalog server is required to complete the logon authentication process for users and computers. In fact, every object authenticating to Active Directory must authenticate with the Global Catalog server.
- **Querying:** Queries concerning directory objects are resolved by the Global Catalog server. This involves all query requests from users, administrators and programs.

Therefore, you might want to place at least one Global Catalog server per site if you have a multi-site environment to reduce inter-site logon traffic and enhance logon and query response time. Keep in mind that adding Global Catalog servers to your network may increase network traffic, as Active Directory information is replicated to all participating Global Catalog servers.

Tip: Use at least one Global Catalog server per site to reduce WAN traffic. Use multiple Global Catalog servers for redundancy.

13.2 Active Directory architecture

In this section we overview the key architectural components of Active Directory.

13.2.1 Objects, schemas, and publishing

A network consist of objects. Therefore the Active Directory stores information as attributes to objects.

To describe objects that share a similar set of attributes schemas were introduced. A schema is a description of an object class. The schema defines the attributes an object must or can have. Every object in the Active Directory belongs to a schema. Schema definitions are also stored as objects.

The process of making objects available in the Active Directory is called publishing. For example, publishing a shared folder object in the directory lets users quickly query for and access the folder.

13.2.2 Sites

Computers in a site are well-connected, for example, by a high-performance local area network (LAN). Sites are separated when the network connection between them is slow. This is typically true for wide area networks (WAN).

The introduction of sites in Windows 2000 eases the management of the IT environment by delegating tasks and grouping resources as well as implements routines to effectively operate distributed IT environments. As an example, configuration of sites ensures that clients can logon to a domain controller in their local network and that replication between domain controllers is optimized for minimum bandwidth consumption.

Be aware that site structure and domain structure are separate.

Tip: Use sites to reduce WAN traffic if you have multiple locations that are connected with slow WAN links.

13.2.3 Domains

The Active Directory is organized in domains. Every domain has at least one domain controller, a directory can consist of one or more domains. As domains in Windows 2000 can consist of millions of objects, the need to introduce multiple domains to manage large amounts of objects (for example, users or computers) is no longer existent — as opposed to NT4.0 with its limitation to about 40.000 objects per domain.

On the other hand, multiple domains might be required for business organization or administrative purposes. For example, if there is an autonomous business unit that needs to be managed separately and that should be in its own security boundary (so that no administrator of that unit can manage IT assets from the rest of the company) then introducing an additional domain is typically the easiest way to achieve this.

Note: Domains can match sites — but typically domains do *not* match sites.

Trees

A tree is a set of one or more domains. So when installing your first domain controller you also create a tree. The domains in a tree form a contiguous namespace, as shown in Figure 13-1.

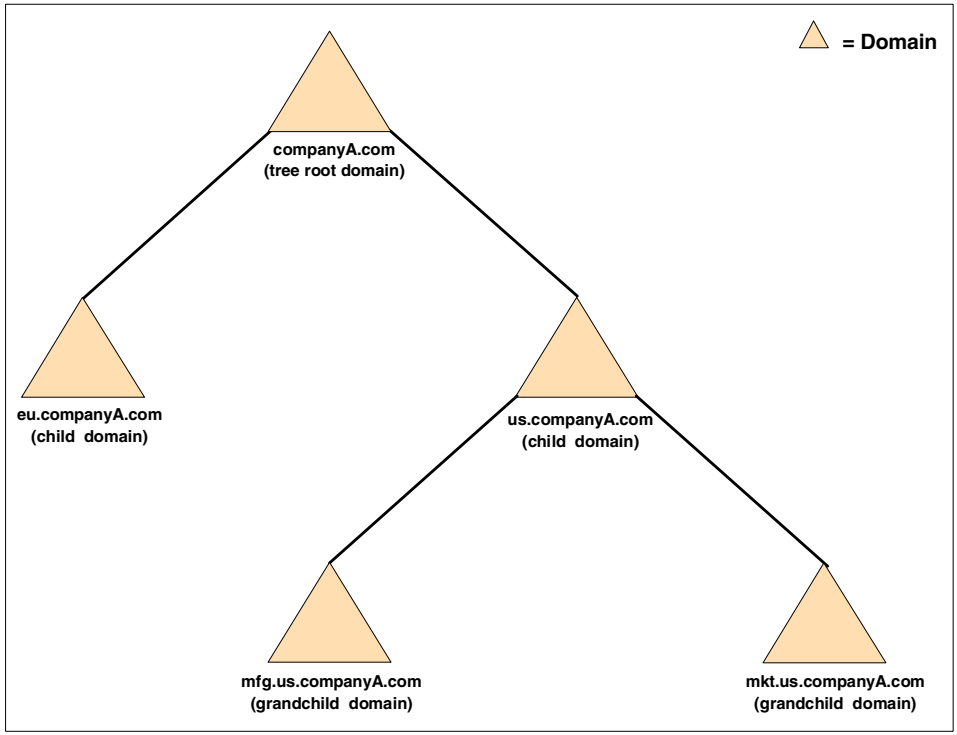

Figure 13-1 Domain tree

The first domain in a tree is the root domain; other domains are child domains. Domains forming a tree are linked by two-way transitive trust relationships as described in “Trusts” on page 474.

Forests

A forest consists of multiple domains that store parts of the forest configuration — the trees (Figure 13-2). The forest distributes the data that needs to be stored across the different domains to optimize access time and reduce network load. All domains in a forest host a copy of the forest configuration and the schema containers. Each domain database stores information about its local environment. All domains in a forest are linked by two-way transitive trusts.

Tip: If you think you need more than one forest — think about it again. In most situations, one forest should be fulfilling your needs.

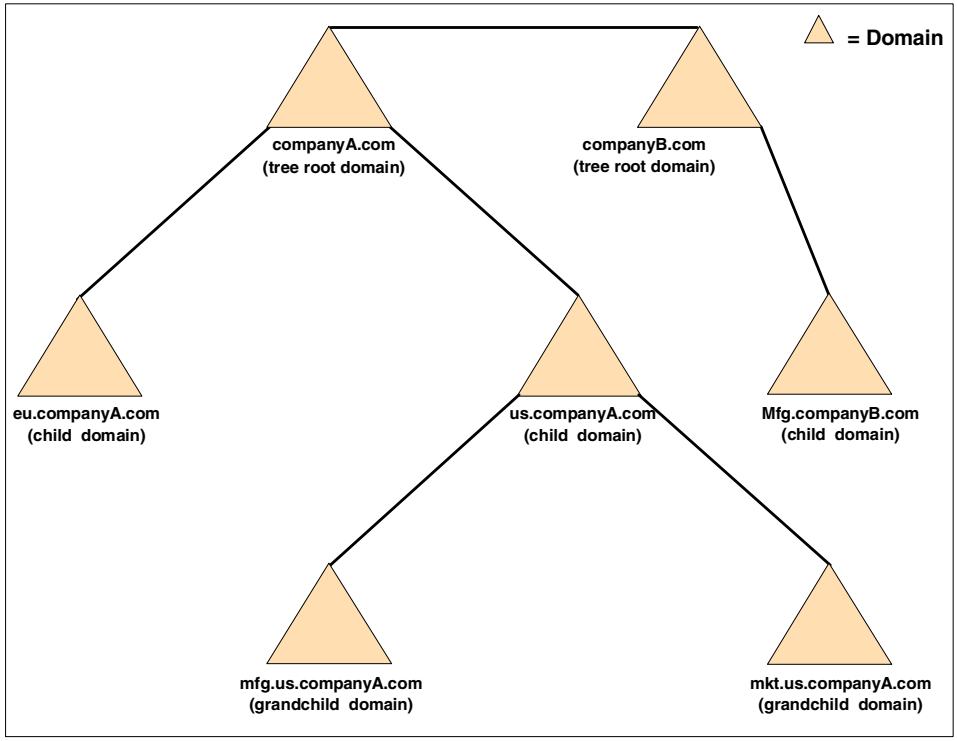

Figure 13-2 Forest with multiple trees

Users see a single directory. So when they query for a resource it is found, wherever in the forest it is located, even across domain and tree boundaries.

There are few reasons to ever introduce more than one forest in any IT environment. Further discussion would be beyond the scope of this paper, so refer to 13.7, “Further reading” on page 494.

Trusts

A trust is a relationship between two domains that allows users that are authenticated by one domain to access data in another domain. In Windows 2000, account authentication is established using two-way transitive relationships.

The term *two-way* describes that authentication is trusted between two domains in both ways. For example, domain A trusts domain B, and domain B trusts domain A (A —> B and B —> A).

The term *transitive* defines a trust that reaches beyond the initial scope of a trust. For example, if domain A trusts domain B, and domain B trusts domain C, then domain A automatically trusts domain C (A —> B and B —> C means A —> C).

There are two additional special trusts:

Shortcut trusts When the authentication path is too long (for example, user in domain A needs to access resource in a domain E with four trusts in between) the administrator can define a shortcut trust between domain A and domain E. This is a one-way transitive trust.

External trusts External trusts are trusts to domains in another Windows 2000 forest or to a non-Windows 2000 domain (for example, Windows NT or Kerberos realm). This enables to authenticate users to external domains. External trusts are one-way transitive trusts.

Note: In Windows NT 4.0 trust relationships are one-way and non-transitive. When upgrading a Windows NT domain to a Windows 2000 domain the existing one-way trust relationships stay this way.

13.2.4 Organizational Units (OUs)

Windows 2000 introduces the Organizational Unit (OU) as a directory object that can be used to store computers, users, groups, shared folders, printers and other OUs. Organizational units are typically used to delegate administrative authority for sets of directory objects. For example, if you want to define a set of administrative accounts to manage your NAS servers, then you would create an OU for that purpose (more on OU design in 13.4.2, “Organizational Unit design” on page 485).

OUs should reflect your IT administration organization, not an end-user or business organization perspective.

Tip: OUs can help you to organize your administrative work. Make use of them.

13.3 Interoperability

In this section we provide an overview of the interoperability features of Active Directory. These include the Lightweight Directory Access Protocol (LDAP) as well as Kerberos. Additionally various sets of Application Programming Interfaces (APIs) are published and can be used by independent software vendors (ISVs) or customers to inter operate with the Active Directory.

Note: If you plan to use heterogeneous systems in your environment, especially when you want to use LDAP and/or Kerberos along with Active Directory, make sure that the vendors support this — and test the interoperability before finalizing the design.

13.3.1 Lightweight Directory Access Protocol

The Lightweight Directory Access Protocol (LDAP) provides access to various directory services, including Active Directory. In fact, Active Directory depends on LDAP for all changes and queries to the Active Directory. Windows 2000 supports LDAP Versions 2 and 3.

13.3.2 Kerberos

Active Directory supports the Kerberos protocol (RFC 1510) for authentication of users and computers, as well as applications. Non-Windows 2000 clients can authenticate in a Windows 2000 domain when they use Kerberos (RFC 1510).

On the other hand, it is possible to provide Windows 2000 Professional clients with a single sign-on to the local workstation and a Kerberos realm at the same time, if such a Kerberos realm already exists. Also, it is possible to establish external trusts between Active Directory and Kerberos realms, so that Kerberos clients can authenticate in a Kerberos realm and use resources in a Windows 2000 domain.

13.4 Active Directory design

In this section, we provide you with some hints and tips on Active Directory design in the context of planning and implementing an IBM TotalStorage NAS environment. However, we do not explain how to design, plan, and implement Active Directory in your IT systems or those of your customers. A wide range of literature is available to assist you in this task.

13.4.1 Domain design

A domain is a security boundary. It allows common security definitions to be applied to all objects contained within it. Domains are also high-level administration units although OUs within domains provide greater granularity when using Windows 2000.

With Windows NT 4.0 domain design, the forest and OU objects were not available to allow for transitive trusts and delegation of administration. Windows 2000 allows for full trust between all domains within a forest, and the forest provides for intra-forest object resolution via the Global Catalog server mechanism.

A common reason for the implementation of multiple domains when using Windows NT 4.0 was that security delegation was not granular enough within a domain. With Windows 2000, OUs are employed for the primary purpose of delegation of administration.

There are three basic domain design choices available when designing an Active Directory:

- The Geographical model
- The Political model
- The Functional model

Which model should you employ? These three models are rarely implemented in their basic state. Instead, based on the structure of the enterprise being modeled, these three basic models are combined to generate new hybrid models. After a brief examination of the three basic models, variations will be examined.

When designing an Active Directory, there are few things to keep in mind:

- First, the design should be modeled on the enterprises political and geographic structure. It makes little sense to build a Windows 2000 AD based on a company's functional structure if the company only performs one major function such as accounting.
- Second, administration requirements will determine the most appropriate basic model to employ. Fully centralized administration may eliminate the need for multiple domains while decentralized administration may drive multiple domain designs.
- Third, replication will drive the number of trees and domains within the forest. A large, distributed network spanning many countries will generally require multiple domains for several reasons. Many countries have laws governing data processing information and intra-country traversal of such data. Also it is common for data security requirements to vary from country to country.
- Finally, replication is an important consideration when designing AD for large networks. A domain provides an important replication partition that reduces the amount of information shared forest-wide therefore countries or areas with inadequate bandwidth will often implement a domain to reduce the need for bandwidth.

Throughout this chapter, various examples are offered to enhance the reader's understanding. These models do not illustrate the implementation of false root domains, but do not preclude the implementation of false root domains either. In addition, OU structures are discussed and modeled as part of overall domain modeling. We further explore OU design in 13.4.2, "Organizational Unit design" on page 485.

The Geographical model

The Geographical model allows the modeling of Active directory to occur based on geographical considerations. This model is very common and allows for centralized or decentralized administration. One of the greatest benefits of the Geographical model is that it is immune to political reorganizations.

Figure 13-3 illustrates a geographical design that incorporates countries as domains and cities or sites within these countries as OUs. This design allows each country to determine its own security requirements such as password length, auditing levels, and event log size. Further, this model allows replication limited to national boundaries.

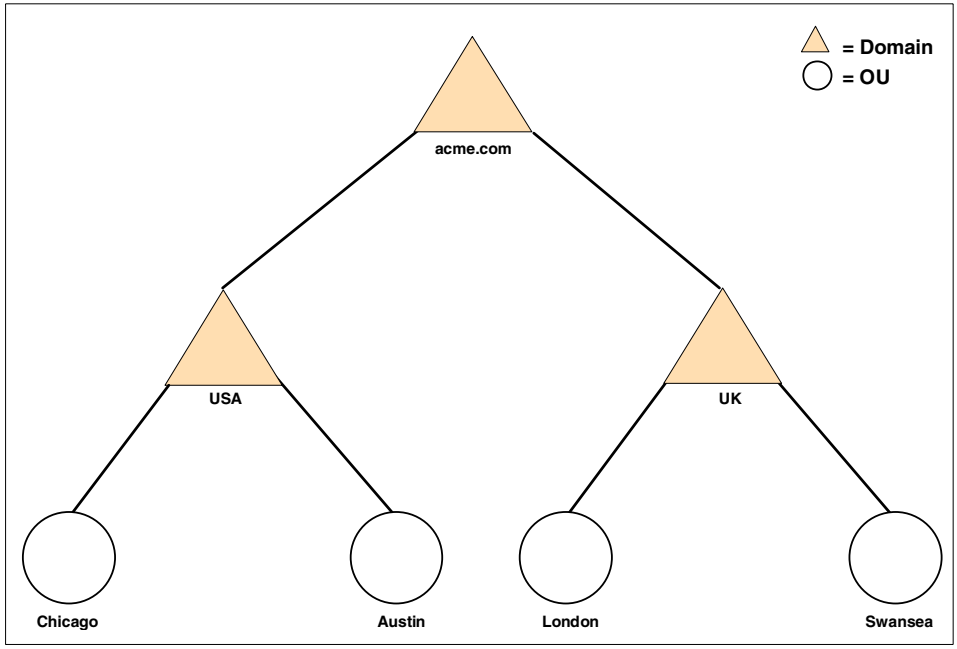

Figure 13-3 Geographical model

Another variation would be to implement regional trees and define specific countries as child domains within these trees. This solution would be useful in situations where there are company locations within a region that don't warrant their own country-based domain. The regional domain can provide a convenient place to house these exceptions as OUs. See Figure 13-4.

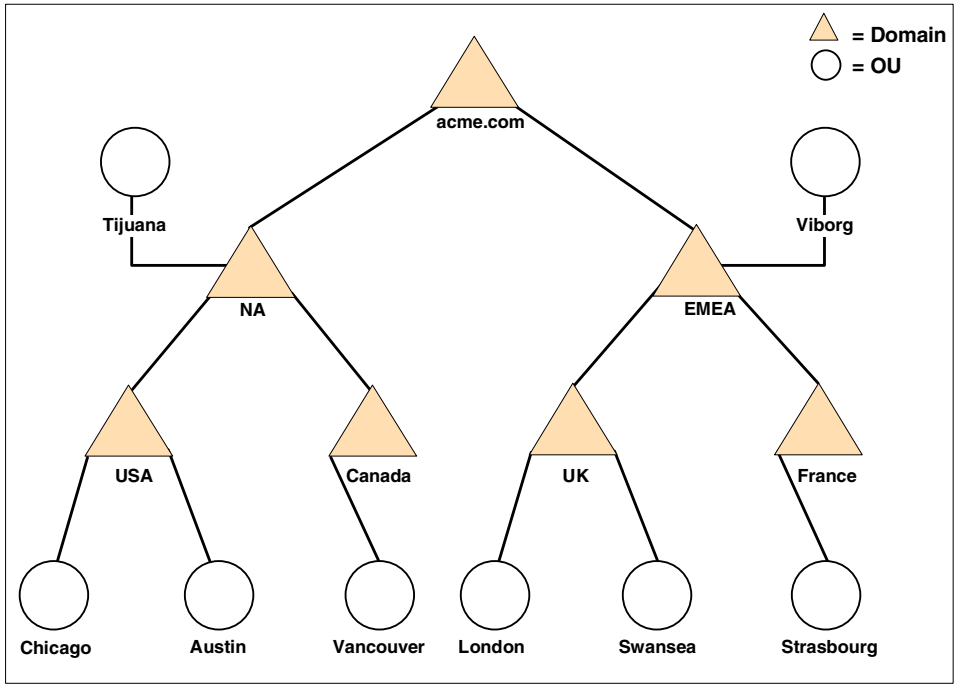

Figure 13-4 Geographical model — regional domains

Figure 13-5 illustrates the implementation of only one domain within the forest and sites are implemented as OUs. This model works very well in enterprises that have fast WAN connectivity and a moderate number of sites. This model is easy to administer and allows for a unified security policy. Administration can be handled in a centralized or decentralized manner.

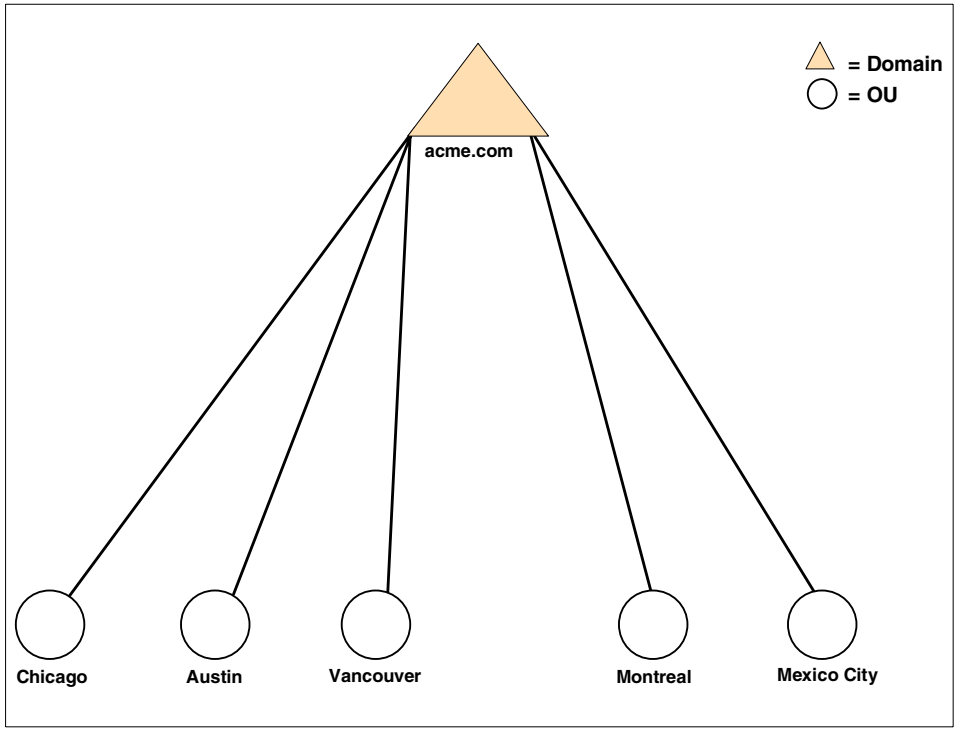

Figure 13-5 Geographical model — cities as OUs

The Political model

The Political model mirrors the political structure of the enterprise. The unfortunate side effect of this is that any political reorganization will require the redesign and implementation of parts or all of the AD. This model also requires very well connected sites. Administration under this model can be either centralized or decentralized. Figure 13-6 on page 481 shows an example.

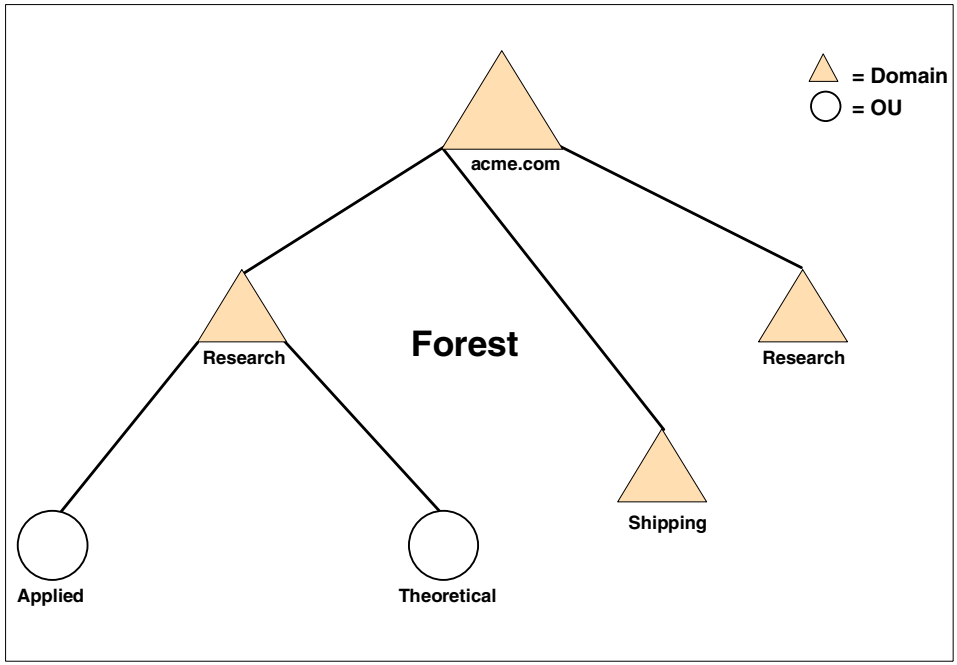

Figure 13-6 Political model

The Functional model

The Functional model is based on the business functions within the enterprise and does not consider political or geographic considerations. This model works well for very small organizations but this model does not scale well. Due to the small size of the organization being modeled, this design will likely not consist of more than one domain. Figure 13-7 on page 482 provides an illustration of the Functional model.

Tip: In most situations, you will find it fitting to use the functional model for the lowest OU levels only, while using different models for the top layer(s). Also see “Other hybrid models” on page 484.

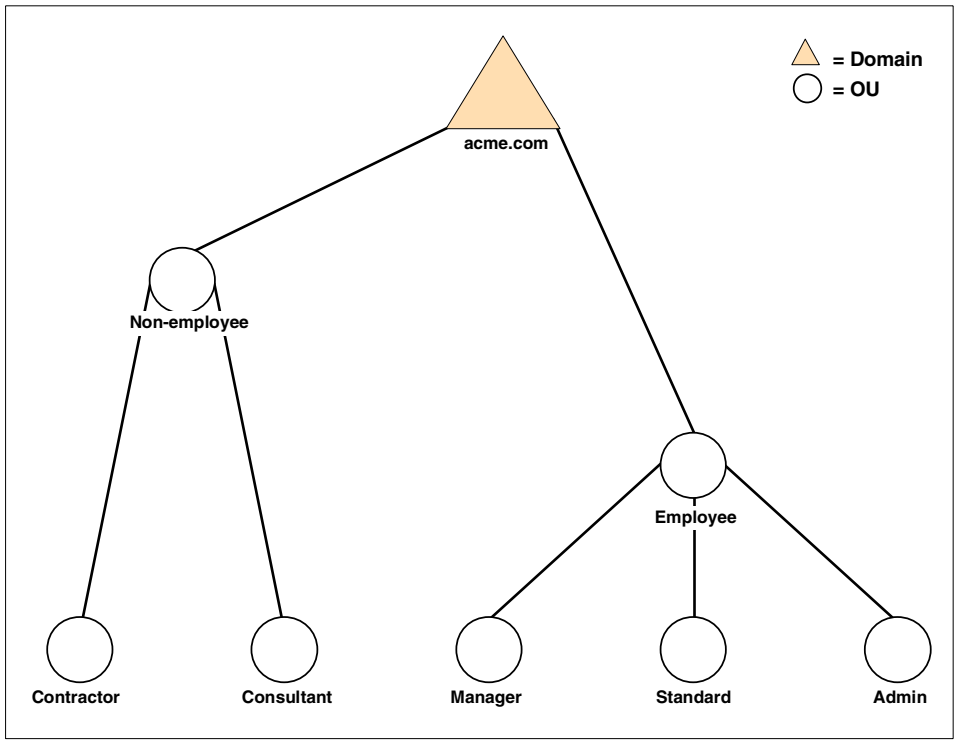

Figure 13-7 Functional model

Hybrid models

By taking the best features of the three basic models and combining them to meet the customer's business requirements, one of two hybrid solutions (the geo-political model or the political-geographic model) is generally arrived at. It is important to note that although the functional aspect of these two hybrid models is not discussed in this chapter, they do provide functional implementation characteristics. These characteristics will be examined in detail in 13.4.2, “Organizational Unit design” on page 485.

Geo-Political model

The Geo-Political model combines components of both the Geographical and Political models at the higher levels of the design to provide a flexible and resilient design that resists being affected by political reorganizations and scales very well. In addition, the Geo-Political model supports somewhat centralized administration (domain-level) but has the flexibility to allow each site or city the ability to administer its own resources.

As Figure 13-8 shows, the Geo-Political model implements geographic-based domains or regional domains at the highest level, then city and/or sites at the next level. Following cities/sites, political units are modeled. Finally, although functional units are not depicted, they would be found below the politically based OUs in the figure.

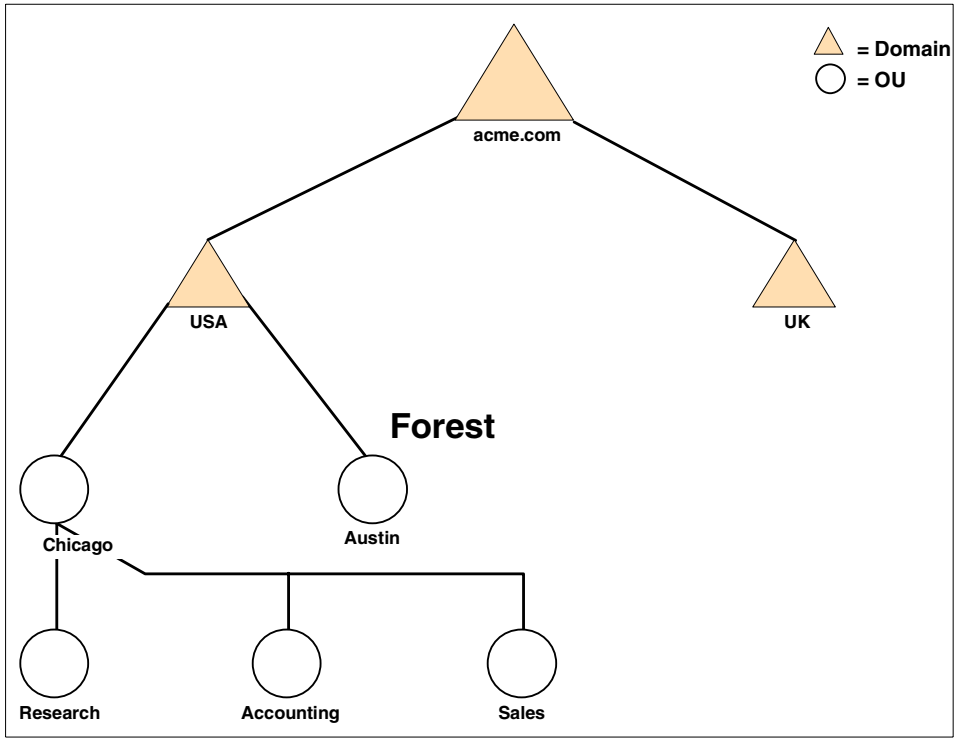

Figure 13-8 Geo-political model

Political-Geographic model

The Political-Geographic model lends itself well to organizations with physically distributed business units and strong divisional boundaries. This model tends to drive decentralized administration although administration could be centralized within each line-of-business. The first and possibly second levels of this design are based on the political boundaries within the enterprise. Next, the geographic sites are defined. This model requires a well developed WAN with plenty of bandwidth. Figure 13-9 shows an example of the political-geographical model.

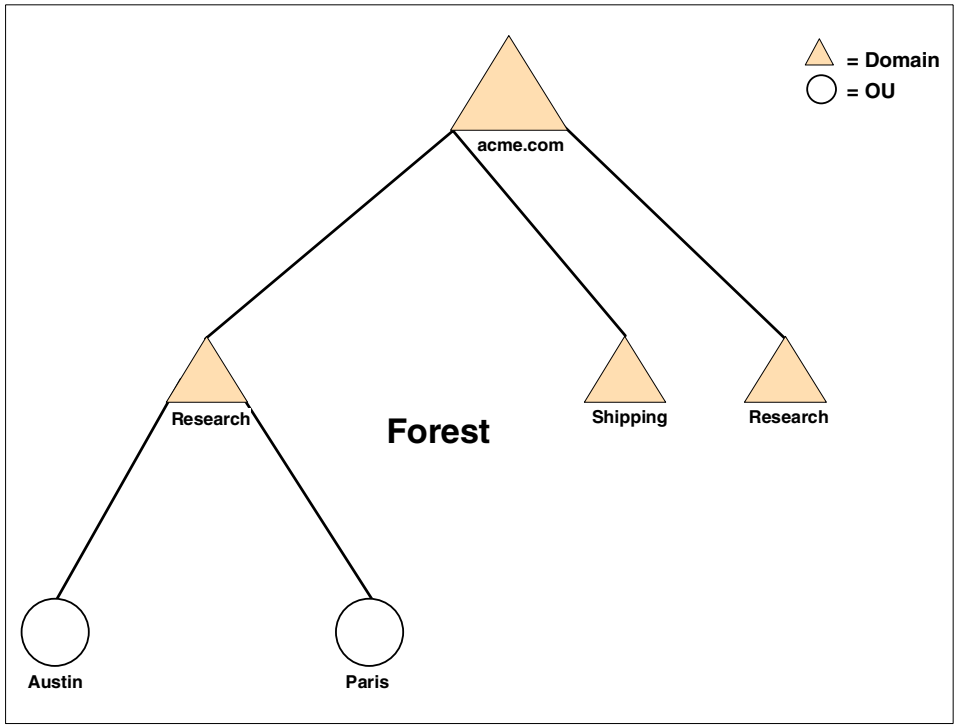

Figure 13-9 Political-Geographic model

Other hybrid models

Other hybrid designs are possible, and you may need to explore various combinations of the three basic models to arrive at a design that works for a specific environment. Keep in mind that, in general, functional characteristics should only be modeled at the lowest levels of the design and as OUs. Another important thing to note is that it is generally best to model cities or sites, but not both. Modeling both cities and sites will complicate the design process unnecessarily. Finally, it is possible to replace domains with OUs, and OUs with domains, with the understanding that OUs can't contain domains. Therefore, three or more levels of domains are possible, and OUs may be nested several layers deep.

13.4.2 Organizational Unit design

An Organizational Unit is simply a container for other objects such as users, groups, computers, and other OUs. OUs are used to simplify the administration tasks associated with Active Directory such as user and group administration, application of group policy objects (GPOs) and software distribution. The default containers and OUs provided by Windows 2000 upon installation are very primitive. In fact, only one OU is defined — Domain Controllers. The rest of the objects are containers, which were also present in Windows NT 4.0.

The three basic models that are employed for forest design also apply to OU design — political, geographic, and functional. Unlike the forest design process, the OU design process involves two stages. First, the determination must be made whether geographic or political OUs are necessary at the top of the directory. In larger organizations with geographically dispersed sites or many divisions, it is likely that one or both OU-types will be implemented. The second step is to define the functional OUs within each geographic or political OU.

Note: Remember — the main reason to have OUs should be to help you to ease and streamline administration of your IT environment. Don't design and create OUs just to have a nice OU design — and keep it simple.

High-level OU definition

A successful OU design will drive the need for fewer domains and will support the customer's administrative model. Further, a correctly designed directory will "feel" right in the way that the enterprise relates to it and the way it naturally fits within existing administrative processes of the enterprise. At the top of the directory, political and geographic OU implementations usually result unless the customer is very small and has only one or two sites. A very simple AD implementation might resemble Figure 13-10.

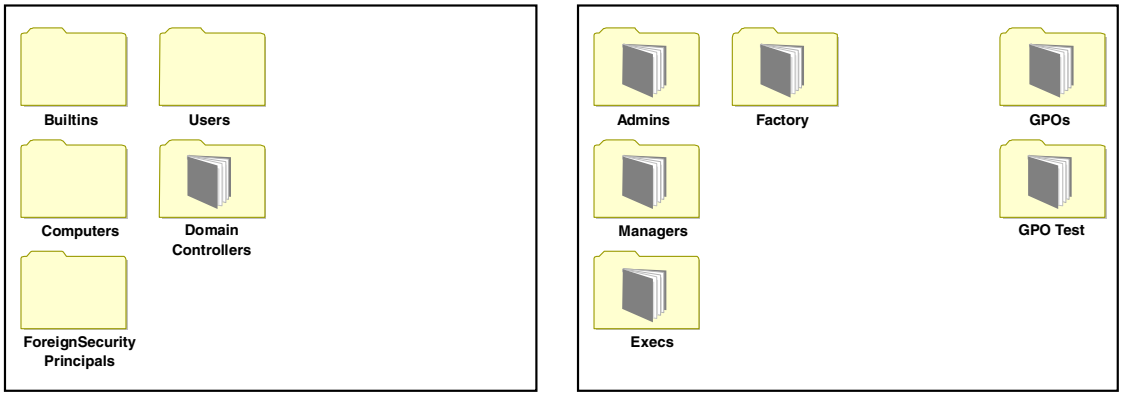

Figure 13-10 Containers and one OU

The geographic and political OU implementation can be implemented within a forest design based on either geography or political domains.

Geographic OU implementations

The implementation of geographic OUs occurs when it is cost prohibitive to implement domains. This usually occurs below the country level. By implementing city, state or regional OUs, administration has the potential to be decentralized to those points.

The most common geographic implementation is based on cities. Figure 13-11 illustrates the OU structure for an enterprise that spans three cities.

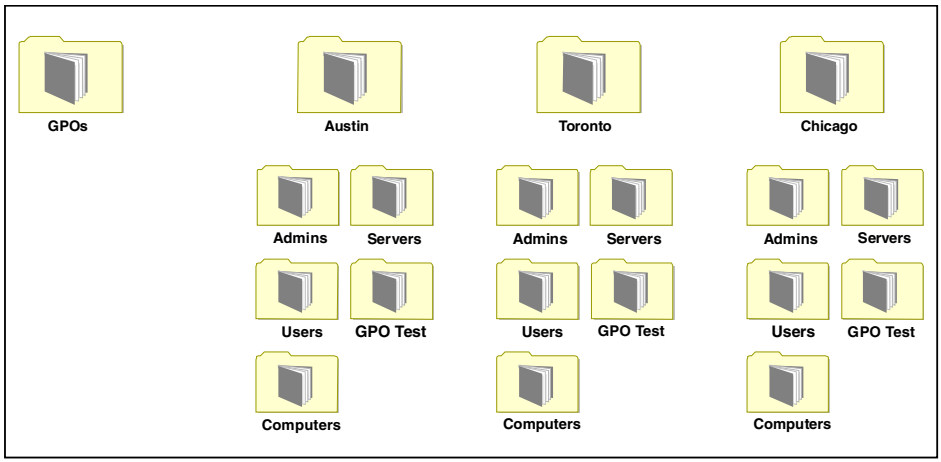

Figure 13-11 OU Structure — three cities

However, there is nothing restricting a more defined regional or state geographic implementation such as is depicted in Figure 13-12. This design allows administration to be centralized within each state if necessary. It also allows another level of group policy (group policy object — GPO) differentiation.

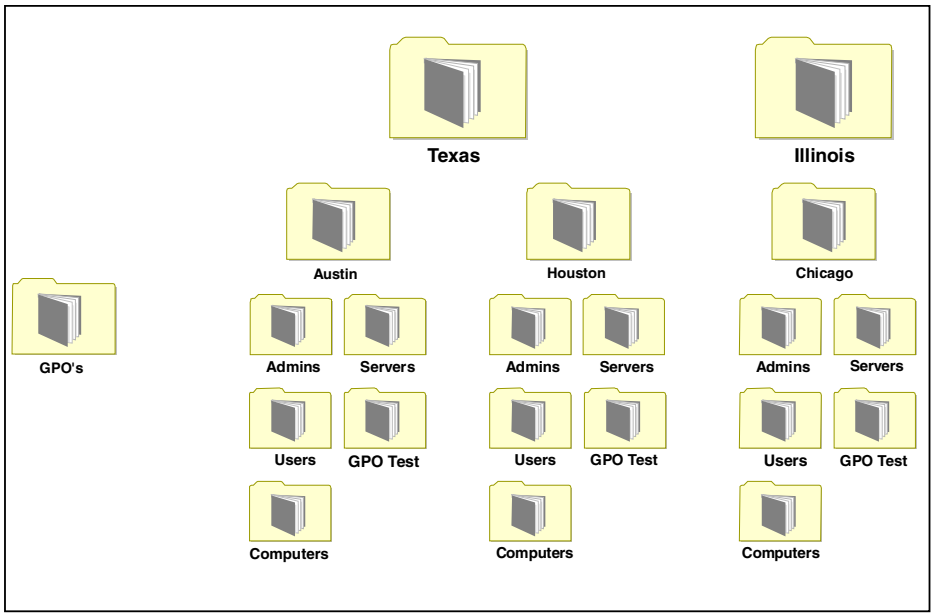

Figure 13-12 OU structure — states and cities

Political OU implementations

Political OU design generally follows the same pattern as geographic OU design except that political elements determine the top OU hierarchy. As with geographic OU designs, no more than two levels are generally implemented. This type of design would be useful within an enterprise that prides itself on decentralization. Perhaps funding for a centralized IT department has never been attempted or is simply not feasible. Figure 13-13 exemplifies a typical political implementation.

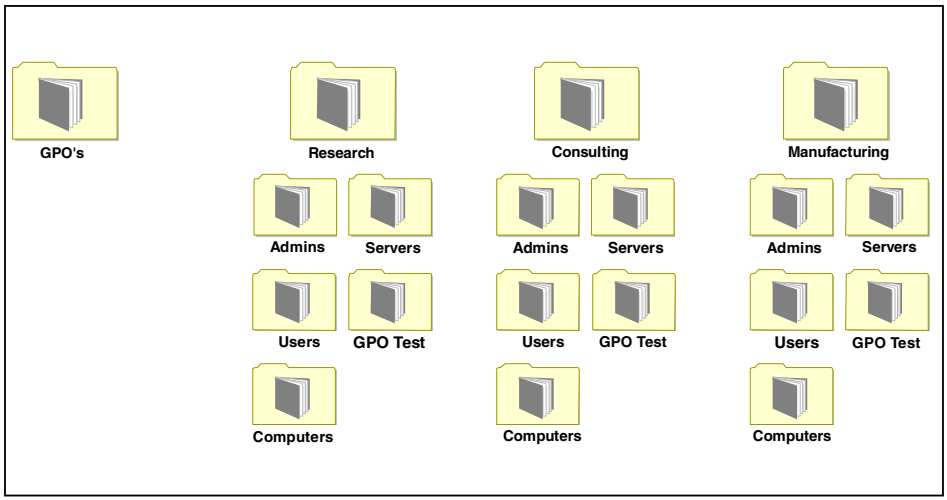

Figure 13-13 Political implementation

This model allows administration to occur along the lines of business versus geography. It is possible to add geographic OUs to further decentralize control but the overhead for definition can be large.

Functional OU definition

The lowest level of OU definition is generally comprised of functional OUs. The higher levels have provided political and/or geographic partitioning but delegation of administration and group policy object assignment are still complicated without further divisions. There is no wrong approach to the definition of functional OUs, only wrong approaches for individual environments.

Most functional OU designs have several things in common including the definition of a computers, users, and servers OU. In addition, Exchange 2000 will drive the need for other OUs to hold items such as distribution lists and contacts. Such a design allows the application of unique group policies to computers, users, and servers. This can be very useful if different security requirements exist for these objects. In addition, GPO-based software distribution will depend heavily on the use of OUs and groups to determine what computers and/or users should receive a specific application.

Figure 13-14 depicts an example OU design that incorporates geography and functional elements.

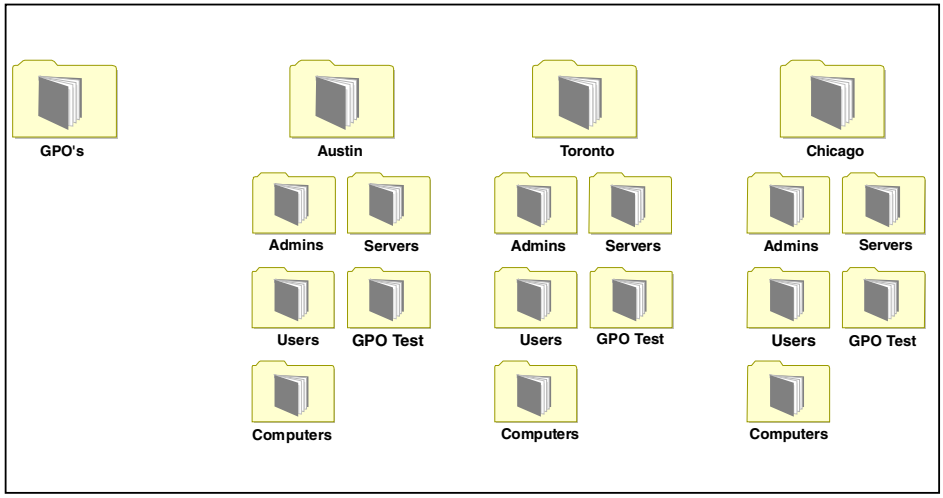

Figure 13-14 Geographical and functional implementation

Note: Of course you have instantly recognized that Figure 13-14 looks just like Figure 13-11 on page 486. This is caused by the fact that in most cases it is useful to place functional OUs (like admins, servers, or users) at the lowest level of the OU design, regardless of whether you use a geographic or political layout on the top layer(s).

Tip: When designing the functional OUs within an AD, keep in mind that each OU definition should relate to an administration need. Each OU should provide a defined point for the application of a group policy or a point of delegated administration.

13.5 Design considerations for IBM TotalStorage NAS

In the beginning of this chapter we provided an overview of Active Directory architecture components and design considerations. But what are the consequences for implementing IBM TotalStorage NAS in a Windows 2000 Active Directory environment? We will approach this subject using the following steps:

1. Placement of the NAS devices
2. Authentication considerations in a domain
3. Administration delegation using organizational units

With the information provided you should be able to successfully implement an IBM TotalStorage NAS system in an Active Directory infrastructure without hassles.

13.5.1 Placement of the NAS devices

The placement of the NAS device needs to fulfill certain requirements:

- Security
- Performance
- Administration

Let us have a more detailed look at those requirements.

Security

First, the physical security should be ensured. This would mean the physical placement of the NAS device in a secure place with restricted access. This could be a dedicated computer room or at least a locked rack.

Next, the placement in the network is important. You should be careful where the NAS device is placed in a logical network environment. For example, it has to be determined whether the NAS device should or must be in the same network environment as the Web servers (with all possible threats from outside), a demilitarized zone (DMZ), or in a more secure zone where you might also place other servers like the database servers.

The more secure the placement of the NAS device, the more likely you will get heavy traffic through your firewalls (for example, when Web servers need to access data on the NAS device) or that you have to open up ports on your firewalls that you would prefer to be closed (like the Network File System (NFS), file transfer protocol (FTP), remote procedure calls (RPC) and others).

Also, it is of high importance to ensure that only valid and authenticated users and applications have access to the data and configuration information of the NAS system. To ensure this, a proper domain and user authentication design of the Active Directory is necessary. Also consider the use of administrative delegation and auditing of your devices. Auditing provides you with information concerning user access and user activity, but you need to make sure that you can actually *identify* a user according to the user name used to log in (for example, using a general administrator account for all administrators would diminish the benefit of auditing). Also keep in mind that auditing decreases performance.

Another aspect to consider is the use of the Encrypted File System (EFS) to safely encrypt data on the NAS system. Remember though that data that is accessed over the network is first decrypted on the NAS device and then send through the network (using CIFS or NFS) to the requesting client machine. So the use of EFS helps only to secure data against local access to the file system, which shouldn't happen anyway (because of other security measures mentioned before). Therefore EFS is normally not used in NAS systems.

Tip: Check physical security (access control), network placement (firewall) and user authentication (individualized accounts, domain design, Kerberos).

Performance

The physical placement of the NAS device can severely affect the performance that users experience when accessing data on the NAS system. This is especially true when you are in a company that uses low-bandwidth lines to connect remote offices. So the best practise is to place the NAS device where the most network traffic to and from the system occurs. For example, in a company with five branch offices you would normally place one NAS system per office to ensure quick data access for the users in the branch offices.

If you use the NAS devices to provide data to application servers (such as Lotus Domino, IBM DB/2 or Microsoft Exchange) then it is of even higher importance to analyze the necessary network bandwidth that is needed between the NAS device and the application servers. Most often you will end up with placing the NAS device and the application servers in a dedicated (for example, backbone) network. Remember that it is possible to place the NAS device in multiple networks, for example, in the user network for user data and in the server backbone for the application data. When analyzing the necessary network throughput, keep in mind that you have to consider peak network usage as well as average network usage.

Tip: Place the NAS device close to the users or applications. Analyze data flow before making a decision where to place the box. It is better to use more devices then absolutely necessary to allow for growth and to accommodate peak usage.

Administration

As the IBM TotalStorage NAS systems support “headless” operations (without keyboard, mouse, monitor attached), the placement for the administration is quite flexible. You need to have a network connection that allows you to use either IBM IT Director, the Web interfaces, or Microsoft Terminal Services client (or any other suitable administration tool you want to use) to administer the system. Keep in mind that remote administration over a WAN connection can be a challenge if the line is congested. Also, you have to ensure that firewalls are configured in a way that you can use the ports required by the respective protocols (for example, HTTP, HTTPS, RDP). Even remote administration using the Internet is possible, although minimally, a virtual private network (VPN) is a must in this case.

Tip: Manage the device remotely. Define a closed set of administrators and the administrative tools to use (Web browser, MMC, Terminal Services client). Delegate administration by use of OUs.

13.5.2 Authentication

Authentication is a key requirement in this environment. It is required that you adhere to common Active Directory design and implementation practices (some have been pointed out earlier) to ensure proper identification and authentication in the directory. Mapping authenticated users to NAS resources (file shares) demands your attention and careful planning and documentation. Regular security reviews of users and shares should be done.

You also need to ensure that authenticated users can be used to track actual individuals or applications accessing data on the NAS device (for example, do not use one general user account for all applications). Also, see “Security” on page 490.

Tip: Use individualized accounts, no generic “administrator”. Change standard user names and make sure the “guest” account is disabled. Re-think AD design and — if required — Kerberos realms.

13.5.3 Administration delegation

Active Directory introduces organizational units (OUs) for administration delegation, so you can ensure that only certain administrators have access to specific administration tasks in the directory. This way, it is possible to delegate the administration of all NAS systems to one group of (NAS) administrators or the administration of one site (including the local NAS systems) to a group of (site) administrators. The placement of NAS devices in your OU design should be already done while designing the OU structure (see 13.4.2, “Organizational Unit design” on page 485 for more information on the design of OUs).

13.6 Bringing it all together — best practice

In case you are tired of all the theory — let us summarize the quick and easy steps to a good Active Directory design and show how to position the NAS devices in your environment:

1. **Design the Active Directory:** Sites, Trees, Forests, Domains, OUs: Keep it simple and straightforward.
2. **Design the Name Resolution:** Consider all possibilities for DNS, DDNS, and WINS, and integrate your name resolution into the AD structure.
3. **Delegate administration:** Decide who manages the NAS devices — and delegate by using OUs.
4. **Think about the data flow:** Decide where the devices should be placed for optimal performance.
5. **Leave room for growth:** Consider the growth in terms of data, users, and applications.

We recommend that you take your time to design your Active Directory. It will be with you for a long time. The amount of work necessary to change an already implemented Active Directory design is immense — and perhaps even impossible to do.

Also remember that Active Directory and Name Resolution go hand-in-hand. One cannot work well without the other.

Using OUs gives you the opportunity to organize the administration of your IT environment and enables a fine-grained security model. How it is implemented will vary from company to company. Some may decide to have administrators manage all NAS devices in all locations, while others might want to delegate all administrative tasks for a location to a group of administrators — including the NAS devices.

Take special care with the data flow. Remember to place the NAS devices as close as possible to the users and applications that need the data stored on them. If these devices are spread across multiple locations, then you may want to consider using dedicated NAS systems for each location and possibly also for specific applications or user groups.

Also leave room for growth. Although IBM provides you with excellent tools to ease the task, upgrading NAS devices is more work than designing the systems from the start so they will fulfill all your anticipated needs. You may have noticed that the amount of data which needs to be stored grows at a tremendous rate.

Also, consider the need to supply more users and more applications with data — or the possibility that more load may be put on the existing application servers. Perhaps the new Web site is more successful then anticipated — in this case you don't want to start planning to purchase new disk drives or an additional NAS device while the Web site is unreachable due to excessive use. You want the device you already have in place to deliver the performance and capacity you need.

13.7 Further reading

There is a wide range of information available on all subjects concerning Active Directory planning, design, and implementation. Here are some documents we would like to point out (these are available without cost):

Microsoft Windows 2000 Server Deployment and Planning Guide (from the Windows 2000 Resource Kit), at:

http://www.microsoft.com/technet/prodtechnol/windows2000serv/reskit/deploy/home.asp

Microsoft's Active Directory Architecture white paper, at:

http://www.microsoft.com/technet/prodtechnol/windows2000serv/deploy/projplan/adarch.asp

Furthermore, a range of information is available concerning IBM TotalStorage NAS systems, at:

http://www.storage.ibm.com/snetwork/nas/index.html

For more information on security practices for IBM TotalStorage NAS systems, refer to the white paper, *Implementing Security Practices for IBM TotalStorage NAS Products*, at:

http://www.storage.ibm.com/snetwork/nas/security_wp.html

Part 4

IBM TotalStorage NAS solutions

In Part 4 of this book we describe how NAS works with applications and explain how it can be used as a business solution:

- First, we show how you can use the NAS appliance with Windows based applications (Chapter 14, “Solutions for Windows based applications” on page 497).
- Finally, after discussing Windows based applications, we switch to UNIX based applications (Chapter 15, “Solutions for UNIX based applications” on page 543).

14

Solutions for Windows based applications

In this chapter we explain how to implement Windows based applications using NAS as a primary storage for the databases. We start with Microsoft SQL Server 2000, in 14.1, “Microsoft SQL Server 2000 on IBM NAS” on page 498. We then show how to set up a MS Exchange 5.5 server with the databases in a NAS appliance, in 14.2, “Microsoft Exchange Server 5.5 on NAS” on page 508.

After covering the Microsoft products, we take a look at IBM DB2 UDB 7.1, in 14.3, “DB2 UDB 7.1 on NAS” on page 514. When finished with the databases, we show how the Lotus Notes Domino Server works with NAS, in 14.4, “Lotus Domino on NAS” on page 525.

Finally, we discuss the SAS software, provided by SAS Institute Inc., which uses SQL or DB2 on NAS. This is presented as a real business solution scenario, shown in 14.5, “SAS on NAS” on page 532.

14.1 Microsoft SQL Server 2000 on IBM NAS

We start with Microsoft SQL Server 2000, using the NAS system as a storage unit for the SQL database.

14.1.1 Installation

The installation of Microsoft SQL Server 2000 is a straightforward process. There are just a few important things to consider while setting up the system.

Attention: While installing the SQL Server, you need to be logged on to the domain. Do **not** log on locally.

While installing SQL, you need to have all the rights and permissions on the application and on the NAS system. We recommend that you install SQL server using the Domain or Active Directory global administrative account. You also need to check that this account has the correct permissions on the NAS share where the database will reside.

Furthermore, we recommend that you create a service account for the SQL server service running on the application server. This service account needs to have to the special right, **log on as a service**. This service account must also have access to the file share on the NAS system. Otherwise, you will not be able to put your database on the NAS system.

Important: The Microsoft SQL Server and the NAS system need to be part of a Domain or Active Directory.

After starting the SQL server installation, we accepted all default values.

For the Windows authentication mode, we choose the **Mixed Mode** option shown in Figure 14-1.

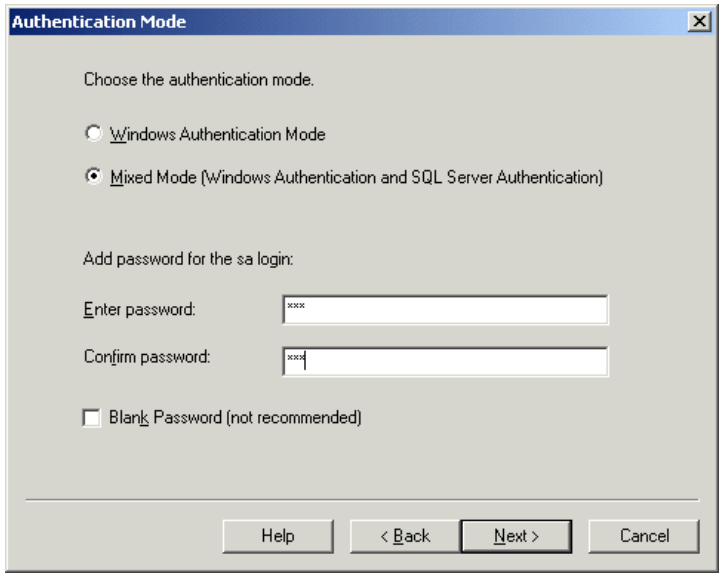

Figure 14-1 SQL authentication mode

During the installation of the Microsoft SQL server, a Services Accounts screen will appear as shown in Figure 14-2. You will need to specify the following options:

- Select **Customize the settings for each service**.
- In the service settings area, you must choose **Use a Domain User account** and provide a service account with the appropriate rights as discussed previously. A local account will not have permissions on the NAS appliance.
- Select the option **Auto Start Service** to automatically start a service when your operating system starts.

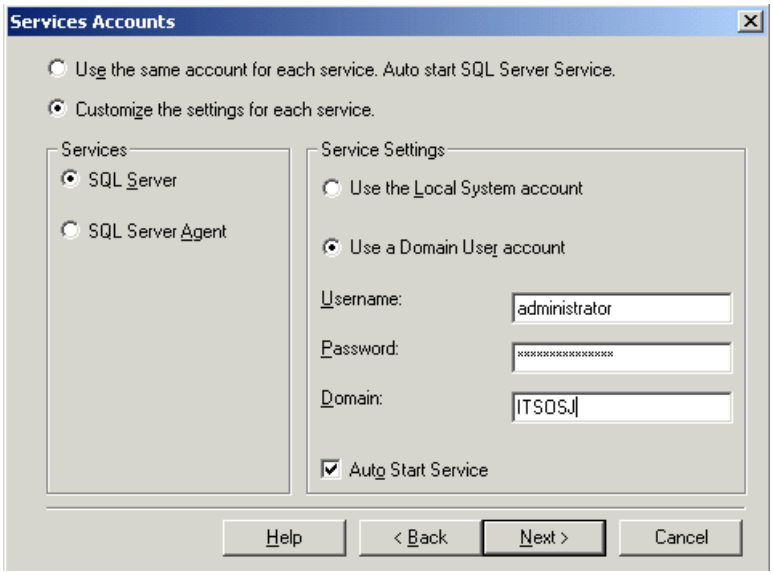

Figure 14-2 SQL services accounts

All program files will be on the SQL server, and the new databases will be created into the NAS server.

We recommend that you use UNC paths, because UNC paths are recognized very early in the boot process, that is, mapping drives; and Microsoft SQL permits use of the UNC.

After the installation is done, go to the SQL Server Service Manager and start the SQL server (Figure 14-3).

Figure 14-3 Start the SQL Server Service

14.1.2 Creating a shared folder

Log in to the NAS server and create a shared folder with sufficient access to the user that you used in installation processes.

The steps are identical to those given in 10.1, “File sharing for Windows clients” on page 354.

However, if you want to use the GUI interface, use the following command to share a folder:

```
net share SQLDATA=E:\MSSQLDATA
```

You will need to set the following share permissions: *read*, *write*, and *change* for the SQL service account.

14.1.3 Creating a database manually using a Query Analyzer

Next, we show how to create a database manually on the NAS system from the SQL Server.

Essentially, only one step is required when creating a database on a NAS system with SQL Server 2000. Prior to issuing any command or performing any operation that uses UNC paths, you must turn on the DBCC trace flag 1807. This flag can be turned on when the SQL Server service is started (making it permanent for the life of the SQL Server session) or, it can be turned on by running a command from the SQL Server 2000 Query Analyzer.

To turn on the DBCC flag when starting up the SQL Server service, the service must be started from a command prompt using the following syntax (substituting proper values for your master database's location and data file extension):

```
sqlservr /dc:\sql\programfiles\msql.mdf /T1807
```

Otherwise, if the flag is to be turned on from within the Query Analyzer, the following statement must be run at the Analyzer:

```
dbcc traceon (1807)
```

Using either method, when the flag is turned on, the SQL Server is informed that you intend to use UNC paths to describe database files. Thereafter, you can use the Enterprise Manager or the CREATE DATABASE command to create the database.

Example 14-1 shows a CREATE DATABASE statement that creates database files on a NAS. The database is called EXAMPLE, and our shared folder is located in \\NAS200\SQLDATA.

Example 14-1 Create sample database

```
DBCC TRACEON (1807)
USE master
      GO
   CREATE DATABASE Example
   ON PRIMARY
   (NAME = Example_data,
      FILENAME = '\\NAS200\SQLDATA\Example_data.mdf',
      SIZE = 10MB,
      MAXSIZE = 20MB,
      FILEGROWTH = 10)
   LOG ON
   (NAME = Example_log,
      FILENAME = '\\NAS200\SQLDATA\Example_log.ldf',
      SIZE = 10MB,
      MAXSIZE = 20MB,
      FILEGROWTH = 20)
   GO
DBCC TRACEON(1807)
```

After you are finished using UNC paths to create your databases on the NAS, you may or may not choose to turn off the DBCC trace flag by issuing the following command from the Query Analyzer:

```
   dbcc traceoff (1807)
```

The screen in Figure 14-4 shows the entire example.

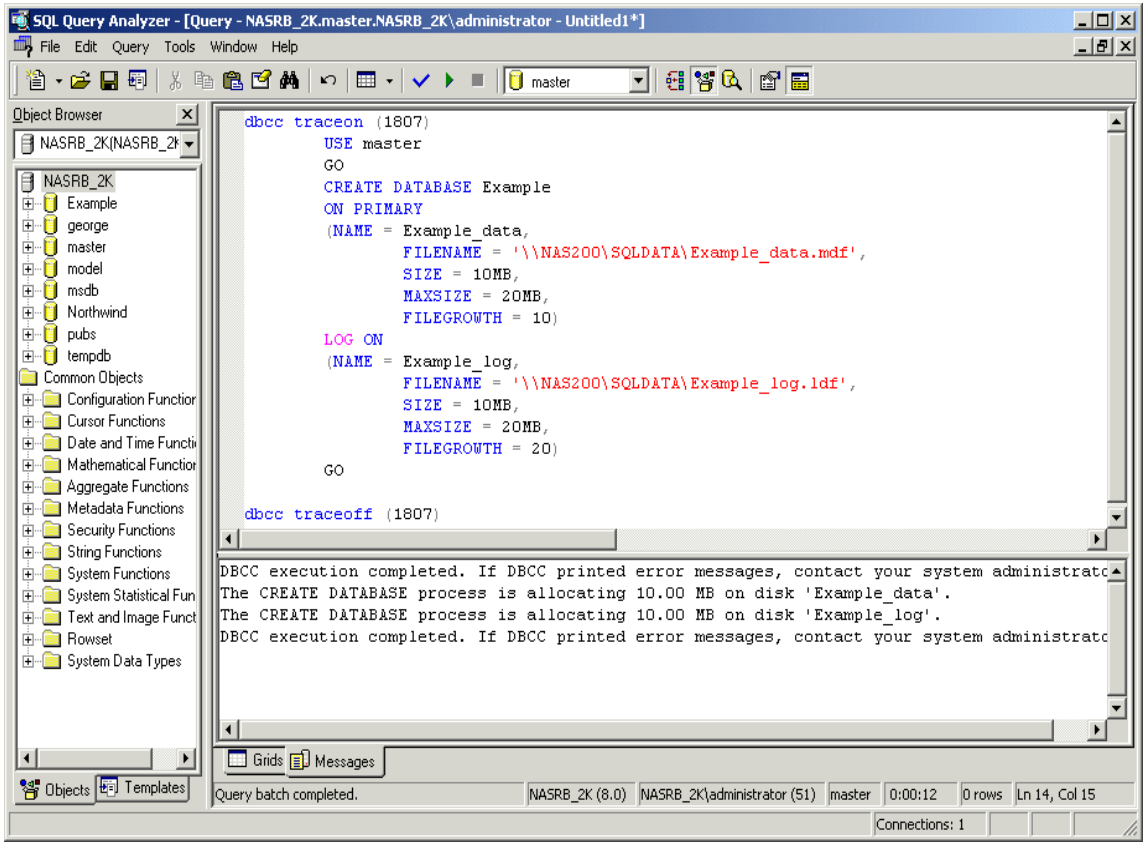

Figure 14-4 SQL Query Analyzer create database query

The database that you have created with UNC paths will work fine at this point. Within the context of SQL Server, everything else you do with the database can be the same as what you would do to a database stored on local disk. Also, rebooting the database server is not an issue, since UNC paths are recognized very early in the NT boot process (unlike mapped drives).

14.1.4 Creating a database using Microsoft SQL Enterprise Manager

You have also the capability to create a database using the SQL Enterprise Manager.

Before you create a database, you must set **`dbcc traceon (1807)`** using the Query Analyzer as done previously.

This example shows how to create a database using the SQL Enterprise Manager. Our database is called EXAMPLE and our shared folder is again located at \\NAS200\SQLDATA.

1. Open the Enterprise Manager and click **Database**, right-click **Databases**, and click **New Database**. On the **General** tab, provide the database name, and click the **Data Files** tab (Figure 14-5).

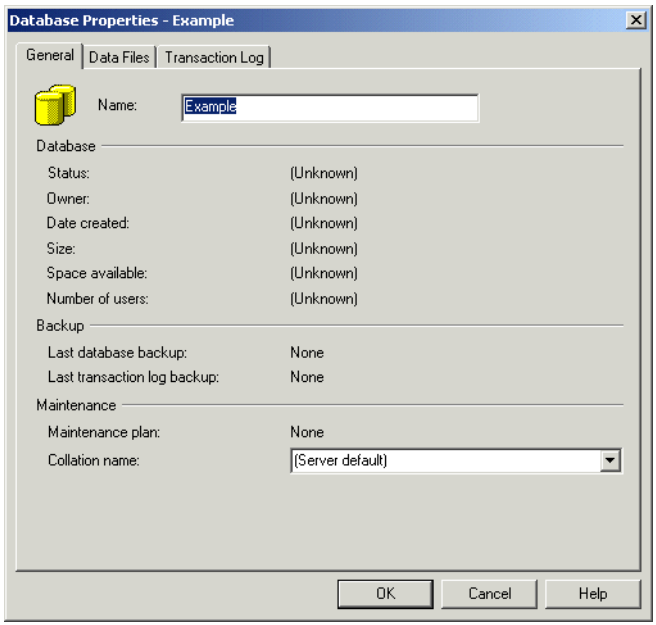

Figure 14-5 SQL Enterprise Manager database definition

2. Change the database location to \\\NAS200\Sqldata\Example_data.mdf and change the size of the database if desired. Click the **Transaction Log** tab (Figure 14-6).

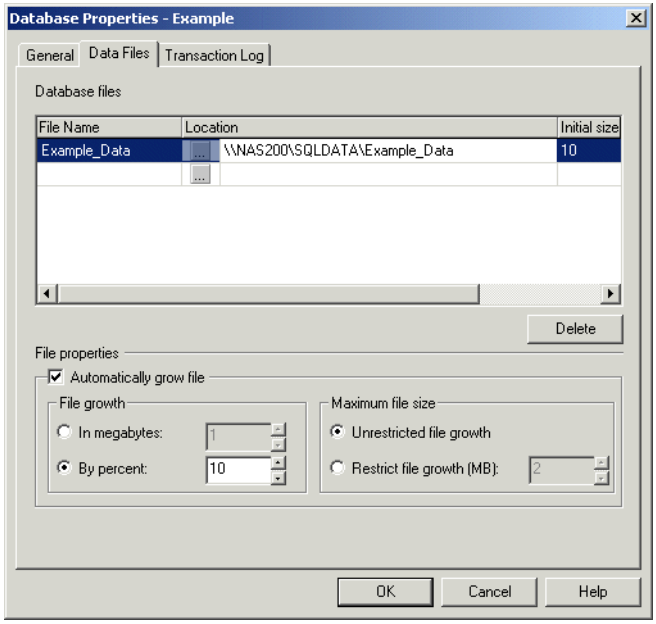

Figure 14-6 SQL Enterprise Manager database files

3. Change the transition log location to \\\NAS200\Sqldata\Example_log.ldf (Figure 14-6).

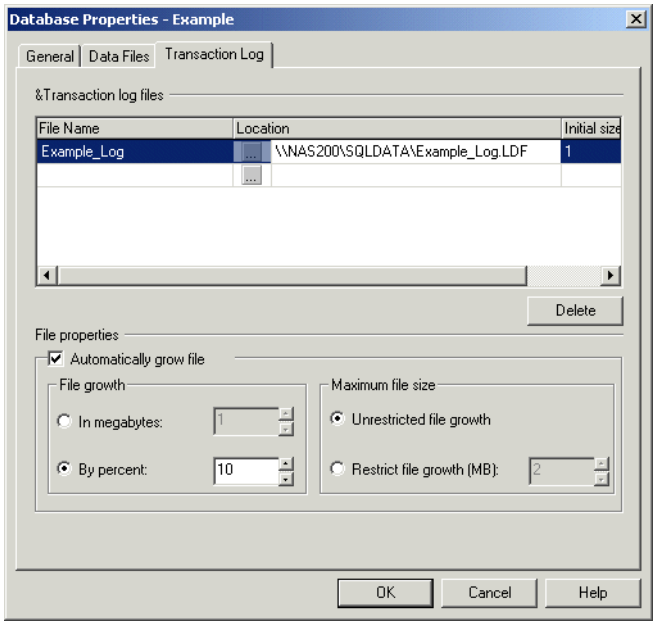

Figure 14-7 SQL Enterprise Manager transition log files

4. Click **OK** to finish and execute this command in the Query Analyzer:

```
dbcc traceoff (1807)
```

14.1.5 Migrating a database from local disk to IBM NAS

Moving a database under SQL Server 2000 can be done utilizing the *sp_detach_db* and *sp_attach_db* stored procedures.

Following are instructions for migrating a database from local disk to NAS storage.

Note: The assumption here is that the locally stored database was originally named *Products*, and its corresponding database files were Products_data.mdf and Products_log.ldf.

1. Ensure that the database that is to be migrated is consistent. To ensure this, once all users have been disconnected from the database, a checkpoint should be issued.
2. Execute the following statement from the Query Analyzer:

   ```
   EXEC sp_detach_db 'Products'
   ```

3. Copy the data (mdf) and log (ldf) files of the database you wish to move.
4. Paste the copied files onto the NAS.
5. In the Query Analyzer, execute the statement shown in Example 14-2 (add additional database files if necessary):

Example 14-2 Move SQL database

```
dbcc traceon (1807)
go
sp_attach_db 'Products', '\\nas200\sqldata\Products_data.mdf',
'\\nas200\sqldata\Products_log.ldf'
dbcc traceoff (1807)
```

6. Delete the copy of the database files from the local disk if desired.

14.1.6 Defining a backup device on IBM NAS

Besides using the NAS system for your database SQL Server database files, you can use it for your backup files as well.

You can create an SQL backup device for the databases shown in Figure 14-8. You can also use the convenient UNC paths in this configuration.

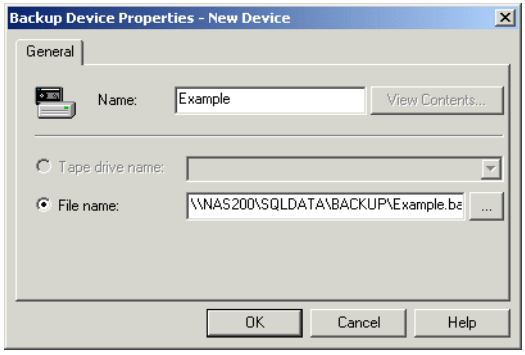

Figure 14-8 SQL backup devices on NAS

14.2 Microsoft Exchange Server 5.5 on NAS

In this section we explain the steps necessary to store the Microsoft Exchange Server 5.5 files in an IBM TotalStorage NAS 200.

Note: Using a NAS appliance for storing the MS Exchange Database Files is not supported by Microsoft or IBM. Creating a test environment to test the solution is strongly recommended. You should evaluate performance, availability, backup, and recovery in your testing.

14.2.1 Pre-installation procedure

There are some steps we need to perform before starting to set up the Exchange Server storage in the IBM TotalStorage NAS. We must have TCP/IP connectivity from the Exchange Server to the NAS appliance. Therefore, the Exchange Server and the NAS server can be connected:

- In the same public LAN (the best choice to exploit resources)
- Using a private LAN or VLAN
- Directly, with a crossover ethernet cable (provides the best performance)

Before setting up a Microsoft Exchange Server, you must create a Windows NT/2000 service account. A service account is a user account that Microsoft Exchange Server services (such as the system attendant and the directory) use to logon.

You should use the same service account for all Microsoft Exchange Server computers in a site. To perform the following procedure, you must be logged on as a domain administrator.

Important: We installed the Exchange Server on a Windows 2000 Advanced Server. The following procedures would change if using a Windows NT server.

- In the `Administrative Tools` group, click **Active Directory for Users And Computers -—> Action -—> New -—> User.**
- Type the name and user logon name of the service account.
- Click **Next**, and in the next dialog type the password, select **User Cannot Change Password and Password Never Expires**.
- Clear the check boxes, **User Must Change Password at Next Logon** and **Account Disabled**, if they are checked.
- Click **Next -—> Finish**.

14.2.2 NAS storage setup

Note: How to set up a NAS device and how to access a shared resource are explained in more detail in 10.2, “Accessing the shares from our Windows clients” on page 370.

The next procedure we will perform is to create a share in the NAS appliance to store the Exchange files:

1. Connect to the NAS device with a browser by typing **`http://nasdevice:1411`** and click “Microsoft Windows 2000 for NAS” to start administering the appliance remotely. Provide the logon information and click **Administer this Server Appliance**.
2. To create a New Share where to store the Exchange Databases, click **Shares -—> Folders**.
3. Choose the Volume in which you will create the share, H:\NAS200DATA, in our case and then click **Share**.
4. In the **General** tab, supply all the information needed as: share name, path to share, protocols allowed for this share.
5. Click the **CIFS** tab (we are going to set up only the CIFS protocol) and give access to the Exchange Service Account user we first created, and limit the access for the others users. Select **OK** to validate.

We can now close the connection. We have configured the NAS Share.

14.2.3 Installing Microsoft Exchange Server 5.5

Next we need to install and configure MS Exchange Server 5.5 as if using the local disks for storage. We perform a simple installation as an example:

1. Run Exchange Setup:
   a. Log on to the Microsoft Exchange Server computer as Administrator. Start the Exchange Setup program.
   b. In the Microsoft Exchange Server Setup dialog box, click **OK** to start the installation process.
   c. Choose the **Complete/Custom** icon for the version that includes the online books, as shown in Figure 14-9.

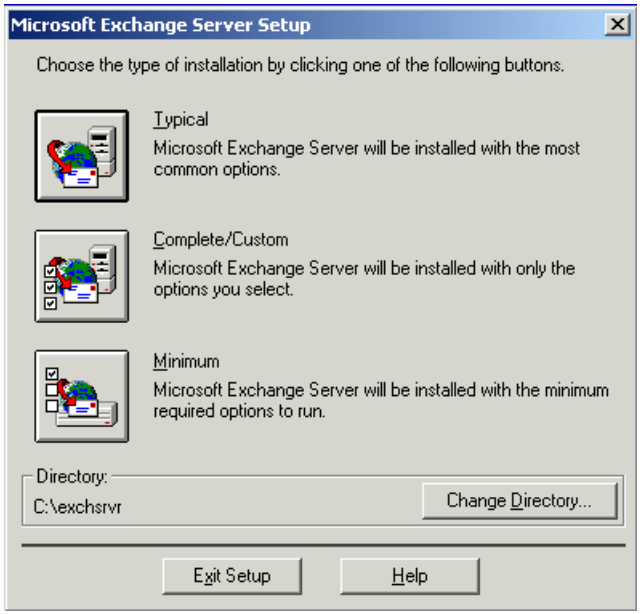

Figure 14-9 Choosing the installation type.

d. Verify that all listed options are selected, and then click **Continue**.

2. Add licenses:
   a. Type the License Key, click **OK**, then click **OK** again.
   b. In the Licensing dialog box, agree to the terms and conditions.
3. Create a new site:
   a. In the Organization and Site dialog box, select **Create a New Site**.
   b. In the Organization Name and Site Name boxes, type the names of your organization and site, and click **OK**.

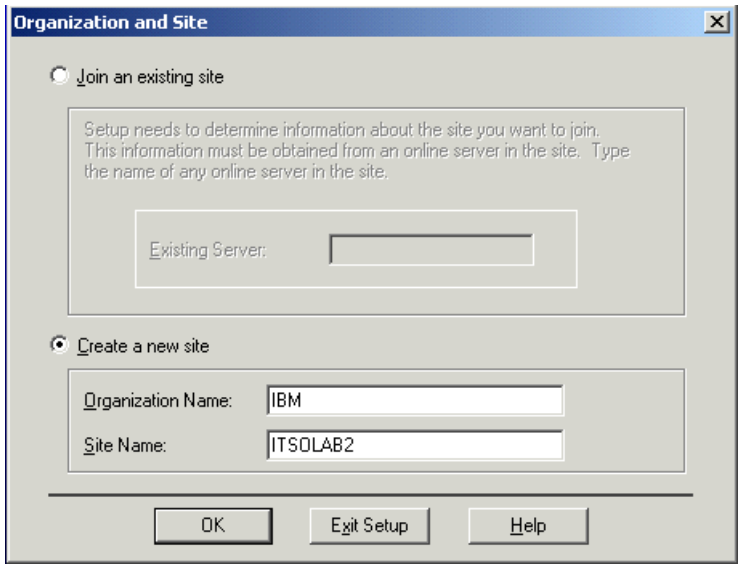

Figure 14-10 Creating a new site

4. When asked if you are sure you would like to create a new site, click **Yes**.
5. Specify the Site Service Account, the Windows user account we created, and the corresponding password. Click **OK**.

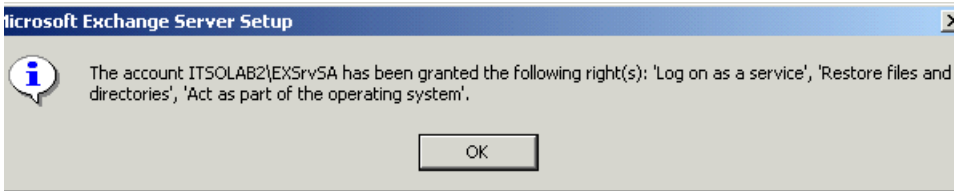

Figure 14-11 Granting the account additional rights

6. After Microsoft Exchange Setup grants the service account additional rights, click **OK**.
7. When prompted to run the Microsoft Exchange Server Performance Optimizer, choose **Exit Setup**.
8. Finally, apply the latest Exchange 5.5 Service Pack.

14.2.4 Moving the files to the NAS Appliance

For this section, we assume that the Exchange 5.5 server is configured using local disks for the storage, as a standard installation. The next step is to move the database files to the NAS device and make it accessible by Exchange. The files we will move as an example are:

c:\dsdata\dir.edb This file stores the Directory Service information.

c:\mdbdata\priv.edb This file stores the private folder information.

c:\mdbdata\pub.edb This file stores the public folder information.

You can also move other Exchange files, including the log files in your particular implementation.

Tip: We recommend that you store the log files in a separate location than the database files. This could be either a local disk or another NAS device.

These are the steps to move the database files:

1. Create the path in the share. We create the *dsdata* and *mdbdata* folders in the share, \\nas200\ExchData. Then we have:

   `\\nas200\ExchData\dsdata`

   `\\nas200\ExchData\mdbdata`

2. Stop the MS Exchange services. Click **Start** -—> **Programs** -—> **Administrative Tools** -—> **Services**.

   These services can vary depending your particular implementation. At least there are:

   – Microsoft Exchange Directory
   – Microsoft Exchange Event Service
   – Microsoft Exchange Information Store
   – Microsoft Exchange Message Transfer Agent
   – Microsoft Exchange System Attendant

   Select each one of these services —> Right-click —> **Stop** to stop it.

3. Move (or copy) the database files to the appropriate path in the share we created.

4. Now we have to provide Exchange with the location for the database files by modifying the path in the registry with the new values. In this example we use the regedit.exe program, as shown in Figure 14-12 on page 513, to modify the file path. This is the most delicate step. We edit the registry keys as follows:

```
HKEY_LOCALMACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\Parameters
Private\DB Path to point to \\nas200\ExchData\mdbdata\priv.edb

HKEY_LOCALMACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeIS\Parameters
Public\DB Path to point to \\nas200\ExchData\mdbdata\pub.edb

HKEY_LOCALMACHINE\SYSTEM\CurrentControlSet\Services\MSExchangeDS\Parameters
\DSA Database File to point to \\nas200\ExchData\dsadata\dir.edb
```

Attention: Using Registry Editor incorrectly can cause serious problems that may require you to reinstall your operating system. Therefore, using the Registry Editor is done at your own risk.

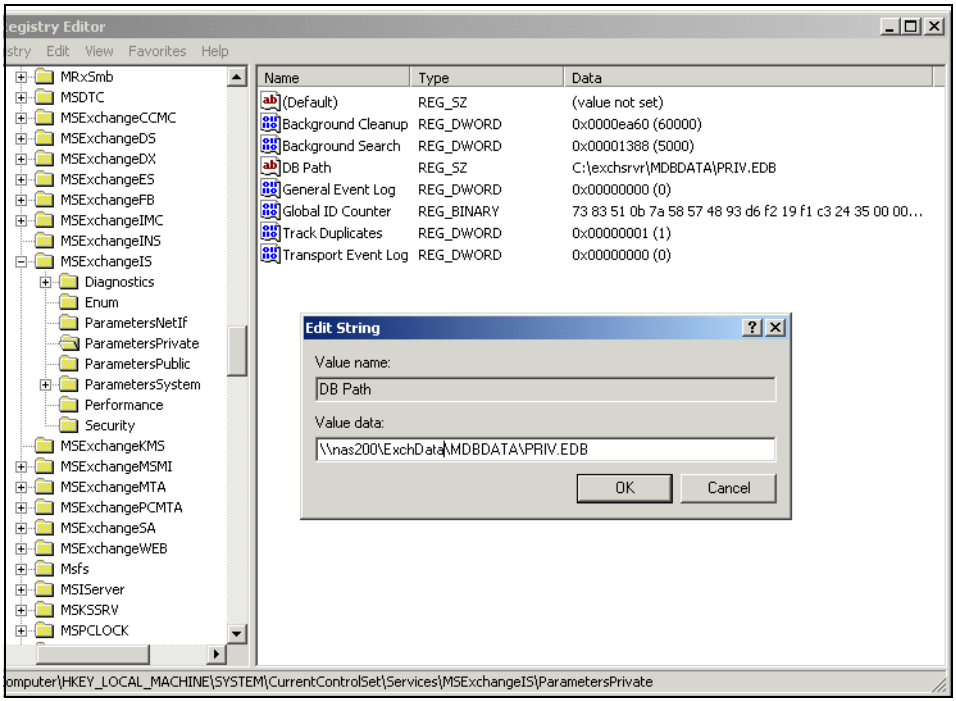

Figure 14-12 Editing the Registry

Note: We can also map a network drive for \\nas200\ExchData and use **autoexnt.bat** to connect to the shared folder on the NAS that contains Exchange Database files. This is shown in Appendix C, “The AutoExNT service” on page 605. We can now substitute the \\nas200\ExchData\ for the mapped drive letter.

5. Start MS Exchange Services. Click **Start -—> Programs -—> Administrative Tools -—> Services**.

   These services can vary depending your particular implementation, and at least include:

   – Microsoft Exchange Directory
   – Microsoft Exchange Event Service
   – Microsoft Exchange Information Store
   – Microsoft Exchange Message Transfer Agent
   – Microsoft Exchange System Attendant

Congratulations! Your Exchange 5.5 Server is now ready to serve clients with the databases stored in an IBM TotalStorage NAS device.

14.3 DB2 UDB 7.1 on NAS

Next, we explain how to install DB2 UDB 7.1 and to create an EXAMPLE database on the NAS system as a storage unit.

14.3.1 Installation

Before we start the installation process, we need to consider some requirements:

- The DB2 UDB 7.1 server needs to be part of a Domain or Active Directory. To satisfy this requirement, the server we are installing, DB2 UDB 7.1, first needs to join the Domain. Later, the Domain security context is used for a DB2 UDB service account.
- You need to set up a DB2 UDB service account that has the special rights, **Act as part of the operating system** and **Log on as a Service**. Refer to the Windows NT or 2000 documentation for using NT Administrative Tools, for instructions on how to grant those special permissions. The service account does not need all the domain administration rights. Also, make sure that the DB2 UDB service account has the right to access the file share on the NAS system.

When you start to install DB2 UDB 7.1, you need to have an account that has administrative rights within your domain.

The installation process needs a domain user account that belongs to the Administrators group on the machine where you are going to perform the installation. You will be asked to specify the installation folder for DB2 UDB. By default, DB2 UDB will be installed in C:\Program Files\SQLLIB, as shown in Figure 14-13.

Figure 14-13 DB2 UDB 7.1 installation

During installation, you are asked to provide a service user account that will be used by DB2 UDB 7.1. By default, the setup program will create a user account with the user *db2admin* (Figure 14-14). You must specify the password. You can accept the default user account or create your own. If you choose to use an existing account, you must use the password that was previously set for this user account.

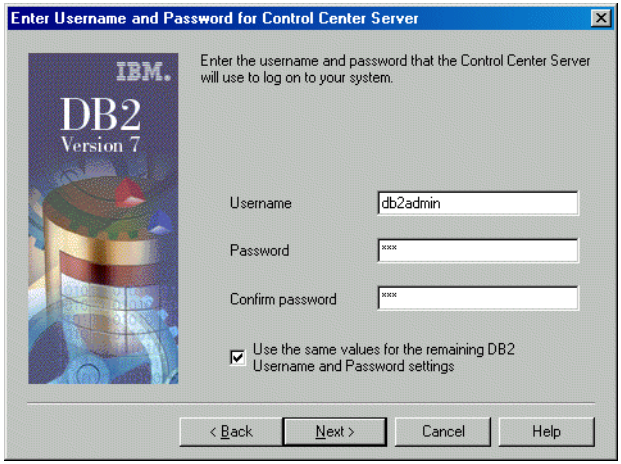

Figure 14-14 DB2 UDB 7.1 service account

Now you are ready to proceed with setting up DB2 UDB 7.1. We chose not to install the OLAP kit, and just finished the installation.

After that, we checked to see that the DB2 UDB services were up and running. You can use **Windows Administrative Tools** to check the status. Start the Computer Management console and select **Services and Applications**. Under **Services**, look for DB2 UDB Services. Their status should be shown as *started*.

14.3.2 DB2 UDB 7.1 object on NAS

Our environment is configured with the DB2 UDB 7.1 files. Programs in the directory **C:\Program Files\SQLLIB,** as well as all the databases, including all containers, will be kept on the NAS server.

To configure the DB2 UDB environment to work with the NAS system, we had to do the following, in this order:

1. Create space on the NAS system.
2. Create a DB2 UDB database.
3. Create the DMS default tablespace.
4. Create user tables.
5. Create temporary and catalog tables.

These steps are discussed in detail in the following sections.

Creating space on the NAS system

DB2 UDB 7.1 needs a mapped network drive when DB2 UDB Services starts up. This network mapping will be created using the AutoExNT service from the Microsoft Windows NT/2000 Resource Kit. For more information, refer to Appendix C, “The AutoExNT service” on page 605.

This example shows how to create a DB2 UDB database using the Wizard from DB2 tools. Our database is called EXAMPLE, our shared folder is located in \\NAS200\DB2DATA, and the network map is S. Here are the steps to perform:

1. Install AutoExNT (refer to Appendix C, “The AutoExNT service” on page 605).
2. Create a shared folder on NAS with sufficient access to the user that you used on installation process.
3. Edit or create an autoexnt.bat file in the directory C:\WINNT\SYSTEM32 and append this line:

```
NET USE S: \\NAS200\DB2DATA
```

Creating the DB2 database

1. From the IBM DB2 program folder, open the Control Center. Open the instance where you are going to create the EXAMPLE database.
2. Right-click the **Databases** folder, then click **Create**, and select **Database Using Wizard** (Figure 14-15).

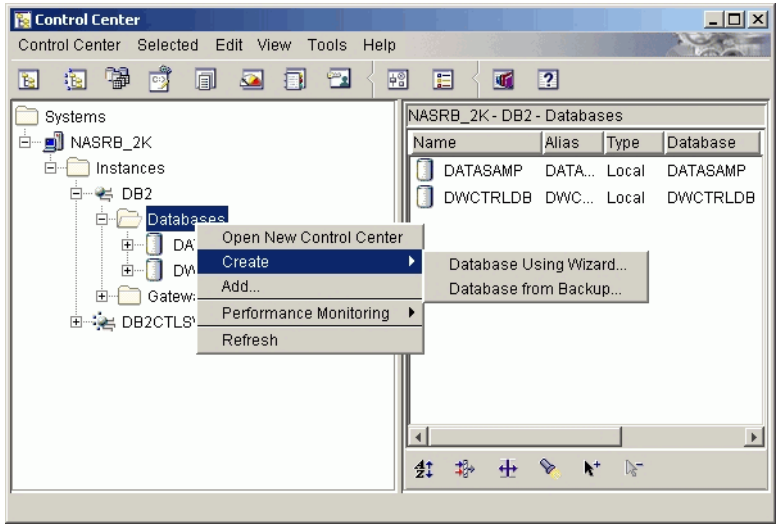

Figure 14-15 DB2 control center launching database wizard

3. For “Database name”, type `EXAMPLE`; for Default drive, choose **S**; for “Alias”, type `EXAMPLE`. See Figure 14-16 for details.

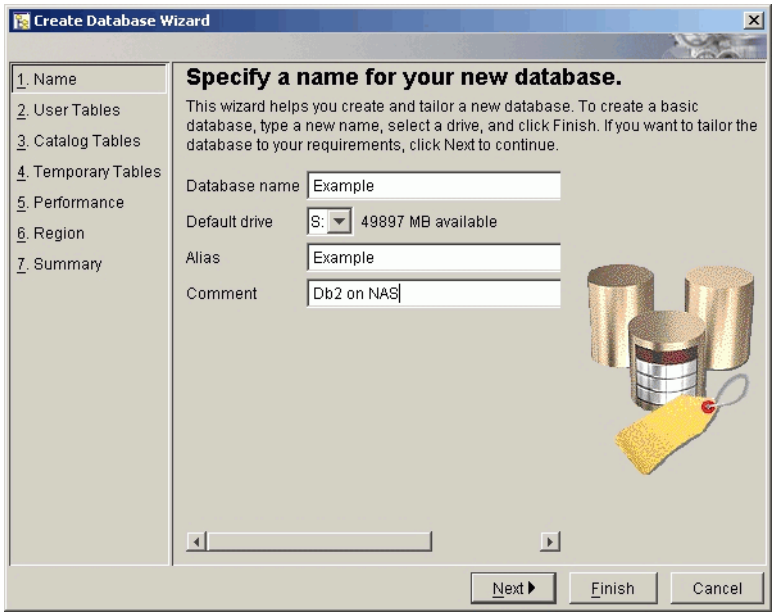

Figure 14-16 DB2 UDB 7.1 create database wizard

Our database will be created using a DMS (Database Managed Space) table space.

NAS works with DMS and SMS (System Managed Space). This allows balancing the needs of administrative simplicity with overall performance. We choose DMS because of performance issues. DMS is faster compared to SMS. DMS supports the raw I/O interface I/O interface. This allows DB2 to write directly to a device or partition without the overhead of a file system. Bypassing the file system in this manner provides the most efficient I/O path possible. One disadvantage of DMS is that it does not grow automatically. Due to this, it needs more administration.

The database manager controls the storage space. The storage model consists of a limited number of devices whose space is managed by DB2 UDB. The Administrator decides which devices to use, and DB2 UDB manages the space on those devices. The table space is essentially an implementation of a special purpose file system designed to best meet the needs of the database manager. The table space definition includes a list of the devices or files that belong to the table space, and in which data can be stored.

Table 14-1 shows the definition of the container sizes and path settings.

Table 14-1 DB2 UDB 7.1 database tables

Database	User Tables	Catalog Tables	Temporary Tables
Container Name	EXCNT_USRTBL	EXCNT_CTLTBL	EXCNT_TMPTBL
Size	20 MB	20 MB	20 MB
Path	S:\Example	S:\Example	S:\Example

4. Finally, on the Create Database Wizard screen, choose **High performance - DMS** and click **Add** to create a container (Figure 14-17).

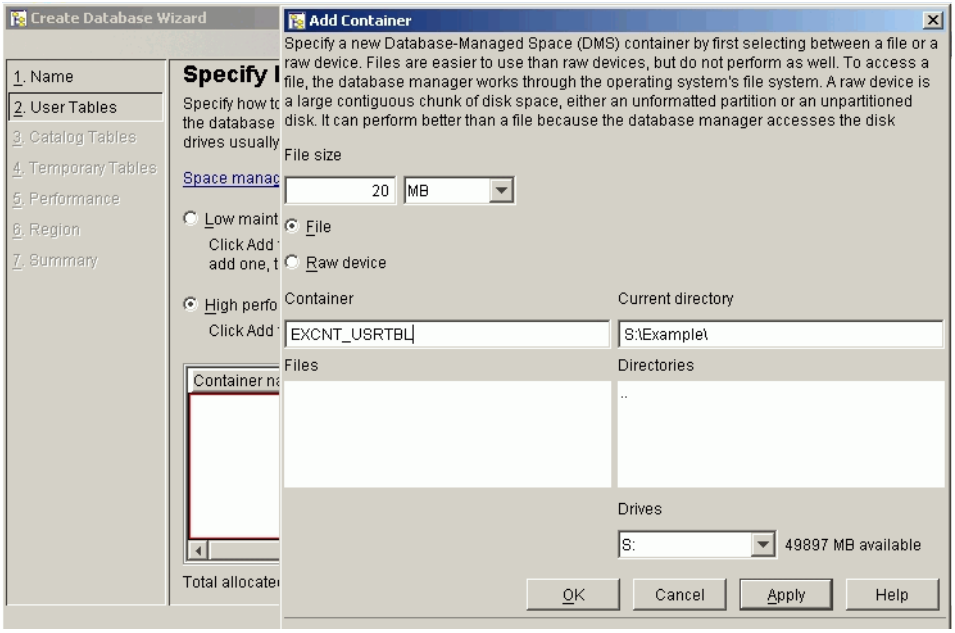

Figure 14-17 DB2 UDB 7.1 creating user tables

Creating user tables

Our user tables container will be the FILE type. We will have 20 MB of space, the name container will be EXCNT_USRTBL, and it will be located in the S:\Example directory. On the Add Container screen (again, refer to Figure 14-17):

1. Click **Apply** to create the container.
2. Click **Cancel**, then click **Next**, to create the user tables.

Creating catalog tables and temporary tables

These are the steps for creating the temporary and catalog tables.

1. First we will create the *catalog tables*:
   a. Choose **High performance - DMS** and click **Add** to create a container. Our catalog tables container will have 20 MB of space, the name container will be EXCNT_CTLTBL, and it will be located on the S:\Example directory (refer to Figure 14-18).
   b. Click **Apply** to create the container.
   c. Click **Cancel**, then click **Next**, to create the catalog tables.

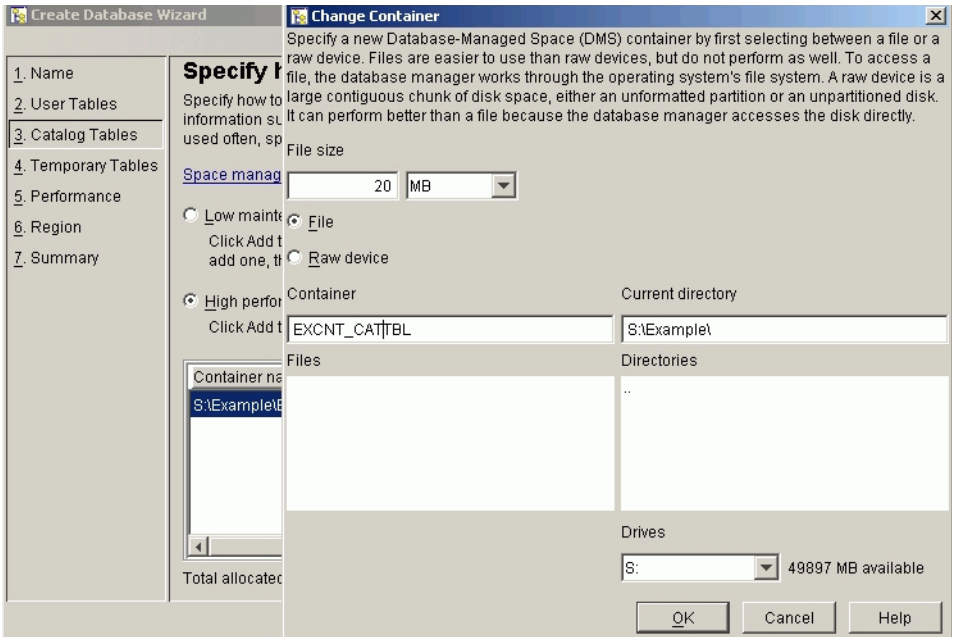

Figure 14-18 DB2 UDB 7.1 creating catalog tables

2. Next we will create the *temporary tables*:

   a. Choose **High performance - DMS** and click **Add** to create a container. Our temporary tables container will be type FILE. They will have 20 MB of space, the name container will be EXCNT_TMPTBL, and it will be located in the S:\Example directory (refer to Figure 14-19).

   b. Click **Apply** to create the container.

   c. Click **Cancel**, then click **Next**, to create the temporary tables.

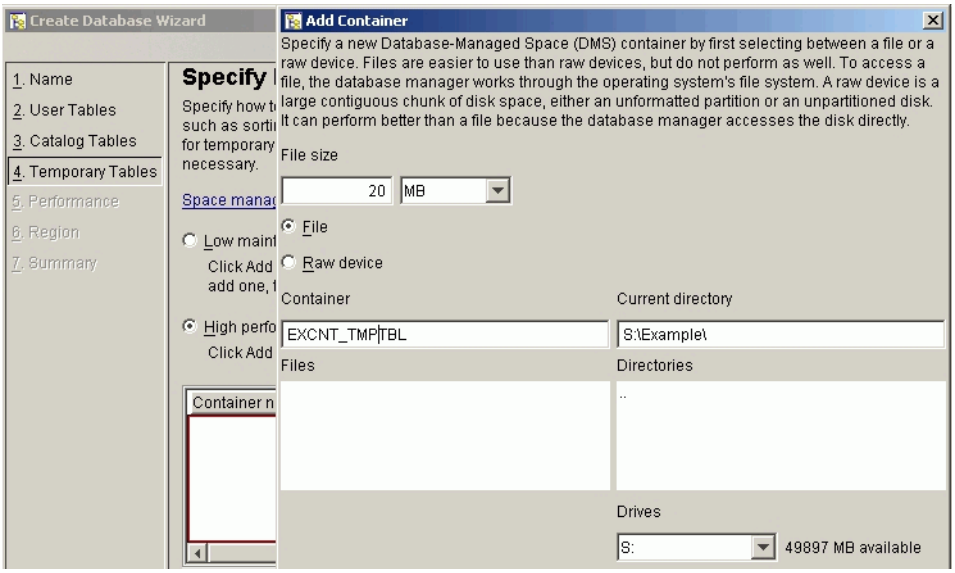

Figure 14-19 DB2 UDB 7.1 creating temporary tables

3. If you wish to create a database manually, see the commands shown in Figure 14-20.

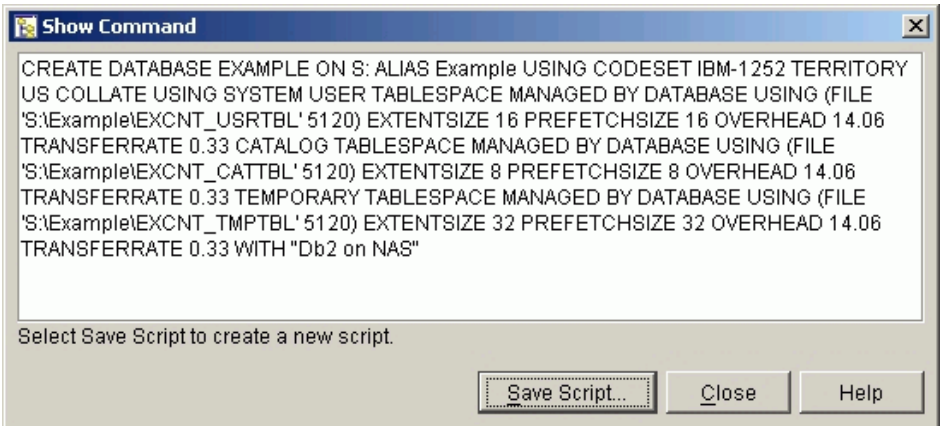

Figure 14-20 DB2 UDB 7.1 commands shown

4. If you want to save your script, then click **Save Script**, and more windows will appear.
5. Select the DB2 UDB instance name.
6. Define a script name, a script description, and a working directory, as shown in Figure 14-21.
7. Click **OK** to save the new script.

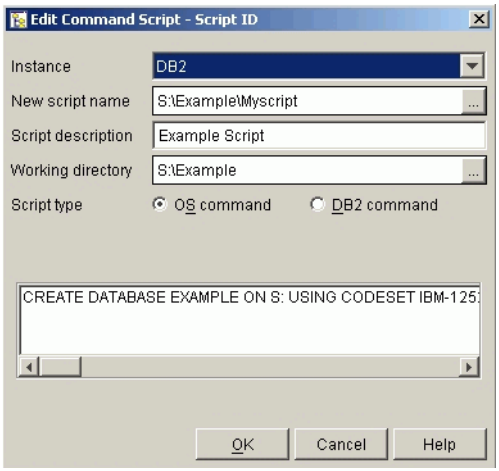

Figure 14-21 Saving a script

Now we can verify the newly created database in the DB2 UDB 7.1 Control Center program (Figure 14-22).

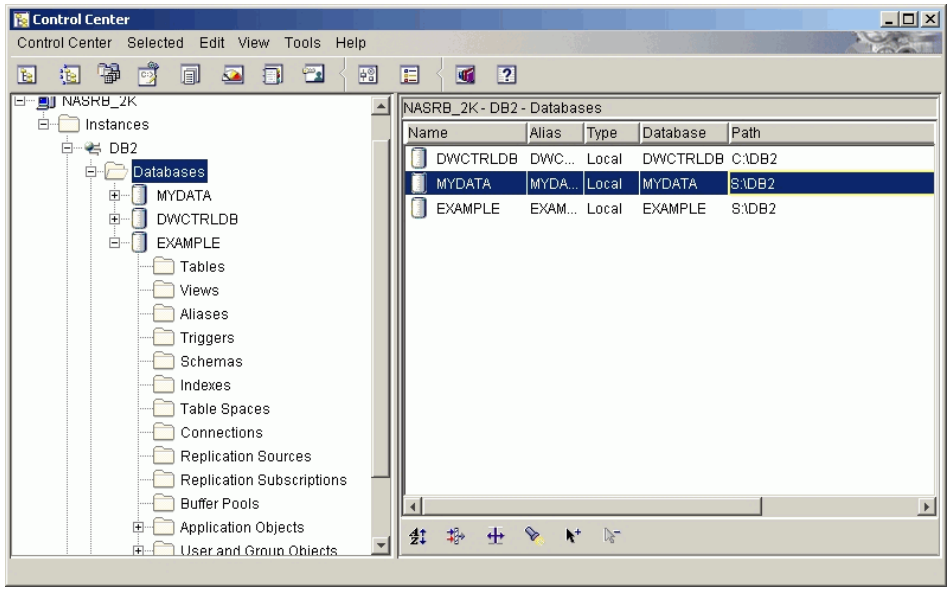

Figure 14-22 DB2 UDB 7.1 Opening the new EXAMPLE database

14.3.3 Connecting to the EXAMPLE database

This procedure shows how a client can connect to our EXAMPLE database. There are three ways to connect to a database: *use a profile*, *search the network*, or *manually configure a connection*.

1. We chose the option, *search the network*:
2. Start the Add Database Wizard and click **Add** on the Client Configuration Assistant's main window.
3. Now click **Add** to create a new connection. The Windows wizard will appear. This wizard defines remote database servers to which your applications can connect and access data. It catalogs nodes and databases while shielding you from the inherent complexities of these tasks. The wizard collects information and creates an entry in the list of available databases. Each entry in the list is identified by a unique local database name. This local name is known as a database alias. It can be used by your applications to access the database.
4. The list of available databases appears in the Client Configuration Assistant's main window. Select the **Source** tab and choose **Search the network** and then click **Next** as shown in Figure 14-23.

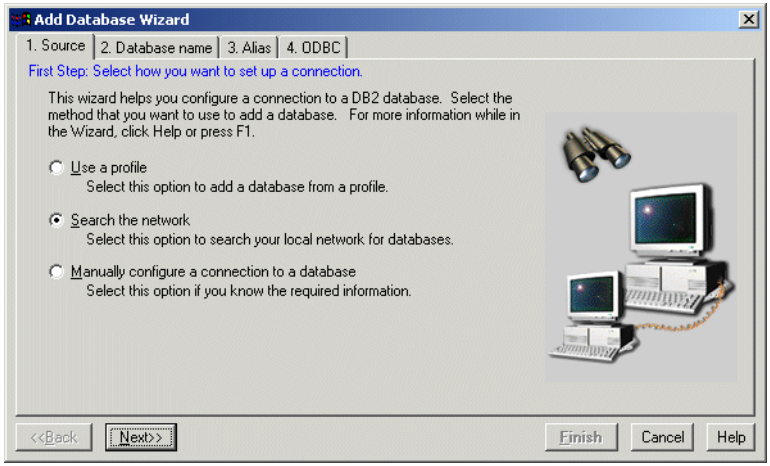

Figure 14-23 Add a database using client configuration assistant

5. To search for databases known to the local machine, expand the object tree until you find the server object and database that you want.
   a. Select the database name from the list and click **Finish**.
   b. If you cannot see the system you want to search, or no systems are listed, click **Add System** to add it to the list of Known Systems (Figure 14-24).

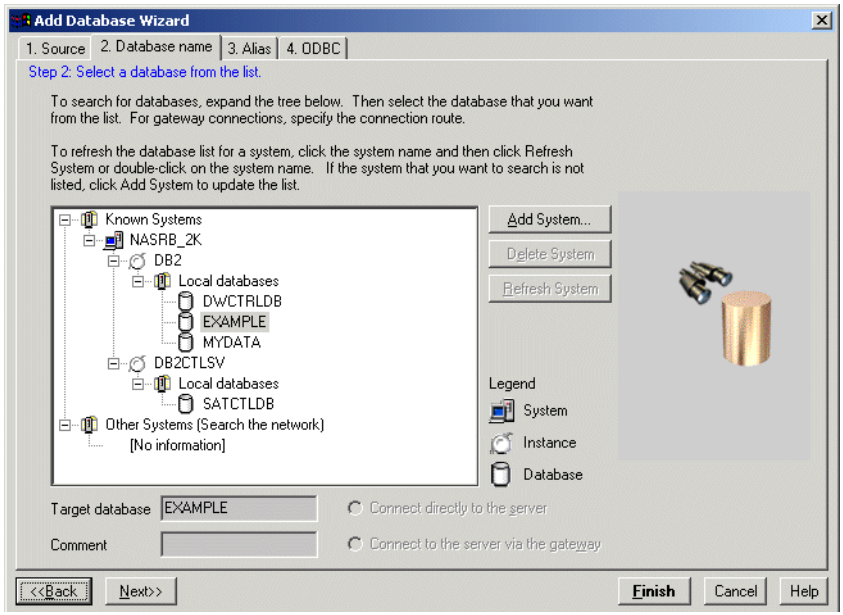

Figure 14-24 Adding Databases using Discovery

6. When you finish adding an entry to the list of available databases, the Confirmation window is displayed (Figure 14-25). You can test the connection, add another entry, or return to the Client Configuration Assistant's main window.

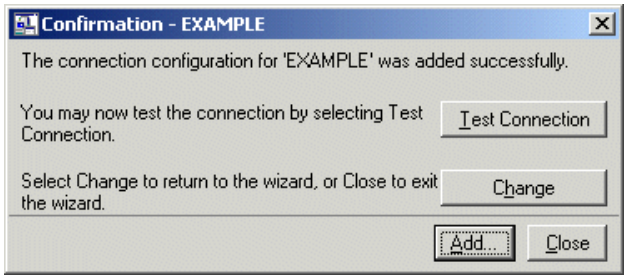

Figure 14-25 Database EXAMPLE added

7. To test the connection, click the **Test Connection** button and the next windows will you ask for a valid username and password for the DB2 UDB 7.1 database. If the connection is successful, the window in Figure 14-26 will appear.

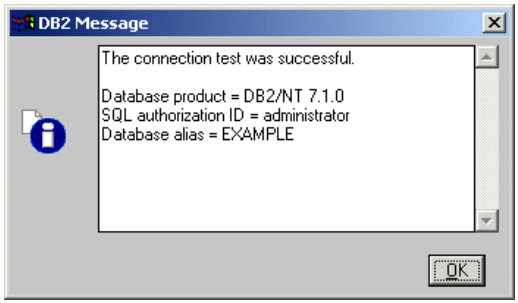

Figure 14-26 Connection successful

14.4 Lotus Domino on NAS

Next we explain how to use the IBM Lotus Domino server with a NAS appliance.

When installing Lotus Domino, you need to perform these steps:

1. Map a network drive from the NAS system.
2. Install the program and data files using any of these combinations:
   a. Program and data files all on NAS
   b. Program files on local drive; data files on NAS (this is the only combination supported if you install Domino as an NT Service)

If the Lotus Domino server is installed as a Windows NT/2000 service, you must do the following extra steps:

1. Install and configure the AutoExNT service.
2. Create an autoexnt.bat batch file.
3. Configure the Domino Service to start manually.

14.4.1 Getting connected to the NAS box

Note: This has been covered previously. If necessary, refer to Chapter 3, “Connecting to the NAS appliance” on page 45.

Simply use the Windows Explorer or the *net use* command to map a network drive from the NAS system. Our shared folder on the NAS system is called \\nas200\domino.

Note: If you plan to install Domino as a Service and use **autoexnt.bat** to connect to the shared folder on the NAS that contains Domino files, you need to uncheck the box *Reconnect at Logon* to prevent double mapping of the same folder. Otherwise, you will get an error message stating that: “The local device name is already in use”.

Our mappings are shown in Figure 14-27.

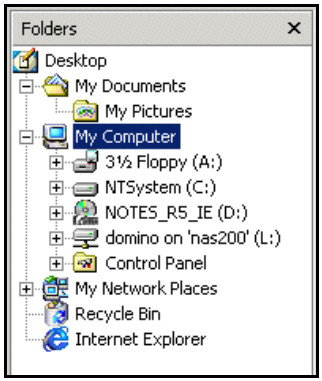

Figure 14-27 Lotus Domino mappings

14.4.2 Installing Lotus Domino

First we will install Lotus Domino with all program and data files on the NAS system. After that, we will install Lotus Domino Server as a Service.

Installing program and data files on NAS

If you chose to install and run everything on the NAS system, follow these steps:

1. Launch the setup program for Lotus Domino.
2. When prompted for the destination folders, specify the mapped network drive on the NAS system (Figure 14-28).

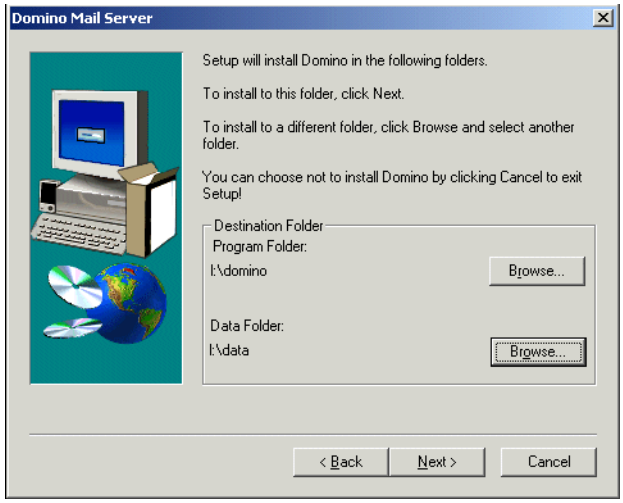

Figure 14-28 Lotus Domino destination folders on mapped drive

3. Follow the on-screen instructions and finish the installation. We went with the default settings.
4. Open Domino and the setup is completed.

Domino Server installed as a Windows NT/2000 Service

If you installed an application as a Service, everything it depends on should be already available, so that there are no problems when the service starts. In the case of Lotus Domino data files being installed on a NAS, the CIFS redirector and related services should have already started, and the drive mapping must have been established for the Domino service to start successfully.

In case the drive mappings are not established, the service will fail to start. To prevent this problem, you need to install and configure another Windows NT/2000 service called AutoExNT. This tool is included in the Windows NT/2000 Resource Kit. For details regarding the installation, see Appendix C, “The AutoExNT service” on page 605.

Follow these steps when you install Lotus Domino as a Windows NT/2000 service:

1. Launch the setup program for Lotus Domino.
2. When prompted for the destination folders, make sure that you specify the local hard disk as the location of the program files, and the NAS folder for the data files, just as shown in Figure 14-29.

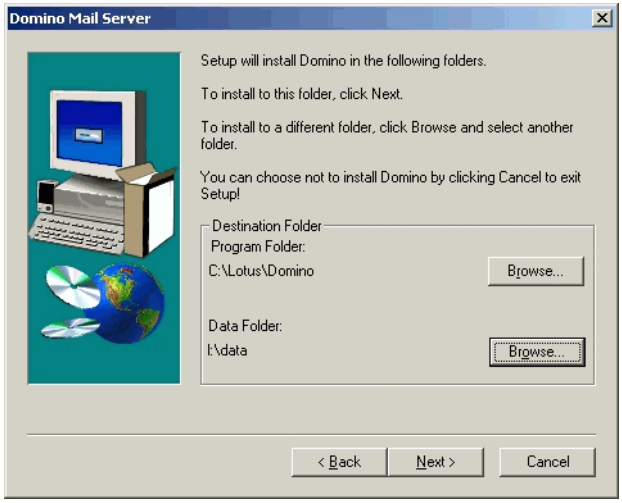

Figure 14-29 Lotus Notes program files on local, data files on NAS

3. When prompted for the components to be installed, check the box beside *Domino as an NT Service*, as shown in Figure 14-30.

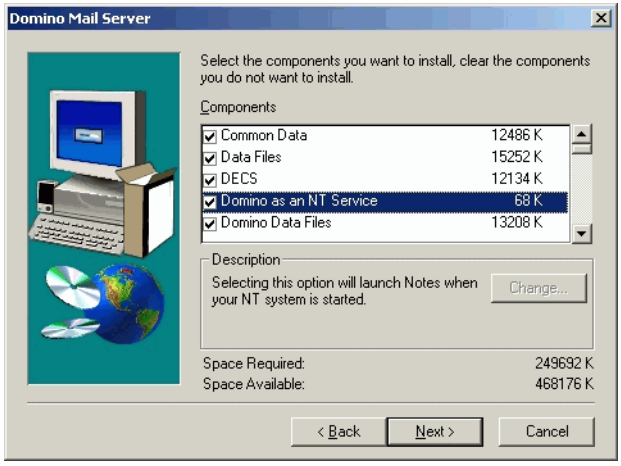

Figure 14-30 Domino installed as an NT Service

4. Follow on-screen instructions and finish the installation. We went with the default settings.
5. Open Domino and complete the setup.
6. Configure Lotus Domino Service to start manually. To do this from your console, right-click **My Computer**, select **Manage**, select **Services and Applications**, and select **Services**. Figure 14-31 shows the service.

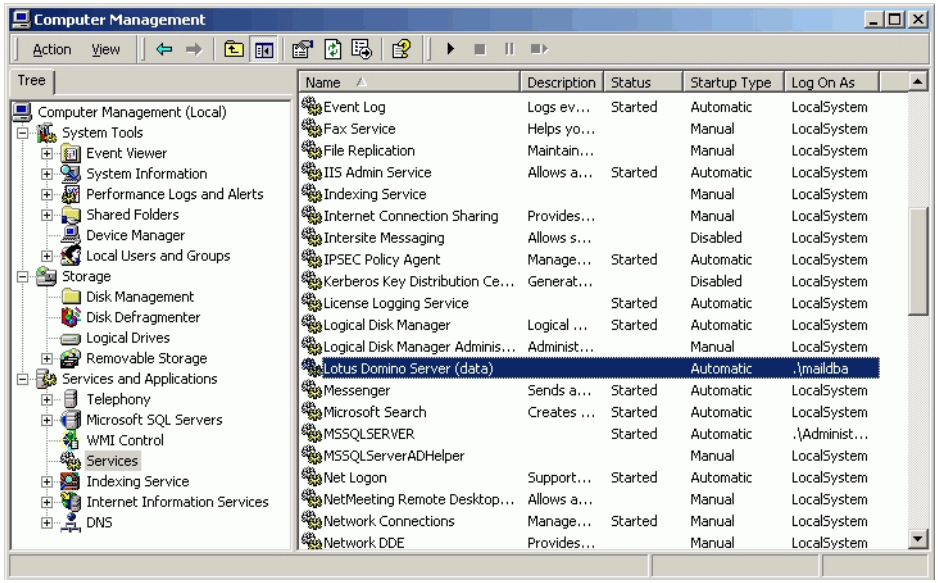

Figure 14-31 Lotus Domino Service configured to start automatically

7. Right-click **Lotus Domino Service**, then select **Properties**. Select the **General** tab. Under **Startup type**, select **Manual** as shown in Figure 14-32.

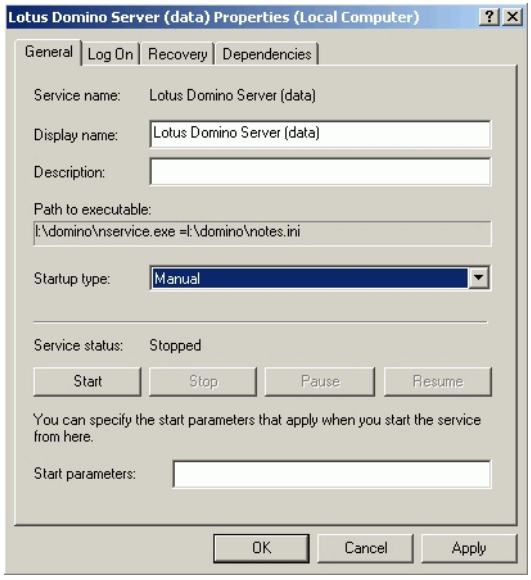

Figure 14-32 Lotus Domino service configured to start manually

8. Select **Log On** tab. Under **Log on as**: click the radio button beside “This account:” and click **Browse**. Select a domain account that has full permissions on the Domino shared folder (for example, ITSOSJ\maildba). Supply the password and confirm it. Click **Apply**. Then click **OK** (Figure 14-33).

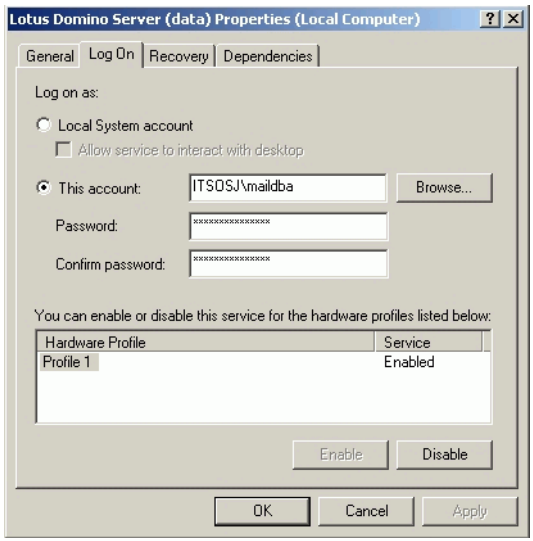

Figure 14-33 Service Logon with permissions to the share

9. Install AutoExNT as discussed in Appendix C, “The AutoExNT service” on page 605.
10. Edit the file named *autoexnt.bat* in the %Systemroot%\System32 directory and type in the following entries:
    a. Specify the drive mapping for the Domino data folder on the NAS (use the same drive letter as the one you used during the installation)

    ```
    c:\winnt\system32\net use l: \\nas200\lotus
    ```

    b. Specify the starting of the Lotus Domino Service

    ```
    c:\winnt\system32\net start “Lotus Domino Server (Data)”
    ```

11. Save the changes in autoexnt.bat (Figure 14-34).

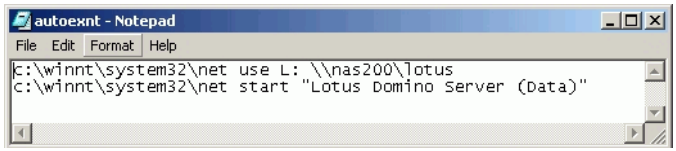

Figure 14-34 Example of autoexnt.bat

12. To test if the setup is working, go to the Services window, right-click the **AutoExNT** service, then select **Start**. Finally, you should see that the status of AutoExNT and Lotus Domino Server has changed to *Started* (Figure 14-35).

Note: If you installed AutoExNT with the */interactive* option, you should see a DOS window pop up, showing the drive mapping commands and the starting of Domino server when you start AutoExNT.

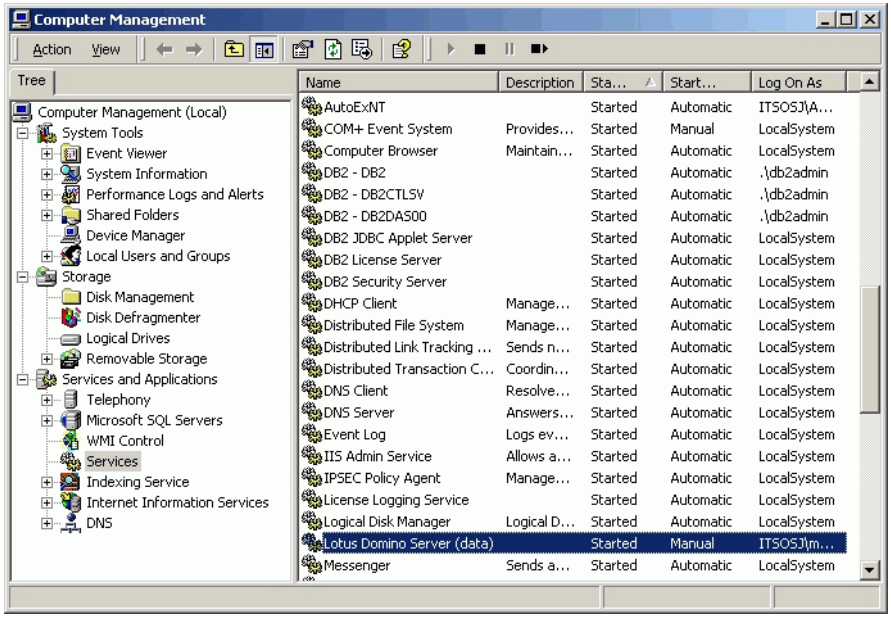

Figure 14-35 AutoExNT and Lotus Domino Server started

Now you have successfully accomplished the task of installing Lotus Domino Server.

14.5 SAS on NAS

In this section we document the network, software, and hardware requirements, as well as the software installation and configuration procedures required, to set up a database application program that demonstrates a data warehouse/data mart scenario.

We assume that the reader is generally familiar with Microsoft SQL Server, DB2 UDB 7.1 and Tivoli Storage Manager software installation and configuration procedures, in addition to the hardware and network device installation and setup procedures required to replicate the demonstration.

14.5.1 What is SAS?

SAS Institute, the world's largest privately held software company, provides software and services that enable customers to transform data from all areas of their business into intelligence. SAS® solutions help organizations make better, more informed decisions and maximize customer, supplier and organizational relationships. Built on award-winning technologies, SAS® software provides the foundation, tools and solutions for data access, data management, data analysis and enterprise-wide information delivery.

Through a modular approach, SAS software offers users the flexibility to choose desired areas of functionality to build customized applications and solutions. SAS products are available for data entry, retrieval, and management; report writing and graphics; statistical and mathematical analysis; business planning, forecasting, and decision support; operations research and project management; statistical quality improvement; computer performance evaluation; and custom applications development. In addition, SAS offers complete business solutions for analytical customer relationship management (CRM), supplier relationship management (SRM), and enterprise performance management (EPM) across a variety of industries.

SAS software solutions offer a tradition of scalability and are designed to handle the ever-growing, massive quantities of data generated from a variety of business channels — including voice, mail, person-to-person, fax, call centers, the Web, and beyond. An enterprise-wide storage infrastructure has to be able to deliver fast and reliable data access to a broad set of users: wherever they are, whenever they need it, and in whatever format their applications call for. SAS software solutions coupled with IBM TotalStorage network storage appliances deliver an optimal combination of performance, scalability and reliability to do just that and more.

You will find the SAS Administrator documentation at:

`http://www.sas.com/service/admin/admindoc.html`

You will find SAS/ACCESS sample programs for UNIX at:

`http://www.sas.com/service/techsup/sample/unix_access.html`

14.5.2 The SAS solution scenario

In this demonstration, we are showing how the NAS 200 and 300 system integrate into a more complex environment. We will use the SAS software to get access to the database residing on the NAS system.

The connection can be seen as a three-step process:

1. Set up the database server (DB2 UDB or SQL Server) and put the database on the NAS system.
2. Set up the SAS server and connect SAS to the database software on the application server.
3. Connect the SAS client from the workstation to the application server.

The last step is optional. All queries could be run on the application server itself. The installation itself is a straightforward process, and is provided in 14.5.3, “SAS installation” on page 537.

But first, here is an overview of our solution scenario:

The heart of the network is a Cisco 3500 XL Gigabit Ethernet switch. The IBM NAS 200 appliance is the data warehouse in the scenario, serving as the data mart.

An xSeries 330 server is used to provide all the network services and the Windows 2000 environment, including Active Directory, DHCP and DNS services. The Tivoli Storage Manager (TSM) software also resides on the server. and an IBM 3583 Tape Unit provides physical backup storage for the network through a SCSI interface.

The other xSeries Server is the application server based on the Windows 2000 Server. DB2 UDB, Microsoft SQL Server, and the SAS server are installed on this machine.

The whole environment is shown in Figure 14-36.

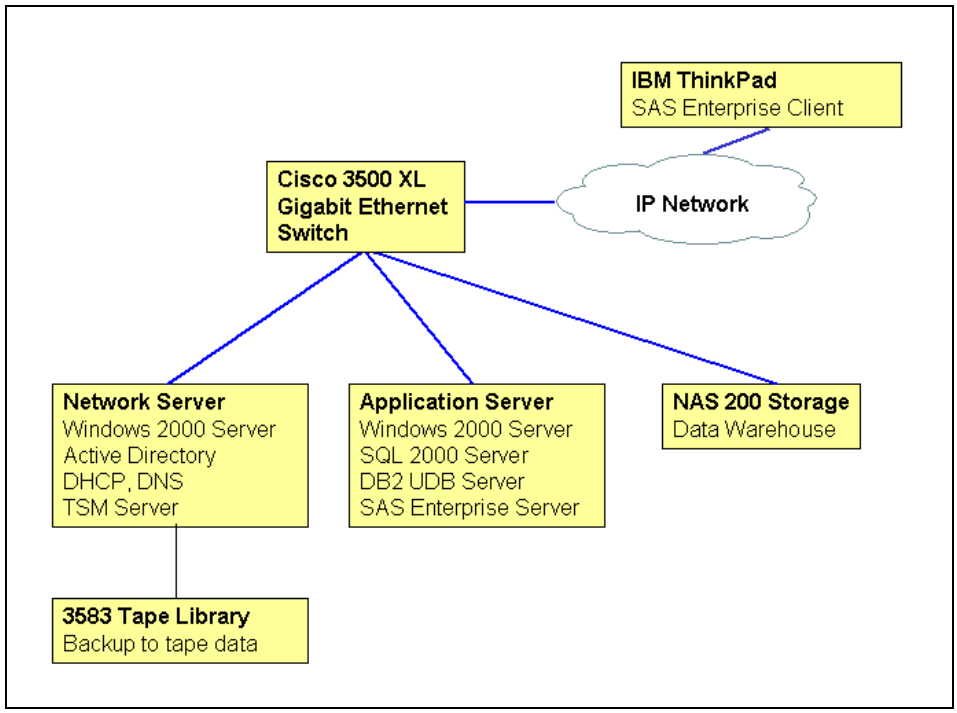

Figure 14-36 SAS solution scenario

You should be familiar with setting up a Windows 2000 environment, including DHCP and DNS name services, since it is beyond the scope of this book to document this in detail. You have already learned how to connect SQL Server and DB2 UDB 7.1 to the NAS system in Section 14.1, “Microsoft SQL Server 2000 on IBM NAS” on page 498 and in Section 14.3, “DB2 UDB 7.1 on NAS” on page 514. You will also need this information to set up SAS and connect it to the database.

Note: Remember, while setting up the database software, that you must have the correct rights and permissions locally, within Active Directory and on the NAS system.

You can start to install DB2 UDB 7.1 or Microsoft SQL Server 2000 as described in the previous chapter on the application server using the NAS server as primary storage for the database.

After a successful installation, check that the database system works with a query analyzer.

The overall process to set up this scenario is shown in Figure 14-37.

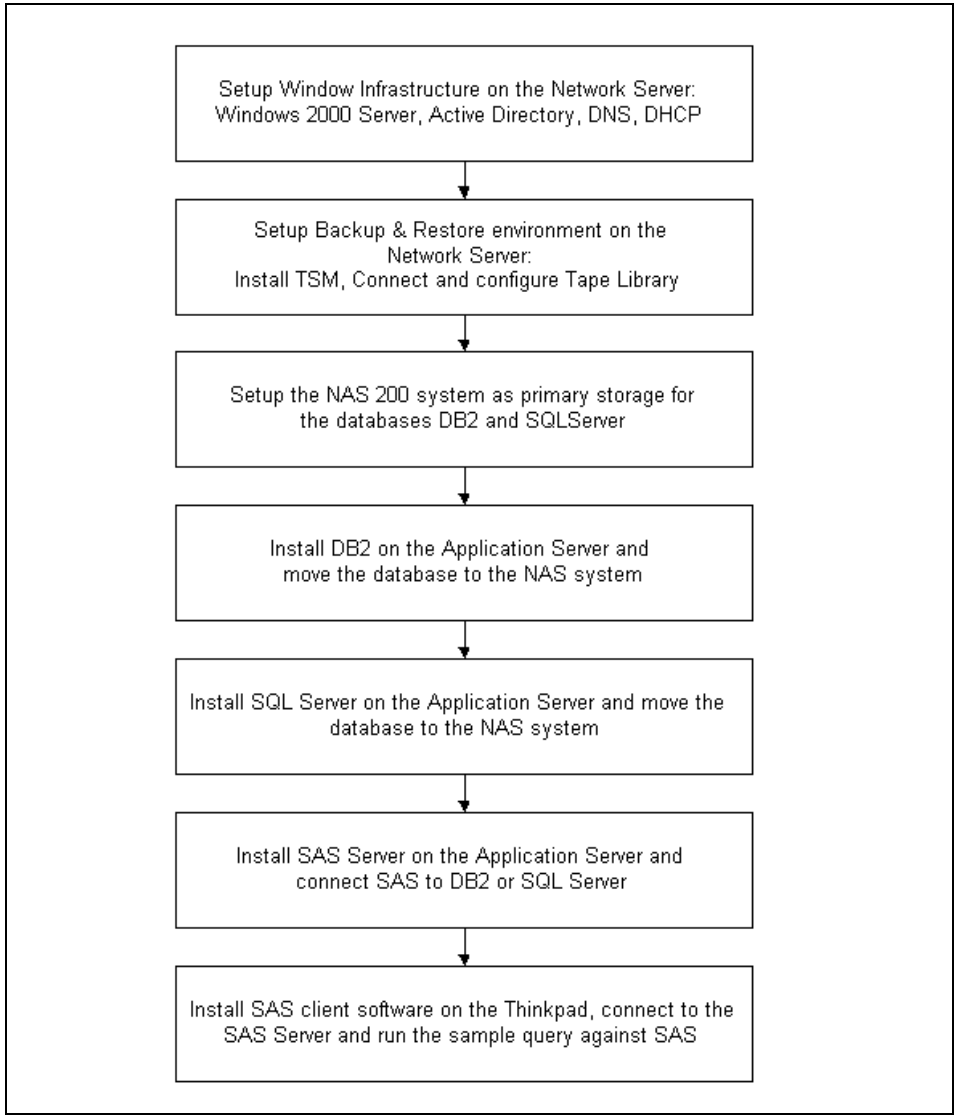

Figure 14-37 SAS solution install process

14.5.3 SAS installation

Next, we cover the SAS specific steps for the setup:

1. Insert the SAS CD and start to install the **SAS System Setup** (Figure 14-38). The SAS system installation wizard will start.

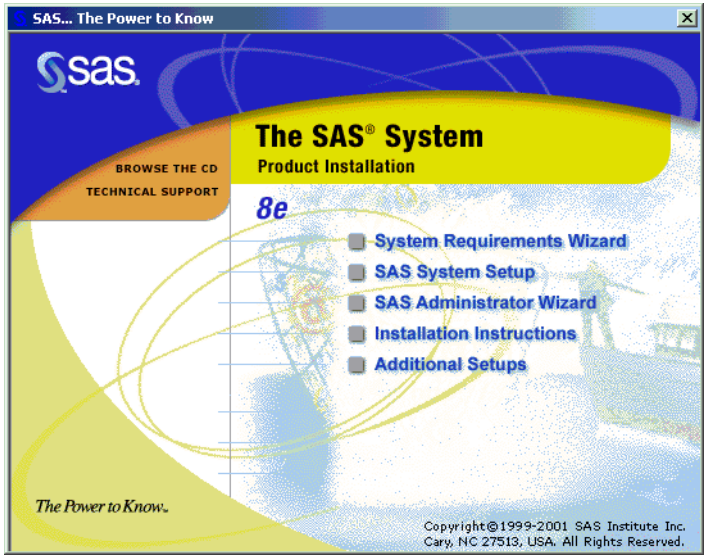

Figure 14-38 SAS setup start

Attention: SAS will need the *Microsoft Runtime Environment*. If you do not have it installed, the SAS installer will do it for you.

2. The SAS installer will prompt you to install the Microsoft Runtime Environment. Follow the instructions if you do not have it installed.
3. Run the **SAS System Setup**.
4. Click **Next**.
5. Select Setup Type = **Complete**.
6. Confirm **Yes** to create new folder for the SAS system files.
7. Accept the destination folder.
8. Select and accept the default for the data files folder.
9. Confirm your destination folder selection.
10. Confirm your temporary files location selection.
11. When prompted, insert the Solution Disk 1.
12. Confirm to update by choosing **Yes**.
13. Click **Finish** and you are done.

Now you can find your SAS installation under **Startup – Programs – SAS system** (Figure 14-39).

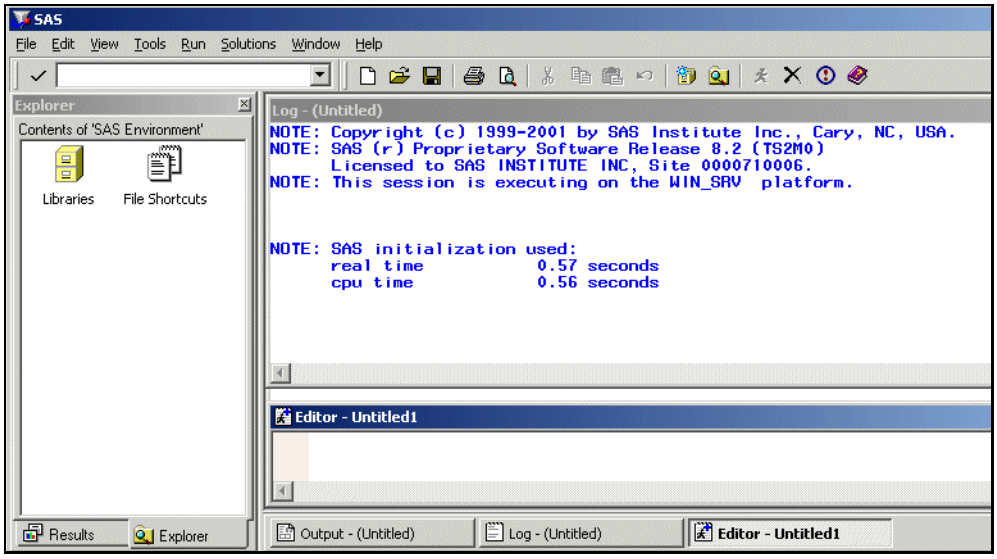

Figure 14-39 SAS initial screen

14.5.4 Connecting SAS to the database

It is important to define which database the SAS software will use. We connected DB2 UDB 7.1 and Microsoft SQL to the system. The next two subsections show what the sample command will look like. For more details, refer to Appendix E, “SAS connections” on page 617 or visit the SAS Web page:

http://www.sas.com

SAS and DB2 UDB 7.1

Example 14-3 shows how to connect DB2 UDB 7.1 to the SAS server.

Example 14-3 DB2 SAS job

```
libname sqldata DB2 ssid=user1  perserve_tab_names=yes preserve_col_names=yes;
```

SAS and Microsoft SQL Server 2000

Example 14-4 shows how to connect SQL Server to the SAS server.

Example 14-4 SQL SAS job

```
libname sqldata SQLSVR ssid=user1  perserve_tab_names=yes
preserve_col_names=yes;
```

Figure 14-40 shows how the SAS database is displayed within the SQL Server Enterprise Manager.

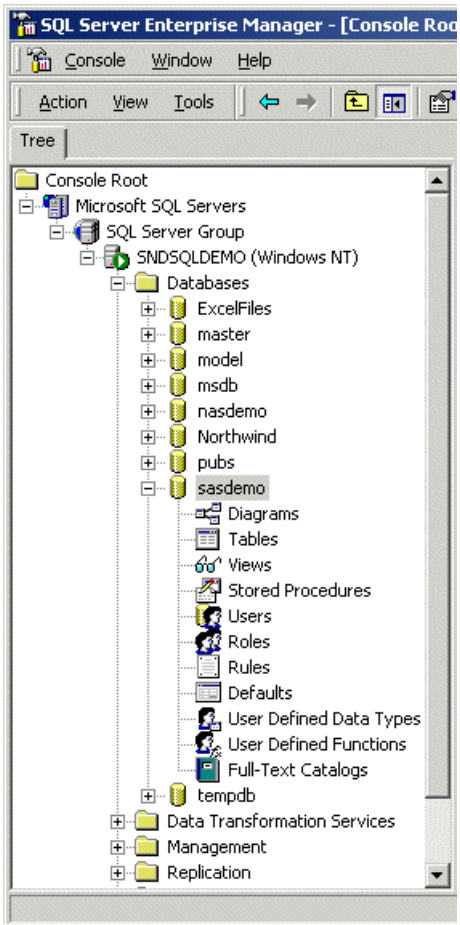

Figure 14-40 SAS in SQL Server Manager view

After that, we looked at some data with the SAS database. The sample is shown in Figure 14-41.

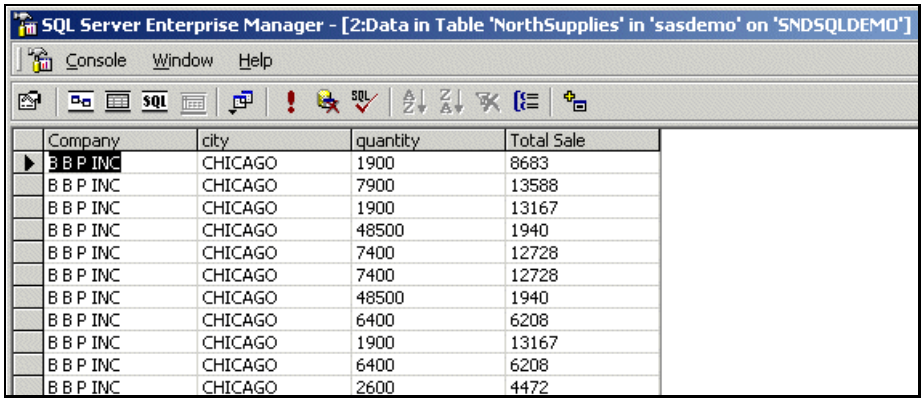

Company	city	quantity	Total Sale
B B P INC	CHICAGO	1900	8683
B B P INC	CHICAGO	7900	13588
B B P INC	CHICAGO	1900	13167
B B P INC	CHICAGO	48500	1940
B B P INC	CHICAGO	7400	12728
B B P INC	CHICAGO	7400	12728
B B P INC	CHICAGO	48500	1940
B B P INC	CHICAGO	6400	6208
B B P INC	CHICAGO	1900	13167
B B P INC	CHICAGO	6400	6208
B B P INC	CHICAGO	2600	4472

Figure 14-41 SAS in SQL Server Manager view of data

14.5.5 Querying SAS and getting the graph

Finally we are able to query the SAS server. The SAS server will use one of our databases, depending on the configuration, and will produce a graph to display the data. The example code is shown in Example 14-5.

Example 14-5 Querying SAS

```
pattern1 color=green;
pattern2 color=red;
pattern3 color=blue;
pattern4 color=cyan;
pattern5 color=brown;
pattern6 color=gray;
pattern7 color=yellow;
pattern8 color=orange;
pattern9 color=pink;

goptions global=reset htext=1 ftext=swissb htitle=1.5;

PROC FORMAT;

  *-GENERATE FORMAT CODES FOR VALUE LABELS-----*;
  VALUE $FUELHEAT
  '0'='Unknown'
  '1'='Gas'
  '2'='LP'
  '3'='Electric'
  '4'='Oil'
```

```
  '5'='Coal'
  '6'='Wood'
  '7'='Solar'
  '8'='Other'
  '9'='None'
  ;
  VALUE $YRBUILT
  '0'='Unknown'
  '1'='1989 or 1990'
  '2'='1985 to 1988'
  '3'='1980 to 1984'
  '4'='1970 to 1979'
  '5'='1960 to 1969'
  '6'='1950 to 1959'
  '7'='1940 to 1949'
  '8'='1939 or earlier'
  ;
run;

libname db2data '/sasdata/census/output2';

title1 h=1.5 c=black f=swissb 'Distribution of Houses in Metropolitan New
York';
title2 h=1   c=black f=swissb 'By Heating Fuel Type';
title3 h=1   c=black f=swissb 'Subgrouped by Year Built';

proc gchart data=db2data.hrecs;
  where msapmsa = '5600';
  vbar fuelheat / subgroup=yrbuilt;
  format fuelheat $fuelheat.
         yrbuilt  $yrbuilt. ;
run;
```

The result of Example 14-5 is shown in Figure 14-42.

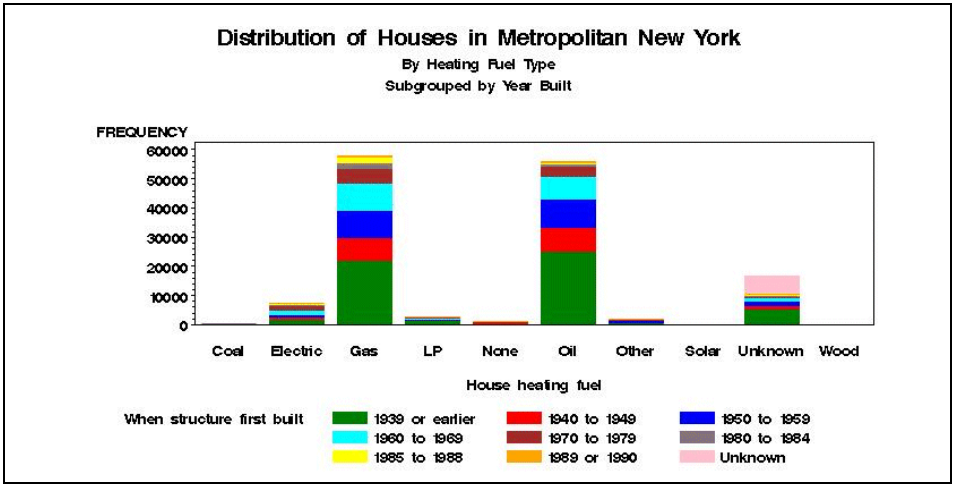

Figure 14-42 SAS sample graph

15

Solutions for UNIX based applications

In this chapter we show how to implement UNIX based applications using NAS as a primary storage for their databases. First, we explain how to mount NFS shares from the NAS 200 or 300 server to UNIX in 15.1, “AIX NFS mount procedure” on page 544. Once the shared folder is mounted with the NFS utility on the UNIX side, we use that shared folder to create our database files.

We start with DB2 UDB and AIX in 15.2, “Implementing DB2 UDB V7.1 on AIX 4.3.3” on page 548. Then we show how Oracle works on AIX and NAS in 15.4, “Implementing Oracle V8.1.7 on AIX 4.3.3” on page 553. Finally, we demonstrate the use of DB2 UDB on Linux in 15.3, “Implementing DB2 UDB V7.1 on RedHat Linux V7.1” on page 550.

15.1 AIX NFS mount procedure

To mount a file system on AIX, you have to be the *root* user, and you can choose from two options: the system management interface (SMIT), or the **mount** command.

1. To mount a file system on UNIX, you have to create a mount point on a UNIX system. You can use the **mkdir** command to create a directory as a *root* (admin ID/authority of UNIX), for example, **mkdir /db2data**.
2. Then enter **smitty nfs** and press the Enter key.
3. Then select **Network File System (NFS)** and press the Enter key.
4. Select **Add File System for Mounting** and press the Enter key.
5. The following three screens show the values for all of the fields that we used. In most cases, we have left the default values, but you can adjust the values based on your environment. On the first screen (Figure 15-1) we have input some fields, which are indicated by highlighting.

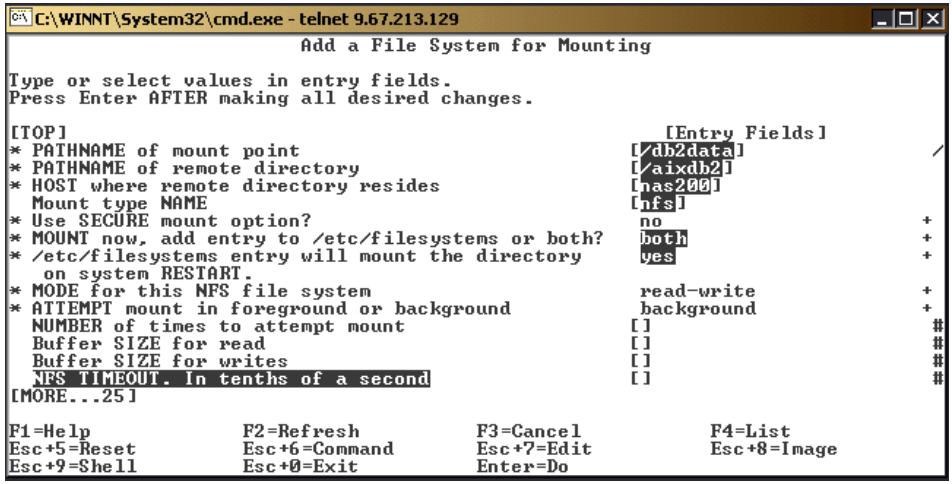

Figure 15-1 Configuring NFS mount screen1

6. The next screen gives us more input options (Figure 15-2).

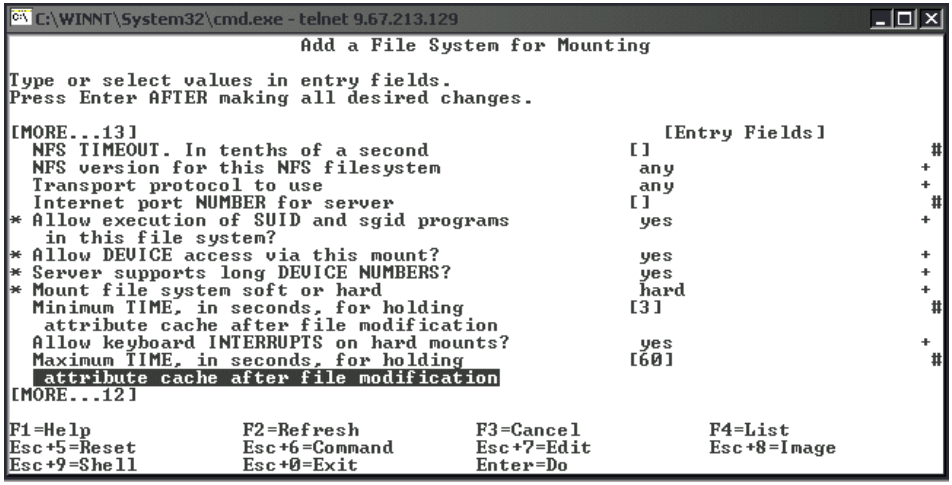

Figure 15-2 Configuring NFS mount screen2

7. Here is the last screen for this task (Figure 15-3).

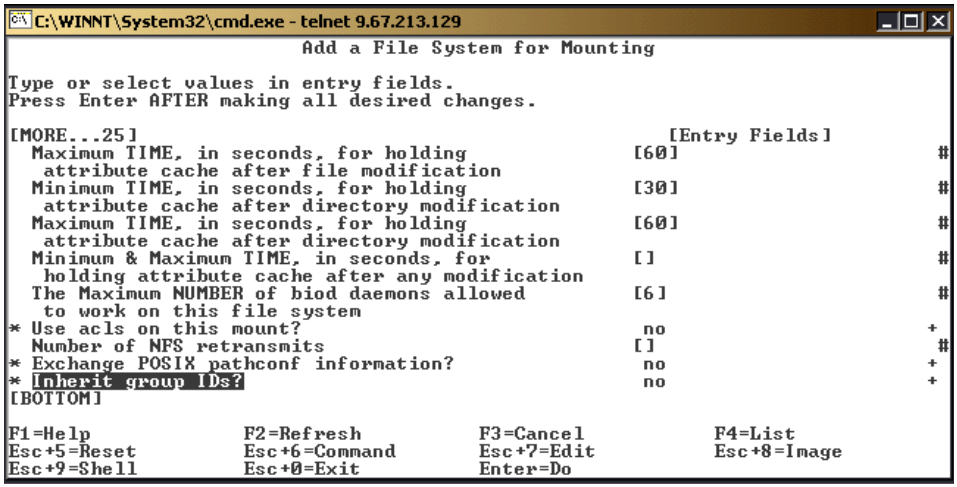

Figure 15-3 Configuring NFS mount screen3

8. Once you have filled out all of the necessary fields, then press the Enter key. You will see successful a completion of the screen, as shown in Figure 15-4.

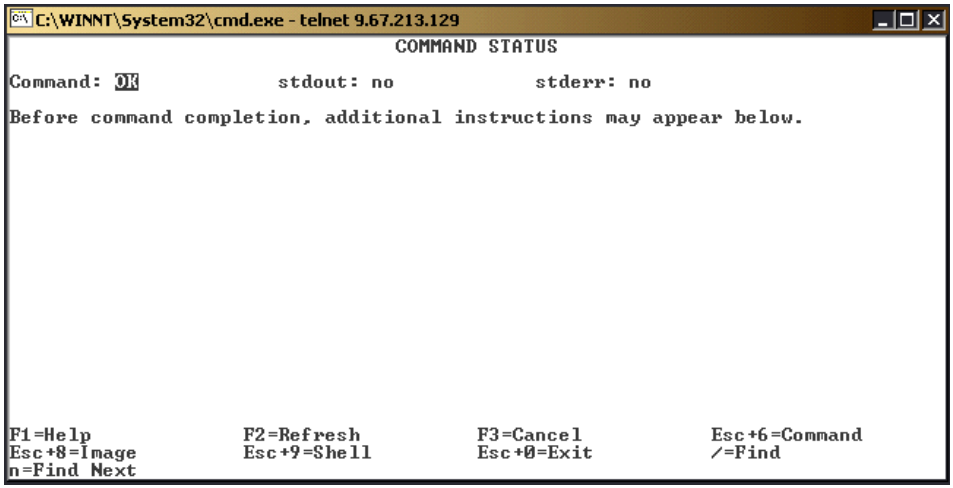

Figure 15-4 Successful NFS mount configuration

9. Alternatively, you can use the command **/usr/sbin/mknfsmnt** with appropriate inputs to mount a file system (see the **man mknfsmnt** command for more detailed options).

10. After successfully mounting the NFS share, the **dk -k** command shows the NFS mounted volume. See the last entry in Figure 15-5.

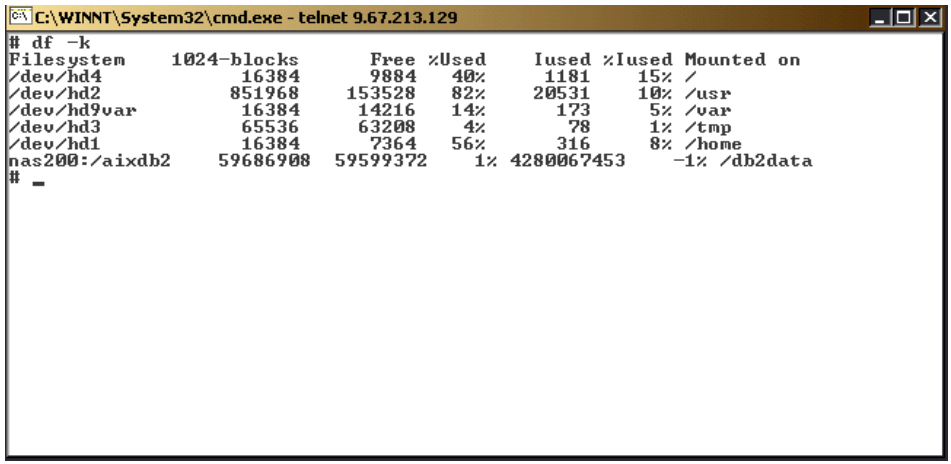

Figure 15-5 File systems display on AIX (df -k command output)

Next, we show how the UNIX user ID/group mapping, with Services for UNIX on the NAS side, works.

1. First we try to create a file or directory with a user called *mfp* belonging to the group *guest*. Notice the authentication errors in Figure 15-6. After that, we mapped that in SFU and see how file permissions looks like. Notice the file (test) and directory (testdir) creation with owner and group ID showing when listing the files with the `ls -l` command.

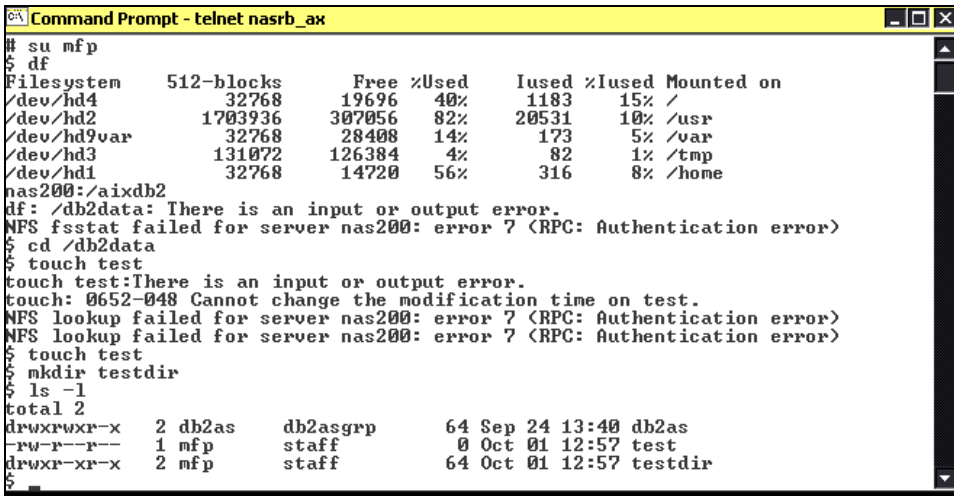

Figure 15-6 UNIX user/group ID mapping effect on NFS shared volume

2. The next example shows what the mount point permissions should look like and how this affects the file and directory creation as a UNIX user. Notice the permission of the mount point prior to the change, then after the change (Figure 15-7).

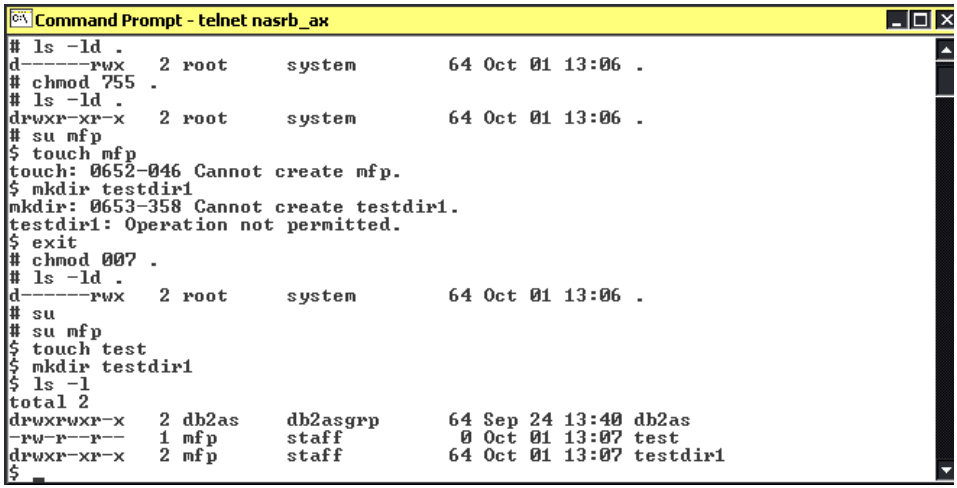

Figure 15-7 NFS mount point permission set for UNIX users

Now our NFS mount from the NAS system /db2data is ready to use by the database software.

15.2 Implementing DB2 UDB V7.1 on AIX 4.3.3

There are two ways you can install DB2 UDB on the AIX server. One way is to do a manual install, or by running the **db2setup** script. We used the **db2setup** script and let it install most of the packages by default, except that we specified where to create the database. We wanted to create the database on the mounted NFS file system.

Once the DB2 UDB installation was completed, we logged on as *db2as* and entered **db2sampl /db2data**, where **db2data** is the mounted NFS volume. You will get a prompt back after successful creation of the database. See the following screen for more details (Figure 15-8).

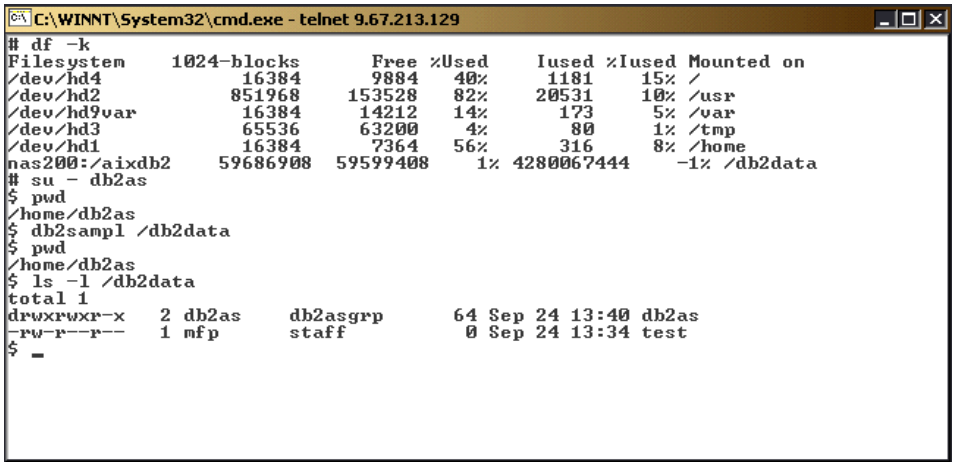

Figure 15-8 Sample database created on NFS mounted volume

Query sample database

Figure 15-9 shows running a query against the database on the mounted NFS volume.

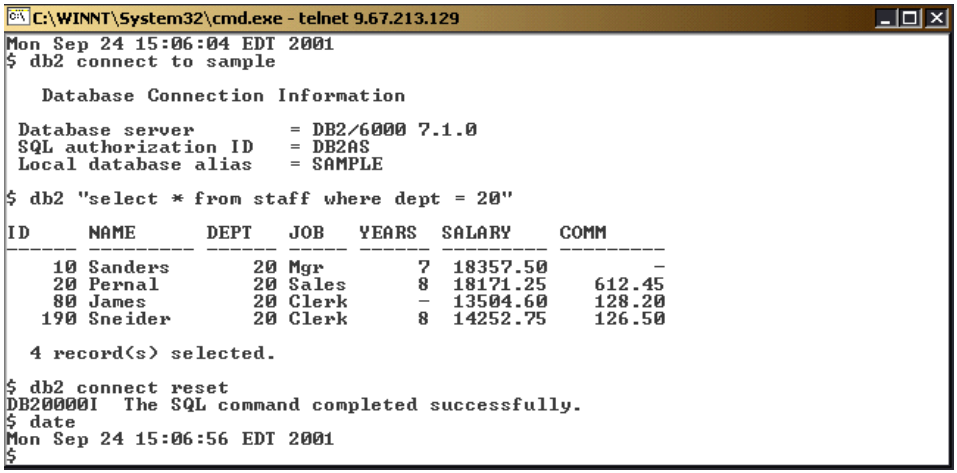

Figure 15-9 DB2 UDB query example against sample database

15.3 Implementing DB2 UDB V7.1 on RedHat Linux V7.1

First of all, we have to mount the NFS shared volume on the Linux server before we can create the database. The DB2 UDB V7.1 database is not officially supported with the database residing on the NAS shared volumes — even though we have successfully created a database on the NAS volume with a workaround.

We have to modify the Windows registry on the NAS200 server to allow DB2 UDB to create data files. The value of EnableSMBLocking has to be **0** in the registry. The full path is *\SYSTEM\Controlset001\Services\NfsSvr\NlmNsm*.

Figure 15-10 shows the entry after the value is changed to **0**. Use either **Regedit** or **Regedt32** started from a command prompt to change the value.

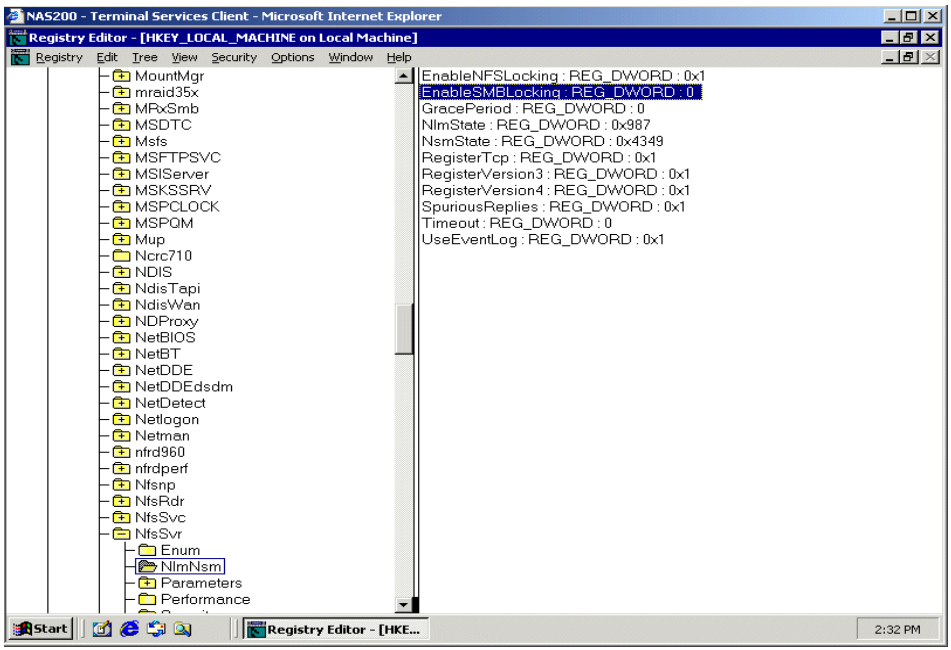

Figure 15-10 Registry entry to change the SMBLocking value

Now you can mount the shared volume on the Linux server and display the file system to see the NFS volume. Figure 15-11 shows the mount and displays the file system.

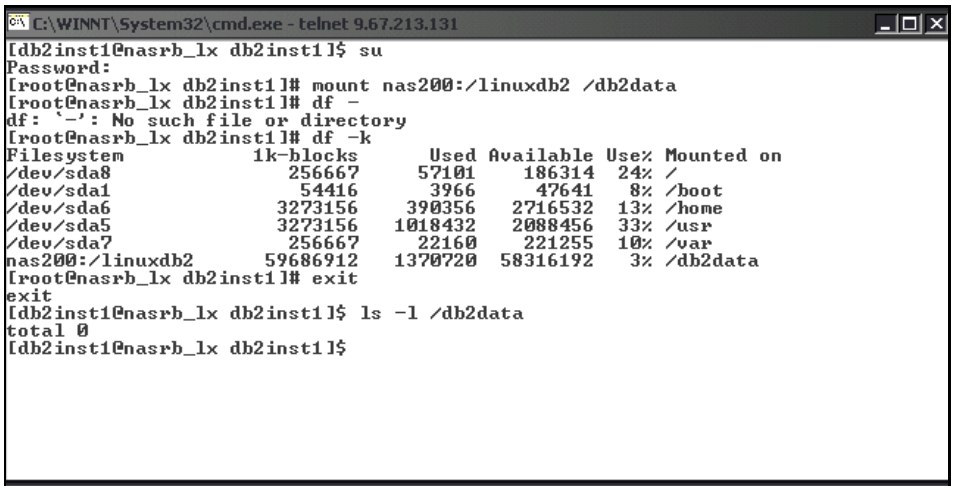

Figure 15-11 NFS mount on Linux

Now we are ready to create a database on the mounted NFS volume.

Note: Keep in mind that all *db2inst1* and *db2as* IDs must be defined on the NAS200 server for them to be able to write and create files on a shared volume. Refer to Chapter 10, “Cross platform storage” on page 353 for more details

Figure 15-12 shows how to create a database, and lists the files created by the *db2inst1* ID.

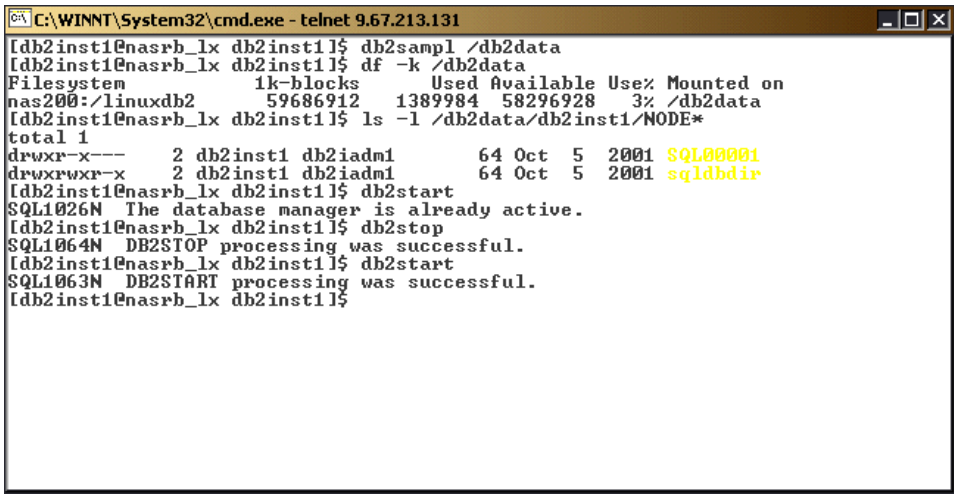

Figure 15-12 DB2 UDB sample database created

Figure 15-13 shows a query example form the databases just created.

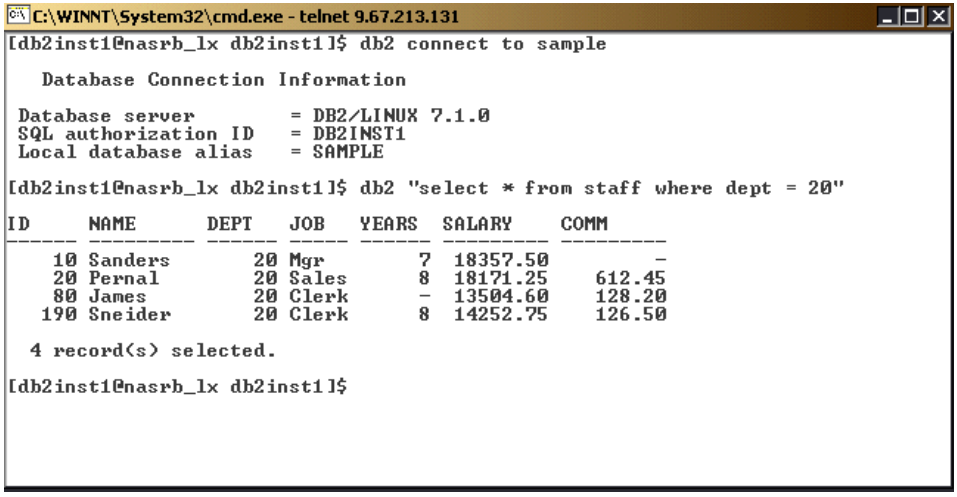

Figure 15-13 Database query from sample database

Figure 15-14 shows another sample of the working DB2 UDB database on the NAS system.

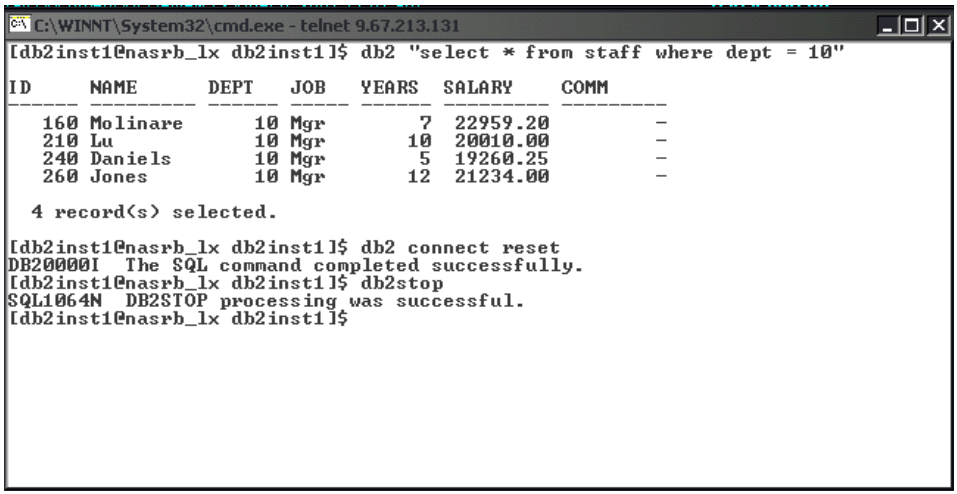

Figure 15-14 Another database query from the sample database

15.4 Implementing Oracle V8.1.7 on AIX 4.3.3

This section describes installing and running the Oracle database software. All database, control, and log files reside on NAS shared storage. We used the same NFS mount utility to mount a shared folder on AIX server. Figure 15-15 shows how the shared volume looks on AIX with the **`df -k`** (display file systems) command.

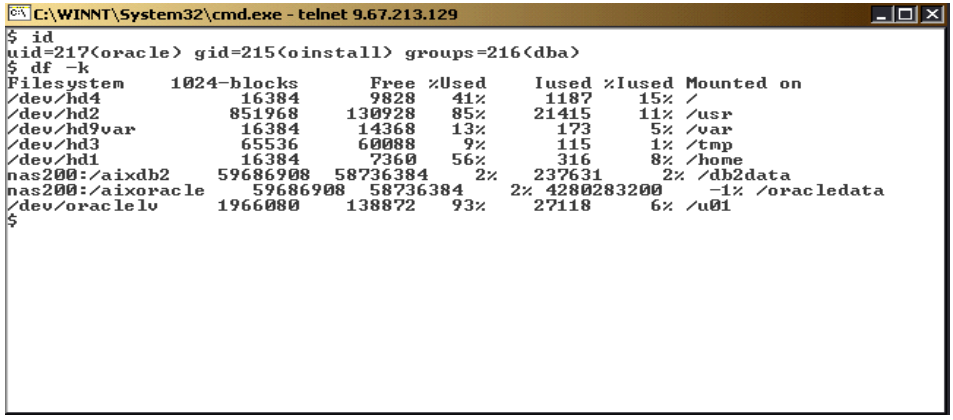

Figure 15-15 AIX — output of command

The NFS mount must have the **mount = true** value in the */etc/filesystems* file. This will mount the NFS file system at the boot time together with all local filesystems. If you do not have this value, or if it is false, the file system will not be mounted.

In the case of our example, the file system is not mounted. Therefore, the database will not start as part of the server boot process, as it will be set up in a normal business environment. Notice the entries started with *nas200:* in Figure 15-16, which represents a NAS shared volume.

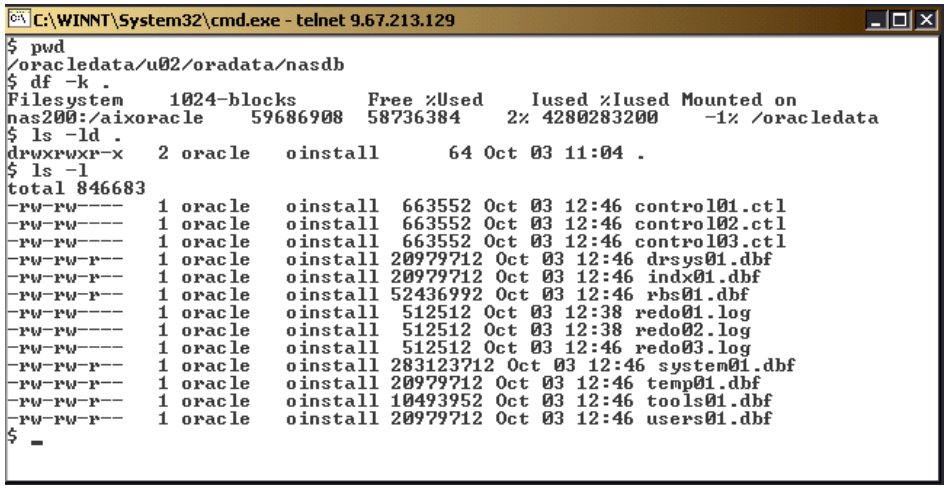

Figure 15-16 List of data, control, and log files created on NAS shared drive

Once you have successfully mounted the NFS file system, you can now create a database on that file system. We installed Oracle 8.1.7. As part of the process, we defined */oracledata* (mount point) and */oracledata/u02* (where to create data, control, and log files). We created the *u02* directory with an Oracle ID.

Make sure that owner (*oracle*) and group (*oinstall*) permissions of the directory */oracledata/u02* are correct before you can create the data files. See Chapter 10, “Cross platform storage” on page 353 for instructions on how to create a shared folder and configure Services for UNIX to allow proper user access and permissions to this shared folder.

We followed the Oracle installation guide, which is available under the download section at:

http://www.oracle.com/

Once we started the Oracle Universal Installer program, we left all the defaults as recommended in installation guide, except we enter */oracledata/u02* when asked where to create the data files. Oracle Universal Installer created all necessary files under the */oracledata/u02/* file system.

Figure 15-16 shows all files created by the Oracle Universal Installer program. Notice the owner and group IDs, which are very important, because they are displayed here as they are assigned on the NAS200 and authorized by the Services for UNIX configuration on NAS200.

Now you should have successfully created all the datafiles. You need to define all Oracle environment parameters. Use the detailed description in the post installation section of the Oracle installation guide.

You can now log on to the AIX system with an Oracle ID to start the database. If you do not have a listener process up, then you can start with the **`lsnrctl start`** command at the command prompt.

To check the listener status, enter the **`lsnrctl status`** command, which should display output as shown in Figure 15-17.

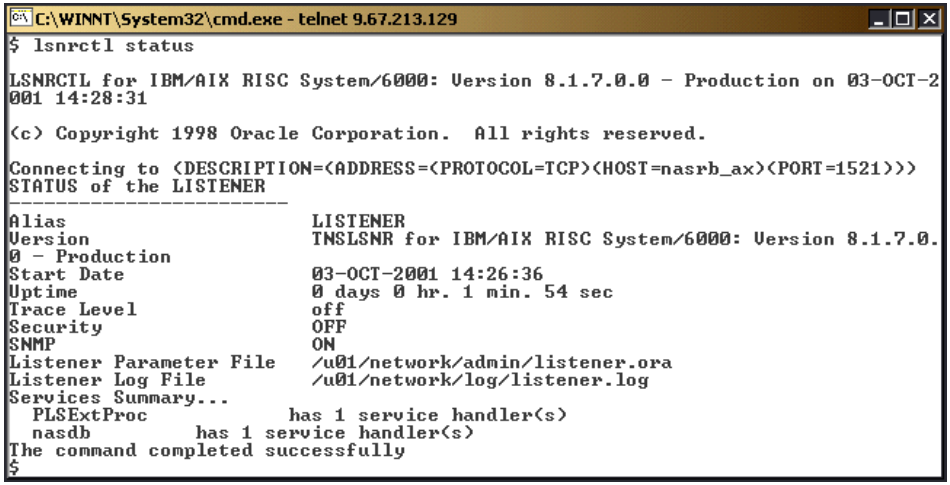

Figure 15-17 Output of lsnrctl status command

You are ready to bring up the database now. You should have all Oracle environment variables defined as part of the post install process. All initialization and configuration files should be created by the Oracle installer program.

Figure 15-18 shows the successful database startup and shutdown from all data, control and log files on NAS shared volume.

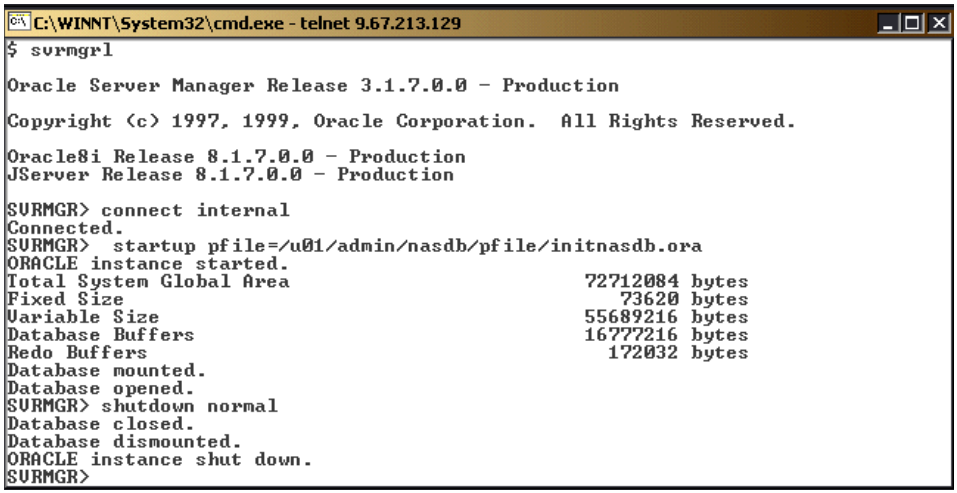

Figure 15-18 Successful database startup and shutdown on NAS

The following figures show some sample output of SQL statements from the database. Figure 15-19 shows a sample *select* statement.

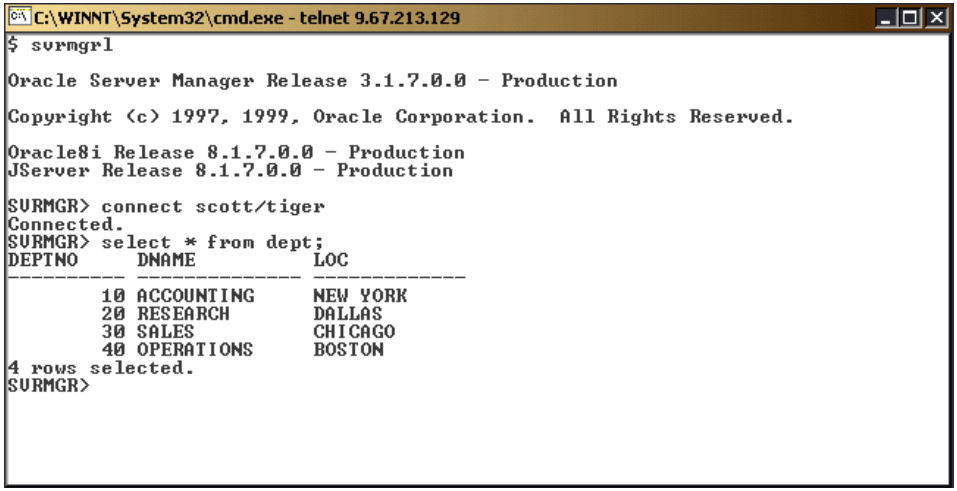

Figure 15-19 SQL output of select statement 1

Figure 15-20 shows another sample output we generated for the database on the NAS system.

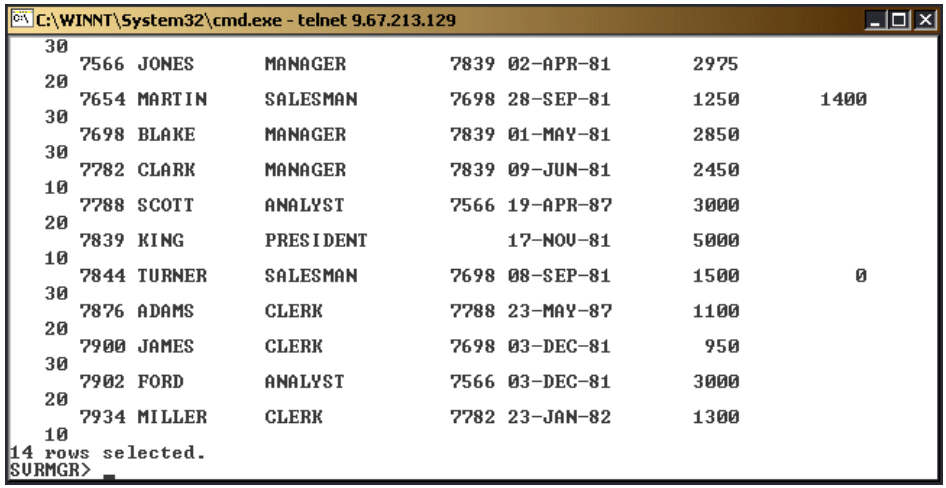

Figure 15-20 SQL output of select statement 2

Part 5

Appendixes

In this part of the book we provide the following supplementary information:

A

Ethernet adapter teaming

In this appendix we describe the steps to enable adapter teaming on the Intel Ethernet adapters (Gb Ethernet SX and 10/100 Ethernet) as well as the Alacritech Ethernet adapters supported in the IBM TotalStorage NAS devices.

Intel adapters

The Intel Ethernet adapters (Gb Ethernet SX and 10/100 Ethernet) that you install in the PCI slots support adapter teaming. With adapter teaming, two or more PCI Ethernet adapters can be physically connected to the same IP subnetwork and then logically combined into an adapter team. Such a team can support one of the following functional modes:

Fault tolerance
Only one adapter in the team is fully active on the Ethernet network (for example, sending and receiving data) at any point in time, while the other adapters are in standby mode (receiving data only). If that adapter detects a link failure or fails completely, another adapter in the team automatically and rapidly takes over as the active adapter, and all Ethernet traffic being handled by the failing adapter is switched to the new active adapter, with no interruption to network sessions (for example, file transfers) in progress at the time of the failover.

Load balancing
All adapters in the team are active, increasing the total transmission throughput over the common IP subnetwork. If any adapter in the team fails (link failure or complete failure), the other adapters in the team continue to share the network transmission load, although total throughput is decreased. Load balancing is only supported for adapter teams consisting of only one type of adapter; different types of adapters cannot be combined in a load-balancing team.

Note: It is strongly recommended that you configure adapter teaming before you set up Microsoft Cluster Server (MSCS) clustering. Additionally, for each team you configure on one node, you must configure an identical team (same type of team, same set of adapters, and so on) on the other node.

To configure adapter teaming, use Intel PROSet II tool, which is preloaded on the NAS device follow these steps:

1. Physically connect the adapters that you want to team to the same IP subnetwork.
2. Access the NAS device desktop by directly attaching a keyboard, mouse, and monitor, or over the network by starting Terminal Services on another workstation.
3. From the NAS device desktop, go to the **Start —> Settings —> Control Panel**.
4. Double-click the Intel PROSet II icon in the Control Panel to start Intel PROSet II. You will see a list of all adapters for each slot and type supported under Network Components.

5. Under Network Components, you will see a list of installed and non-installed adapters for each slot and type supported. Drivers are preset for all supported adapter configurations but will be loaded only for installed adapters.
6. Identify which adapters you are going to team. Left-click the adapter under Network Components, and select one of the adapters that will be part of the teaming.
7. Right-click the adapter, select **Add to Team —> Create New Team**.
8. Select the type of team to create.
9. Select the adapters to add to the team from the list, and then select **Next**.
10. Verify that these settings are correct, and then select **Finish**.
11. If you are configuring a clustered NAS device as a NAS300 or a NAS 300G, perform steps 1 to 10 for the other node.

This procedure creates a device named *Intel Advanced Network Services Virtual Adapter*. It also binds all network protocols that were bound to the physical adapters that were added to the team to this virtual adapter, and unbinds those protocols from the physical adapters. If you delete the team, the settings will return to the state prior to creating the team.

For complete help on adapter teaming, from Intel PROSet II, click **Network Components —> Help**.

Note: The integrated Ethernet controller on each Model 326 node is dedicated to the clustering interconnection between it and another node and cannot be used for teaming.

Alacritech Quad-Port 10/100 Ethernet Adapter

The Quad-Port 10/100 Ethernet Adapter supports the Cisco Fast EtherChannel (FEC) technology across the four ports on the adapter.

Fast EtherChannel (FEC) is a proprietary technology developed by Cisco. With FEC, you can create a team of two to four ports on the 4-port 10/100 Ethernet PCI Adapter in to increase transmission and reception throughput. The FEC might also be referred to as load balancing, port aggregation, or trunking.

When you configure this feature, the adapter ports comprising the FEC team or group create a single high-speed, fault-tolerant link between the engine and the Ethernet switch sharing one IP address. With FEC, fault tolerance and load balancing is provided for both outbound and inbound traffic, unlike other load-balancing schemes that only balance outbound traffic.

Note: FEC requires an Ethernet switch with FEC capability. The FEC implementation on the 4-port 10/100 Ethernet Adapter does not support the optional Port Aggregation Protocol (PAgP) feature of FEC-capable Ethernet switches.

To configure the adapter, follow these steps:

1. Click **Control Panel**.
2. Select **Network** and **Dial-Up**.
3. Select **Adapter**.
4. Select **Properties**.
5. Select **Alacritech SLIC Team Configurator**.
6. Select **New Team**.

B

NAS 100 troubleshooting and BIOS access

This appendix covers advanced topics for the NAS 100 appliance. We offer some help for troubleshooting and show how to access the BIOS.

Troubleshooting

In this section we provide basic troubleshooting information to help you resolve common problems that might occur with your appliance. The following tools are available to help you identify and resolve hardware-related problems:

- **Error messages:** Error messages and logs indicate successful test completion or the detection of a problem. See “Error messages” on page 566 for more information.
- **Temperature checkout:** Cooling of the system is important for correct operation and system reliability. See “Temperature checkout” on page 567 for more information.
- **LEDs:** The LEDs help you identify problems with appliance components. See “Identifying problems using LEDs” on page 567 for more information.
- **SCSI adapter test:** This procedure tests the SCSI adapter of your appliance. See “SCSI adapter test” on page 573 for more information.
- **Accessing the BIOS:** Corrupted CMOS data can prevent the appliance from booting or connecting to the network. Resetting the BIOS to its default value can correct the problem. Because of the complexity and importance of the BIOS access we created a separate section for it. See “Accessing the BIOS” on page 576.
- **Maintenance tools:** These commands are for maintenance purposes only. See “Maintenance tools” on page 575 for more information.

Error messages

All error messages and logs are recorded in the Windows event log.

Note: If you receive a critical-level message, you must take action. If the critical message reports an HDD failure, see “Hard disk drive LED problem determination” on page 571. For all other critical-level messages, call the IBM Support Center.

Perform the following steps to access hardware errors in the Windows event log by means of Terminal Services:

1. From the Start menu in the PC task bar, click **Program —> Administrator Tools —> Event Viewer**.
2. From the Tree panel on the left, click **Application Log**.
3. Click **View —> Filter**.

4. Type **NAS100Svc** in the Event Source field and click **OK**. The Windows event log opens.

Temperature checkout

Cooling of the system is important for correct operation and system reliability. Make sure that:

- Each of the drive bays has a drive installed.
- The top cover is in place during normal operation.
- There is at least 50 mm (2 in.) of ventilated space at the sides of the appliance and 100 mm (4 in.) at the back of the appliance.
- A removed hot-swap drive is replaced within 2 minutes of removal.
- The fans are operating correctly and the air flow is good.

Identifying problems using LEDs

The NAS 100 has three types of LEDs to help you identify problems with hardware failure and some software errors and system status. Two types of LEDs are located on the front bezel and the third type is located on the back of the chassis:

- System status LEDs are located on the operator information panel near the power button.
- Hard disk drive (HDD) status LEDs are located along the bottom of the bezel.
- Ethernet port status LEDs are located on top of the Ethernet ports (LAN 1 and LAN 2).

Table B-1 on page 568 describes the system status LEDs, Table B-2 on page 570 describes the hard disk drive status LEDs, and Table B-3 on page 571 describes the Ethernet port status LEDs.

Operator panel LEDs

System status LEDs are located on the operator information panel near the power button.

Table B-1 Operator panel LEDs

Symbol	Function	Color	Description	Action
	Power	Green steady	AC and DC power on	No action is necessary.
		Green blinking	AC Power on and DC power off	System on standby. Push the power button to switch on the system.
		No light	AC power unavailable	Connect the power cable.
	LAN 1	Green steady	Link OK.	No action is necessary.
	LAN 2	Green steady	Link OK.	No action is necessary.

Symbol	Function	Color	Description	Action
	Warning	Amber steady	If a Hard Disk Drive Status LED is also amber steady, that hard disk drive has failed.	Replace failed hard disk drive.
			If no Hard Disk Drive Status LED is amber steady, a system component has failed.	Replace the NAS100.
		Amber blinking, 4/sec. cycle	The port is configured to use DHCP and failed to retrieve an IP address from the DHCP server.	Contact the network administrator and report that the IP address could not be retrieved.
			The port is configured to use static IP and is set to 192,168.0.1.	Configure a new IP address that is different from 192.168.0.1. Refer to Chapter 4, “Implementing the IBM TotalStorage NAS 100” on page 57
		Amber blinking, 1-sec. cycle	Network device driver error or BIOS failed to boot.	Remove HDDs 1 and 2 and reboot. If system reboots, run system recovery. If system does not reboot, see “Hard disk drive LED problem determination” on page 571

Hard disk drive LEDs

HDD status LEDs are located on the front of the chassis, along the bottom of the bezel.

Table B-2 Hard disk drive LEDs

Location	Function	Color	Description	Action
Left	Status	Green steady	OK.	No action is necessary.
		Green/amber blinking, 4/sec., on all drives in array	RAID rebuilding.	No action is necessary.
		Green/amber blinking, 1-sec. cycle	OS booting from an HDD other than HDD1.	See “Hard disk drive LED problem determination” on page 571
		Green/amber blinking, 2-sec. cycle, on the three drives that need to be rebuilt	RAID needs to be rebuilt to provide fault tolerance.	Rebuild the RAID. Refer to 12.1, “Recovering the NAS 100” on page 454
		Amber steady	Hard disk drive failure.	See “Hard disk drive LED problem determination” on page 571
Right	Access	Amber steady	System is accessing the HDD for data.	No action is necessary.

Ethernet port status LEDs

Ethernet port status LEDs are located on the back of the chassis, on top of the Ethernet ports (LAN 1 and LAN 2).

Table B-3 Ethernet port LEDs

Location	Function	Color	Description	Action
Left	Connection	Amber steady	A valid LAN connection exists.	No action is necessary.
Right	Activity	Green blinking	System is sending or receiving network data.	No action is necessary.

Hard disk drive LED problem determination

All HDDs have the capability to boot the operating system by using the OS boot failover mechanism. When the Status LED of an HDD other than HDD1 is blinking green/amber on a 1-second cycle, the OS is booting from that HDD.

During normal operation, HDD1 is used for the system boot. If HDD1 fails to boot after three tries, the booting device is switched automatically to HDD2 (mirrored system volume). If HDD2 also fails to boot after three tries, the recovery system on HDD3 boots. To determine what actions you must take to correct the problem, refer to Table B-4.

Note:

1. Before replacing and rebuilding an HDD:
   a. If possible, back up all user data.
   b. Clear the CMOS data (see “Clearing CMOS data” on page 25).
2. If you need to order a replacement HDD, contact your IBM sales representative or your place of purchase.
3. If you appliance displays any other combination of LEDs, call the IBM Support Center.

Table B-4 HDD LED problem determination

Warning LED	HDD1 Status LED	HDD2 Status LED	HDD3 Status LED	HDD4 Status LED	Action
Off	Green	Green	Green	Green	No action is necessary
Off	Green	Green	Green/amber blinking, 1-sec. cycle	Green	Run system recovery.
Off	Green	Green	Green	Green/amber blinking, 1-sec. cycle	Call the IBM Support Center.
Amber steady	Amber	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Replace HDD1.
Amber steady	Green/amber blinking, 2-sec. cycle	Amber	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Replace HDD2.
Amber steady	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Amber	Green/amber blinking, 2-sec. cycle	Replace HDD3.
Amber steady	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Amber	Replace HDD4.
Amber steady	Amber	Amber	Green/amber blinking, 1-sec. cycle	Green/amber blinking, 2-sec. cycle	Call the IBM Support Center.
Amber steady	Green/amber blinking, 2-sec. cycle	Amber	Green/amber blinking, 1-sec. cycle	Green/amber blinking, 2-sec. cycle	Replace HDD2.
Amber steady	Amber	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 1-sec. cycle	Green/amber blinking, 2-sec. cycle	Replace HDD1.
Amber steady	Amber	Amber	Green/amber blinking, 1-sec. cycle	Amber	Call the IBM Support Center. See **Note**.

Warning LED	HDD1 Status LED	HDD2 Status LED	HDD3 Status LED	HDD4 Status LED	Action
Amber steady	Green/amber blinking, 2-sec. cycle	Amber	Green/amber blinking, 1-sec. cycle	Amber	Replace HDD2 and HDD4. See **Note** below.
Amber steady	Amber	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 1-sec. cycle	Amber	Replace HDD1 and HDD4. See **Note** below.
Amber steady	Green/amber blinking, 2-sec. cycle	Green/amber blinking, 2-sec. cycle	Amber	Green/amber blinking, 1-sec. cycle	Call the IBM Support Center.
Amber blinking, 1-sec. cycle	Amber	Amber	Amber	Amber	Call the IBM Support Center.
Amber blinking, 1-sec. cycle	Green	Green	Green	Green	Call the IBM Support Center.

Note: Your user data is already lost at this point.

SCSI adapter test

Before starting this procedure, check the following hardware and make sure that:

- The SCSI adapter is physically installed in the PCI expansion slot.
- The SCSI cable is correctly connected between the SCSI adapter connector and the tape device.
- The tape device is powered on and no error shows on the tape device's control panel.

Now proceed as follows:

1. Access the NAS 100 by means of Windows Terminal Services.
2. Click **Start —> Settings —> Control Panel**. The Control Panel window opens.
3. Double-click **System**. The System Properties window opens.

4. Click the **Hardware** tab, and then click **Device Manager**. The Device Manager window opens.
5. Double-click **SCSI and RAID controllers**.
6. If you:
   - Are able to locate **Symbios Ultra3 PCI Adapter; 53C1010-66 Devic**e, continue with step 7.
   - Cannot locate **Symbios Ultra3 PCI Adapter; 53C1010-66 Device**, go to step 10.
7. Make sure that there is no error mark on **Symbios Ultra3 PCI Adapter; 53C1010-66 Device**. If there is, either the SCSI adapter or the system board is defective. Contact your IBM service representative for further action.
8. If you want more information, double-click **Symbios Ultra3 PCI Adapter; 53C1010-66 Device**, click the **General** tab, and check the message **This device is working properly.** in the Device Status window.

Note: The location of this device must be PCI bus 1, device 13, and function 0

9. Go to step 12.
10. Click **SCSI and RAID controllers**, and then click **Action —> Scan for hardware changes** to install the correct device driver.
11. If the device is not working correctly, click the device and click **Action —> Enable** to enable the device. If the device still does not work, either the SCSI adapter or the system board is defective. Contact your IBM service representative for further action. Otherwise, continue with step 12.
12. In the Device Manager window, double-click **Tape drives**.
13. If you:
    - Locate **IBM XXXXXX-XXX SCSI Sequential Device**, where *XXXXXX-XXX* is the name of the tape device, continue with step 14.
    - Cannot locate **IBM XXXXXX-XXX SCSI Sequential Device**, go to step 17.
14. Make sure that there is no error mark on **IBM XXXXXX-XXX SCSI Sequential Device**. If there is, either the SCSI adapter, the SCSI cable, or the tape device is defective. Contact your IBM service representative for further action.
15. If you want more information, double-click **IBM XXXXXX-XXX SCSI Sequential Device**, click the **General** tab, and check the message **This device is working properly.** in the Device Status window.

16. Go to step 19.
17. Click **SCSI and RAID controllers**, and then click **Action —> Scan for hardware changes** to install the correct device driver.
18. If the device is not working correctly, click the device and click **Action —> Enable** to enable the device. If the device still does not work, either the SCSI adapter, the SCSI cable, or the tape device is defective. Contact your IBM service representative for further action.
19. If both devices are working correctly, the SCSI adapter is working normally. If the hardware tests complete successfully, but the problem persists during normal appliance operations, a software error might be the cause. If you suspect a software problem, refer to the *IBM TotalStorage NAS 100 User's Referenc*e, included on the documentation CD.

Maintenance tools

Use the following commands for maintenance purposes only.

bootchg command

This command changes the boot order of the HDDs on Windows Powered OS.

1. Access the NAS 100 by means of Windows Terminal Services.
2. Click **Start —> Program —> Accessories —> Command Prompt**. The Command Prompt window opens.
3. At the prompt, type **`cd \ibm\nas100`** and press Enter.
4. Type **`bootchg [r] [0-3] [q]`**, where:
   - r is read current boot order
   - 0-3 is set boot disk (0: HDD1 - 3: HDD4)
   - q is quiet mode

   Now press Enter.

configCMOS command

This command sets the PXE boot to enable or disable on Windows Powered OS.

1. Access the NAS 100 by means of Windows Terminal Services.
2. Click **Start —> Program —> Accessories —> Command Prompt**. The Command Prompt window opens.
3. At the prompt, type **`cd \ibm\nas100`** and press Enter.
4. Type **`configCMOS pxe=[e/d/r] [q]`**, where:
   - e is set PXE boot enable
   - d is set PXE boot disable
   - r is read current PXE boot setting

– q is quiet mode

Now press Enter.

Accessing the BIOS

In some cases it will be important to access the BIOS of the NAS100. Because it is a headless device, the only way to access the BIOS is via RS232 (COM) and HyperTerminal. In this section we provide helpful information on how to handle the BIOS and configure it to your needs.

Clearing CMOS data

Clearing the CMOS data resets all the system counters, so usually this is not recommended. However, if the CMOS data in the NAS 100 is corrupted, the appliance might not boot up or connect to the network. In this case, resetting the CMOS data can solve the problem. You can also use this procedure to reset the supervisor password if it is lost.

The Clear CMOS button is located in a small hole to the right of the LAN 2 port on the back of the chassis. You can press it by using a pointed implement, such as a paper clip. See Figure B-1.

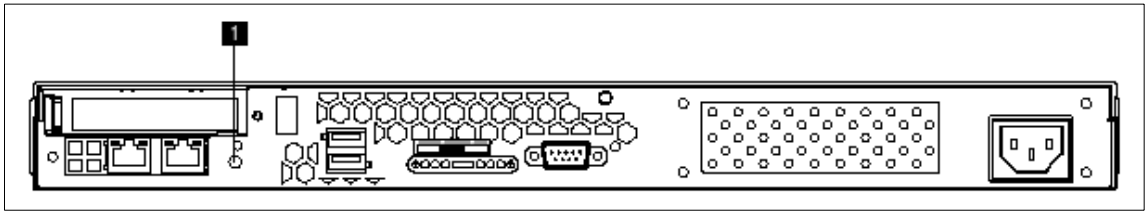

Figure B-1 Clear CMOS button (1)

To reset the CMOS data:

1. If the appliance is rack-mounted:
   a. If necessary, remove it from the rack and place it on a level, stable surface

Attention: Before removing the appliance from the rack:

a. Power off the appliance

b. Disconnect the power cord and all cables from the appliance.

   c. Reconnect the appliance.

d. Power on the appliance.

2. Simultaneously press both the Power button on the front of the chassis and the Clear CMOS button on the back of the chassis for one second.
3. Release the Power button, but keep pressing the Clear CMOS button until a beep sounds (about 10 seconds). This clears the CMOS data and reloads the BIOS to its default value.

Note: If you press the Power button for more than four seconds, the appliance will power off.

4. After you clear the CMOS data, reset the date, time, and, if appropriate, the supervisor password in the BIOS. See “Security” on page 595.

Note: When you clear CMOS data, the supervisor password is reset to the default of **001san** (not case-sensitive).

Preparing to use the remote BIOS setting function

The NAS 100 is a “headless” appliance; it does not have a keyboard, mouse, or monitor directly attached to the appliance. Therefore, to access the BIOS you must have a PC loaded with Windows 2000 and ServicePack2.

1. Power off your appliance.
2. Connect the COM port of the PC to the serial port of the appliance using a 9-pin serial cable (RS-232C female crossover cable). See Table B-5 for the signal and pin assignments.

Table B-5 RS-232C female crossover cable

Pin number	Signal name
1	Data carrier detect
2	Receive data
3	Transmit data
4	Data terminal ready
5	Signal ground
6	Data set ready
7	Request to send
8	Clear to send
9	Ring indicator

3. Open HyperTerminal on your PC by clicking **Start —> Program —> Accessories —> Communications —> HyperTerminal**. The HyperTerminal start window (Figure B-2) will appear.

Figure B-2 HyperTerminal start screen

4. After that the Connection Description window opens (Figure B-3).

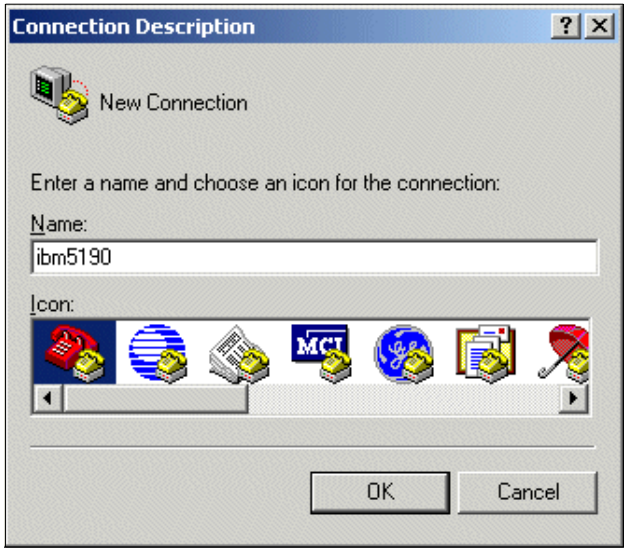

Figure B-3 Connection description window

5. Type a name for the connection and then click **OK**. The Connect To window opens (Figure B-4).

Figure B-4 Connect o window

6. Click the COM port that has the serial cable attached **(COM1-4)** in the Connect using field and click **OK**. The COM Properties window opens (Figure B-4).

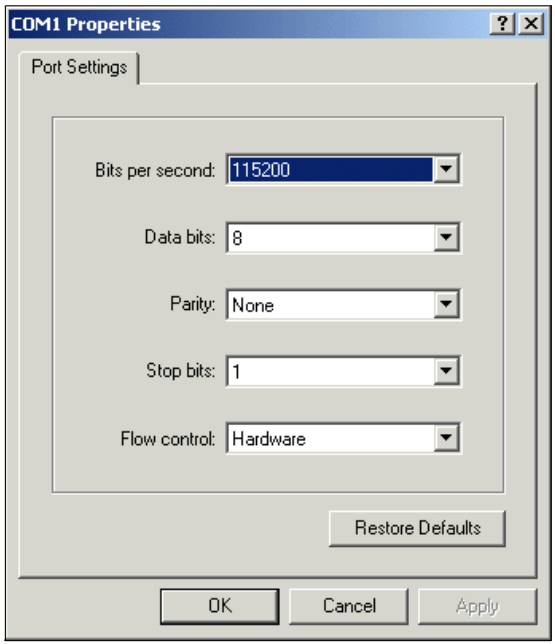

Figure B-5 COM properties window

7. Select the following values in the Properties window:
   - Bits per second: 115200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: Hardware

   Now click **OK**. The HyperTerminal window opens (Figure B-6).

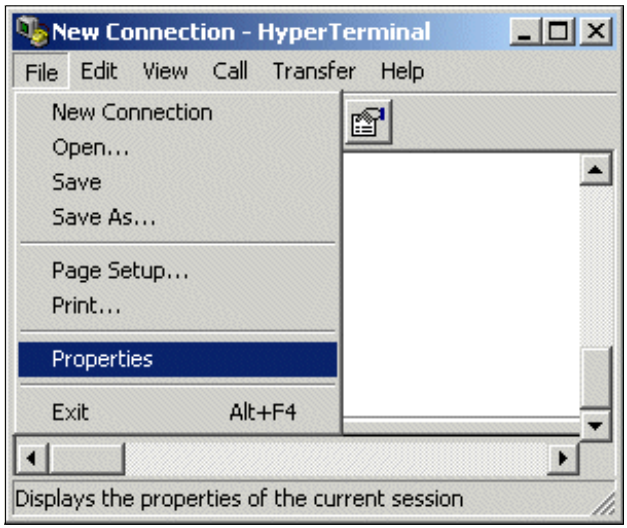

Figure B-6 HyperTerminal window with pull down menu to properties

8. Click **File —> Properties**. The Properties window opens () (Figure B-7).

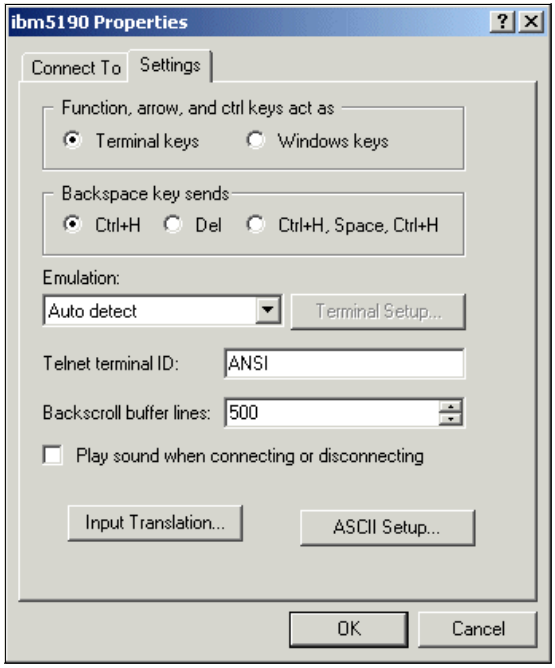

Figure B-7 Property window

9. Enable the function keys of HyperTerminal by clicking the **Settings** tab, making sure that the radio button for the **Terminal keys** option is selected, and then click **OK**. The HyperTerminal window opens again.

10. Click **Call —> Wait for a call** (Figure B-8).

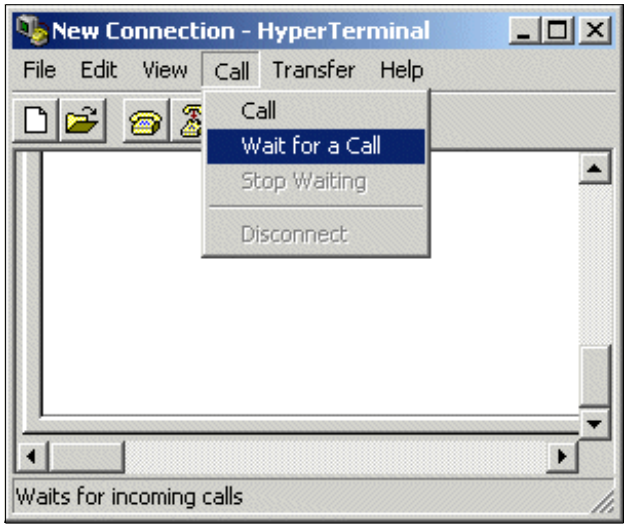

Figure B-8 Wait for call tab

11. If you want to:
    - Upgrade the BIOS, go to step 4 on page 603.
    - Update the BIOS, continue with step12.

12. Power on the appliance.

13. BIOS POST messages should appear on the HyperTerminal screen (Figure B-9).

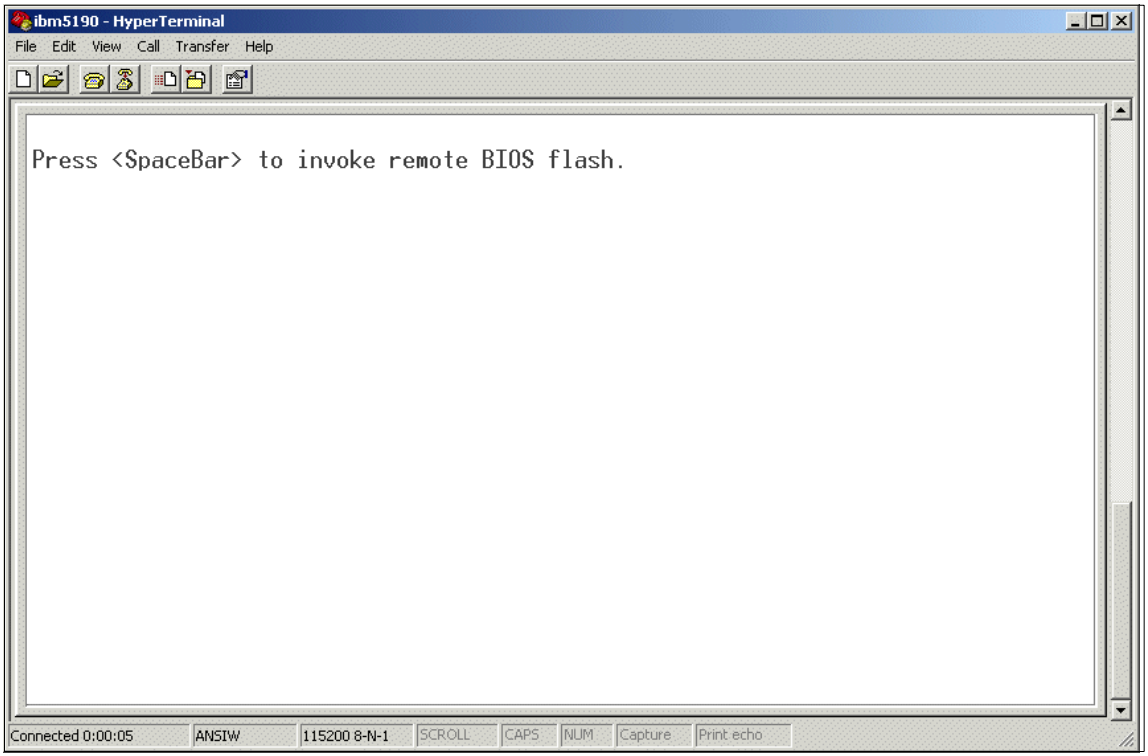

Figure B-9 Remote BIOS flash access

14. When the message **Press F1 to enter setup** opens, press F1 (Figure B-10).

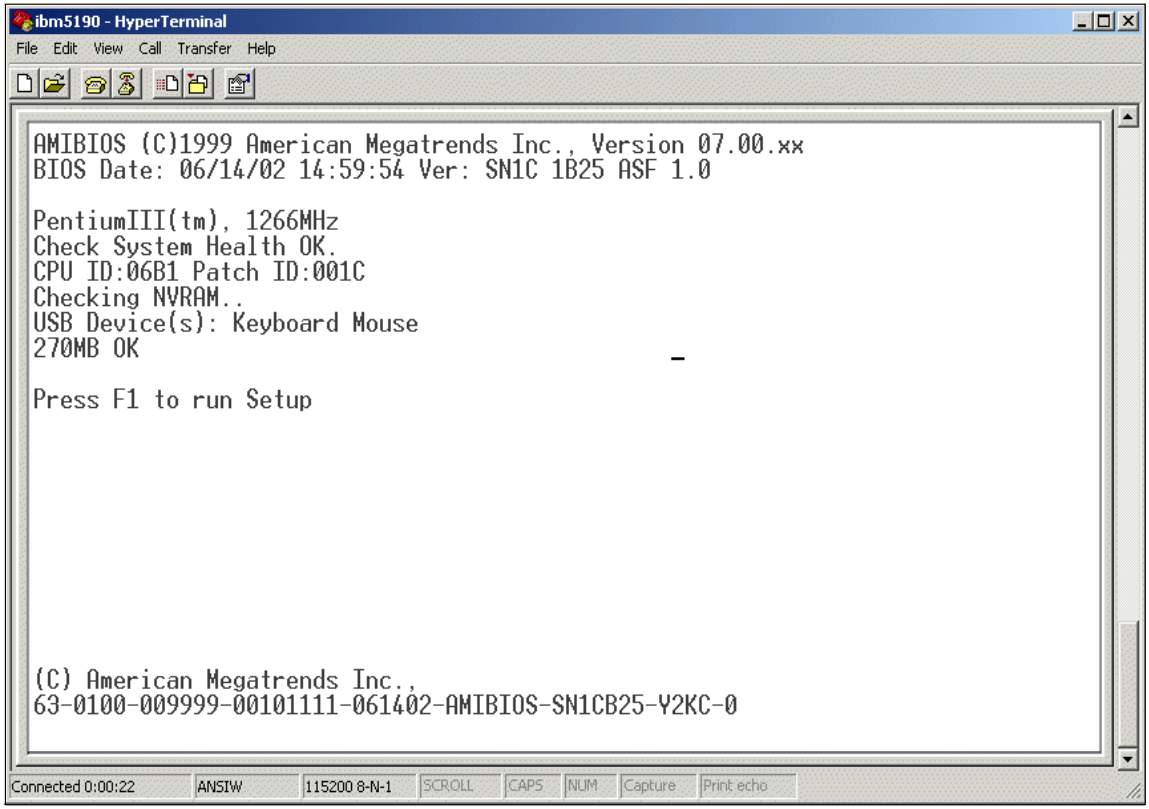

Figure B-10 F1 to enter setup

Note: If the messages do not appear:
- Check the settings and try again.
- Terminate the HyperTerminal program start the procedure again.
- Ensure that the serial cable meets the correct specifications.
- Ensure that Windows 2000 is correctly installed.

15. The appliance will detect its hard disks (Figure B-11).

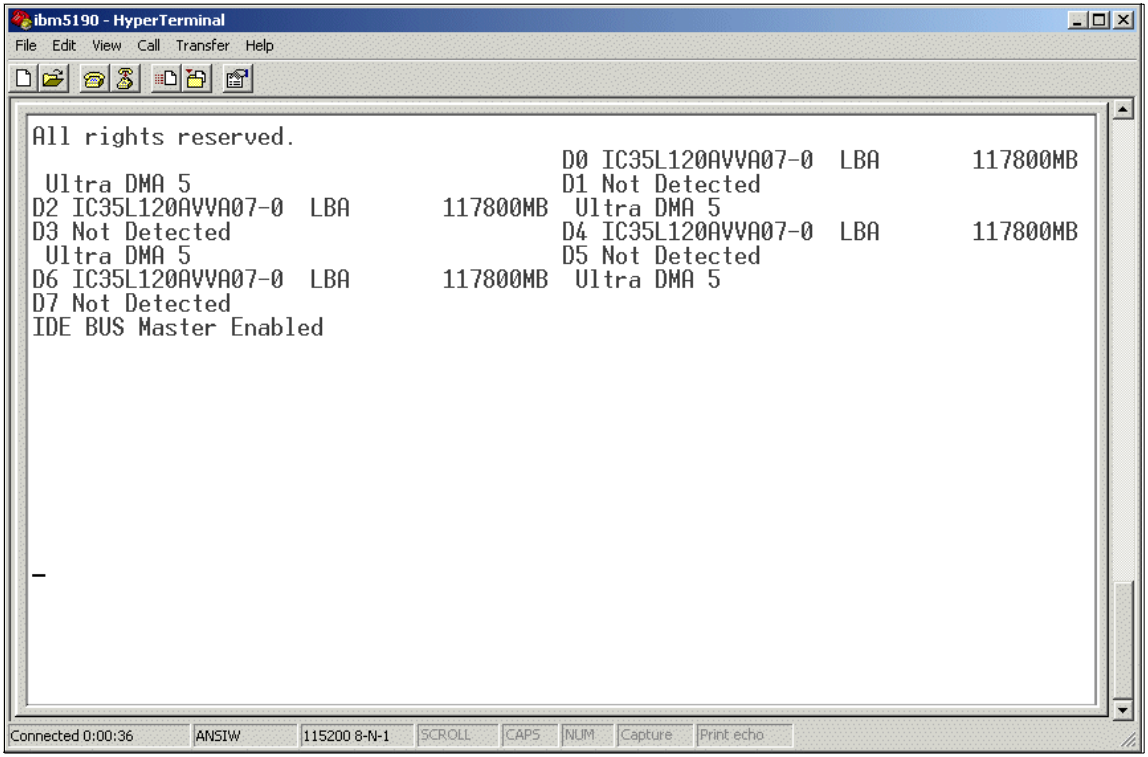

Figure B-11 Detected disks

16.Now the ethernet adapters are shown (Figure B-12).

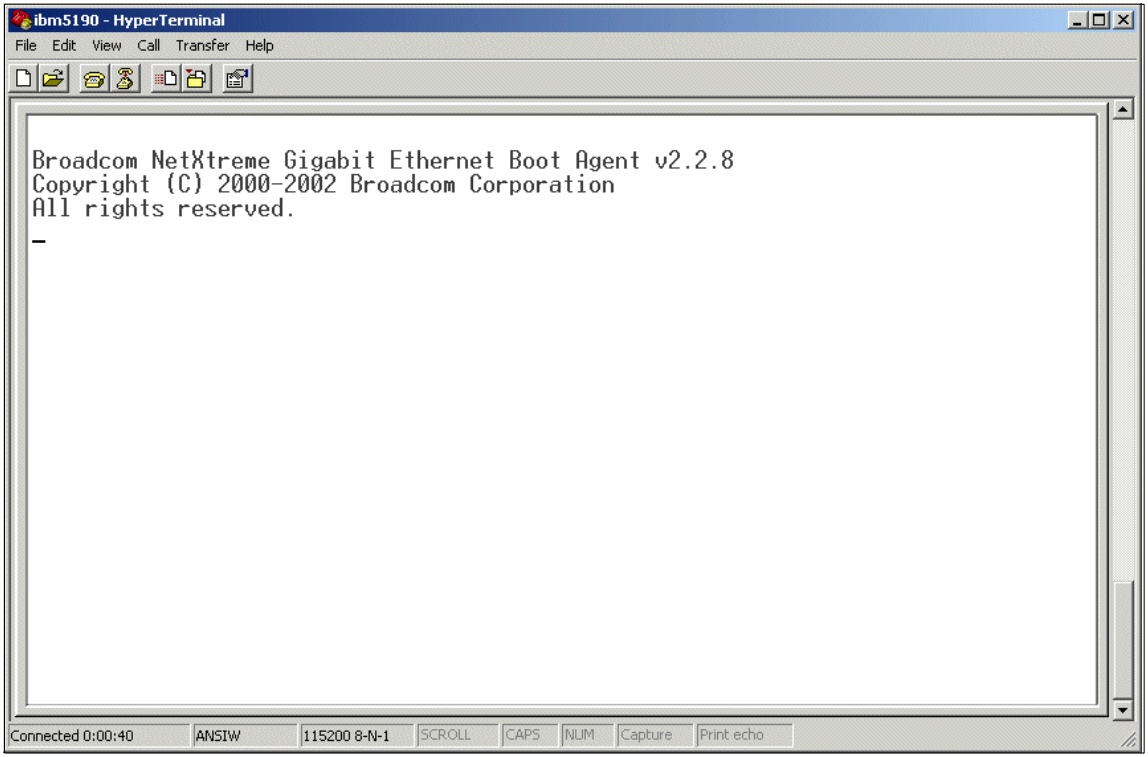

Figure B-12 Detecting ethernet adapters

17. The message **Enter current password** opens (Figure B-13). Type the password and press Enter.

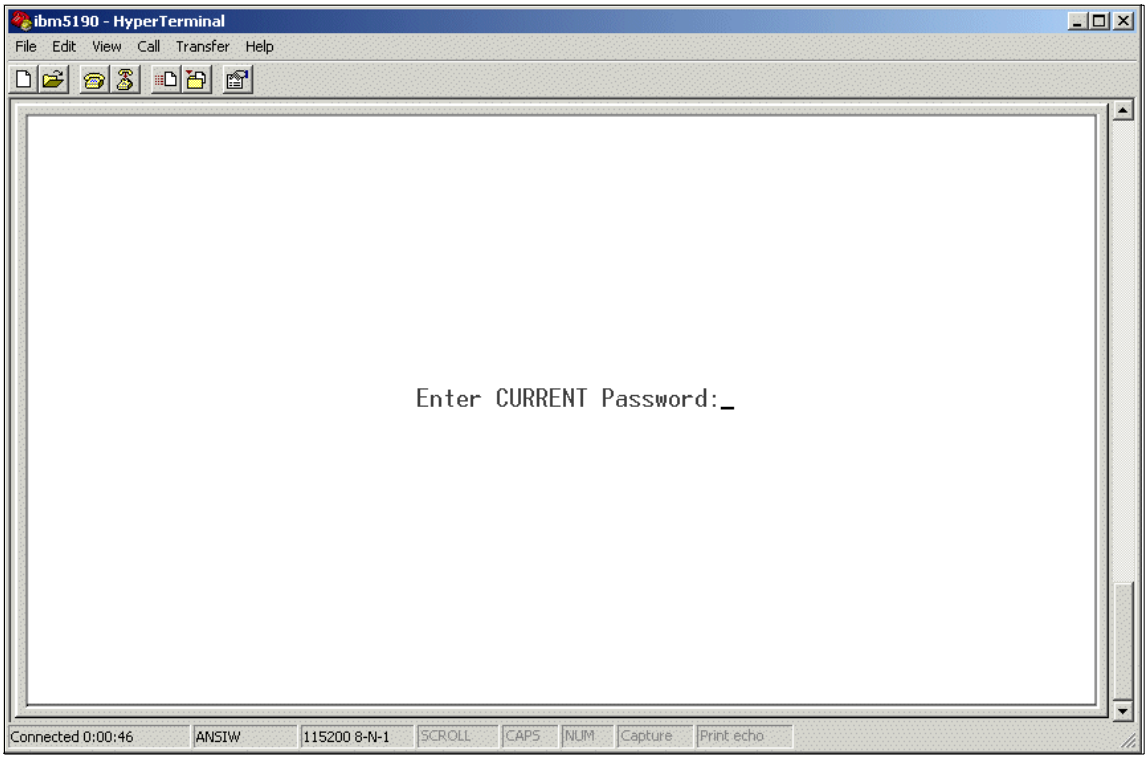

Figure B-13 Entering Setup password

Note: The default password is **001san** (not case-sensitive).

18. The BIOS setup window opens (Figure B-14).

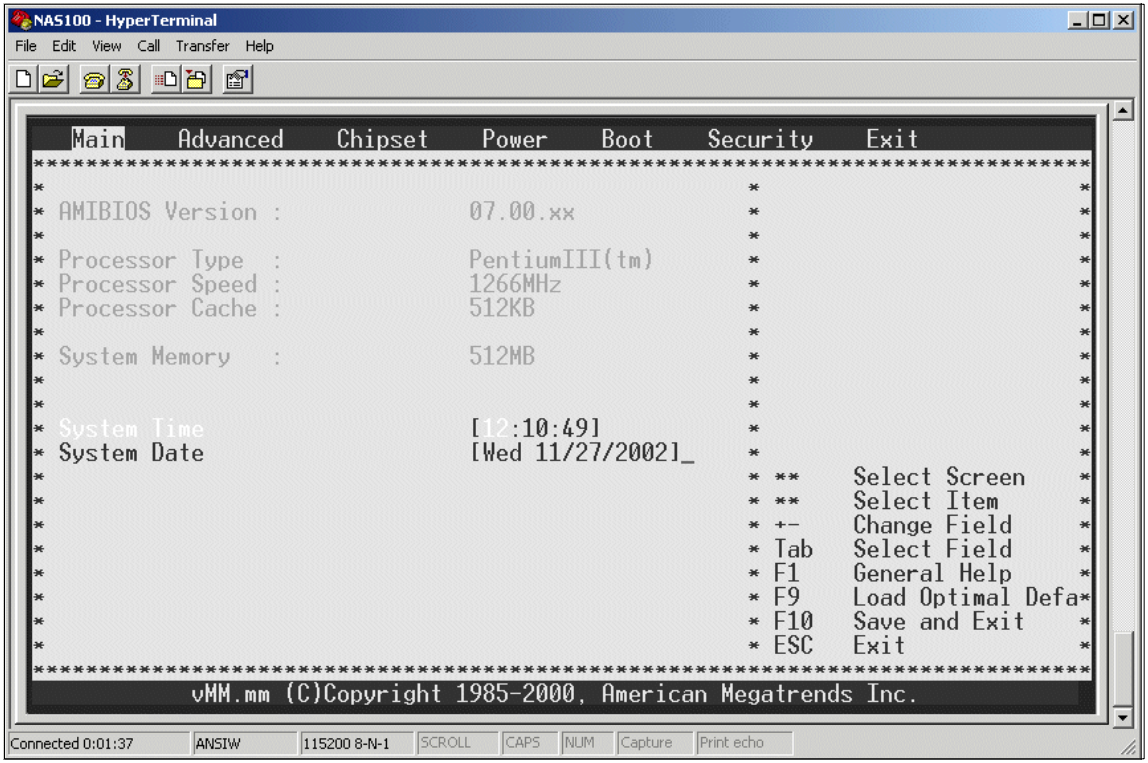

Figure B-14 BIOS Main screen

Note: In the BIOS setup window, the following function keys are disabled:
- F7
- F8
- F9
- F10
- Page Up
- Page Down

19. Continue with “Making changes to the BIOS”.

Making changes to the BIOS

Use the right and left arrow keys to move the cursor between the tabs.

1. Select one of the following tabs in the BIOS setup window.
   - Main
   - Advanced
   - Chipset/Power/Boot
   - Boot
   - Security
   - Exit

Attention: Be very careful when changing BIOS settings. Incorrect settings might cause the system to malfunction. Change only the following items.

Main

The Main window displays some system information and provides the means to set Date/Time. When you have cleared the CMOS data (see “Clearing CMOS data” on page 576), you must reset Date/Time.

Setting System Time

1. Type the appropriate numbers in the **Time**, **Minutes**, and **Seconds** fields. You can use the **Tab** key to move through the fields. You can modify the values in white.
2. Use the up and down arrow keys to move the cursor between the System Time and System Date fields.

Setting System Date

1. Type the appropriate numbers in the **Date** (Month), **Day**, and **Year** fields. You can use the **Tab** key to move through the fields. You can modify the values in white.

Note: Use two-digit numbers for the **Date** (Month), and **Day** fields a four-digit number for the **Year** field (for example, 04/12/2003).

2. Press Enter. The day of the week field is set automatically.

Advanced

1. Use the arrow keys to move the cursor to **VPD Data Configuration** (Figure B-15).

Attention: Although the Advanced window displays six items that can be changed, do not select any configuration except VPD Data Configuration.

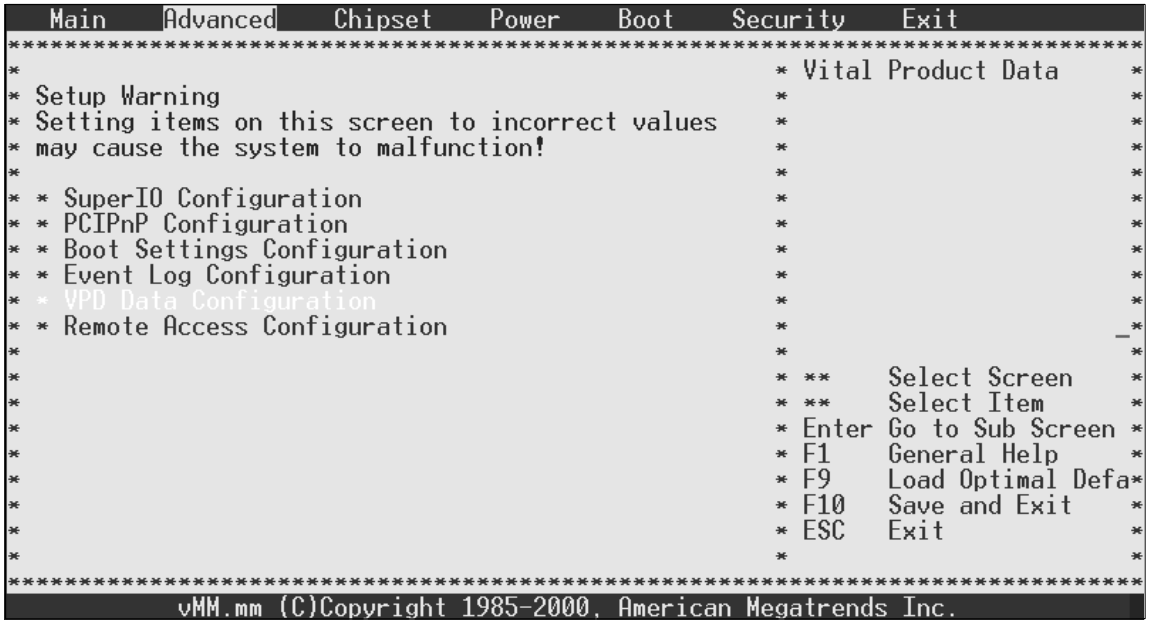

Figure B-15 BIOS advanced screen

The VPD Data Configuration window opens, displaying the following VPD data:

- BIOS build date
- BIOS ID
- Machine type
- Model name
- Machine serial number
- Original machine serial number
- System part number
- Motherboard serial number
- Universal Unique ID (UUID)

You can change the machine serial number.

Attention: Do not change any other settings in this window.

Changing Machine Serial Number:

1. Using the arrow keys, move the cursor to **Change Machine Serial Number** and press Enter. The Change Machine Serial Number window opens (Figure B-16).

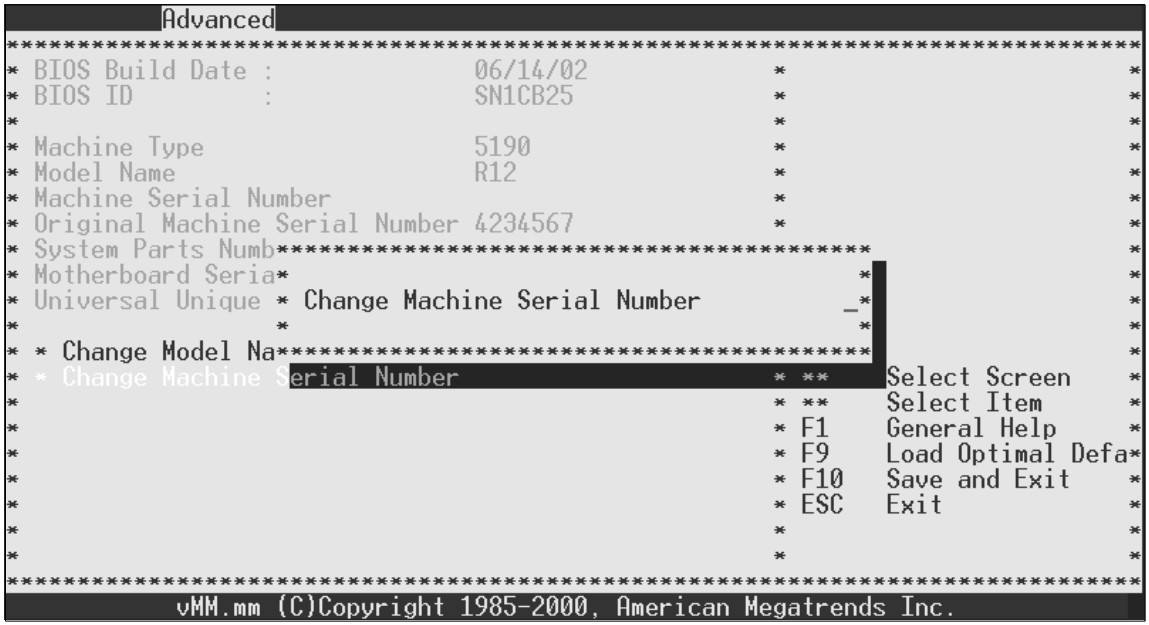

Figure B-16 Change machine serial number

2. Type the new machine serial number and press Enter. The VPD Serial Number Installed panel opens.
3. Press Enter.
4. Press Esc to return to the Advanced window.

Enabling USB support

To get the remote control from IBM Director working on the NAS100 appliance you have to enable the USB support (See 9.4.10, “Usage tips for the NAS 100” on page 329).

1. Using the arrow keys, move the cursor to **PCIPnP Configuration** and press Enter. The PCIPnP Configuration window opens (Figure B-17).

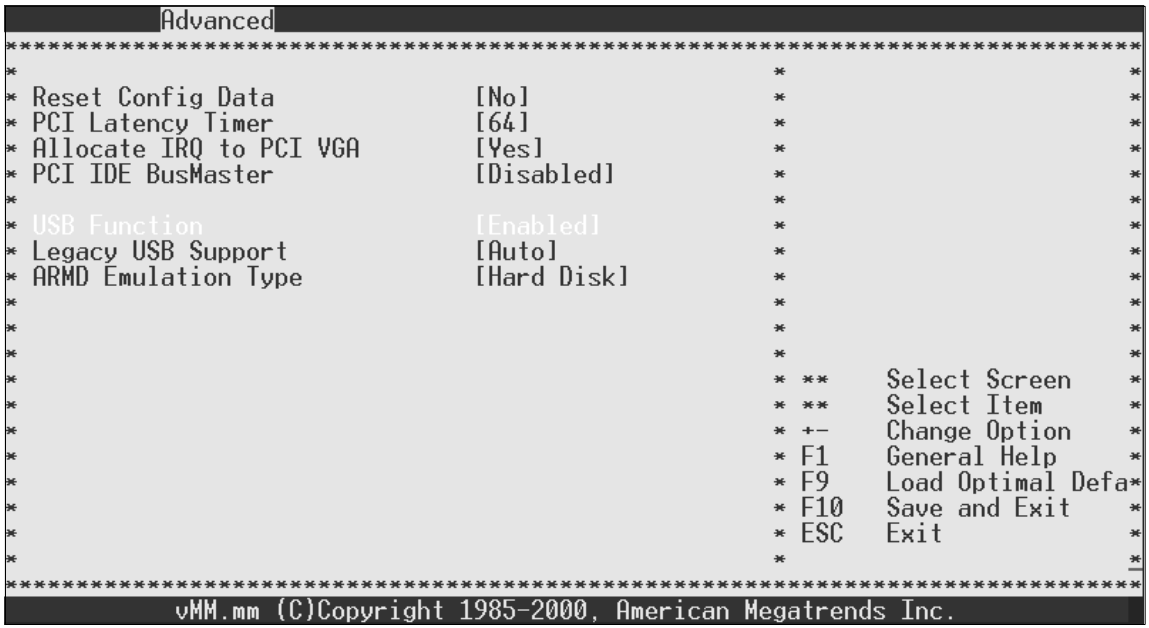

Figure B-17 Enable USB function in PCIPnP Configuration screen

2. Move to USB Function with the arrow keys and choose **Enabled**.
3. Press Escape to exit the PCIPnP Configuration.

Chipset

Do not change any settings in this window (Figure B-18)!

Figure B-18 Chipset screen

Power

Do not change any settings in this window (Figure B-19)!

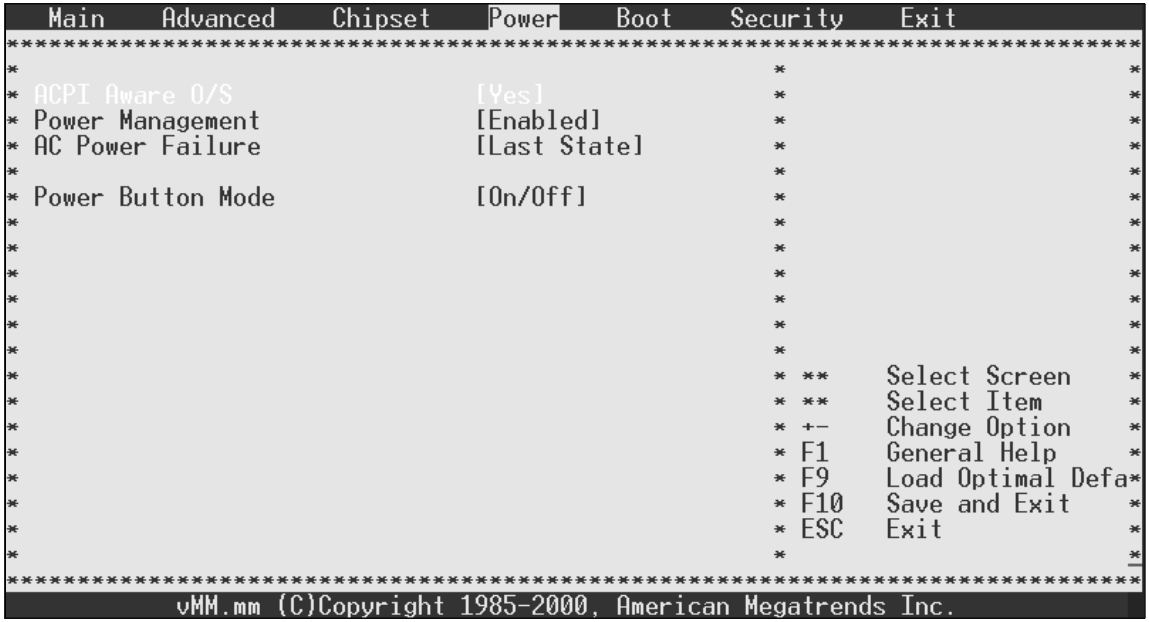

Figure B-19 Power screen

Boot

The Boot window displays the list of bootable devices. Boot priority order is:

1. USB Floppy (if enabled)
2. USB CDROM (if enabled)
3. PXE (if enabled)

You can change the setting of the devices in this window with the following actions:

1. Using the arrow keys, move the cursor to the bootable device that you want to change and press Enter. The Options panel opens.
2. Move the cursor to either **Enabled** or **Disabled**, and press Enter.

A USB CD-ROM or floppy drive is not yet supported.

If there is no need, do not change any settings in this window (Figure B-20)!

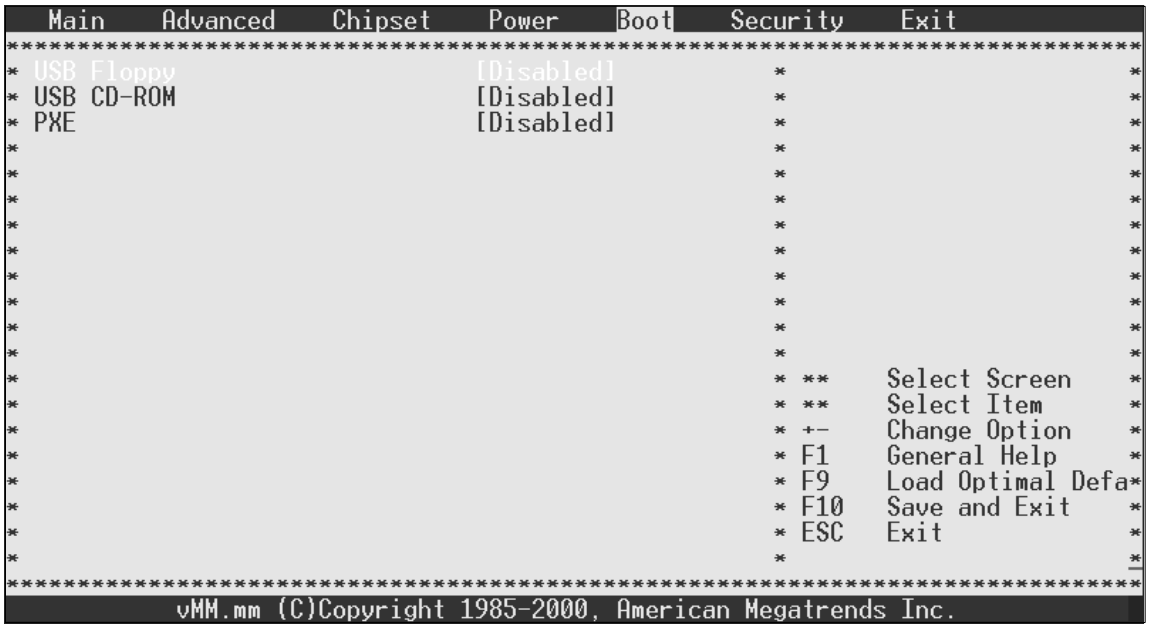

Figure B-20 Boot order screen

Attention: After you have enabled a specific device and booted, you must change its setting back to Disabled.

Security

You can change the supervisor password in this window.

Note: When you clear CMOS data, the supervisor password is reset as the default **001san** (not case-sensitive).

1. Press Enter. The Enter New Password window opens.
2. Type the new, six-character password. The new password is displayed on the panel as ****** (Figure B-21).

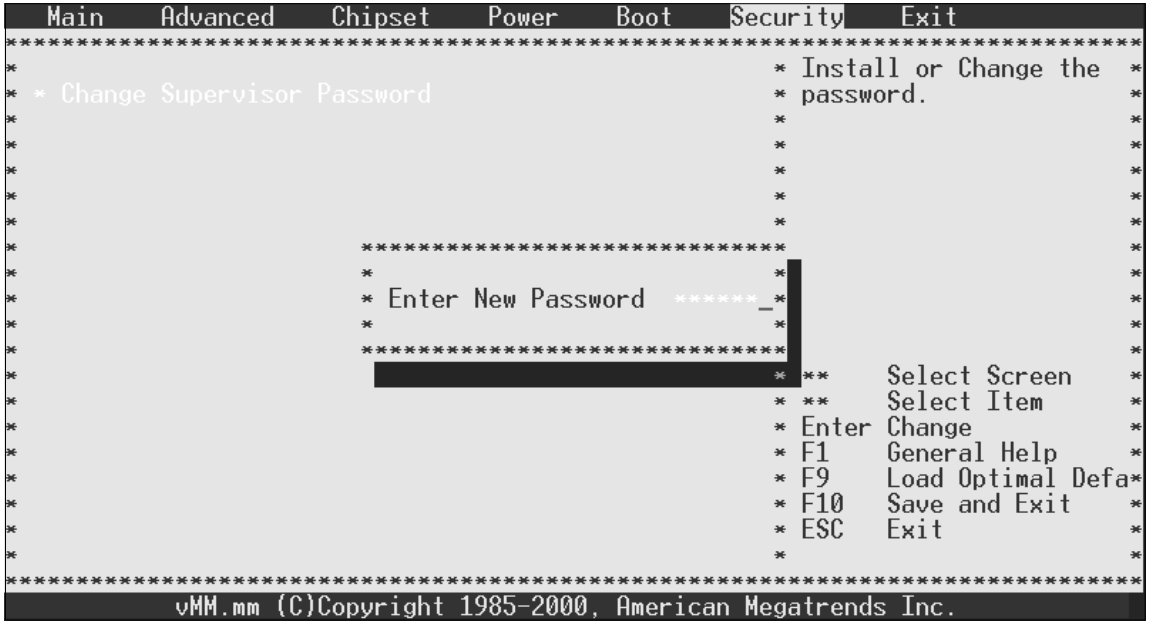

Figure B-21 Enter new password screen

3. Press Enter.
4. Retype the new password and press Enter (Figure B-22).

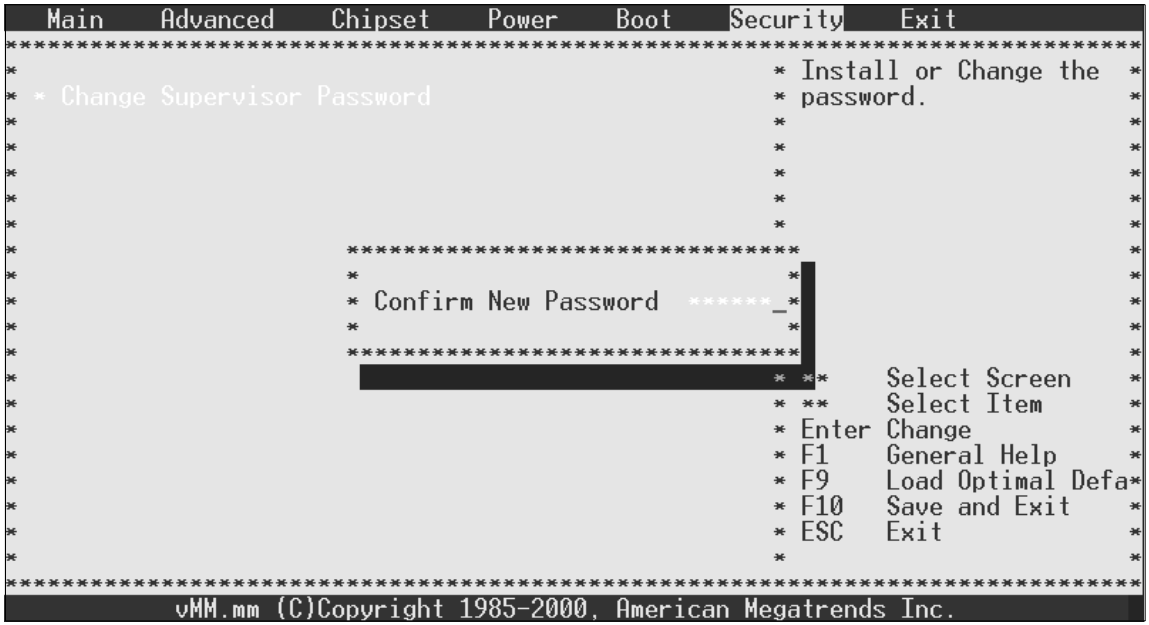

Figure B-22 Confirm password screen

5. The Password installed panel opens (Figure B-23).

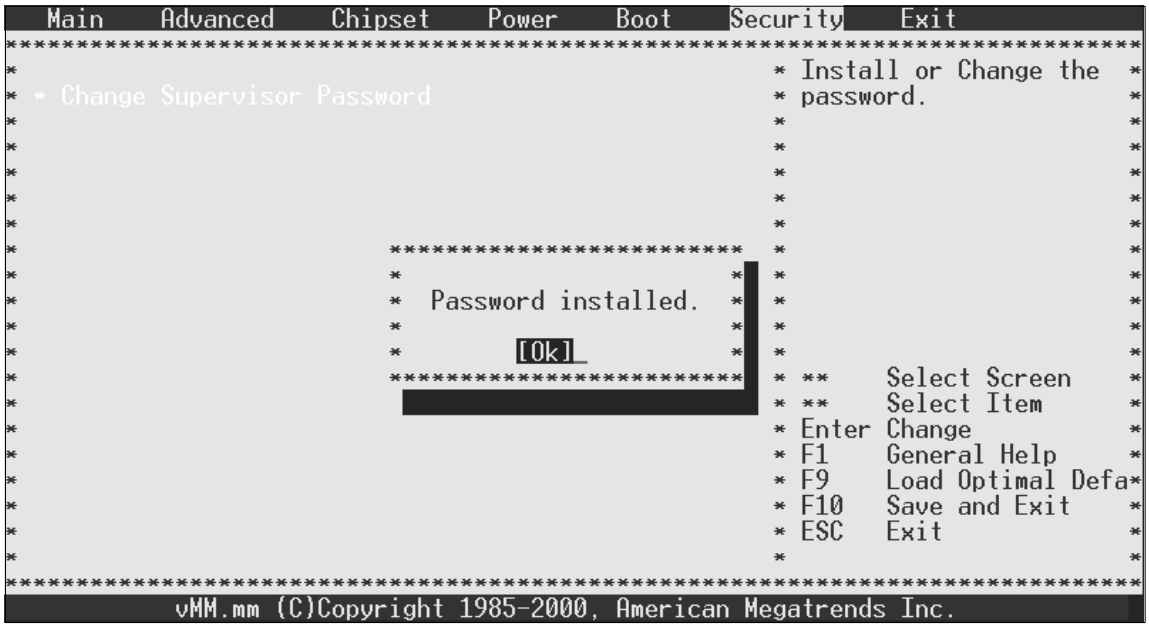

Figure B-23 Password installed screen

Exit

This window allows you to exit from the BIOS setup window.

Note: Although this window specifies that you press F9 and F10 for some actions, these keys are disabled. Use the arrow keys to move through each menu option.

Exit Saving Changes

This allows you to exit and save your changes:

1. Using the arrow keys, move the cursor to **Exit Saving Changes** (Figure B-24) and press Enter. The Save configuration changes and exit now? panel opens.

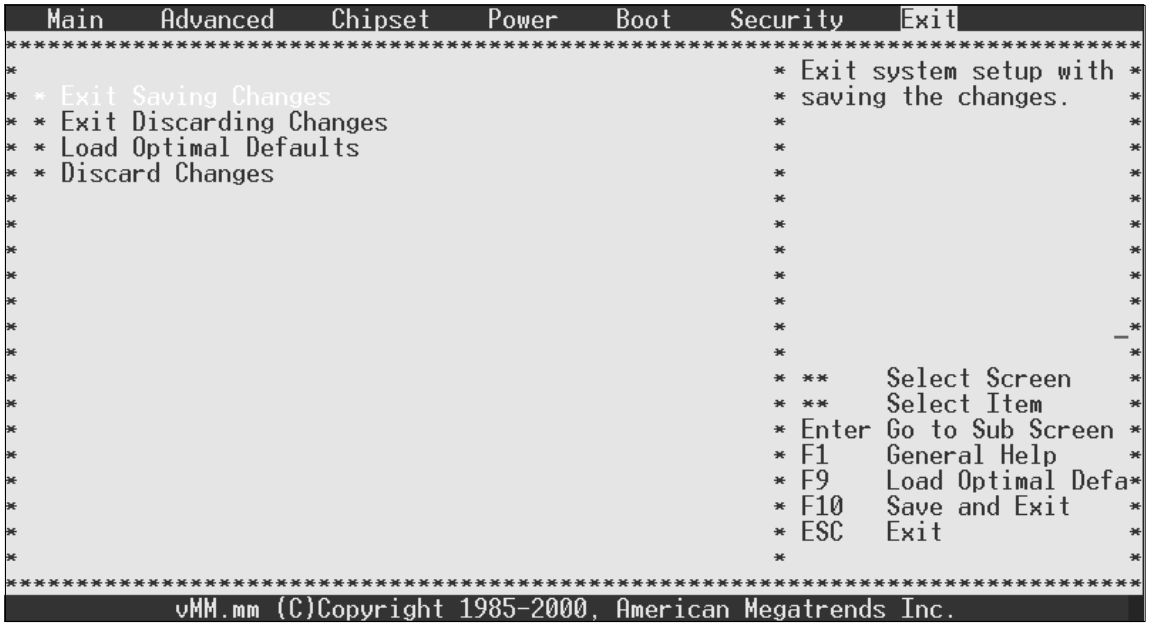

Figure B-24 Exit saving changes

2. If you want to:
   - Save the changes and exit the BIOS setup, select **OK** and press Enter.

Note: When you exit the BIOS setup, the system reboots automatically.

   - Save the changes but not exit the BIOS setup. Select **Cancel** and press Enter to return to the Exit window.

Exit Discarding Changes

This allows you to exit and discard your changes:

1. Using the arrow keys, move the cursor to **Exit Discarding Changes** (Figure B-25) and press Enter. The Discard changes and exit setup now? panel opens.

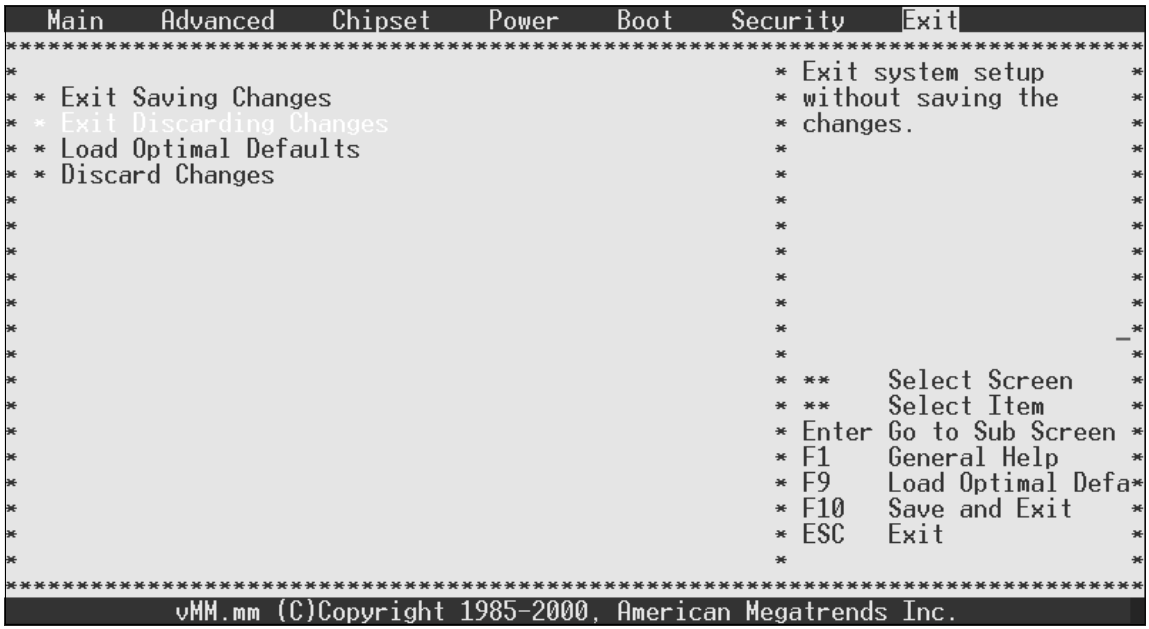

Figure B-25 Exit discard changes

2. If you want to:
   - Discard the changes and exit the BIOS setup, select **OK** and press Enter.

   Note: When you exit the BIOS setup, the system reboots automatically.

   - Save the changes but not exit the BIOS setup, select **Cancel** and press Enter to return to the Exit window.

Load Optimal Default

This offers you the following option:

1. Using the arrow keys, move the cursor to **Load Optimal Defaults** (Figure B-26) and press Enter. The Load optimal defaults? panel opens.

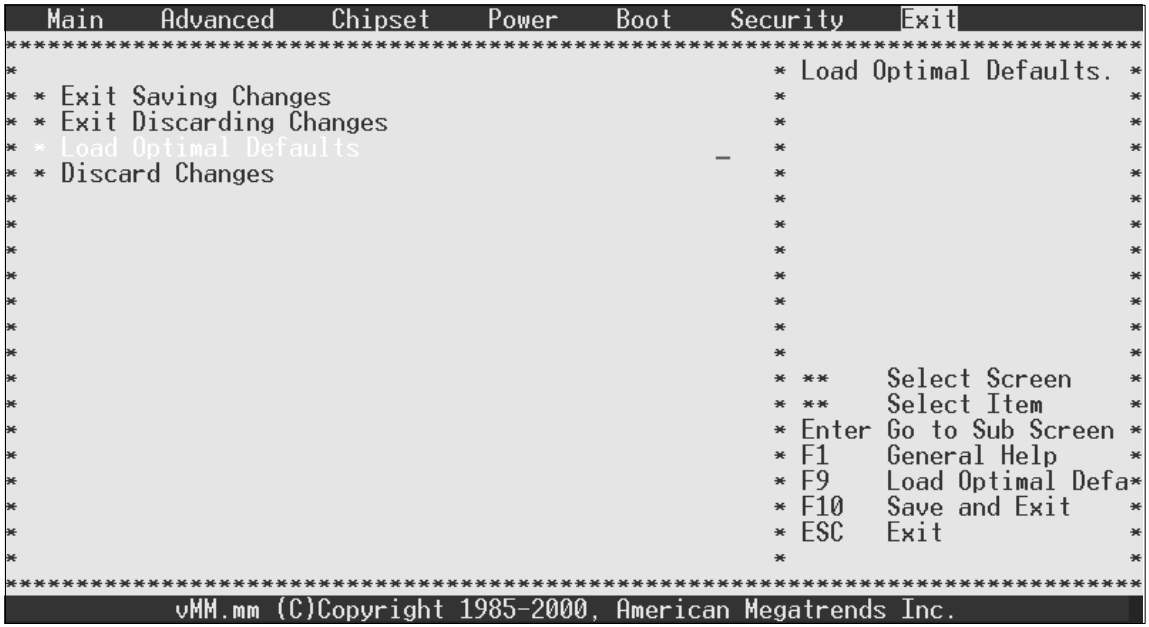

Figure B-26 Load optimal defaults

2. If you want to:
   - Set all values to their defaults, select **OK** and press Enter to reset the values to their defaults and to return to the Exit window.
   - Not set all values to their defaults, select **Cancel** and press Enter. The panel disappears.

Discard Changes

This offers you the following option:

1. Using the arrow keys, move the cursor to **Discard Changes** (Figure B-27) and press Enter. The Discard changes? panel opens.

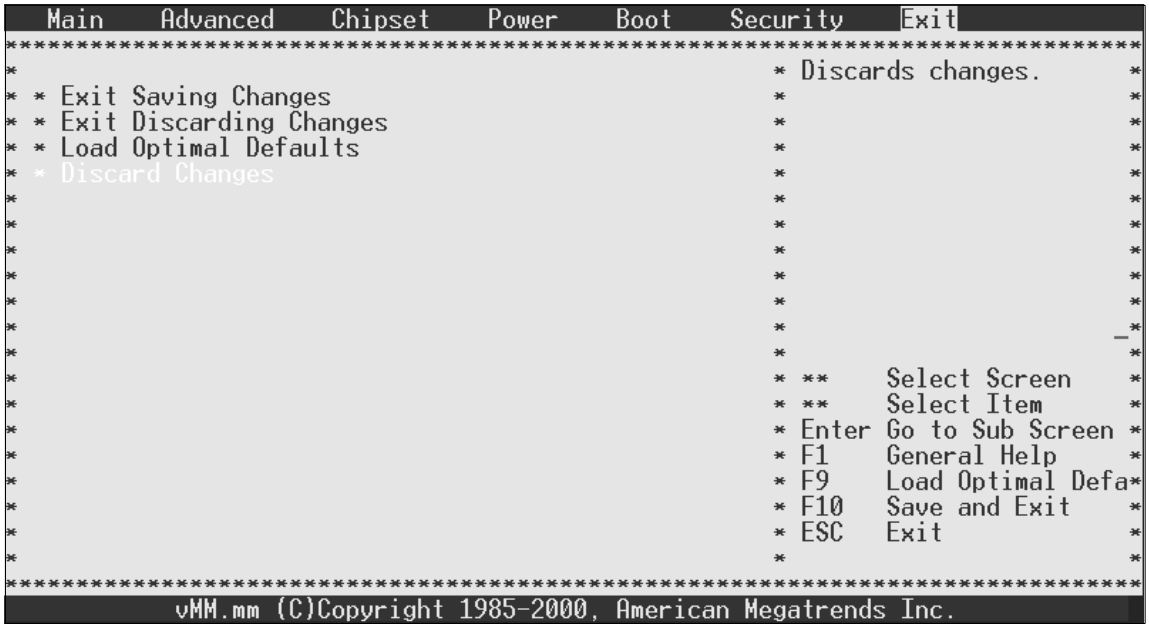

Figure B-27 Discard changes

2. If you want to:
   - Discard the changes, select **OK** and press Enter to return to the Exit window.
   - Not discard the changes, select **Cancel** and press Enter. The panel disappears.

Upgrading the BIOS

You can upgrade the BIOS by means of the serial port.

Attention: Upgrading the BIOS must be done with extreme caution. Make sure that the correct BIOS file is selected as the new BIOS when you transfer the file. If an incorrect file is loaded, you will not be able to reboot your appliance again.

To upgrade the BIOS:

1. Clear the CMOS data (see “Clearing CMOS data” on page 576).
2. Connect your appliance to a PC by means of the serial port, go to:

   http://www.ibm.com/storage/support

   Then download a new BIOS file to the PC.
3. Follow the procedure in “Preparing to use the remote BIOS setting function” on page 577 to set up HyperTerminal.
4. Power on your appliance.
5. When the message **Press <spacebar> to invoke remote BIOS flash** opens, press the spacebar.

Perform steps 5 through 8 *very quickly*.

6. The message **Begin remote BIOS flash? Y/N** opens. Type Y. The message **Starting remote flash** opens, and then the message **Update new BIOS file using Xmodem protocol**.
7. When random characters appear on the screen, click **Transfer** from the HyperTerminal menu, and then click **Send file**.
8. Specify the filename field **New BIOS file** path in your PC, set the protocol to **Xmodem**, and click **Send**.

Note: If you do not perform this step quickly enough, BIOS might proceed to the next step. If so, the message **Aborting remote flash** opens. You will need to power off the appliance and repeat the upgrading procedure more quickly.

After you send the BIOS file, the message **New BIOS received OK! Writing new BIOS to flash-Do not power DOWN or RESET!** opens.

After the new BIOS is updated in the flash, the appliance beeps four times and starts rebooting automatically.

After the appliance has rebooted, reset the date and time. See “Making changes to the BIOS” on page 589 for the procedure.

C

The AutoExNT service

In the distant past, there was something called Windows 3.1. Perhaps you can remember first starting DOS, and after that, starting Windows. The autoexec.bat file was one of the main configuration files. This file was executed at system start and could do all kinds of useful batch jobs.

Now, with the Windows NT/2000 Operating System as a multiuser OS, this possibility no longer exists. Fortunately, a programmer at Microsoft missed this batch file too and created the AutoExNT service, which works in the same manner as the autoexec.bat under Windows NT/2000. No user needs to be logged on to the system to use this great tool.

AutoExNT is a utility under Windows NT/2000 that included in the Resource Kit. In our scenario it can be used to start the redirector, establish drive mappings, and start the applications installed as an NT Service. It is a Windows NT/2000 service that executes a batch file called **autoexnt.bat**.

Installation instructions for AutoExNT

Here are the steps to install AutoExNT:

1. From the Resource kit, copy the following files from the \ntreskit folder to the %Systemroot%\System32 directory:

   a. autoexnt.exe
   b. instexnt.exe
   c. servmess.dll

2. Install AutoExNT by typing the following on the command prompt from the %Systemroot%\System32 (Figure C-1):

```
instexnt install
```

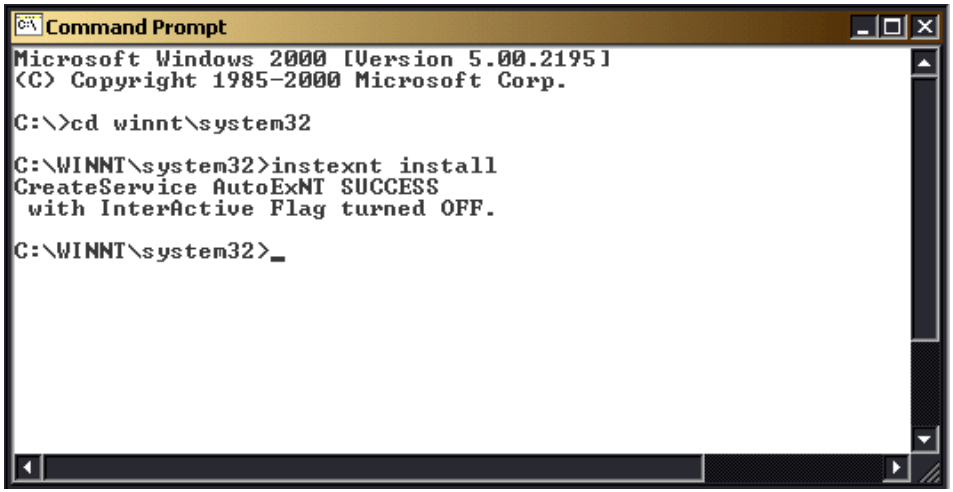

Figure C-1 AutoExNT service installation

You can also use the */interactive* switch to see the commands that are executed by AutoExNT when the system starts and a user has logged on:

1. Right-click **My Computer**, then select **Manage**. Select **Services and Applications**. Double-click **Services**.
2. Right-click **AutoExNT**, then select **Properties**. Select the **General** tab. Under *Startup type:* make sure that it is set as **Automatic** as shown in Figure C-2 (it should be set as **Automatic** by default).

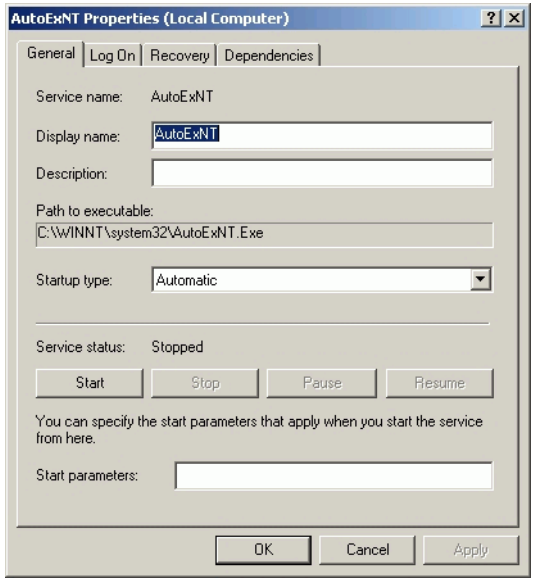

Figure C-2 AutoExNT set to start automatically

3. Select the **Log On** tab. Under *Log on as:* click the button beside *This account:.* Click **Browse...**
4. Select the domain administrator account (for example, ITSOSJ\Administrator). Supply the password and confirm it. Click **Apply**. Then select **OK** (Figure C-3).

Note: We used the administrator account in our scenario, but we recommend that you create a special service account with the special rights: **log on as service** and **act as part of operating system**.

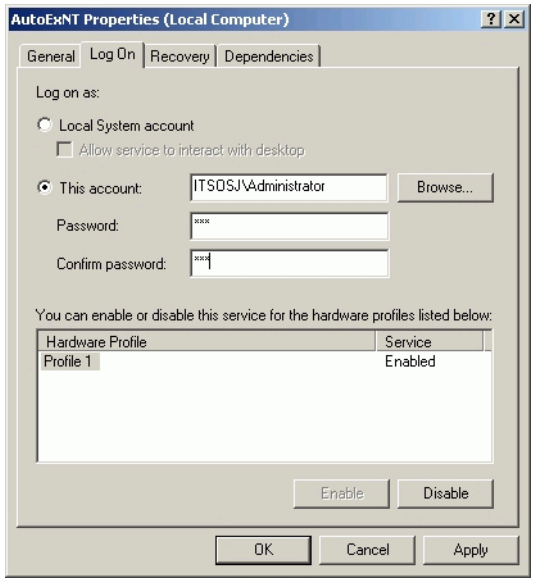

Figure C-3 Use the domain administrator account for AutoExNT

5. Create a file named **autoexnt.bat** and save it on the directory *%systemroot%\system32*. You can leave it empty for the moment. The actual content will depend on the applications that you are going to use and the drive mappings that you will need before starting the applications (Figure C-4).

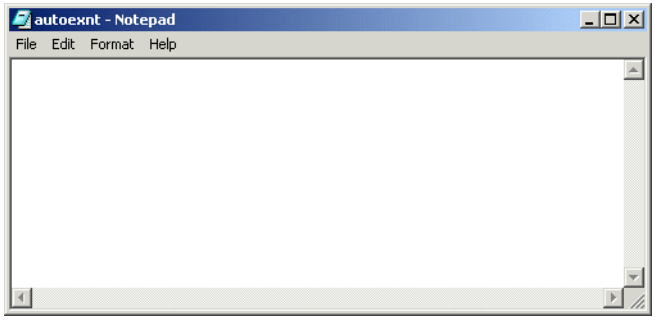

Figure C-4 The autoexnt.bat file is initially empty

Now you have successfully installed the AutoExNT service.

D

IBM TotalStorage NAS file system integrity features

In this appendix we describe the checkdisk (Chkdsk) tool and IBM TotalStorage Network Attached Storage.

When considering Network Attached Storage (NAS) solutions, an important subject is the file system. With Windows OS Powered IBM NAS appliances, the functionality of the NT file system (NTFS) Version 5, and especially the need to use Chkdsk to verify the file system integrity, might be questioned.

We start with some general considerations concerning file systems, discuss NTFS and Chkdsk, and then provide a quick overview of specific features provided by IBM to ensure system and data availability in IBM TotalStorage NAS systems.

You will see that there are good reasons to use NTFS in an IBM TotalStorage NAS system.

Introduction to the file writing process

IBM TotalStorage NAS appliance systems have a Windows powered OS that is based on Windows 2000 Advanced Server. It works with files in exactly the same way that Windows 2000 does.

When the operating system (OS) or application wants to write a file to disk, the file has to go through the file system that issues the write commands to the disk driver or disk controller driver. This driver is aware of the physical structure of the disk (cylinders, tracks and sectors) and sends the low-level commands (for example, SCSI commands) to the disk controller or the disk, as shown in Figure D-1.

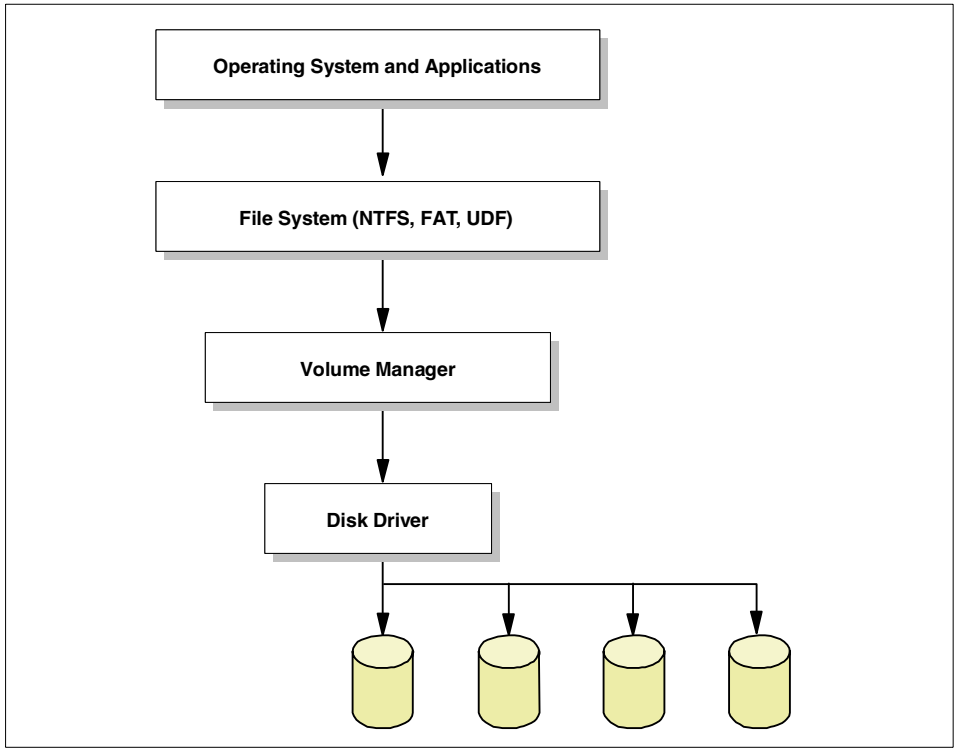

Figure D-1 Windows powered OS file writing process

File systems

Several different file systems are supported for Windows 2000 (and Windows powered OS as well). FAT, FAT32, NTFS, CDFS, and UDF. CDFS and UDF are file systems related to CD-ROMs. Regarding hard disk file systems, NTFS is a much more advanced and complete file system than FAT or FAT32. For a more detailed description of NTFS, see “NTFS” on page 612.

Caching

Some devices, like disks, are slow in comparison to other computer components such as memory and processors. To improve the data flow between fast and slow devices, caches were implemented. A cache is fast temporary data storage where the data to be transmitted can wait until the device is ready to receive it.

Disk controllers implement the cache in one of two ways when writing a file to disks:

Write Through (WT) Write Through cache waits until the data is actually written to the disk before this write is confirmed.

Write Back (WB) Write Back cache confirms the writing of the data immediately when data is in cache, although the data is not yet written to disk.

The advantage of a Write Back cache is the improved performance, while a Write Through cache is more secure.

Here is how you can configure a disk controller cache to ensure data and file system integrity:

- **Configure cache as Write Through**: This will avoid any committed, but unwritten transactions to the disk.
- **Configure cache as Write Back**: In this case, we highly recommend battery backup cache. In the case of a power outage or operating system hang, committed but not yet written data is kept in the cache and can be written to the disks as soon as the system is restarted.

The cache file systems concerned can work in similar ways as disk controllers when writing data to disks.

As we explain in “NTFS” on page 612, NTFS incorporates the advantages of both caching methods. As a recoverable file system, NTFS maintains a log file of changes to the volume. NTFS logs each transaction, and if the system fails, can then undo or redo any change that has not been completed. This keeps NTFS in a consistent state and ensures that the volume structure does not become corrupted.

NTFS and Chkdsk explained

In this section, we describe key functionalities of NTFS and Chkdsk as a basis for understanding the benefits that the combination of NTFS and Chkdsk can have on your NAS environment.

Note: Windows powered OS uses NTFS Version 5. Therefore, the information in this section refers specifically to NTFS Version 5.

NTFS

NTFS is the most commonly used file system in Windows 2000 environments. There are many reasons, which include the following:

- Recoverability
- Enhanced partition size
- Data compression
- Support for disk quotas
- Support for Encrypted File System (EFS)
- Volume mount points
- File change journal
- Support for Active Directory

Let us have a more detailed look at the recoverability features of NTFS, as used in Windows 2000.

In general, file systems can be divided into two generic types:

Careful write A careful write file system waits until the data is actually written to the disk before this write is confirmed.

Lazy write A lazy write file system confirms the writing of the data immediately, although the data is not yet written to disk.

The advantage of a lazy write file system is the improved performance, while a careful write file system is more secure.

So what about NTFS? NTFS takes the best of both worlds: the performance of lazy write file systems along with the security of careful write file systems. So, as soon as the operation is logged, this operation can be confirmed as the file system can recover from potential failures.

Similar to database log files, NTFS keeps track of changes in the file system by using a log file. This log file can be used to recover volume structure information when inconsistencies occur. This does not protect the user data though, it just makes sure that the data can still be accessed.

To ensure the volume structure integrity, NTFS can roll back (undo) or roll forward (redo) transactions to the last commit point, just like a database using a transaction log file. So how does this work?

Every five seconds NTFS sets a commit point and resets the log file by sending all modified data to disk. So, NTFS need only roll back transactions to the last commit point in order to recover consistency within the file system.

In the event of a system crash (for example, caused by a power failure), NTFS performs three passes to ensure the integrity of data in the volume:

1. **Analysis pass:** NTFS analyzes inconsistencies between data and log file entries.
2. **Redo pass:** Logged transactions are redone.
3. **Undo pass:** Incomplete transactions are undone.

Here is an example: If, during a write process, NTFS is adding data to a file, and the write process is interrupted after NTFS has extended the file length, but *before* actually adding the data to the file, then the transaction log is used to shorten the file again to the original length. This would be an undo operation. If a roll forward operation is possible, NTFS will use the information in the log file to redo the operation in question, for example, a file deletion.

This way, the consistency of the file system is ensured.

For more detailed information, refer to the Microsoft Knowledge Base Article “Q101670” currently available at:

`http://support.microsoft.com/default.aspx?scid=kb;en-us;Q101670`

Chkdsk

Windows 2000 based operating systems use Chkdsk to verify the logical integrity of a file system. Chkdsk takes actions to repair logical inconsistencies in the file system if they occur.

Reasons for such inconsistencies are most likely caused by power or hardware failures, or operating system hangs.

For example, in the case of a power outage while writing data to the disk, NTFS will detect an inconsistency in the file system at boot time and will initiate an Autochk process, which invokes Chkdsk code to be executed while booting.

Note: In clustered environments (as in the IBM TotalStorage NAS 300 and NAS 300G products that make use of Microsoft's Cluster Services), the situation is somewhat different. Volumes that are shared between the two cluster nodes cannot be checked using Autochk at boot time, and therefore, must use Chkdsk after the owning node has booted. If a volume is inconsistent (has the "dirty" bit set) a Chkdsk is performed before the volume is brought online, to ensure volume and data integrity.

Chkdsk activity is split into three major passes:

1. Examining file record segments (FRS) in the volume´s master file table (MFT).
2. Examining the volume´s indexes (directories) and their associated file entries.
3. Examining the security descriptors for the files and directories in the volume.

An optional fourth pass can be forced by the use of the /R command line switch. In this pass, every sector in the volume´s free space is checked for bad sectors. Actually, this switch is not widely used, as bad sectors in use are detected by the *normal* Chkdsk, and if unused, they are detected and re-mapped when used for the first time.

For more detailed information, refer to these Microsoft Knowledge Base Articles:

For "Q187941", see:

http://support.microsoft.com/default.aspx?scid=kb;en-us;Q187941

For "Q218461", see:

http://support.microsoft.com/default.aspx?scid=kb;en-us;Q218461

IBM TotalStorage NAS systems

IBM offers a range of NAS devices that answer the current and growing demand for Network Attached Storage. As an example, we provide you with some details on the IBM TotalStorage NAS 300, available at the time this was written.

The NAS 300 is a member of IBM's TotalStorage NAS family of products, which are specifically designed, configured, and packaged to provide solutions to help overcome the challenges of cost effectively sharing, managing, and protecting data within complex network infrastructures.

These features are of particular interest for reliability and data availability:

- FC-based RAID Controller with battery backup cache

  The battery backup cache helps ensure data integrity, as described in Chapter , “Caching” on page 611.

- Clustering for fast failover in case a NAS engine fails
- IT Director for easy administration and error discovery
- Tivoli Storage Manager (TSM) client for fast and reliable backup

E

SAS connections

In this appendix we show how easy it is to connect the SAS software to DB2 UDB and Microsoft SQL Server. We describe general connection setup possibilities and specifications for those two databases.

Engine name and connection options for DB2 UDB

The SAS/ACCESS **engine name** for the interface to DB2 UDB UNIX/PC is **DB2**. You use this engine name in your LIBNAME statements, for example:

```
libname db2data db2;
```

There are several ways to connect to DB2 UNIX/PC when using the LIBNAME statement. Use ***only one*** of the following methods for each connection, since they are mutually exclusive:

- Specify USER=, PASSWORD= and DATASRC=, or
- Specify COMPLETE=, or
- Specify NOPROMPT=, or
- Specify PROMPT=, or
- Specify REQUIRED=.

Definitions of these **engine connection options** and a usage example are provided here:

USER=<'>*username*<'>

Enables you to connect to a DB2 database with a user ID that is different from the default ID.

USER= is optional. If you specify USER=, you must also specify PASSWORD=. If USER= is omitted, your default user ID for your operating environment is used.

PASSWORD=<'>*password*<'>

Specifies the DB2 password that is associated with your DB2 user ID.

PASSWORD= is optional. If you specify USER=, you must specify PASSWORD=.

DATASRC=<'>*data-source-name*<'>

Specifies the DB2 data source or database to which you want to connect.

DATASRC= is optional. If you omit it, you connect by using a default environment variable.

DSN= is an alias for this option.

COMPLETE=<'>*CLI-connection-string*<'>

Specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connect string is returned in the SYSDBMSG macro variable.

If you do not specify enough correct connection options, you are prompted with a dialog box that displays the values from the COMPLETE= connection string. You can edit any field before you connect to the data source.

This option is not available on UNIX platforms. See your DB2 documentation for more details.

NOPROMPT=<'>*CLI-connection-string*<'>

Specifies connection options for your data source or database. Separate multiple options with a semicolon.

If you do not specify enough correct connection options, an error is returned (no dialog box is displayed).

PROMPT=<'> *CLI-connection-string*<'>

Specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connect string is returned in the SYSDBMSG macro variable.

PROMPT= does not immediately attempt to connect to the DBMS. Instead, it displays a dialog box that contains the values that you entered in the PROMPT= connection string. You can edit values or enter additional values in any field before you connect to the data source.

This option is not available on UNIX platforms.

REQUIRED=<'>*CLI-connection-string*<'>

Specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connect string is returned in the SYSDBMSG macro variable.

If you do not specify enough correct connection options, a dialog box prompts you for the connection options. REQUIRED= only allows you to modify required fields in the dialog box.

This option is not available on UNIX platforms.

In Example E-1, the LIBREF MYDBLIB uses the DB2 engine to connect to a DB2 database by using the NOPROMPT= option. PROC PRINT is used to display the contents of the DB2 table CUSTOMERS.

Example: E-1 SAS DB2 example

```
libname mydblib db2
   noprompt="dsn=userdsn;uid=testuser;pwd=testpass;"

proc print data=mydblib.customers;
   where state='CA';
run;
```

Engine name and connection options for SQL Server

The SAS/ACCESS **engine name** for the interface to SQL Server under UNIX hosts is **SQLSVR**. You use this engine name in your LIBNAME statements, for example:

```
libname mydblib SQLSVR user=testuser password=testpass;
```

In the preceding example, USER= and PASSWORD= are **engine connection options**. There are multiple ways that you can connect to SQL Server when using the LIBNAME statement. Use ***only one*** of the following methods for each connection since they are mutually exclusive:

- Specify USER=, PASSWORD=, and DATASRC=, or
- Specify COMPLETE=, or
- Specify NOPROMPT=, or
- Specify PROMPT=, or
- Specify REQUIRED=.

Definitions of these engine connection options and usage examples are provided here:

USER=<'>*username*<'>

This enables you to connect to SQL Server with a user ID that is different from the default ID.

USER= is optional. UID= is an alias for this option.

PASSWORD=<'>*password*<'>

This specifies the SQL Server password that is associated with your user ID.

PASSWORD= is optional. PWD is an alias for this option.

DATASRC=<'>*SQL Server-data-source*<'>

This specifies the SQL Server data source to which you want to connect. For PC platforms, data sources must be configured by using the SQL Server icon in the Windows Control Panel. For UNIX platforms, data sources must be configured by modifying the **.ODBC.ini** file.

DSN= is an alias for this option that indicates that the connection is attempted using the ODBC SQLConnect API, which requires a data source name. Optionally, a user ID and password (described below) can be used in conjunction with DSN=. This API is guaranteed to be present in all drivers.

COMPLETE=<'>*SQL-Server-connection-options*<'>

This specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connect string is returned in the SYSDBMSG macro variable.

If you do not specify enough correct connection options, you are prompted with a dialog box that displays the values from the COMPLETE= connection string. You can edit any field before you connect to the data source.

See your driver documentation for more details.

NOPROMPT=<'>*SQL-Server-connection-options*<'>

This specifies connection options for your data source or database. Separate multiple options with a semicolon.

If you do not specify enough correct connection options, an error is returned (no dialog box is displayed).

PROMPT=<'> *SQL-Server-connection-information*<'>

This specifies connection options to the data source.

This specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connect string is returned in the SYSDBMSG macro variable.

PROMPT= does not immediately attempt to connect to the DBMS. Instead, it displays a dialog box that contains the values that you entered in the PROMPT= connection string. You can edit values or enter additional values in any field before you connect to the data source.

REQUIRED=<'>*SQL-Server-connection-options*<'>

This specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connect string is returned in the SYSDBMSG macro variable.

If you do not specify enough correct connection options, a dialog box prompts you for the connection options. REQUIRED= only allows you to modify required fields in the dialog box.

In Example E-2, the LIBREF MYDBLIB connects to a Microsoft SQL Server database using the NOPROMPT= option.

Example: E-2 SAS SQL example

```
libname mydblib SQLSVR
   noprompt="uid=testuser;pwd=testpass;dsn=sqlservr;"
   stringdates=yes;

proc print data=mydblib.customers;
   where state='CA';
run;
```

Glossary

A

Agent A software entity that runs on endpoints and provides management capability for other hardware or software. An example is an SNMP agent. An agent has the ability to spawn other processes.

AL See arbitrated loop.

Allocated storage The space that is allocated to volumes, but not assigned.

Allocation The entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set.

Arbitrated loop A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate. See also Fibre Channel Arbitrated Loop and loop topology.

Array An arrangement of related disk drive modules that have been assigned to a group.

B

Bandwidth A measure of the data transfer rate of a transmission channel.

Bridge Facilitates communication with LANs, SANs, and networks with dissimilar protocols.

C

Client A function that requests services from a server, and makes them available to the user. A term used in an environment to identify a machine that uses the resources of the network.

Client authentication The verification of a client in secure communications where the identity of a server or browser (client) with whom you wish to communicate is discovered. A sender's authenticity is demonstrated by the digital certificate issued to the sender.

Client-server relationship Any process that provides resources to other processes on a network is a server. Any process that employs these resources is a client. A machine can run client and server processes at the same time.

Console A user interface to a server.

D

DATABASE 2 (DB2) A relational database management system. DB2 Universal Database is the relational database management system that is Web-enabled with Java support.

Device driver A program that enables a computer to communicate with a specific device, for example, a disk drive.

Disk group A set of disk drives that have been configured into one or more logical unit numbers. This term is used with RAID devices.

E

Enterprise network A geographically dispersed network under the backing of one organization.

Enterprise Storage Server Provides an intelligent disk storage subsystem for systems across the enterprise.

Event In the Tivoli environment, any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task. Examples of events are: the normal starting and s ping of a process, the abnormal termination of a process, and the malfunctioning of a server.

F

Fabric The Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is often used to describe a more complex network utilizing hubs, switches, and gateways.

FC See Fibre Channel.

FCS See Fibre Channel standard.

Fiber optic The medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel A technology for transmitting data between computer devices at a data rate of up to 1 Gb. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

Fibre Channel Arbitrated Loop A reference to the FC-AL standard, a shared gigabit media for up to 127 nodes, one of which can be attached to a switch fabric. See also arbitrated loop and loop topology. Refer to American National Standards Institute (ANSI) X3T11/93-275.

Fibre Channel standard An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. Refer to ANSI X3.230-199x.

File system An individual file system on a host. This is the smallest unit that can monitor and extend. Policy values defined at this level override those that might be defined at higher levels.

G

Gateway In the SAN environment, a gateway connects two or more different remote SANs with each other. A gateway can also be a server on which a gateway component runs.

H

Hardware zoning Hardware zoning is based on physical ports. The members of a zone are physical ports on the fabric switch. It can be implemented in the following configurations: one to one, one to many, and many to many.

HBA See host bus adapter.

Host Any system that has at least one internet address associated with it. A host with multiple network interfaces can have multiple internet addresses associated with it. This is also referred to as a server.

Host bus adapter (HBA) A Fibre Channel HBA connection that allows a workstation to attach to the SAN network.

Hub A Fibre Channel device that connects up to 126 nodes into a logical loop. All connected nodes share the bandwidth of this one logical loop. Hubs automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

IP Internet protocol.

J

Java A programming language that enables application developers to create object-oriented programs that are very secure, portable across different machine and operating system platforms, and dynamic enough to allow expandability.

Java runtime environment (JRE) The underlying, invisible system on your computer that runs applets the browser passes to it.

Java Virtual Machine (JVM) The execution environment within which Java programs run. The Java virtual machine is described by the Java Machine Specification which is published by Sun Microsystems. Because the Tivoli Kernel Services is based on Java, nearly all ORB and component functions execute in a Java virtual machine.

JBOD Just a Bunch Of Disks.

JRE See Java runtime environment.

JVM See Java Virtual Machine.

L

Logical unit number (LUN) The LUNs are provided by the storage devices attached to the SAN. This number provides you with a volume identifier that is unique among all storage servers. The LUN is synonymous with a physical disk drive or a SCSI device. For disk subsystems such as the IBM Enterprise Storage Server, a LUN is a logical disk drive. This is a unit of storage on the SAN which is available for assignment or unassignment to a host server.

Loop topology In a loop topology, the available bandwidth is shared with all the nodes connected to the loop. If a node fails or is not powered on, the loop is out of operation. This can be corrected using a hub. A hub opens the loop when a new node is connected and closes it when a node disconnects. See also Fibre Channel Arbitrated Loop and arbitrated loop.

LUN See logical unit number.

LUN assignment criteria The combination of a set of LUN types, a minimum size, and a maximum size used for selecting a LUN for automatic assignment.

LUN masking This allows or blocks access to the storage devices on the SAN. Intelligent disk subsystems like the IBM Enterprise Storage Server provide this kind of masking.

M

Managed object A managed resource.

Managed resource A physical element to be managed.

Management Information Base (MIB) A logical database residing in the managed system which defines a set of MIB objects. A MIB is considered a logical database because actual data is not stored in it, but rather provides a view of the data that can be accessed on a managed system.

MIB See Management Information Base.

MIB object A MIB object is a unit of managed information that specifically describes an aspect of a system. Examples are CPU utilization, software name, hardware type, and so on. A collection of related MIB objects is defined as a MIB.

N

Network topology A physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

N_Port node port A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

NL_Port node loop port A node port that supports arbitrated loop devices.

O

Open system A system whose characteristics comply with standards made available throughout the industry, and therefore can be connected to other systems that comply with the same standards.

P

Point-to-point topology It consists of a single connection between two nodes. All the bandwidth is dedicated for these two nodes.

Port An end point for communication between applications, generally referring to a logical connection. A port provides queues for sending and receiving data. Each port has a port number for identification. When the port number is combined with an Internet address, it is called a socket address.

Port zoning In Fibre Channel environments, port zoning is the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones. See also LUN masking and subsystem masking.

Protocol The set of rules governing the operation of functional units of a communication system if communication is to take place. Protocols can determine low-level details of machine-to-machine interfaces, such as the order in which bits from a byte are sent. They can also determine high-level exchanges between application programs, such as file transfer.

R

RAID Redundant array of inexpensive or independent disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

S

SAN See storage area network.

SAN agent A software program that communicates with the manager and controls the subagents. This component is largely platform independent. See also subagent.

SCSI Small Computer System Interface. An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface utilizes a SCSI logical protocol over an I/O interface that configures attached targets and initiators in a multi-drop bus topology.

Server A program running on a mainframe, workstation, or file server that provides shared services. This is also referred to as a host.

Shared storage Storage within a storage facility that is configured such that multiple homogeneous or divergent hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. You need to design the programs to handle the effects of concurrent access.

Simple Network Management Protocol (SNMP) A protocol designed to give a user the capability to remotely manage a computer network by polling and setting terminal values and monitoring network events.

SNMP See Simple Network Management Protocol.

SNMP agent An implementation of a network management application which is resident on a managed system. Each node that is to be monitored or managed by an SNMP manager in a TCP/IP network, must have an SNMP agent resident. The agent receives requests to either retrieve or modify management information by referencing MIB objects. MIB objects are referenced by the agent whenever a valid request from an SNMP manager is received.

SNMP manager A managing system that executes a managing application or suite of applications. These applications depend on MIB objects for information that resides on the managed system.

SNMP trap A message that is originated by an agent application to alert a managing application of the occurrence of an event.

Software zoning Is implemented within the Simple Name Server (SNS) running inside the fabric switch. When using software zoning, the members of the zone can be defined with: node WWN, port WWN, or physical port number. Usually the zoning software also allows you to create symbolic names for the zone members and for the zones themselves.

SQL Structured Query Language.

Storage administrator A person in the data processing center who is responsible for defining, implementing, and maintaining storage management policies.

Storage area network (SAN) A managed, high-speed network that enables any-to-any interconnection of heterogeneous servers and storage systems.

Subagent A software component of SAN products which provides the actual remote query and control function, such as gathering host information and communicating with other components. This component is platform dependent. See also SAN agent.

Subsystem masking The support provided by intelligent disk storage subsystems like the Enterprise Storage Server. See also LUN masking and port zoning.

Switch A component with multiple entry and exit points or ports that provide dynamic connection between any two of these points.

Switch topology A switch allows multiple concurrent connections between nodes. There can be two types of switches, circuit switches and frame switches. Circuit switches establish a dedicated connection between two nodes. Frame switches route frames between nodes and establish the connection only when needed. A switch can handle all protocols.

T

TCP See Transmission Control Protocol.

TCP/IP Transmission Control Protocol/Internet Protocol.

Topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, arbitrated loop, and switched fabric are all Fibre Channel topologies.

Transmission Control Protocol (TCP) A reliable, full duplex, connection-oriented, end-to-end transport protocol running on of IP.

W

WAN Wide Area Network.

Z

Zoning In Fibre Channel environments, zoning allows for finer segmentation of the switched fabric. Zoning can be used to instigate a barrier between different environments. Ports that are members of a zone can communicate with each other but are isolated from ports in other zones. Zoning can be implemented in two ways: hardware zoning and software zoning.

Other glossaries:

For more information on IBM terminology, see the IBM Storage Glossary of Terms at:

http://www.storage.ibm.com/glossary.htm

For more information on Tivoli terminology, see the Tivoli Glossary at:

http://www.tivoli.com/support/documents/glossary/termsm03.htm

Abbreviations and acronyms

ABI	Application Binary Interface
ACE	Access Control Entries
ACL	Access Control List
AD	Microsoft Active Directory
ADSM	ADSTAR Distributed Storage Manager
AFS	Andrew File System
AIX	Advanced Interactive eXecutive
ANSI	American National Standards Institute
APA	All Points Addressable
API	Application Programming Interface
APPC	Advanced Program-to-Program Communication
APPN	Advanced Peer-to-Peer Networking
ARC	Advanced RISC Computer
ARPA	Advanced Research Projects Agency
ASCII	American National Standard Code for Information Interchange
ATE	Asynchronous Terminal Emulation
ATM	Asynchronous Transfer Mode
AVI	Audio Video Interleaved
BDC	Backup Domain Controller
BIND	Berkeley Internet Name Domain
BNU	Basic Network Utilities
BOS	Base Operating System
BRI	Basic Rate Interface
BSD	Berkeley Software Distribution
BSOD	Blue Screen of Death
BUMP	Bring-Up Microprocessor
CA	Certification Authorities
CAL	Client Access License
C-SPOC	Cluster single point of control
CDE	Common Desktop Environment
CDMF	Commercial Data Masking Facility
CDS	Cell Directory Service
CERT	Computer Emergency Response Team
CGI	Common Gateway Interface
CHAP	Challenge Handshake Authentication
CIDR	Classless InterDomain Routing
CIFS	Common Internet File System
CMA	Concert Multi-threaded Architecture
CO	Central Office
COPS	Computer Oracle and Password System

CPI-C	Common Programming Interface for Communications
CPU	Central Processing Unit
CSNW	Client Service for NetWare
CSR	Client/server Runtime
DAC	Discretionary Access Controls
DARPA	Defense Advanced Research Projects Agency
DASD	Direct Access Storage Device
DBM	Database Management
DCE	Distributed Computing Environment
DCOM	Distributed Component Object Model
DDE	Dynamic Data Exchange
DDNS	Dynamic Domain Name System
DEN	Directory Enabled Network
DES	Data Encryption Standard
DFS	Distributed File System
DHCP	Dynamic Host Configuration Protocol
DLC	Data Link Control
DLL	Dynamic Load Library
DS	Differentiated Service
DSA	Directory Service Agent
DSE	Directory Specific Entry
DNS	Domain Name System
DTS	Distributed Time Service
EFS	Encrypting File Systems
EGID	Effective Group Identifier
EISA	Extended Industry Standard Architecture
EMS	Event Management Services
EPROM	Erasable Programmable Read-Only Memory
ERD	Emergency Repair Disk
ERP	Enterprise Resources Planning
ERRM	Event Response Resource Manager
ESCON	Enterprise System Connection
ESP	Encapsulating Security Payload
ESS	Enterprise Storage Server
EUID	Effective User Identifier
FAT	File Allocation Table
FC	Fibre Channel
FDDI	Fiber Distributed Data Interface
FDPR	Feedback Directed Program Restructure
FEC	Fast EtherChannel technology
FIFO	First In/First Out
FIRST	Forum of Incident Response and Security
FQDN	Fully Qualified Domain Name
FSF	File Storage Facility
FTP	File Transfer Protocol
FtDisk	Fault-Tolerant Disk
GC	Global Catalog
GDA	Global Directory Agent
GDI	Graphical Device Interface

GDS	Global Directory Service
GID	Group Identifier
GL	Graphics Library
GSNW	Gateway Service for NetWare
GUI	Graphical User Interface
HA	High Availability
HACMP	High Availability Cluster Multiprocessing
HAL	Hardware Abstraction Layer
HBA	Host Bus Adapter
HCL	Hardware Compatibility List
HSM	Hierarchical Storage Management
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines Corporation
ICCM	Inter-Client Conventions Manual
IDE	Integrated Drive Electronics
IDL	Interface Definition Language
IDS	Intelligent Disk Subsystem
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IIS	Internet Information Server
IKE	Internet Key Exchange
IMAP	Internet Message Access Protocol
I/O	Input/Output
IP	Internet Protocol
IPC	Interprocess Communication
IPL	Initial Program Load
IPsec	Internet Protocol Security
IPX	Internetwork Packet eXchange
ISA	Industry Standard Architecture
iSCSI	SCSI over IP
ISDN	Integrated Services Digital Network
ISNO	Interface-specific Network Options
ISO	International Standards Organization
ISS	Interactive Session Support
ISV	Independent Software Vendor
ITSEC	Initial Technology Security Evaluation
ITSO	International Technical Support Organization
ITU	International Telecommunications Union
IXC	Inter Exchange Carrier
JBOD	Just a Bunch of Disks
JFS	Journaled File System
JIT	Just-In-Time
L2F	Layer 2 Forwarding
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network
LCN	Logical Cluster Number

LDAP	Lightweight Directory Access Protocol
LFS	Log File Service (Windows NT)
LFS	Logical File System (AIX)
LFT	Low Function Terminal
JNDI	Java Naming and Directory Interface
LOS	Layered Operating System
LP	Logical Partition
LPC	Local Procedure Call
LPD	Line Printer Daemon
LPP	Licensed Program Product
LRU	Least Recently Used
LSA	Local Security Authority
LTG	Local Transfer Group
LUID	Login User Identifier
LUN	Logical Unit Number
LVCB	Logical Volume Control Block
LVDD	Logical Volume Device Driver
LVM	Logical Volume Manager
MBR	Master Boot Record
MCA	Micro Channel Architecture
MDC	Meta Data Controller
MFT	Master File Table
MIPS	Million Instructions Per Second
MMC	Microsoft Management Console
MOCL	Managed Object Class Library
MPTN	Multi-protocol Transport Network
MS-DOS	Microsoft Disk Operating System
MSCS	Microsoft Cluster Server
MSS	Maximum Segment Size
MSS	Modular Storage Server
MWC	Mirror Write Consistency
NAS	Network Attached Storage
NBC	Network Buffer Cache
NBF	NetBEUI Frame
NBPI	Number of Bytes per I-node
NCP	NetWare Core Protocol
NCS	Network Computing System
NCSC	National Computer Security Center
NDIS	Network Device Interface Specification
NDMP	Network Data Management Protocol
NDS	NetWare Directory Service
NETID	Network Identifier
NFS	Network File System
NIM	Network Installation Management
NIS	Network Information System
NIST	National Institute of Standards and Technology
NLS	National Language Support
NNS	Novell Network Services

NSAPI	Netscape Commerce Server's Application
NTFS	NT File System
NTLDR	NT Loader
NTLM	NT LAN Manager
NTP	Network Time Protocol
NTVDM	NT Virtual DOS Machine
NVRAM	Non-Volatile Random Access Memory
NetBEUI	NetBIOS Extended User Interface
NetDDE	Network Dynamic Data Exchange
OCS	On-Chip Sequencer
ODBC	Open Database Connectivity
ODM	Object Data Manager
OLTP	OnLine Transaction Processing
OMG	Object Management Group
ONC	Open Network Computing
OS	Operating System
OSF	Open Software Foundation
OU	Organizational Unit
PAL	Platform Abstract Layer
PAM	Pluggable Authentication Module
PAP	Password Authentication Protocol
PBX	Private Branch Exchange
PCI	Peripheral Component Interconnect
PCMCIA	Personal Computer Memory Card International Association
PDC	Primary Domain Controller
PDF	Portable Document Format
PDT	Performance Diagnostic Tool
PEX	PHIGS Extension to X
PFS	Physical File System
PHB	Per Hop Behavior
PHIGS	Programmer's Hierarchical Interactive Graphics System
PID	Process Identification Number
PIN	Personal Identification Number
PMTU	Path Maximum Transfer Unit
POP	Post Office Protocol
POSIX	Portable Operating System Interface for Computer Environment
POST	Power-On Self Test
PP	Physical Partition
PPP	Point-to-Point Protocol
PPTP	Point-to-Point Tunneling Protocol
PReP	PowerPC Reference Platform
PSM	Persistent Storage Manager
PSN	Program Sector Number
PSSP	Parallel System Support Program
PV	Physical Volume

PVID	Physical Volume Identifier
QoS	Quality of Service
RACF	Resource Access Control Facility
RAID	Redundant Array of Independent Disks
RAS	Remote Access Service
RDBMS	Relational Database Management System
RFC	Request for Comments
RGID	Real Group Identifier
RISC	Reduced Instruction Set Computer
RMC	Resource Monitoring and Control
RMSS	Reduced-Memory System Simulator
ROLTP	Relative OnLine Transaction Processing
ROS	Read-Only Storage
RPC	Remote Procedure Call
RRIP	Rock Ridge Internet Protocol
RSCT	Reliable Scalable Cluster Technology
RSM	Removable Storage Management
RSVP	Resource Reservation Protocol
SACK	Selective Acknowledgments
SAK	Secure Attention Key
SAM	Security Account Manager
SAN	Storage Area Network
SASL	Simple Authentication and Security Layer
SCSI	Small Computer System Interface
SDK	Software Developer's Kit
SFG	Shared Folders Gateway
SFU	Services for UNIX
SID	Security Identifier
SLIP	Serial Line Internet Protocol
SMB	Server Message Block
SMIT	System Management Interface Tool
SMP	Symmetric Multiprocessor
SMS	Systems Management Server
SNA	Systems Network Architecture
SNAPI	SNA Interactive Transaction Program
SNMP	Simple Network Management Protocol
SP	System Parallel
SPX	Sequenced Packet eXchange
SQL	Structured Query Language
SRM	Security Reference Monitor
SSA	Serial Storage Architecture
SSL	Secure Sockets Layer
SUSP	System Use Sharing Protocol
SVC	Serviceability
TAPI	Telephone Application Program Interface
TCB	Trusted Computing Base

TCP/IP	Transmission Control Protocol/Internet Protocol
TCSEC	Trusted Computer System Evaluation Criteria
TDI	Transport Data Interface
TDP	Tivoli Data Protection
TLS	Transport Layer Security
TOS	Type of Service
TSM	Tivoli Storage Manager
TTL	Time to Live
UCS	Universal Code Set
UDB	Universal Database
UDF	Universal Disk Format
UDP	User Datagram Protocol
UFS	UNIX File System
UID	User Identifier
UMS	Ultimedia Services
UNC	Universal Naming Convention
UPS	Uninterruptable Power Supply
URL	Universal Resource Locator
USB	Universal Serial Bus
UTC	Universal Time Coordinated
UUCP	UNIX to UNIX Communication Protocol
UUID	Universally Unique Identifier
VAX	Virtual Address eXtension
VCN	Virtual Cluster Name
VFS	Virtual File System
VG	Volume Group
VGDA	Volume Group Descriptor Area
VGSA	Volume Group Status Area
VGID	Volume Group Identifier
VIPA	Virtual IP Address
VMM	Virtual Memory Manager
VP	Virtual Processor
VPD	Vital Product Data
VPN	Virtual Private Network
VRMF	Version, Release, Modification, Fix
VSM	Virtual System Management
W3C	World Wide Web Consortium
WAN	Wide Area Network
WFW	Windows for Workgroups
WINS	Windows Internet Name Service
WLM	Workload Manager
WOW	Windows-16 on Win32
WWW	World Wide Web
WYSIWYG	What You See Is What You Get
WinMSD	Windows Microsoft Diagnostics
XCMF	X/Open Common Management Framework
XDM	X Display Manager
XDMCP	X Display Manager Control Protocol
XDR	eXternal Data Representation
XNS	XEROX Network Systems
XPG4	X/Open Portability Guide

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 640.

- *IBM TotalStorage NAS 100 Integration Guide*, SG24-6913
- *Implementing the IBM TotalStorage NAS 300G, High Speed Cross Platform Storage and Tivoli SANergy!*, SG24-6278
- *IBM TotalStorage NAS Backup and Recovery Solutions*, SG24-6831-00
- *Managing IBM TotalStorage NAS with IBM Director*, SG24-6830-00
- *IP Storage Networking: NAS and iSCSI Solutions*, SG24-6240
- *Implementing IBM Director Management Solutions*, SG24-6188.
- *A Practical Guide to Tivoli SANergy*, SG24-6146
- *Tivoli SANergy Administrator’s Guide*, GC26-7389
- *Tivoli Storage Management Concepts*, SG24-4877
- *Getting Started with Tivoli Storage Manager: Implementation Guide*, SG24-5416
- *Using Tivoli Storage Manager in a SAN Environment*, SG24-6132
- *Tivoli Storage Manager Version 4.2: Technical Guide,* SG24-6277
- *Red Hat Linux Integration Guide for IBM eServer xSeries and Netfinity*, SG24-5853
- *AIX 5L and Windows 2000: Side by Side*, SG24-4784
- *Migrating IBM Netfinity Servers to Microsoft Windows 2000*, SG24-5854
- *Using TSM in a Clustered NT Environment*, SG24-5742
- *ESS Solutions for Open Systems Storage: Compaq Alpha Server, HP and SUN*, SG24-6119
- *Backing Up DB2 Using Tivoli Storage Manager*, SG24-6247-00

Other resources

These publications are also relevant as further information sources:

- Larry Peterson and Bruce Davie, *Computer Networks - A Systems Approach*, Morgan Kaufmann Publishers, 1996, ISBN 1558603689
- A. S. Tanenbaum, *Computer Networks*, Prentice Hall, 1996, ISBN 0133499456
- M. Schwartz, *Telecommunication Networks: Protocols, Modeling and Analysis*, Addison-Wesley, 1986, ISBN 020116423X
- Matt Welsh, Mathias Kalle Dalheimer, and Lar Kaufman, *Running Linux (3rd Edition)*, O'Reilly, 1999, ISBN 156592469X
- Scott M. Ballew, *Managing IP Networks with CISCO Routers*, O'Reilly, 1997, ISBN 1565923200
- Ellen Siever, et al., *Linux in a Nutshell (3rd Edition)*, O'Reilly, 2000, ISBN 0596000251
- Andreas Siegert, *The AIX Survival Guide*, Addison-Wesley, 1996, ISBN 0201593882
- William Boswell, *Inside Windows 2000 Server*, New Riders, 1999, ISBN 1562059297
- Paul Albitz and Cricket Liu, *DNS and BIND (4th Edition)*, O'Reilly, 2001, ISBN 0596001584
- Gary L. Olsen and Ty Loren Carlson, *Windows 2000 Active Directory Design and Deployment*, New Riders, 2000, ISBN1578702429
- *Microsoft Windows 2000 Professional Resource Kit*, Microsoft Press, 2000, ISBN 1572318082
- D. Libertone, *Windows 2000 Cluster Server Guidebook*, Prentice Hall, 2000, ISBN 0130284696
- *Microsoft Services for UNIX Version 2 white paper*, found at:

  `http://www.microsoft.com/WINDOWS2000/sfu/sfu2wp.asp`
- C. J. Date, *An Introduction to Database Systems (7th Edition)*, Addison-Wesley, 1999, ISBN 0201385902
- George Baklarz and Bill Wong, *DB2 Universal Database V7.1*, Prentice Hall, 2001, ISBN 0130913669

Referenced Web sites

These Web sites are also relevant as further information sources:

- IBM Storage
  http://www.storage.ibm.com/
- IBM TotalStorage
  http://www.storage.ibm.com/ssg
- IBM NAS
  http://www.storage.ibm.com/snetwork/nas/index.html
- IBM NAS reference information
  http://www.storage.ibm.com/snetwork/nas/library.html
- IBM TotalStorage NAS 100
  http://www.storage.ibm.com/snetwork/nas/100/index.html
- IBM TotalStorage NAS 200
  http://www.storage.ibm.com/snetwork/nas/200/index.html
- IBM TotalStorage NAS 300
  http://www.storage.ibm.com/snetwork/nas/300/index.html
- IBM TotalStorage NAS 300G
  http://www.storage.ibm.com/snetwork/nas/300g_product_page.htm
- Microsoft Technical Library
  http://www.microsoft.com/windows2000/techinfo/default.asp
- Microsoft Services for UNIX
  http://www.microsoft.com/WINDOWS2000/sfu/default.asp
- Tivoli
  http://www.tivoli.com/
- Tivoli Sanergy Support
  http://www.tivoli.com/support/sanergy
- Storage Networking Industry Association
  http://www.snia.org/
- Sysinternals Microsoft Tools
  http://www.sysinternals.com/
- Linux Documentation
  http://www.linuxdoc.org/
- Linux Kernel Resource
  http://www.kernel.org/
- Red Hat Linux
  http://www.redhat.com/

- SUSE Linux
  http://www.suse.com/index_us.html

How to get IBM Redbooks

Search for additional IBM Redbooks or Redpieces, view, download, or order hardcopy from the Redbooks Web site:

ibm.com/redbooks

Also download additional materials (code samples or diskette/CD-ROM images) from this IBM Redbooks site.

Redpieces are IBM Redbooks in progress; not all Redpieces become IBM Redbooks, and sometimes just a few chapters will be published this way. The intent is to get the information out more quickly than the formal publishing process allows.

IBM Redbooks collections

IBM Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the IBM Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W

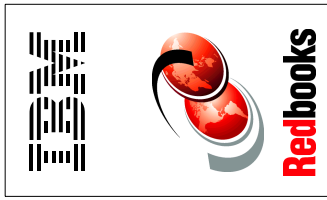

The IBM TotalStorage NAS Integration Guide

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

The IBM TotalStorage NAS Integration Guide

Learn all about the IBM TotalStorage NAS 100, 200, and 300

Implement high speed cross platform storage solutions

Configure NAS using this hands-on guide

This IBM Redbook describes how to integrate, install, and configure the very latest IBM TotalStorage Network Attached Storage 100, 200, and 300 in heterogeneous environments. Concrete scenarios for IBM DB2 UDB, Lotus Domino Server, Microsoft SQL Server, Microsoft Exchange Server, Oracle, and SAS illustrate the integration process for storage solutions.

The NAS 100, 200, and 300 are innovative Network Attached Storage (NAS) appliances that connect clients and servers on an IP network to storage. Their value is enhanced by their support of multiple protocols, allowing seamless file sharing across dissimilar platforms. They provide excellent Microsoft Windows performance that enhances client productivity while simultaneously protecting a customer's data and business continuity. This book shows how to integrate the units and explains how a company may benefit by utilizing these innovative solutions.

This easy-to-follow guide describes and explains NAS installation, ease-of-use, remote management, Microsoft Active Directory integration, expansion capabilities, high availability (clustering), and backup and recovery techniques. Other concepts, such as cross platform storage and methodologies for common data sharing for Linux/UNIX and Windows NT/2000 environments, are also covered.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6505-02 ISBN 0738429341