SG24-4786-00

OS/2 Warp Server, Windows NT, and NetWare: A Network Operating System Study

December 1996

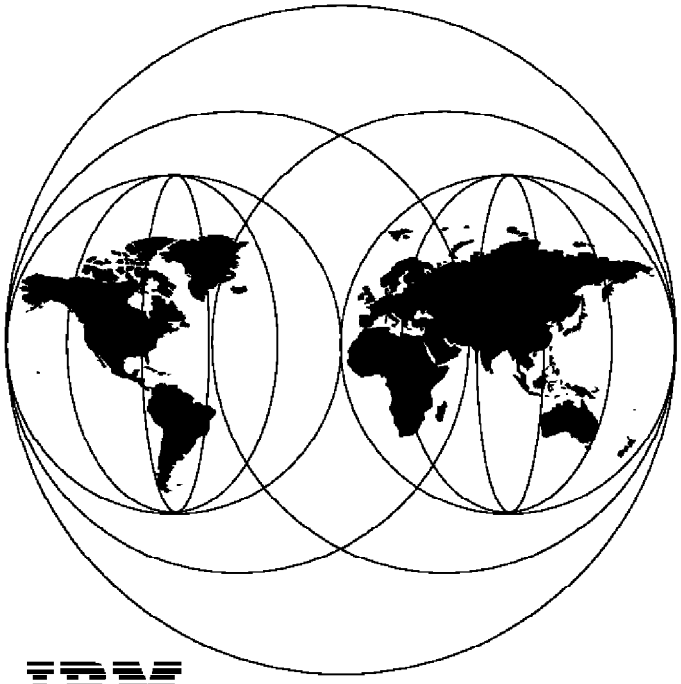

IBM

International Technical Support Organization

Austin Center

SG24-4786-00

International Technical Support Organization

OS/2 Warp Server, Windows NT, and NetWare: A Network Operating System Study

December 1996

IBML

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special Notices” on page 529.

First Edition (December 1996)

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. JN9B Building 045 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

Contents

Figures

Tables

Preface

With the release of OS/2 Warp Server, IBM has taken another step forward in the network operating system (NOS) race. In comparison to former LAN Server versions, IBM's inclusion of a cornucopia of management tools is a bold move by IBM to redefine just what network managers expect to find in a server.

IBM's highly reliable, highly available, and highly serviceable high-performance Warp Server network operating system has two main competitors, Microsoft's Windows NT Server 4.0 and Novell's NetWare 4.1. A blue ribbon group of IBM system engineers was invited to investigate these products and render an informative comparative analysis of their functions, features, and usability.

The goal of this examination was to produce a redbook designed for all kinds of readers. Whether you are an experienced network administrator of one platform who needs an overview of other network operating systems, a support person, office guru, or repair technician who might need more information about a particular network operating system, or just someone who is involved in network-related discussions as a sales person, consultant, or decider, you will find a lot of information about the three major Intel-based network operating systems, how they differentiate from each other, and how they interoperate with each other.

How This Redbook Is Organized

This redbook contains 550 pages and is organized as follows:

1. Part 1, “Theory and Design”
   - Chapter 1, “Architecture and Concepts”

     This chapter describes the three network operating system from an architectural point of view.
   - Chapter 2, “Principles of Administration”

     This chapter describes the pyramid of administration tasks.
   - Chapter 3, “Dynamic TCP/IP”

     This chapter gives an overview of dynamic TCP/IP solutions.
   - Chapter 4, “Systems Management”

     This chapter shows aspects of systems management.
   - Chapter 5, “Remote Access”

This chapter shows aspects of remote access services.

- Chapter 6, “Backup and Restore”

  This chapter introduces the backup and restore strategies.

2. Part 2, “Common Information”
   - Chapter 7, “Logon Interoperability: File and Print”

     This chapter provides information about clients and severs and their interoperability.

   - Chapter 8, “Security Issues”

     This chapter introduces security aspects for a network.

   - Chapter 9, “File Systems”

     This chapter discusses the most common file systems installed at clients and servers.

   - Chapter 10, “Performance, Scalability and Availability”

     This chapter introduces fault tolerance and other methods to increase availability of data and machines.

3. Part 3, “Configuring and Using Features”
   - Chapter 11, “Warp Server”

     This chapter discusses Warp Server-specific features that were discussed in theory in Part 1 of this book.

   - Chapter 12, “Windows NT”

     This chapter discusses Windows NT-specific features that were discussed in theory in Part 1 of this book.

   - Chapter 13, “Novell Directory Services”

     This chapter discusses NetWare-specific features that were discussed in theory in Part 1 of this book.

   - Chapter 14, “Introduction to Directory and Security Services”

     This chapter introduces the Warp Server add-on product DSS, which makes Warp Server more appealing to large account customers and to DCE fans.

4. Appendix A, “SOCKS (SOCKet Secured) Server”

   This chapter describes enabling SOCKS support. SOCKS support originally comes with Warp 4, but technically, it can be installed on Warp Server as well.

5. Appendix B, “Understanding Bridging and Filtering”

This chapter describes bridging and filtering functions built into the Remote Connection Services of OS/2 Warp Server.

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, Austin Center.

Uwe Zimmermann is an advisory technical support representative at the International Technical Support Organization, Austin Center. His areas of expertise include LAN/Warp Server, Windows NT, NetWare, and dynamic TCP/IP.

Rolf Berger is an advisory systems engineer at IBM Germany. He has been with IBM for 24 years. His areas of expertise include AIX, Windows NT, OS/2, and DOS/Windows.

Narendra Chhana is an information technology specialist at IBM South Africa. He has four years of experience in programming, supporting and administering large networks.

Franz-Stefan Hinner is a technical specialist at IBM Germany. With IBM for more than 10 years, his areas of expertise include desktop and server machines, DOS and OS/2, and defect and non-defect support.

Hermann Pauli is an advisory systems engineer at IBM Germany. He has eight years of experience in programming REXX, supporting OS/2 Warp Server and all OS/2-related products, such as Warp 4 and communications software.

Lucia Ronzoni is an information technology specialist at IBM Italy. She has four years of experience in networking and communications software. Her areas of expertise include OS/2 and OS/2-related products such as Communications Manager/2 and LAN Server.

Diana Scalmani is a systems specialist at IBM Italy. She has three years of experience in OS/2 and OS2-related products such as DB2 for OS/2 and LAN Server.

Karl-Heinz Wollert is a systems engineer at IBM Germany. He has 19 years of experience in large, mid-range, and PC systems. His areas of expertise include customer support, second-level support, PC dealer support, and large account pre-sales marketing support.

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

- IBM Austin

  Tom Arbuckle
  Marcus Brewer
  Oscar Cepeda
  Steve Dobbelstein
  Alexander Gregor
  Gary Hunt
  Steven King
  Susan Roberts
  Jim Wendelken
  Chuck McKelley

- IBM Raleigh

  Paul Chenger
  Barry Nusbaum

- Other IBM sites and other companies.

  Brice Bartec, Network TeleSystems, Sunnyvale, California
  Steven Perricone, Network TeleSystems, Sunnyvale, California
  Susan Richards, Vinca Corporation, Utah

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

redbook@vnet.ibm.com

Your comments are important to us!

Part 1. Theory and Design

Chapter 1. Architecture and Concepts

In this chapter we give an brief overview about Microsoft's Trusted Domain concept, NetWare's Directory Service, OS/2 Warp Server's Domain concept and the concept of IBM's Directory and Security Server for OS/2 Warp. A short comparison of these products is added. The goal of this chapter is to help readers to understand the principles and main differences between these concepts, rather than to discuss deep technical details.

This chapter assumes the reader is familiar with the basic domain concepts in a server network and with the principles of the Distributed Computing Environment (DCE).

1.1 The Domain Concept

In any server environment, a domain is the basic unit of security and centralized administration. It is one or a group of servers running a Network Operating System (NOS) that, in many ways, function as a single system. These servers form one administrative unit by centralizing and managing a master database of account information in one server (the domain controller).

In other words, a domain is a named network consisting of a group of workstations linked together to share resources such as directories, printers, modems and plotters. You logon to a server domain and gain access to shared resources which may be located on a number of server workstations in the domain. To the user it appears as though they are connected to a single server and they are unaware that they are accessing resources which may be located on different servers. They are presented with a *single system image*.

Each domain consists of the following types of workstations:

- There is always one, and only one, primary server workstation called the *domain controller* that maintains the master copies of the user and group definitions. Access permissions Access Control Lists) reside on each server where the shared resource is. The domain controller dictates how the shared resources are made available to the users when they log on.

  The domain controller processes user's logon requests and may also share it's own resources.

- Optionally, additional server workstations may be installed to provide shared resources and to serve as backup domain controllers. Backup

domain controllers support the domain controller in processing logon requests and can take over the role of the domain controller should it fail.

- Requesters, from which users can access shared resources on the server workstations.

This is a high-level overview of domains. For a more detailed discussion, please refer to the OS/2 Warp Server publication named *Up & Running!*, also see the Redbooks titled *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602, and *Inside OS/2 Warp Server, Volume 2: System Management, Backup/Restore, Advanced Print Services*, SG24-4702 (in press).

If an environment requires multiple domains, a user needing access to a broad range of resources would require a separate user account in each domain. This would also force the user to log on separately into each domain where the required resources exists. All of these additional user accounts would greatly add to the complexity of the administrator's job of maintaining appropriate user access and privileges. However, OS/2 Warp Server has an unique way of handling resources that reside on servers in different domains. Users can use those resources without being logged on to these domains. The only prerequisite that exists is that the users who need to work with external resources must be defined with the same user ID and password at the domains where the external resources reside.

Today's networks are becoming more and more complex. The resources, data, and information have to be shared among users on heterogeneous networks at multiple sites, on multiple platforms, and from multiple vendor applications. Networks have been installed independently by departments, workgroups, and business units. But the complexity has intensified with today's frequent mergers and acquisitions. Connections with external organizations such as vendors, customers, and business partners have increased the requirements for interoperability and security. On the other hand, impact to users should be minimized, and the productivity should be maximized.

The answer to this challenge is an Enterprise Directory Service. There are some specific requirements for such an directory service:

- **Single point of logon:** User accounts should only have to be created once, and users should be able to log in from anywhere on the network, being validated once by the network, and gaining access to resources wherever these resources are located.
- **Graphical management tools:** Pictorial tools should be included to help handle the complexity of a typical corporate directory.

- **Hierarchical directory:** It should be possible to place network resources or objects in groups.
- **Flexible directory structure:** It must be easy to make changes to corporate organizational structures or resource allocations.
- **Logical object naming:** Network resources should be defined logically, not in terms of location.
- **Consolidated object databases:** Strategically located global directory databases containing regularly updated information from all distributed servers are necessary to reduce the heavy network traffic generated by directory searches.
- **Distributed architecture:** Each disparate network group should be able to manage its own users and resources. Directory information should be stored throughout the LAN, rather than on a single central server.
- **Third-party software support:** The ultimate justification for adopting Enterprise Network Directory Services is a steady supply of directory-enabled applications.

To overcome these problems and to fulfil these requirements, NOS providers have developed different solutions, such as Microsoft's Trusted Domain concept, Novell's Directory Service and IBM's Directory and Security Server. These solutions are described in the following sections.

IBM's Directory and Security Server is an add-on feature for OS/2 Warp Server. OS/2 Warp Server itself has a certain kind of directory services as far as the use of aliases is concerned. Aliases point to the resource and to the resource's server. Using aliases makes it simple to access resources without knowing the physical location of the resource. OS/2 Warp Server is not discussed in too much detail in this chapter, and only OS/2 Warp Server, with IBM's Directory and Security Server installed, is compared to competitive products.

1.2 OS/2 Warp Server Domain

OS/2 Warp Server Version 4 is IBM's one-size-fits-all server operating system solution for customers ranging from small and medium-sized businesses to large enterprises. It combines a foundation for application serving with integrated file and print sharing, and offers an easy-to-use graphical user interface for drag-and-drop administration.

Following on the heels of OS/2 Warp Connect and OS/2 Warp 4, IBM's network client operating systems, OS/2 Warp Server, combines the market-proven quality of OS/2 Warp and LAN Server 4.0 with a wealth of

functional enhancements in systems management, backup and recovery, remote access, enhanced TCP/IP support, advanced print function, and LAN Internet access. All services are integrated into the product, eliminating the time and cost of having to separately install each component. However, services such as file and print can be selectively installed, allowing users to customize OS/2 Warp Server to meet their specific needs. The installation procedure also includes auto-detection of devices such as network interface cards.

OS/2 Warp Server inherits from LAN Server 4.0 a sophisticated set of network capabilities, including an easy-to-use drag-and-drop administration model that allows network administrators and resellers to quickly install, set up, configure, and manage a network. It offers tight security that is flexible enough to be customized to the needs of any business by assigning various privileges down to specific files on the server. OS/2 Warp Server also uses a powerful, high-performance file system and includes a NetWare migration utility that allows an organization to migrate NetWare 2.x and 3.x users and information onto an OS/2 Warp Server environment by using a graphical user interface.

OS/2 Warp Server possesses the same 32-bit, preemptive multitasking capabilities of IBM's powerful and battle-tested OS/2 Warp operating system, and comes Internet-ready with IBM's popular Internet Access Kit and WebExplorer. It offers reliable crash protection, runs OS/2 and DOS applications, and contains IBM's WIN-OS/2 code, which provides support for 16- and 32-bit Windows applications.

1.2.1 Resource Sharing

The main goal of a local area network is to share common resources within a workgroup. In OS/2 Warp Server this is done by File and Print Sharing Services.

The File and Print Sharing Services component of OS/2 Warp Server is a Local Area Network (LAN) application that is functionally equivalent to OS/2 LAN Server 4.0 with Service Pack IP08152 applied. It allows you to share hardware and software resources that are located on a server workstation.

You may share the directories (and the applications and files contained in them) and the printers and serial devices (such as a modem or plotter) that are connected to the server workstation. These shared resources are also referred to as network resources.

From a workstation, after you have connected to a network resource, you may use that resource in the same way you use local resources.

Servers can be grouped to domains. Once the user has logged on to a domain, he/she has transparent access to all resources within the domain. With this solution, every user in a workgroup has access to common used resources.

With this workgroup design, there are several limitations, which can be severe in an enterprise environment. Some of these limitations are:

- The administrator can only administer a single domain.

  Exception: Warp Server's cross-domain administration as described in 11.8, “Multiple Domain Administration” on page 230.

- The client has only access to a single resource domain.

  Exception: Warp Server's use of external resource aliases makes cross-domain access to resources possible as described in 11.9, “Creating Cross-Domain Resource Definitions” on page 232.

- Only 255 groups can be defined per server.
- Only 128 servers can be defined per domain.

This is no problem in a single LAN environment, but in an enterprise environment there is a need to access recourses in other domains. That can be resolved, for example, by registering every client in every domain, but that causes security problems and will be administrator's nightmare. See Figure 1 on page 8 for the current implementation of OS/2 Warp Server.

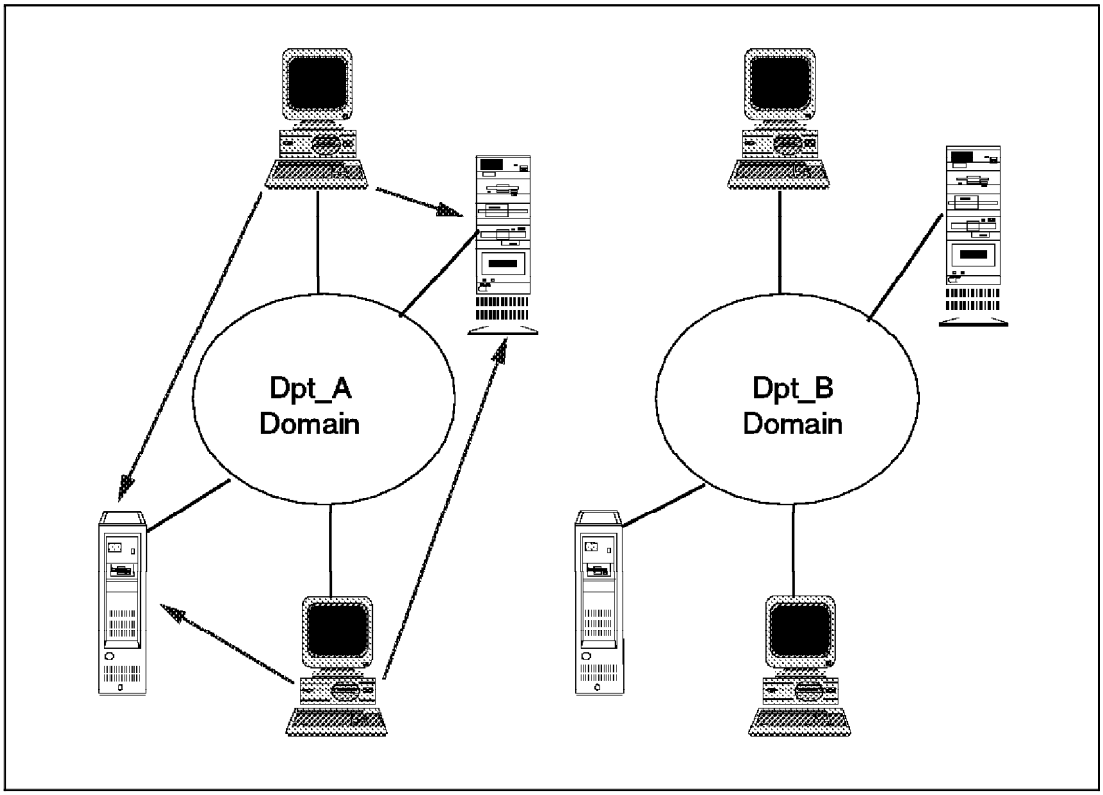

Figure 1. LAN Server and Warp Server Workgroup Design

OS/2 Warp Server supports the use of external resources. This means that a resource in a domain other than the home domain can be used. The usage of external resources requires that the user is defined in the other domain and that access rights are defined for that user. Alternatively the GUEST access can be chosen if, for example, the external resource is a networked printer and public access rights are defined for that resource. A user can log on to different servers, for example to a Warp Server domain and to a Novell NetWare server, and other services on the network, such as main-frame and mid-range computers, but this usually requires multiple logons and manual tracking of passwords by the user. To make life a little bit easier in this environment, OS/2 Warp Server offers the Network SignON Coordinator, which is described in the following section.

1.2.2 Password Coordination

The Network SignON Coordinator (NSC) Client provides the end user a way to perform a signon/signoff operation. The NSC Client can operate on either an OS/2 or DOS platform and manage passwords and logons in:

- OS/2 Warp Server Domains
- NetWare Servers
- Hosts

- Local facilities

These operations are specified in a user-configured ASCII file (NSC.INI) that contains the location definitions that may contain a user ID to be used for request processing. The same location can be defined as many times as is necessary to include all the user IDs that you have at that location.

Options that can be used in NSC are:

- A different user ID than the one the end user inputs can be used.
- You can specify that the user is to be logged on to a specific domain.
- You can specify an Exit Routine to be executed after Network SignON Coordinator performs the sign on/sign off operation. Network SignON Coordinator allows a user to sign off all services with a single command.
- You can change password across all defined domains in one operation. If the user selects the option to change passwords, the user is prompted to enter and confirm the new password. The password change is then initiated at all services defined in their Network SignON Coordinator configuration file.

Network SignON Coordinator provides additional functions and options to allow users to tailor the system to fit their needs. These functions and options include:

- Queueing requests to LAN Server domain controllers when they are not available
- The ability to specify different user IDs on each system while using the same signon password on every system
- An OS/2 API and toolkit that supports all of the functions of Network SignON Coordinator while bypassing the user interface
- User Exits for additional coordination or synchronization of signons, changing passwords, and signoffs
- Configuration options for user ID character set, minimum/maximum user ID length, and minimum/maximum password length

To summarize, Network SignON Coordinator is a tool for end users who, by entering their user ID and password once at a menu, have their signon requests processed at any number of OS/2 Warp Server domain controllers.

Network SignON Coordinator is not a security product; it is a productivity aid. However, since it does help the user manage passwords, some care has been taken to avoid creating additional security exposures for the user.

Caution

Review the following with respect to your security requirements. If any of these possible exposures is unacceptable, you should not use Network SignON Coordinator.

- Network SignON Coordinator assumes the user has the same password at all systems. If a user's password is compromised, the security exposure may be greater since all systems can be accessed with that password.
- Network SignON Coordinator can remember the user's password once it has been entered only if the SAVEPW option is configured. The default operation requires the user to reenter the password each time it is required.

  The password is always discarded when Network SignON Coordinator is terminated, even if SAVEPW is configured.
- Network SignON Coordinator does not keep passwords in the clear in memory except when necessary to call external application programming interfaces. The password is masked and distributed using a simple reversible algorithm designed to prevent casual viewing of the password.
- Network SignON Coordinator does not send passwords from Network SignON Coordinator Clients to Network SignON Coordinator Servers in the clear. The password is masked to prevent casual viewing of the password via network analyzers.
- Products supported by Network SignON Coordinator send the passwords across the network using different techniques. For information on how passwords are communicated by these products, consult the product information for that product.
- Network SignON Coordinator provides no function for restricting access to locations. Access to other locations is controlled by each location's own security facility.
- Network SignON Coordinator performs no encryption and is therefore not subject to any export restriction related to encryption. However, Network SignON Coordinator uses masking techniques instead.

1.3 Microsoft Trusted Domains

Windows NT Server 4.0 is an extension of Microsoft's earlier networking product, LAN Manager. In LAN Manager, each domain is an independent,

nonhierarchical database of account information. It does not have a mechanism to tie multiple, independent domain databases together. Microsoft attempted to overcome this design limitation by introducing the concept of *trust relationships* between domains in Windows NT. Trust relationships permit cross-domain administration, allowing users and groups in one domain to be assigned rights to resources in other domains that "trust" the user's home domain.

Microsoft calls this concept *Windows NT Server Directory Services*. This expression is a little bit misleading because Microsoft's Directory Services integration is not state-of-the-art, when compared with other solutions in the marketplace.

When trust relationships are properly established between the domains of a network, they allow a user to have only one user account yet access the entire network. All the computers on the network can recognize the user account. A user needs to log on and provide a password only once to access any computer on the network, meaning any Windows NT Server in the network.

A trust relationship does not grant users access to resources in trusting domains. Rather, it permits an administrator in the trusted domain to grant access rights to resources in the trusting domain. Only after a trust relationship is established between domains can an administrator grant rights to resources in the trusting domain. The principles of trust relationships are discussed in the following paragraphs.

1.3.1 Models for Building Domains

Four basic domain models can be created using the trust relationship facilities. Each Windows NT Server network originates from one of these four basic designs or models:

- **Single Domain model**

  As the name implies, this configuration consists of only one domain. There is only one domain controller with potentially multiple domain servers. In this case this single domain functions as both the account and resource domain.

Table 1. Advantages and Disadvantages of Single Domains	
Advantages	**Disadvantages**
• This is the best model for companies with few users and resources • It provides centralized management of user accounts. • There is no management of trust relationships necessary. • Local groups need to be defined only once.	• Poor performance occurs if the domain has too many users and groups. • There is no grouping of users into departments. • There is no grouping of resources. • Browsing is slow if the domain has a large number of servers.

- **Master Domain model**

  This configuration consists of several domains, one of which is the account domain. All user accounts reside within this domain. All other domains act as resource domains and trust the single *master* domain. Each domain has its own domain controller and may also contain multiple domain servers.

Table 2. Advantages and Disadvantages of Master Domains	
Advantages	**Disadvantages**
• It is the best choice for companies that do not have too many users and must have shared resources split into groups. • User accounts can be centrally managed. • Resources are grouped logically. • Department domains can have their own administrators who manage the resources in the department. • Global groups need to be defined only once (in the master domain).	• Poor performance occurs if the domain has too many users and groups. • Local groups must be defined in each domain where they are to be used.

- **Multiple Master Domain model**

  In the Multiple Master domain model, there are several master account domains, each containing a subset of the user account database. A complete trust relationship between these master account domains ensures that users may log on anywhere in the network. All other domains act as resource domains and will trust some or all of the

master domains. Each domain has its own controller and may also contain multiple domain servers as well.

Table 3. Advantages and Disadvantages of Multiple Master domains

Advantages	Disadvantages
• Good for many users and a centralized support department. • Resources are grouped logically. • Department domains can have their own administrators, who manage the resources in the department.	• Both local and global groups may have to be defined multiple times. • There are more trust relationships to manage. • Not all user accounts are located in one domain.

- **Complete Trust model**

  Finally, there is the Complete Trust model, which consists of several domains with each domain performing its own administration. No single domain exerts any control over the others, and they can all function as both account and resource domains. In this model, each domain has its own controller.

Table 4. Advantages and Disadvantages of Complete Trusts

Advantages	Disadvantages
• They are scalable to networks with a large number of users. • Each department has full control over its user accounts and resources. • Both resources and user accounts are grouped into departmental units.	• Because there is no central management of users, this model is not practical for companies with a central support department. • There is a very large number of trust relationships to manage. • Each department must trust that other departments do not put inappropriate users into global groups.

1.3.2 One-Way Trust and Two-Way Trust

Trust relationships can only be established between Windows NT Server domains. There are only two types of trust relationships possible between Windows NT Server domains:

- One-Way
- Two-Way (Reciprocal)

The basic *trust relationship* is only a one-way relationship. In a one-way trust relationship only one domain trusts the other. They do not both trust each other. Remote user accounts and global groups may be used from only one of the domains, the trusted domain. The one-way trust relationship in Figure 2 would allow a user residing within the trusted (account) domain to use resources such as files or printers available within the trusting (resource) domain.

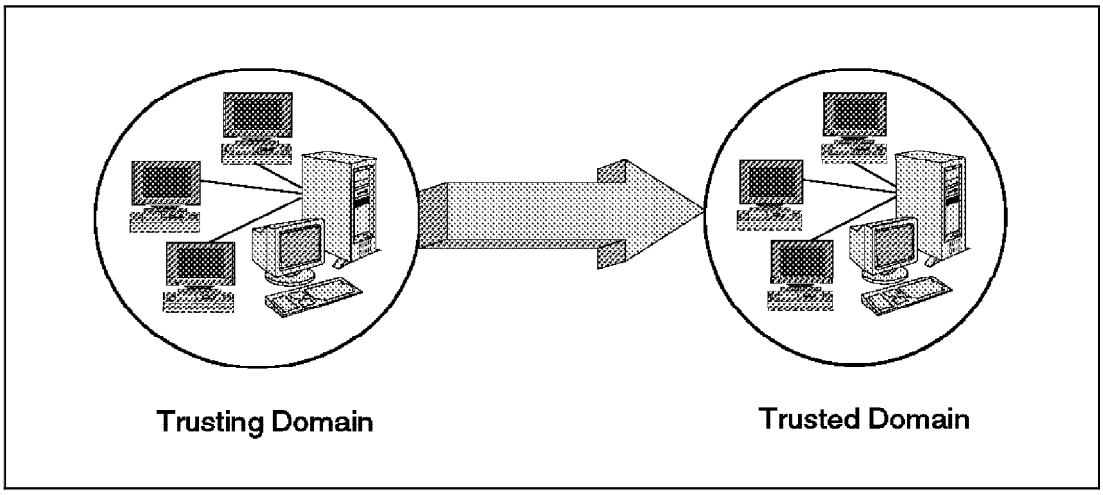

Figure 2. One-Way Trust

An account domain contains the user account database. Its primary job is to authenticate the logons, access rights, and privileges of users identified on the user account database. On the other side of a trust relationship is a resource domain. A resource domain contains resources such as workstations, file and print servers, or data directories. Most multi-domain environments will divide up domains into these two basic functional groups, but it is also possible for any single domain to function as both an account and a resource domain at the same time. This is possible when a two-way relationship is established.

The two-way trust is nothing more than two one-way trusts. In a two-way trust both domains trust each other equally. This allows users to log on from either domain to the domain that contains their account. Using this implementation, each domain can have both accounts and resources, and remote user accounts and global groups may be used from either domain to grant rights and permissions to resources in either domain. In other words, both domains are trusted domains. Figure 3 on page 15 shows the two-way trust.

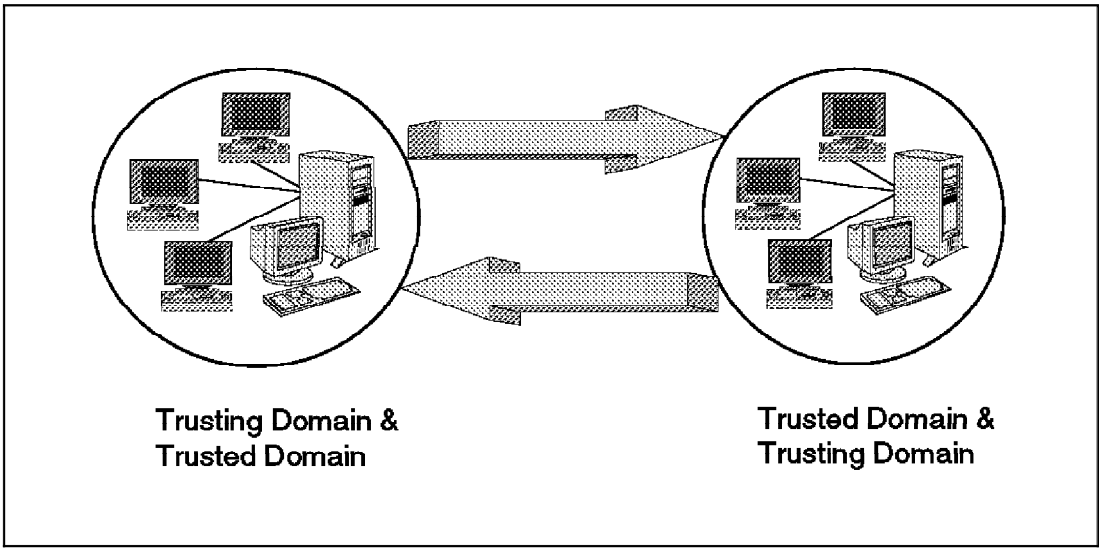

Figure 3. Two-Way Trust

An administrator can implement multiple trust relationships between multiple domains. Several resource domains can trust one account domain so that the single domain may contain all the user accounts, or one resource domain can trust several account domains so that the user accounts are spread among several account databases. In all cases, it is the functionality of the domains and not the physical location that is important. Users can log on *from* any trusting domain as long as they log on *to* a trusted account domain in which they have a valid account. This is true because of pass-through authentication.

1.3.3 Pass-Through Authentication

Pass-through authentication makes it possible for users to log on from machines or domains in which they have no user account. When a user logs onto a resource (trusting) domain, an access token containing the user's SID (security identifier) is passed on to the account (trusted) domain. Authentication of both the user's identity and password actually takes place within the account domain, hence the name pass-through authentication. This mechanism effectively allows a user to have an account in only one domain and yet access the entire network using trusted domains.

Pass-through authentication occurs in one of two circumstances:

- At initial logon from a workstation when a user is logging on to a trusted domain
- When connecting to a resource in a trusted domain

The trusted domain logon process is done in the following sequence (refer to Figure 4 on page 16):

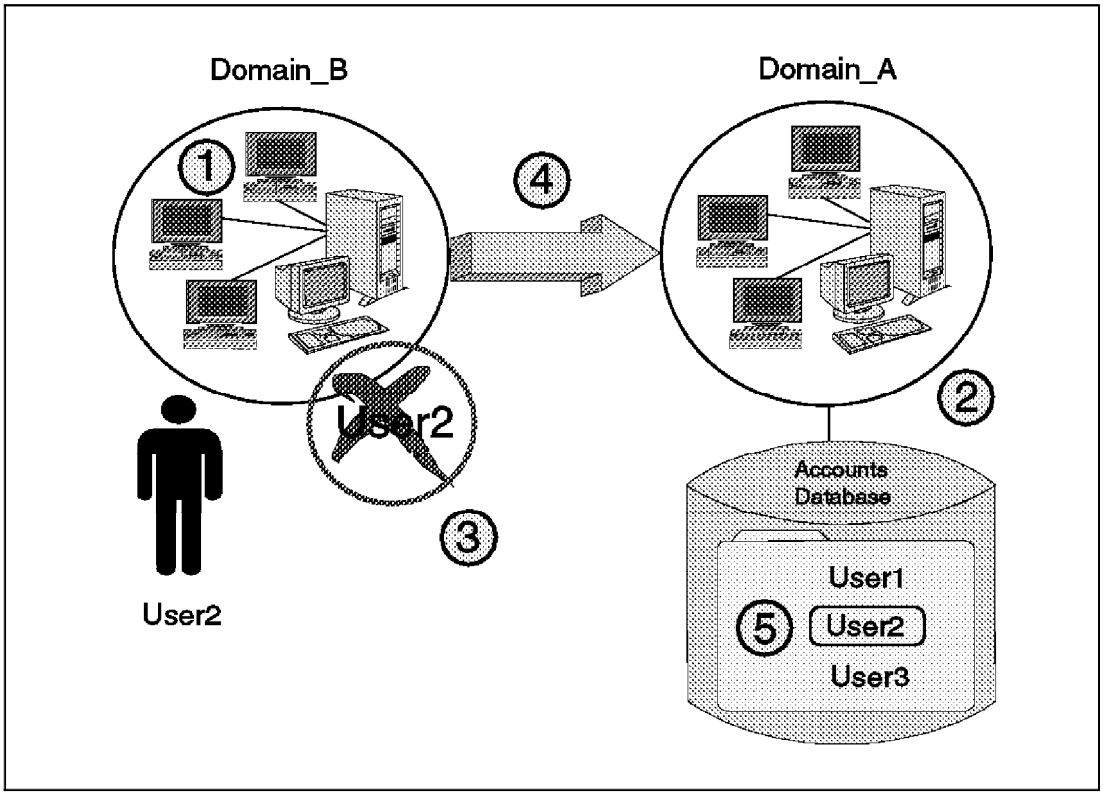

Figure 4. Pass-Through Authentication

1. The client's Net Logon service starts; it performs the process of locating a domain controller in its domain (Domain_B).
2. User2 attempts to log on at a computer in Domain_B with a user account from Domain_A by changing the **From...** entry in the logon dialog box to indicate Domain_A.
3. The domain controller in Domain_B cannot authenticate the request because the request is for a Domain_A user account.
4. The authentication request is passed through the trust to a domain controller in Domain_A. This domain controller checks Domain_A's account database for the existence of User2's account and for correct password information.
5. The domain controller in Domain_A authenticates User2's request and passes SID and group information about User2 back to the domain controller in Domain_B. The domain controller in Domain_B then passes the information to User2's client.

One very important point to consider is that pass-through authentication is not transitive. This means, it can only be used where a trust relationship exists. This has been designed into the system so that an administrator can

not accidentally establish a circular trust relationship that was never intended.

For example, in Figure 5 if Domain_A trusts Domain_B, and Domain_B trusts Domain_C, then Domain_A does not automatically trust Domain_C. A user with an account in Domain_C who attempts to log on while physically located in Domain_A will not be authenticated because pass-through authentication will not occur between Domain_A and Domain_C. Domain_C and Domain_A have to specifically set up their own trust relationship before pass-through authentication can occur between them.

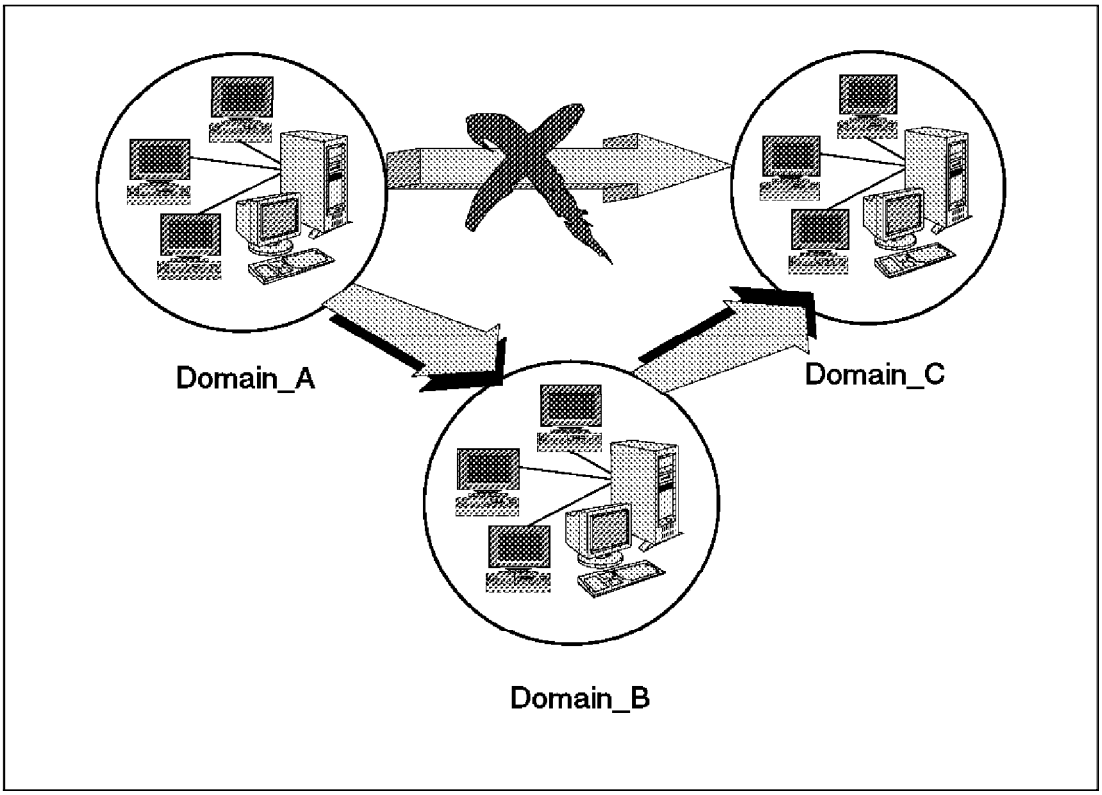

Figure 5. Intransitive Trust

1.3.4 The Complete Trust Domain Model

In the Complete Trust model, every domain on the network trusts every other domain. Each department manages its own domain and defines its own users and global groups. These users and global groups can be used on all domains in the Complete Trust model.

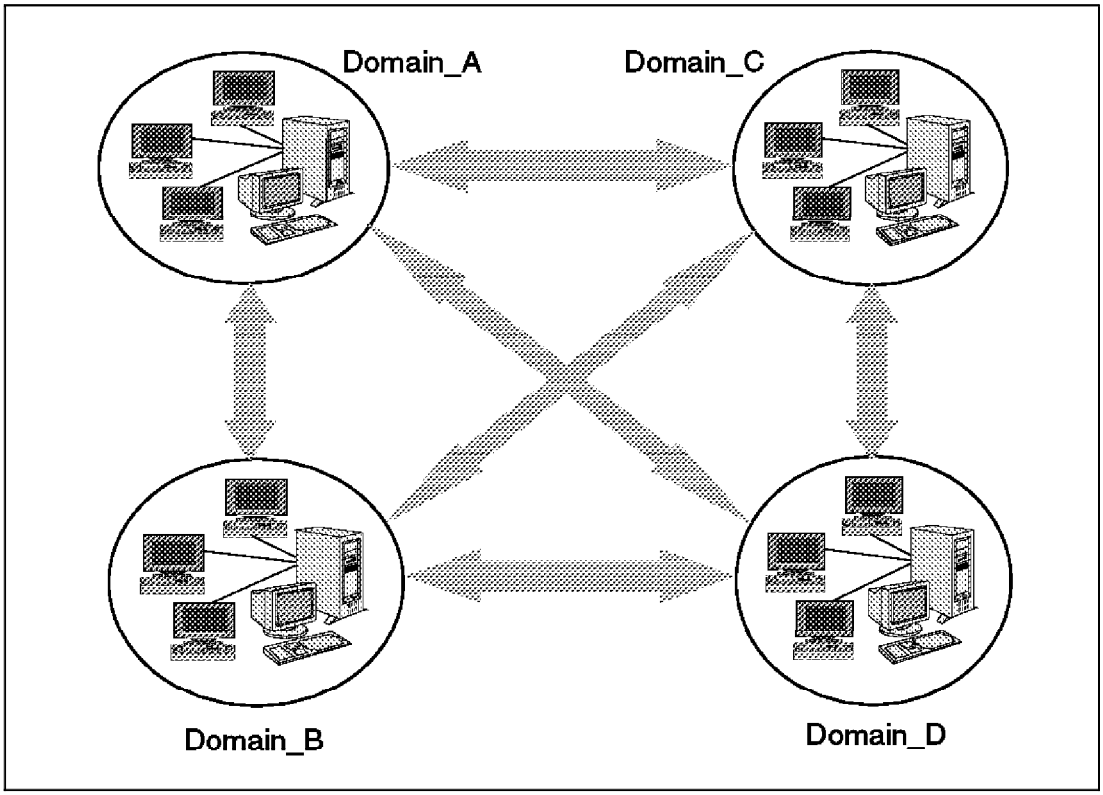

Figure 6. The Complete Trust Model

The Complete Trust model implies that there is no centralized security authority in the environment. Each domain is to be administered independently. Resources in one domain are granted based on mutual trust relationships. Unfortunately, the number of trust relationships that must be established and maintained grows geometrically as additional domains are added.

The number of trust relationships required for a company with *n* domains is:

```
n * (n-1)
```

For example, 4 domains require 12 trust relationships, and 20 domains require establishing 380 new trust relationships. Adding a new domain to an existing network of 10 domains requires establishing 20 new trust relationships.

Because there is no central administration in the Complete Trust model, users from other domains with access to resources could pose a security risk. This model requires a high degree of confidence in global groups from other trusted domains. Each department must trust that the other departments will not put inappropriate users into global groups.

1.4 Novell Directory Services

Novell Directory Services (NDS) is an information database service in NetWare 4 that organizes network resources such as users, groups, printers, volumes and other physical network devices, into a hierarchical tree structure. NDS has facilities for storing, accessing, managing, and using information about network resources and provides global access to all network resources regardless of where they are physically located, forming a single information system.

NDS treats all network resources, users, groups, printers, and volumes as individual objects in a distributed database known as the NetWare Directory Infobase. The database organizes resources in a hierarchical tree structure, independent of their physical location. Users and supervisors can access any network service without having to know the physical location of the server that stores the service. "Directory" means the global database provided by NetWare 4 servers.

The Directory replaces the bindery, which served as the system database in previous versions of NetWare. While the bindery supports the operation of a single NetWare server, NDS supports an entire network of servers. Instead of storing all information on one server that can be a single point of failure, information is distributed over a global database and accessed by all servers. NDS helps in managing directory resources such as NetWare users, servers, and volumes. Graphical and text utilities provide administrative functions for NDS and the file system.

This single, network-wide directory, a superset of the X.500 standard, is accessible from multiple points by users and supports multiple applications. NDS is an object-oriented implementation that allows users and administrators to build sophisticated naming schemes and databases across small or large networks.

It is very important to understand that the NetWare 4 directory structure and the file system (directories, files, applications) are separate, distinct hierarchical structures. Files and directories are *not* objects and are *not* in the NDS database. For example, trustee rights that are assigned in the directory to a Volume object *do not* flow down to directories and files in that volume.

Instead of logging in or attaching to individual servers, NDS users log in to the network. Users need only one password to gain access to all network resources available to them. Once the NDS name context is properly set, users can log in to the network by typing

```
LOGIN LSCOTT
```

instead of

```
LOGIN servername/LSCOTT
```

When a user accesses resources on the network, background authentication processes verify that the user has rights to those resources. Authentication allows a user to access any servers, volumes, printers, and so on in the network to which the user has rights. User trustee rights in the directory restrict the user's access within the network. Authentication is a means of verifying that a user is authorized to use both directory and file system resources. Authentication works in combination with the Access Control List (ACL) to provide network security.

1.4.1 NDS Structure

NDS allows an organization with multiple file servers to control user access to resources as if all resources were subsidiary to a single file server. However, instead of all resources being subsidiary to one file server, they are all subsidiary to one root of an NDS tree. The purpose of NDS is to organize all users and resources into an easy-to-manage hierarchical system. Figure 7 and Figure 8 on page 21 show the relative difference in network organizations under NetWare 3.x and NetWare 4.1.

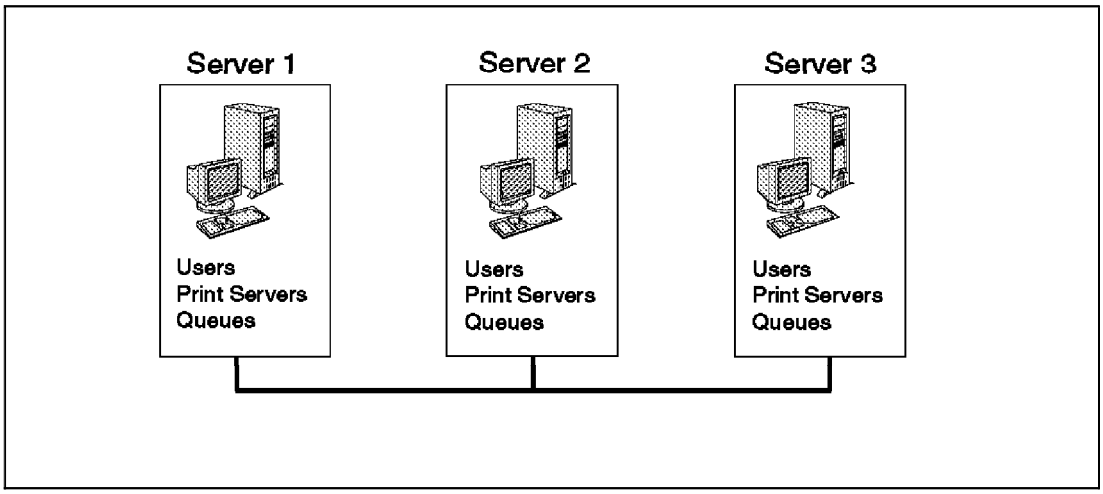

Figure 7. NetWare 3.x

In Figure 7, NetWare 3.x servers are organized as peers, in a flat relationship. Users defined on one server have access to objects defined in that server's bindery. For a user to access objects defined in another server, another user account must be defined in the second server's bindery.

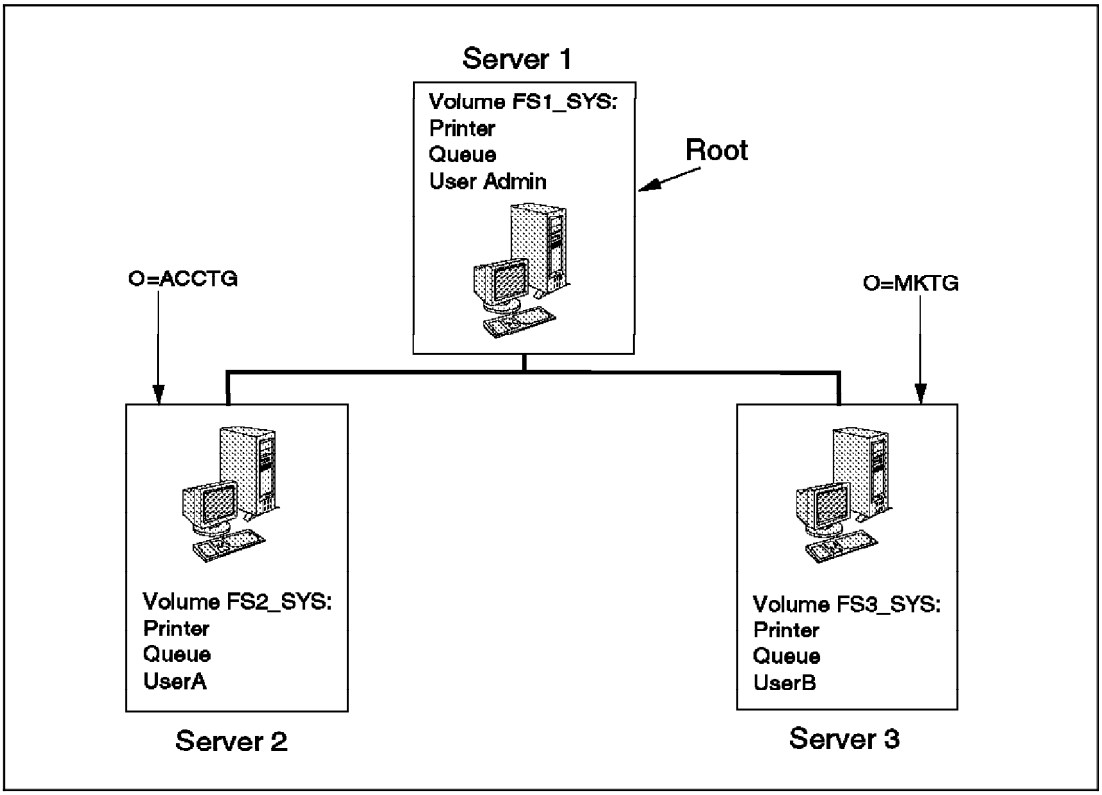

Figure 8. NetWare 4's NDS

In Figure 8, NetWare 4.1 servers are just one of several objects and are organized in an NDS tree superstructure. UserA is defined in the "context" `O=ACCTG` and can access only objects defined in the same context. Access to servers, volumes, printers, and print queues in other contexts, such as `O=MKTG`, is not permitted.

The Supervisor, however, can access all objects because Supervisor's context includes the ROOT, plus all directories `(O=MKTG, O=ACCTG)` lower in the tree. In this system, the user does not need to attach to multiple servers, nor must user accounts be defined on each server. The user is defined in the Global Replicated Database, and the user logs in to a context.

1.4.1.1 NDS Objects

NDS is *object oriented*. Physical devices are represented by *objects* or logical representations of physical devices. Users are logical user accounts and are one type of NDS object. One of the benefits of working with an object-oriented system is that moving a device does not change the object's definition. This makes system administration much easier.

The key terms used in NDS are:

- **Objects**

Objects are logical objects, representations of physical resources, users, and user-related entities, such as groups. For example, a User object is one of over 20 different NDS object types; a Printer object is another type of NDS object.

- **Object properties**

  Object properties are different types of information associated with an NDS object. For example, a User Login Script is one of 59 object properties associated with a User object.

- **Property values**

  Property values are simply names and descriptions associated with the object properties. For example, HP3 might be the property value for the Printer Name object property, which is in turn associated with the Printer object.

1.4.1.2 Tree Structure

NDS uses a hierarchical tree structure to organize the various objects. Hence the structure is referred to as the *NDS tree*. The tree is made up of these three types of objects:

- The [Root] object
- Container objects
- Leaf objects

The location in which objects are placed in a tree is called the *context* or *name context* (similar to a pointer in a database). The context is of key importance. To access a resource, the User object must be in the same context as the Resource object. A User object has access to all objects that lie in the same directory and in child directories.

The [Root] object is the top of a given Directory tree. Branches are made up of container objects, within them are leaf objects (see Figure 9 on page 23).

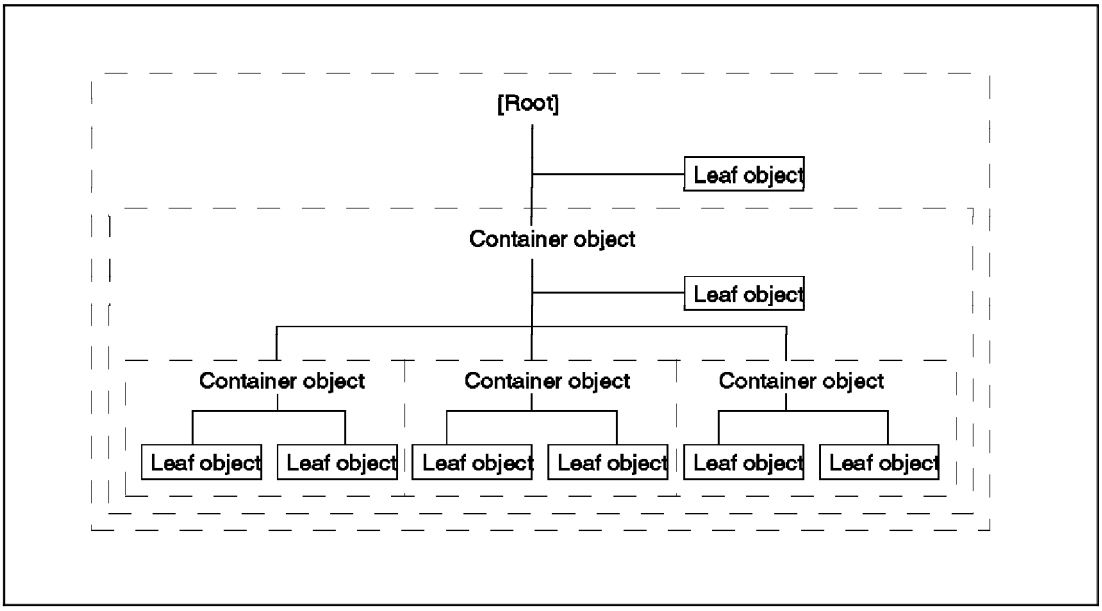

Figure 9. The NDS Tree

The [*Root*] *object* is created automatically when NDS is installed. It cannot be renamed or deleted. There can be only one [Root] object in a given NDS tree.

Container objects provide a way to logically organize other objects in the NDS tree. A Container object can house other Container objects within it. The top container is called the *Parent object*. Objects contained in a container object are *Child objects*.

There are three types of parents, or containers:

- Organization (O=)
- Organizational Unit (OU=)
- Country (C=)

There must be at least one *Organization* object within the NDS tree, and it must be placed one level below [Root]. The Organization object is usually used to denote a company or main organization.

Organizational Units are optional. If they are used, they must be placed one level below an Organization object or below another Organizational Unit. They can be used to denote divisions or departments within a company.

Organizational Units can be defined within Organizational Units to configure a deeper organizational structure. At a higher level, an Organizational Unit may represent a division of a company. Organizational Units contained in

the divisional Organizational Unit may represent departments within a division. An example is given in Figure 10 on page 24.

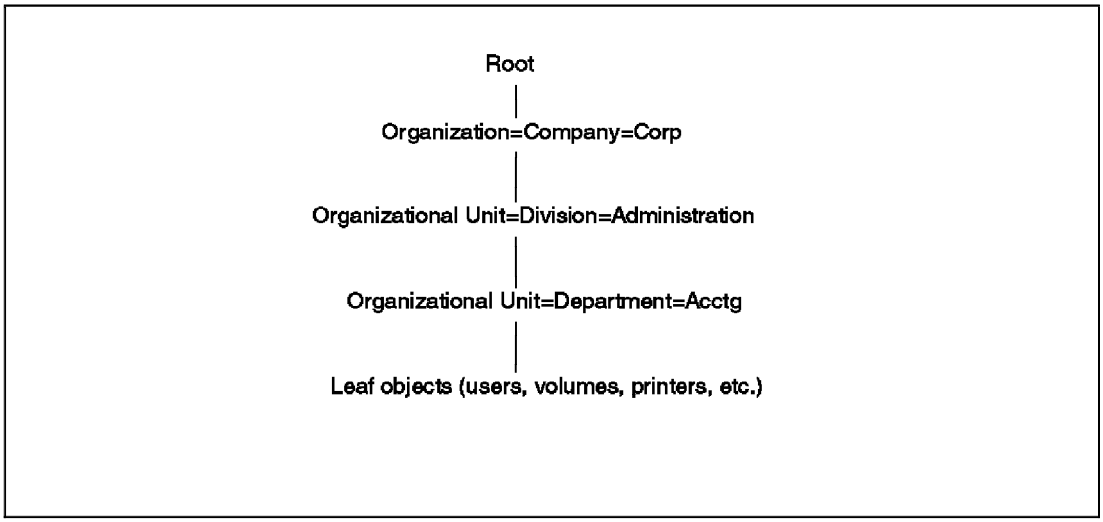

Figure 10. A Typical Company Tree

Because NDS is based on the CCITT X.500 specification, the use of *Country* Container objects (C=) is also supported. Country containers are located below [Root] and above Organization Container objects. Country Containers are useful for a multinational company.

Leaf objects are single-entity objects. They do not contain other objects. They correspond to actual physical entities such as users, servers, and printers. A Leaf object is denoted by CN= (Common Name).

Objects are either user-related or resource-related. Objects are all intended to provide users (user-related objects) with access to resources (resource-related objects). One Leaf object is put in a container to provide access to another Leaf object. Container objects are provided for the purpose of organizing Leaf objects.

Associated with each object is a set of *object rights*. Depending on the object rights assignment, a user may or may not have access to certain parts of the tree. Specifically, he may or may not have access to network resources (such as printers) in those parts of the tree. Object rights are NDS-based rights. When users are given access to a Volume object, they still must be granted file-system trustee assignments, which are separate from NDS rights.

Also associated with each object is a set of object properties. There are also rights associated with these properties, which are known as the *object*

property rights. These rights determine rights to view or edit the properties of objects.

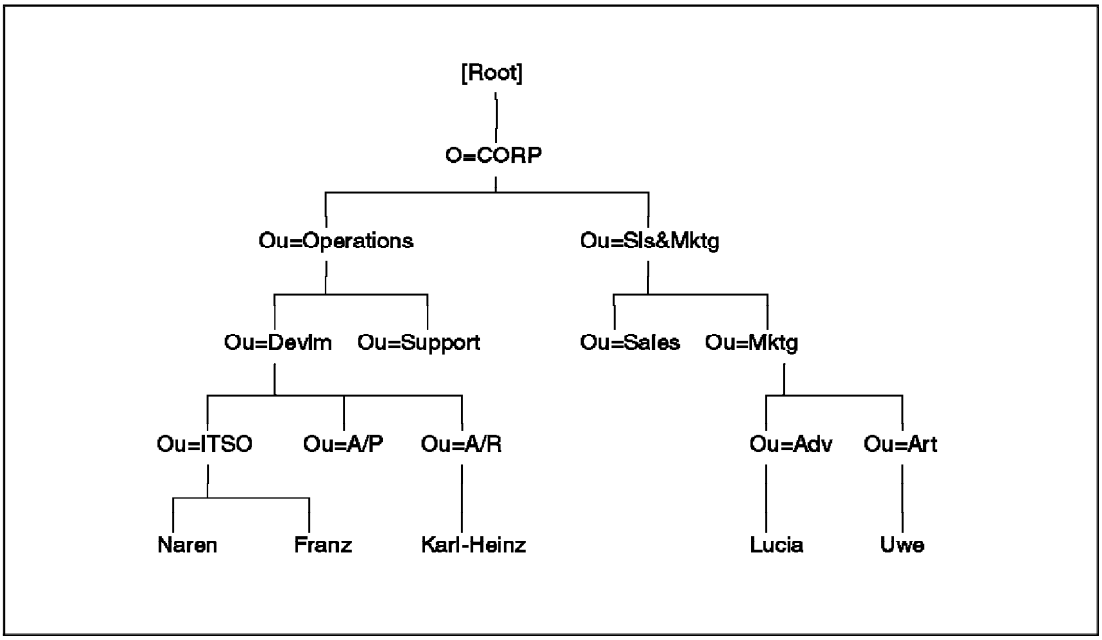

Figure 11. A Sample NDS Tree

1.4.1.3 Partitions

In large networks, the size of the complete directory services database can be too large to place on individual servers throughout the network. In some cases, it does not make sense to distribute or replicate the complete directory services database to other servers on the network. For example, directory services information for the engineering division may not need to be replicated to servers with information from the sales division. The database can therefore be divided into partitions.

Partitions in NDS are logical divisions of the NDS's Global Replicated Database. Each partition is a distinct portion of the Global Replicated Database and can be stored in different servers on the network. Each tree can have one or more partitions. The purpose of partitioning is to provide faster searches and more reliability.

An example of partitions is given in Figure 12 on page 26. The first server is installed in `OU=Development`. Because this is the first server in the NDS tree, a Root partition is created and a Master *replica* is placed on server LAB_SERVER. When the second server is installed in `OU=ITSO`, a new partition is created because the server is placed in a different context. The BUILDING49 partition is created. A Master replica is stored on BUILDING49.

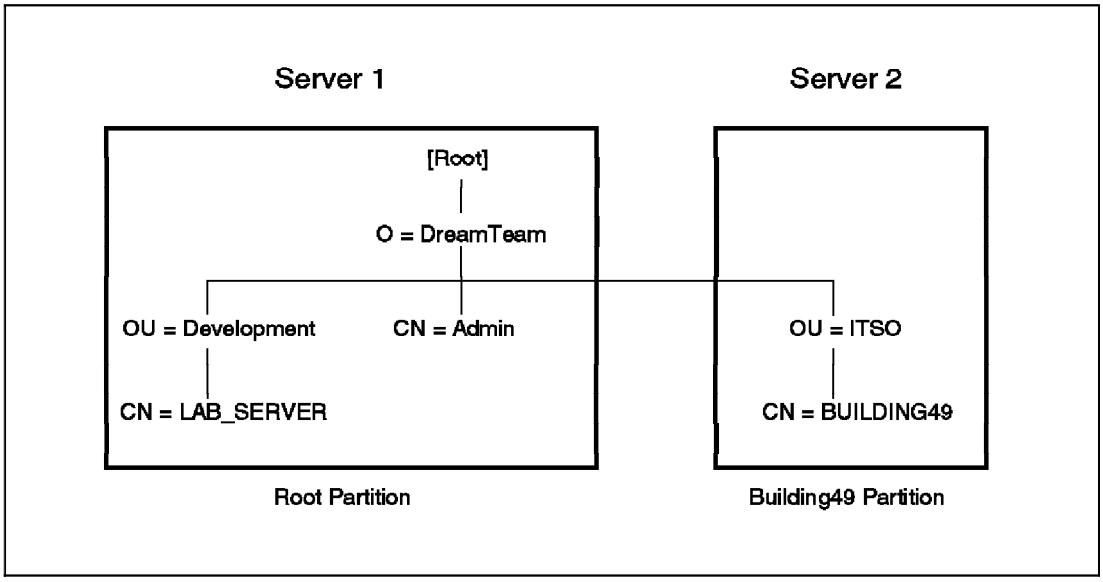

Figure 12. Default Partitions with Two Servers

By default, when a partition is created, a Read/Write copy (replica) of this partition is placed on the server containing the new server's parent context. In the example shown in Figure 12, a Read/Write replica of the BUILDING49 partition is placed on LAB_SERVER.

1.4.1.4 Replicas

Partition information that is copied and placed on other servers is called a *replica*. The group of replicas that exists for a particular partition is called the *replica list*. It is important to remember that a partition is a logical structure, and the replica is a physical instance, or copy, of the partition. There is no limitation to the number of replicas that a server can have. However, the goal is to place partitions in locations that make sense for quick authentication and reduced WAN traffic. Replicas also provide a layer of fault tolerance by allowing multiple servers to store copies of the same partition. If one server fails, another server can provide the directory information to clients.

A partition can be replicated in one of four forms:

- **Master replica**

  NDS designates one replica of a partition to be the master replica. Only the master replica can be used to create child partitions. It can also be used to create, modify, and delete other replicas.

- **Read/write replica**

  Clients can use this replica to create, modify, and delete directory entries.

- **Read-only replica**

  This replica responds to user requests. Clients can not use it to create, modify, or delete entries, but it synchronizes with the other replicas.

- **Subordinate reference**

  This replica links a parent partition and a child partition.

1.4.1.5 Synchronization

Directory synchronization is the process of replicating directory changes throughout the tree so that all directory services partitions always have the latest information. This process has two important aspects: maintaining consistency of information and granularity on synchronization. Maintaining consistency involves being able to ensure that information is synchronized even if a server that has a portion of the database is unavailable.

Synchronization granularity determines the level of information that is synchronized when the directory information is updated. For example, when an administrator changes a user's password, either the password information or the entire replica can be updated. Efficient granularity would dictate that only the password information be updated.

Because both Master and Read/Write replicas can alter the information in the directory, NDS must maintain synchronization for consistency among replicas. NDS ensures the integrity of the partition by circulating new and modified information among the replicas.

A distributed database has no typical model of consistency. NDS does not guarantee the consistency of replicas throughout the directory at any single point in time; however, all replicas eventually converge over time. This concept of replicas synchronizing over time is called loose consistency. The directory provides a *loose consistency* to accommodate high levels of partitioning and replication.

If updates are circulating through the system, clients that query the database can get different answers depending on the replica partition that responds. This is unlikely, however, since the database does not change frequently and because synchronization generally occurs within seconds of a database change.

Loose consistency also helps optimize directory performance. Because synchronization occurs less frequently than updates, changes that occur over a period of time are sent together, reducing wire traffic.

Another benefit of the NDS synchronization scheme is that there is no single type of replica (as with a master domain) that a client has to communicate

with to make an update. For example, an administrator that wants to disable a user account could do so on either a Master or Read/Write replica.

1.4.2 Summary

OS/2 Warp Server is a superior server system with an excellent domain concept for a workgroup environment. However, the problems in an heterogeneous environment are severe due to the lack of a real directory system. Therefore, IBM has developed the Directory and Security Server for OS/2, which is described in the following section.

1.5 IBM Directory and Security Server for OS/2 Warp

The IBM Directory and Security Server for OS/2 WARP (DSS) is a product that both extends the IBM OS/2 LAN Server from the workgroup environment to the distributed environment and delivers an OS/2 Warp implementation of the Open Software Foundation (OSF) Distributed Computing Environment (DCE), a state-of-the-art set of distributed services for building distributed applications across multiple platforms.

In cases where the only requirement is to install a "pure" DCE network for distributed computing in a heterogeneous environment, DSS can fill that requirement in a way that makes it easy to remotely manage the network from an OS/2 Warp Graphical User Interface (GUI), regardless of the platforms upon which the various DCE components reside.

When the need is to extend an existing OS/2 LAN Server 4.0 or OS/2 Warp Server network to allow the use of a global directory and global security, DSS can provide a powerful enhancement to OS/2 LAN Server 4.0 and OS/2 Warp Server, which enables them to take advantage of DCE's global directory and security services. DSS delivers this enhancement to OS/2 LAN Server 4.0 and OS/2 Warp Server as an add-on feature. DSS allows untouched, existing OS/2 LAN Server clients and servers to take advantage of the directory and security enhancements. This allows existing OS/2 LAN Server installations to be migrated on the customer's schedule.

The DSS OS/2 Warp implementation of DCE is fully-compliant with OSF DCE Version 1.1. It provides distributed directory, security, and time services as well as a remote procedure call interface for distributed application development. In addition, the DSS OS/2 LAN Server integration server feature, for OS/2 LAN Server 4.0 and OS/2 Warp Server domain controllers and additional servers, replaces the OS/2 LAN Server directory and security services with open, scalable DCE directory and security services. This allows existing OS/2 LAN Server and OS/2 Warp Server clients on any

supported platform to seamlessly access resources across domain boundaries by using a single identification and password. Administrators no longer need to maintain a separate definition for each user in every domain that they must access.

DSS contains several components that can each be installed on the same or different computers. The following components are all part of the DSS package:

DCE Directory Server: This is a standard OSF DCE 1.1 Cell Directory Server (CDS). CDS enables DSS clients and DCE application programs to locate objects in a DCE or DSS network. This component can be installed on the OS/2 Warp Server version of OS/2 Warp plus the latest OS/2 Warp FixPak applied to it. If it is installed on a single machine in combination with the DCE Security Server, the OS/2 LAN Server integration server feature and OS/2 LAN Server 4.0, it can be installed on previous versions of OS/2 Warp which has the latest OS/2 FixPak applied to it.

DCE Security Server: DCE uses a Kerberos third-party security service for authenticating both application clients and application servers. In addition, DCE Access Control Lists (ACLs) allow the owners of resources to determine who is allowed to access the resources. DSS provides a security server with full OSF DCE 1.1 support, including the use of extended registry attributes, to make it easy to integrate existing applications with DCE. This component can be installed on the OS/2 Warp Server version of OS/2 Warp plus the latest OS/2 FixPak applied to it. If it is installed on a single machine in combination with the DCE Directory Server, the OS/2 LAN Server integration server feature, and OS/2 LAN Server 4.0, it can be installed on previous versions of OS/2 Warp, plus the latest OS/2 FixPak applied to it.

DCE Client: The DCE client provides the Remote Procedure Call (RPC) interface that enables distributed application support on heterogeneous platforms. Both application clients and application servers run on the DCE client. The DSS DCE client is fully OSF DCE 1.1-compatible. In addition, it has an enhanced installation that allows a slim version of the DCE client to be installed on machines that only run application clients. This feature does not affect the interoperability of the DCE client. Machines that run application servers contain the full OSF DCE client. The DCE client also includes a Graphical User Interface (GUI) that can be used to administer any DCE component in the system. The administration GUI makes it easy to administer DCE components from IBM and other vendors and can be used to remotely administer DCE components on platforms other than OS/2 Warp. This component can be installed on OS/2 Warp plus the latest OS/2 FixPak applied to it.

DFS Client: This is a fully OSF DCE 1.1-compatible implementation of the DCE Distributed File System (DFS) client. It is compatible with DFS server implementations on AIX and other platforms. This component can be installed on OS/2 Warp plus the latest OS/2 FixPak applied to it plus the DSS DCE client.

OS/2 LAN Server Integration Server Feature: The OS/2 LAN Server integration server feature allows the OS/2 LAN Server directory and security services on either OS/2 LAN Server 4.0 or OS/2 WARP Server domain controllers or additional servers to be replaced with the DCE directory and security services. This component can be installed on an OS/2 Warp Server domain controller or an additional server plus the latest OS/2 FixPak applied to it. It can also be installed on an OS/2 LAN Server 4.0 domain controller or an additional server plus OS/2 Warp plus the latest OS/2 FixPak applied to it. See Figure 13 for a schematic view of the integration feature.

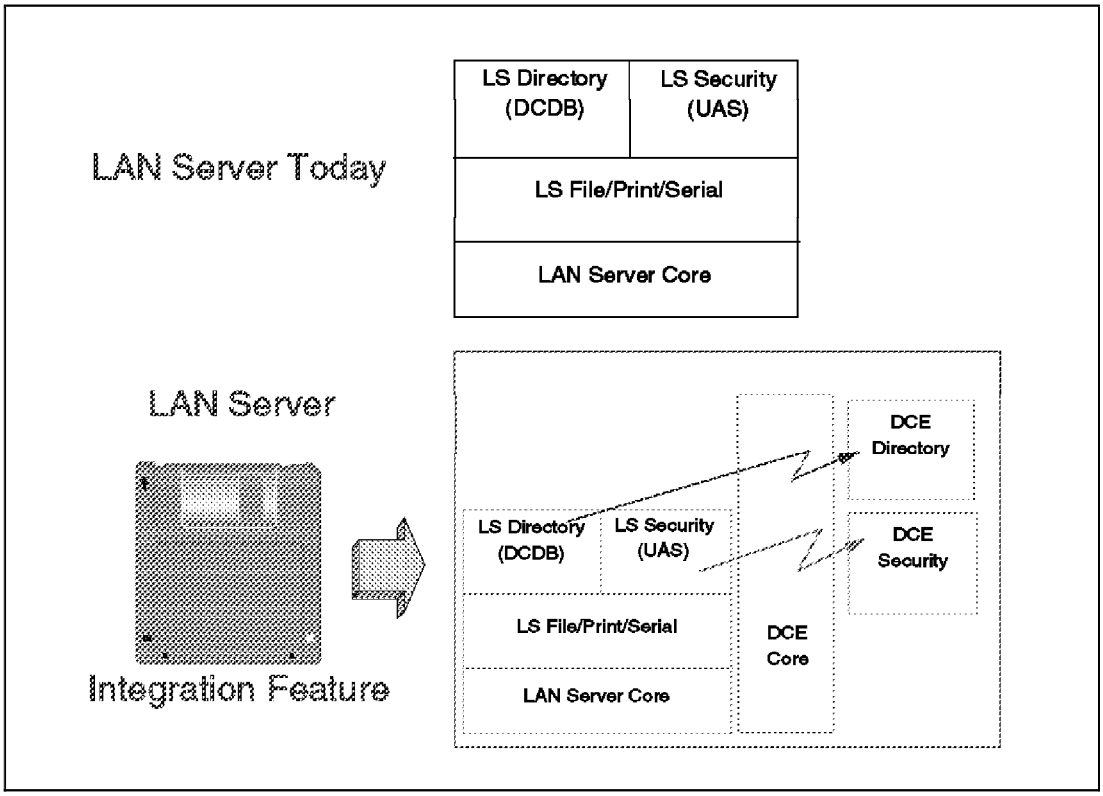

Figure 13. LAN Server Integration Feature

DSS Client: This component contains an OS/2 Warp Server OS/2 client and a Directory and Security Server slim DCE client, as well as function to allow them to work together. The DSS client also contains an administration GUI. This GUI is a superset of the DCE administration GUI, and it can be used to administer both the DCE and OS/2 LAN Server components of Directory and Security Server. There is also a slimmer user GUI that can be used to

perform tasks such as password change on DSS clients that are not used for full administration functions. This component can be installed on OS/2 Warp plus the latest OS/2 FixPak applied to it.

The Directory and Security Server delivers a great deal of distributed system power to the OS/2 LAN Server and OS/2 Warp Server networks. With this power, system and network administrators can perform some new tasks that were not possible prior to the introduction of DSS. These tasks involve the following new concepts and new servers which are described in greater detail later in this chapter.

1.5.1 The Directory and Security Server Cell

The key Directory and Security Server concept is that of the DSS *cell*. The cell is the basic DSS administrative unit. Although it can be compared to an OS/2 LAN Server domain, the DSS cell is much larger in scope and capacity than an OS/2 LAN Server domain. While an OS/2 LAN Server domain is usually used to define all of the resources and users for a single workgroup, a DSS cell can be used to define all of the resources and users for a line of business, a region, or an entire company.

1.5.2 Directory and Security Server Core Servers

DSS introduces the following new server types that are used to perform the basic cell functions, such as security and directory services. At least one of each of these core servers is part of every DSS cell.

Security servers: These servers are used to perform security operations, such as authenticating users to ensure that they are who they say they are. DSS security servers contain the database of user identities for the cell. Because cells can be spread over a wide geographic area, the DSS registry database can be *replicated* on multiple security servers to improve performance and availability.

For example, the main corporate office might contain the primary or *master* database, while each branch office might contain a local copy or *replica*. These replicas are read-only copies of the database. Updates must be made to the master.

Directory servers: Clients must be able to find servers and resources that they want to use, no matter where they exist in the cell. Directory servers allow them to do that. The directory server contains a database of definitions for all of the resources and services in the cell. When a client asks to use a resource or service, the directory server provides the address. Like the registry database, the directory database can be replicated. That is, local directory servers (those which are physically close to the client) can

have a copy of the directory database. In addition, the directory database can be *partitioned* or divided into sections so that each directory server need only have a copy of the definitions for those resources that exist locally. Even if a local directory server contains only a partition, users can still transparently access resources that are defined on other directory servers when the need arises.

Time servers: These servers are used to ensure that operations performed on multiple computers are synchronized so that the first operation is performed before the second and so on. They do this by synchronizing the system time on all of the computers in the network.

1.5.3 OS/2 LAN Server and the DSS Cell

When existing OS/2 LAN Server installations are converted to Directory and Security Server installations, many OS/2 LAN Server domains are combined to form a single DSS cell. Once this is done, clients can seamlessly access any resources for which they have authorization in any domain in the cell. This is possible because users are defined only once, in the DSS *cell registry*. The cell registry contains the user definitions and passwords for every user in the cell. Users have only one password to change, and administrators have only one definition to maintain for each user, regardless of the number of OS/2 LAN Server domains that are combined in the cell. Prerequisite for this is that all domain controllers and servers have DSS installed on their systems. If you also have clients with DSS installed, you can run DCE applications on all machines, regardless whether they are domain controllers with DSS installed, or servers, or clients (also AIX/UNIX workstations) with DSS or DCE installed.

For added administrative flexibility, the Directory and Security Server cell can be divided into several smaller administrative units that can have a hierarchical administrative relationship. These smaller administrative units are called *resource domains*. Resource domains can be used to contain the OS/2 LAN Server resources (for example, printers, modems, shared files) for a single department or branch office. They also make it easier to combine existing OS/2 LAN Server domains into a DSS cell.

When existing OS/2 LAN Server installations are migrated into Directory and Security Server cells, each existing OS/2 LAN Server domain becomes a resource domain in the cell. The administrator of the OS/2 LAN Server domain can continue to administer the resources that were part of that domain by becoming an administrator of the resource domain to which the OS/2 LAN Server domain was migrated. Because resource domains can have a hierarchical administrative relationship, however, new and more powerful and flexible administrative structures are now possible.

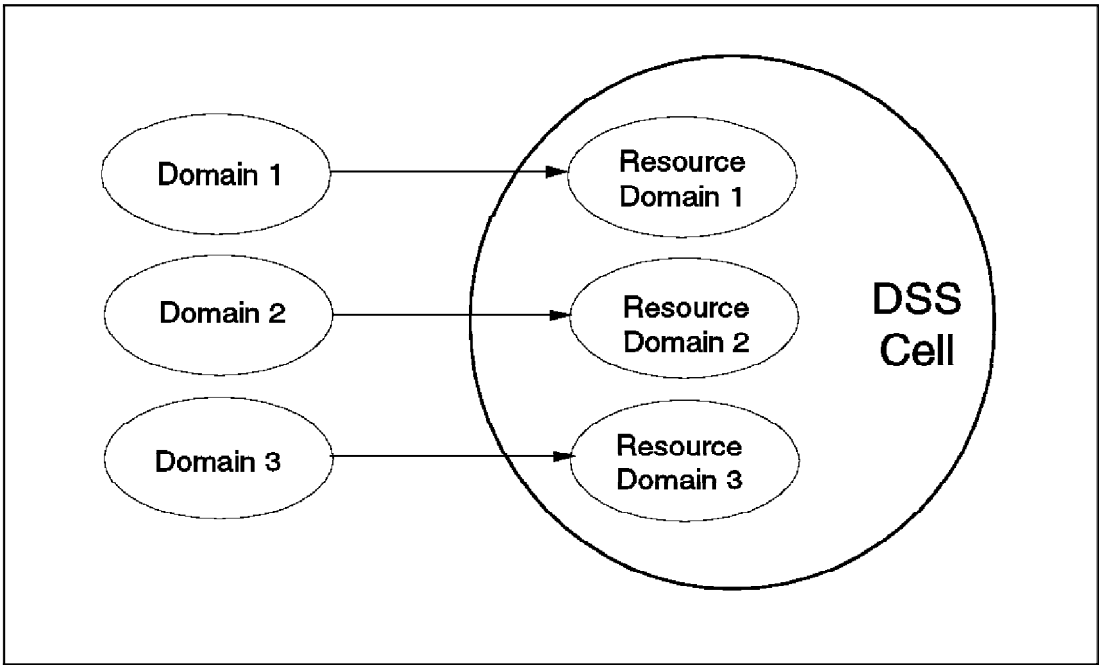

Figure 14. Domain Consolidation

In Figure 14 multiple domains are consolidated into one cell. The domains become resource domains in the cell. There is no change from the user's point of view; he/she can access resources as he/she always has. But now he/she can have access to all resources in the cell. Most advantages are in the administration area. Resource domains can include other resource domains. The administrator can administer single domains (no change to the Workgroup concept) and can administer the global domain. Figure 15 on page 34 gives an example. There are two resource domains (Region_A and Region_B) that include other resource domains.

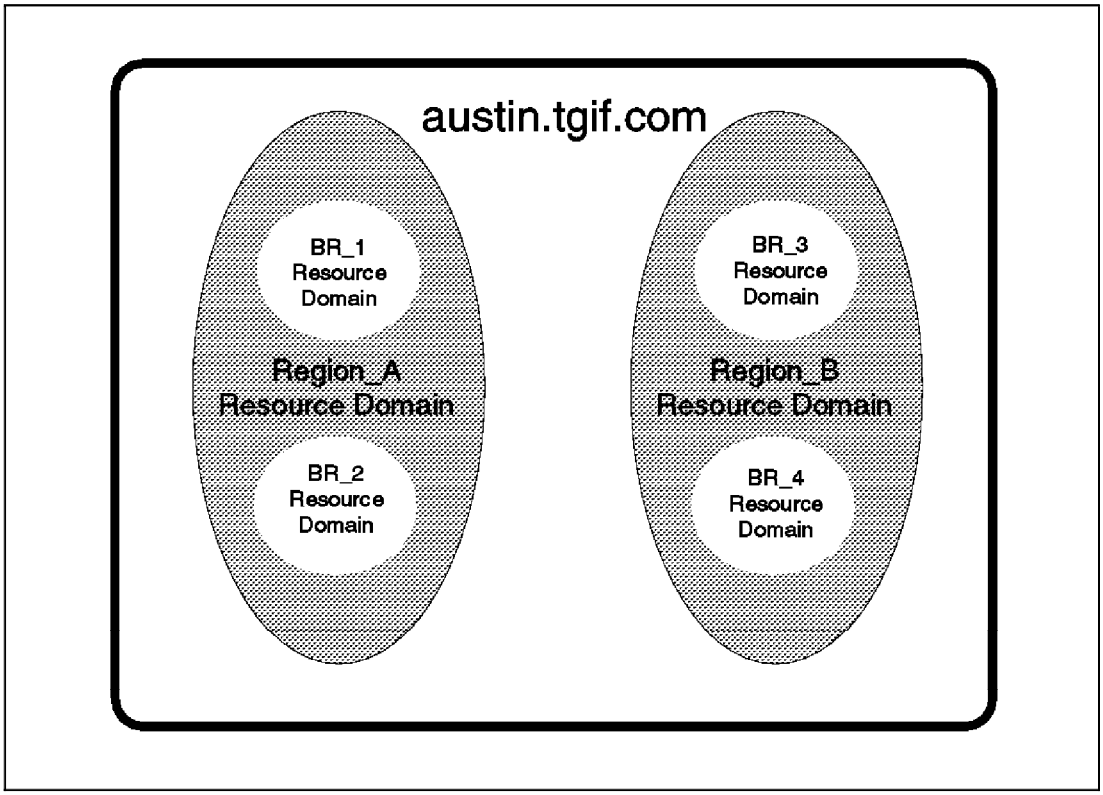

Figure 15. Resource Domains

There are several advantages with this concept:

- Local administration of local resources
- Central administration of users and groups
- Hierarchical administration relationships
- Familiar paradigm for OS/2 Warp Server administrators and users
- Single ID/password works in all resource domains

For example, in a situation where a company has several branch offices, each of which was an OS/2 LAN Server domain with its own administrator, it is now possible to migrate each of those domains into resource domains in a single DSS cell. The branch office administrators can continue to administer the branch office resource domains. If they need help or if there is a need to manage the branch office resource domains remotely on off-shifts or weekends, an administrator at a regional office or central site can transparently take over that responsibility without affecting the ability of the branch office administrator to administer the resource domain during normal business hours.

1.5.4 The Role of the Directory and Security Server Domain Controller

In OS/2 LAN Server, the domain controller is used to allow a client to log on to a domain, rather than to a single server, and to access resources to which the client is authorized that are located anywhere in the domain. This works because the domain controller contains a database (NET.ACC) that holds all of the user definitions for a single OS/2 LAN Server domain and another database (the domain control database) that contains all the resource definitions for the OS/2 LAN Server domain. Other file and print servers in the domain, *additional servers*, depend upon the domain controller to provide them with user identification information.

The NET.ACC file is similar to the Directory and Security Server registry database that contains all of the user definitions for the DSS cell. In the same way, the domain control database is similar to the DSS directory database.

When an OS/2 LAN Server domain is migrated to a Directory and Security Server cell by installing DSS on the domain controller, the additional servers can be unaffected. They still ask the domain controller for user identification and resource information. Only the domain controller knows that the master database has been moved to DCE and the requested information is really in the DSS databases instead of the NET.ACC file and the domain control database. For this reason, the DSS domain controller contains functions to synchronize the DSS registry and directory databases with the OS/2 LAN Server NET.ACC and domain control database. This allows unchanged OS/2 LAN Server additional servers to participate in the DSS cell. This synchronization is done on a resource domain boundary. For each resource domain, the administrator can decide whether that resource domain should be synchronized. The administrator can also decide which user and group definitions should be synchronized.

1.5.5 Installation Configurations

One of the key design points for the Directory and Security Server was to ensure that existing OS/2 LAN Server and OS/2 WARP Server installations can be migrated at the customer's pace. For this reason, it is possible to migrate existing OS/2 LAN Server and OS/2 Warp Server domains to a DSS cell by simply installing DSS on top of OS/2 Warp Server, or OS/2 LAN Server 4.0 plus OS/2 Warp, on each domain controller that is to be migrated.

The Directory and Security Server allows customers to migrate OS/2 LAN Server and OS/2 Warp Server installations in the manner that works best for each individual customer situation. All of the domain controllers can be migrated at once, with additional servers migrated later as needed, or all of the servers in a domain, both the domain controller and the additional

servers, can be upgraded when the domain is migrated to a DSS cell. Similarly, existing OS/2 LAN Server and OS/2 Warp Server clients can be used indefinitely, or they can be migrated to OS/2 DSS clients.

1.5.5.1 Domain Migration Recommendations

It is recommended that administrators:

- Migrate existing LAN Server domains to the DSS cell. Migration to a DSS cell requires that the domain controller of each domain be upgraded to a DSS domain controller. This migration creates a *single-user image* environment that greatly simplifies user, group, and password management.
- Migrate and maintain existing LAN Server domains in separate resource domains. This creates an environment similar to the existing LAN Server environment. Changes to the resource domain structure can be made after you become familiar with the DSS features and functionality.
- Set the resource domain name and broadcast address to that of the migrated LAN Server domain; this is the default during DSS domain controller installation. These settings ensure that existing LAN Server clients and servers continue to function without change and ensure that DSS clients and existing LAN Server clients reference the resource domain with the same name.

Two migration examples are given in the following paragraphs.

1.5.5.2 DSS on Domain Controllers

In Figure 16 on page 37 DSS is only installed on the Domain Controllers; Additional Servers and clients are unchanged.

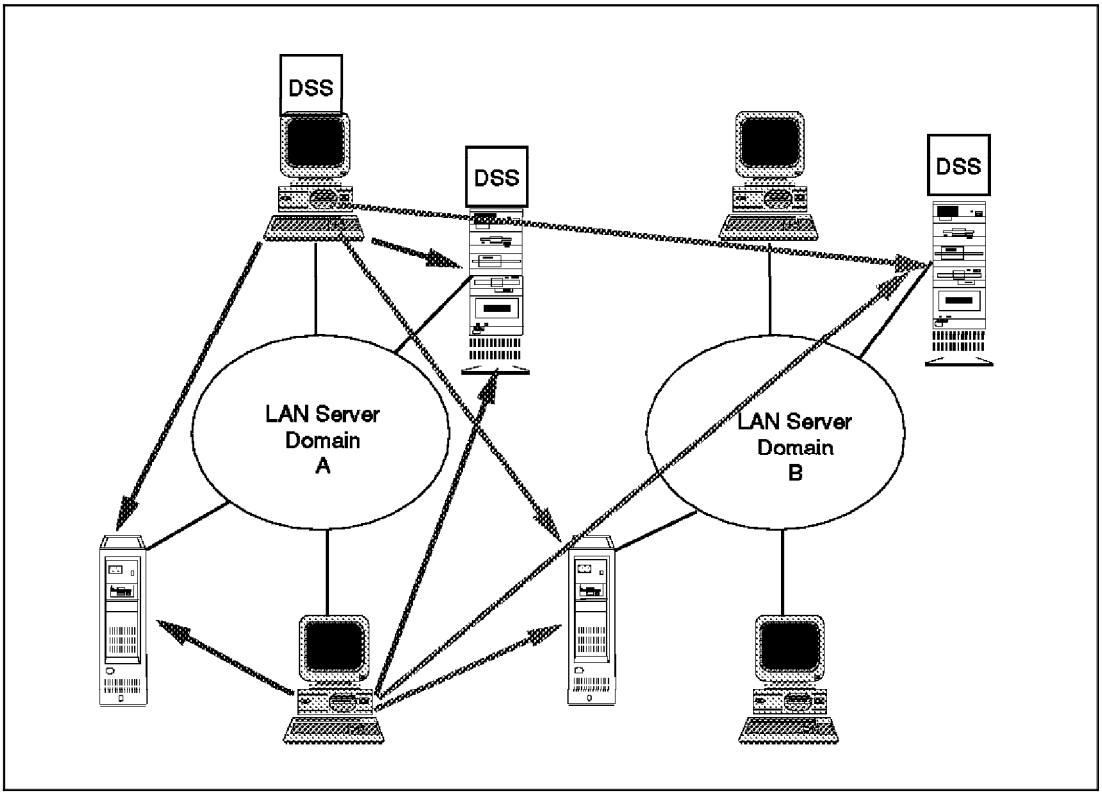

Figure 16. DSS on the Domain Controllers

Clients and additional servers remain unchanged although it is necessary to install at least one DSS client for administration. This configuration results in the following benefits:

- Users now have a single identity and a single password, which allows them to access any resources (files, printers, and so on) that they are authorized to access in any domain in the cell.
- Administrators can now remotely administer users on a cell basis rather than administer separate IDs in each domain that a user must access.
- The domain controllers now use DCE access control lists (ACLs) rather than OS/2 LAN Server ACLs. This results in a finer level of granularity with respect to access control because DCE ACLs support more security classes than the explicit user ID and guest classes supported by OS/2 LAN Server. In addition, the domain controller can now make its OS/2 LAN Server resources available across cell boundaries to users of the DSS client. The unchanged additional servers still use OS/2 LAN Server access control lists.
- The domain controllers are no longer subject to OS/2 LAN Server limits such as 256 groups per server. The unchanged additional servers are still subject to the OS/2 LAN Server limits.

- Users of the DSS client can seamlessly access resources on DSS servers across cell boundaries, still using only a single ID and password defined in their home cell.
- Users of the DSS client can use DCE security's Kerberos security server for end-to-end third-party security on DSS servers.
- Directory and registry databases can be replicated, where needed, throughout the enterprise to provide multiple directory and registry services, automatically load-balanced by DCE, for fast directory and registry lookups across LANs and WANs. The directory and registry servers can be any OSF DCE Version 1.1-compliant directory and security server from any vendor.
- The cell directory can be partitioned to allow performance tuning and to gain high capacity with relatively small machines.

1.5.5.3 DSS on Domain Controllers and Additional Servers

If DSS is installed on additional servers as well as on the domain controllers, then DCE ACLs are used on the additional servers. The additional servers are no longer subject to the OS/2 LAN Server group limits. Also, the additional servers use DCE directory and security services directly; so there is no need for DSS to synchronize the DCE and OS/2 LAN Server registries which results in less network traffic. In addition, they can make their resources available across cell boundaries.

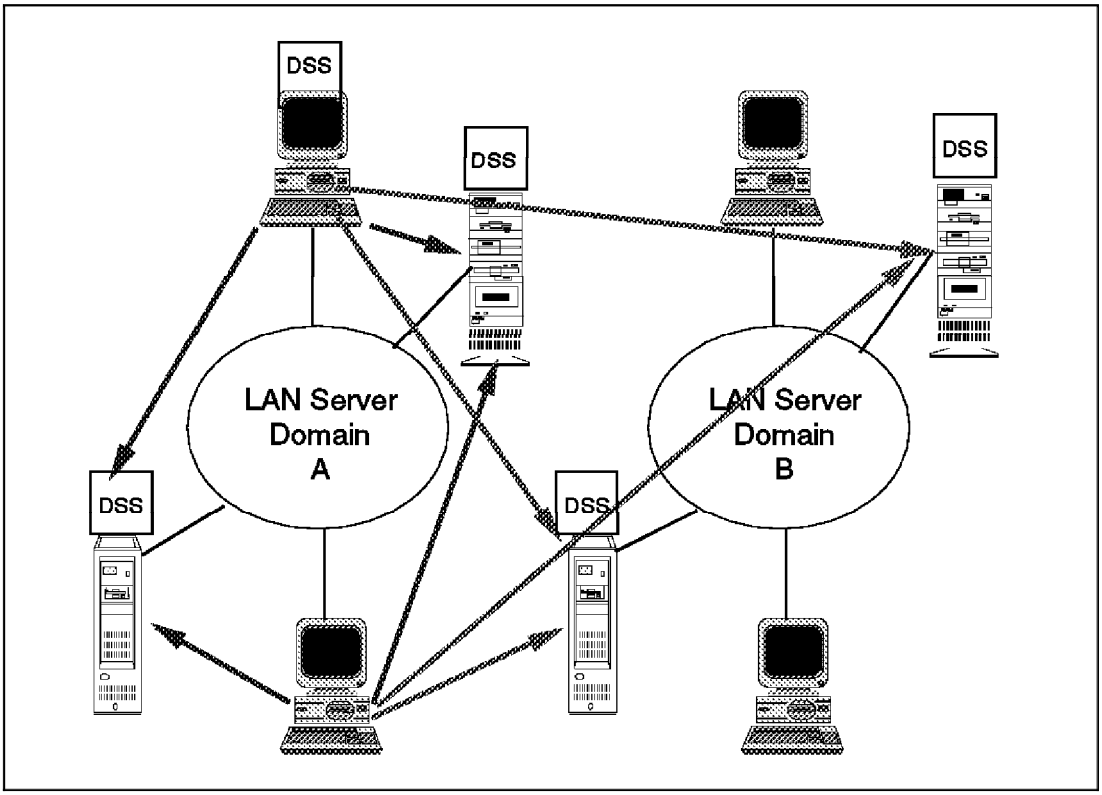

Figure 17. DSS on Domain Controllers and Additional Servers

The advantages in this scenario are:

- Updated servers
  - Use DCE Access Control Lists (ACLs)
  - The server resources are available to all users in the DSS cell and are not limited to the users explicitly defined in the resource domain
  - Ability to cross-cell access server resources
  - Remove the LAN Server limits. There are now more than 256 groups per server and more than 128 servers per domain
  - No need for directory synchronization
- Resource domain administration becomes much simpler because users do not need to be explicitly defined to a resource domain to access resources on its servers
- Administrators can administer multiple domains and all resources in the cell
- Clients have multiple domain resource access
- Stronger security, third-party authentication with DSS clients

- Reliability of the cell improves because user and group information need not be synchronized to DSS servers

1.5.6 More About the Directory and Security Server

In any discussion of the Directory and Security Server, it is helpful to understand how the DSS servers are related.

The directory servers use the security server's authentication service to prove to each other that they are indeed valid servers rather than imposters attempting to mine information from the directory database. They use the security service's Access Control List (ACL) facility to restrict directory server administration to only those user IDs that have been registered as administrators. Entries in the directory server database are time-stamped using the facilities provided by the time server.

The DSS domain controller uses the directory server to locate objects, both in the home cell and in other cells. It also stores the DSS resource domain information in the directory server's database. The DSS domain controller also uses the directory server to obtain resource information for propagation to unchanged additional servers via the OS/2 LAN Server Domain Controller Database.

The time server registers itself as an object in the directory database and uses the directory server to find other time servers in the network. It uses the security service to ensure that it is communicating with a valid time server rather than an imposter.

The DSS domain controller uses the time server to set the system time on all domain controllers and additional servers in the network.

The security server registers itself with the directory server as an object in the directory database. The security service uses the time synchronized by the time service to ensure that passwords and the service tickets that allow clients to use Directory and Security Server services are properly time-stamped and voided when they expire.

The DSS domain controller uses the security server to authenticate unchanged OS/2 LAN Server clients and servers. In addition, the DSS domain controller uses the security server to obtain security information for propagation to unchanged additional servers via the NET.ACC file. The DSS client uses the security server for authentication and to obtain permission to use Directory and Security Server services.

1.5.7 Directory

The directory server, sometimes called the *cell directory service (CDS)* is used by the DSS services to locate objects in the network. It stores information about objects such as domain controllers, additional servers, printers, and file system directories, in a common database that represents a common cell *name space*. A name space is simply the total collection of names shared by the DSS directory servers and other DCE directory servers. This frees users from knowing the real location of these objects. Clients can transparently access these objects (also known as *resources*) in their local resource domain or in any other resource domain in the cell without knowing or caring about the real physical location of the objects.

The database in which the name space is stored is known as a *clearinghouse*. Clearinghouses contain *replicas*, which are just physical copies of the database or of portions of the database. One of the powerful features of the Directory and Security Server is that the directory database can be replicated on several directory servers throughout the network. This allows remote sites (for example, branch offices) to maintain a local copy of the directory to improve performance and to ensure that the directory is available even when the communication line to the master directory server is down. Another powerful directory feature is the ability to partition the directory database so that servers need only manage that portion of the database that is most relevant to the local clients.

1.5.8 Security

The basic role of the security server is to allow DSS clients and servers to prove their identity. The security server includes several services:

The Registry Service: This service manages the cell registry database, which holds all of the user definitions for the entire cell. One of the benefits of DSS is that users are defined once in a home cell. This means that users of unchanged clients can access resources in any resource domain in the cell using a single user ID and password. Users of the DSS client can use this ability across cell boundaries.

The Authentication Service: This service allows a DSS user to positively identify themselves to the network.

The Privilege Service: This service manages the privilege attributes associated with users and groups. It is these privileges that determine which resources a user can access.

The Access Control List (ACL) facility: This facility, in combination with the Privilege Service, ensures that resource owners can grant privileges to only those users and groups that the resource owner feels have a need for

access to the resource. One of the major differences between unchanged additional servers in a DSS cell and additional servers that have been upgraded with DSS is that those that have been upgraded can use DCE ACLs, while the unchanged additional servers still use the less granular OS/2 LAN Server ACLs. This allows the upgraded servers to make their resources available to DSS clients across cell boundaries and to allow greater control over who is allowed to access the resources.

The Login facility: This facility authenticates users to the network. In DSS, users log in to DCE and OS/2 LAN Server with a single, integrated login.

Like the directory server, the security server supports the use of replicas for performance enhancement and fault tolerance.

1.6 Comparative Conclusions

This analysis is the correlation and interpretation of information collected from a number of sources, including consultant reports, the trade press, and sources on the Internet.

First, each networking product is described in brief, followed by a comparison of the advantages and disadvantages of that specific product. For a detailed discussion of each product, see the preceding sections. After that, a table compares the functions and features of each product.

1.6.1 IBM Directory and Security Server

DSS provides a unified directory for a single log-in access to networks on multiple servers, a single global view of all resources in all cells. The global directory of users and resources is divided into cells. This increases the number of workgroups and users within a workgroup that can be defined under a directory naming scheme, and it it easy to administer.

DSS links OS/2 Warp Server systems with other DCE implementations from HP, Sun, and DEC, as well as IBM mainframes, AS400, and RISC 6000 servers. It provides a build-in Kerberos security for authentication, encryption, and authorization. DSS offers administrative tools via the OS/2 graphical, object-oriented, point-and-click, drag-and-drop interface.

Table 5. Advantages and Disadvantages of IBM DSS	
Advantages	**Disadvantages**
• It is scalable from the desktop to the mainframe. • It provides good performance with a very large number of objects. • It provides interoperability with other DCE platforms. • It uses open DCE APIs. • It offers multi-cell support, providing customers with the ability to develop cross-company directories, which allows access to specific resources. • It is a snap-on to LAN Server, easy to install, consistent format with LAN Server. • By upgrading only the server, legacy clients gain single login and single password. • It uses Kerberos authentication.	• It has minimal Application Integration with IBM Servers/Products and third party support. • There is no Windows Client; it can not administer Windows from DSS. • It has a larger footprint (RAM/DASD) due to the robustness of DCE.

1.6.2 Novell NetWare 4.1

Novell Directory Services provides a fully distributed directory, but logically appears as a single database. It has a hierarchical tree structure to store objects and their associated attributes. The NDS database can be partitioned; this improves the networks reliability and accessibility. It provides a graphical utility, which enables the use of drag-and-drop operations to change the way partitions and replicas are set up. Because it is object-oriented, it can create directory-map objects that point to physical volumes. New installations have to be planned very carefully, and a migration from 3.x to 4.1 is not trivial.

Table 6. Advantages and Disadvantages of Novell NetWare 4.1

Advantages	Disadvantages
• Two years lead in availability, providing NDS enhancements to overcome initial shortcomings • Application Integration with Novell products such as GroupWise, ManageWise, NEST, Tuxedo • Slowly adding third party support for NDS-enabled applications in areas such as communications, databases, printers/fax machines, virus protection, software distribution • Robust attribute search and schema capabilities	• Proprietary APIs • Currently limited to NetWare and UNIXWare platforms • Performance problems with replicas of 25,000 objects (which seems to be solved with NetWare 4.11) • Migration from NetWare 3 to NetWare 4 is significant

1.6.3 Microsoft NT 4.0

Microsoft currently supports the Domain Naming Service instead of a directory. It uses physical names; so if the name changes, all references to it must be changed (for example, in the logon-script). Microsoft provides only a simple mapping of network names to addresses; they do not posses true directory capabilities such as location-independent extensive querying and searching capabilities, which means users must log in separately to each server and resource. To provide single login and global access of resources, trust relationships must be established manually between every domain on the network, a cumbersome task, especially in an enterprise network.

Trusted domains do not provide a single view of the enterprise network. Domains only apply to users and groups. It does not extend to other objects on the network such as file servers, application servers, or printers. There is no way to centrally manage other resources.

Table 7. Advantages and Disadvantages of Microsoft NT 3.51	
Advantages	**Disadvantages**
• Easy to install • Availability of applications • Developer support • Market share	• Must establish trusts between all the domains • Trusted domains apply only to users and groups, not to printers, file servers, application servers, and so on • Can not centrally manage all network resources • Windows-centric • Proprietary APIs • Does not provide location-independent extensive querying and searching capabilities. • No Cross-platform support • Weak Scalability

1.6.4 Function and Feature Comparison Summary

The following table (Table 8) compares the functions and features of the networking products OS/2 Warp Server, IBM DSS, NDS, and Microsoft Windows NT Server. Mainly networking functions are compared, but there are also some functions that are not directly related to networking or directory services, but are sometimes useful for reference. Nevertheless, the table does not claim to be complete and things may change with technical progress.

Table 8 (Page 1 of 3). Function and Feature Comparison				
Function or Feature	**OS/2 Warp Server**	**IBM Directory and Security Server**	**Novell NetWare 4.1**	**Microsoft Windows NT 4.0**
Partitioned directory database	No	Directory: Yes, Security: No	Yes	No
Replication	Yes	Yes	Yes	Yes
Synchronization	Yes	Yes	Yes	Yes

Table 8 (Page 2 of 3). Function and Feature Comparison

Function or Feature	OS/2 Warp Server	IBM Directory and Security Server	Novell NetWare 4.1	Microsoft Windows NT 4.0
Federation support, multi cell	No	Yes	No	No
Thin (Slim) Client	No	Yes	Yes	No
Attribute Search	No	No	Yes	No
Schema Support	No	Yes	Yes	No
Yellow Pages/ Catalog Services	No	No	No	No
ACLs	Yes	Yes	Yes	Yes
API Set	Propriety	Open	Propriety	Propriety
Central administration	Limited	Yes	Yes	Limited
Global (Alias) Names	Yes	Yes	Yes	No
Hierarchical	No	Yes	Yes	No
X500-based	No	Yes	Yes	No
Single login to network or services	Yes, via NSC	Yes	Yes	Yes, via Trusted Domains
Location transparency	Yes (see note 10)	Yes	Yes	No
TCP/IP support	Yes	Yes	Limited	Yes
Dynamic DNS Support	Yes	Yes	Add-on, NetWare NFS Services	No
DHCP support	Yes	Yes	Yes	Yes
NBNS support	Standard	Standard	Not applicable	Propriety
Platform support	OS/2	DCE platforms	NetWare, SCO UnixWare	Windows NT
Scalability and Performance	Add Warp servers	Add Directory/ Security servers	Add NetWare servers	Add NT servers

Table 8 (Page 3 of 3). Function and Feature Comparison

Function or Feature	OS/2 Warp Server	IBM Directory and Security Server	Novell NetWare 4.1	Microsoft Windows NT 4.0
Administration from Worldwide Web	Limited	No	No	Limited
Unicode	No	No	Yes	Yes
Moving a user account	Drag and Drop	Drag and Drop	Drag and Drop	Point and Click
Mobil user support	Yes	Yes	Yes	Limited
C2 security	No	No	No	Yes
Multi Purpose Server	Yes	Yes	No	Yes
Symmetric Multiprocessing	Yes	Yes	Yes	Yes
Server Console Management	Yes	Yes	No	Yes
FTP support	Yes	Yes	No	Yes
Installation planning	Easy	Easy	Difficult	Easy

Note:

1. C2-Security applies only if the entire network consists of Windows NT workstations only.
2. OS/2 Warp Server SMP has been available since September 1996.
3. MPTS ServicePak WR08210 has to be installed for h-node support or Warp 4's MPTS.
4. This means standard implementation of RFC 1001/1002.
5. See 3.14, “Name Resolution for Microsoft Windows Networking” on page 91 and 3.15, “TCPBEUI Interoperability with Microsoft” on page 99 for more information.
6. Applies only when Network SignON Coordinator is installed and configured.
7. If installed as an OS/2 application, NetWare servers gain multi purpose functionality (Only applies to NetWare 4.02).
8. Domain Controller Database information is replicated from the PDC (Primary Domain Controller) to BDCs (Backup Domain Controller). User and group account information is replicated from the PDC to all BDC and additional servers.
9. DHCP is provided with NetWare/IP 2.2.
10. When using aliases you get location transparency.

Chapter 2. Principles of Administration

This chapter briefly describes the principles of administration and the tasks involved from a general point of view. It highlights tasks that should be done by or coordinated by administrators. In the chapters of the practical part of this book we show how to do basic administration of all three network operating systems. Each graphical user interface is also discussed.

Before administration can take place, it is important to know what kind of tasks must be done and how to plan those tasks. It is an advantage to have a basic idea of planned extensions of a growing network as well as an idea of how to realize solutions for network users.

Making a decision for a particular network operating system may end up being fatal, when you have not given a though of scalability. A solution that suits well for 500 users does not necessarily mean that it suits for a 1000 user network as well.

2.1 Administrator Responsibilities

Administrators can face a series of responsibilities in managing a network. Those responsibilities can be structured into nine categories that can be reflected by using the model of a pyramid as shown in Figure 18.

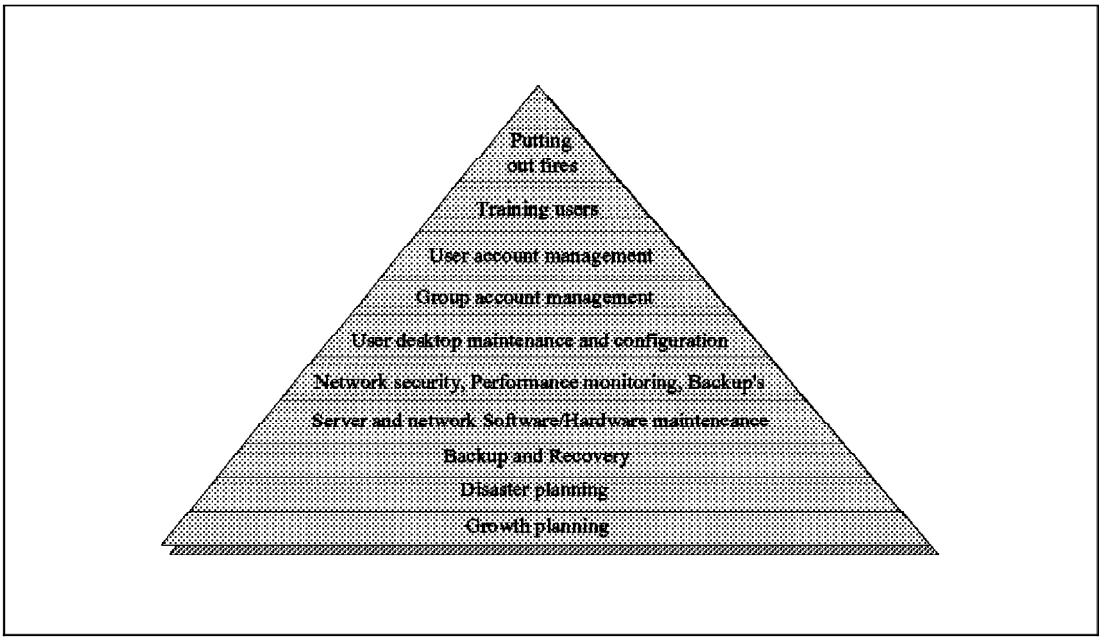

Figure 18. Pyramid of Responsibilities for Network Administrators

The bottom reflects responsibilities consuming most administration time; whereas the top reflects the least time being consumed.

However, also to be considered is the fact that these tasks usually are continuous activities. The tasks reflected on top of the pyramid are those that do not occur very often, but they can consume an entire day.

Chapter 3. Dynamic TCP/IP

TCP/IP management is the cross that all TCP/IP network administrators must bear. TCP/IP has opened your network to the Internet and intranets, but until recently, the consequence has been IP address assignment and management headaches. The truth is network administrators are human. We all make mistakes. And no wonder, most of us assign IP addresses by hand. Without a network manager's constant diligence, many pieces of IP management can fall through the cracks. For example, you can duplicate IP addresses or incorrectly enter default routers and subnet masks. Just the thought of tracking down yet another IP address conflict or planning another move causes most of us in the pen-and-paper world to cringe. DHCP, the Dynamic Host Configuration Protocol, can dramatically assist network administrators with many of their woes.

This section describes the purpose of Dynamic IP and the benefits that can be derived from it. We will also introduce the Dynamic IP components and give an overview of the design concepts to help you understanding what advantages NetBIOS name resolution has in such an environment.

Note on the Abbreviation's Meaning of DNS

Be not confused when you see "Domain Name Server" as well as "Domain Name System" abbreviated as "DNS". DNS used to be called Domain Name Server but the name has changed to "Domain Name System". Anyway, both names mean the same. They are just used differently by companies.

3.1 TCP/IP Configuration Parameters

To add a new workstation to an IP network, several parameters and a variety of information is required to configure the TCP/IP software. Network components, such as a domain name server, are also required. A new TCP/IP host would normally require the following information:

1. IP address
2. IP subnet mask
3. Default router address
4. Local host name
5. Domain name
6. Domain Name Server address
7. NetBIOS Name Server address

Additional parameters, such as other server addresses, time zones, or protocol-specific configurations, may be necessary in some cases.

Keeping track of that information in a large TCP/IP network may not always be an easy task for network administrators, especially if users or machines, or both, change their location frequently. IP address lists and domain name server databases have to be updated manually in order to keep track of any changes in the network.

From a user's point of view, a system administrator would have to be called to provide the necessary information in order to install a TCP/IP system. If the user moves to another location, this information must not be taken; the user will have to be assigned at least a new IP address if not a new hostname as well. Smart users may, therefore, cause potential disorder in a TCP/IP network.

Even if workstations will be automatically installed by software distribution techniques, the TCP/IP configuration parameters have to be preassigned per distribution client.

The Bootstrap Protocol (BootP), as described in RFCs 951 and 1497, was introduced to the TCP/IP community in 1985 to provide automatic assignment of some TCP/IP configuration parameters to a new TCP/IP host. A table has to be maintained at BootP servers to enter information specific to any client that has been planned for installation. Typically, clients are identified by their LAN adapter's hardware address, which has to be known to the system administrator in charge of a BootP server before he/she can prepare a new client entry in the database. Even though some manufacturers nowadays put the adapter hardware address on a label on the backplate of their LAN adapters, this ends up being a tedious process if many hosts have to be installed in a short period of time.

3.2 Objectives and Customer Benefits of Dynamic TCP/IP

To overcome the problems of having to manually update any centrally maintained information files and of having a user manually configure a TCP/IP workstation, the Dynamic Host Configuration Protocol (DHCP) has been designed and is described in an IETF DHC working group Internet draft and in RFCs 1533, 1534, 1541, and 1542. A DHCP server need not be preconfigured with a workstation's LAN address in order to submit the necessary TCP/IP configuration to it.

With DHCP in place, the assignment of IP addresses has become a lot easier. One problem still persists — how would a domain name server learn about those dynamically assigned IP addresses and hostnames so it

can update its database accordingly? This can be solved by the Dynamic Domain Name Services (DDNS), as proposed by an IETF DNSIND working group Internet draft.

Having DHCP and DDNS available gives system administrators the advantage of a high degree of flexibility and automation, and users do not have to worry about TCP/IP configuration parameters anymore. Persons in charge of information technology investment budgets may also prefer to spend their money on open standards, which will give them the assurance that products from different vendors will coexist in their TCP/IP networks.

IBM is actively participating in the designs and implementations of DHCP and DDNS, and it has coined the term *Dynamic IP*. To summarize, the objectives of Dynamic IP and its benefits to TCP/IP system administrators and users are as follows:

- Provides automatic IP network access and host configuration
- Simplifies IP network administration
- Leverages existing IP network products and infrastructure
- Employs only open standards
- Allows customers to administer site-specific host environments
- Enables customized, location-sensitive parameter setups

The following sections discusses the DHCP and DDNS protocols in more detail and give examples of their implementations.

Dynamic IP Components: Table 9 gives a brief description of the four types of network components that comprise Dynamic IP.

Table 9 (Page 1 of 2). DHCP Server Configuration

System Component	Description
Dynamic IP Hosts	Dynamic IP hosts contain DHCP client software and Dynamic DNS client software. Together, they discover and cooperate with their DHCP and Dynamic DNS server counterparts in the network to automatically configure the hosts for network participation.
DHCP Servers	DHCP servers provide the addresses and configuration information to DHCP and BootP clients on the network. DHCP servers contain information about the network configuration and about host operational parameters, as specified by the network administrator.

Table 9 (Page 2 of 2). DHCP Server Configuration	
System Component	**Description**
DDNS Servers	Dynamic DNS servers are a superset of today's static DNS BIND servers. The dynamic enhancements enable client hosts to dynamically register their name and address mappings in the DNS tables directly, rather than having an administrator manually perform the updates.
BootP Relay Agents (or BootP Helpers)	BootP relay agents may be used in IP router products to pass information between DHCP clients and servers. BootP relays eliminate the need for having a DHCP server on each subnet to service broadcast requests from DHCP clients.

3.3 Dynamic Host Configuration Protocol (DHCP)

The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts (clients as well as servers) on a TCP/IP network. DHCP is based on an enhancement to the Bootstrap Protocol (BootP), adding the capability of automatic allocation of reusable network addresses and additional configuration options. DHCP captures the behavior of BootP relay agents, and DHCP participants can interoperate with BootP participants.

In contrast to BootP, DHCP offers the possibility to assign an IP address to a client for a limited amount of time, and it also offers a way to supply all required configuration parameters for a client. This is not possible with BootP.

For a more detailed explanation, please see the latest version of the IETF DHC Internet draft, which is available online on the World Wide Web at the following URL:

`http://www.ietf.cnri.reston.va.us/ids.by.wg/dhc.html`

It should also be noted that other NOS providers also have DHCP solutions. Therefore, a table is included here for your reference that shows the functionality of the DHCP implementations of OS/2 Warp Server, Windows NT, and Novell NetWare.

Table 10 (Page 1 of 2). DHCP Functionality

Function	Subfunction	OS/2 WARP Server	Windows NT 4.0	Novell NetWare /IP 2.2
Address Pool Management	Support multiple pools of address	Yes	Yes	Yes
	Support multiple pools on subnet	Yes	No	No
	Can add block of addresses to pool	Yes	Yes	Yes
	Can add single addresses to pool	Yes	Yes	Yes
	Can modify pool of addresses	Yes	Yes	Yes
	Can modify single node in pool	Yes	No	No
	Can modify block of addresses in pool	Yes	Yes	Yes
	Can remove single addresses from pool	Yes	No	Yes
Reserved Address Management	Support for permanent, non-expirable leases	Yes	Yes	No
	Support for reservations by MAC address	Yes	Yes	Yes
	Support for reservations by host name	Yes	No	No
	Support for reservations by Client ID	Yes	No	No
	Support for reservations by Class ID	Yes	No	No
	Support for static BootP assignments	Yes	No	Yes
	Support for dynamic BootP assignments	Yes	No	Yes
	Can deny addresses to specific nodes	Yes	No	Yes

Table 10 (Page 2 of 2). DHCP Functionality

Function	Subfunction	OS/2 WARP Server	Windows NT 4.0	Novell NetWare /IP 2.2
Extended Address Management	Allows modifications to Global DHCP	Yes	Yes	No
	Allows modifications to pool-specific	Yes	Yes	Yes
	Allows adding new DHCP options	Yes	Yes	No
	Can manage multiple servers for single point	No	Yes	No
	Integrated with DNS	Yes	No	No

3.4 Dynamic Domain Name Services (DDNS)

DNS uses a distributed name space to maintain and supply name-to-IP address resolution. DNS is used throughout the Internet and allows a name to be used instead of an IP address when establishing a data conversation. DNS queries are performed with the DNS protocol.

Each site (university department, campus, company, or department within a company, for example) maintains its own database of information and runs a server program that other systems across the Internet (clients) can query.

Today's Domain Name System (DNS) servers support only queries on a statically configured database. The Dynamic DNS (DDNS) protocol defines extensions to the Domain Name System to enable DNS servers to accept requests to update the DNS database dynamically. These extensions provide support for adding and deleting a set of names and associated resource records within a single zone automatically.

The extensions assume that DNS security extensions, as defined by the IETF DNSSEC working group, have been implemented, but are not necessarily in use. DNS security extensions are used in DDNS to authenticate hosts that request to enter or change entries in the DDNS server database.

Without client authentication, another host, with perhaps malicious intent, may impersonate an unsuspecting host by remapping the address entry for the unsuspecting host to that of its own. After the remapping occurs, data (for example, logon passwords!) intended for the unsuspecting host is effectively intercepted by the malicious, spoofing host. IBM implements fail-safe RSA public-key digital signature technology to secure the DNS

database updates and eliminate the possibility of spoofing. IBM is the first company to introduce products that support Dynamic DNS and associated DNS security extensions.

For a more detailed explanation, please see the latest version of the IETF DNSIND and DNSSEC Internet drafts, which are available online on the Worldwide Web at the following URLs:

`http://www.ietf.cnri.reston.va.us/ids.by.wg/dnsind.html`

and

`http://www.ietf.cnri.reston.va.us/ids.by.wg/dnssec.html`

3.4.1 Generic Domains

Table 11 lists the normal classification of the seven generic domains.

Table 11. Seven Generic Domains

Domain	Description
com	commercial organizations
edu	educational institutions
gov	other U.S. governmental organizations
int	international organizations
mil	U.S. military
net	networks
org	other organizations

3.5 NetBIOS over TCP/IP for File & Print

Due to the fact that both OS/2 Warp Server and Microsoft Windows NT server are Server Message Block (SMB) servers, and therefore cannot use TCP/IP as their native protocol for file and print services, NetBIOS needs to be imbedded, packed, or converted (translated) into TCP/IP protocol. However, the clients and servers are communicating via NetBIOS names rather than an IP addressing scheme.

This chapter discusses NetBIOS Name resolution issues for the OS/2 Warp Server environment with OS/2 LAN Server, OS/2 LAN Requester and DOS LAN Services (DLS) client as well as the solutions for Microsoft Windows NT server.

Even though Novell is not using the NetBIOS protocol, a brief overview of Novell's participation in an TCP/IP network is given at the end of this chapter.

3.6 Overview of NetBIOS Name Resolution over TCP/IP Network

Clients and servers need to know how to find one another in order to share information. The NetBIOS conventions built into DOS, Windows and OS/2 clients/servers use 16 byte NetBIOS names which refer to one another by name. Different applications on the same PC uses different names to represent their applications.

NetBIOS names, like *Steve's_PC* or *Printer_HP1* can be built into programs or solicited from humans with relative ease. NetBIOS names can be used as unambiguous identifiers even if a station is moved to another location. However, to send one another packets of information, the TCP/IP protocol drivers of the respective PCs must refer to one another by IP address. The problem exists, then, of having to translate NetBIOS Names into IP addresses in order to effect PC-to-PC communication on an IP network.

To date, this translation has been handled in one of two ways: by use of static tables residing on each client and server, or by use of (dynamic) broadcast queries (packets sent to every client and server) asking in effect *Where is Steve's_PC?*

The problem with static tables is that they must be continually updated and maintained, an activity far more troublesome than the maintenance of IP addresses alone. Every time any new station is added to the network, all of its applications' names must be added to the static table of each other station that wants to send it data. And with static entries, though the name is always mappable, there is no telling whether the named application is actually active at the time interaction is desired by another station. The problem with broadcast queries is that IP networks cannot propagate broadcasts beyond a single (logical) cable segment. Resources located on the other side of a router from the broadcasting station will not receive the query. Every station on the same side of the router will be pestered with queries for which it doesn't know the answer.

A *NetBIOS over TCP/IP* protocol has been defined by the governing TCP/IP standards body, the Internet Engineering Task Force (IETF), which overcomes each of these problems. The IETF standard describes how NetBIOS stations may interact with a NetBIOS Name Server in order to dynamically register their own application names and to learn the name-to-address mappings of other applications.

There are several defined classes or modes of NetBIOS over TCP/IP implementations specified by RFCs 1001 and 1002.

The NetBIOS over TCP/IP modes include the following:

- b-node, which uses broadcasts to resolve names
- p-node, which uses point-to-point communications with a name server to resolve names
- m-node, which uses b-node first (broadcasts), then p-node (name queries) if the broadcast fails to resolve a name
- h-node, which uses p-node first for name queries, then b-node if the name service is unavailable .

For DHCP users on a network, the node type is assigned by the DHCP server. A client computer gets NetBIOS-names-to-IP address resolution by establishing a point-to-point communication (p-node) with a NetBIOS Name Server, assuming NetBIOS Name Server (NBNS) is in place on the network. If there is no NetBIOS Name Server in place, b-node broadcasts will occur to resolve names. This is explained in the following sections. To understand the different nodes, we describe it in a little bit more detail:

3.6.1 B-Node

The b-node mode uses broadcasts for name registration and resolution. That is, if PC1 wants to communicate with PC2, it will broadcast to all machines that it is looking for PC2 and then wait a specified time for PC2 to respond. B-node has two major problems:

- In a large environment, it loads the network with broadcasts.
- Routers do not forward broadcasts; so computers that are on opposite sides of a router will never hear the request.

3.6.2 P-Node

The p-node mode addresses the issue that b-node does not solve. In a p-node environment, computers neither create nor respond to broadcasts. All computers register themselves with the NetBIOS Name Server. The NetBIOS Name Server is responsible for knowing computer names and addresses and for ensuring no duplicate names exist on the network. All computers must be configured to know the address of the NetBIOS Name Server.

In this environment, when PC1 wants to communicate with PC2, it queries the NetBIOS Name Server for the address of PC2. When PC1 gets the appropriate address from the NetBIOS Name Server, it goes directly to PC2 without broadcasting. Because the name queries go directly to the NetBIOS Name Server, p-node avoids loading the network with broadcasts. Because broadcasts are not used and because the address is received directly, computers can span routers.

The most significant problems with p-node are the following:

- All computers must be configured to know the address of the NetBIOS Name Server (although this is typically configured via DHCP).
- If for any reason the NetBIOS Name Server is down, computers that rely on the NetBIOS Name Server to resolve addresses cannot get to any other system on the network, even if they are on the local network.

3.6.3 M-Node

The m-node mode was created primarily to solve the problems associated with b-node and p-node. This mode uses a combination of b-node and p-node. In an m-mode environment, a computer first attempts registration and resolution using b-node. If that is successful, it then switches to the p-node. Because this uses b-node first, it does not solve the problem of generating broadcast traffic on the network. However, m-node can cross routers. Also, because b-node is always tried first, computers on the same side of a router continue to operate as usual if the NetBIOS Name Server is down.

M-node uses broadcasts for performance optimization because in most environments local resources are used more frequently than remote resources.

3.6.4 H-Node

The h-node mode is also a combination of b-node and p-node that uses broadcasts as a last effort. Because p-node is used first, no broadcasts are generated if the NetBIOS Name Server is running, and computers can span routers. If the NetBIOS Name Server is down, b-node is used; so computers on the same side of a router continue to operate as usual.

The h-node mode does more than change the order for using b-node and p-node. If the NetBIOS Name Server is down so that local broadcasts (b-node) must be used, the computer continues to poll the NetBIOS Name Server. As soon as the NetBIOS Name Server can be reached again, the system switches back to p-node. Also, h-node can be configured to use local hostlists after broadcast name resolution fails.

The h-node mode solves the most significant problems associated with broadcasts and operations in a routed environment. The h-node has replaced the m-node.

3.7 NetBIOS over TCP/IP in OS/2 Warp Server

Several components of OS/2 Warp Server can use NetBIOS for communications, but they can also use other protocols like TCP/IP or IPX.

File and Print Sharing Services remains the only OS/2 Warp Server component that can only use NetBIOS as a programming interface.

The original NetBEUI protocol along with the NetBIOS API has some specific characteristics that limit its use in certain wide area network environments:

- NetBIOS uses a flat name space.
- NetBIOS relies on the broadcast technique to register/find a name.
- NetBIOS cannot be routed.

One solution to overcome these limitations can be found in RFCs 1001 and 1002. They describe the standard way to implement the NetBIOS services on top of the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). Multi Protocol and Transport Services provides a full TCP/IP protocol stack and a TCPBEUI API conversion layer, which is a ring 0 implementation of RFC 1001/1002.

Note: In order to use NetBIOS over TCP/IP, you do not need to install the TCP/IP Services of OS/2 Warp Server since the support for this combination of protocols is fully included in Multi Protocol and Transport Services.

TCP/IP Services of OS/2 Warp Server means TCP/IP applications on top of the TCP/IP protocol, such as File Transfer Protocol (FTP), Line Printer Requester (LPR), Dynamic Host Configuration Protocol (DHCP) and Dynamic Domain Name System (DDNS) Server etc.

Another solution of routing NetBIOS is to use the NetBIOS over the Internet Packet Packet Exchange (IPX) protocol that is also supplied with Multi Protocol and Transport Services.

The capability of running NetBIOS applications over routable protocols offers new flexibility when designing OS/2 Warp Server networks. OS/2 Warp Server systems, Warp Connect Peer workstations, LAN Servers, and LAN Requester workstations can be on remote LAN segments connected by IP routers. This also means that such systems can be introduced into existing TCP/IP networks without introducing an additional network protocol.

OS/2 TCPBEUI (TCP/IP NetBIOS extended user interface; IBM's implementation of NetBIOS over TCP/IP) is a high performance, ring 0, implementation of NetBIOS over TCP/IP. TCPBEUI provides the LM10 protocol driver interface. It is the same LM10 functionality that is also provided by NetBEUI (NetBIOS extended user interface) Figure 19 on page 63 shows this interface.

Mapping of NetBIOS API Calls

TCPBEUI maps NetBIOS API calls into the TCP/IP protocol.

NetBIOS over TCP/IP contains enhancements over the b-node standard that improve system performance by decreasing broadcast frames and by expanding communications over routers and bridges. These enhancements, described in 3.9, “Reducing Broadcast Frames with TCPBEUI” on page 67, are transparent to NetBIOS applications and do not interfere with other b-node implementations that lack similar functions.

RFC 1001/1002 is not an encapsulation technique; it builds special packets and sends them out via UDP and TCP. For example, once a NetBIOS session has been established, TCPBEUI will use sockets-send commands over a TCP connection to send NetBIOS session data. TCPBEUI builds a four-byte session header that precedes the actual user data. Thus, a NetBIOS Chain Send of 128 KB would have an overhead of only four bytes.

TCPBEUI allows peer-to-peer communication over the TCP/IP network with other computers that have compatible services. Figure 19 on page 63 shows the relationship between the NetBIOS, NetBIOS over TCP/IP, and TCP/IP protocol stacks as implemented in Multi Protocol and Transport Services.

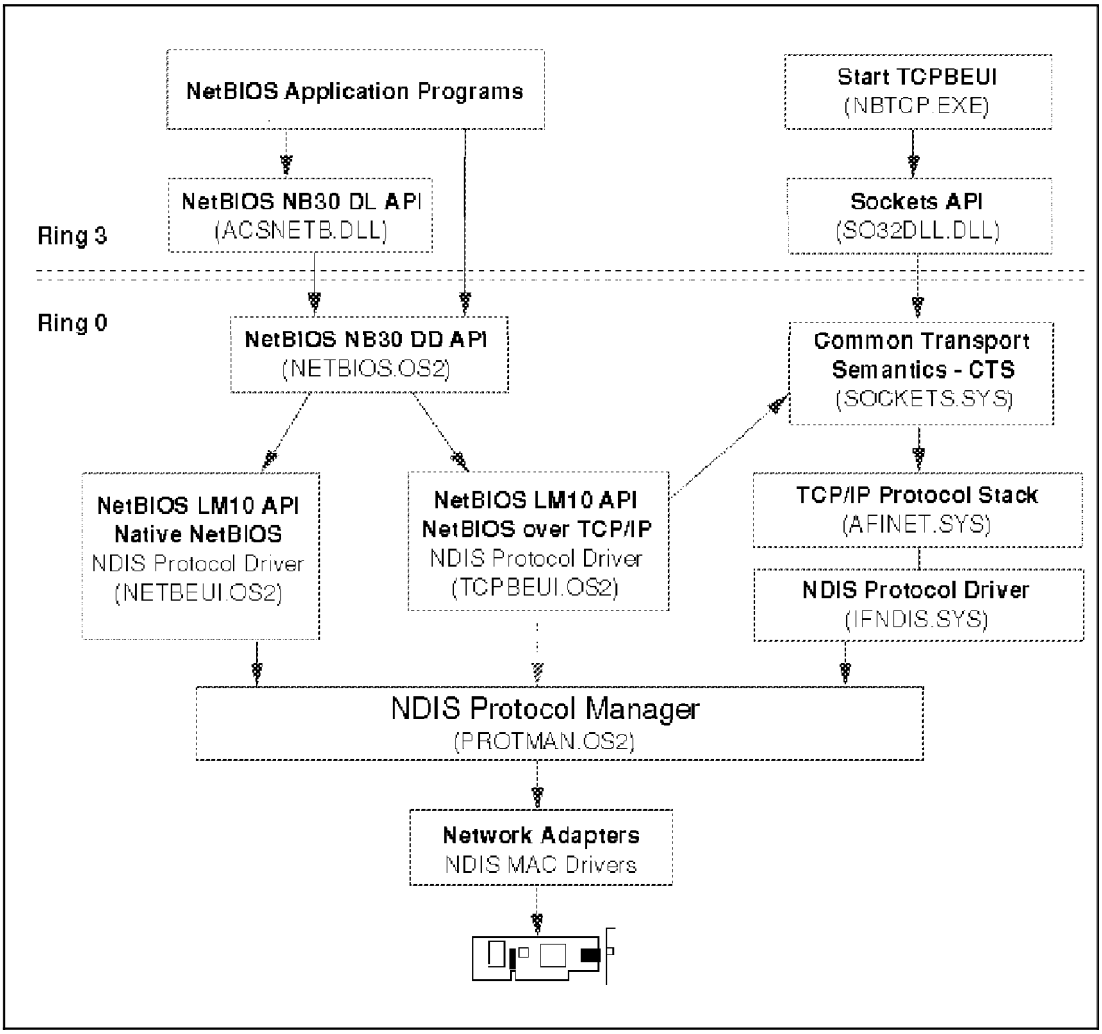

Figure 19. NetBIOS, NetBIOS over TCP/IP and TCP/IP Structure

Unlike NETBEUI.OS2, the TCPBEUI.OS2 program doesn't directly communicate with the NDIS interface. The dotted line in the figure indicates TCPBEUI has a BINDINGS statement in the PROTOCOL.INI file, but a bind process is only required in order to create a control block area.

Figure 19 also illustrates how NetBIOS applications can use both NETBEUI and TCPBEUI protocol stacks. ACSNETB.DLL provides the ring 3 NetBIOS DLL API for application programs. Ring 3 NetBIOS commands are sent to NETBIOS.OS2 for processing. NETBIOS.OS2 provides the ring 0 NetBIOS DLL API for applications and other device drivers to use, and it binds to one or more LM10 (LAN Manager 1.0) transport protocol drivers.

The LAN redirector component of File and Print Sharing Services (NETWKSTA.200) and HPFS386 use the LM10 interface.

Support for NetBIOS over TCP/IP can easily be added to the existing NetBIOS structure. For example, you can have one NETBEUI and four TCPBEUI (total of five) protocols bound to one physical adapter. It is provided by having NETBIOS.OS2 bind to TCPBEUI.OS2. Although the file and print install program of Warp Server is limited to defining 4 logical adapters, you can manually add NICs and protocols to your MPTS configuration after the initial install by using the `MPTS` command from an OS/2 command line. After configuring MPTS, do not forget to add `NETx` (where `x` is 1, 2, 3, 4, 5, 6, 7, 8, for example) statements to your IBMLAN IBMLAN.INI file. Figure 24 on page 67 shows an example with four `NETx` entries.

To enable NETWKSTA.200 to use TCPBEUI, there must be a `NETx` statement in the IBMLAN.INI file configured appropriately (see Figure 22 on page 66).

Data transfer to LAN is handled by a Medium Access Control (MAC) device driver, for example the IBMTOK.OS2 device driver.

3.8 TCPBEUI Coexistence with NetBEUI

Multi Protocol and Transport Services provide the capability of configuring NetBIOS applications, especially File and Print Sharing Services, with both NetBEUI and TCPBEUI on the same network interface card. This dual protocol stack configuration allows local sessions to continue running with NetBEUI performance while also providing wide area network connectivity with NetBIOS over TCP/IP.

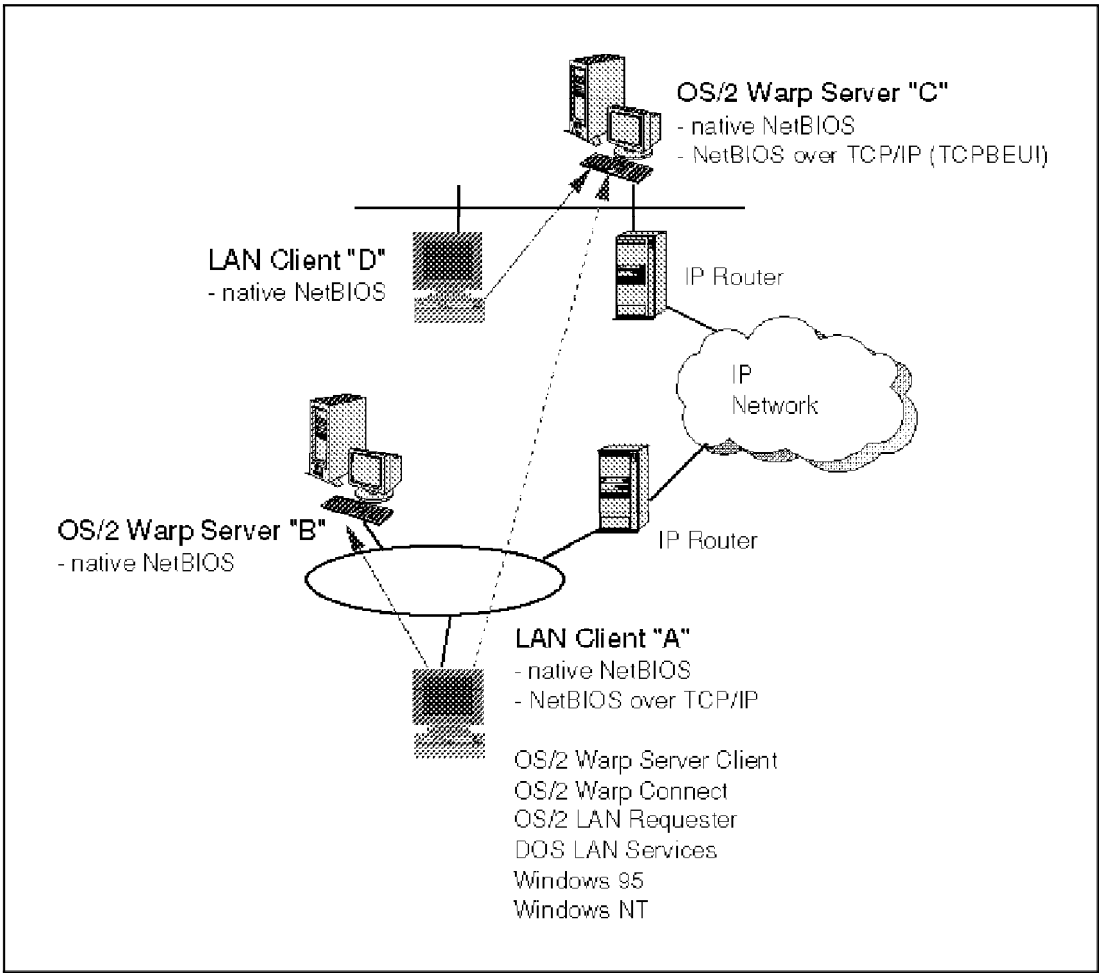

Figure 20. TCPBEUI Coexistence

Figure 20 shows an example scenario with both TCP/IP and NetBIOS protocols being used and TCP/IP Services installed on a server. In this example, LAN Client A is able to access File and Print Sharing Services resources on the OS/2 Warp Server B on the local LAN segment via NetBIOS. LAN Client A accesses the OS/2 Warp Server C on the remote LAN segment across the IP network via TCPBEUI. In addition, it is able to use the TCP/IP applications provided by TCP/IP Services to access local and remote TCP/IP hosts via the native TCP/IP protocol.

Multi Protocol and Transport Services provide the TCP/IP protocol capability with or without TCP/IP Services installed, but with only a limited set of TCP/IP functions and services. These functions basically enable you to configure IP interfaces and routes and to test the TCP/IP protocol for proper operation:

- ARP
- HOST

- HOSTNAME
- IFCONFIG
- IPTRACE
- IPFORMAT
- NETSTAT
- PING
- ROUTE

When configuring Multi Protocol and Transport Services for both NetBEUI and TCPBEUI, even though a single LAN adapter is present in the workstation, the two protocols need to be configured on different logical adapters. The Current Configuration window on the MPTS Configuration panel should be changed as follows:

```
IBM Compatible Token-Ring Network Adapter (IBMTOK.OS2) ...
  0  - IBM IEEE 802.2
  0  - IBM OS/2 NETBIOS
  1  - IBM OS/2 NETBIOS OVER TCP/IP
  0  - IBM TCP/IP
```

Figure 21. MPTS Configuration Panel. Single token-ring adapter bound to IEEE 802.2, NetBIOS, TCPBEUI, and TCP/IP.

Note: The logical numbers of the protocol drivers must be set differently although only one physical LAN adapter is present.

File and Print Sharing Services handles this configuration as if there were two adapters present. Therefore, two `NET` entries will be made in IBMLAN.INI file:

```
[networks]

  net1 = NETBEUI$,0,LM10,102,175,14
  net2 = TCPBEUI$,1,LM10,102,175,14

[requester]

  wrknets = net1,net2

[server]

  srvnets = net1,net2
```

Figure 22. IBMLAN.INI for Two NetBIOS Networks. NetBIOS and TCPBEUI bound to a single LAN adapter (Extract).

When configuring Multi Protocol and Transport Services for NetBEUI and TCPBEUI with two LAN adapters present in the workstation, the two protocols can be configured as shown in Figure 23 on page 67.

```
IBM Compatible Token-Ring Network Adapter (IBMTOK.OS2) ...
  0  - IBM IEEE 802.2
  0  - IBM OS/2 NETBIOS
  1  - IBM OS/2 NETBIOS OVER TCP/IP
  2  - IBM OS/2 NETBIOS OVER TCP/IP
  0  - IBM TCP/IP
IBM Compatible Token-Ring Network Adapter (IBMTOK.OS2) ...
  3  - IBM OS/2 NETBIOS
```

Figure 23. MPTS Configuration Panel. Two token-ring adapters; first one is bound to IEEE 802.2, NetBIOS, two times TCPBEUI, and TCP/IP; second one is bound to NetBIOS.

Note: The logical numbers of the protocol drivers must be set differently.

With the MPTS configuration shown in Figure 23 your server can support more than 500 NetBIOS users and more than 500 NetBIOS over TCP/IP users. Just make sure, that the IBMLAN IBMLAN.INI's [networks] section contains the right number of `NETx` statements, for example, as shown in Figure 24.

```
[networks]

  net1 = NETBEUI$,0,LM10,102,175,14
  net2 = TCPBEUI$,1,LM10,102,175,14
  net3 = TCPBEUI$,2,LM10,102,175,14
  net4 = NETBEUI$,3,LM10,102,175,14

[requester]

  wrknets = net1,net2,net3,net4

[server]

  srvnets = net1,net2,net3,net4
```

Figure 24. IBMLAN.INI. NetBIOS and TCPBEUI bound to two physical LAN adapters (Extract).

3.9 Reducing Broadcast Frames with TCPBEUI

NetBIOS over TCP/IP, or TCPBEUI, provides an extension to b-node. It is called Routing Extensions. It supports b-node name resolution with name

resolution extensions that help to reduce name resolution broadcast traffic. This section discusses these topics and also gives you information on how to use an existing Domain Name Server (DNS) in a TCPBEUI environment. With all these settings, you can reduce TCP/IP broadcast frames on the network.

3.9.1 Routing Extensions

Three of the enhancements to TCPBEUI are in the form of *routing extensions*. These extensions allow communication between networks and over IP routers and bridges. The following subsections describe these routing extensions:

3.9.1.1 Names File

A names file consists of pairs of NetBIOS names and an IP addresses. NetBIOS over TCP/IP conducts a prefix search of the names file before broadcasting on the network. The prefix match succeeds if the entry in the names file matches the given name, up to the length of the entry. The first match is used; therefore, the order in which NetBIOS names are listed in the names file is important.

To enable this routing extension, set the `NAMESFILE` parameter in the `TCPBEUI` section of PROTOCOL.INI to a nonzero integer that represents the number of names file entries.

3.9.1.2 Domain Name Server (DNS)

A network administrator can maintain pairs of NetBIOS names and IP addresses in a DNS. If a name query fails, NetBIOS over TCP/IP can append the NetBIOS Domain Scope String to the encoded NetBIOS name and issue a request to the DNS to look up an IP address for that NetBIOS name. The Domain Scope String is defined by the PROTOCOL.INI parameter `DOMAINSCOPE`.

For more information on how to set up the DNS with the NetBIOS names, see 3.9.4, “Storing NetBIOS Names on the Domain Name Server (DNS)” on page 70.

3.9.1.3 Broadcast File

A broadcast file contains a list of host names, host addresses or directed broadcast addresses. It is read at startup, and each valid address is added to the set of destination addresses for broadcast packets. Remote nodes included in the broadcast file are then treated as if they were on the local network. Use of a broadcast file has the effect of extending a node's broadcast domain to its own subnet and to any other subnets listed in the broadcast file. A maximum of 32 broadcast file entries are supported, each

of which could include additional subnets, thus extending the node's broadcast domain.

If your routers support directed broadcasts (that is, you can ping the broadcast address of a distant IP subnet, and get back a response from all the stations on that subnet), then you can place the broadcast address for each subnet in the server's broadcast file. Also enable the TCPBEUI name cache described in 3.9.3, “Name Cache and Name Discovery Algorithm” on page 70. This greatly reduces broadcast traffic and eases administration. (The clients still need to know the IP address and NetBIOS name of each server and peer server.)

3.9.2 Configuring TCPBEUI Routing Extensions

Use the MPTS configuration program (which is explained in detail in *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602) and add the IBM OS/2 NetBIOS over TCP/IP protocol to an adapter. Then double-click on **IBM OS/2 NetBIOS over TCP/IP** to invoke the following menu:

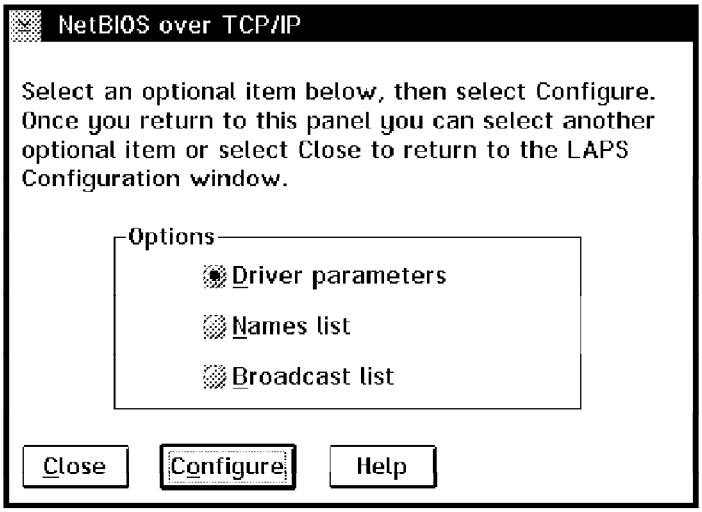

Figure 25. TCPBEUI Configuration

On this menu, select:

Driver parameters to configure the parameters for the TCPBEUI protocol

Names list to configure the names file (IBMCOM RFCNAMES.LST)

Broadcast list to configure the broadcast file (IBMCOM RFCBCST.LST)

When you make changes to the names or broadcast files while TCPBEUI is active, you can reinitialize TCPBEUI with the new files using the RFCADDR.EXE program.

3.9.3 Name Cache and Name Discovery Algorithm

Another enhancement NetBIOS over TCP/IP provides is a *name cache* for storing remote names that have been discovered. Since TCPBEUI uses broadcasting as a mechanism for name discovery, by checking the cache first, broadcast traffic can be reduced. This cache is enabled by setting the `NAMECACHE` parameter in the `TCPBEUI` section of the PROTOCOL.INI to a nonzero integer that represents the number of names stored in the directory.

The information in the remote name cache (or directory) is also stored on disk (in the IBMCOM RFCCACHE.LST file) and periodically updated. When the system is restarted, this information can be preloaded into the cache at bootup time. Preloading can reduce the amount of broadcast frames on the network since NetBIOS will not have to rediscover names for remote workstations. To preload the remote names cache, set the `PRELOADCACHE = YES` in the `TCPBEUI` section of the PROTOCOL.INI file.

When NetBIOS over TCP/IP is searching for a name, the following name discovery algorithm is used:

1. Check the local name cache first.
2. If not found, check the local names file.
3. Next, issue `GetHostByName()` to the domain name server. The tcpip etc hosts file is checked if the `GetHostByName` to the DNS fails.
4. Finally, issue a broadcast using the broadcast file's entries.

It is recommended that when running NetBIOS over TCP/IP in a wide area network (WAN), you should turn name caching on at the server (for instance, setting it to a value of 100).

3.9.4 Storing NetBIOS Names on the Domain Name Server (DNS)

In a larger network where a DNS already exists, you can use the DNS database to store NetBIOS names and IP addresses pairs, thereby eliminating the need for maintaining a broadcast file or names file on each client. In each client PROTOCOL.INI file, you must only ensure that the `DOMAINSCOPE` parameter is set to the TCP/IP domain name. TCPBEUI will then know to search that domain's DNS for the IP address of the requested server.

Notes:

1. The solution described in this section assumes that the server is already set up as a TCP/IP machine with a host name/IP address pair that is registered in the DNS database.
2. If you do not have a DNS, you can set up the local node's hosts file (tcpip etc hosts) in the same way we describe here. That is, the NetBIOS names must be encoded in the hosts file just as they must be in the DNS. TCPBEUI first looks for the requested server IP address in the DNS; if one does not exist or the address is not specified in the DNS, TCPBEUI checks for the local hosts file.

The servers' NetBIOS names must be added to the DNS database in an *encoded* format. The encoding is necessary because NetBIOS names are 16 bytes of *any* bit pattern, and a TCP/IP DNS only accepts host names in the character set *A to Z* and *0 to 9*.

For example, if you have specified

```
DOMAINSCOPE=austin.ibm.com
```

in the PROTOCOL.INI file and the NetBIOS name you have requested is not found in the local names cache or the local names file, then a sockets `GetHostByName(netbios_name.austin.ibm.com)` call will be made. TCPBEUI translates the 16-byte NetBIOS name into a 32-byte reversible, half-ASCII biased encoded format, such as:

```
GetHostByName(GCHCGJGDGFCACACACACACACACACACACA.austin.ibm.com)
```

and sends it to the DNS. If the DNS knows this name, it sends back the IP address to TCPBEUI. For this to work, the administrator must store the NetBIOS names in the DNS in the encoded format.

How do you encode NetBIOS names and store them in the DNS database? You must encode the 16-byte name into a 32-byte string by using the MAPNAME utility, which is located in the APPLETS directory of MPTS diskette 5 (MPTSAPLT.ZIP). This utility can also be found on the OS/2 Warp Server CD-ROM under the CID SERVER IBMLS IBM500N5 subdirectory. Then, you store the names in the DNS database so that they point back to the original host name, where the TCP/IP address is already listed. We will take you through an example of how to do this.

For each server, there will be at least three entries in the DNS database in addition to the initial host name entry. (Remember, we are assuming that the LAN Server is already set up as a TCP/IP host with a host name/IP address pair that is registered in the DNS database.) The three entries are necessary because LAN Server issues a NetBIOS `NCB.AddName` call three

times, using the Computername parameter specified in the IBMLAN.INI's [Requester] section and ending each with a unique hex value as the sixteenth byte. The hex values used as the sixteenth byte are 0x20 (blank or null), 0x00 and 0x03. If the server is a domain controller, there must be a fourth entry, the encoded domain name with the sixteenth byte of 0x00.

Let's say that we have a DNS already set up on our network. We have installed LAN Server 4.0 on the domain controller and one additional server. Both machines also have TCP/IP for OS/2 installed, and their host names are registered in the DNS database. We want to configure for TCPBEUI so that clients can access servers across our IP router without requiring a broadcast file or names file at each client. To do this, we take advantage of the DNS, and add the appropriate DOMAINSCOPE entry to each client's PROTOCOL.INI file.

In this example, our domain name is ITSCAUS, and our two servers are configured as follows:

Domain controller Computername: ITSCSV00
TCP/IP host name: ITSCWK00

Additional server Computername: ITSCSV01
TCP/IP host name: ITSCWK01

Note: The Computername refers to the IBMLAN.INI parameter. This is also referred to as the server name or machine ID.

Here's an extract from our DNS database *before* we add the encoded NetBIOS names:

```
ITSCWK00                        86400  IN  A      129.35.144.210
                                   IN  HINFO  DC HOST NAME
;
ITSCWK01                        86400  IN  A      129.35.144.211
                                   IN  HINFO  AS HOST NAME
;
```

Figure 26. Sample DNS Database File before Adding Encoded NetBIOS Names. The TCP/IP host names are listed with the workstation IP addresses.

The HINFO keyword specifies comment information. In this case, we have indicated that ITSCWK00 is the TCP/IP host name for the domain controller, and ITSCWK01 is the host name for the additional server. TCP/IP looks up the host name in the DNS database and finds the actual IP address.

Now we want to use TCPBEUI and take advantage of the DNS database. To do this, we must encode the server NetBIOS names using the MAPNAME

utility. Typing MAPNAME by itself will give you help on how to use the command. The utility converts NetBIOS names to RFC-encoded names and vice versa. Using our example, the following steps show you how to encode your server NetBIOS names.

MAPNAME Requires Uppercase NetBIOS Names

When using MAPNAME, be sure to type any NetBIOS names in *uppercase* letters because this is a case-sensitive utility. If you type names in lowercase, the output will be incorrect.

1. Use the MAPNAME utilty with the /RB parameters to specify that you want the output to be in RFC format and padded with blanks for up to 16 characters.

   ```
   MAPNAME ITSCSV00 /RB
   ```

   The following 32-byte encoded name is displayed:

   ```
   RFC name: EJFEFDEDFDFGDADACACACACACACACACA
   ```

   This is the first of the four encoded names you need for the domain controller. Here, the sixteenth byte, CA, is null (0x20). The following command would give us the same result, but since null characters are the default, the L20 is unnecessary.

   ```
   MAPNAME ITSCSV00 /RBL20
   ```

2. This time, also use the L parameter to specify that you want the last character of the output to be 0x00, as follows:

   ```
   MAPNAME ITSCSV00 /RBL00
   ```

   The result is:

   ```
   RFC name: EJFEFDEDFDFGDADACACACACACACAAA
   ```

   AA is hex 0x00.

3. Again, use the L parameter to specify the last character of the output to be 0x03, as follows:

   ```
   MAPNAME ITSCSV00 /RBL03
   ```

   You receive this output:

   ```
   RFC name: EJFEFDEDFDFGDADACACACACACACAAD
   ```

   AD is hex 0x03.

4. Because this is the domain controller, you must also specify the encoded domain name with the sixteenth byte of 0x00, as follows:

   ```
   MAPNAME ITSCAUS /RBL00
   ```

The encoded name is:

```
RFC name: EJFEFDEDFDFGDADACACACACACACAAA
```

5. Now we go through the first three steps for the additional server, ITSCSV01, to get the following output (Do not encode the domain name for additional servers):

```
MAPNAME ITSCSV01 /RB
RFC name: EJFEFDEDFDFGDADBCACACACACACACACA

MAPNAME ITSCSV01 /RBL00
RFC name: EJFEFDEDFDFGDADBCACACACACACACAAA

MAPNAME ITSCSV01 /RBL03
RFC name: EJFEFDEDFDFGDADBCACACACACACACAAD
```

6. Edit the DNS database to add the entries for the domain controller and additional server. Use the DNS `CNAME` keyword to point back to the host name entry for the machine where the actual IP address is already specified. In other words, the encoded names we have generated are *aliases* for the host names ITSCWK00 and ITSCWK01.

   Note: You cannot have two entries pointing to the same IP address; so you must use the `CNAME` keyword to create aliases.

   The following example shows how our DNS database file looks *after* adding the NetBIOS encoded names. Again, we use `HINFO` to designate comments.

```
ITSCWK00                        86400  IN  A     129.35.144.210
                                       IN  HINFO  DC HOST NAME
;
EJFEFDEDFDFGDADACACACACACACACACA 86400  IN  CNAME  ITSCWK00
                                       IN  HINFO  ITSCSV00 (0x20 in byte 16)
;
EJFEFDEDFDFGDADACACACACACACACAAA 86400  IN  CNAME  ITSCWK00
                                       IN  HINFO  ITSCSV00 (0x00 in byte 16)
;
EJFEFDEDFDFGDADACACACACACACACAAD 86400  IN  CNAME  ITSCWK00
                                       IN  HINFO  ITSCSV00 (0x03 in byte 16)
;
EJFEFDEDFDFGDADACACACACACACACAAA 86400  IN  CNAME  ITSCWK00
                                       IN  HINFO  ITSCAUS (0x00 in byte 16)
;
ITSCWK01                        86400  IN  A     129.35.144.211
                                       IN  HINFO  AS HOST NAME
;
EJFEFDEDFDFGDADBCACACACACACACACA 86400  IN  CNAME  ITSCWK01
                                       IN  HINFO  ITSCSV01 (0x20 in byte 16)
;
EJFEFDEDFDFGDADBCACACACACACACAAA 86400  IN  CNAME  ITSCWK01
                                       IN  HINFO  ITSCSV01 (0x00 in byte 16)
;
EJFEFDEDFDFGDADBCACACACACACACAAD 86400  IN  CNAME  ITSCWK01
                                       IN  HINFO  ITSCSV01 (0x03 in byte 16)
;
```

Figure 27. Sample DNS Database File after Adding Encoded NetBIOS Names. The encoded NetBIOS names point back to the TCP/IP host names (using CNAME*), where the workstation IP addresses are specified.*

For the domain controller (ITSCSV00), there are four encoded entries, three for the server name (computername) and one for the domain name (ITSCAUS). For the additional server (ITSCSV01), there are three encoded entries for the server name. The encoded entries are all aliases that point back to the host names.

7. On your clients, be sure that you set the DOMAINSCOPE parameter to point to the correct TCP/IP domain, for example:

   ```
   DOMAINSCOPE=austin.ibm.com
   ```

   This enables TCPBEUI to use the DNS to find the NetBIOS name/IP address pairs, eliminating the need for a broadcast file or names file at each client.

Notes:

1. It does not make any difference if you are using an existing DNS server or if you are using a new dynamic DNS server that is a part of TCP/IP Services of OS/2 Warp Server. Since the dynamic DNS server cannot determine the difference between a TCP/IP host name and an RFC-encoded NetBIOS name, you still have to add those resource records manually.
2. The RFCs 1001/1002 also specify a NetBIOS name server and NetBIOS datagram distribution server functions. Apart from returning IP addresses when queried with NetBIOS names, those servers also allow clients to register, update, and delete their NetBIOS names and IP addresses with the server dynamically. RFC NetBIOS servers also take care of proper NetBIOS datagram delivery throughout a TCP/IP network. Such functions are not implemented in OS/2 Warp Server.

For further information on the Domain Name Server, please refer to the *DNS Administration Reference* and *Dynamic DNS Implementation Guide*, available as online books with OS/2 Warp Server. They are located in the DDNS Server Services folder inside the TCP/IP folder.

3.10 Configuring TCPBEUI to Support 1000 Clients

The new TCPBEUI protocol driver that is included in OS/2 Warp Server can be bound to one adapter up to four times, thus providing the capability to support 1000 client workstations by using the NetBIOS over TCP/IP protocol. The following should be considered before setting up this kind of configuration:

1. One adapter with four TCPBEUIs

   This is the only method possible to provide TCPBEUI support for 1000 clients. This is because TCPBEUI can only be used with the lan0 TCP/IP interface.

2. Four adapters on different IP subnets

   This configuration allows you to start the server or requester without any errors, but no clients or peers can connect from any IP subnets other than the one used with the lan0 interface. This configuration is therefore neither recommended nor supported.

3. Four adapters on the same IP subnet

   This configuration is not possible since TCPBEUI will detect a NetBIOS name conflict. This configuration is therefore neither recommended nor supported.

4. WRx8210 MPTS FixPak/Refresh Warp Server

   This FixPak is needed to run TCP/IP on multiple adapters even if all TCPBEUI protocols are on lan0 (There used to be a missing filtering function).

Figure 28 illustrates how TCPBEUI can be used four times over a single LAN adapter in order to support 1000 NetBIOS clients from a single OS/2 Warp Server system.

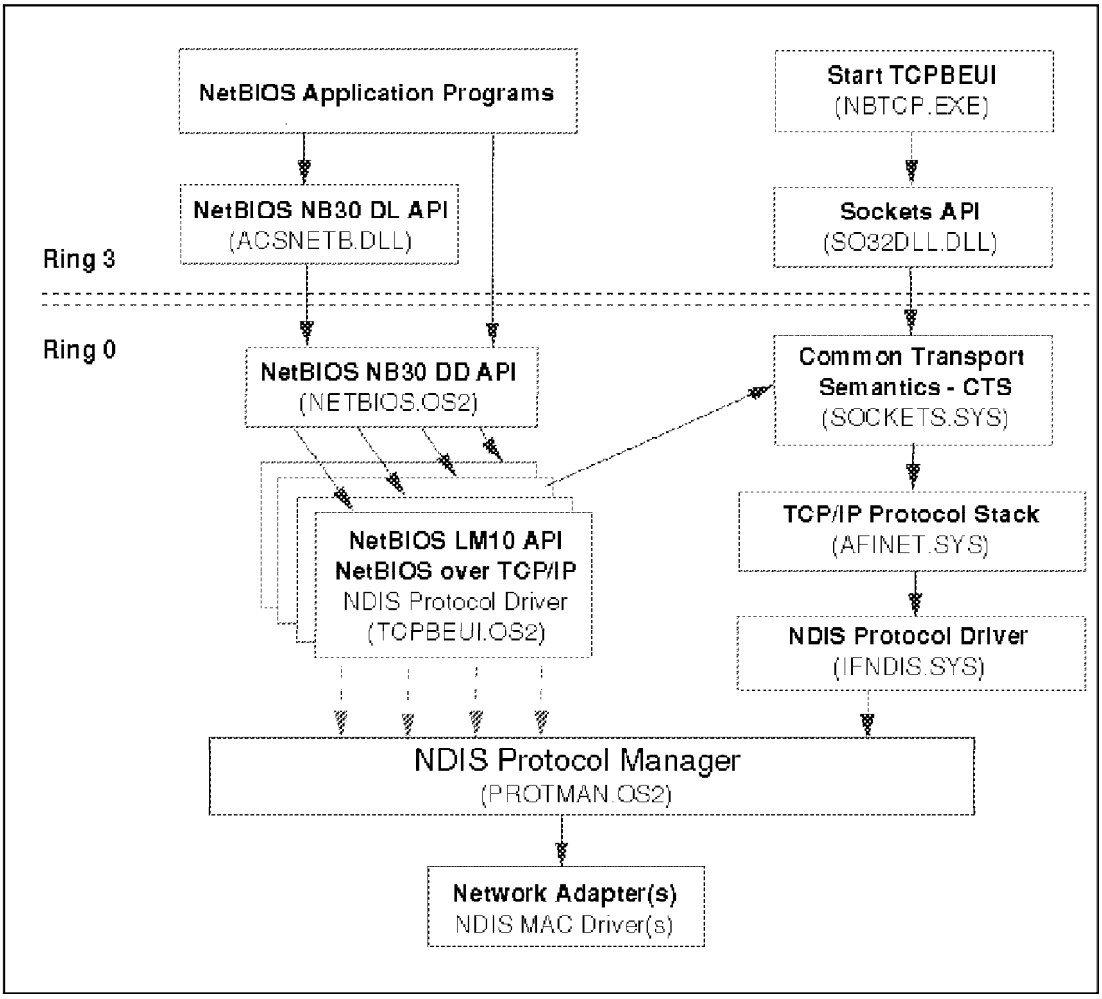

Figure 28. TCPBEUI Configuration for 1000 Clients

The following lines are extracted from the PROTOCOL.INI file to reflect what drivers must be loaded in order to support the configuration as shown above:

```
...

[NETBIOS]

   DriverName = netbios$
```

```
   ADAPTER0 = tcpbeui$,0
   ADAPTER1 = tcpbeui$,1
   ADAPTER2 = tcpbeui$,2
   ADAPTER3 = tcpbeui$,3

[tcpbeui_nif]

   DriverName = tcpbeui$
   Bindings = IBMTOK_nif,IBMTOK_nif,IBMTOK_nif,IBMTOK_nif
   OS2TRACEMASK = 0x0
   SESSIONS = 130
   NCBS = 225
   NAMES = 21
   SELECTORS = 15
   USEMAXDATAGRAM = "NO"
   NETBIOSTIMEOUT = 500
   NETBIOSRETRIES = 2
   NAMECACHE = 100
   PRELOADCACHE = "YES"
   NAMESFILE = 50
   DATAGRAMPACKETS = 20
   PACKETS = 150

[tcpip_nif]

   DriverName = TCPIP$
   Bindings = IBMTOK_nif

[IBMTOK_nif]

   DriverName = IBMTOK$
   MAXTRANSMITS = 6
   RECVBUFS = 2
   RECVBUFSIZE = 256
   XMITBUFS = 2
   XMITBUFSIZE = 4224
```

3.11 Using TCPBEUI with Dial-Up Connections

The purpose of NetBIOS over TCP/IP is to allow applications to use the NetBIOS protocol in a wide area network (WAN). So far, we have discussed the usage of TCPBEUI in LAN configurations only. This implies that a router must be available somewhere in order to actually access the TCP/IP WAN. What if the gateway to the WAN should be the OS/2 Warp Server itself, and the remote client does not have any LAN attachment? The following points should be considered for that kind of configuration:

1. As we have seen before, TCPBEUI will only work with the lan0 TCP/IP interface.
2. Since TCPBEUI is implemented as an NDIS protocol driver, it must be bound to an adapter driver in PROTOCOL.INI.
3. There is no Network Driver Interface Specification (NDIS) loopback MAC driver supplied with OS/2 Warp Server in order to fake a lan0 interface for TCP/IP and TCPBEUI.

Therefore, a dial-up connection for TCPBEUI will only work if a physical LAN connection exists for the systems on either end of the WAN link. Hence, we recommend using Remote Access Services in this case.

3.12 Performance Considerations for TCPBEUI

The performance when using TCPBEUI is generally slower than using native NetBIOS due to the additional overhead of mapping NetBIOS API calls to TCP/IP. (However, using OS/2 Warp Server over TCPBEUI is significantly faster than using LAN Server 3.0 with the TCP/IP 2.0 NetBIOS kit because there is no longer a transition overhead from ring 3 to ring 0.) The performance difference can range widely depending on the environment.

Some environmental factors that can affect performance are the type of client (OS/2 or DOS), the server CPU workload, the type of network operations being performed, the network media, network congestion, and communication line speeds. We've observed the performance of NetBIOS over TCP/IP being anywhere from eight percent slower to as much as up to twenty percent slower than NetBEUI. For example, having a NetBIOS name server in the network increases TCPBEUI performance dramatically (see 3.13, “Using NetBIOS Name Server” on page 82 for more details).

One of the environments in which performance tests were conducted was a medium-sized LAN on a 16 Mbps token-ring with no WAN connections. We ran a set of industry standard business applications on TCPBEUI clients and again on OS/2 NetBEUI clients. In this environment, NetBIOS over TCP/IP was 20 percent slower than NetBEUI. The performance of DOS LAN Services (DLS) NetBIOS over TCP/IP clients was significantly better than that of the OS/2 clients.

Performance Difference Not Noticeable

The user may not notice a difference in performance in the two protocols.

Database applications generally use small records when accessing shared databases residing on the server. Often these small records are retrieved from the file system cache with no physical disk access being required. The performance of this type of application on NetBIOS over TCP/IP may be noticeably slower than if the application were run using NetBEUI. However, if the number of database accesses of this type in performing a typical operation is in the order of hundreds, not thousands, the user may not notice a difference in performance in the two protocols.

Speed of WAN Connections

Most WAN connections today are made over relatively low-speed communication lines when compared with a LAN speed of 4 to 16 Mbps. This fact is more important to be aware of than the discussion about efficiency of the different protocols.

It may be necessary to periodically update client applications or other files by copying them from the server disk. DCDB replication from a domain controller to a remote additional server also generates I/O operations, sometimes known as file transfers. This type of file I/O activity over a network shows little or no performance difference between NetBEUI and TCPBEUI due to protocol characteristics. One should be aware, however, that most WAN connections today are made over relatively low-speed communication lines when compared with a LAN speed of 4 to 16 Mbps. File transfer operations over WAN communication lines will probably be slower than over LANs but most likely not due to the network protocol.

3.12.1 Tuning Considerations for TCPBEUI

If you're using NetBIOS over TCP/IP in a token-ring environment, file transfer performance might be improved by increasing the Maximum Transfer Unit (MTU) size. We have seen up to a 20 percent increase in performance of large file transfers by using an 8 KB packet instead of the default 1500 bytes. The default of 1500 was chosen because of Ethernet's packet size limitation and prevalence in TCP/IP environments. The MTU size can be changed with the `IFCONFIG` command in the MPTN BIN SETUP.CMD file.

Set the `MTU` size to the desired packet size plus 40 bytes, the maximum TCP/IP header size. The desired packet size should be a multiple of 2048. Your network adapter must be configured to support transmission of buffers that are at least the size specified for the MTU. On an IBM 16/4 Token-Ring Adapter, this would be accomplished by setting the `XMITBUFSIZE` parameter in the token-ring section of the PROTOCOL.INI file.

Note: If you use LAN adapter cards that need a system memory area below 1 MB to map buffer space (memory-mapped I/O), make sure the adapter RAM is set to at least 16 KB before you increase the `XMITBUFSIZE` and `MTU` size parameters.

Check your network interface card documentation for information on configuring your adapter.

It is also recommended that you use the INETCFG program to change the default `keepalive` value from the default of 120 minutes to a lower value. The example of the command input is:

```
inetcfg keepalive=3
```

The reason for this is that a TCPBEUI server is *not* informed of a TCP/IP connection breaking for a period of two hours. Thus, a TCPBEUI server could accumulate a large number of *ghost* sessions. By issuing the `inetcfg keepalive=3` command, TCP/IP will inform TCPBEUI after 3 minutes that a TCP/IP connection is broken (that is, a remote client has gone down).

If you are experiencing difficulties accessing a remote server over a slow WAN connection, try gradually increasing the `NETBIOSTIMEOUT` parameter in PROTOCOL.INI.

When using both NetBEUI (for LAN access) and TCPBEUI (for WAN access), it is best to have both `net1=NETBEUI$` and `net2=TCPBEUI$`, as shown in Figure 22 on page 66. In this dual protocol environment, it is recommended that you decrease `NETBIOSRETRIES` to 2 or 3 (from the current default of 8). Also, be aware that if the `NETBIOSTIMEOUT` parameter is set too high, some local LAN functions, such as logon or `NET USE` command, may take significantly longer.

When TCPBEUI is configured for more than 250 sessions, it is recommended to increase the value of the `PACKETS` parameter to 150.

Recommendation - Dual Protocol Stack

Because there may be a performance difference in a particular environment, it is recommended to configure and use NetBEUI in the local area network (LAN) environment and to use NetBIOS over TCP/IP in the wide area network (WAN) environment. The Multi Protocol and Transport Services shipped OS/2 Warp Server provide the capability of configuring your server with both NetBEUI and TCPBEUI on the same network interface card.

The dual protocol stack can be configured through the installation/configuration program. When selecting protocols, install logical adapter 0 with NetBEUI and logical adapter 1 with TCP/IP and NetBIOS over TCP/IP (on the same physical adapter). This dual protocol stack configuration allows local sessions to continue running with NetBEUI performance while also providing WAN connectivity with TCPBEUI.

3.13 Using NetBIOS Name Server

In the marketplace, there are some vendor products that support implementation of the RFC 1001/1002 standard NetBIOS Name Server (NBNS). In this section, we cover Microsoft's WINS and Network TeleSystems (NTS) product called Shadow. How to configure each NetBIOS Name Server is covered in the practical part of this book (Part III).

An IPX NetWare implementation has no need of the service offered by a NetBIOS Name Server; so NetWare systems will not be considered further in this section.

Figure 29 illustrates how NetBIOS Name Server works for server and clients over a TCP/IP network.

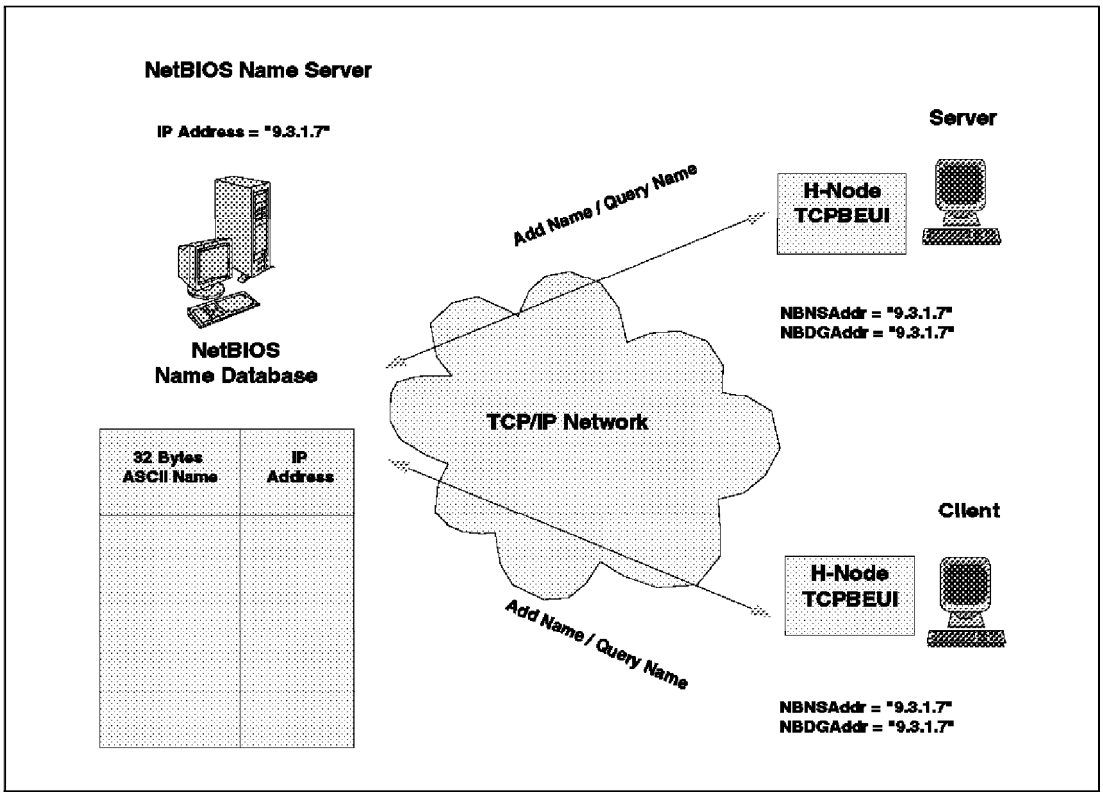

Figure 29. How Server and Client Work with NetBIOS Name Server

In Figure 29, the server registers its domain name (NetBIOS group name) and server name (unique NetBIOS name) to the NetBIOS Name Server (NBNS hereafter). OS/2 TCPBEUI should have h-node support, and PROTOCOL.INI should have a pointer to NBNS, which is the IP address 9.3.1.7 in the figure. The OS/2 client should have the same TCPBEUI code and PROTOCOL.INI file. When the OS/2 requester starts, it registers the requester name to NBNS.

When a user logs on from the client, it must find an IP address of the domain. The requester queries NBNS with the target domain name and NBNS returns the server's IP address by looking up its database. Then the requester establishes the session with the server through the TCP/IP network. Figure 30 shows the NetBIOS name registration and query process. All of these operations can be done through different subnets and without any broadcast operation.

First, the DHCP client sends out a message asking for a DHCP server. This is done by sending out a datagram. This datagram is a BootP broadcast message. Usually, routers do not forward broadcast messages. The BootP standard got around this by defining an RFC 1542, a specification whereby routers would recognize BootP broadcasts and would forward them to other subnets. This feature must be implemented in the router's software, and it is commonly known as *BootP forwarding*.

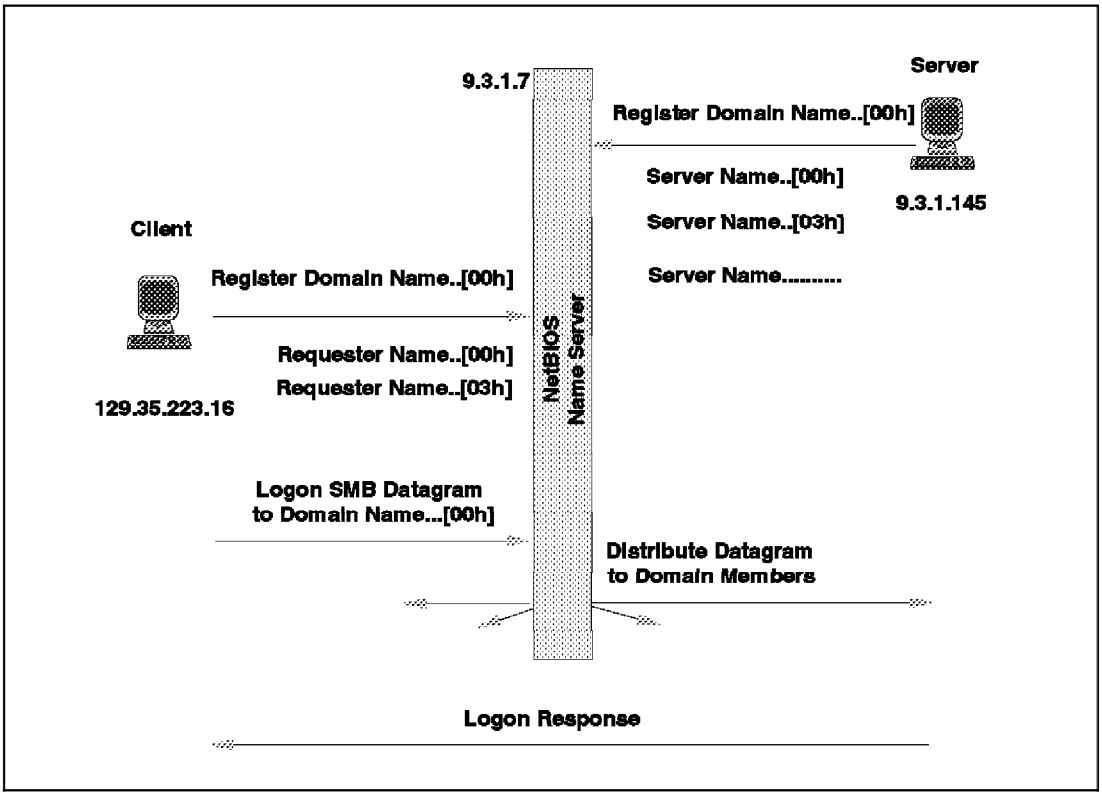

Figure 30. Detail Flow of Server/Client to NetBIOS Name Server

When a SMB datagram is sent from a client to a real NetBIOS local network, it is a datagram to a group name. It means there might be more than one member in a group. Requesters are the members if the `DOMAIN =` parameter in the IBMLAN.INI file has the same domain name. More important domain members are additional servers and backup domain controllers. In a real NetBIOS network, backup domain controllers receive the same logon SMB

datagram; so in case the primary domain controller is down, the backup will respond to process the logon.

In the TCP/IP network this process is defined as a datagram distributor. Without a datagram distributor function on the network, NetBIOS over TCP/IP has less function than a real NetBIOS network. Datagrams are sent to all the TCP/IP hosts on a network or subnetwork. A datagram sent to the broadcast address is received by all the hosts on the network and processed as if the datagram was sent directly to the host's IP address.

With this datagram distributor function of NetBIOS Name Server, we can have a backup domain controller somewhere in the TCP/IP network. With the DNS name resolution technique described in 3.9.4, “Storing NetBIOS Names on the Domain Name Server (DNS)” on page 70, we cannot have a backup domain controller with the same TCP/IP hostname. One DNS domain file entry must have only one IP address associated with it.

NTS's NetBIOS Name Server (Shadow) supports a full datagram distributor function. Microsoft Windows NT server has a Windows Internet Name Service (WINS) server function, and it is similar to NTS's NetBIOS Name Server. However WINS does not have a datagram distributor function; so we cannot recommend WINS as our NetBIOS Name Server, especially when it comes to large networks.

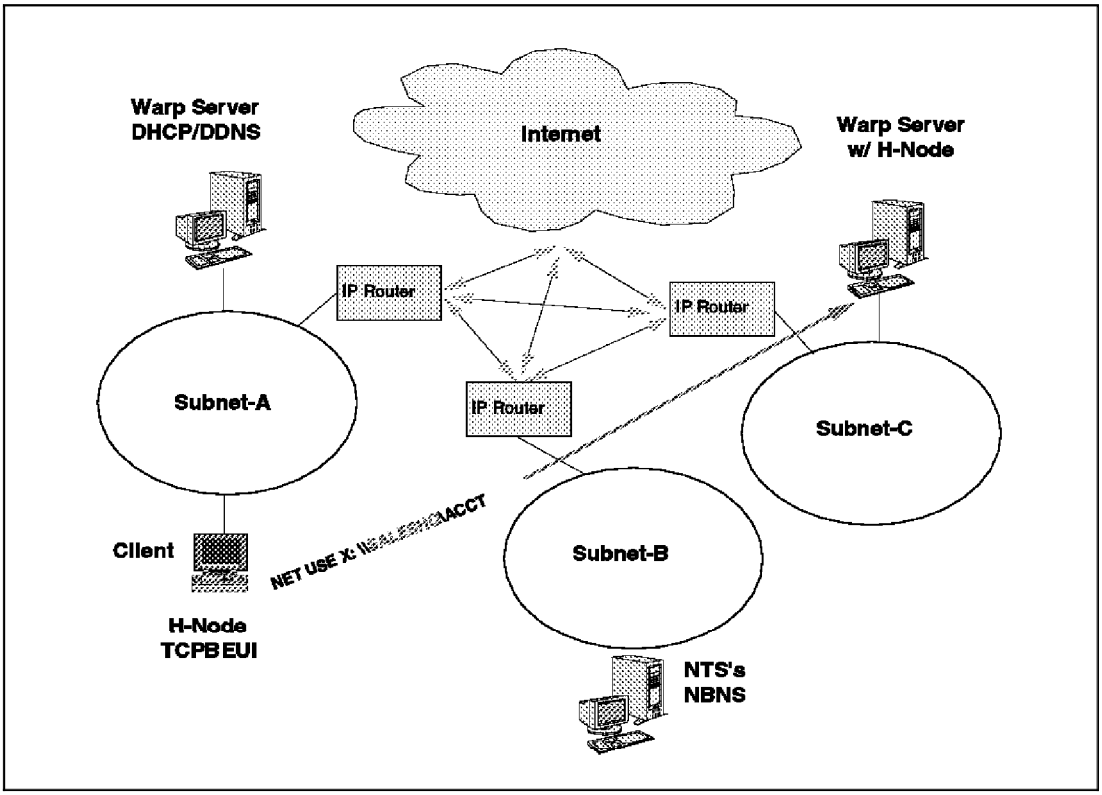

Figure 31. Ideal Solution with DHCP/DDNS plus NBNS

Figure 31 shows the ideal solution for TCP/IP applications such as a Web browser and LAN Server/Clients over TCP/IP network.

3.13.1 NBNS in a Dynamic IP Environment

The NetBIOS name resolution procedure in an dynamic IP environment is shown in Figure 32 on page 86.

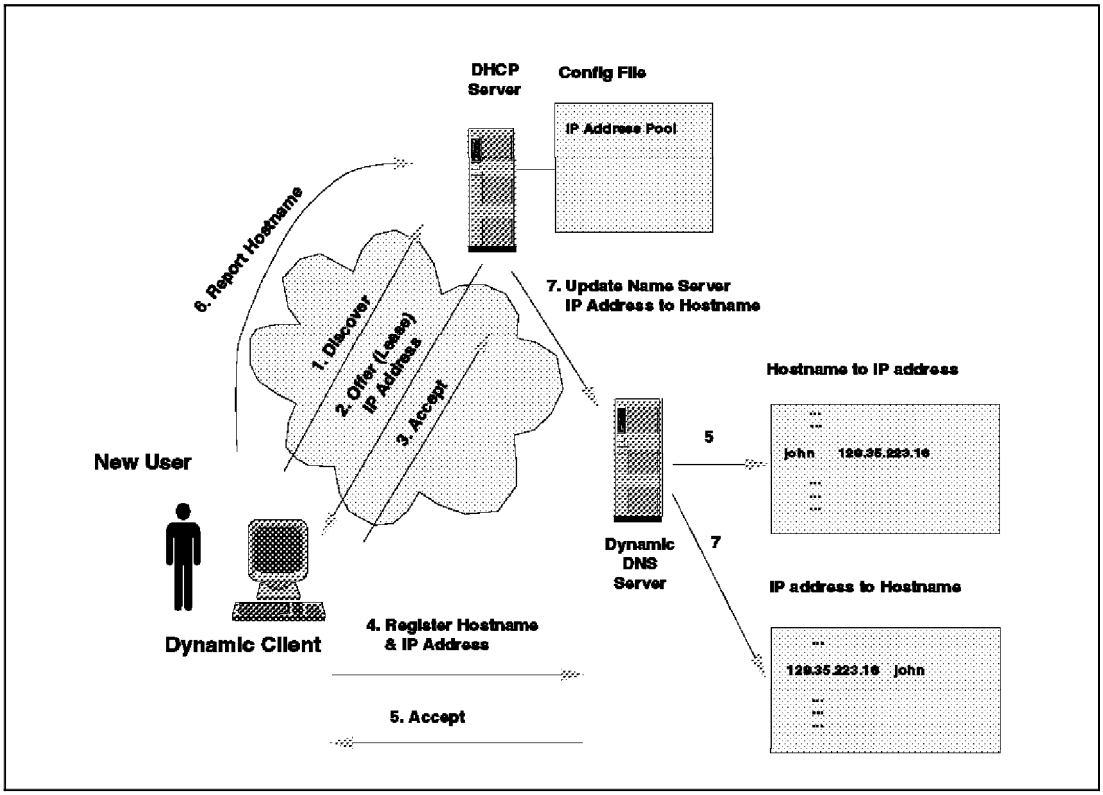

Figure 32. NetBIOS Name Resolution in Dynamic IP Environment

The data flow in this environment is as follows:

1. The LAN client is broadcasting a DHCPDISCOVER datagram to discover an DHCP server and asking for an IP address.
2. The DHCP server is sending back an DHCPOFFER, offering an IP address. If there is more than one server, the client receives several offers.
3. The client selects one of the offerings and sends the selected IP address via DHCPREQUEST to the selected DHCP server. The DHCP server is verifying if this address is unique in the network (via ping) and if this is the case, it is sending an acknowledgement and additional configuration information back to the client. If the address is not unique in the network, the whole process starts again.
4. The client is notifying the DDNS server, giving host name to IP address mapping information.
5. The DDNS server is registering the hostname to IP address mapping and sends an acceptance message to the client.
6. The client reports its hostname to the DHCP server.

7. The DHCP server is sending an update of the IP address to hostname mapping to the DDNS server, which is registering this mapping in its database.

There are several advantages in a dynamic IP environment when compared to a static IP environment:

- The IP address is given to the client on a leasing base. It can be leased for minutes, hours, days and so forth. It is reusable after the lease is expired. This is good for a mobile environment. With the static solution, different IP addresses are to be assigned to mobile users, and therefore they have to use different hostnames in different locations.
- All the administration work is done automatically after the initial configuration. In a static environment, the administrator has to assign all mappings manually.
- The client is supplied with all TCP/IP configuration information like:
  - IP address
  - IP subnet mask
  - Default router address
  - Domain name
  - Name server address
  - Local host name

  In a static environment, the configuration of each client has to be done by the administrator manually.

3.13.2 Shadow

Shadow, developed by Network TeleSystems (NTS) is both a NetBIOS Name Server and a NetBIOS Datagram Distributor (NBDD) built according to the Internet Standards defined in RFCs 1001 and 1002. Both NBNS and NBDD functions are needed in a NetBIOS-over-TCP/IP system if it is to allow both NetBIOS connections and NetBIOS transaction services (datagrams) uniformly throughout a large enterprise net. Shadow also provides the platform on which NTS will implement additional enterprise services, such as the Dynamic Host Configuration Protocol, BootP (the Bootstrap Protocol, and Domain Name-to-IP address mapping.

Shadow operates on a stand-alone, dedicated PC. The software, which is written as 32-bit, protected-mode code provides high performance for large enterprise networks that may have thousands of NetBIOS entries that must be resolved "on the fly" without degrading the performance of the network.

Shadow supports a wide range of functions which are essential to the administration of a NetBIOS and TCP/IP enterprise environment. The primary features are listed below.

- NBNS Services
- NBDD Services
- DHCP/BootP Services

Services include static entries, dynamic entries, replicated (backup) servers, and distributed servers. Shadow can be configured and operated as a DHCP/BootP server for TCP/IP client stations.

More than one Shadow can be used to provide naming and datagram distribution services in a large enterprise network. Shadows can be used in pairs, as backups, which replicate one another's' database for purposes of redundancy and load balancing, or can be used in groupings as large as 16 coservers, which maintain disjoint portions of the composite database for efficiency purposes. Each Shadow server can maintain up to 65536 NetBIOS names, thus allowing a total name space in a single Shadow system of over one million names.

3.13.2.1 Simple Network Management Protocol (SNMP) Services

A limited set of SNMP functions has been incorporated into Shadow. Remote SNMP managers may query Shadow for information regarding the status of the server.

3.13.2.2 Requirements

NTS's NetBIOS Name Server (Shadow) has the following requirements:

- Intel 80386 or 80486
- AT-compatible EISA or ISA bus
- IDE or Enhanced IDE hard disk
- 8 MB RAM for 16000 NetBIOS Names or 16 MB for 64000 NetBIOS Names
- FAT 16 File System
- DOS 6.3 or later
- Color Monitor, VGA
- LAN adapter such as IBM's Auto 16/4 Token-Ring ISA Adapter, P/N 92G7632 or an Eagle NE2000 Ethernet adapter, available from Microdyne.

Shadow runs on DOS, but it runs just like a network operating system and effectively uses the LAN adapter card interface and hard disk interface as fast as possible.

3.13.2.3 Why does Shadow only run on a Stand-alone System?

A NetBIOS Name Server should be compared to a router. Enterprise routers are built for speed. They need to handle packets at a rate that the networks offer them. This performance can never be achieved by an application program sitting atop a commercial operating system. A NetBIOS Name Server should not be slower than an enterprise router. It should be able to handle every packet arriving on one or more network adapters. At 10 Mbps, an Ethernet can carry over 10000 NetBIOS-sized packets per second. An NetBIOS Name Server deployed in an ATM environment needs to be able to handle packets delivered at 155 Mbps on one or more links. The only architecture that can sustain this type of packet flow is one built specifically for that purpose.

But there are some disadvantages when using special hardware. It is not very flexible, and it is expensive. PC platforms offer an advanced scalability and pricing flexibility. But if a PC platform is used and the design points must be met, the server has to be a dedicated system. The server software then can be designed to support the hardware for its special needs, without taking care of operating system limitations. This also improves the reliability. A comparison of the Shadow architecture to WINS architecture, as can be seen in Figure 33 on page 90, makes this very clear.

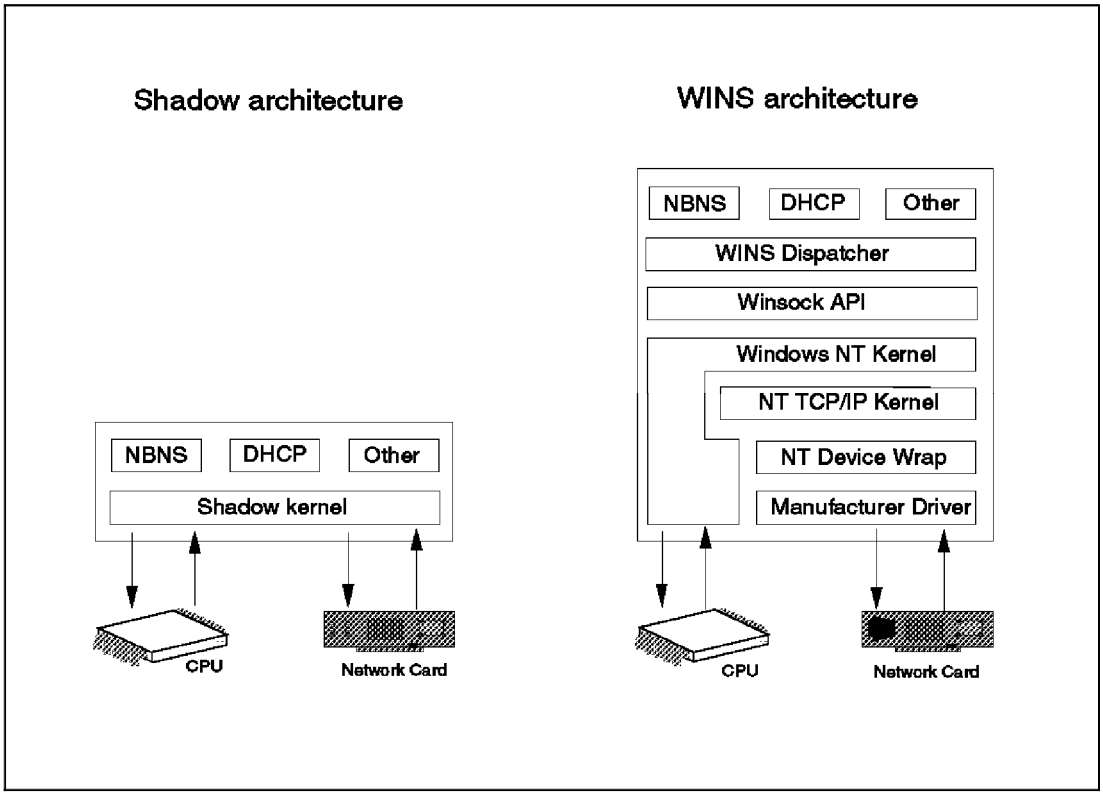

Figure 33. Comparison between Shadow and WINS Architecture

There can be a remote control station for Shadow running on top of Windows with NTS's TCPPro program or IBM OS/2 TCPBEUI workstation's WIN-OS/2 program. The following screen capture shows an example of an NBNS remote system manager.

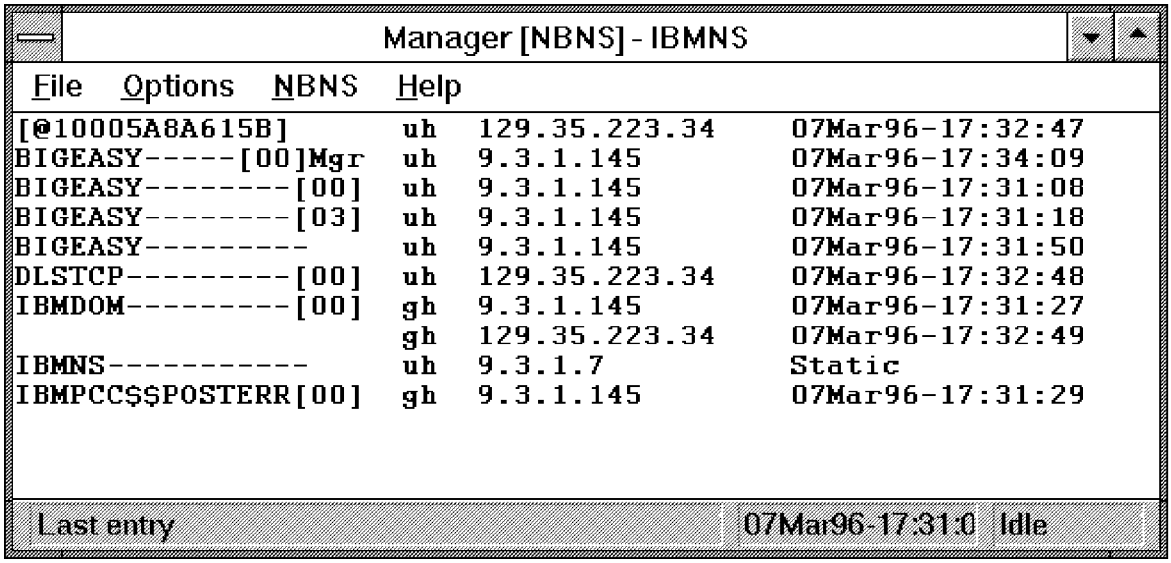

Figure 34. Example of Remote System Manager for NTS NBNS

3.14 Name Resolution for Microsoft Windows Networking

Because Windows NT Server is a Server Message Block (SMB) server that uses NetBIOS, it cannot use TCP/IP as its native protocol. Therefore it is necessary to perform name resolution tasks in the preferable same way they have done with OS/2 Warp Server.

Microsoft is hiding the fact, that Windows NT is using NetBIOS over TCP/IP. In the TCP/IP panels (see Figure 35) there is no item for configuring this task (as opposed to OS/2 Warp Server, see Figure 25 on page 69). Instead, if TCP/IP service is installed, all components for NetBIOS over TCP/IP are installed behind the scenes.

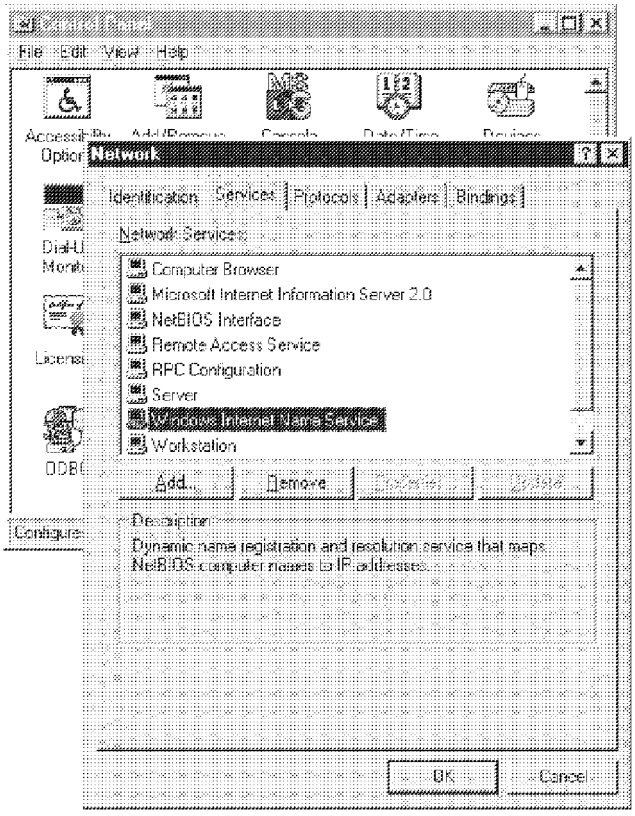

Figure 35. Selection of the WINS Service in Windows NT 4.0

Windows NT computers can use one or more of the following methods to ensure accurate name resolution in TCP/IP networks:

- Windows Internet Name Service

  Windows NT computers can use WINS if one or more WINS servers are available that contain a dynamic database mapping computer names to IP addresses. WINS can be used in conjunction with broadcast name

resolution for an internetwork where other name resolution methods are inadequate. This is a p-node operation.

- Broadcast name resolution

  Windows NT computers can also broadcast name resolutions, which is a b-node operation. This method relies on a computer making IP-level broadcasts to register its name by announcing it on the network. Each computer in the broadcast area is responsible for challenging attempts to register a duplicate name and for responding to name queries for its registered name.

- DNS name resolution

  The Domain Name Server (DNS) provides a way to look up name mappings.

- An LMHOST file to specify the NetBIOS computer name and IP address mappings, or a HOSTS file to specify the DNS name and IP address.

3.14.1 Name Resolution with HOST Files

For computers located on remote subnets where WINS is not used, the HOSTS and LMHOSTS files provide mappings for name to IP addresses. The HOSTS file can be used as a local DNS equivalent. The LMHOSTS file (similar to the NAMESFILE in Warp Server) can be used as a local WINS equivalent. Each of these files are also known as a *host table*. These files can be edited using any ASCII editor.

Microsoft TCP/IP can be configured to search HOSTS, the local host table file, for mapping of remote host names to IP addresses. The LMHOSTS file is a local text file that maps IP addresses to NetBIOS computer names for Windows-networking computers that communicate outside of the local subnet. The LMHOSTS file is read when WINS or broadcast name resolution fails, and resolved entries are stored in a system cache for later access.

When the computer uses the replicator service and does not use WINS, LMHOSTS entries are required on import and export servers for any computers on different subnets participating in the replication. LMHOSTS is also used for small-scale networks that do not have servers.

3.14.2 B-node in Combination with LMHOSTS

Another variation is also used in Microsoft networks to span routers without a WINS server and p-node mode. In this mode, b-node uses a list of computers and addresses stored in the LMHOSTS file. If a b-node attempt fails, the system looks in LMHOSTS to find a name and then uses the associated address to cross the router. However, each computer must have this list, which creates an administrative burden in maintaining and

distributing the list. Both Windows for Workgroups 3.11 and LAN Manager 2.x used such a modified b-node system. Windows NT uses this method if WINS servers are not used on the network. In Windows NT, some extensions have been added to this file to make it easier to manage, but modified b-node is not an ideal solution.

Some sites may need to use both b-node and p-node modes at the same site. Although this configuration can work, administrators must exercise extreme caution in doing so, using it only for transition situations. Because p-node hosts disregard broadcasts and b-node hosts rely on broadcasts for name resolution, the two hosts can potentially be configured with the same NetBIOS name, leading to unpredictable results. Notice that if a computer configured to use b-node has a static mapping in the WINS database, a computer configured to use p-node cannot use the same computer name.

3.14.3 Windows Internet Name Service (WINS)

WINS consists of two components: the WINS server, which handles name queries and registrations, and the client software, which queries for computer name resolution.

Windows networking clients can use WINS directly. Non-WINS computers on the network that are b-node-compatible, as described in RFCs 1001 and 1002, can access WINS through proxies, which are WINS-enabled computers that listen to name query broadcasts and then respond for names that are not on the local subnet or are p-node computers. This is because WINS is using a non-standard implementation of the p-node and h-node specifications. Other products, like Shadow, support a full datagram distributor function and therefore don't need proxies.

Figure 36 on page 94 shows a small internetwork with three local area networks connected by a router. Two of the subnets include WINS name servers that can be used by clients on both subnets. WINS-enabled computers, including proxies, access the WINS server directly, and the computer using broadcasts access the WINS server through proxies. Proxies only pass name query packets and verify that registrations do not duplicate existing systems in the WINS database. Proxies, however, do not register b-node systems in the WINS database.

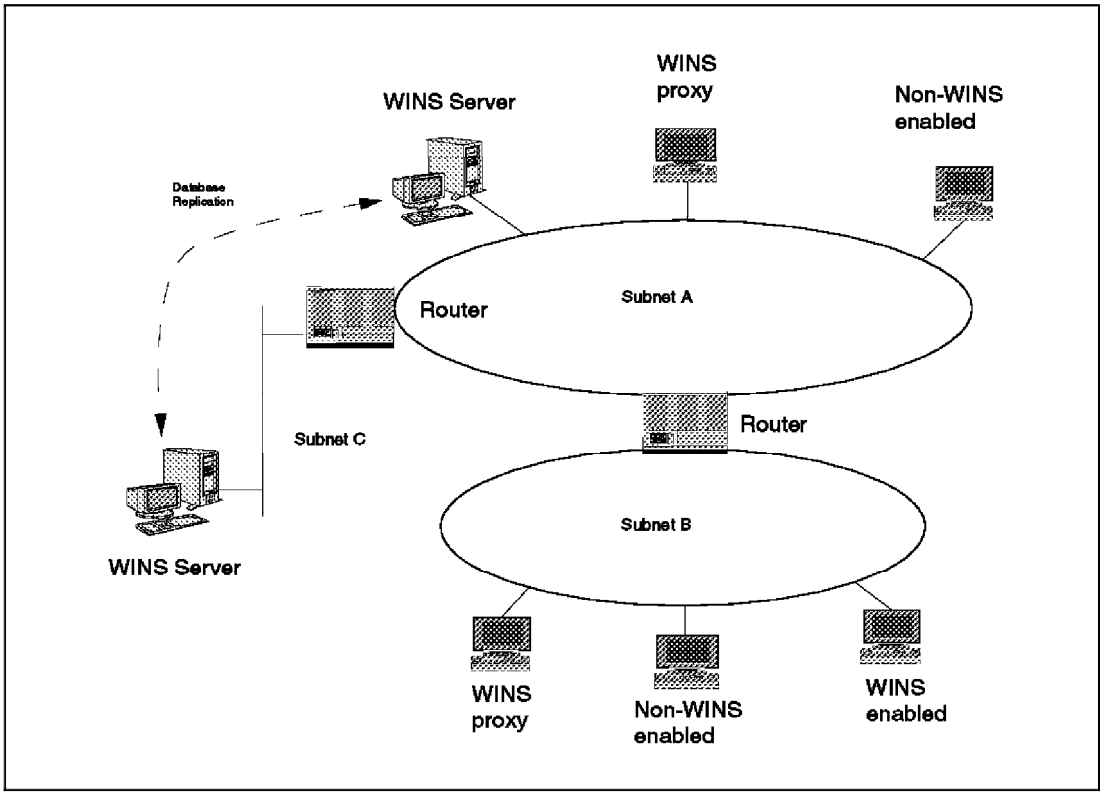

Figure 36. Example of an Internetwork with WINS Servers

The proxy communicates with the WINS server to resolve names (rather than maintaining its own database) and then caches the names for a certain time. The proxy serves as an intermediary by either communicating with the WINS server or supplying a name-to-IP address mapping from its cache. The following picture shows the relationships among WINS servers and clients, including proxies for non-WINS computers and the replication between WINS servers.

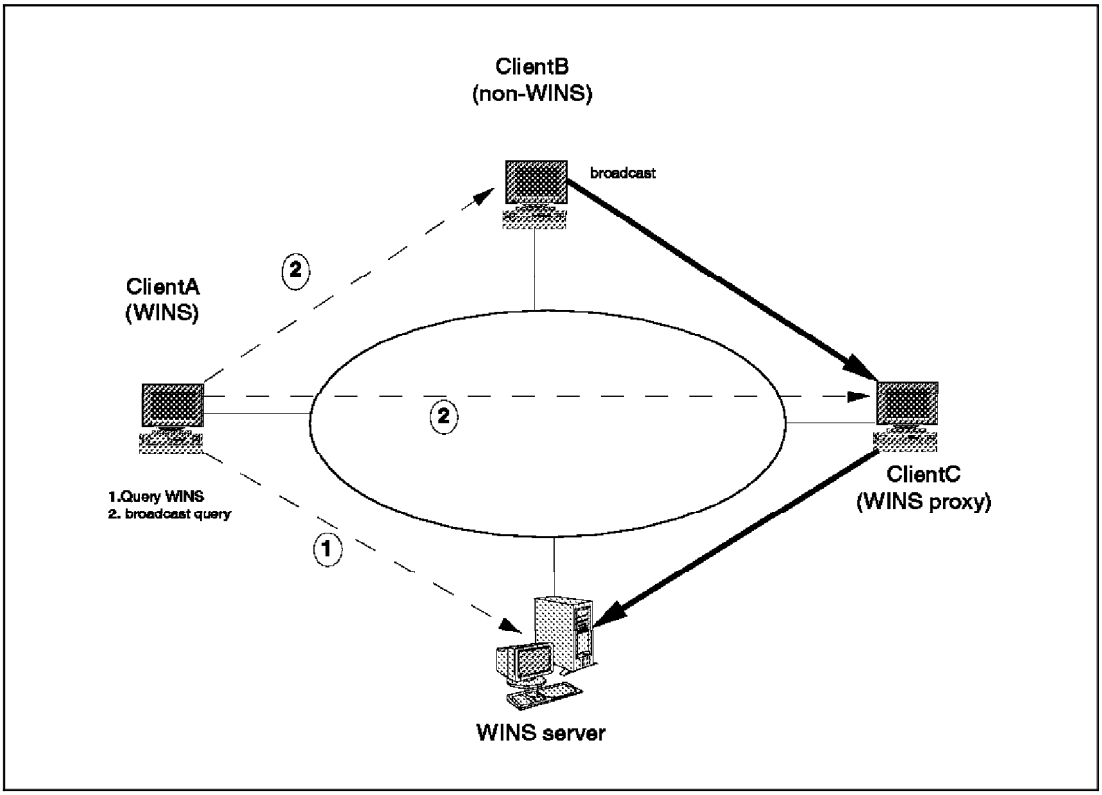

Figure 37. Example of Clients and Servers Using WINS

In Figure 37, ClientA can resolve names by first querying the WINS server and, if that fails, then using broadcast name queries. ClientB, which is not WINS-enabled, can only resolve names using broadcast name queries, but when ClientC receives the broadcast, it forwards the request to the WINS server and returns the address to ClientB.

However, a complex environment presents additional problems. For example, an internetwork might consist of two subnets with all the computers belonging to DomainA attached to Subnet1, all the computers in DomainB attached to Subnet2, and computers from DomainC attached to either of the subnets. In this case, without WINS, DomainA computers can browse Subnet1; DomainB computers can browse Subnet2, and DomainC computers can browse both subnets as long as the primary domain controller for DomainC is available. With WINS, computers from all domains can browse all subnets if their WINS servers share databases.

If the Windows NT client is also DHCP-enabled and the administrator specifies WINS server information as part of the DHCP options, the computer will usually be automatically configured with WINS server information. The WINS setting can be manually configured:

- To enable WINS name resolution for a computer that does not use DHCP, specify `WINS server addresses` in the TCP/IP Configuration dialog box.
- To designate a proxy, check the `Enable WINS Proxy Agent` option in the Advanced Microsoft TCP/IP Configuration dialog box.

With WINS servers in place on the internetwork, names are resolved using two basic methods, depending on whether WINS resolution is available and enabled on the particular computer. Whatever name resolution method is used, the process is transparent to the user after the system is configured.

3.14.3.1 If WINS is Not Enabled

The computer registers its name by broadcasting *name_registration_ request* packets to the local subnet via UDP datagrams. To find a particular computer, the non-WINS computer broadcasts *name_query_request* packets on the local subnet, although this broadcast cannot be passed on through IP routers. If local name resolution fails, the local LMHOSTS file is consulted. These processes are followed whether the computer is a network server, a workstation, or other device.

3.14.3.2 If WINS is Enabled

The computer first queries the WINS server, and if that does not succeed, it broadcasts its name registration and query requests via UDP datagrams (h-node) in the following series of steps:

1. During TCP/IP configuration, the computer's name is registered with the WINS server, and the IP address of the WINS server is stored locally so the WINS server can be found on the internetwork. The WINS database is replicated among all WINS servers on the internetwork.

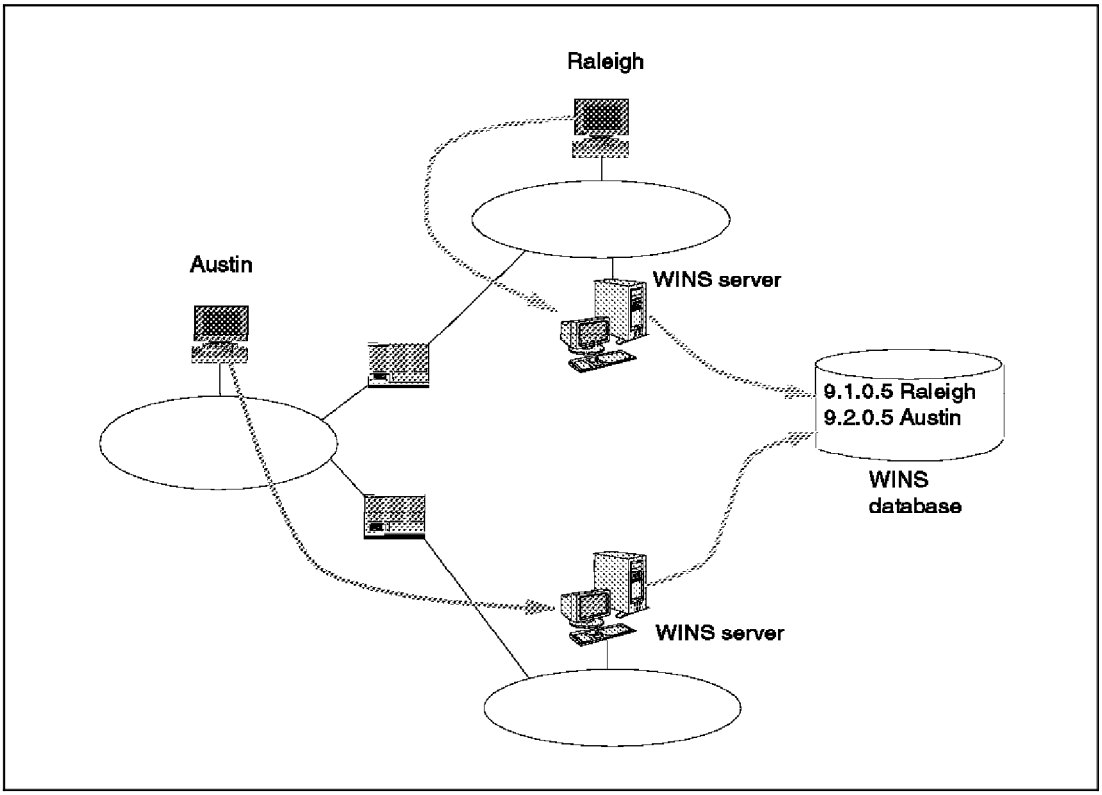

Figure 38. IP Address Registration

2. A *name_query_request* is send first to the WINS server, including requests from remote clients that are routed through an IP router. This request is a UPD datagram. If the name is found in the WINS database, the client can establish a session based on the address mapping received from WINS.

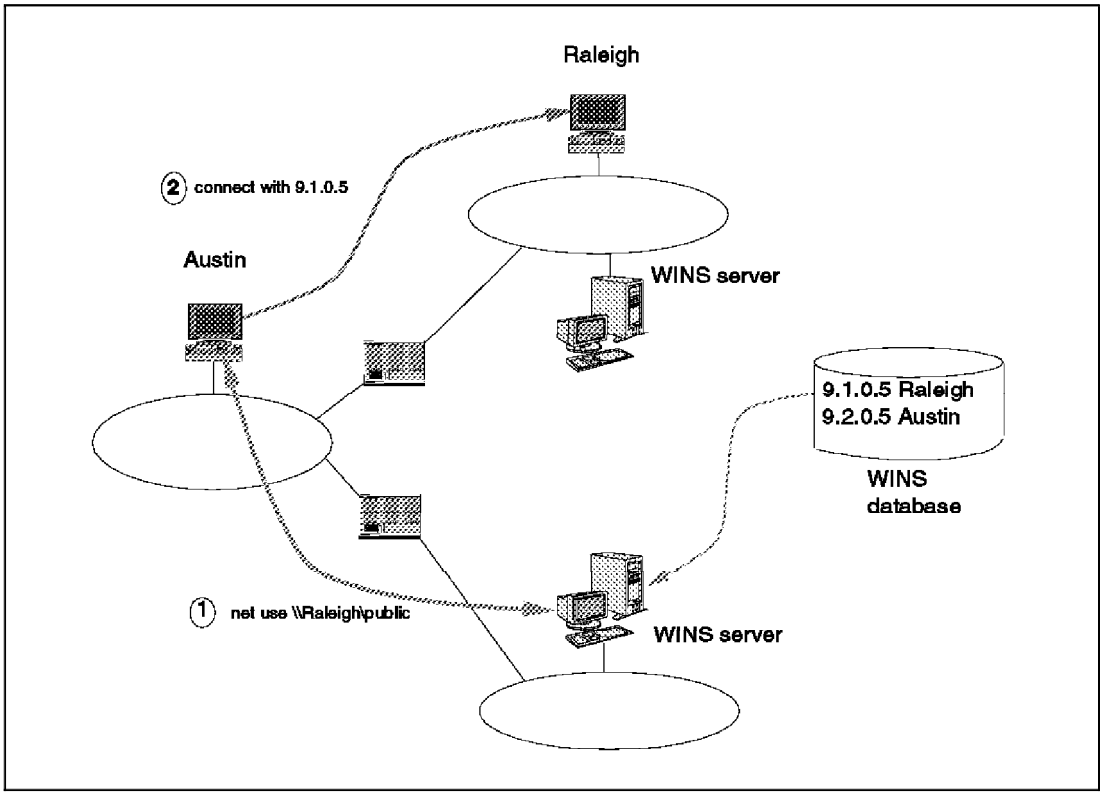

Figure 39. Request and Resolving of IP Address Mapping

3. If querying the WINS server does not succeed and if the client computer is configured as an h-node, the computer broadcasts *name_query_request* packets in the same manner as a non-WINS-enabled computer.
4. If all of the above methods fail, the local LMHOSTS file is checked. This also includes a search of any centralized LMHOSTS files referred to in #INCLUDE statements.

WINS servers accept and respond to UDP name queries. Any name-to-IP address mapping registered with a WINS server can be provided reliably as a response to a name query. However, a mapping in the database does not ensure that the related device is currently running, only that a computer claimed the particular IP address and it is a currently valid mapping.

3.14.3.3 WINS in an Enterprise Environment

Microsoft is using a proprietary implementation of RFC 1001/1002. Only a subset of these RFCs is used. This is a unique implementation of the domain name registration. The domain can have members. Those members are called the Internet Group. The Domain Name is registered with '1C'h as 16th byte. However, the standard defined in RFC 1001/1002 is '00'h as 16th byte. WINS does not treat it as an Internet Group name. This

leads to the fact, that domains registered in the WINS fashion do not have members. Figure 40 on page 99 shows the registration of Domain Names in WINS. A standard implementation is shown in Figure 30 on page 83.

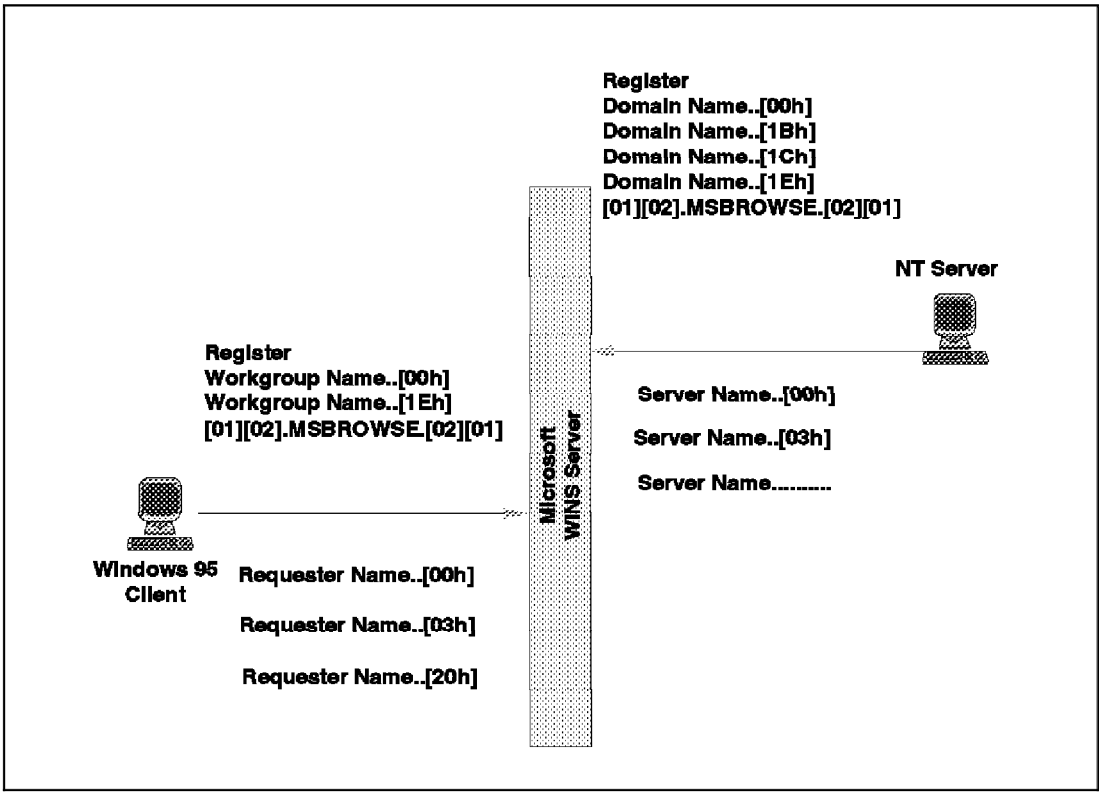

Figure 40. WINS Registration

This can cause problems in a heterogeneous enterprise environment. Windows clients are looking for '1C'h; they cannot logon to domains, which are registered with '00'h. Other clients, which are looking for '00'h, cannot logon to a Windows NT Domain (because it is registered with '1C'h), and they cannot logon to other domains (because they have no members).

3.15 TCPBEUI Interoperability with Microsoft

Because Microsoft is using a non-standard implementation of the domain registration, we will have a closer look to some scenarios in a mixed environment with Windows clients, OS/2 clients, OS/2 Warp Server, Windows NT server, and whether or not they can work together.

3.15.1 Using WINS as a NetBIOS Name Server

In the first scenario in Figure 41 on page 100, there is a IBM LAN Server and a WINS Server. There are only Microsoft clients, which are requesting IP addresses. The LAN Server Domain Name is registered at the WINS

Server. Remember, that byte 16 of the domain name is set to '00'h. The MS client tries to logon to the IBM LAN Server Domain. The client is issuing a Name Query and is looking for '1C'h in the 16'th byte. Because the IBM LAN Server Domain Name is registered as '00'h, the logon will fail. After that, it broadcasts the query in its subnet. The client can only log on, if the server is in the same subnet and has a LMHOST file provided with static configured Host to IP-address definitions. But this also means, that no logon is possible via routers.

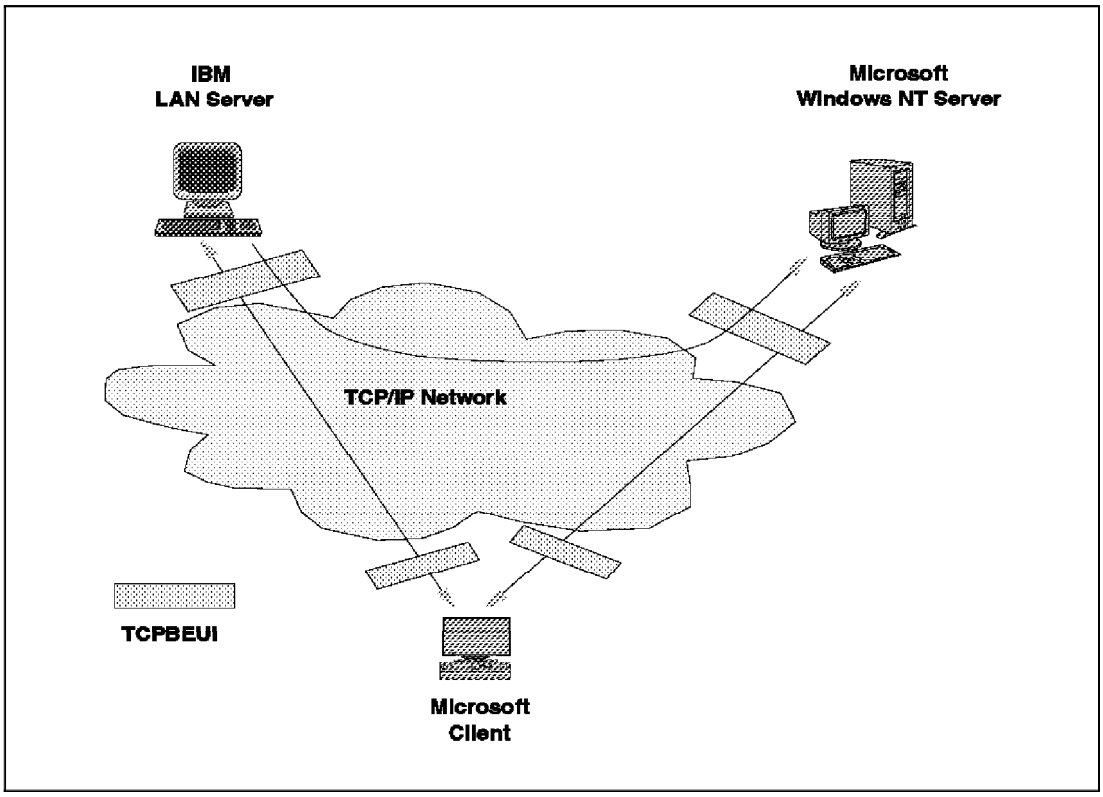

Figure 41. Using WINS as NetBIOS Name Server, Scenario 1

In the second scenario in Figure 42 on page 101, there is only a WINS server, but different IBM clients. To connect to WINS, h-node support is required. Warp 4 clients or OS/2 clients with the MPTS ServicePak IPx8210 installed support h-node. OS/2 clients and servers cannot log on to the NT server or to the IBM LAN Server due to the failing name registration at WINS. OS/2 LAN Server or OS/2 Warp Server are not "visible' to OS/2 clients anymore.

The Warp Server DOS LAN Services (DLS) client can register at the WINS server. H-node support is configurable by PROTOCOL.INI parameters. Because DLS supports Domain Names with '1C'h as 16'th byte, it can logon to an NT Server, but not to an IBM LAN Server because the Group Name of LAN Server has no members.

If WINS is used as NetBIOS Name Server, only IBM DLS clients can log on to NT Server; no IBM client can log on to IBM LAN Server.

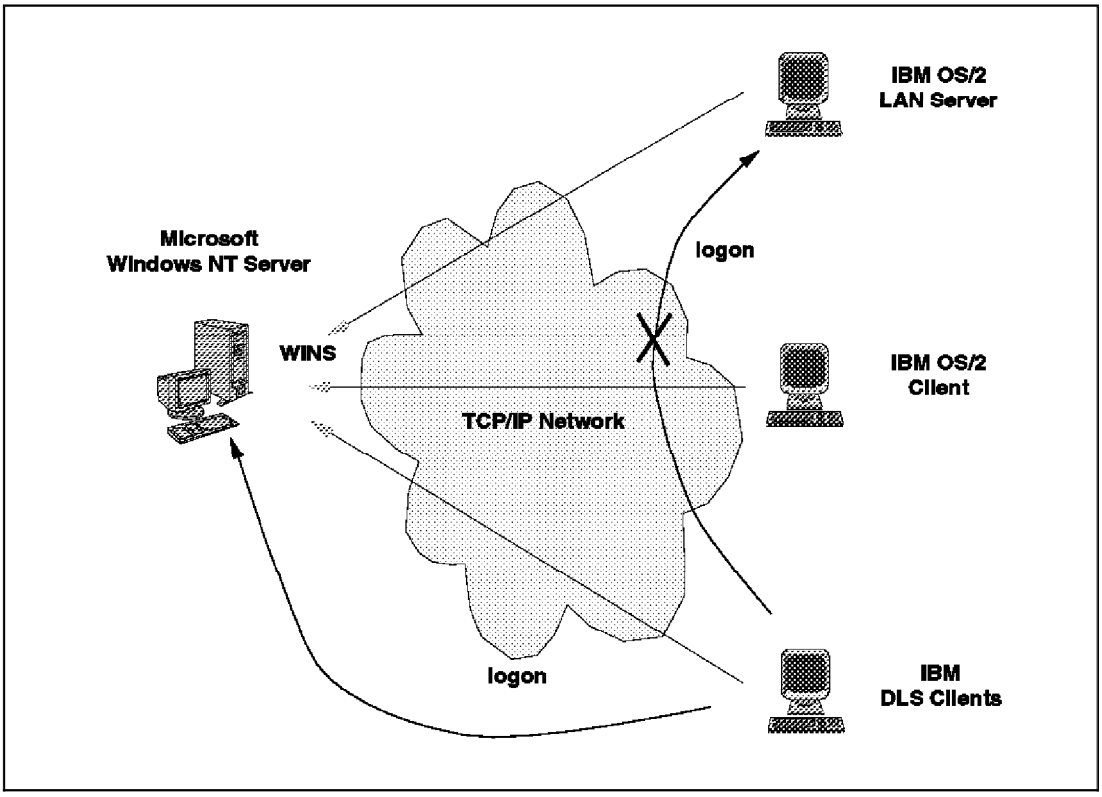

Figure 42. Using WINS as NetBIOS Name Server, Scenario 2

3.15.2 Using NTS's NetBIOS Name Server (Shadow)

In the scenario in Figure 43 on page 102, NTS's NetBIOS Name Server Shadow is used. Here all IBM Clients can logon to IBM servers. Warp Server DLS client can log on to Windows NT. Microsoft clients can log on to Microsoft Servers.

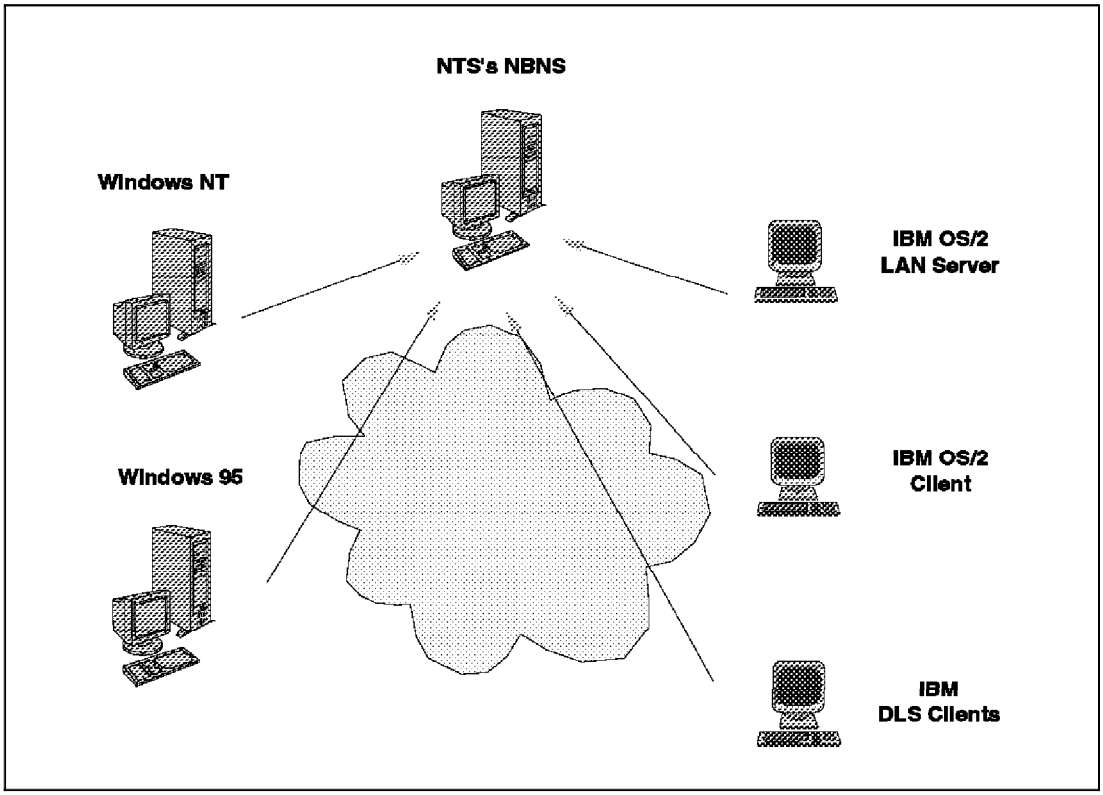

Figure 43. Using NTS's NetBIOS Name Server

Network TeleSystems' NetBIOS Name Server, Shadow, is 100 percent RFC 1001/1002 compliant. One area of RFC 1001 that is especially important to Warp Server is in the handling of NetBIOS group names. For a NetBIOS Name Server (NBNS) to be fully RFC 1001 compliant in this area, it should keep a list of IP addresses for all nodes that have added a particular group name. Additionally, an NBNS should keep the list of IP addresses for members of a group name regardless of the the name itself (regardless of the 16th byte).

Suppose your network consists of three subnets, A, B, and C, connected by routers. Subnet A contains a Warp Server with a computer name of `AUSTIN` and a domain name of `TEXAS`. Subnet B contains a Shadow NBNS. Subnet C contains clients of any type (DOS, WFW, OS/2, and NT). When the Warp Server starts, it will add three unique NetBIOS names to the NBNS database:

```
AUSTIN---------[00]
AUSTIN---------[03]
AUSTIN---------[20]
```

The Warp Server will also add its domain name as a NetBIOS group name to the NBNS database:

TEXAS----------[00]

Note: Remember, a NetBIOS name consists of 16 bytes of any bit pattern. The name AUSTIN---------[00] consists of six ASCII characters followed by nine blanks represented by the dashes and a hex 00 in the 16th byte.

When any client on subnet C attempts to log on to the Warp Server on subnet A, the client will need to resolve the Warp Server's domain name, TEXAS----------[00], into an IP address. The clients accomplish this by sending an RFC 1001 query packet to the Shadow NBNS. Shadow will respond with a list of IP addresses for all members of the TEXAS domain. The client can now proceed with the logon process.

NT servers can also use Shadow as its NetBIOS Name Server (NBNS). Also any Microsoft platform can use Shadow as its NBNS. Under the WINS configuration, change the IP address to point to Shadow. (Also ensure that the WINS service is not started. An NT Server will continue to use WINS if WINS is started even though the WINS IP address is pointing to Shadow).

The data flow is essentially the same. Both Microsoft clients and IBM clients have to resolve a file server's NetBIOS domain name into an IP address. When Shadow is used as the NBNS, it will keep a list of IP addresses for any NetBIOS group name. This includes Microsoft file servers' domain name with 0x1c in the 16th byte and IBM file servers' with 0x00 in the 16th byte. Therefore, clients of any type can contact file servers of any type. Microsoft clients query the NBNS for file server domain names with both 0x1c and 0x00 in the 16th byte. This allows them to contact both Microsoft servers and IBM servers.

Note:

Find a detailed example in the practical part (Part III) of this book.

3.15.3 Conclusion

In a pure Microsoft Windows environment WINS is a good solution for resolving the IP administration problems. In an heterogeneous environment WINS only has full functionality for the Windows client part. For other vendor clients, other NetBIOS Name Servers, like NTS's Shadow, have to be installed. Due to the proprietary implementation in WINS, the full datagram distribution function is not supported. WINS experiences performance problems if more than 300 users are in the network. NTS's Shadow is designed to support 64000 users without a significant performance decrease. Therefore, we cannot recommend WINS in an heterogeneous environment.

IBM OS/2 Warp Dynamic IP, on the one hand, is designed to serve all IP hosts and to update a host's name and IP address information in the universally-deployed Domain Name Server (DNS). Microsoft's DHCP/WINS system, on the other hand, is designed to serve Microsoft clients using NETBIOS-over-IP protocols only and to update a host's NetBIOS name in a NetBIOS Name Server. In other words, Dynamic IP is a general IP networking solution that has broad application and that scales easily to the Internet network. The Microsoft DHCP/WINS system is a limited networking system that is difficult to scale and that is useful only in the context of NetBIOS-over-IP workgroup networks.

3.16 NetWare and TCP/IP

Novell NetWare is using IPX/SPX as its networking protocol. But, as mentioned before, TCP/IP is becoming more common in large networking environments. Therefore, Novell had to offer a possibility to participate in a TCP/IP network.

Novell Inc.'s NetWare/IP 2.2 is a software option that is free for NetWare 4.1 customers and provides TCP/IP networking. The TCP/IP services built in to NetWare 4.1 offer simple IP services but no DHCP or DNS capabilities. NetWare/IP 2.2 fills this gap by providing both DHCP and DNS servers and is described in 3.16.6, “NetWare/IP 2.2” on page 110.

Usually, the TCP/IP NLM is loaded on a NetWare server where it can convert incoming TCP/IP datagrams. The last versions of NetWare have allowed simultaneous support of TCP/IP and IPX/SPX on the same network card because of the use of the Link Support Layer (LSL). TCP/IP fits into that layered structure in two differed ways, depending on the manner in which the protocol is implemented. TCP/IP can occupy a layer of its own (essentially replacing NetWare's IPX/SPX as the network protocol), or it can work with the NetWare protocols. This is the most commonly used method because it allows TCP/IP and NetWare to be used. Figure 44 on page 105 shows how TCP/IP works within the NetWare architecture to allow TCP/IP to function over a NetWare-based network.

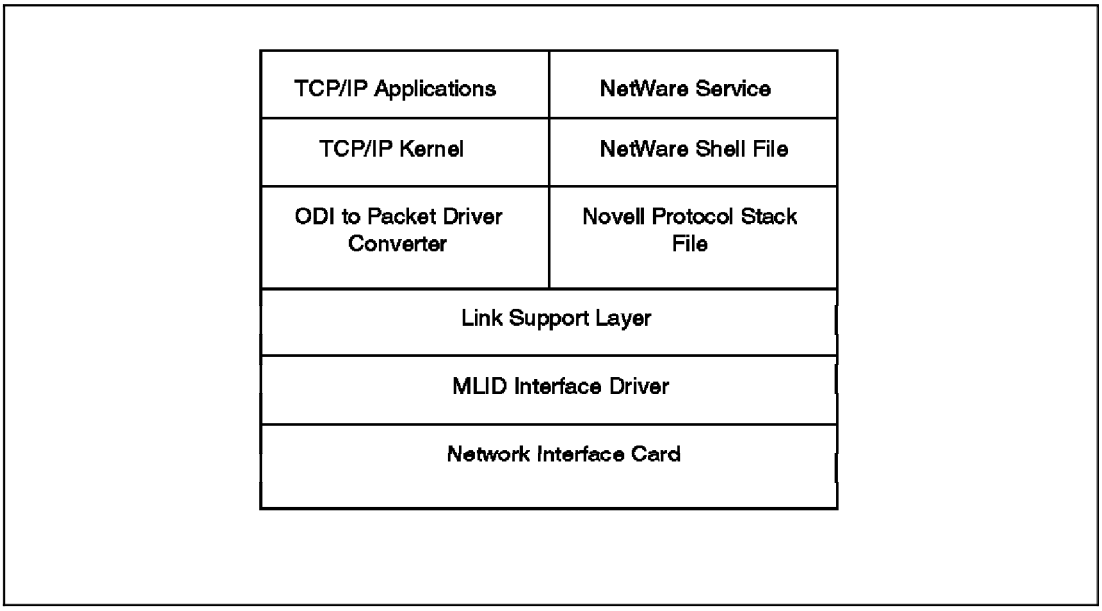

Figure 44. TCP/IP and NetWare Sharing Layers

The Open Data Link Interface (ODI) driver is used to provide compatibility between NetWare and other protocols that use the Novell IPX system. It acts to convert TCP/IP datagrams to Novell packets. The reverse is simply a matter of replacing the converter and the lower layers.

The ODI driver sits on top of the Link Support Layer (LSL), which acts as an interface to the Multiple Link Interface Driver (MLID). The LSL determines to which stack (TCP/IP or NetWare) to pass incoming packets. The MLID sends and receives packets from the network itself.

With Novell's NetWare server product, TCP/IP protocol can be used in several different ways as a full replacement or in combination with NetWare's network protocol. The following are the most common methods:

- Replacing NetWare's IPX with IP
- Using both TCP/IP and IPX/SPX at a workstation
- Using a TCP/IP gateway
- Tunneling IPX in IP packets

The following sections discuss each method in more detail.

Note: A disadvantage of using this method is the decrease of performance. By Novell's own admission, the NetWare IP product runs approximately 10 percent to 15 percent slower than NetWare with IPX alone.

3.16.1 Replacing IPX with IP

Novell allows IP to be used as the network protocol instead of IPX. In this case, IPX packets are still retained but are encapsulated by an IP header, with all NetWare servers and clients using IP. Novell offers NetWare IP specifically as a replacement for IPX.

NetWare stores the configuration files for its TCP/IP capabilities in the SYS: ETC directory. The files are unstructured ASCII format and can be edited with any ASCII editor. There are four important files:

- HOSTS
- NETWORKS
- PROTOCOL
- SERVICES

The HOST file contains a list of IP addresses and machine names (including any aliases) that users can connect to. The NETWORKS configuration file is similar to the HOSTS file except dedicated to describing network masks (the network identify in the IP address). The PROTOCOL file contains parameters about the protocols known to the network. The SERVICE file lists all the IP services available. This TCP/IP basics are explained in more detail in the Redbook *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602.

3.16.2 Workstation Support for TCP/IP

The Workstation-based implementation of TCP/IP uses a workstation with a TCP/IP stack (for example, a UNIX or OS/2 Workstation), to load and support IP locally. All communications to the workstation are through IP. Figure 45 on page 107 shows such an configuration. The UNIX workstation cannot communicate directly with the NetWare server, which uses only IPX/SPX, but the workstation can go through the PC, which runs both TCP/IP and IPX/SPX, and then to the server if the PC can perform a conversation.

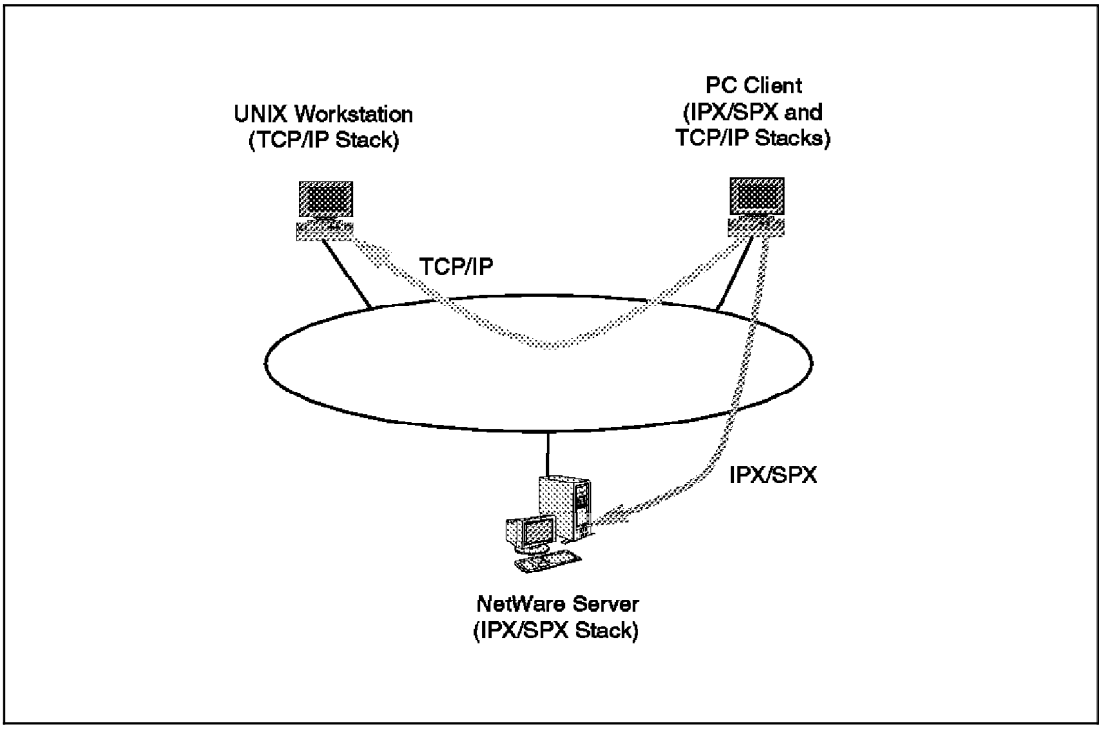

Figure 45. IPX/SPX and TCP/IP Stacks on one Machine

Each PC that wants to communicate with the workstation must have IP as well as the IPX/SPX protocol stacks for communicating with the NetWare server. This approach to implementing TCP/IP and IPX/SPX is useful when only a few machines must use IP yet maintain ready access to other machines on the network. A disadvantage is that one or more machines must have dual protocol stacks with the overhead that entails.

To operate between IP and IPX, you must install software that can handle both protocols simultaneously. Most of this software is third-party although Novell offers the LAN WorkPlace software product to perform this function.

3.16.3 Using TCP/IP Gateways

Using a dedicated TCP/IP gateway allows the IP protocol stack at each workstation to be eliminated by depending on the gateway to perform conversations between IP and IPX/SPX. An example is given in Figure 46 on page 108.

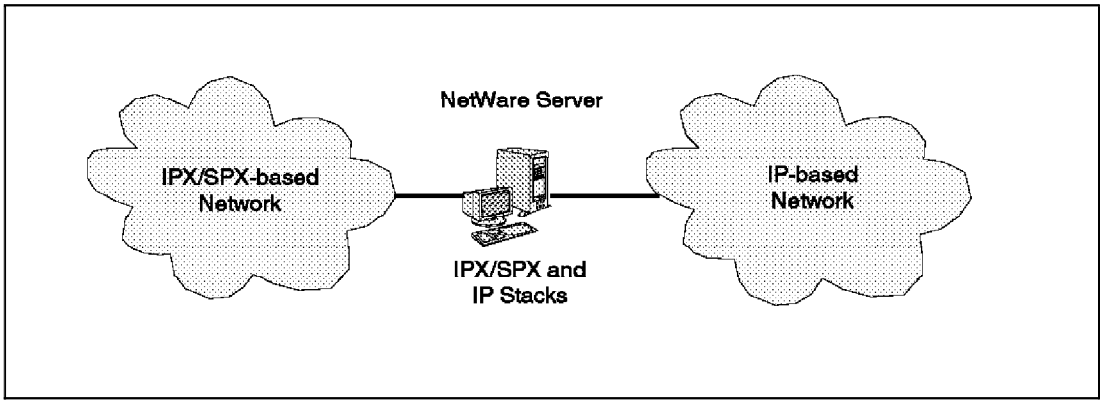

Figure 46. Different Network Protocols Connected via Gateway

This approach is useful when a mix of machines must be supported with only occasional cross-network traffic. The disadvantage of this approach is that traffic between the networks increases to the point that the gateway becomes a bottleneck.

3.16.4 Tunneling IPX within IP

In some instances, IP must be supported without IPX coexisting. To provide IPX support across an internetwork, the IPX packets must be encapsulated into IP packets, a process called *tunneling*. At the receiving device, the IP packet header is stripped off and the IPX packet rerouted into the destination network.

An example is given in Figure 47 on page 109. The gateways to each NetWare network have IP and IPX protocol stacks and perform the encapsulation of IPX packets into IP packets.

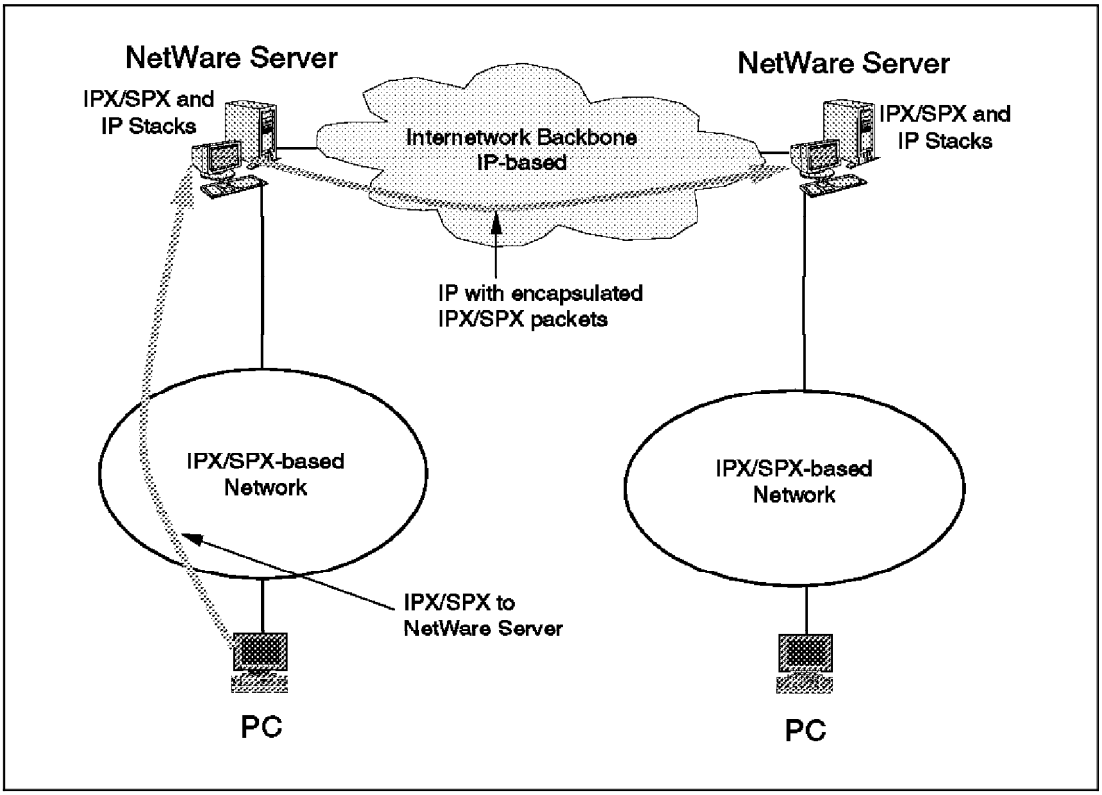

Figure 47. Tunneling IPX within IP

This approach allows a mixture of network protocols to be used in a large internetwork. Dedicated network gateways are not required unless there will be a lot of traffic.

IP tunneling relies on a gateway into each NetWare IPX-based network that also runs the IP stack protocol. The gateway, often a NetWare file server, is called a *router* or on *IP Tunnel Peer*.

Establishing this type of network is straightforward at the NetWare servers. The steps involved using the `LOAD IPTUNNEL` command to establish the tunneling implementation of IP.

3.16.5 Other TCP/IP Solutions

In addition to the Novell solutions for TCP/IP support described in the preceding sections, a number of third-party companies provide TCP/IP solutions for NetWare. These solutions are divided into two categories: NetWare server-based solutions (NLMs) and dedicated gateway solutions.

The NetWare server solutions consists of NLMs that run on the NetWare server to provide TCP/IP gateway functionally. The dedicated gateway

solutions require a workstation dedicated to providing gateway functionally. Some also require an OS/2 or UNIX operating system.

These products allow workstations connected to the IPX network to communicate with the gateway (dedicated or server-based NLMs). The gateway then converts the request into an IP transmission and sends it to the destination. Inbound communications (to the workstation) are addressed to the gateway, which then translates the packet into IPX format for transmission to the workstation. Thus, a TCP/IP host appears to be communicating with another IP host, and all translation is handled by the gateway.

3.16.6 NetWare/IP 2.2

NetWare/IP 2.2 provides both DHCP and DNS servers. The DHCP server is configured by a NetWare Loadable Module (NLM) that allows IP address-range creation and management. It also lets you reserve DHCP and BootP leases. The interface is NLM-like and offers little in the way of extended functionality. In some cases, such as exclusion ranges, the configuration tool is quite limited.

The DNS server can be used as a name-resolution server for any client capable of DNS lookups. The `Unicorn` administration utility provides easy configuration and management of the DNS server. However, the DNS server is quite static — it cannot perform any form of dynamic DNS updating. This simply means that using the DNS server for anything other than centralizing your DNS server with your DHCP server is unrealistic.

NetWare/IP is now part of IntranetWare, Novell's latest version of NetWare that includes Internet and intranet capabilities. The Internet technologies in IntranetWare do not require NetWare/IP, but it allows users to access NetWare-specific services such as the NetWare Core Protocol. NetWare/IP is backward-compatible with NetWare-based applications, including IPX applications.

For those of you Netware 4.1 users who do not have NetWare/IP, it is available for free from Novell's Web site at the following URL:

`http://www.novell.com/corp/offices/san_fran.us/download.html`

Chapter 4. Systems Management

Wouldn't it a good idea to have a systems management solution that would entail:

- An open, highly scalable, cross-platform management environment that imparted new business benefits to customers while allowing them to leverage their already massive IT investments?

As organizations move to distributed client/server systems, one of the biggest challenges is managing those systems. Large businesses need to maintain thousands of desktop PCs in numerous locations worldwide. Systems management is a vast topic and this chapter is intended to provide an overview of the systems management disciplines together with a few tools provided on the different platforms.

4.1 Meaning of Systems Management

Systems management covers the management of systems, networks, hardware, data, applications, transactions, voice , and so on. These are all resources that need to be managed. The result of the systems management processes is that it increases the availability of information system services to users in the network.

Organizations are relying to a greater extent than ever before on local area networks (LANs) to run their critical business applications. As a result systems management has become a key customer requirement for keeping the network and the workstations working efficiently.

Systems and network management has long been the Achilles heel for organizations attempting to move to client/server. A recent report from Forrester Research stated: "Companies have adequate tools to manage corporate servers, databases and network devices. But managing the client/server applications that run the business is a nightmare."

The issues facing IT departments are complex and recognized as:

- Heterogeneity

  Managing multi-vendor and multi-platform environments

- Scalability

  Everything from LANs to global networks with hundreds, thousands, or tens of thousands of users

- Standardization

Setting consistent policies for managing resources in the network and processes, such as operations administration and data, software, and resource deployment

- Integration

  The ability to view all of the elements of the infrastructure-networks, systems, middleware, databases, and applications, regardless of where they are located.

IT wants to select "best-of-breed" products; mix and match with legacy equipment and to simply integrate emerging technologies such as the Internet.

4.2 Hidden Costs of Managing a Network

Systems management is already a major focus for many businesses to reduce the cost of network ownership. Forrester Research, Inc. has estimated that out of the total cost of 1270 dollars spent per user annually to support a 5,000 user network, 750 dollars (or almost 60%) is spent on LAN administration. Four hardware and software costs pale in comparison to the five year personnel costs of managing a network.

4.3 Introducing TME 10

TME 10 is an integrated comprehensive multi-platform product set that addresses the five disciplines of systems and network management, permitting IT professionals to get a common view of essential management applications across all enterprise resources:

- Systems
- Databases
- Networks
- Internet connections
- Applications

The five disciplines include:

- Deployment management

  Speeding the delivery of information resources in conjunction with other products; configuration and change management, inventory and asset management, and software distribution.

- Availability management

Maintaining mission-critical service levels through proactive analysis of the entire computing environment; centralized system monitoring, automated actions, and performance management.

- Security

  Ensuring user access while safeguarding corporate information assets; establishing unified user views and controls across multiple systems, databases, services, and applications.

- Operations and administration

  Maintaining operational integrity while minimizing overhead by means of automated facilities for scheduling, help desk, backup and restore, output management, and other utility functions.

- Integration

  Meshing seamlessly with third-party systems and applications management products through 10/Plus Association; providing transparent extensions and interfaces to previously disparate functions and technologies.

Within this architecture, IT can apply core Tivoli management applications, customization and extension toolkits, and third-party applications.

Tivoli has regularly released new platform versions of TME, including TME for Windows NT. Today, Tivoli's framework and applications support more than 20 platforms, including 14 types of UNIX, Windows, Windows NT, OS/2, OS/400, NetWare and OS/390.

Now back to the question about having a systems management solution that would be open, highly scalable, and cross-platform. Isn't it that exactly what IBM's strategy Tivoli is offering today? - In the practical part of this book, we will briefly introduce systems management functions that are delivered with the different network operating systems discussed in this book.

Chapter 5. Remote Access

There are basically three types of remote LAN technologies in the market place:

1. Remote Client
2. Remote Control
3. Remote Node

Remote client is the addition of dial-in capability to a client/server application. Lotus Notes, Lotus cc:Mail, or Microsoft NT Server are good examples of remote client support.

Remote control is typically thought of as screen mapping. There is a host PC connected to your local area network that sends screens to a remote PC. The remote PC displays exactly what is on the host PC. You are able to control any application on the host PC from your remote keyboard. PC Anywhere, Carbon Copy, Reach Out, IBM Distributed Console Access Facility (DCAF), and Tivoli TME 10 NetFinity Remote Workstation Control (RWC) are examples of remote-control technology.

Remote node technology essentially extends your LAN connection across the wide area network to a remote PC. Applications running on the remote PC "think" they are connected to the Ethernet or token-ring LAN, when in reality, all information is running through one of the communication (COM) ports.

All three technologies have their place in a LAN design. For example, Warp Server exploits all three technologies where appropriate. The remote client and remote control support is shipped with Warp Server's systems management. Remote node comes with Warp Server's Remote Connection Services. System management's remote client capabilities allows access directly by a dial-up connection.

Server's remote access service utilizes remote node technology and is complimentary to the other two technologies, adding function and versatility.

The remote access service is a software-only solution. It will run over a standard communication (COM) port on a PC. If you are running any dial out software today, you can utilize that same port and modem to provide dial-in access to your LAN.

For example, Warp Server remote access has two components:

1. A remote node that essentially turns your remote communications adapter (COMx) into a logical LAN Adapter
2. A Connection Server that allows access to the network for one to 128 concurrent remote clients (256 concurrent remote clients in Windows NT).

The Connection Server, besides providing connectivity, enforces security, provides audit trails, and LAN-to-LAN bridging connectivity.

Supported LAN topologies are:

- Ethernet
- Token-Ring

For wide area networks, the following connections are supported:

- ISDN
- X.25 (ISO standard for interface to packet switched communications services)
- Computer-Controlled Branch Exchange (CBX)
- Asynchronous and synchronous

All major protocols are supported. This includes:

- NetBEUI with NetBIOS
- IEEE 802.2
- TCP/IP
- NDIS
- NDIS/ODI

Novell's IPX is written to ODI; NT Server utilizes NetBEUI. TCP/IP and most SNA communications applications are written to 802.2. The result is that most applications run over this remote access solution.

Warp Server's remote access is supported by Star Gate, Eicon, Cubix, and several other vendors.

5.1 Principle of Remote Access

A simplified view of a remote workstation attaching to a LAN is shown in Figure 48 on page 117. The remote workstation, as its name suggests, is a workstation that is not local to the LAN. When the remote workstation links

into the LAN, it forms a Wide Area Network (WAN). The link between the remote workstation and the OS/2 Warp Server is known as the WAN link.

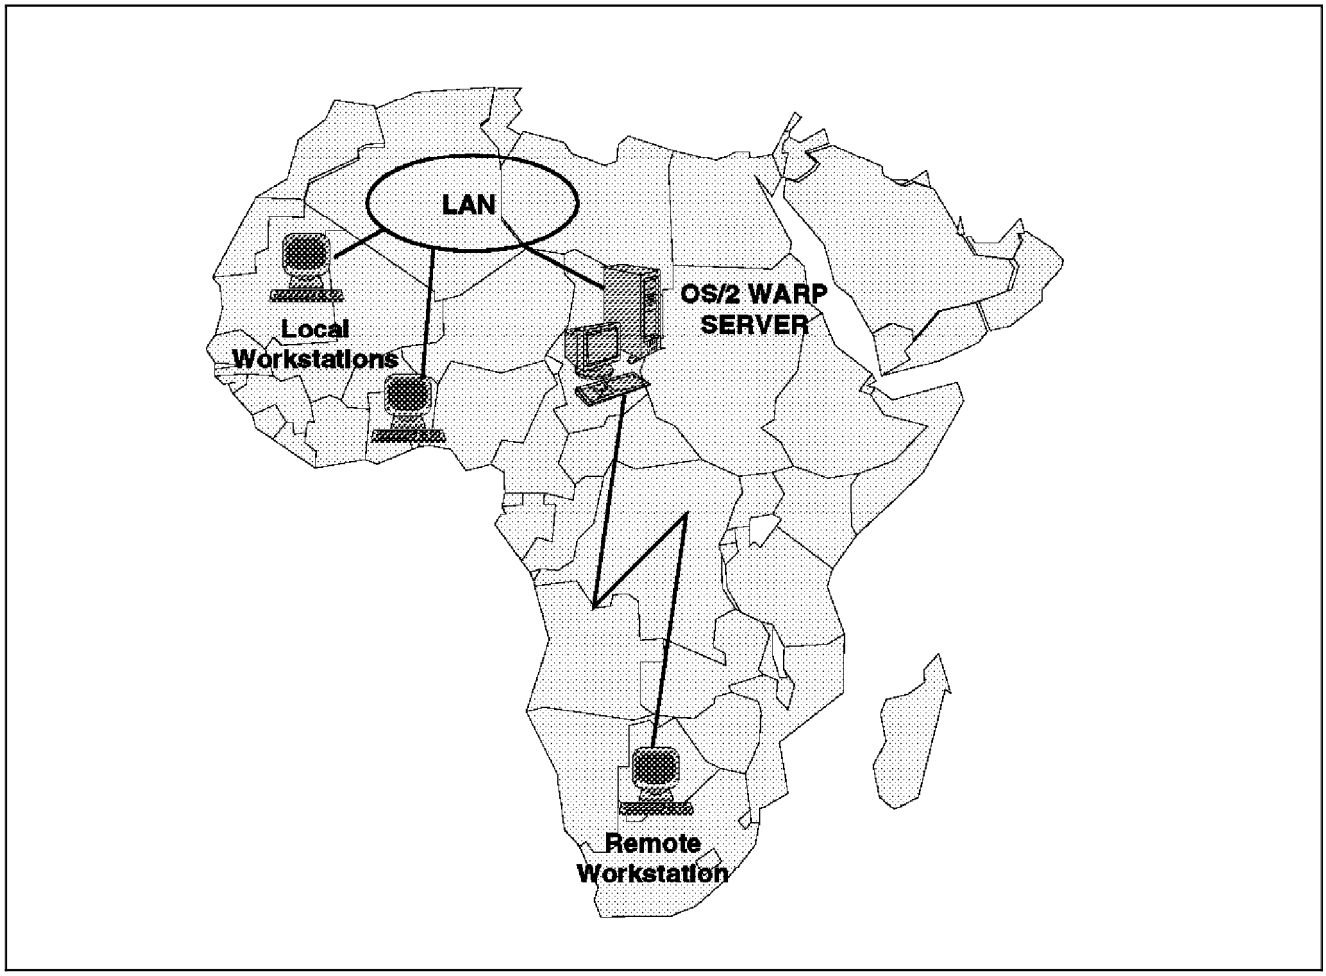

Figure 48. Remote Access Services Overview

In this case, all that the remote workstation require is a communications (COM) port and modem. The Remote Access Services must be connected to both the LAN, via a LAN adapter, and to the remote workstation via the communications link.

Remote workstations can access the Remote Access Services via a number of communications methods, including asynchronously and synchronously over switched and non-switched telephone lines and ISDN Basic-Rate switched connections. Remote workstations can connect to token-ring LANs and Ethernet LANs. The Remote Access Services also supports access to X.25 networks through asynchronous modems with X.25 Packet Assembler Disassembler (PAD) capabilities.

5.2 Remote Access Services Environments

The Remote Access Services supports the four types of remote LAN access environments described in Figure 48 on page 117:

- **Remote-to-LAN or LAN-to-Remote**

  Probably the most common use of the Remote Access Services is to provide this type of access. The Remote-to-LAN environment is a flexible solution for users requiring access to resources from remote locations, such as home or while traveling. Users can dial the Remote Access Services on their office LAN and run the same applications remotely that they use in the office.

  Alternatively, LAN-attached workstations can request the Remote Access Services server to establish a connection with a remote OS/2 workstation. This could possibly be used when someone in a central office needs to send an updated file to a number of remote workstations. They could dial out through the Remote Access Services and copy the update to each remote workstation in turn.

- **Remote-to-Remote**

  Two remote Remote Access Services clients can establish a WAN connection, as shown in the Remote-to-Remote environment, to form a virtual LAN. The Remote-to-Remote environment is a simple, low-cost solution for stand-alone workstations that require direct access to resources on other stand-alone workstations. For example, the Remote-to-Remote environment can be used in a local office environment in lieu of expensive LAN cabling or by traveling employees who need access to their office workstations.

- **LAN-to-LAN**

  You can establish a connection between two Remote Access Services to form a casual bridge between two LANs as shown in LAN-to-LAN Environments. LAN workstations on LAN A can use the Remote Access Services connection to access LAN resources on LAN B as if they were physically attached to LAN A. Similarly, LAN workstations on LAN B can access resources on LAN A.

- **Remote-to-Central Server (No LAN)**

  The Remote Access Services can be installed as a stand-alone server to support remote access workstations. No LAN hardware is necessary on the server, and the remote workstations can access all the resources at the Connection Server.

Find in Figure 49 on page 119 a graphical representation of the different Remote Access Services environments.

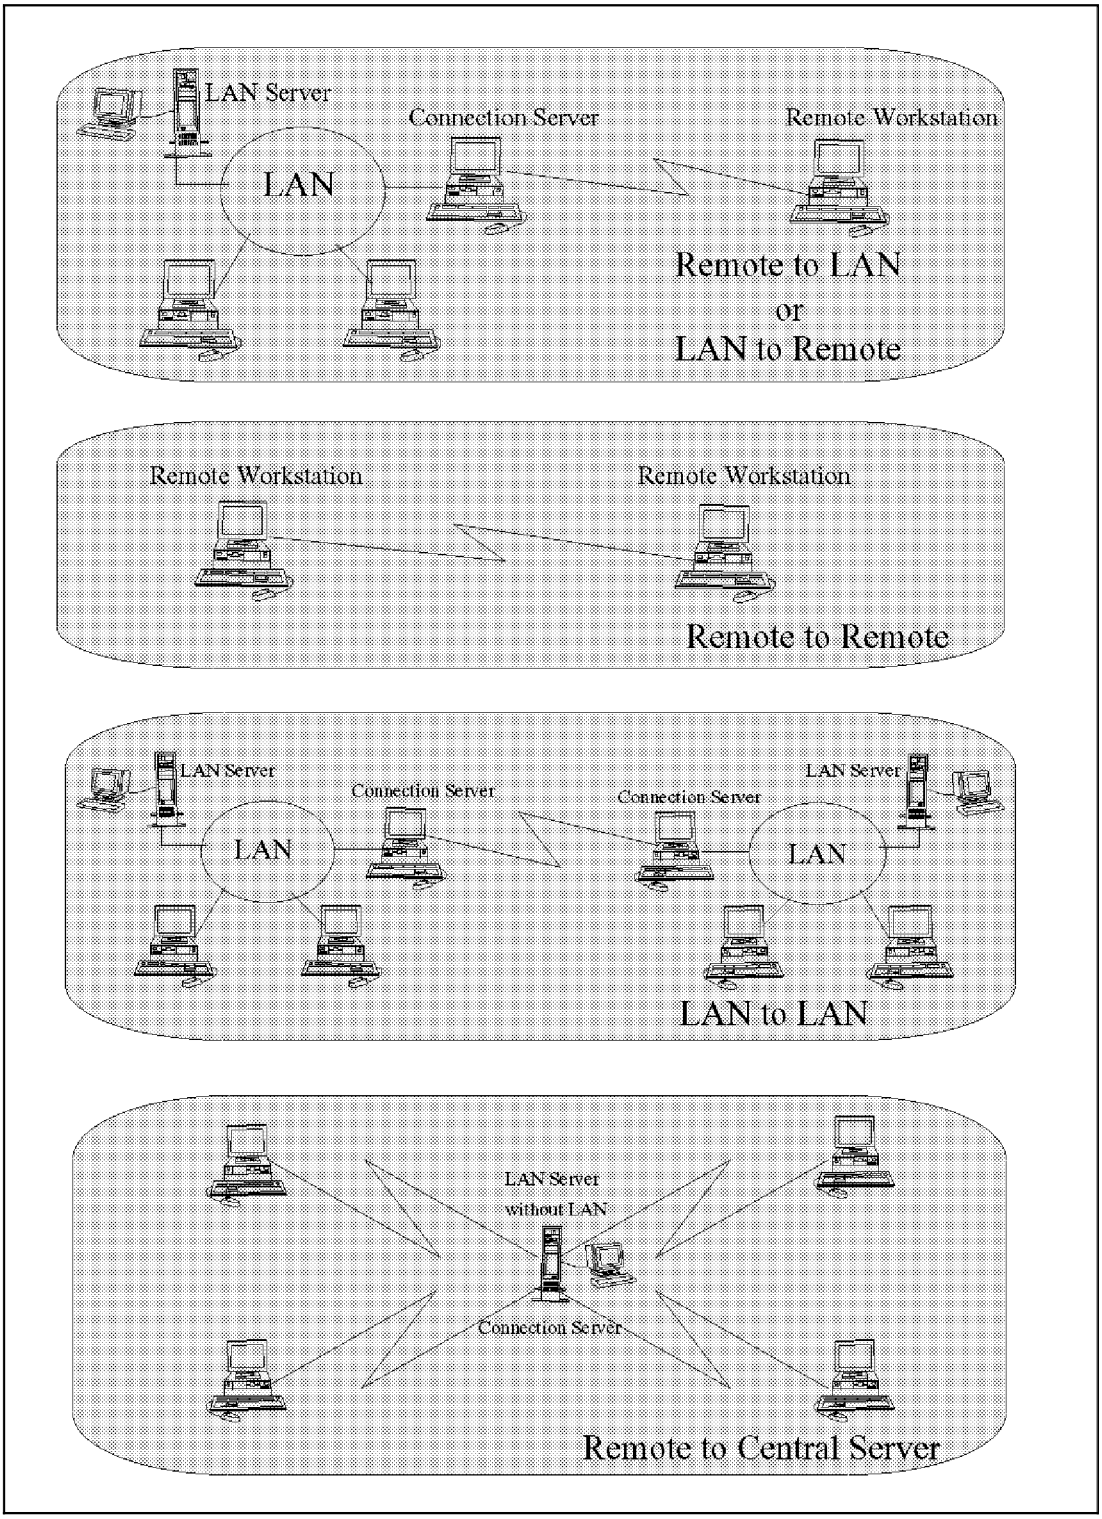

Figure 49. Remote Access Services Configurations

5.3 RAS Physical Connection Options

You may connect the RAS client to the RAS server over any of the following:

Phone Lines: Most telephone lines are analog lines, and a modem is required at both the client and the server to convert the digital signals to analog signals used by the phone lines.

ISDN: ISDN or Integrated Services Digital Network is a line installed by a phone company at the customer site. An ISDN card is needed at both ends of the communication, which can be a problem as the place you are calling may not have ISDN service. ISDN was designed from the beginning for computer use as it is digital and provides extremely high throughput. Several levels of service are available and multi-megabit speeds are possible. Two such levels of service are Basic Rate Interface (BRI) and Primary Rate Interface (PRI).

The BRI provides two digitized channels for user information at 64 Kbps and one data channel for signal control purposes. The 64 Kbps channel is referred to as the B (Bearer) channel, and the data signal channel is called the D channel. The D channel operates at 16 Kbps. With PRI, you can transmit up to 128 Kbps.

The PRI provides twenty-three B channels at 64 Kbps and one D channel operating at 64 Kbps. With PRI, you can transmit up to 1.544 Mbps. These larger data rates makes ISDN very attractive.

X.25: X.25 is an older, reliable, better established, and widely available wide area network service that is based on packet switching. In most cases, companies lease the use of a commercial X.25 network and gain access through a dedicated phone line or a dial-up port.

Computers can access an X.25 network in two ways:

- Through a modem using a device called a packet assembler/disassembler (PAD) which is provided by the X.25 provider. The PAD performs the task of formatting user messages into packets that can travel through the network and then reassembling messages from the received packets. The connections to the PAD can be dial-up connections that are established only for the duration of the call. This makes it suitable for occasional access to the remote computers. If permanent access is required, a direct X.25 connection is supported by X.25 PADs. This, however, is more costly.

- Through a direct connection using a device called an X.25 smart card. These X.25 smart cards act as modem cards and have a PAD embedded in the card itself.

RS-232C Null Modem: A null modem is a cable that crosses wires to enable two RS-232 ports to connect as though there were modems between them. If a network attachment is not available between two computers, you can physically connect them by using an RS-232 null modem cable.

5.4 Security Options in Remote Access Services

Whenever discussing remote access to a network, security should be the first topic of discussion. For example, Warp Server's Remote Connection Server access security is excellent. It uses the DES encryption model when transmitting login information to ensure password security. Instead of a single password, passphrases are supported. A passphrase can contain from four to 32 mixed-case characters including blanks. A valid passphrase is: My name is Skywalker.

The security administrator can determine key characteristics of the passphrase such as preventing duplicates, aging limits, and number of login attempts. Once the number of login attempts is exceeded, the user ID is disabled. Only the security administrator can re-enable the ID. It should be noted that once a login session has been established, there is no data encryption. Therefore, unless the application encrypts passwords, further login information is flowing in the clear. It is highly recommended that the remote access passphrase be different from other passwords used on the system. By setting the minimum passphrase length, a security administrator can enforce this policy.

Warp Server also supports other security features. With fixed callback, once a login occurs, the system will disconnect and dial the user at a predefined phone number. For home users, this is an excellent method to tighten security. For users with high access rights, such as administrators or security officers, this is an excellent means of verifying their identity. A mobile callback can also be performed. During login, the user is prompted for a phone number for callback. The mobile callback can prove useful if reversal of telephone charges is needed.

The security administrator can also specify up to eight LAN logical adapter network addresses per user account. For instance, a user may have a terminal at home and a laptop for travel. In this instance when a user login occurs, a check will be done to ensure that the user ID matches the appropriate workstation address before allowing access to the network. If no logical adapter addresses are specified, a user can log in from any

remote access client. Network access for users can be restricted to specified hours of the day or days of the week. Access during early morning hours can be restricted when backups need to take place or when no one is available to monitor dial-in activity.

Chapter 6. Backup and Restore

Network administrators have a very pragmatic approach to openness. To them, the definition of openness is: It works with what I have. By this definition, industry backup and recovery applications are not very open. Few vendors consider the impact of their individual backup strategies not fitting into an overall corporate backup strategy. Because of this lack of cooperation, many companies are forced to implement two different backup strategies on their networks.

Mainframes and mini's still monopolize the high-speed tape backup capability, while network servers are left to their own means. This results in a patchwork quilt backup strategy. Sometimes two strategies are implemented: one strategy for mainframes and a second for network servers.

6.1 Backup and Recovery Strategies

The following scenarios demonstrate how a company can grow through four backup and recovery strategies. Start with a server in a small stand-alone environment and grow to a large corporate, heterogeneous environment.

- Stand-alone backup strategy

  This is usually a small stand-alone network with a single server providing most of the services on the network. In this case, a server stand-alone backup strategy simply writes the data to tape or to a second disk for safekeeping. This can be done during off-shift hours with scheduling features if provided by the backup and restore software. As the network grows, the customer may choose to implement separate stand-alone backup strategies for each server.

- Cloning backup strategy

  As a company grows, it keeps doing what works. Servers are added to the network. It is easiest to order servers with tape drives and clone the existing strategy to these new servers. Quick, easy, and it works.

- Push backup strategy

  As the company expands, the time comes when having separate, stand-alone server backup strategies becomes unmanageable. The administrator may find it impossible to "make the rounds" of all the servers on multiple floors or remote sites. The cost of tape drives becomes prohibitive. The company's high server availability requirements may force a small backup window. A "push" backup strategy can then be implemented. A centralized server can be a

dedicated backup server. Each remote server "pushes" its data quickly to this dedicated server. This dedicated server is backed up while the other servers go back into service. As for Warp Server, all this is can be accomplished without utilizing precious LAN drive letters, as you would do in a Windows NT environment.

- Pull backup strategy

  The company has grown significantly. Instead of a "push" strategy, there is a need for a full-time backup specialist. This specialist must be able to establish and control a central server backup strategy and "pull" data centrally if necessary. The ADSM API set integrates Warp Server into this environment. Centralized control provides the highest productivity, reliability, and speed.

In the practical part of the book, we introduce each platform's backup and restore solution. However, Novell does not come with a backup/restore solution. Third-party products have to be installed in order to obtain this function in NetWare 4.1.

6.2 Comparison Backup/Restore Features

Find in Table 12 a comparison about the backup/restore features of each network operating system.

Table 12 (Page 1 of 2). Backup/Restore Comparison Features

Operating Systems	**OS/2 Warp Server**	**Windows NT**	**Novell 4.1**
Built-in feature	Yes	Yes	No - Third Party
Devices	SCSI-I and SCSI-II devices (such as tape drives and optical drives), diskette, local hard disk, LAN drives, ADSM	Tape with SCSI interface, depending on hardware compatibility	
GUI Interface	Yes	Yes	
Drag and Drop	Yes	No	
Disaster Recovery	Yes	Yes	

Table 12 (Page 2 of 2). Backup/Restore Comparison Features			
Operating Systems	**OS/2 Warp Server**	**Windows NT**	**Novell 4.1**
Time Setting for Backup	Yes	Yes	
Sound Feature	Yes	No	
Backup Locked Files	Yes	No	

Notes:

1. It works creating a set of three bootable diskettes that will then give you the possibility to restore an entire Backup Set to recover your system.
2. You have to set up your Emergency Disk Repair, inserting three diskettes of the operating system setup. In case of changes from the set up of the diskettes to your last successful boot, it could be that the Emergency Disk Repair will fail.
3. If your machine has a sound card, you can play audio files. It is possible to have sound features for Starting Automatic Backup, Starting Manual Backup, Starting Restore Feature, Successful Backup, Unsuccessful Backup, Successful Restore, or in case of need to change tape device, tape cartridge, or WORM (Write Once Read Multiple/Many).
4. The OS/2 Warp Server feature to back up locked files will be available with the first Warp Server Service Pack or Warp Server SMP.

Part 2. Common Information

Chapter 7. Logon Interoperability: File and Print

In this section, only interoperability between Windows NT Workstation/Server and Warp Server will be discussed. Both IBM's Warp Server and clients and Microsoft's Windows NT Server and clients are based on the SMB (Server Message Block) protocol and therefore use NetBIOS, either native or over TCP/IP, as login flow. No native TCP/IP communications occurs here.

On the other hand, Novell Netware does not use the NetBIOS APIs for file and print communications and therefore is not discussed here as far as logon interoperability is concerned. There are Netware clients available on both platforms that address the login issue.

7.1 Logon Interoperability: Windows NT to Warp Server

What happens to customers who are attempting to roll-out Windows NT Workstation and Warp Server together? Many people have been chasing this particular problem, and below there are a couple of notes that explain the whole scenario. Even though there is more technical information available, the result is the same: NT Workstation will not log on to Warp Server or any other SMB server. Ironically, NT Workstations cannot even log on to Microsoft's own LAN Manager.

Microsoft has intentionally not published their logon API flows; their strategy is to provide a proprietary solution. They don't want their NT clients to log on to IBM servers (AIX, OS/400, and so forth) or other SMB servers. This does not fit their strategy.

We could speculate about why and how, and if we will get a resolution from Microsoft, but this would not help at the moment. However, customers do not want to be locked in, and that is exactly what is happening here.

We have discovered that the NT Workstation default log-on procedure, which is used to perform basic server access functions, uses encrypted, proprietary flows between the NT Workstation client and the NT Server. No AIX server, OS/400 server, or Warp Server can match this flow. Only NT Server can match these flows. Therefore NT Workstations can only log on to NT Servers for basic functions, such as security, file and print, remote access, and so forth.

7.2 Logon Flow of a Windows NT Workstation

1. Basic NT Workstation client support presumes connection to an SMB server for file and print capability and for base client support such as security. LS/400, AIX Connect, and Warp Server all provide this level of SMB server support.
2. When a user "logs on" with NT Workstation, the default access that occurs is to find and log on to an SMB server. (Or to a collection of SMB servers called a "domain.")
3. Unfortunately, what should work does not work. In all customer situations we have found so far, and in our test labs, when an NT Workstation attempts to log on to any of IBM's SMB servers, (for example, Warp Server, LS/400, LS/X, AIX Connection), the logon fails. We know the NT Workstation can, indeed, "find" our SMB servers and complete the actual log on, but special processing by NT Workstation occurs after log on that stops NT Workstation from working with our servers.
4. Our tests have shown what is happening is that the Windows NT client uses an encrypted, proprietary FAP (Formats and Protocol) to further identify itself to the NT Server after logon.
5. After the initial logon sequence, NT Workstation executes these proprietary flows, which are only supported by NT Server. The NT Workstation then erroneously sends an error message to the user stating it cannot find a server or domain to log on to, and the logon fails.

7.2.1 Workaround for Logon Interoperability

It is possible to "fool" NT Workstation to avoid the failure by having the Microsoft logon sequence only apply to the local workstation. That is what IBM Software Servers do to support Windows NT clients on AIX and OS/2. However, that does cripple some of Windows NT Workstations' basic client to server security.

Our investigations have uncovered that this logon failure occurs with all other servers, including NetWare, HP, and even Microsoft's previous LAN Manager servers. We suspect it will fail with other server types as well, including Banyan Vines, Pathworks, and others.

7.3 Introducing Warp Server Windows NT Client

In the first quarter IBM has planned to introduce a native Windows NT client that will enable an NT client to log on to a Warp Server and to gain security (although it is not using the Microsoft undocumented flows). All typical

Warp Server file and print functions will work, such as logon assignments, the use of alias resource names, and so forth.

Chapter 8. Security Issues

One of the most important features of the safe sharing of network resources is security. On a stand-alone machine, where you can add, delete, or change files and directories at your own will, you have complete access to all disk contents because the disk is exclusively yours. On the network the situation is much different; you can share access to the files and directories on network disks, but sharing must be accomplished by some well-defined rules. If this is not the case, remember that your files or someone else's files can be deleted or changed, accidentally or intentionally.

Without any kind of control, a network-based accounting system is subject to deliberate tampering or outright destruction. In this chapter we will differentiate between local and network security among OS/2 Warp Server, Windows NT 4.0, and NetWare 4.1. As for Warp Server we will discuss DSS also; for Windows NT we will go into C2, and for NetWare we will talk about NDS.

8.1 Warp Server Security Services

The OS/2 Warp Server file and print security model is built upon the domain concept to provide a Single System Image for user logon and resource access. Generally, a Warp Server network may consist of one or more domains depending on the size and number of workgroups that exist in an organization. These domains are independent of each other and as such are separately administered and controlled.

The key element of Warp Server's security system is the User and Group Accounts Database, the NET.ACC file, which holds all the user and group definitions for a single Warp Server domain. The master copy of NET.ACC is maintained on the server designated as the Primary Domain Controller (PDC), and user and group information from NET.ACC is propagated to all additional servers and Backup Domain Controllers (BDC) that exist in the domain.

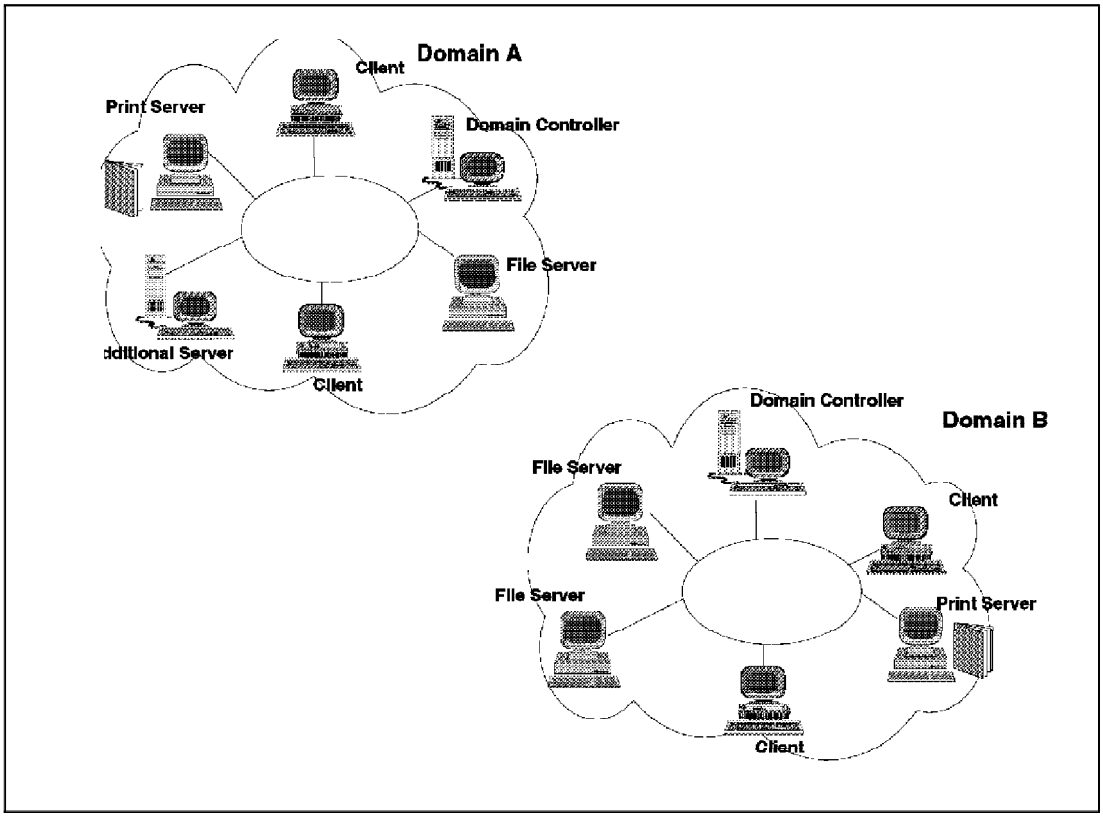

Figure 50. LAN Server Workgroup Environment

For migration and interoperability purposes, the internal format of NET.ACC will be described. This will help in understanding the structure and use of NET.ACC as Warp Server's security database and also assist in planning for the migration of a Warp Server environment to DSS. The NET.ACC file is broken into three main sections:

- Header Information
  - Also referred to as the modals containing global settings; most of which can be viewed with the `NET ACCOUNTS` command. Parameters listed in this section include the integrity information, update counter on Domain Controller (DC) and additional server (for NET.ACC synchronization), minimum password length, computer name, and server role.
- User/Group Definitions
  - List of Group names, user IDs and names, encrypted passwords, group membership and statistics counters.
- Access Control Profiles
  - For non-HPFS386 resources, for example, serial devices, printers, named pipes and FAT drives. This section consists of resource

names and a list of users and their access rights (Read, Write, Create, Delete, Execute, Attribute, and Permission). This information is server-specific, and as such it is not replicated to other servers in the domain. HPFS386 drives store their access control profiles as an integral part of their file system.

One of the main functions of the NETLOGON service in Warp Server is as an announcer, notifying additional servers of NET.ACC changes to enable them to submit update requests to the domain controller. To protect against unauthenticated intruders, all member servers in the domain belong to a group called SERVERS, with each server having a unique ID and an associated password defined. This internal password is exchanged in authenticating sessions between member servers and their domain controller in order to guard against intruders. The NET.ACC synchronization process is essential to the integrity of LAN Server's security system as NET.ACC is used for the session authentication at the initial user logon and for subsequent authorization when accessing network resources. However, as NET.ACC is unique to a domain, cross-domain operations require additional administration to enable a user access to resources outside of their logon domain. This results in having to define the user in each external domain where resource access is required.

A domain logon involves a data flow at various transport levels:

1. After the `Session_Initialized` frame and the `Session_Confirm` commands are completed at the NetBIOS layer, the requester sends the `Negotiate` SMB command. In this frame, the requests tells the server the SMB dialects that it can accept. Although there are several dialects listed, there are two main categories: user-level security with password encryption (which requires a valid user ID/password combination) and share-level security.
2. The domain controller handling the logon will respond to the requester by selecting the SMB dialect which offers the highest security mode, for example LANMAN2.1. The domain controller also sends the client a random encryption key.
3. The LAN Requester logon program uses an encryption algorithm to encrypt the password entered by the user. The resulting encrypted password should match the one stored in NET.ACC on the domain controller.
4. The LAN Requester logon program then encrypts the already encrypted user password with the random encryption key supplied by the domain controller. This doubly encrypted password is called a "proof".

5. The LAN Requester logon program then sends the proof to the domain controller.
6. The LAN Server domain controller retrieves the encrypted user password stored in NET.ACC and encrypts it with the same random key it sent to the requester. Then this doubly encrypted password is compared to the "proof" received from the requester. If they match, the user is authenticated.
7. Since the domain controller has authenticated the user, it sends a response to the requester indicating that session setup is complete.
8. After the session has been set up, the DCDB information is processed and a to give the user their logon environment (logon assignments, public application, and so forth).
9. Once the user is given their working LAN environment and they are ready to access servers in the domain, the login authentication session with the domain controller is terminated.

   Note: Setting up subsequent authenticated sessions with resource servers is similar to the initial domain logon. The requester will negotiate the SMB protocol to use with each additional server it needs to connect to. In turn additional servers supply the requester with the random number used to compute the doubly encrypted password (proof). Once the proof is verified as valid, the session setup completes and the network request (for example NET USE) is handled. Since the passwords are synchronized between servers in a domain, this logon succeeds most probably.

8.1.1 Warp Server Access Control Model

To understand the difference between DSS's authorization mechanism and that of Warp Server's, the basics of the access control structure in Warp Server is first described.

Warp Server supports what is referred to as a sparse ACL model, where ACLs are not defined for each object. Instead, ACLs exist only on the physical resources and not on object definitions. An example of an object is a public application definition or a user or group account, whereas a physical resource could be a printer or a file directory. Authorization in cases where ACLs do not exist is assumed or determined by an account's administrative authority or operator rights. For example, an administrator will by default have total control over the domain, which means he/she can create, delete, and modify all account, group, resource, and application definitions. However, an account with print operator privileges can create, delete, and change only print resource definitions.

This model is not limited to object definitions; it also extends to the physical resources defined in a domain. This effectively means that an account with administrator privileges will automatically have total rights over all domain resources regardless of whether they are defined in the Access Control List for the resource alias or not. In LAN Server, there are two main account types that can be defined, administrator and user. In addition, an account with type of "user" can be granted special operator privileges.

- Administrator
  - There are no limits imposed on this account regarding all aspects of domain administration relating to user/group and resource definitions.
- User
  - By default, this type of account does not have administrative authority in the domain. However, there are four operator privileges that can be granted to a "user" account to allow:
  - Print queue management
  - Group and user account management
  - Serial device resource management
  - Shared resource management

For physical resources that exist in a domain, there is an alias definition and an associated Access Control Profile (ACP) that comprises a list of entries containing a user ID or group ID and the permission. The ACP is used to either restrict or grant access to accounts requesting to use a specific resource. The LAN Server authorization mechanism proceeds in the following order:

1. An Access Control Profile is first sought for the resource the user is trying to access.
2. If no such profile is found, LAN Server will check for an Access Control Profile for the current (parent) directory.
3. If no profile is found for the current directory, the program searches for an Access Control Profile for the root directory.

8.1.2 HPFS386 ACL Behavioral Differences

In HPFS386, access control information is stored in the file system rather than in the NET.ACC, as is the case with non-HPFS partitions and resources. This results in some behavioral differences that are due to the object-oriented system design of HPFS386. One difference is creation of directories and files locally on a server disk, which results in ACL

inheritance of the parent directory. Other differences can be illustrated by the MOVE and COPY file operations.

When a file is moved, the path in the file system is modified to reflect the file's new location but there is no actual relocation of the file contents. Since no file creation takes place, the MOVE command causes the original ACL to move as well. However, with the COPY operation, the target file is created, and as such it inherits the ACL of the target directory.

HPFS386 also distinguishes between implicit and explicit Access Control Lists. An implicit ACL is usually the implied or inherited ACL of the parent directory and as such does not display with the NET ACCESS command. However, explicit ACLs can be displayed using NET ACCESS because the ACL exists for the resource and is not derived by parent inheritance. An example of a implicit ACL becoming explicit is during a file MOVE operation. This ensures that the file's ACL is that of the old parent's directory and not assumed to be that of its new target directory.

For local security, HPFS386 protects files on formatted partitions on the server regardless of whether LAN Server (or file and print services in Warp Server) is running. When a local user tries to access a file on an HPFS386 partition on a server, local security checks the permissions for the file. Access is allowed only if the user has the required permissions. Local security also checks file permissions when you run a program that accesses files. File-access auditing is also improved on a server with local security installed. Auditing of files and directories on HPFS386 partitions on the server begins when the workstation is started, even if the server service is not started.

8.1.3 Local Security

Permissions that you set for a resource usually apply only to remote users accessing the resource from different workstations. LAN Server Advanced as well as Warp Server Advanced provide local security for the HPFS386 (referred to as local security). Local security extends access restrictions to local users working at the server. It protects all files on HPFS386 partitions of the server from unauthorized local access. The files that are stored on a FAT file system partition are not protected by local security. However, they are still protected from a remote access by unauthorized users; administrators set permissions for all files on the server and are not subject to access control permissions. If you are working on the server with local security, you can access a file only if you have adequate permissions for that file. If you run a program, the program can access a file only if you have permission to access the file. Your permissions, not those of the program, control access to files. Local NetAPI calls are also subject to

privilege checking when local security is active. It is also possible by an administrator can also start a program and give it special privileges.

8.1.3.1 Starting a Server with Local Security

When you start a server on which local security is installed, many programs are started. To be sure that the operating system and programs run successfully when nobody is logged on to the server, give the special group *Local* read (R) and execute (X) permissions to the files needed for the operating system. Only administrators can start services on servers with local security. To help you set up the right permissions for system files, LAN Servers displays the following prompt when a server with local security is started:

```
LAN Server Local Security has started.

Press ESC to log on now, or press ENTER to start the workstation
with no one logged in.
```

The prompt will shortly go away, allowing the server to start automatically without intervention. This should allow OS/2 and its desktop to start with nobody logged on. When the logon prompt is displayed, log on using a user ID that has administrator privileges on the server. It is also possible to adjust the server file permissions for the group Local if you want the workstation to start with no one logged on. After you grant permissions for system files, log off so that no one else can use the administrator privilege on the server. When you log on locally to a server with local security and then you logon to the network and then you decide to log off from the network, you will only loose that connection; you are still logged on to the local server.

8.1.3.2 Accessing Files with Local Security

On servers with local security, LAN Server creates a special group called Local. No users are members of this group, and its not possible to add them either. The Local group is a special group ID whose permissions are granted to the system or to users using the local system when nobody is logged on. The Local group is not shown in the User Profile Management (UPM) window, but you can see it in the LAN Server Administration GUI.

Users working at the server workstation when no one is logged on receive the permissions granted to the Local group. If you log on with user privilege, you have the permissions granted to your own user ID and to all the group where you are a member, in addition to the permissions granted to the Local group. If you are logged on with the administrator privilege, you can access all files on the server.

Here is an example of some permissions for a user working at the console of a server with local security:

- Local user not logged on; permissions granted to the group Local but no access to the network
- Local user logged only to a local server; permissions granted to the group Local and the user ID but no access to the network
- Local user logged on to the network; permissions granted to the group Local, the user ID, and access to the network

Accounts Database

Both the user ID and the corresponding password must be in the accounts database for the server running local security before LAN Server will grant network permissions. If LAN Server does not find a valid user ID, LAN Server grants only the permissions of the Local group at the local server.

8.1.4 Remote Security

Adding remote access capabilities to your LAN can make your LAN and its resources vulnerable to unauthorized remote access. The security features provided by the Remote Access Services product control access to the connection server and help prevent LAN access by unauthorized users.

The Remote Access Services security subsystem provides two main services:

1. It protects LAN from casual, unauthorized external access. When an external WAN circuit is established, the security service checks that until the caller is authenticated:
   - No LAN frames are transferred onto a WAN circuit.
   - No WAN frames are transferred onto the LAN.

   In addition, if there are already several external users that have been authenticated and are currently accessing an application server on the LAN, the security subsystem ensures that a new caller does not see any of the traffic between the LAN and any of these other users until the new caller has been authenticated.
2. Continuous validation of remote requests: When a Connection Server receives a request for service, it can determine whether the:
   - Request was sent by an unauthorized user
   - Request received has not been modified in transmission

- Current message is not a copy of prior message

Before a remote workstation sends a request to a secured Connection Server, the user at the remote workstation must first be authenticated by the Connection Server.

8.1.4.1 Security Features

There is a configuration option, Remote Access Services security feature, that can be enabled on a remote workstation as well as on the Connection Server. This function is not available on Windows workstations (however if it is enabled on the Connection Server, then both OS/2 and the Windows requester must supply a user ID and password. If security is disabled, any person can access the configuration interface at the Connection Server and enable its security option. However, once security is enabled, only a user designated as a security administrator can log on to the secured workstation and disable the security subsystem. It is only possible to enable and disable security at a remote workstation as a local operation only; it can not be performed remotely.

Passwords: To minimize the possibility of password detection, the security database supports up to 32 case-sensitive characters that can even be used to build password phrases. The passphrase is one-way encrypted using a *halsh algorithm*; this is a special method of transforming a source key to an object key, and it is extremely difficult to derive the source key from an object key.

User Permissions Types: On each remote workstation the user account database is independently maintained. It is possible to configure the Connection Servers to operate independently or to use a shared database. This database contains information on each user such as:

- User Accounts Database
- Accounts, User
- Database, User Accounts

  This database contains information on each user, such as the user ID, password key, and user type. The three user types are:

  - User:

    This is the lowest security classification. A user of this type can also view and change selected information (for example, user description and user passphrase) within the user's own account at a secured workstation.

  - Administrator

This user type has the same privileges as a user type and is able to perform the following tasks:

  - Create and maintain dialing and answering specifications
  - Manage connections
  - Manage ports
  - Resolve errors situations

- Security Administrator

  This user type has the same privileges as an administrator and in addition is authorized to maintain the *security policy* (for example, maximum number of logon attempts permitted during a single call). A security administrator can view, add, and delete user accounts within the User Account Database. This user type can change any of the account information contained in other users' accounts.

Single Logon: A user is required to log on and be authenticated by each Connection Server before accessing the server's services. For example, a user that has been authenticated can:

- Use dialer services
- Use management services
- Access the target LAN wire

However, a user need only be involved in a single logon task (that is supplying a user ID and passphrase) provided the user has the same user ID and passphrase at each of the secured Connection Server workstations that the user subsequently attempts to access. The user ID and password key used during the first logon are saved (in memory only) by the workstation security component and used first for each of the following logon attempts at the other secured Connection Servers. The user is required to participate in a second logon only if the user ID or passphrase is different at the next secured Connection Server.

If a Connection Server has security enabled, then remote workstation users, both OS/2 and Windows, are prompted for the user ID and passphrase after they dial and establish a link with the Connection Server. If an OS/2 remote workstation has security enabled, then the user at that workstation must also log on *locally* before accessing local services, such as Settings. In addition, users at remote workstations, both OS/2 and Windows, attempting to access an OS/2 remote workstation where security has been enabled must first log on to that remote workstation, just as they would to a secured Connection Server. This additional function is not available for Windows Remote Workstation users.

If security is enabled at an OS/2 remote workstation and if the user ID and the passphrase match between the OS/2 remote workstation and the Connection Server, then the user is prompted for only one logon, the first local logon; an implicit logon occurs after a connection is established. If the user ID and the passphrase do not match between the remote workstation and the Connection Server, the user is prompted to log on again to the Connection Server after the link has been established.

After the remote logon and filtering has completed, it is the responsibility of the LAN-based applications, such as OS/2 LAN Server, to provide security for their own applications. The logon to these applications are separate from the remote logon.

8.2 Security Standards

Requirements for secure system are articulated by the U.S. Department of Defense's National Computer Security Center (NCSC) in the publication titled *Trusted Computer System Evaluation Criteria*, also known as the *Orange Book*. All systems, whether they are network operating systems or stand-alone operating systems, are evaluated under the criteria set forth in the *Orange Book*.

NCSC established ratings that span four hierarchical divisions: D, C, B, and A, in ascending order. For each division, one or more classes are defined for a total of seven classes. The ratings reflect increasing provisions of security:

- Division D - No Guaranteed Security
  - Class D1: No Security
- Division C - Discretionary Access Control
  - Class C1: Discretionary
  - Class C2: Controlled Access Protection

    Note: This is a machine (not network) based security.
- Division B - Mandatory Access Control
  - Class B1: Labeled Security Protection
  - Class B2: Structured Protection
  - Class B3: Security Domains
- Division A - Verified Model
  - Class A1: Verified Design

8.2.1 RedBook

The NCSC's Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria, or *RedBook*, is an interpretation of the *Orange Book* security requirements as they would be applied to the networking component of a secure system. The *RedBook* does not change the requirements, it only indicates how a network should operate in order to meet *Orange Book* requirements for a C2 secure system.

8.2.2 C2 Security

Some of the most important requirements of C2-level security are:

- Discretionary Access Control

  The owner of a resource (a file) must be able to control access to the resource.

- Object Reuse

  The operating system must protect data stored in memory for one process so that it isn't randomly reused by other processes. Another important feature is that when a file is deleted, users must not be able to access the files data.

- Identification and Authentication

  Each user must uniquely identify. The system should be able to track the users activities.

- Auditing

  System administrators must be able to audit security-related events and the actions of individual users. The access to the audit data must be limited to authorized administrators.

8.2.2.1 C2 Security in Windows NT Server

Windows NT is designed to be a complete, secure solution that includes the desktop, server, and network. Windows NT Server was recently posted to the NCSC's Evaluated Products List as *Orange Book* evaluated and is currently in the formal phase of the *RedBook* interpretations. After the first listing, the NCSC will continue evaluating additional components of the Windows NT operating system and add them to the list of evaluated products.

NCSC has found the core components of Windows NT to be C2 compliant, and customers can use Windows NT as a component in building their C2-certifiable systems. Windows NT server is also being evaluated as the networking component of a secure system, the *RedBook* interpretation of the *Orange Book* guidelines.

Windows NT C2 Certification

Windows NT is C2 certified. However C2 has hardware and other requirements in order to comply with C2 certification, and it should be installed and used within the C2 environment. Unless the above-mentioned environment is used, Microsoft does not guarantee a C2 secure environment. Carefully read 8.2.2.2, “C2 Implementation in Windows NT” on page 145 for a better understanding what C2 Security really means.

8.2.2.2 C2 Implementation in Windows NT

The Windows NT Server implementation is entirely software-based. So there is no need to install additional hardware on either their servers or clients to meet C2-level security requirements. Windows NT (both the server and the workstation), were designed from the ground up to be C2 secure. This means that every process and feature was designed with C2-level security.

However, C2 security means local security. As soon as the server is connected to a network, the C2 security is not applicable. A stronger security level, B2 for example, is necessary. In comparison to Windows NT, IBM's main frame operating system OS/390 is B2 certified. C2 security in a network would apply to one domain of only NT servers and NT workstations, as long as these are running NTFS. Machines running DOS, Win 3.10/3.11, Win95, OS/2, and NT systems with other file systems than NTFS are not allowed in this C2 secured domain. In addition, diskette readers must be rendered inaccessible. Other domains must be protected using passwords.

What is also important is that if an NT Server is connected in a network, it must be the master server. It can not be operated from anything else in the industry.

Building a secure network operating system required careful planning. Security features must be included throughout the system. The file system, user account directory, user authentication system, memory management, environment subsystems, and other components all require special design consideration if the system is to be secure.

User accounts are also managed centrally. The administrator can specify group memberships, logon hours, account expiration dates, and other user account parameters via easy-to-use graphical tools. The administrator can also audit all security-related events, such as user access to files, directories, printers and other resources, and logon attempts. The system can even be set to "lock out" a user after a prescribed number of failed

logon attempts. Administrators can also force password expiration and set password complexity rules so that users are forced to choose passwords that are difficult to discover. From the user perspective, Windows NT Server security is complete, yet easy to use. A simple password-based logon procedure gives users access to the appropriate network resources by typing a unique logon name and password before being allowed to access the system.

What the user does not see are processes such as the system-level encryption of their password so that it is never passed over the wire. This encryption prevents unauthorized discovery of a user's password through wire "sniffing."

Users are also able to define access rights for any resource they own. For example, if a user needs to share a specific document with other users, he or she can specify exactly who has read and write access to that document. These rights are easily assigned through the familiar Windows File Manager. Of course, access to organizational resources is fully managed only by authorized administrators.

An even deeper example of Windows NT Server security capabilities is its protection of data, even while that data is in a machine's physical memory. Windows NT Server allows only authorized programs to access data. When such a program accesses data, that data is placed in physical memory. Despite the fact that the data is no longer only on the disk, Windows NT Server still protects it from unauthorized access. No unauthorized program will be able to access that data while it is in memory. Therefore, it is impossible for a rogue application to take advantage of another application's use of data while that data is in the physical memory of a machine.

8.3 NetWare Security

NetWare 4.1 provides the tools to prevent uncontrolled sharing. Each network will have at least one user that has supervisory control, and this network administrator has the responsibility to use NetWare's security tools to create a safe network environment.

8.3.1 Controlling Logins and Passwords

Login and passwords are the very first and most basic security control tools for governing which users receive access to what. The administrator creates a logon name for each user permitted to log into the server. It could be that the users are even requested to use passwords everytime they log in. Logging in is the first step for using the resources of the

network. If you do not have a login name, you can not access any server-based files; you can not run any programs stored on the server disk volumes and you can not even send jobs to the printer. To log in you have to switch to the new drive created when you loaded the requester and type `LOGIN`. You will see a message `Enter your Login name`. Type in your login name, and in case you use a password, type it in also. At this point you are logged into the network, but you do not have full access to each server's disk content. Usually only network administrators have full access to the servers they manage. Because server disks are shared by many different users, NetWare allows network administrators to grant users various levels of access to different areas of every server disk.

- NetWare file server hard disks, are divided into directories.
- Each directory stores a group of related files.
- Directories themselves can be divided further into subdirectories.

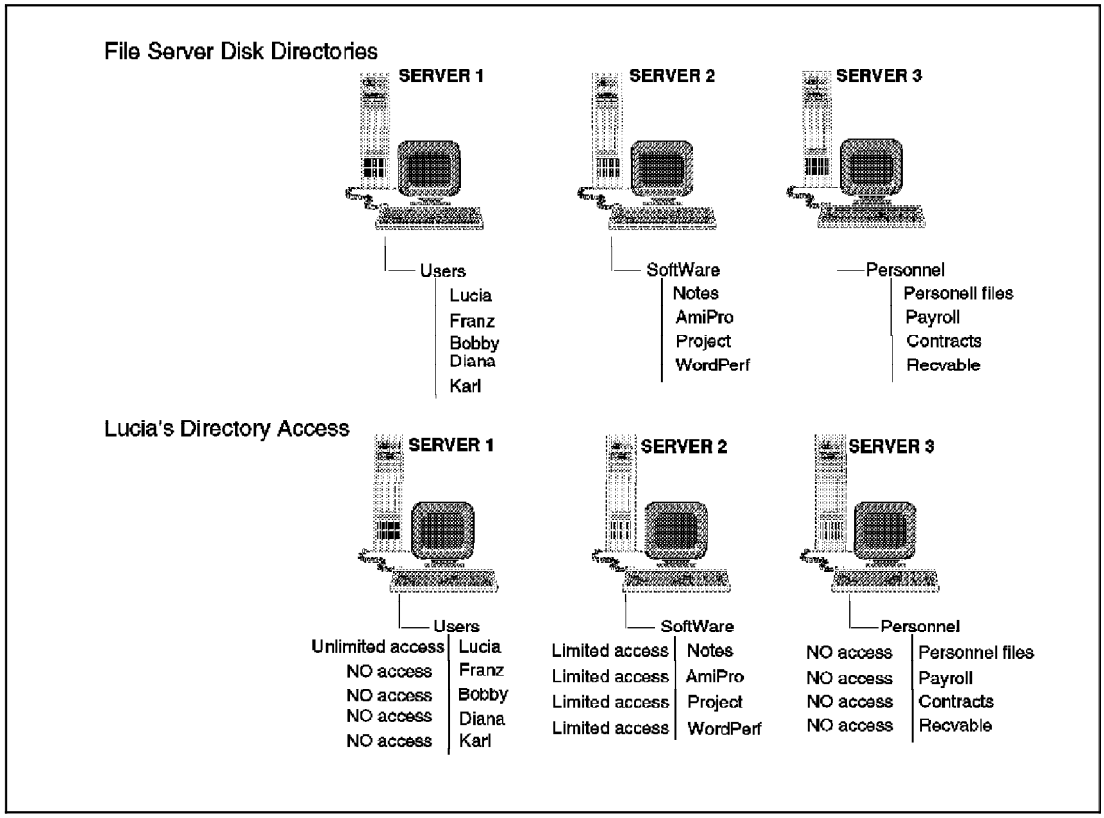

Figure 51. NetWare 4.1. Security

As we can see in Figure 51, all users do not need the same kind of access to each directory. Let's bring up an example:

- In Server 1 we have a directory named users where each user is given a separate subdirectory for personal storage. In the personal

subdirectories in Users you need unlimited access, with the possibility to create delete or modify files as much as you want.

- In Server 2 we have a directory named software and subdirectories with some useful products. In the Software directory you should have limited access to be able to run the programs; you are not supposed to delete or modify a program.
- In Server 3 we have a directory named personnel that may store the company's Personnel information. In the personnel directory none of the users should have access to the personnel server where confidential information is stored.

Trustee Security

There is a NetWare feature called *Trustee Security*. The administrator makes that user a Trustee to that directory. As a Trustee, the user can see and use the files in a particular directory. Referring back again to the screen capture, you can see that the first user is a trustee to her subdirectory under the Users directory and a Trustee to the Software directory (which will give her limited access to each subdirectory under Software), and she is not a Trustee in the Personnel directory or to the subdirectories of the other users.

Table 13. Trustee Security

Access Right	**Allows Users To ...**
READ	Read from (or run) files
WRITE	Write to existing files
CREATE	Create files
ERASE	Delete Files
ACCESS CONTROL	Act as a "mini supervisor" to that directory, granting rights to other users
FILES SCAN	Search the directory's file list
MODIFY	Change file attributes and names
SUPERVISOR	Automatically grant all the above rights and make it impossible for anyone but a network supervisor to take away these rights

The NetWare administrator can further customize the accesses by limiting the way the users can use the files. For example, in the Software subdirectories the user for now is not able to delete or change anything in the AmiPro subdirectory, but when a user becomes a Trustee to a directory,

NetWare allows the administrator to give or withhold eight different directory and file rights. By deciding which rights to grant, the administrator can selectively grant access to the programs in the Software directory.

8.3.2 NetWare Directory Services (NDS)

As already mentioned in 1.4, “Novell Directory Services” on page 19, with NetWare 4.1 there is a new feature called Novell Directory Services, known as NDS, that lets you log in to the entire network, not to individual servers. NDS provides a single security system for your entire network; it is not necessary anymore to provide a logon name and password for each server you need to use.

NDS enables the management of objects such as users, groups, servers, printers, volumes directories and other network resources without the need to know where the resources are located. NDS is an object-oriented resource-management system that links network resources, and it's structure is a hierarchical tree structure, much like a file system directory tree. NDS allows an organization with multiple file servers to control user access to resources as if all resources were subsidiary to a single file server. However, instead of all resources being subsidiary to one file server, they are all subsidiary to one root of an NDS tree. NDS is needed to organize all users and resources into an easy-to-use hierarchical tree structure. So, the structure is referred to as the NDS tree. When designing your NDS tree, remember the following points:

- **A naming convention:** It is important to have a naming in convention so that you will not have duplicated objects names. Naming conventions also help you and your users to locate a particular resource.
- **Good design:** As with a database, a well-planned NDS tree is easy to manage and expand. A good design will even help you better to minimize network traffic.
- **Good resource placement within the tree:** Consider how the various resources are to be used by your users. Put resources and their users in the same container whenever possible. This arrangement makes access much easier and requires less management work in setting up access rights.
- **Logical grouping:** Whenever possible, set up user groups so that access rights are assigned only once. This also eases your manager workload.

The tree is constructed of the following three kind of objects, as you can see in Figure 52 on page 150.

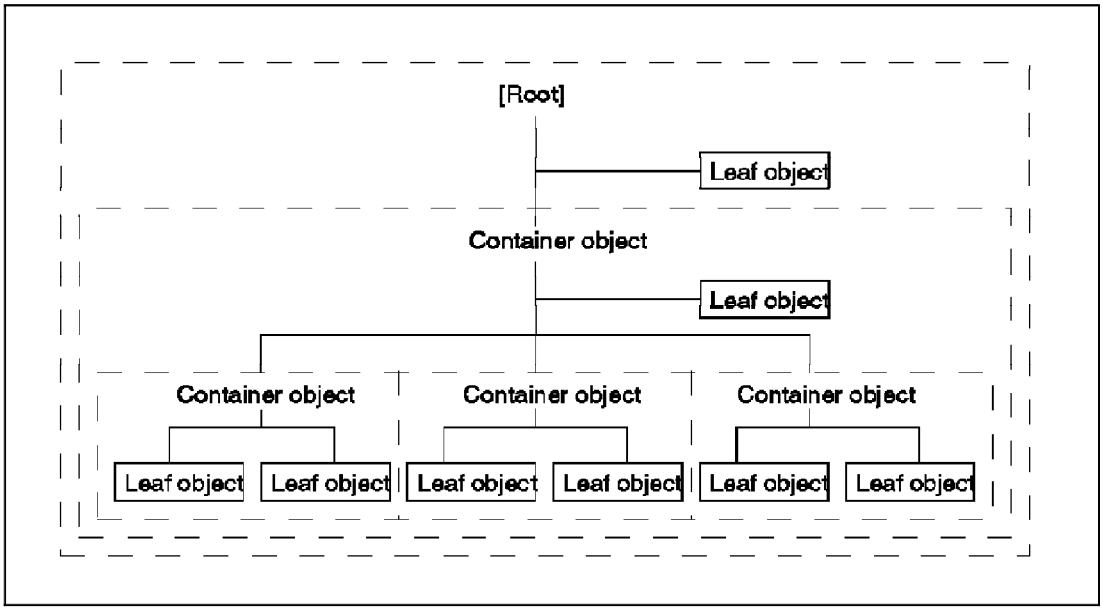

Figure 52. A Schematic View of NDS Tree

- *The* [*Root*] *object*

  The [Root] object is the top a directory tree. You have branches that are made up of container objects, and within them are leaf objects. The [Root] object is created directly when you first install NDS; it cannot be renamed or deleted, and there can be only one [Root] object in a given NDS tree. All objects are either user-related or resource-related. Objects are all intended to provide users (user-related objects) with access to resources (resource-related objects)

- *Container objects*

  - Container objects provide a way to logically organize other objects in the NDS tree. A container object can house other container objects within it. The top container is called the *parent object*. The objects that are contained in a container object are said to be *child objects*

    There are three types of containers:

    - Organization (O=)

      You must have at least one Organization object in the NDS tree, and it must be placed one level below [Root]. The Organization object is usually used to denote a company or main organization.

    - Organizational Unit (OU=)

The Organizational Units are optional; if you use them, you must place them one level below an Organization object or below another Organizational Unit. You can use OU to denote divisions or departments within a company.

- Country (C=)

  The use of Country container objects (C=) is also supported. Country containers are located below [Root]. and above Organization container objects.

Country Objects

It is recommended that you do not use Country objects in your NDS tree. Country objects can make moving around within the NDS tree more difficult.

- Leaf objects
  - Leaf objects are single-entity objects. They do not contain other objects. They correspond to actual physical entities such as:
    1. Users
    2. Group
    3. NetWare Server
    4. Volume
    5. Printer
    6. Print Queue
    7. Print Server

The location is which objects are placed in a tree is called the *Context or Name Context* (similar to a "pointer" inside a database). The Context is very important; we can call it a Key importance. To access a resource, the User object must be in the same context as the Resource object. A user object has access to all objects that lie in the same directory and in child directories.

8.3.2.1 Partitions

NetWare 4.1 encourages you to divide your network's directory tree into partitions that are stored on multiple servers on your network. If you have different locations with many servers, and you are designing a directory tree for a network, consider the possible consequences of storing the entire directory tree database for this network on a single server. If the server that stores that directory goes down, you can not do any of the network features, like connect to the printers, login, or use the network's disk

volumes, because they can't use the directory tree. So when the directory tree is unavailable, so is your network. You can divide the directory tree into partitions starting with any *container object* A partition is simply a logical division of the tree that allows you to create copies (replicas) of that part of the database. By dividing your network directory tree into smaller sections, you can make replicas more easily. When you make a container object, the starting point for a partition is automatically included in the partition.

8.3.2.2 Replicas

NetWare 4.1 also allows you to create replicas of directory tree partitions. A replica is a duplicate of a partition and is stored on a different server. It is possible to create as many replicas for a partition as you want, and NetWare will keep them synchronized. The more replicas there are, the more reliable your directory tree becomes. A good performance of this feature is that a user will surely find the directory tree information he or she needs on a nearby server. The only problem of having many replicas is that there will be more background network traffic because directory tree information must be sent to all replicas. We can divide replicas in three types:

- Master
  - Splits the partition into two new partitions. A partition can have only one master replica.
- Read/Write
  - Read/Write partitions can be updated with new directory tree information, like adding a new user. This information will in turn be replicated to the master partition.
- Read-Only
  - This feature is only to view directory tree information, such as when you want to locate another network resource.

8.3.3 Guidelines for NDS

NDS is very flexible; there can be many different ways to design a directory tree, and they are all good. We can summarize a few guidelines to follow that are applicable in most situations.

Keep Object Names Brief: As we have seen, your complete network name consists of your object name followed by each of you container objects. So if each user's context is three or four levels deep and you've used long object names, the users must know and sometimes type in the full network name when they need to log in. Remember that there is a way to configure a network workstation with the possibility to avoid the full typing of the

name, but when he/she needs to log in from a different workstation on another part of the network, the name has to be a full network name.

Try to Keep Directory Tree to Minimum Levels: It is recommended that you keep your directory's tree depth to no more than three or four levels. This is because the more layers you use, the longer each user's name becomes. This happens because each layer becomes a part of the user's full network name.

Use Consistent Naming Schemes for Particular Object Types: It is reliable to adopt a standard method for naming particular types of objects. A good choice can be to use short informative names. Names should not be longer than eight characters; it is easier to remember them.

Organizational Units Are Based on Company Divisions, Geography, or Both: Organizational Units are the container objects that shape your directory tree according to how your company is organized. You should decide whether the Organizational Units objects in your directory tree should correspond to company divisions and departments or to your company's geography.

Use Alias Objects When an Object is Needed in More Than One Place: If an object in one context needs to also appear in another context, use an alias object to create a virtual copy of that object in the new context.

Agreement of Naming Scheme in Advance: With NDS, it is possible to assign multiple users to manage different parts of the tree; so if different network supervisors are going to add objects to the tree, it is desirable for all supervisors to agree to a naming convention or to a scheme for naming objects.

8.4 Warp Server and Directory and Security Server

Windows NT Server ought to be compared to Warp Server with the Directory and Security Security Server (DSS) installed on it. IBM's Warp Server plus DSS forms one entity. Once DSS is installed on top of Warp Server, it cannot be removed, unless Warp Server is reinstalled. The same thing also applies to if DSS is installed on top of LAN Server 4.0. DSS can interact with Warp Server and LAN Server 4.0.

DSS provides a complete network security today. For DCE applications using DSS, application data can be encrypted using DES (52bit encryption) or a weaker, exportable version of DES called CDMF (40bit encryption). DSS requires that all DSS OS/2 clients and all DSS servers prove to a third-party authentication server that they are who they say they are. This is called

Kerberos, a security system developed by MIT, also used in IBM's TCP/IP products. This prevents masquerading servers and clients from being inserted into the network. Kerberos enables automatic end-to-end security across the network.

The DCE security service also allows file/data encryption; however, this isn't really interesting if you aren't using DCE application programs throughout. LAN Server/Warp Server file access is not encrypted with DSS. Only authentication flows and DCE application program data transfers are encrypted.

Notes:

NT Server security flows are built in and encrypted in all domain login scenarios between NT and NT systems. Encryption methods are proprietary.

Existing LAN Server and Warp Server clients use existing LAN Server/Warp Server authentication flows. However, the primary domain controller does a DCE login on behalf of the client to ensure that the user is still valid in the DCE security registry.

DSS implements DCE, and uses DCE Access Control Lists to limit access to files. These apply to all files within a DCE cell, which can be as large a a whole enterprise. In addition, if the user is running a DSS client, he can access resources across DCE cell boundaries. This applies even if the files are on another DCE compliant server on another platform.

In addition, there is no industry agreement on the format of an Access Control List, so DCE compliance may or may not win in the customer's mind. Note, that NT has networked ACLs also, so if ACLs are an issue, NT has them. NT also works for the whole enterprise. However, it becomes painful because MS recommends a maximum of 10-15000 users per domain. More than that, then trust relationships must be established and maintained. It is the trust relationships which are complicated and turn administrators off.

DCE access control applies only to files on servers upon which DSS has been installed. Legacy additional servers and backup domain controllers which do not have DSS installed on them use LAN server ACLs. Nevertheless, installing DSS on each primary domain controller in the cell allows users to seamlessly access files on both DSS and legacy servers, in any domain in the cell, using the single user ID and password in the DSS cell registry.

8.4.1 Single Sign-On

When DSS is used, a user with a DSS client can seamlessly access resources on DSS servers in his own cell and across cell boundaries using a single ID and password. This requires that DSS is installed on all servers in all cells in the network and OS/2 DSS clients are the only clients that are used. Legacy clients cannot access resources across cell boundaries. Legacy servers cannot export resources across cell boundaries.

On the other side, to implement single sign-on in a network with more than one NT master domain requires the administrator to establish trust relationships between the domains. These allow one domain to "trust" another domain to authenticate the user. In non-trusted domains, multiple logons are required.

8.4.2 Servers in a Network

When Windows NT is running in a large network (more than 10-15000 user accounts), the network must be split up into several NT master domains. There is no automatic mirroring of the master directory in this case, so to access files in another domain, a trust relationship must have been set up, or a new logon must be done.

Maintaining more than one master domain is cumbersome, since keeping the master directories in sync must be done manually. In addition, when passwords are changed, they will flow in the clear, enabling a masquerading server to pick them up.

DSS directories and registries can be replicated. This does not happen automatically. The administrator has to explicitly set things up this way. Assuming this is done, you not only get fault tolerance but also load balancing because DSS will spread queries out over the available servers. No passwords flow across the wire during login. Obviously, passwords do flow across the wire when they are changed. Masquerading primary domain controllers can't be inserted into the network because the domain controllers have to prove to the security server that they are who they say they are and, in order to do that, they have to go through the Kerberos login sequence in a similar way to what clients must do.

8.4.3 Introduction to DSS Security Services

DSS builds upon Warp Server's security model to extend it from the workgroup environment to distributed enterprise-wide networking systems. These increased capabilities in the areas of interconnectivity and interoperability present security challenges in guarding against unauthorized access. DSS addresses these issues by taking full advantage of the DCE architecture and providing an implementation fully compliant with

DCE OSF V1.1 support. The main three elements which form the basis for the DSS security model are:

- Authentication (confirming the identity of clients and servers)
- Secure communication (guaranteed data integrity and privacy)
- Authorization (confirming the privileges and access permissions for a particular identity)

As well as achieving these three main objectives for existing Warp Server environments, the DSS security server offers additional benefits through its use of a single database registry that stores all the user definitions in the cell. While this database is similar to the NET.ACC file, it provides a powerful function for administrators and users by eliminating the need to maintain multiple user IDs and passwords for cross-domain resource access. In DSS, only one user definition is maintained, and using a single logon, the user can access resources within their cell (across multiple resource domains within the cell) and external resources in foreign cells. Other benefits offered by the DSS security services include:

- Increased scalability in a distributed environment due to DCE's ability to support large numbers of users and groups without the need to replicate the registry database to every server in the cell.
- Option for replicating the master registry database to remote sites to improve performance and increase system availability.
- Highly granular administration authority due to DSS implementing the DCE Access Control List Facilities.
- DSS uses very sophisticated protocols based on DCE and Kerberos standards for authenticating and establishing sessions with clients and servers.
- Interoperability and coexistence with LAN Server's security model provide support for incremental migration, making the transition to a full DSS environment easier to manage and implement.
- Interoperability extends to non-OS2 and non-IBM security services due to DSS's use of the open, standard DCE architecture. You can find more detailed information in 1.5, “IBM Directory and Security Server for OS/2 Warp” on page 28.

8.4.4 Introduction to DSS Time Services

DSS implements DCE time servers, which use the security servers to synchronize time between themselves. The security server uses the time services to issue service tickets for file access. These tickets are time stamped, and become void when they expire.

In addition to the DCE function, DSS can automatically reissue the service tickets for DSS and legacy clients. This was done to try to keep the model as close to LAN Server as possible so that users do not have to deal with the situation where they try to access a resource but their ticket has expired.

No similar services exist for Windows NT today.

8.4.5 Introduction to DSS Directory Services

The DSS directory and registry data bases can be replicated among several servers throughout the network. DSS does load balancing for fast directory and registry lookup across LANs and WANs. In addition the directory and registry servers do not have to be from IBM: any OSF DCE 1.1 compliant servers, even from other vendors, can be used. Mirroring of information between the servers is automatic.

Chapter 9. File Systems

In this chapter you find information about file systems that come with the network operating systems discussed in this book. Our interest in file systems is not in all the details, but it may give you hints and tips to decide which file system suits you best for your environment. We also give you information about the partitioning of hard disks.

9.1 Directory Structure

The directory structure and methods for organizing a partition is called a File System. Different File Systems reflect different operating system requirements or different performance assumptions. UNIX, for example, has the convention that lowercase and uppercase are different in file names, so "sample.txt" and "Sample.txt" are two different files. DOS and the systems that descend from it (Windows 95, OS/2, and Windows NT) ignore case differences when finding file names. Some File Systems work better on small machines, others work better on large servers.

Each partition is assigned a type (in the Master Boot Record for primary partitions, in the Extended Partition directory for logical volumes). When the partition is formatted with a particular File System, the partition type will be updated to reflect this choice.

The same hard disk can have partitions with File Systems belonging to DOS, OS/2, NT, AIX (or other UNIX clones). Generally, an operating system will ignore partitions whose type ID represents an unknown file system type. It is fairly easy (given a big enough disk) to install all of the different operating systems and all of the File System types. There are a few rules to make things simple.

Each File System is described in detail in a separate section.

- FAT File System

  The FAT File system is used by DOS and is supported by all the other operating systems. It is simple, reliable, and uses little storage.

- VFAT

  VFAT is an alternate use of the FAT file system available in Windows 95 and Windows NT 3.5 and higher. It allows files to have longer names than the "8.3" convention adopted by DOS. VFAT stores extra information in the directory that older DOS and OS/2 systems can ignore.

- HPFS

  HPFS is used by OS/2 and is supported by Windows NT up to version 3.51 (It is not supported by Windows NT 4.0 anymore). It provides better performance than FAT on larger disk volumes and supports long file names. However, it requires more memory than FAT and may not be a reasonable choice on systems with only 8 MB of RAM. The 32bit version of HPFS is called HPFS386. This file systems increases I/O performance dramatically and therefore is a good choice for file servers.

- NTFS

  NTFS provides everything. It supports long file names, large volumes, data security, and universal file sharing. A departmental NT file server will probably have all its partitions formatted for NTFS. Because the other operating systems cannot use it, NTFS is less attractive on personal desktop workstations or portables.

9.2 File Systems and Disk Letters

DOS and Windows 95 can only boot from the C: disk. Technically, the C: letter will be assigned to the first Primary Partition on the first hard disk that has a FAT file system. In no case can DOS boot from a second hard disk or a logical volume in the extended partition. However, if as the system comes up, the DOS boot sector and DOS files turn out to be on the second Primary Partition on the first hard disk, then this will not be a problem so long as the first partition has a non-FAT file system. DOS simply ignores primary partitions that are formatted for other operating systems.

Some people exploit this feature. They put an HPFS or NTFS file system on the first Primary Partition, and a FAT file system on the second. This can produce confusion. When the other operating system boots up, it will now assign letter C: to its first partition, and the disk that DOS calls "C:" will become "D:" on the other system. If the two systems share application programs, it becomes very difficult to configure INI files as the drive letter keeps changing back and forth. It is a simpler and safer strategy to accept the view that the first Primary Partition on the first hard disk should be formatted with the FAT file system and should be the C: drive in every operating system.

9.3 Partitioning a Hard Disk

Why should you partition your hard drive? There are several good reasons to to this:

- To optimize your hard disk for storing the maximum amount of information

- To run multiple operating systems from one PC
- To increase the degree of security and integrity of the information you store on the hard disk

First, creating smaller partitions on a large (larger than 512 MB) hard drive will force your system to save its files more efficiently. Huge hard drives have very large minimum cluster sizes. To explain what this means, we will use an analogy involving a very large file cabinet.

Every drawer of this cabinet can contain only one file. If it fills the drawer, you are using the drawer efficiently. However, if you just put several sheets of paper into that drawer, you will waste tremendous amounts of space. Computer files can span several clusters, with the wasted space 'living' in the last cluster occupied. Cluster size is determined by the size of the hard drive. Use the table shown in Table 14 to see where you stand on wasting space when using the FAT file system.

Let us assume the following. We have a computer with a 1.2 GB hard drive. From the chart shown in Table 14, you can see at this size, our cluster size is 32 KB. This means that if you are saving a 8 KB word processing document, you are saving 8 KB in a 32 KB "drawer". You can not save anything else in the drawer; so you are wasting approximately 75 percent of the space available.

If you partition your large drive into two partitions of 600 MB each, your cluster size drops to 16 KB. Storing a 8 KB file in a 16 KB cluster improves your storage situation by percent. Now you are wasting only 50 percent of a cluster.

To maximize your situation, you can repartition your drive into ten partitions of 127 MB and have clusters of just 2 KB. That way your 8 KB file will occupy four clusters, with nothing left over. If this file you were storing were a 9 KB file, it would occupy 5 clusters, with 1 KB wasted space in the fifth cluster. The file is only wasting 10 percent which is significantly better than the 32 KB sectors.

Table 14. Clustering of FAT Partitions depending on Partition Size

Total Partition Size	Minimum Cluster Size	Space Wasted
16-127 MB	2 KB	2 percent
128-255 MB	4 KB	4 percent
256-511 MB	8 KB	10 percent
512-1023 MB	16 KB	25 percent
1024-2048 MB	32 KB	40 percent

9.4 File Allocation Table (FAT)

FAT file system advantages:

- Supported by all operating systems
- Minimal memory use
- Simple and reliable

FAT file system disadvantages:

- filename.ext (8/3) file names
- Not efficient on large partitions
- Not suitable for file servers

The File Allocation Table (FAT) was designed and coded in February 1976. It was a version for Basic that could store programs and data on floppy disks. The FAT design was incorporated by Tim Patterson in an early version of an operating system for the Intel 8086 chip. Bill Gates bought the rights to the system, then rewrote it to create the first version of DOS.

The FAT file system is simple and reliable. It does not lose data because the computer crashed in the middle of an update. It does not use a lot of memory. It does do a lot of extra administrative I/O to different areas of the partition.

However, FAT was a good fast method to access files on a diskette drive. At this time diskette drives had a capacity of 128 KB, 360 KB, and in a Evolution Version 1.2 MB. To get performance advantages the FAT was read into memory. This was a big advantage against CP/M where allocation information was scattered over the whole disk.

The described advantages turned into big disadvantages with the upcoming hard disks. This shifted the proportion. With the growing hard disk space, the proportion between memory and hard disk space grow from 1:5, in a PC with 64 KB RAM and a 360 KB Diskette, to 1:160, in a PC with 640 KB and 100 MB Hard disk. The result was that FAT, in the flow of this evolution, could only partly read into memory. This means that when working with files, the hard disk's head had to do unnecessary head movements that slowed down the whole system.

If you look at the history, there was another gap that DOS had to deal with in Version 4.0. This version was the first version that supported partitions greater than 32 MB. Theoretically, it is possible to support partitions up to 2048 GB (32-bit pointer to 512 byte sectors). As a practical matter, FAT is not

useful for partitions greater than 127 MB because files with a length of 1 byte will reserve 4 KB of disk space. This value will double at partition sizes of 256 MB, 512 MB, 1024 MB, and so on.

This phenomenon applies because of the use of 16-bit pointers to the clusters of a file. At the maximum, 16-Bit pointers can address 65536 different values. This means this is the number of clusters that could be given to different files. In this case a cluster can include different sectors.

Note: This means that the bigger the clusters, the more ineffective they are for small files.

At the start of the FAT-Filesystem partition, you find the so-called Bootsector. This contains, after a short load sequence (IBMDOS.COM), the Basic-Input-Output-System (IBMBIOS.COM). Behind the Bootsector follows a depict of the cluster for the partition. For every cluster, you find a 12- to 16-bit entry, depending on the DOS version. This depiction is the File-Allocation-Table.

The directory is allocated at the start of the partition, and it contains the table of free space. To write a new dataset, or to add data to an old one, the disk arm must be constantly moved between the location of the directory and the place where the data is being written.

When the system crashes, no data is lost. However, a FAT system may have removed disk area from the chain of free space, but may not yet have assigned it to any permanent new dataset. The `CHKDSK` (or on newer systems the `SCANDISK`) command examines the FAT table to determine the status of every record on disk. The records that are not part of any dataset may be returned to the free space chain. After `CHKDSK` finds unallocated sectors, it asks you whether they should be turned into files.

By design, FAT supports a maximum of 64 KB allocation units. When the disk partition is 32 megabytes or less, then an allocation unit is a 512-byte sector. However, as the disk gets larger, the units get larger. A 64-megabytes disk partition has 1 KB allocation units. A 128-megabytes partition has 2 KB units. A 256-megabytes partition has 4 KB allocation units.

Each file occupies one or more allocation unit. As the allocation units get large, any large number of small files wastes a lot of disk space. The classical FAT directory structure limits file names to eight characters with a three-character extension. This 8-to-3 naming convention was borrowed from earlier DEC minicomputers.

The FAT structure also maintains for each file a set of attributes:

- System dataset
- Hidden dataset
- Archived next time the disk is backed up
- Read-only

There is also a data and time stamp when the file was last changed. OS/2 allows a FAT file to have additional Extended Attributes. Since there is no room for these attributes in the FAT directory, OS/2 creates a separate hidden file on the disk volume named "EA DATA. SF" and stores the information there.

For directories, DOS stores 32 bytes in the FAT. But when the directories are sharing the same space on the disk as the files, only the Root-Directory has a fixed position. This means the directories are stored unordered, together and between the files. The effect is that the operating system has the whole disk for deep directory structures and has to reposition the disk head many times. This slows down the overall system performance.

If you look at the disk structure, you see at the beginning of the disk the bootsector, followed by the FAT, and a FAT Copy. After that, you have the fixed position of the root directory that is followed by the file/directory area. For a better understanding, look at the following graphical representation:

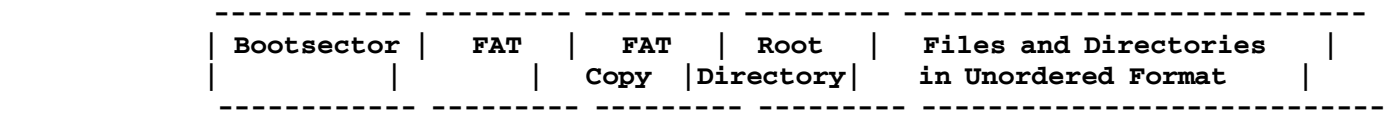

Figure 53. Structure of the File Allocation Table

9.5 Virtual File Allocation Table (VFAT)

VFAT Filesystem advantages:

- Supports long file names on FAT under Windows 95 and NT

VFAT Filesystem disadvantages:

- Long file names are not supported by DOS, OS/2, and AIX.
- Long file name structure could be destroyed by DOS, old Windows 16-bit Programs, OS/2, and AIX.
- Proprietary because it is only useable by Microsoft Operating Systems

Windows 95 and Windows NT support VFAT. Technically, VFAT is not a new File System. It uses the same directory structure, format, and partition type as ordinary FAT.

VFAT is simply a way to store more information in the FAT directory. The most important feature of VFAT is the ability to store long file names. Since it is built on ordinary FAT, each file has to have an 8-character name and 3-character extension.

However, VFAT then allocates additional directory blocks to hold a longer file name. Programs running in DOS, OS/2, AIX, and old 16-bit Windows programs will not see the longer file name. Only WIN32 programs running in NT or Windows 95 can make use of the longer name.

Because VFAT uses the old FAT directory to add some unusual new entries, the VFAT additions can be damaged if the disk is manipulated by a DOS, AIX, old 16-Bit Windows programs, or by an OS/2 disk utility that does not understand the new structure. Even a simple `ERASE` command under OS/2 for a FAT dataset with a long file name can leave the extra blocks in the directory.

9.6 High Performance File System (HPFS)

HPFS advantages:

- Supports long file names with mixed case
- Directory positioned throughout the disk
- Much faster to create new files
- Support for Windows NT up to Version 3.51, Linux, and OS/2 with long file names
- Allocates individual sectors (512 bytes)

HPFS disadvantages:

- Uses more memory, bad for small PCs with less than 8 MB of RAM

With the appearance of OS/2 Version 1.2, there was a new file system delivered and called High-Performance-File-System (HPFS). This file system was developed by IBM and Microsoft and should overcome the limitations of the FAT file system design.

Technically, with HPFS, IBM followed a concept implementing new file systems by making the filesystems loadable over so-called Installable-File-System (IFS) driver.

There are many myths around HPFS. An example is that HPFS is defragmenting the hard disk by itself. This is not correct, by design HPFS has a better structure and by this it prevents the data from being fragmented. Technically, the tables that describe the location of files and free space are positioned at regular intervals throughout the dataset. New datasets are written where there is a large amount of free space. This reduces fragmentation and keeps the disk arm from jumping around. In addition the directory band is placed in the middle of the disk; that makes it possible to faster find directories and prevent the disk head from unnecessary head movements.

Structure of HPFS: The data filing works much more effectively because the disk is managed on sector level. HPFS maintains a 512-byte allocation unit (clustersize) no matter how large the volume becomes. With this action, HPFS minimizes the disk space that is lost when filing small files; this means a 1-byte file occupies 512 bytes no matter how big the disk is. The HPFS directory allows file names to be long, to have multiple periods, and to have lowercase letters, unlike NT and Windows 95.

The structure on the disk is enhanced to meet the requirements of effective head movements and performance. Therefore HPFS splits file-bands and bitmap-bands over the disk. The file-bands are surrounded by the bitmap-bands. So to write a file, the system has to look into the bitmap-band and write the file to the free space in the closest file-band. To illustrate this, see the following graphical representation:

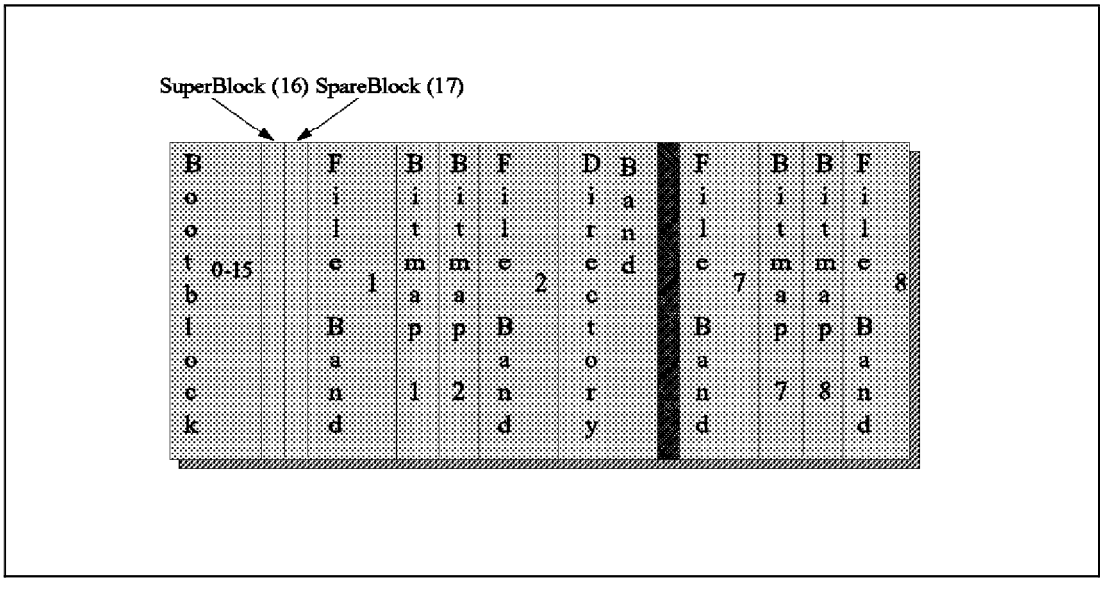

Figure 54. Structure of High Performance File System

The first 16 sectors, from sector 0 to 15, of an HPFS partition represent the Bootblock. Inside the Bootblock is the partition name, a 32-bit big partition-ID and the BIOS-Parameterblock, a Bootstrap program.

Sector 16 is the location of the Superblock. The Superblock consists of pointers to the bitmap-band, a list of defect blocks, the directory-bands and the root directory. Within the Superblock is also recorded when the hard disk was last checked with `CHKDSK /F` to find inconsistencies.

Caching and Delayed Write of HPFS: Decisive for the high performance of HPFS is the technology HPFS uses to delay write accesses to the hard disk. The "LAZY-WRITE" design implemented delays the writing of data to the disk and maintains a balance between high and low load. This means a text program already has the confirmation that the write operation is successfully done though HPFS and it keeps information in a disk-cache area of memory until it needs to be written to disk. The design has also a function to switch off this feature.

The caching of write access in this case is a problem. Because the system can only pinpoint defect sectors when the file system actually writes to the disk or when something like a power failure occurs, the data is not written. To be sure that all the information has been properly written to disk, a user should try to shut the system down properly rather than just turning the power off.

To address the defective sector problem HPFS prepares to so-called Hot Fix within the Spareblock. The data will be diverted to these sectors when a problem comes up. All read and write access first checks the sector numbers and remappings with the Hot Fix list. In addition HPFS sets the dirty flag inside the Spareblock to indicate that not all physical write accesses were completed. This indication bit lets the system recognize that and check the hard disk automatically.

Fixing and Recovering Data after System Crash: The system will occasionally crash. Any HPFS volume that was in use during a crash is marked "dirty". Before it can be used, the next boot of the operating system must run `CHKDSK` to examine the chains of free space and file locations to correct any problems. As the disk volumes get larger, running `CHKDSK` after a crash can take some time.

`CHKDSK` can recover serious problems in the filesystem through the architecture. Every data object is double connected, has a 32-bit signature and the FNodes are aware of the starting letters of the files and directories.

But CHKDSK not only has the task to recover inconsistent drives, it also has the task of building files that are saved in the Hot Fix sector back into the file system and making room for the next emergency.

Compatibility to Other Operating Systems: OS/2 regards HPFS as its native file system and provides full support. Windows NT up to version 3.51 has the capability of reading and writing files from to to an HPFS partition. It can not format partitions using the HPFS file system.

Notice for DOS Users

There is a very nifty driver named AMOS that provides read-only access to HPFS volumes. AMOS3 is installed as a terminate and stay resident (TSR) program by the AUTOEXEC.BAT file. It then supports access to the HPFS file structure much as the MSCDEX utility provided read-only access to files on the CD-ROM.

9.6.1 High Performance File System 32-Bit (HPFS386)

Advantages of HPFS386:

- Access Control Lists (ACLs) based on the directory structure and file structure
- Directory and user-based disk space limitations (DASD limits)
- Local security
- Software RAID 1 (fault tolerance)
- Fault tolerance functionality for VINCA support (see 10.1.1, “Vinca StandbyServer for OS/2 Warp Server” on page 179)
- HeapSize configured depending on the cache
- Cache size minimum 256 KB, maximum based on memory room (HPFS386 = 320 MB).
- Efficient for large disks

Disadvantages of HPFS386:

- Access on hard disks with ACLs only with HPFS386 boot diskettes
- No access with other operating systems

The HPFS386 file system is highly optimized and designed for Pentium Pro, Pentium, 80486, 80386, 80486SX, and 80386SX-based platforms with large disk systems. HPFS386 provides extremely fast access to large disk volumes and optimizes performance in the server environment, in which many files are open simultaneously. HPFS386 is an enhancement of the

regular HPFS that is part of OS/2 1.2, 1.3, 2.0, 2.1, 2.11, Warp 3 and Warp 4. It represents the logical evolution of LAN Server technology. The server consists of an optimized ring 0 server tightly coupled with a bootable Installable File System (IFS) and customized device drivers to accelerate network I/O.

The 386-specific version of the HPFS is disk-format compatible with the OS/2 Standard Edition version. The existing HPFS partitions do not require reformatting when HPFS386 is installed.

The major changes from HPFS386 are the following enhancements:

- Cache addressing capability beyond 16 MB memory
- Increased maximum number of open files from 8192 to 65536
- Increased maximum number of file finds from 3072 to 8192
- Increased maximum number of file searches from 1024 to 6144

Access Control System: The user IDs, group IDs, and passwords for all users within a server's domain are stored in a user accounts database (NET.ACC) on the server. On a LAN Server Entry (or Warp Server) workstation with either the FAT or HPFS file system, the access control profile information is stored in the NET.ACC file. On a LAN Server Advanced (or Warp Server Advanced) workstation with HPFS386 installed, the access control profiles for the HPFS386 files and directories are stored within the file system.

This means:

- Every file or directory on a HPFS386 volume is anchored on a fundamental file system object called Fnode. The Fnode is a fundamental object in an HPFS volume and is the first sector allocated to a file or directory.
- Each Fnode contains control and access history information used internally by the file system.
- The Fnodes contain a new allocation information about the access control list.

The access control profiles for all other resources (for example, FAT files, print spooler queues, and serial device queues) and drive-level access control profiles for HPFS386 drives are stored in the NET.ACC file. Up to 8192 access control profiles can be stored in the NET.ACC file. A 386 HPFS workstation stores an unlimited number of access control profiles for directories and files residing on the HPFS386 drives.

The access control list consists of different entries, called Access Control Entries (ACE). For every access restriction of an ID you can find an entry. The entry consists of the user ID or the group and the permission. The entries are created when an access profile is created for a special file or directory. Each ACE is sometimes called as Access Control Profile (ACP).

Inherited Access Control: For file aliases, an Access Control Profile usually must be created before users can use this resource. However, an Access Control Profile is inherited automatically if the files resource is either created remotely or resides on an HPFS drive and the HPFS386 is installed on the server.

When you create a directory either locally or remotely on an HPFS386 server, the newly created directory inherits the Access Control Profile information of the parent directory. Because of the way the File Allocation Table (FAT) works, you can inherit only a remotely created directory's Access Control Profile on a FAT file server. You must have access to the Access Control Profile on the server to be able to inherit it. You must be logged on with an ID that is allowed access to the parent access control profile. If successful, a new profile is created with the same permissions as the parent of the new directory.

Effects Of Renaming Or Deleting Directories: If you rename a directory, you must manually delete and recreate any Access Control Profiles for subdirectories under the directory. Renaming a directory does not automatically update Access Control Profiles for the subdirectories. This only applies to HPFS or FAT file systems. The HPFS386 Access Control Profiles remain with the renamed directory.

If you delete a directory on a local drive, the associated Access Control Profile is not deleted. However, if you delete a directory on a local HPFS386 drive, the Access Control Profile is deleted. If you delete the directory of a redirected drive, the Access Control Profile is always deleted, whether it is on a local HPFS386 drive or not. Check the list of Access Control Profiles on each server periodically, and delete those that have no existing files resource.

Backing Up Access Control Information: Access control information cannot be copied when backing up systems to tape because the ACL (Access Control List) is now in the file node. It is recommended that the `BACKACC` utility be used to back up the ACLs in the HPFS386 into a file, and then back up the whole disk using the tape backup program.

Existing Drives Or Other File Systems: When you install the HPFS386 option, it installs a device driver that automatically enables you to make use of the HPFS386 technology. The data integrity remains the same (a format is not required if the partition or drive is already formatted HPFS). With the HPFS386 driver loaded, an HPFS drive remains an HPFS drive unless the IBM OS/2 Warp Server Advanced code uses/touches that drive. For example, if you create an alias or issue a `NET SHARE` command against a drive (which creates an Access Control Profile), the drive becomes an HPFS386 drive. Therefore, a particular partition may not be a HPFS386 drive unless you have used that drive in association with IBM OS/2 Warp Server Advanced file and print functions.

Note that if you need to keep certain partitions FAT or simple HPFS, it is important to not have IBM OS/2 Warp Server Advanced deal with that particular partition (which you are using for another file system). If the IBM OS/2 Warp Server Advanced system does use that partition, it will become HPFS386-"formatted" and the other file system will not be able to utilize that drive again.

Hot Fix: With IBM OS/2 Warp Server Advanced and HPFS386, there is a feature called "Hot Fix", a segment of the disk is reserved for moving data from faulty disk areas. The "Hot Fix" feature is a part of the HPFS file system, and since HPFS386 has the same basic format as HPFS, this feature is also a part of the HPFS386 file system. The "Hot Fix" feature will write data to a reserved part of a disk when it attempts to write to a bad sector, instead of writing to the bad sector. This reserved "Hot Fix" area takes up approximately 10 percent of total disk space.

CACHE386.EXE reports the number of times it has to write to the "Hot Fix" area, and thus the number of bad disk sectors since bootup.

9.7 Windows NT File System (NTFS)

NTFS advantages:

- Supports long file names
- Access control by directory or file
- Can compress individual files or directories
- Efficient for large disks
- Add space when partition fills up

NTFS disadvantages:

- Access only possible with NT

- Proprietary file system only used from Microsoft's NT

NTFS is a new file system for Windows NT; it includes different beneficial functions of other file systems. In addition to that, this 64-bit file system provides support for file sizes of up to 17 billion GB and corresponding partition sizes. Also the security design is included in NTFS. However, NTFS in incompatible to all other operating systems until now.

NTFS is not allocating individual sectors (512 bytes); it is cluster oriented. This means, that the cluster size grows bigger and holds more individual sectors, as the hard disk gets bigger. NTFS can compensate this a bit with the 64-bit filesystem structure. The sector/cluster size doubles every 512 MB. For the factors, see the table shown in Table 15:

Table 15. Relation between Cluster Size and Partition Size

Sector per Cluster	**Cluster size**	**maximal Partition size**
1	0,5 KB	512 MB
2	1 KB	1 GB
4	2 KB	2 GB
8	4 KB	4 GB

This is especially bad when saving small files because NTFS is managing clusters and not sectors, a 1-byte file will use 2 KB of space on a 2 GB hard drive. The unused space can not be allocated from other files.

The NTFS file system supports long UNICODE file names. In theory, an NTFS file can have its name in Chinese or Hebrew characters. At the same time, NTFS maintains an 8-to-3 name for the file so that it can be used by a DOS program or old 16-bit Windows program, but only within a Windows NT session. NTFS also supports case-sensitive file access, for UNIX programs and case-insensitive file access for DOS, OS/2, and Windows programs.

The recorded information is completely saved in the form of files. The important information is hold in the Master File Table (MFT). This self-organized file has the structure of a table that holds the information about the contents of the entire drive, including the information about the contents of the MFT itself. The structure of an NTFS hard disk looks as shown in Figure 55 on page 173.

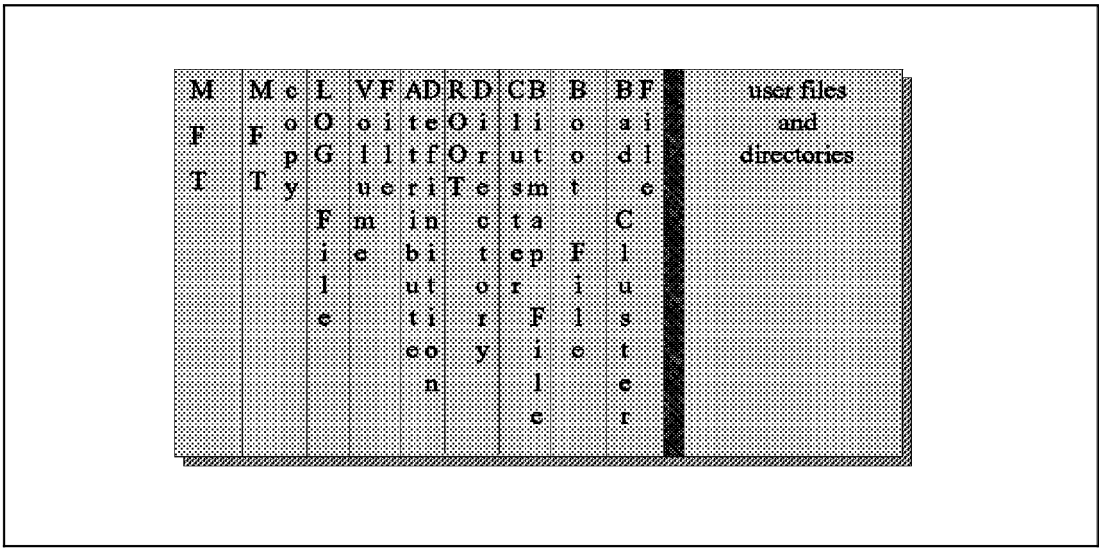

Figure 55. Structure of File Records for NTFS Metadata Files in the MFT

In NTFS, all data stored on a volume is contained in a file, including the data structures used to locate and retrieve files. The first 16 sectors of the MFT are reserved for special information. The first file holds the information about the contents of the volume, and the second file holds a copy of the first file. This is the "real" MFT. For the additional functionality, NTFS offers special files, reserved for different information. These are the Log file, the Volume file, the Attribute-Definition file, the Root Directory file, the Cluster Bitmap file, the Boot file (when it is a bootable media), and the Bad Cluster file. All this is followed by a section for user files and directories. Storing everything in files allows the data to be easily located and maintained by the file system, and each separate file can be protected by a security descriptor.

9.7.1 Transaction and Cache Concepts

The concept of the transaction orientation starts with the Log file, which logs the transactions taken in the file system. Because NTFS works with transaction orientation, every change on a file or directory is seen as a transaction, and the changes of the file system will be logged. In addition there is information on what has to be done to make these modifications and how to cancel the transaction. After a transaction is successfully done, a confirmation will be received. Unsuccessful transactions can simply be canceled. The same could be done when the system crashes.

These functions are needed to fix the problems that could appear when using a Lazy-Write file system concept. The file system raises its performance with timely, delayed write-back processes. This means an application could already have received a confirmation that a file is successfully written, without ever being written on the disk. The time delay

has another effect. When a file can not be written, there must be emergency measures to avoid a data loss. Also there must be a way, in case of a power loss, to recreate the volume and file structure and to recover the data.

The Log file holds all the information to redo transactions when writing on a disk fails due to a writing error. Compared with the HPFS architecture, this is a time intensive solution, especially with big files, because the whole transaction has to be done again after a write error is detected.

Attributes and Extents: As noted before, the MFT contains records for every file on the disk. This also includes a file for the MFT itself. So a record in the MFT contains either all attributes for the file, this is called *Residents*, or a Virtual Cluster Number (VCN) to Logical Cluster Number (LCN) mapping of where to find the additional attributes. Those attributes not stored with the file stored attributes are called *Extents*.

So additional information to the file would not be stored in Extended Attributes as it is in HPFS386; it would be stored with the file itself as it is in HPFS386, and saved as Extents only when the Attributes are getting too big.

Security- and User-Model*: NTFS supports a variety of multiuser security models. There is native Windows NT security established by file manager based on the groups to which a user's account belongs.

Windows NT server also supports a Macintosh security model that simulates an Apple File Server. UNIX applications will see security that obeys the POSIX model. A Windows NT Server cannot share disk space with Macintosh computers unless the volume is in NTFS format.

Special NTFS Functions: NTFS supports "volume sets" where a single disk letter is associated with a "volume" created from a number of separate free-space areas scattered across several disks; this means NTFS in not only supporting software RAID 1, like HPFS386, it also supports software RAID 5. NTFS has built-in software capabilities that provides support of RAID 0 up to RAID 5. So if an NTFS volume fills up, it can be dynamically expanded by adding an extra chunk of free disk space from the same or from another hard disk. From the performance point of view, this RAID 5 functionality can not replace an hardware RAID array. However, it might be an acceptable solution for small networks.

While DOS has drivers that allow an entire disk to be compressed, NTFS allows infrequently used files or directories to be individually selected for compression. Files are automatically decompressed as they are used, and

new files are compressed if they are stored in a compressed directory. Frequently used files can be left uncompressed to avoid slowing the system.

9.7.2 Cracking NTFS

Alarmed Windows NT sites are scrambling to assess the implications of a shareware program that cracks NT's file protection mechanism.

Corporate users and security specialists who have downloaded the utility program NTFSDOS from the Internet warn that it compromises NT security (visit the World Wide Web at `http://www.ntinternals.com/ntfsdos.htm` for more information about this utility).

The program is an NTFS network file system redirector. Creators Mark Russinovich and Bryce Cogswell describe it as a redirector based on Linux code which is designed to recognize and mount NTFS drives for transparent access.

But inserting a DOS boot disk containing the utility in an NT workstation or server allows NT's mechanism for limiting access to user files to be circumvented.

Files stored under NTFS are protected by a security bit related to the file. Setting the bit causes access requests to be directed to NT's password-enabled user management system. However, NTFSDOS, loaded onto an NT workstation or server via a DOS boot disk, permits the security bit to be disregarded and all user files in that disk partition to scrutiny.

The current version of the utility lets files be read or copied to another disk, which raises obvious concerns about password files. It does not permit deletions or writing to files nor can the utility see compressed drives.

Chapter 10. Performance, Scalability and Availability

Symmetric Multiprocessing (SMP) and Clustering are two technologies being worked on in the Intel market space to improve performance, scalability and availability of applications and data. Customers, implementing either technology, should carefully consider server workloads, network and system bottlenecks and system application requirements.

All network operating systems studied in this book come with fault tolerance which is a good way, for example, to enhance availability of data, applications, and resources. However, increasing availability of nodes is a different story. This can be done by establishing clusters. Different scenarios are described in 10.5, “Introducing Clustering Technology by IBM” on page 195 and shown in Figure 62 on page 196.

Fault tolerance allows drive mirroring and duplexing as well as error logging, alerting, and monitoring of disk activity, thereby improving data integrity. Fault tolerance works together with a compatible disk device driver and the file systems to detect and correct disk faults. When a fault is detected, it is logged. Drive fault monitoring is possible regardless of number of disks in the system and whether they are mirrored or duplexed. Fault tolerance protects against a single failure, but not against multiple failures.

10.1 Fault Tolerance in Warp Server Advanced

The fault tolerance feature is only available with IBM OS/2 Warp Server Advanced, and it must be installed from File and Print Sharing Services Services. Two disks are mirrored when they are on the same disk controller, and they are duplexed when they are on two different disk controllers. When two disks become a mirrored pair, an extra level of protection against disk failure is put into place. If one disk in the mirrored pair fails, no data is lost because the other disk in the pair contains the same information. Also, no interruption in server availability occurs because the server continues to operate. When two disk become a duplexed pair, the level of redundancy extends not only to the disks themselves but also to their controllers. With a mirrored disk pair, a failure of the controller they share results in server failure and possible data loss or corruption. With a duplexed pair, the probability of a controller failure causing data loss is reduced because duplicates of each disk controller exist. So as you can see in the following steps:

- **Drive Mirroring**

Drive mirroring is the duplication of a single logical drive or volume on two partitions that do not reside on the same physical disk. If the data on the two partitions differs, the drives are synchronized through drive verification; that makes sure that the mirrored or duplexed drives are identical.

- **Drive duplexing**

  Drive duplexing is a special type of drive mirroring, with the additional advantage that the two disks on which the two partitions reside are controlled by two different disk controllers. Drive duplexing provides protection against errors caused by a faulty controller and improves read performance.

Each mirrored/duplexed pair of partitions must reside on disks handled by the same device driver. If you want to use drive mirroring or duplexing, you must have a minimum of **two physical disks** on the workstations. For mirroring/duplexing the following additional constraints must be addressed:

- There must be enough free space for `FTSETUP` to create the secondary drive, or for the back-up drive.
- You can only mirror into a logical drive.
- Recovery is a special feature of mirroring; so you cannot recover into a primary partition.
- If you want to mirror your boot drive, the boot drive should be created as a logical drive using `FDISK`. First create the Boot Manager partition, then a C: drive configured as logical drive.

  Note: Do not create any other primary partitions. If you do so, different drive letters will be assigned to your logical drives which could cause the system from not being able to start up anymore.
- You can mirror drives on disks of different physical sizes on the condition that there is enough contiguous free space available. You can mirror a logical drive located on a 320 MB disk to a 160 MB as long as the logical drive is less than 160 MB. Of course you cannot mirror a 200 MB logical drive to a 160 MB disk.

Drive duplexing also requires that these disks be on two different controllers. Each mirrored or duplexed pair of partitions must reside on disks handled by the same disk device driver.

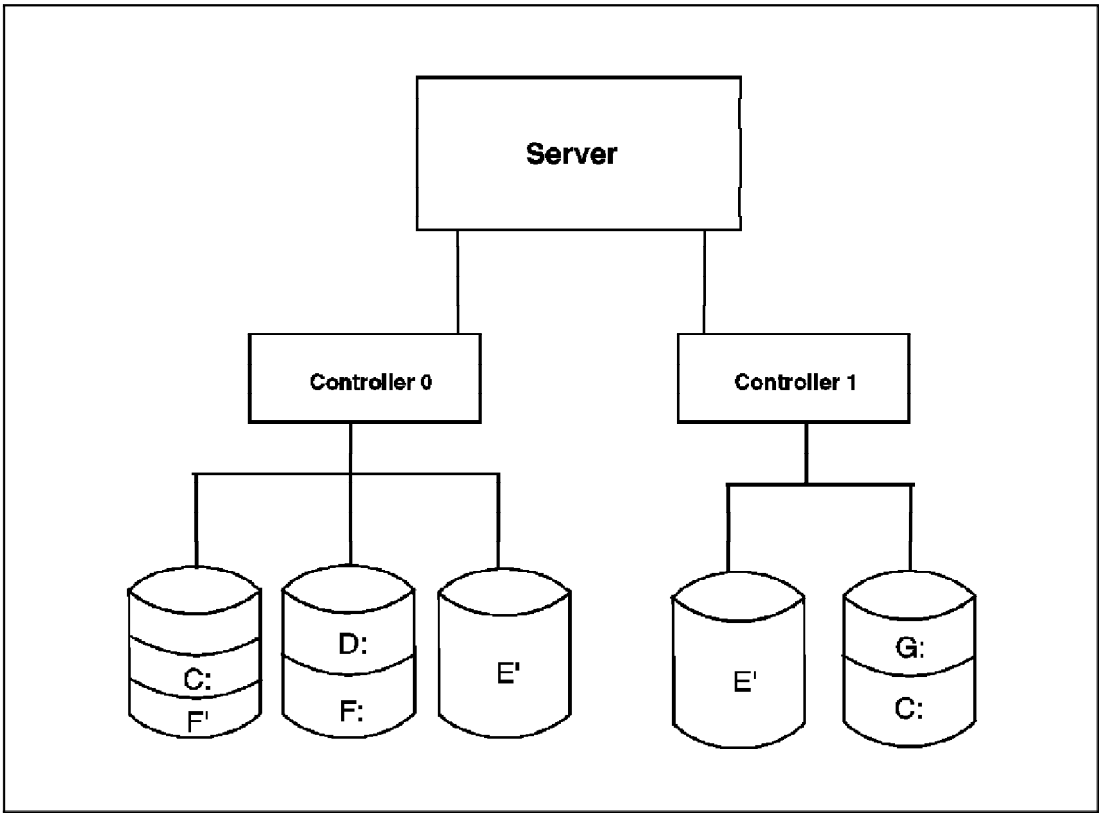

Figure 56. Server Configuration with Mirrored and Duplexed Drives

Figure 56 illustrates a server configuration with five physical disks. Because of mirroring, the actual storage capacity is three disks. Drives C: and E: are duplexed drives, F: is a mirrored drive.

10.1.1 Vinca StandbyServer for OS/2 Warp Server

Another feature for fault tolerance is *Vinca StandbyServer*. StandbyServer is a software product consisting in device drivers, services, and utilities. It works cooperatively with OS/2, Windows NT and NetWare, mirroring the disk drives of two servers over a dedicated link enabling one server to watch out for the other. In case of failure, it takes over control of the other. The hardware consists of a pair of network interface cards and a length of cable providing the dedicated link over which the software mirrors the data. For networks running OS/2 Warp Server, *Vinca StandbyServer* is the premier fault-tolerant mirroring system.

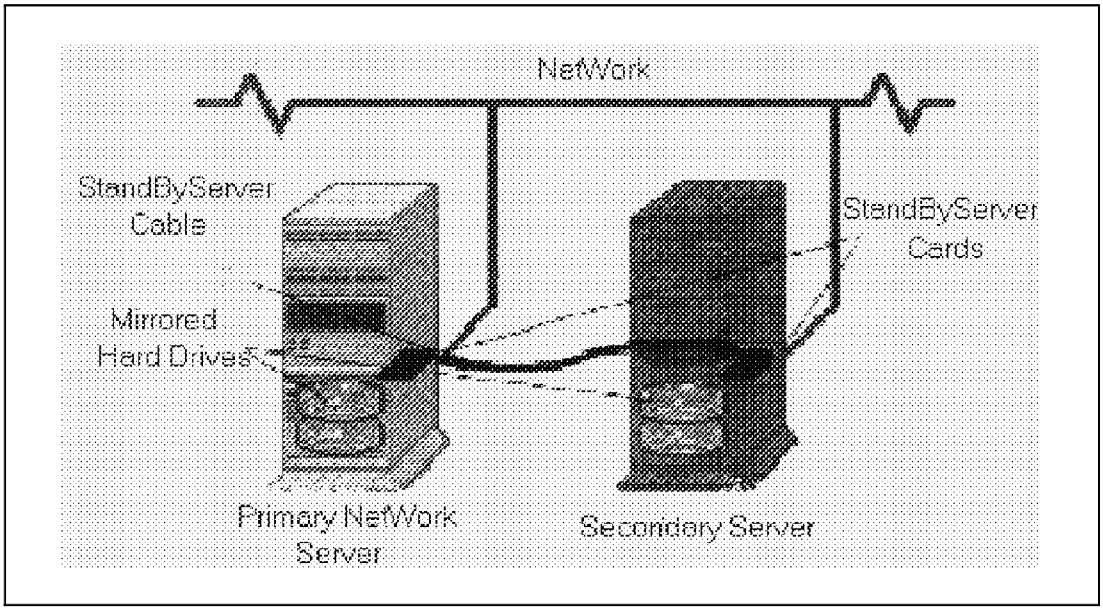

Figure 57. Vinca Fault Tolerance in OS/2 Warp Server

It gives you the possibility to connect a hot online replacement server to the main server that takes over instantly when the main server fails. Data is automatically mirrored between the machines via a high-speed dedicated link. *Vinca StandbyServer* for OS/2 Warp Server was developed specifically for mission-critical application servers. It ensures data availability in the case of software or hardware failure. *Vinca StandbyServer* automatically switches to the standby machine when the main server fails, no manual intervention is required.

So, when OS/2 Warp Server, plus Vinca's software solution, is running on a standby server and the main server fails either from hardware or software corruption, LAN Server is started. This guarantees that the same NetBIOS names are used; so it is not possible to have them temporarily up because there will be duplicate NetBIOS names. It also works with IBM's TME 10 NetFinity systems management software to send alerts to a local or remote client workstation; so network managers using TME 10 NetFinity can easily monitor the following information from anywhere in the network:

- Status of connection of the main server
- Disk drive identification and capacity
- Disk requests on the standby machine
- Real-time data indicating if and when the standby machine was loaded and how long it was connected to the main server

Vinca StandbyServer features and benefits for OS/2 Warp:

- All data is fully protected and is available through a redundant file server.
- With mirroring information being passed over high-speed Vinca components, the system causes no network traffic.
- Switching to the standby machine is automatic with only a momentary delay for users. This is because the connection information is stored in the client workstation without the need of a new login.
- It does not require identical servers.
- Customizable remote notification features are compatible with TME 10 NetFinity.
- Compatible with:
  - OS/2 2.x
  - OS/2 Warp 3
  - OS/2 Warp Connect and OS/2 Warp 4
  - OS/2 Warp Server

StandbyServer for OS/2 Warp was designed to take advantage of OS/2 Warp Server mirroring or Vinca's mirroring software that is included with the product. StandbyServer extends these mirroring capabilities to a second machine. For OS/2, this second machine appears to be nothing more than an additional storage resource similar to an external disk subsystem. The standby constantly monitors the primary server, if the primary server fails due to hardware or software error, the standby machine automatically initializes and takes over as the primary server. At this point, the mirrored data on the standby machine is then available to users on the network.

StandbyServer for OS/2 Warp can include two high-performance interface boards, a cable, Vinca's software, and IBM's NetFinity, or it can be purchased as software only. An interface board is installed in both the primary server and in the standby machine. Both machines are connected to the network wiring, but all data mirroring activity happens over a high-speed link capable of transferring data bi-directionally without any additional traffic being added to the network. Both machines must have OS/2 loaded and running. If the primary server is used as a file and print server running Warp Server, then Warp Server must be installed on the standby machine but configured not to autoload. Any additional applications that will be running on the primary server must also be installed on the standby machine, but they have to be configured not to start automatically. StandbyServer device drivers are loaded on both machines. In addition, Vinca's application software is also loaded on the standby machine. With these hardware and software components installed and configured,

mirroring can be set up on the primary server using OS/2 Warp Server's FTSETUP utility. Any data written to the primary server is also written to the standby machine.

10.1.1.1 Installation of Vinca StandbyServer for OS/2 Warp

The basic installation of Vinca's StandbyServer for OS/2 Warp is a simple step-by-step process, each of the following steps must be done on both the machines:

1. Install StandbyServer hardware and create a physical link between the primary and standby machines. This requires the installation of a network interface card in each machine, plus compatible, dedicated cable between the machines.
2. Install or modify OS/2 and ensure that OS/2 is set up on both primary and standby machines to recognize the disk device configuration required for StandbyServer.
3. Install the operating system on the primary and standby machines with the appropriate user accounts, domains, replication services, and machine names.
4. Install StandbyServer software on both the primary and standby machines.
5. Install TME 10 NetFinity Manager on the standby machine and TME 10 NetFinity services on the primary machine to facilitate automatic switching of the server function between the primary and the standby machines.
6. Install mirroring

   At this point you can test the system under controlled conditions to ensure that it works properly and that data is accessible after a system failure.

Vinca's Design

All requirements for OS/2 Warp Server mirroring or Vinca's mirroring must be met when installing and using StandbyServer for OS/2 Warp. StandbyServer for OS/2 Warp does not mirror the first physical drive in each system. This allows different system configuration files to be kept on each system so that different hardware configurations may be used for the primary and the secondary machine. The first drive in each system needs only to be large enough to hold the system software and OS/2 Warp.

10.1.1.2 Vinca Requirements Package

1. Two Vinca adapter cards.
2. All necessary software.
3. Server running OS/2 2.x, Warp, Warp Connect, Warp Server.
4. Standby machine running OS/2 2.x, Warp, Warp connect, Warp Server.
5. There must be at least two disk devices in each machine.
6. Cable between 25 and 50 foot, 20 meters, maximum distance between the two machines.
7. The two servers do not have to be identical, but both must capable of running OS/2 2.x, Warp, Warp Connect, Warp 4, and Warp Server.

Table 16. Models and Specifications for Vinca

Model #	Host Bus	Primary OS	Secondary OS	Max Servers Distance	Autoswitch
SBOS-E	EISA	OS/2	OS/2	50' (20 meters)	YES
SBOS-M	MCA	OS/2	OS/2	50' (20 meters)	YES
SBOS-E/M	EISA/MCA	OS/2	OS/2	50' (20 meters)	YES
SBOS-1	ISA	OS/2	OS/2	50' (20 meters)	YES

10.2 Fault Tolerance in Windows NT Server

Even if you make regular backups, it can happen that you have to go through the time and aggravation of finding your last backup and restoring the data from tape, and the last work from your backup will never be recoverable. A disk error on a file server is different.

1. Several different users will be saving their data to the disk, multiplying the amount of data you will have to reconstruct, because it was created after backup.
2. Some users will not be technical enough to know exactly what they have done that day.
3. It could be that some data is really irreplaceable.

Windows NT has several Fault Tolerance integrity features designed to allow your server to continue running right through a disk error, up to and including a drive crash.

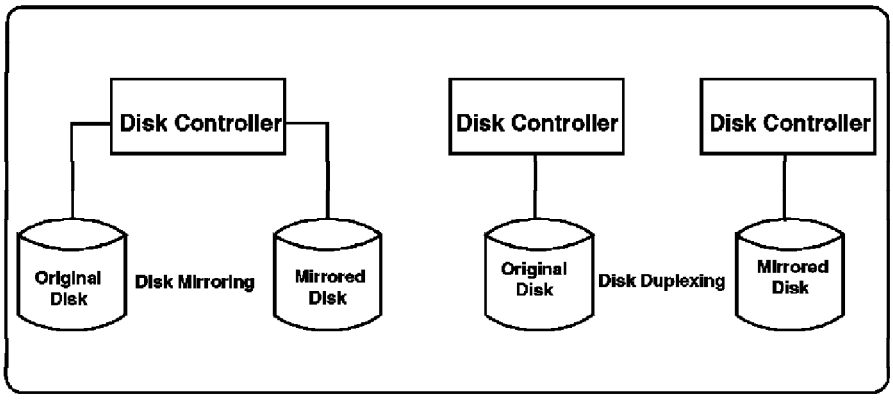

Figure 58. Disk Mirroring versus Disk Duplexing

These fault tolerance features are provided with Windows NT:

- **Data Mirroring**

  Data mirroring is the simplest of NT Server's fault tolerance features that is designed to manage with a failed disk drive.

  - If you configure two 500 MB disks to be mirrored pair, they will from that point look to users like a single 500 MB drive. Windows NT Server will duplicate all disk writes so that it writes the same data to both drives.
  - When one of the drives has an error reading some data, NT Server will simply read the data from the other disk and place a message in the Event Log. Users at their own workstation won't even know that something went wrong, the other file server will use the other drive.
  - Windows NT Server data mirroring does not require you to mirror the entire disk drive. It is possible to mirror partitions. This gives you the flexibility of having different-sized drives. You just have Windows NT mirror, what can be mirrored disk-space wise. The remaining hard disk space can be defined as an additional partition so that you can use the extra space.
  - If possible, you should put the two partitions of a mirrored pair on drives that are connected to different disk controllers or SCSI host adapters. Using multiple controllers protects your data from:
    - Drive Failure
    - Controller Failure

10.2.1.1 How to Create a Mirrored Set

1. Use Disk Administrator to create the first partition of the pair.
2. Select that partition, and CTRL-click on an unused area of disk space on another disk drive so that both the unused disk area and the first partition are selected.
3. Select **Establish Mirror from Disk Administrator's Fault Tolerance menu** Disk Administrator will create a disk partition out of the selected free space.

This new partition will be the same size as the original partition and will contain a copy of any data on the original partition. If one of the drives that you are mirroring fails, or if you need to add additional drives to your server, or if you need to rearrange your partitions, it could be necessary to break one or more of your mirrored sets. Breaking a mirrored set does not destroy the data on any partition, it only stops the process of duplicating data between the two partitions.

10.2.1.2 How to Break a Mirrored Set

- Select **Disk Administrator**
- Select **Fault Tolerance - Break Mirror**
- Confirm your decision of breaking a mirrored set

  Breaking a mirrored set reduces the data integrity of the server so that you are prompted to confirm your decision with an informal message appearing on your screen: `This will end mirroring and create two independent partitions`.

  Note: If you break a mirror set that is working perfectly, each half becomes a primary partition with its own drive letter. The original partition keeps its original drive letter, and the backup partition gets the next one available.

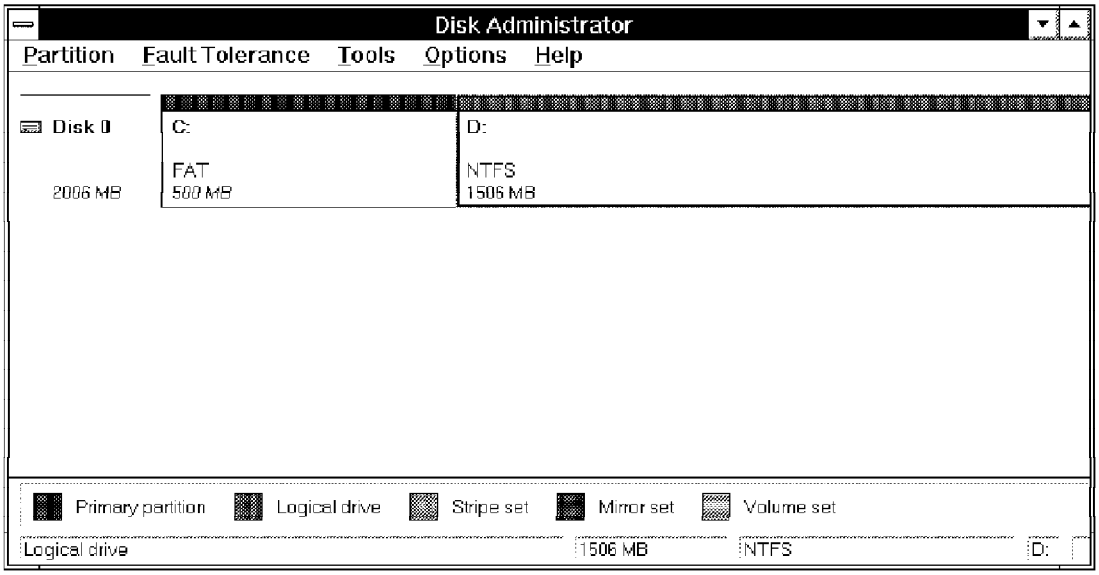

Figure 59. Mirrored Set

10.2.1.3 Things that are Important to Remember

1. Mirroring to drives from the same device controller does not protect your data from drive controller failure. If there is any kind of failure, you will just get your data back if you are mirroring to a disk run from a separate controller.
2. For higher disk-read performance and greater fault tolerance, use a separate disk controller for each half of the mirror set.
3. Disk mirroring effectively cuts your available disk space in half. Don't forget that when you are figuring out how much drive space you have on a server.
4. Disk mirroring will slow down writes because the data must be written in two places every time, but will speed up reads because the I/O controller has two places to read information from.

10.2.2 Vinca StandbyServer for Windows NT

StandbyServer for Windows NT is a fault-tolerant server-mirroring system that gives you the ability to connect a warm online secondary server directly to the main file server. Data is automatically mirrored to the standby machine over an industry-standard dedicated link that does not add traffic to the network. When the main server fails, StandbyServer for Windows NT switches automatically to the secondary machine. Users experience no downtime, and they don't have to attach the network again.

StandbyServer for Windows NT is a transaction-based server fault-tolerant solution specifically designed for mission critical-application servers. It ensures data diagnostic availability in the case of software or hardware failure while preserving diagnostic failure information in the primary server. StandbyServer for Windows NT connects two servers with a standard, dedicated, high-speed point-to-point link. By installing StandbyServer for Windows NT, the network operating system can access disk drives located in a machine other than the primary server, as you can see in Figure 60 on page 188.

Windows NT reads and writes to these remote disks as if they were located in the primary server. The network administrator can define which disk space is mirrored; all disk space does not need to be mirrored to the StandbyServer machine. The StandbyServer machine is constantly monitoring the status of the primary machine. If the primary server fails, the Vinca software running on the standby machine recognizes the failure and automatically initializes the second one as the primary one. Network services are automatically restarted on the new primary server, and any server-based application can be automatically reloaded as well. The entire switchover typically takes less than one minute, all user account information and defined shares are transferred and reestablished on the new server. Users can continue to access network data and services.

The StandbyServer for Windows NT is a software product. The software consists of NT services, drivers, and utilities that help you manage the installation, management, and carrying out the functionality of the product. Operation of StandbyServer for Windows NT requires that a dedicated communications link be installed between the primary server and the standby machine. The link can be implemented using a pair of industry-standard Network Interface Cards (NICs) or, alternatively, a pair of Vinca-proprietary boards. The standby machine must be equipped with an additional disk drive besides its own system disk. This will be the drive onto which the primary server's data will be mirrored. The drives do not have to be identical.

10.2.2.1 Installation

The installation is very easy. You have to install Windows NT Server on both machines; they must be both preconfigured with memory, Network Interface Cards, disk host adapters, and disk drives. Connect the standby machine to the same LAN as the primary server. You have to physically install one of the industry-standard Network Interface Cards (or Vinca-proprietary alternative) in each of the two machines. Connect them with the cable. This is the dedicated link over which disk mirroring and monitoring of the primary server will occur. Install the Vinca StandbyServer software on both machines. Then run the StandbyServer configuration on

the primary server. If the connections all are correct, the program will allow configuration of the StandbyServer for Windows NT application. At this point, the mirrored drive of the standby machine is "exported" by the standby machine and "imported" by the primary server. The next step is to define under what conditions a switch-over should occur and following what time-out duration. Mirroring can now be configured using Windows NT graphical disk administrator, which takes advantage of Windows NT fault-tolerant capabilities. It is Windows NT that will cause the mirroring, not the StandbyServer software. Vinca's StandbyServer software will redirect mirrored data to the standby machines drives. The mirroring will start as soon as the configuration part has ended.

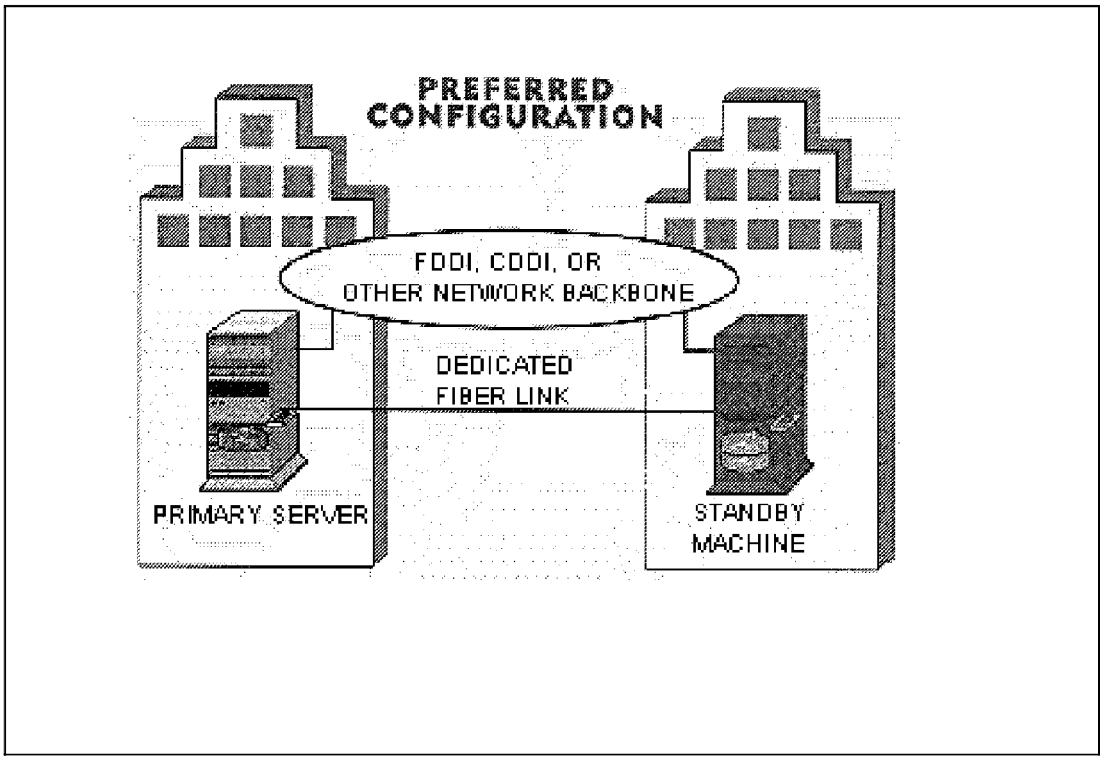

Figure 60. Vinca for Windows NT

10.2.2.2 StandbyServer Features and Benefits for Windows NT

- All data and server functions are fully protected and made available through a redundant server.
- Supports notification and monitoring features of third-party network.
- Does not require identical servers.
- Operates a simple device driver. Does not require any modifications to the network operating system.

Requirements Package:

1. Two servers running Windows NT Server 3.5 or above. Hardware does not need to be identical
2. All necessary software
3. StandbyServer for Windows NT users
4. Two Network Interface Cards to function as high-speed data transfer link between the servers
5. Dedicated cabling to connect servers

Table 17. Models and Specifications for Vinca

Model #	Host Bus	Connector	Cable Type	Max Servers Distance
PILA8465B	PCI	RJ-45	CAT 5 UTP	250' (100 meters)

10.3 Fault Tolerance System in NetWare 4.1

Novell's System Fault Tolerant Level III, known as SFT III, mirrors file servers. SFT III guards the file server from potential loss of file server operation. It separates hardware-related operating system functions on the mirrored servers so that a fault on one hardware platform does not affect the other. It is a server operating system designed to work in tandem with two servers. Both servers receive all updates through a special link called Mirrored Server Link (MSL), using adapters (hardware) and drivers that are specially dedicated for this purpose. MSL provides a physical link between the servers at high data-transfer rates. The servers also communicate over normal LAN communications they share in common, so that one knows if the other has failed. This happens even if MSL has failed. When a failure occurs, the second server automatically takes over without interrupting communications in any end-user-detectable way. SFT III also uses redundancy to ensure reliability. Each server monitors the other server acknowledgments to see that all requests are serviced and that all Operating Systems are constantly maintained in a mirrored state.

NetWare 4.1 provides two features that protect your server from the data loss and interruption that can result from disk failure. *Hot Fix* is an automatically implemented feature that protects against small-scale and gradual disk failure, and disk mirroring and disk duplexing are the optional features that protect your server from a complete disk failure.

Hot Fix: Every physical disk write can be read-after-write verified. If a block of data written to the disk cannot be read back and verified accurately, a media defect is detected. When this problem occurs, the data in the dirty

cache buffer is written to a special volume reserved specifically for this purpose. The Hot Fix volume is allocated during installation for each NetWare volume and serves as a safety net for file blocks written to bad sectors on a disk surface. A block is the standard unit of information that NetWare 4.1 reads or writes to disk. You choose a block size of 4 KB, 8 KB, 16 KB, 32 KB, or 64K during installation. The NetWare 4.1 Hot Fix feature is designed to minimize the problems caused by this gradual decline. After writing to a disk block, NetWare or the disk controller reads the block to determine whether the information was written accurately. If the disk block fails the test, the block is marked up as unusable, and the data still in memory is redirected to a block in the Hot Fix redirection area.

Disk Mirroring/Duplexing: Disk mirroring provides mirrored physical write requests to redundant drives. When two drives are mirrored, all files stored on one drive are also stored on the second drive. This activity is directly done by NetWare server operating system and is completely transparent to applications and to end-users. All configuration is done at the server; users can see no evidence that mirroring is in process, and in case of disk failure there will be no loss of data or downtime. When the drive will be replaced and repartitioned, NetWare will automatically remirror the active drive to the new drive in a background mode. User requests are processed as normal, with remirroring in the background.

Once mirroring is installed, a low-priority background copy occurs, with no delay of user I/O requests. No NetWare screen information indicates any difference between mirroring and duplexing. If the two drives are physically connected to the same adapter, they are mirrored; if they are connected to separate adapters, they are duplexed. Disk mirroring is simple to implement. The following procedure demonstrates the process of adding a second disk and mirroring:

1. The disk physically needs to be installed.
2. Load `INSTALL.NLM` and create a NetWare partition.
3. The two drives have to be paired.

Mirroring commences in the background. Check the server console screen for messages. It may take a long time to complete the mirror copy, especially if the server is busy. Remirroring does not affect user disk access because it is assigned a very low priority. When the drives are synchronized, a message appears at the console indicating that the drives are synchronized. A problem of Disk Mirroring means that when a drive goes down, you will probably not detect it; so we suggest you frequently monitor your server console or use a third-party utility that enables remote monitoring. In disk mirroring if the controller fails, both drives are down.

This is why duplexing is more reliable. When a controller fails it could be due to different problems that are intermittent. Controller problems have been known to do the following:

- Deactivate drives
- Lose volumes
- Corrupt data
- IDE drives

Disk Mirroring Failure

If a disk controller fails, before destroying the data on the disk, you can try to install a new identical disk controller. Your data may still be intact. Mirroring IDE drives may not provide the type of fault tolerance you expect. Under normal conditions, when one mirrored IDE drive fails, both drives go down. If this problem happens to you, you must separate the drives, and bring up the remaining good drive as a single drive.

Disk Duplexing means the two physical drives are located on separate controllers or host adapters. It is more reliable and more fault tolerant than mirroring. The only difference in implementing disk duplexing is the fact that the drives are connected to separate adapters. Duplexing adds more redundancy with duplicated disk adapters and adds an additional element of disk I/O performance. Because two drivers must be loaded, one for each disk adapter, NetWare automatically splits, enabling concurrent read/write access to both drives. For duplexing to improve performance, at least one of the two drive subsystems must be SCSI. SCSI drives have the capability to physically read and write at the same time. Since each disk has it's own controller, read and write procedures are performed simultaneously.

All that is required to both mirror or duplex disks is that the two drives have to have the same amount of disk space remaining after partitioning and configuring Hot Fix. The option to turn mirroring on is a menu selection in the server `Install` utility.

10.3.1 Vinca StandbyServer 32 for NetWare

StandbyServer 32 for NetWare, Vinca's hardware/software solution, is very similar to Novell's SFT III, but without some hardware restrictions. Using a high-speed connection between two servers, StandbyServer 32 uses Novell's disk mirroring feature to mirror the drives of a primary server to a secondary, standby server. StandbyServer then creates a copy of the primary server, including the Bindery, NDS, and other, subsequent data onto the secondary server, creating a sort of twin.

StandbyServer 32 is a fault-tolerant, server-mirroring system that gives you the ability to connect a warm online secondary server directly to the main server. Data is automatically mirrored to the standby machine via a high-speed, dedicated link that does not add traffic to the network. All data is written to the disk subsystem on the standby machine using standard NetWare mirroring. There is a current copy of all data on the standby machine at all times because StandbyServer 32 is a real-time disk mirroring solution with no latency, not a file copy solution. When the main server fails, StandbyServer 32 switches automatically to the secondary machine. Users are back online in a matter of minutes. The switch, however, is not totally transparent to the end-users because the secondary server must reset itself to become a primary server, clients must reboot to reconnect. Users logon and access data as they usually do. StandbyServer 32 is a complete NetWare-to-NetWare solution that takes advantage of all the benefits of this market-leading 32-bit operating system on both machines, including performance, reliability, and multitasking.

StandbyServer 32 provides automated recovery from any single failure, software or hardware. With StandbyServer 32, you can immediately recover from a software crash while preserving the fault on the primary server for diagnosis. The secondary server does not have to consist of identical hardware; it only requires that disks are greater than or equal to the one on the primary server and that the servers have similar buses.

Vinca's new Autoswitch feature takes care of swapping in the standby server without intervention from the network administrator. In case the primary server fails, Autoswitch detects the failure and automatically brings the secondary server up as the primary server. StandbyServer 32 connects two NetWare servers with a high-speed, point-to-point communication link. All data that is written to the primary disk subsystem on the standby machine using standard NetWare mirroring.

There is a current copy of all data on the standby machine at all times because StandbyServer 32 is a real-time disk-mirroring solution. StandbyServer 32 provides a monitoring screen on the secondary server with information regarding connection status, CPU utilization, and transaction processed.

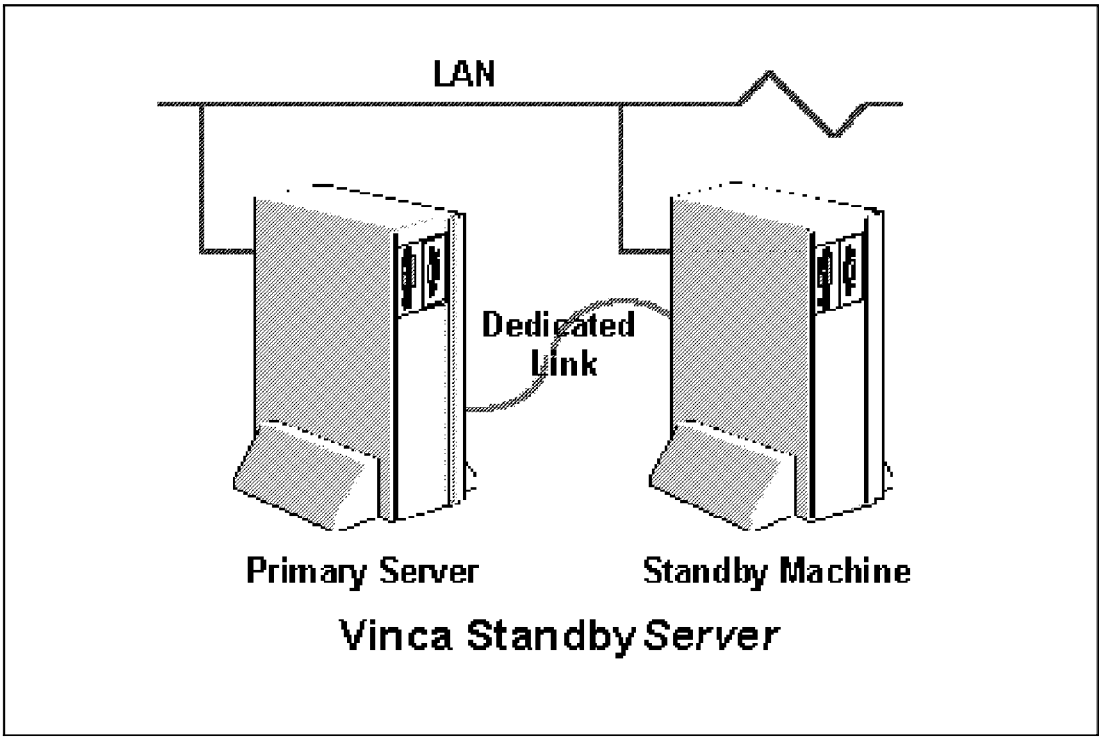

Figure 61. Vinca Solution for Netware

10.3.1.1 StandbyServer Features and Benefits

- All data is protected and made available through a redundant file server.
- Automatically switches to the StandbyServer machine when the main server fails.
- Does not require identical servers, allowing companies to extend the life of older equipment.
- Data travel's over Vinca's high-speed link; no traffic is added to the network.
- Dedicated cabling to connect servers.

Requirements Package:

1. NetWare 3.11, 3.12, or 4.x file server.
2. Server running NetWare or Runtime NetWare, provided with StandbyServer 32.
3. The two servers do not have to be identical, but they must have Novell-certified adapter cards.

4. The devices in the standby machine must be supported by NetWare drivers. It is not necessary to be identical to the device of the main server.
5. Two adapter cards.
6. All necessary software.

10.4 Comparison of Vinca's Solution

Table 18 (Page 1 of 2). Comparison Features

StandbyServer for	**OS/2 Warp Server**	**Windows NT**	**Novell 4.1**
Requires Identical Servers	Yes	Yes	Yes
Requires additional operating system license	Yes	Yes, requires only a one user license	No, NetWare runtime included with product
NOS compatibility	OS/2 2.11 and above plus LAN Server 3.0 and 4.0	Windows NT 3.51	NetWare 3.12 or higher and 4.x
Requires dedicated link	Yes	Yes	No, but it is highly recommended
Hardware configurations	IBM PCI 100/10 Ethernet adapter and cable	IBM PCI 100/10 Ethernet adapter and cable	Any card and cable using IPX protocols
Available now	Yes	Yes	Yes
Protects against hardware and software failure	Yes	Yes	Yes
Routable and bridgeable	Yes	Yes	Yes, only with StandbyServer 2.0 for NetWare

Table 18 (Page 2 of 2). Comparison Features

StandbyServer for	OS/2 Warp Server	Windows NT	Novell 4.1
Maximum distance between servers	100 meters or approx. 300 feet	100 meters or approx. 300 feet	StandbyServer 32/50 ft. Campus StandbyServer 2 kilometers StandbyServer 2.0 whatever distance support by cards and cable
Automatic client reconnect to server	Yes	Yes	Yes, if client is using the 32 bit Client from Novell
Handles open files	No	No	Yes, when using SnapShot server add-on module
Installs from one server/workstation	No	No	Yes, when using SBS 2.0
Standby machine can be used for other network services	Yes	Yes	Yes

10.5 Introducing Clustering Technology by IBM

Warp Server clustering, as demonstrated at PC Expo in June of 1996, demonstrated that clustering provides other key features beyond SMP capabilities. Clusters are collections of interconnected computers that appear to a user as a single system.

Clusters allow performance and capacity increases by simply adding additional servers to the cluster. Clusters facilitate the deployment of 7X24 business-critical applications. Clusters ease maintenance by keeping critical applications and data available during routine maintenance.

Because of the finely tuned file system that Warp Server Advanced utilizes, many applications that may reach an I/O bottleneck first (like file and print

applications) may get relief from a clustering solution since the CPU utilization of Warp Server is minimal in this environment. A scenario demonstrated at the PC Expo in June of 1996 is shown in Figure 62 on page 196.

SMP, in contrast, may yield little benefit for file and print solutions. Clustering, in addition to providing higher throughput, provides fault tolerance, a significant consideration depending on the size of the business and network availability requirements.

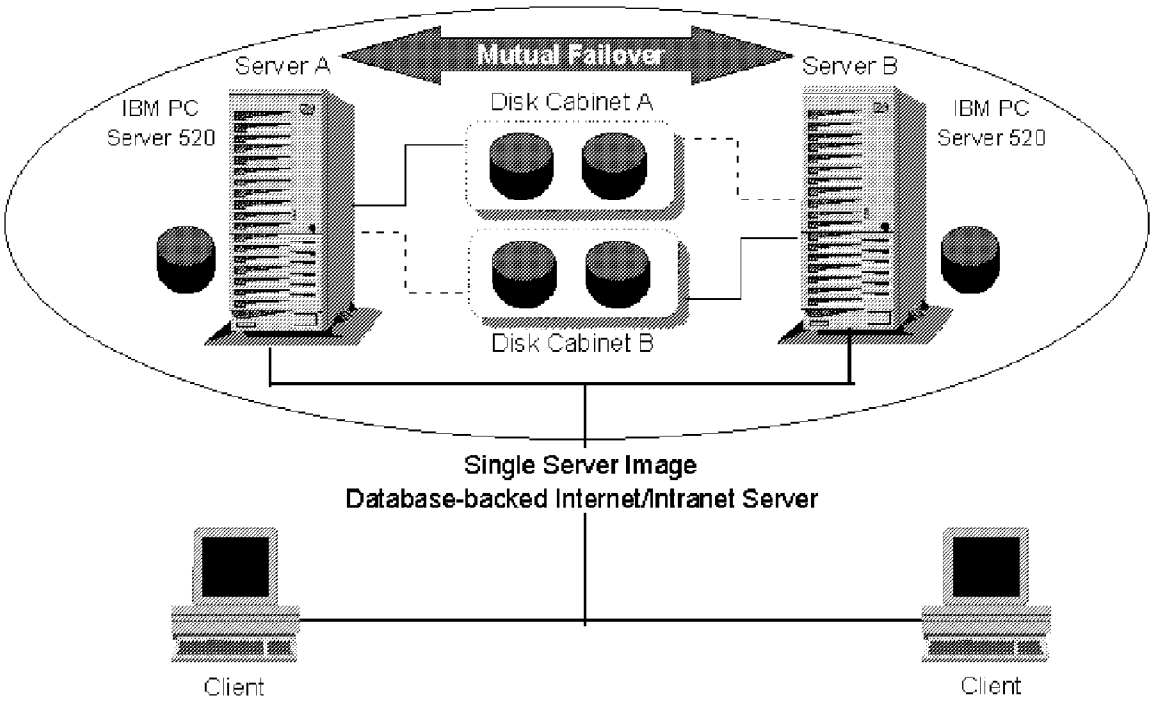

Figure 62. Warp Server Clustering Demonstration

10.5.1 Warp Server Clustering Demonstration — Key Features

1. Recovery from Application Failures

   - The cluster can monitor specific applications to detect failures. In the event of an application failure, the application can be restarted on the same or a different node of the cluster. Studies show that application failures account for the greatest portion of downtime. Detection and recovery of these types of failures minimizes this downtime and loss of productivity, especially when the application can be automatically recovered and reintegrated into the cluster.

2. Recovery from Node Failures

- The high availability cluster uses sophisticated heartbeat mechanisms to determine when one of the servers in the cluster has failed. When a node failure is detected, the remaining server automatically restarts the highly available application and makes it available to the client, with no end-user disruption in most cases. Detection and recovery of common system failures keeps customers' business-critical applications up and running.

3. Operator Initiated Failover
   - The system administrator can selectively initiate the failover of a highly-available application from one node to the other. This gives the system administrator more control over scheduled downtime for maintenance, effectively allowing real-time system maintenance and repair.
4. Reintegration of a Failed Node
   - Systems can rejoin the cluster after a shutdown or failover event. Highly available applications that were running on the node when it failed can be automatically migrated back once the node has rejoined the cluster. This allows the cluster to be returned to full productivity quickly and automatically.

Although this scenario was just a technology show, it proved that IBM is capable of providing clustering technology on the Intel platform. Clustering technology will be part of Warp Server in the fourth quarter of 1997. Network Computing (NC) clients will take advantage of high availability clusters as well.

Part 3. Configuring and Using Features

Chapter 11. Warp Server

This chapter provides an overview of the graphical user interface (GUI) for file and print that is included with Warp Server.

11.1 Warp Server Graphical User Interface for File and Print

This user interface is object-oriented to allow the user or system administrator to configure and use the Warp Server environment via the manipulation of visual objects. It is just as easy as drag-and-drop. The paradigm used has a consistent look and feel with the Workplace Shell (WPS).

Note: We have selected a few tasks to show you that are performed by using the GUI. For more information and step-by-step instructions for using the GUI, see the *Up and Running* documenting that comes with the product.

11.2 Where to Find the LAN Server Graphical User Interface

Besides the fact that the OS/2 Warp Server GUI is very memory-intensive, you typically do not start it directly from your server but instead from an administrators workstation. This not only gives you the freedom to choose to work, independently and away from where your server is located but also saves memory at the Warp Server, which is better used for its original purpose of providing file and print services.

In this chapter we used the LAN Server GUI provided with OS/2 Warp Version 4.0. If you are using the LAN Server GUI on a previous version of OS/2 or on the Warp Server itself, the location of the LAN Server Administration icon differs slightly but there is no difference in terms of look and feel and usage.

On the OS/2 Warp Version 4.0 you will find the the LAN Server Administration icon by double-clicking on the **Connections** folder. Move on to the Network folder, and double-click on the **Network Services Folder**. The Network Services folder contains all server-relevant icons.

The OS/2 Warp Center introduced with OS/2 Warp Version 4.0 gives you the opportunity to browse through your desktop by clicking on the OS/2 Warp wave logo. It will show you the content of your desktop and the contents of all folders located on the desktop. If you click on a folder within this list, it will show you the contents of this folder and so forth. This mechanism of the OS/2 Warp Center, originally invented by Lotus several years ago, fails

to show you the Network Folder. Therefore you will be able to browse any folder and subfolder of your desktop, except the Network Folder (see Figure 63 on page 202). There is a workaround. Just drag and drop the Network Services folder onto the desktop or create a shadow of that folder somewhere outside of the Network Folder, and you will be able to browse it via the Warp Center.

With OS/2 Warp Version 4.0, IBM also changed the naming of the Settings Notebook. It is now called Properties Notebook. If you see the term 'Properties' used hereafter and you are using an earlier OS/2 version just read it as a synonym for 'Settings'.

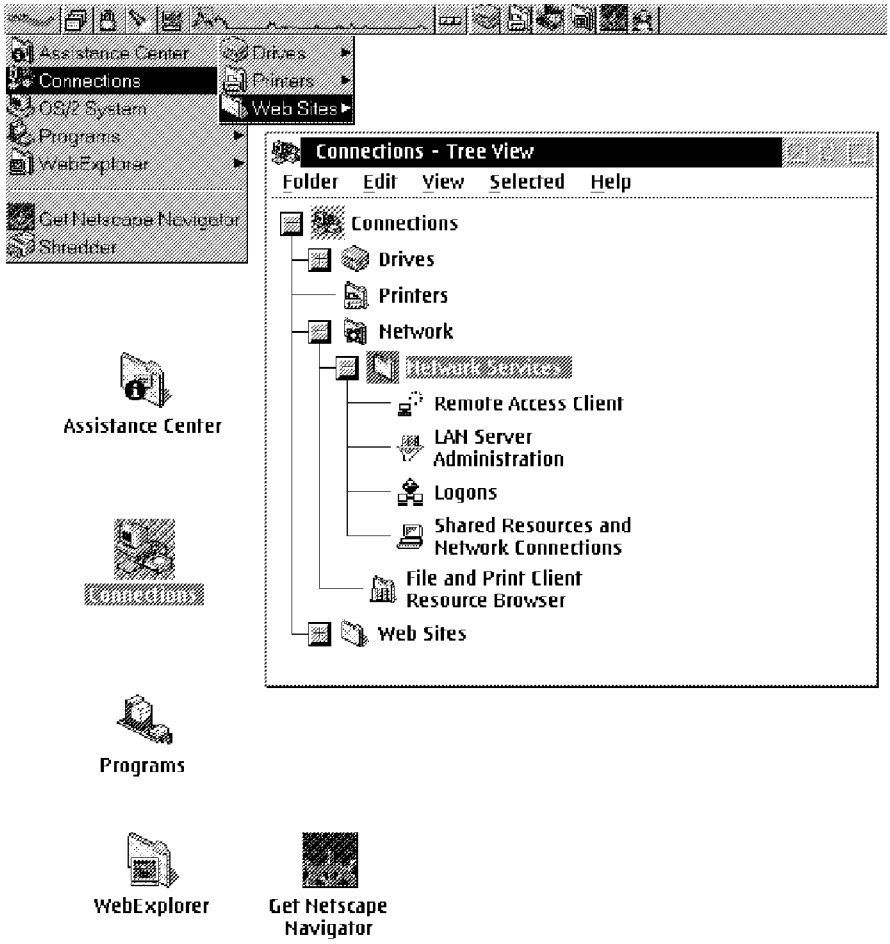

Figure 63. Location of the LAN Server Administration Icon

11.3 Drag and Drop of Objects

With the new object-oriented graphical user interface, the user is able to use the same techniques like, *drag-and-drop*, he/she already uses in other OS/2

applications and of course in OS/2 itself. Drag-and-drop, which draws on how humans ordinarily treat real objects, is a major function of graphical user interfaces in general. Warp Server gives you the choice of using the method you prefer, either the more state-of-the art drag-and-drop approach or the old-fashioned pull-down menus.

In the example shown in Figure 64, you can see two opened folders:

1. The LAN Server Administration folder
2. The domain contents, opened by double-clicking on the domain object (here named ITSCAUS)

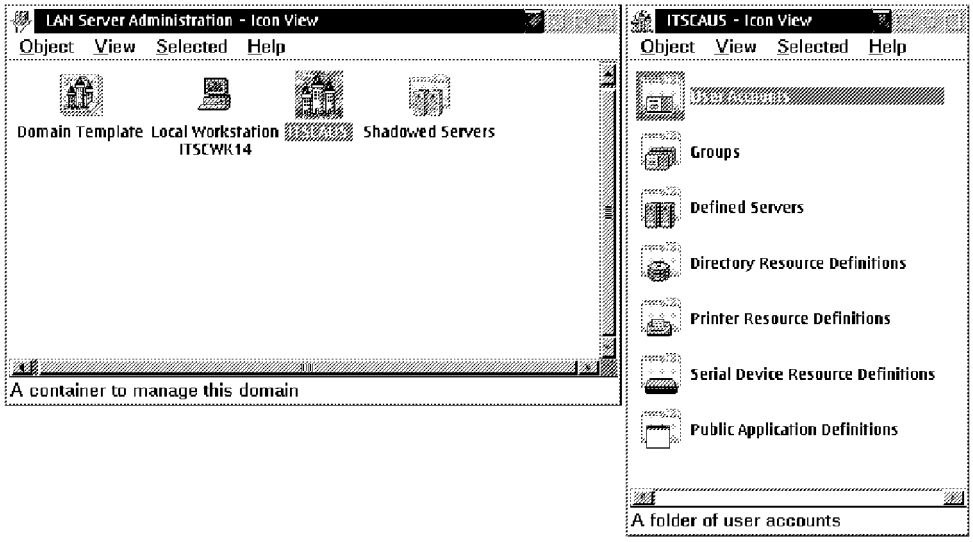

Figure 64. LAN Server Administration and Domain Contents

Figure 65 on page 204 gives you an example of a possible arrangement of the different folders you can open from the Domain Contents folder:

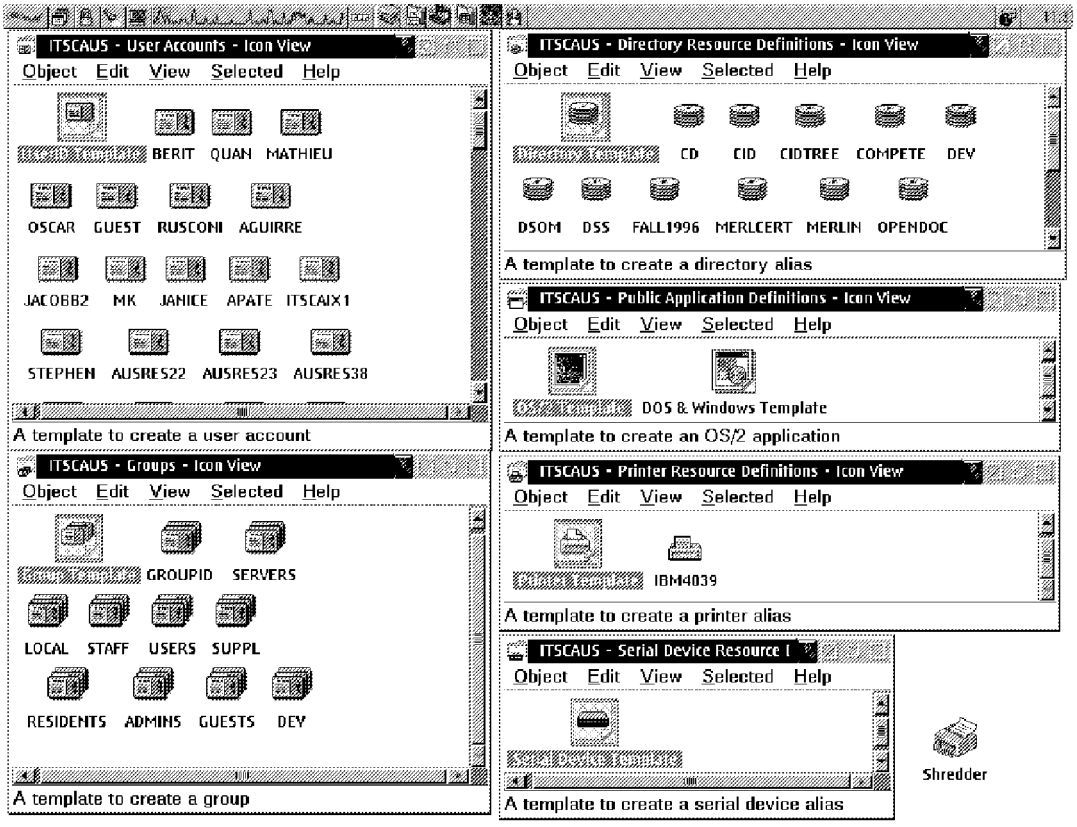

Figure 65. Arrangement of Domain Content Folder

- User Accounts
- Groups
- Directory Resource Definitions
- Public Application Definitions
- Printer Resource Definitions
- Serial Device Resource Definitions

11.3.1 How to Create a User ID

The Users Accounts folder contains at least one UserID-Template (see Figure 65). To create a new user just drag a copy of the User ID Template to an open area in the User Accounts folder (see Figure 69 on page 207). Immediately after releasing the right mouse button, the User Account-Create notebook is displayed, and you can begin to complete the fields.

11.3.2 How to Clone a User ID

The LAN Server graphical user interface offers two ways to clone a user ID:

- Create a new user ID by copying an existing user ID
- Create a new user ID by using an User ID template

To clone a user ID by using an already existing user account, do the following steps:

1. In the User Accounts folder, select the user account you want to clone, and press on the right mouse button to open the object's Context menu as shown in Figure 66.

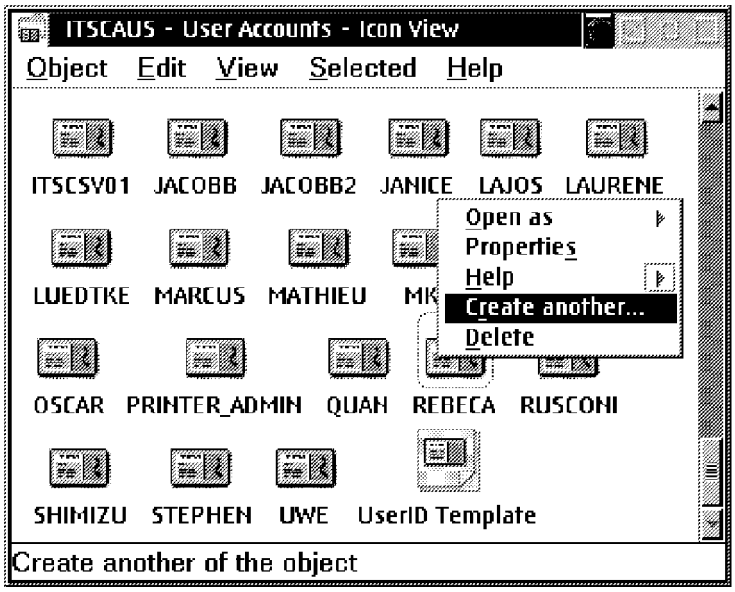

Figure 66. User Object's Context Menu

2. For example, select the user account **REBECA** and press the right mouse button. Select **Create another...**. The User Account-Create notebook will be opened for you.

   All logon information but the user account name and the password will be inherited from the original user account.

3. In the Identity page, type in the new user account name. The password must be specified in the Password page.

4. Adjust home directory information in the Home Directory page. By default, the newly cloned user account will have the same home directory location as the original user account.

To clone a user account by using an user ID-template, you first must have a user account with the Template checkbox marked, as shown in Figure 67 on page 206.

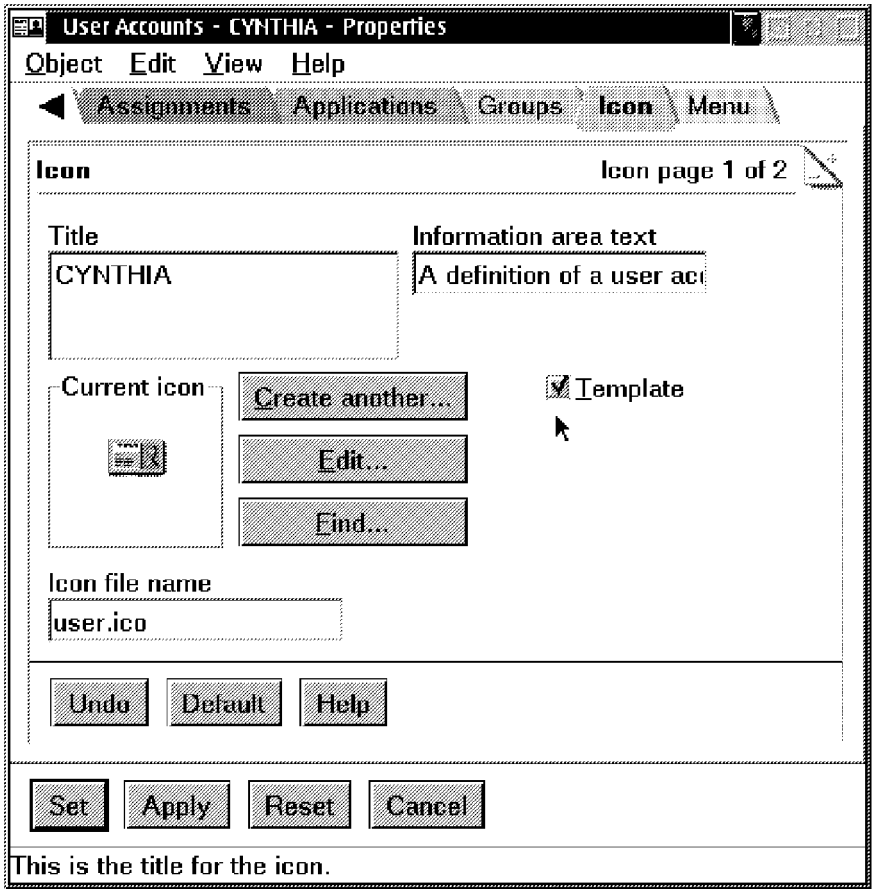

Figure 67. User Object's Notebook Icon Page

For example, the Template checkbox of the user account CYNTHIA is marked; therefore you will note a different user account object in the User Account folder. An example of a user ID template is shown in Figure 68.

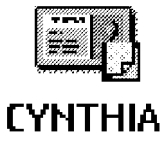

Figure 68. Template User Account Object

To clone a user account by using this user ID template, do the steps described in 11.3.1, “How to Create a User ID” on page 204.

A great advantage of cloning an user account by using a user ID template is that all information but the user account name is given to the newly created user account. That means password information is inherited as well.

11.3.3 How to Change a User ID's Attributes

To change a user ID's attributes, do the following steps:

1. In the User Accounts folder, select the user account whose attributes you want to change and double-click on the user account's object. The User Account-Properties notebook will be opened for you.
2. Make necessary changes and select **Set** to save changes and exit.

   Note: The Apply button saves the changes and keeps the notebook open for you.

11.3.4 How to Create a Group, Printer, Serial Device, and Directory Aliases

Using the same drag-and-drop concept used to create a user ID, you can easily create new groups with the Group Template in the Groups folder. The same procedure applies to new Printer, Serial Device or Directory aliases with the templates in the Resource Definitions folder (see Figure 69).

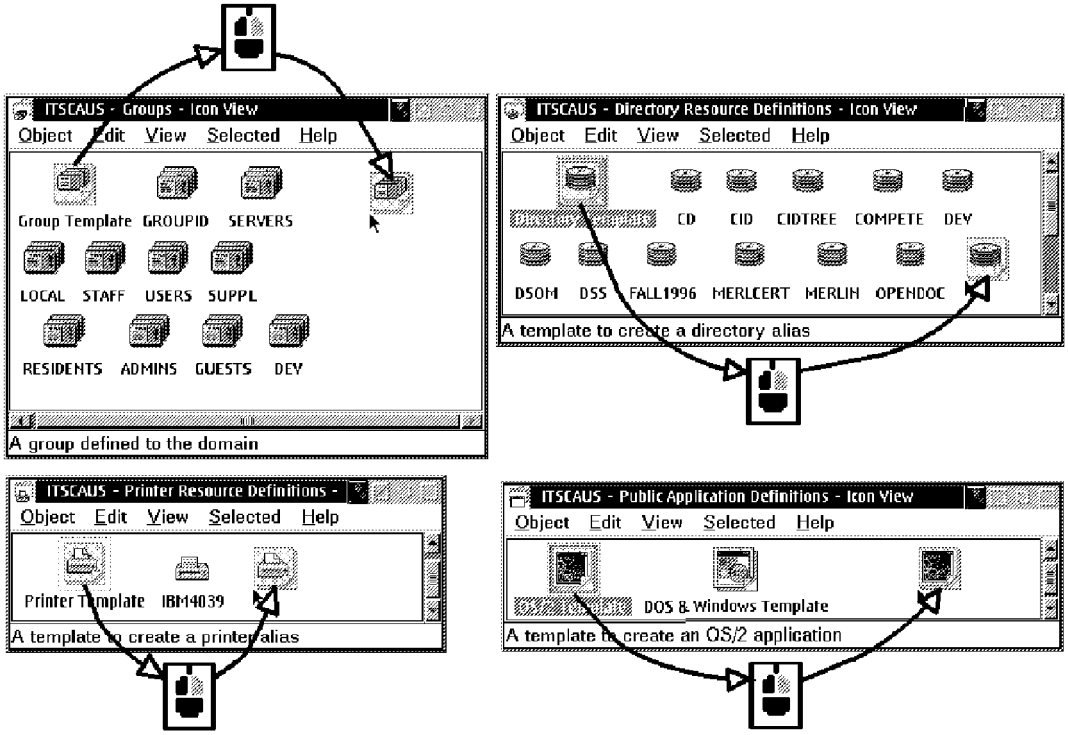

Figure 69. Creating Groups, Printer, Serial Devices, or Directory Aliases

11.3.5 How to Assign Users to a Group

In the Groups folder, we have created the Marketing group as shown in Figure 70 on page 208.

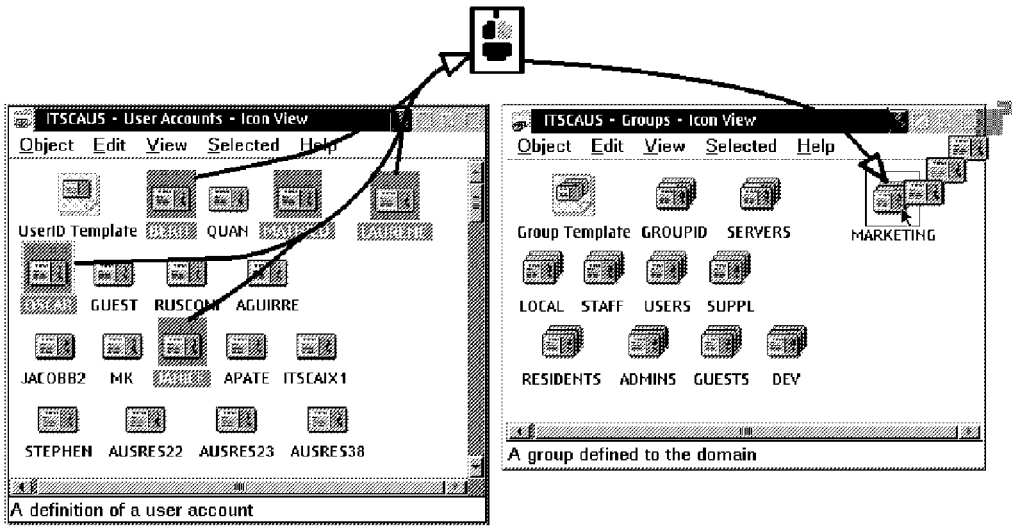

Figure 70. Assigning Users to a Group

If we now would like to include the users BERIT, MATHIEU, LAURENE, OSCAR, and JANICE in that group, we simply drag and drop each user object to the MARKETING group object. You can do this with one action by pressing the **Ctrl** key while selecting all users and then dragging and dropping them on the MARKETING group object.

11.3.6 How to Assign Logon Assignments and Access Controls to User Accounts and Groups

In the Resource Definitions folder, we have defined a printer resource called IBM4079. To assign that resource, for example, to the user ID LAURENE, simply drag and drop the IBM4079 object to the user ID object LAURENE. You will then have the chance to add or replace access permissions for the user in the Grant Access to Resource panel and also to select the logical port that will be assigned to this resource when the user logs on in the Administer Logon Assignments panel.

The result is that the user LAURENE now has the resource IBM4079 as a logon assignment. Similarly, you could drag this resource to *many* different users or drag *many* resources (that you have highlighted by pressing the **Ctrl** key while selecting) at once to a user.

You can also drag and drop that resource to one or more group IDs or vice versa (see Figure 71 on page 209).

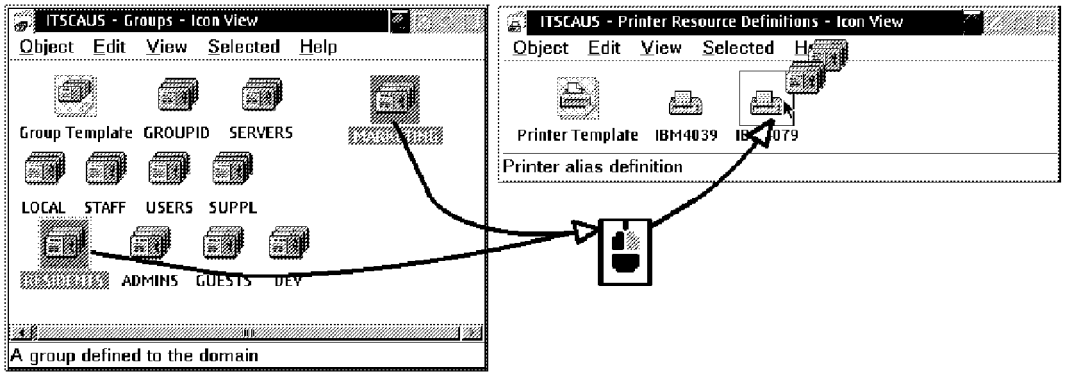

Figure 71. Assigning a Printer to Several Groups

The result is that the resource is then assigned to each user in that group.

Note: After dragging a resource to a group and making it a logon assignment, you are able to go back later and either change or delete the assignment for the group.

To do this, drag the resource back to the group (just as you did when you made the assignment), and select **Continue** in the Grant Access to Resource panel to go back to the Administer Logon Assignments panel, where you can change the assignment or delete it.

11.4 User Account Create Notebook

With OS/2 LAN Server 3.0, User Profile Management (UPM), an interface separate from the LAN Server full-screen interface, was used to create and manage users and groups. With LAN Server, you can still use UPM to manage users and groups. However, not all features that come with LAN Server will be offered in UPM. For example:

- Creating user IDs (type **User**, not Administrator) with specific privileges that allow:
  - Managing printer queues
  - Managing groups and users
  - Managing serial devices
  - Managing shared resources
- Assigning home directories to user IDs
- Administering logon assignments and public application assignments

If you want to get all LAN Server user/group management features in one interface, you should use the User Account Create notebook to manage users and the Groups Create notebook to manage groups. You will get this

notebooks by clicking on an existing user or group icon, or it will be automatically presented to you at creation time.

Figure 72 shows you what you can do with the User Account Notebook and gives you an idea of why it may suit your needs better than UPM for managing LAN Server users and groups.

Notes:

If you are using OS/2 Warp Version 4.0, all notebook tabs are placed at the top of the notebook. To go directly to a specific notebook you can either browse through the notebook pages by clicking on the appropriate arrow in the right- or left-hand corner, or just click with the right mouse button on any notebook tab. You will get a pop-up window showing you all tabs at a glance to select from (see Figure 72).

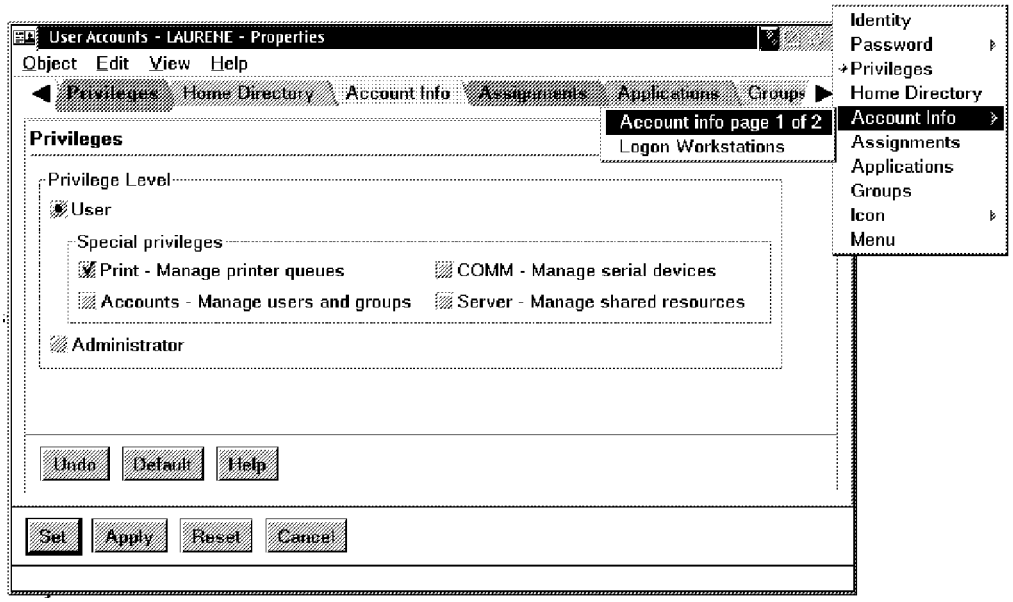

Figure 72. User Accounts Notebook

Within the User Account properties notebook, you can browse, create, delete, and modify the settings for a user account. These settings include the identity, password, account information, home directory, logon assignments, public applications, and groups that this user belongs to.

Notes:

1. The user account name can be up to 20 characters. However, if you have Database2 for OS/2 (up to Version 1.2) or Communications Manager/2 installed, or DBCS is installed on FAT partitions, you should make the User account name no more than eight characters for compatibility reasons.
2. If the user account name is more than 15 characters, it cannot be added to the network as a messaging name, and the user cannot send or receive messages.
3. In the User account name field, you cannot use the National Language Support (NLS) characters 130 though 139 (using the ALT-numbers technique), since the 13 is interpreted as a carriage return, and you are now taken to the next page. Also some characters (from 140 - 150) are not entered correctly.
4. Be aware that if you have LAN Server for Macintosh clients, there is currently a problem if you use names more than eight characters in length. If there are LAN Server for Macintosh clients in your network, do not use more than eight character names for your names.

In Figure 72 on page 210 you can use the following notebook pages to change or display different settings for a user account:

- Select the **Identity** notebook page to display or change identifying information about the user account.
- Select the **Password** notebook page to change the password for the user account.
- Select the **Privileges** notebook page to display or change information about the privilege levels.
- Select the **Home Directory** notebook page to display or change home directory information for the user account.
- Select the **Account Info** notebook page to display or change information about a user's account options.
- Select the **Assignments** notebook page to display or change the logon assignments assigned to the user account.
- Select the **Applications** notebook page to define public applications for the user account.
- Select the **Groups** notebook page to display or change the groups of which the user account is a member.

11.5 Logon Assignments and Logon Profiles

In this section, we discuss general logon assignment and logon profile considerations.

In some cases, especially in previous full-screen-interface LAN Server versions, administrators preferred using particular logon profiles for each user to assign printer aliases and file aliases. Since LAN Server now has a graphical user interface, there is not as much of a need for individual profiles. However, existing profiles will still run under LAN Server.

The DOS user's profile is named PROFILE.BAT, and the OS/2 user's profile is named PROFILE.CMD. Both files reside in the
IBMLAN DCDB USERS userID subdirectory on the domain controller.

A typical PROFILE.CMD for OS/2 users could look like Figure 73.

```
/************************/
/* User Profile in REXX */
/************************/

'@ECHO OFF'
trace o;

'NET USE LPT1: \\ITSCSV00\IBM4079 >NUL'
'IF EXIST X:\LANUSER.CMD CALL X:\LANUSER.CMD'

exit 0;
```

Figure 73. PROFILE.CMD for OS/2 Users

Note: REXX is only available to OS/2 and PC DOS 7.0 users.

A typical PROFILE.BAT for DOS users could look like Figure 74.

```
@ECHO OFF

NET USE LPT1: \\ITSCSV00\IBM4079 >NUL
IF EXIST X:\LANUSER.BAT CALL X:\LANUSER.BAT
```

Figure 74. PROFILE.BAT for DOS Users

One of the major advantages of having profiles is flexibility. However, with LAN Server interface, you have got at least the same flexibility now. If global changes have to be made, you simply can use the GUI.

Note: You cannot make network application assignments using PROFILE.CMD and PROFILE.BAT.

For example, let's say you would like to change printer assignments for the MARKETING group. The users belonging to the MARKETING group have the printer alias IBM4079 assigned to LPT1. Because the users now have local printers attached to LPT1, you want to assign the network printer to their logical LPT2 port.

You can do this as follows:

1. Drag and drop the IBM4079 printer object to the MARKETING group as you would assign the printer to the group. The Grant Access to a Resource properties notebook is opened.

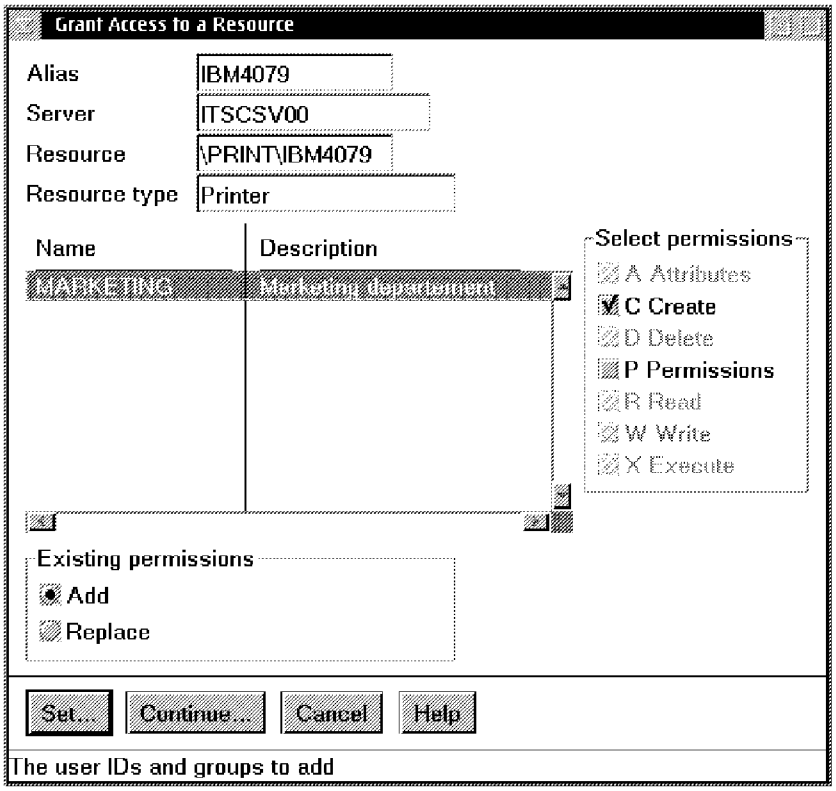

Figure 75. Grant Access to a Resource

   Select **Continue...**.

2. In the Administer Logon Assignments window shown in Figure 76 on page 214, you easily can make the change needed. In this example, select the **Add assignment** button.

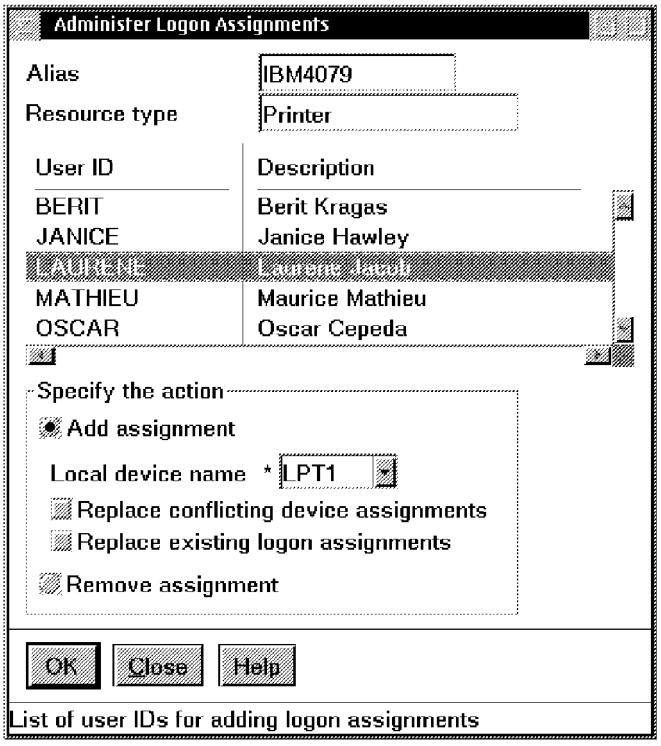

Figure 76. Administer Logon Assignments Window

3. Change **Local device name** from LPT1 to LPT2. Place a check mark in the appropriate box, for example **Replace existing logon assignment** and/or **Replace conflicting device assignments**. You also may delete the printer assignment if you need to do so.
4. Select **OK** to complete the change.

Using logon profiles, you would have changed the line for the printer for each user, or you would have changed a command in a file called by the logon profile (such as LANUSER.CMD and LANUSER.BAT, as shown in the figures on page 212).

11.6 Access Control Profile Creation

In previous IBM LAN Server versions, the administrator had to remember to define an Access Control Profile after the definition of an alias. Then, to propagate that profile down the directory tree, the administrator had to use the Apply function from the full-screen interface.

With LAN Server, you are prompted to create an Access Control Profile after the creation of an alias. After creating the profile, you are prompted to propagate it down the tree.

In Figure 77 on page 215, you can see the system tells you that an Access Control Profile does not exist for the resource you have just defined with an alias. By taking the default and clicking on the **OK** button, the window in Figure 78 on page 216 is displayed.

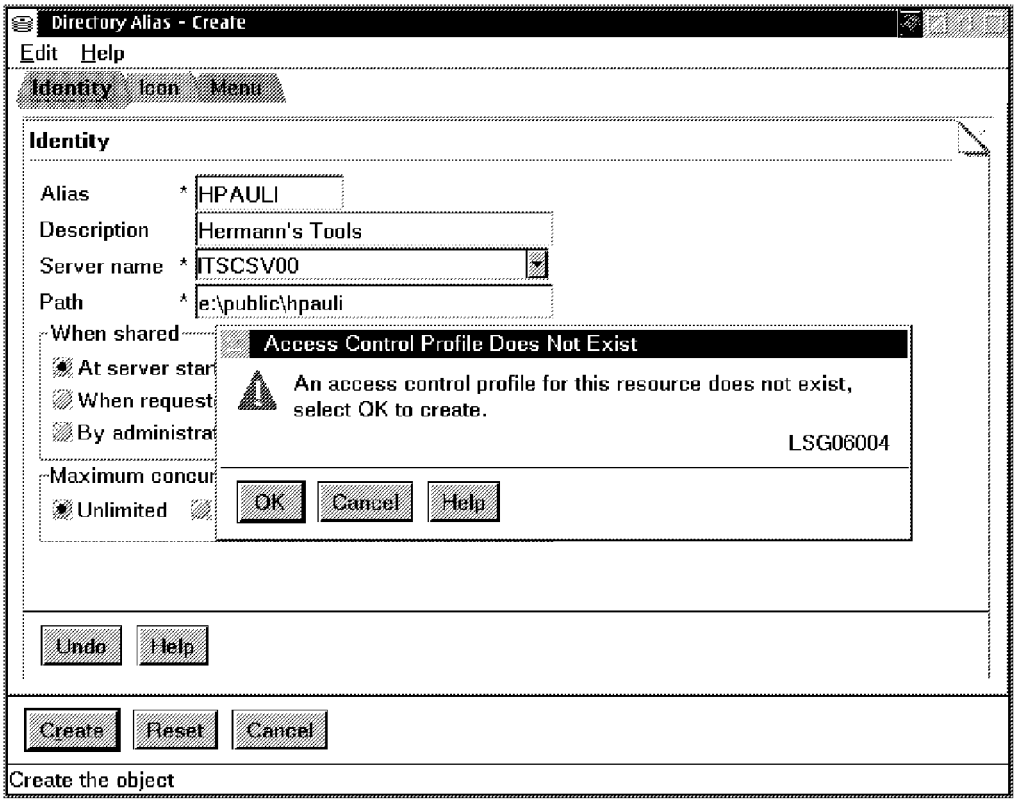

Figure 77. Access Control Profile Does Not Exist Window

In the Access Control Profile - Settings View notebook shown in Figure 78 on page 216, you can now set the permissions and the auditing settings for the alias.

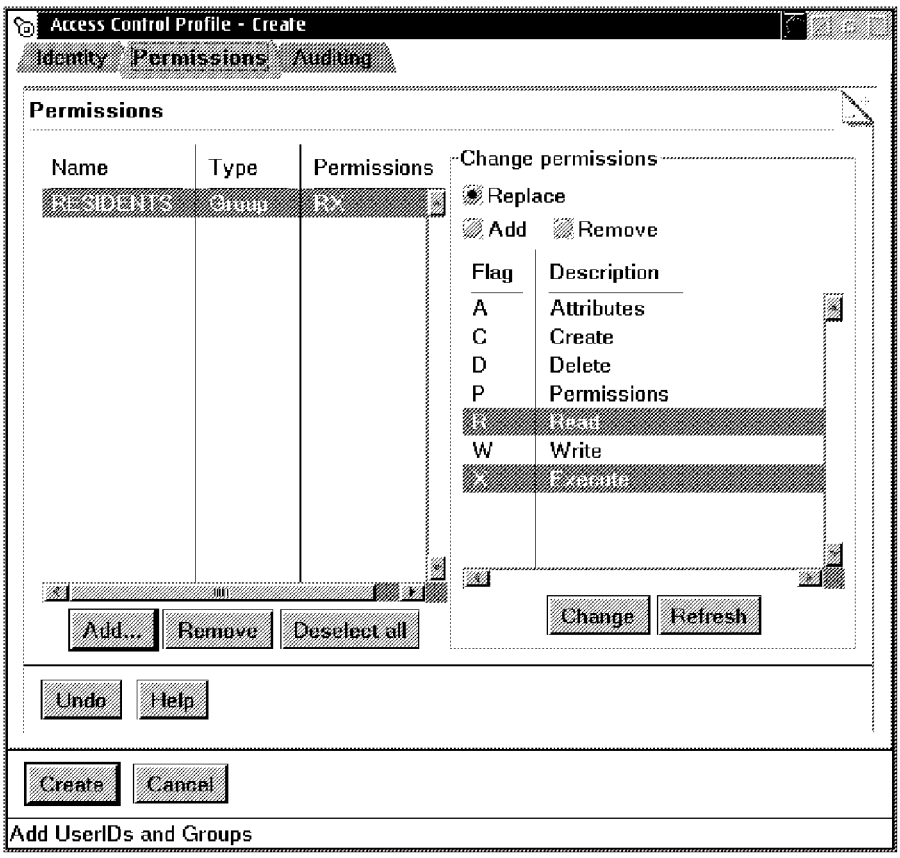

Figure 78. Access Control Profile - Settings View Notebook

After pressing the **Set** button, the Propagate Access Profile to Subdirectories window is displayed as shown in Figure 79 on page 217. By taking the default and pressing the **OK** button, the Access Control Profile is propagated to all of the resource's subdirectories.

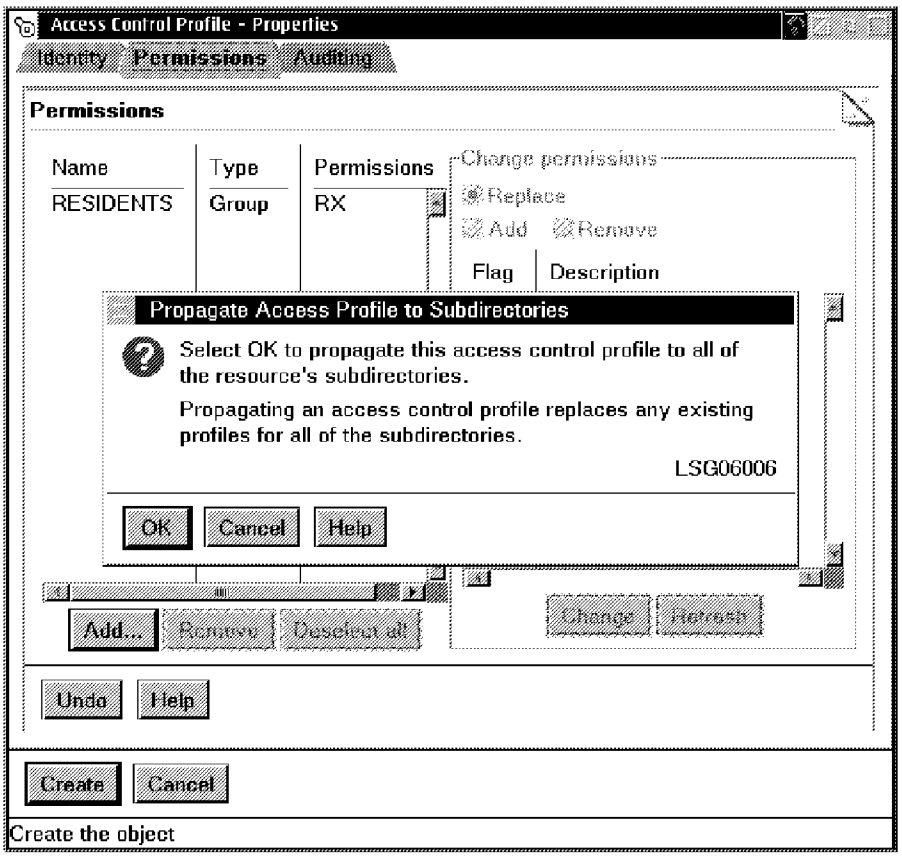

Figure 79. Propagate Access Profile to Subdirectories Window

11.7 Network Applications

This section provides examples of installing OS/2, DOS, and Windows network applications running on an OS/2 Warp Connect or OS/2 Warp 4 workstation.

A new feature of file and print services in Warp Server is the ability to represent the public and private applications by their actual product icons in a single folder named Network Applications (at the client's workstation). This is because the RASx.EXE programs are no longer used to invoke applications. Now, all programs are started directly by the actual application's executable file.

Also, you can now define DOS and Windows applications for DOS and for OS/2 users using DOS Templates. Since OS/2 reads information from the EXE header, it is able to start the correct environment, so that, for example, the execution of the Windows version of AmiPro really starts a Windows environment at the OS/2 client. Because applications icons are stored in

the EXE file as well, the real icons are shown in the Network Applications folder.

In the past, to define DOS and Windows application for OS/2 users, you had to create OS/2 command files, which were invoked from inside an OS/2 application definition. OS/2's ability to read the EXE header was used here.

Note: Make sure you fulfill all license agreements when you install applications on the server. Applications you share in your network must be enabled to run over a network.

11.7.1 Installing an OS/2 Public Application

The following steps show you how to install Lotus 1-2-3 for OS/2 as an OS/2 network application located at a server:

1. Make a decision where to install the application on the server.
2. Install the application on the server. Follow the instructions that come with the product. For Lotus 1-2-3, for example, you may use the path `D: 123G`.
3. In the LAN Server Administration folder, double-click on the **Domain** object.
4. Before you actually create a public application for Lotus 1-2-3, you first have to create a directory alias that points to the subdirectory on which Lotus 1-2-3 was installed (in our example it is D: 123G).
5. Double-click on the **Resource Definitions** object.

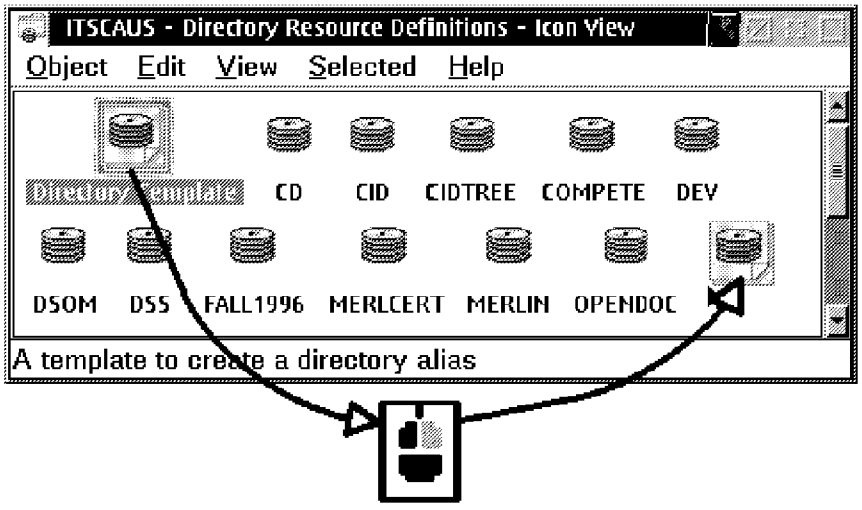

Figure 80. Resource Definitions Folder

6. In the Resource Definitions folder (as shown in Figure 80), drag and drop a copy of the **Directory Template** to an open area in the Resource

Definitions folder. The Directory Alias - Create properties notebook is opened, as shown in Figure 81 on page 219.

In the Directory Alias - Create properties notebook, complete the **Identity page**.

Consider license control. In our example, we purchased 10 licenses of Lotus 1-2-3. So the value in the **Number of connections** field should be set to `10`. LAN Server ensures that no more that 10 users will be able to start the application.

Note: Some network-enabled applications come with their own license control. If this is the case, carefully read instructions that come with the product. In most cases, it is enough to assign a so-called *work directory* (see Figure 86 on page 224) that points to the subdirectory in which license control files reside.

7. Select **Create**.

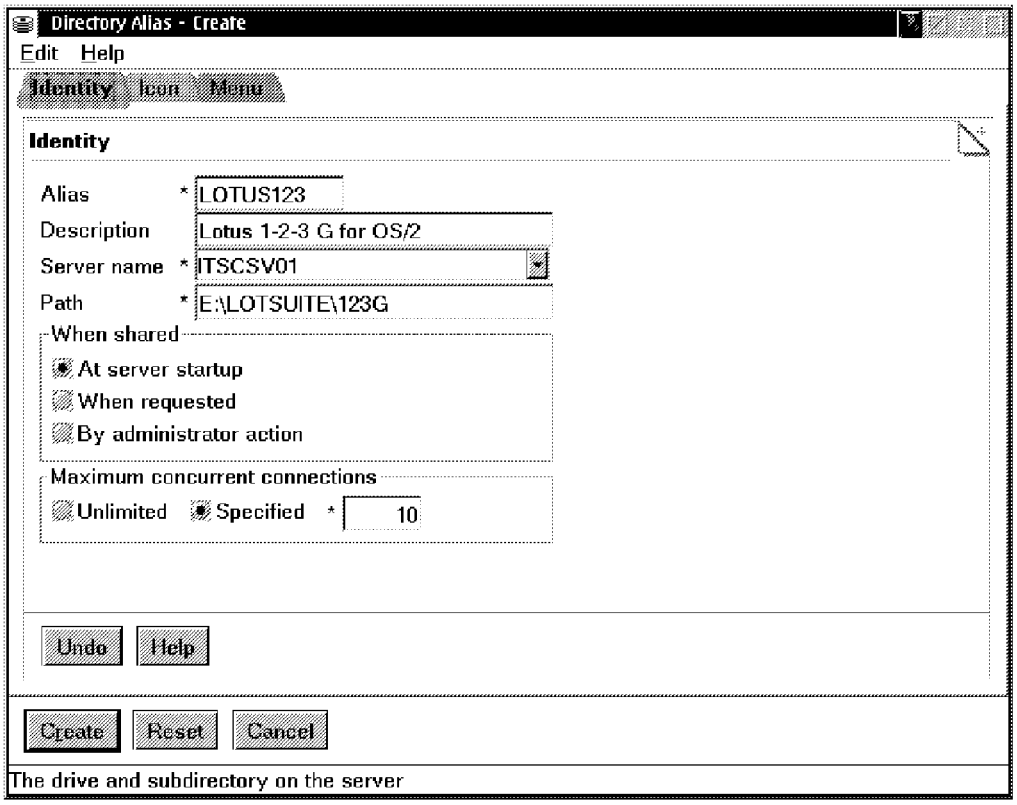

Figure 81. Directory Alias - Create Notebook

The following window is displayed:

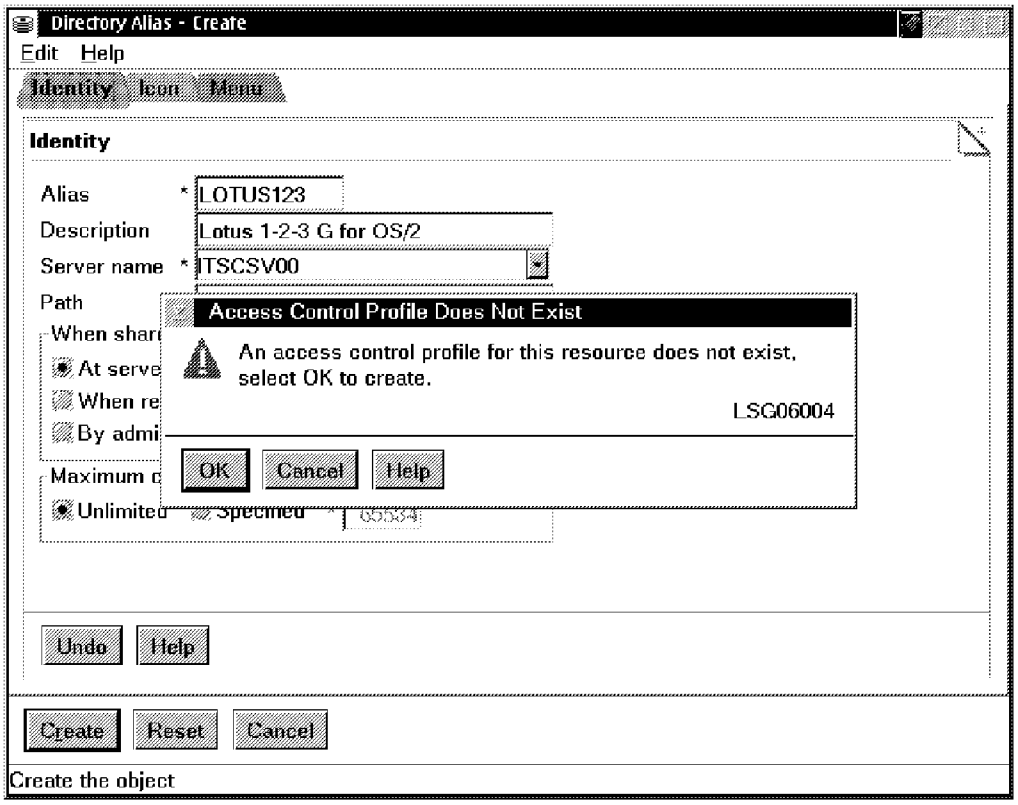

Figure 82. Access Control Profile Does Not Exist Window

8. Select **OK**.

   The Access Control Profile - Settings View notebook is displayed.

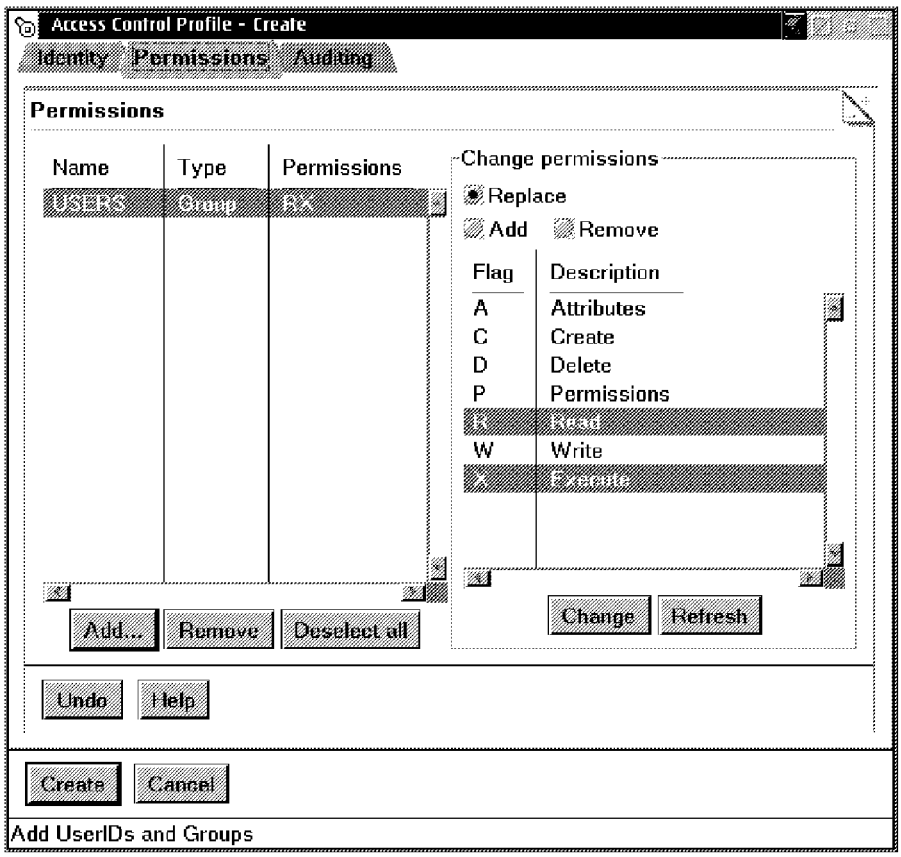

Figure 83. Access Control Profile - Settings View Notebook

9. Complete the permissions page as shown in Figure 83. In this case, all defined users will get Read and Execute rights. As in previous LAN Server releases, the USERS group is a special group to which all defined users belong.
10. Select **Create**.

    The following window is displayed:

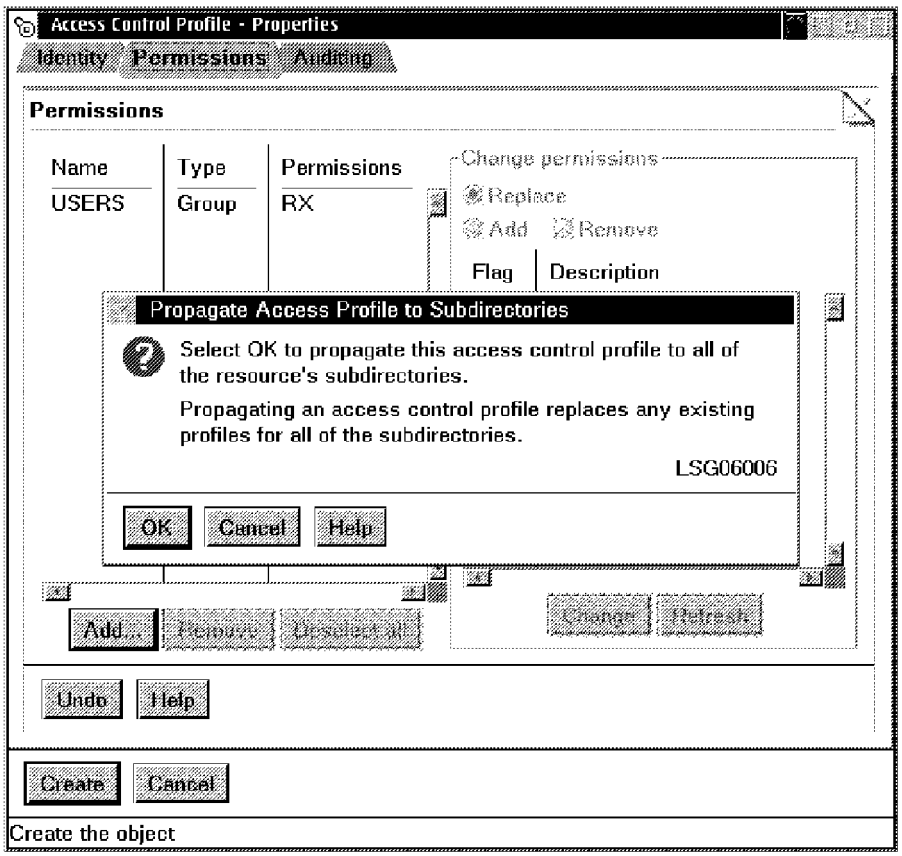

Figure 84. Propagate Access Profile to Subdirectories Window

11. Select the **OK** button.

    Since you just created the file alias for Lotus 1-2-3, you can now continue by creating a public application (also called network application) for Lotus 1-2-3.

12. To do so, double-click on the **Public Application Definitions** object.

    The Public Applications Definitions window will be opened. It is shown in Figure 85 on page 223.

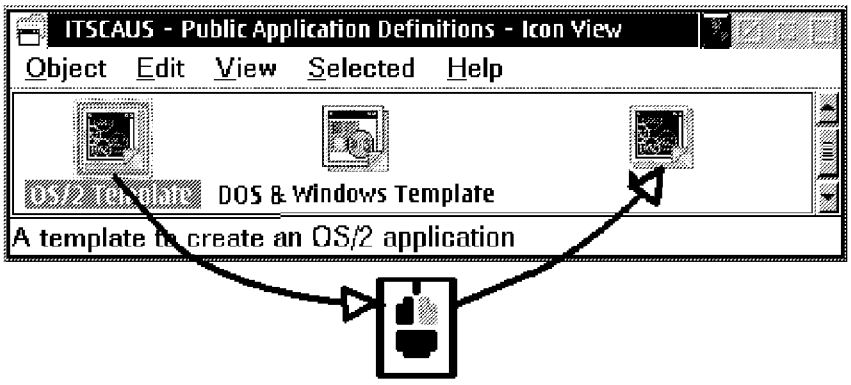

Figure 85. Public Application Definitions Folder

13. In the Public Applications Definitions folder, drag and drop an **OS/2 Template** to a free area of the folder.

    The OS/2 Application Definition - Create properties notebook is displayed as shown in Figure 86 on page 224.

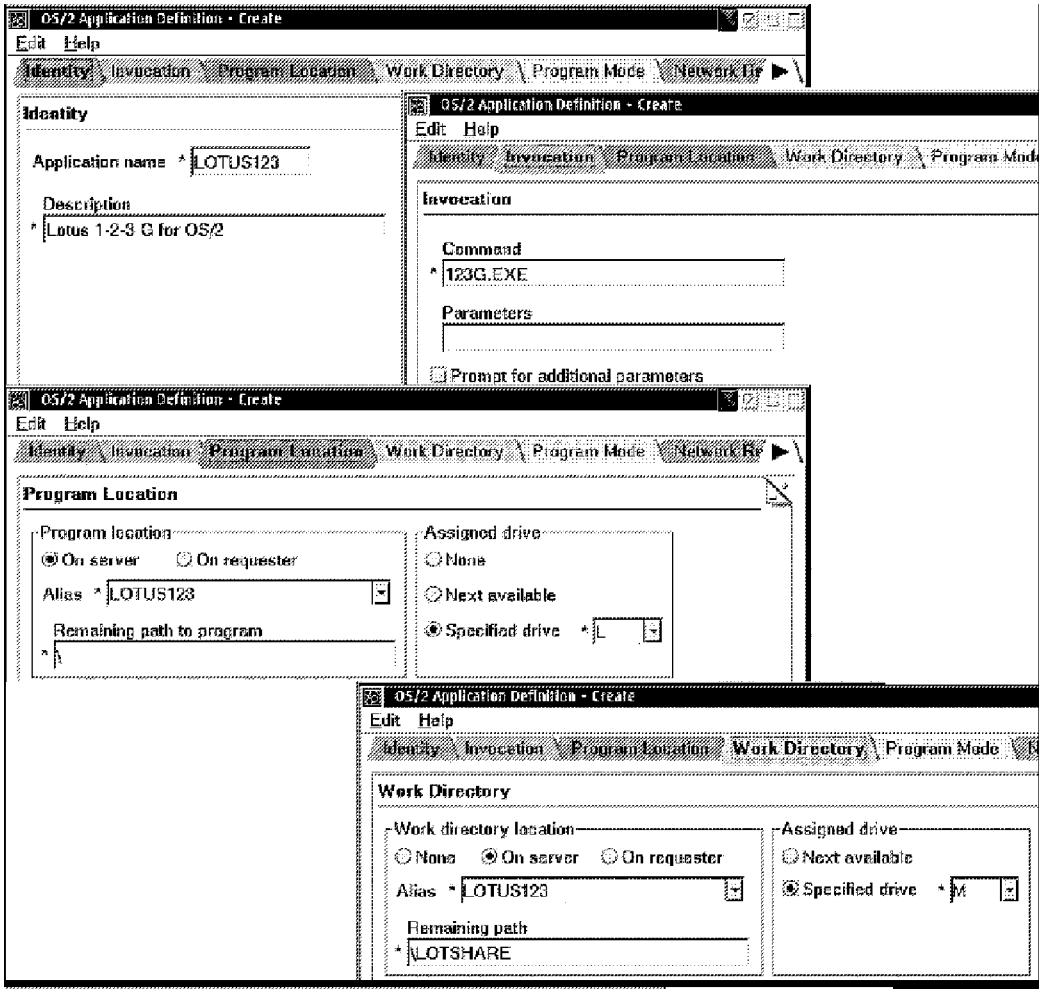

Figure 86. OS/2 Application Definition - Create Settings Notebook

14. In the OS/2 Application Definition - Create properties notebook, complete the Identity, Invocation, Program Location and Work Directory pages. The Program Location in this Lotus 1-2-3 example is on the server. The directory alias LOTUS123 (created in the steps before) has to be used.

    Note: For other applications that come with their own license control, you probably need to assign a drive that points to the license control software. After having created a file alias for this particular license control software, you may assign this alias by using either the Work Directory or Network Resources page.

15. Select **Create**.

    Now that you have created a network application, you may assign it to a user or group.

16. In the Domain Contents folder, double-click on the **User Accounts** folder and the **Groups** folder.

17. Drag and drop the newly created Lotus 1-2-3 application object to the user or to the group that has the Lotus 1-2-3 product selectable from its Network Applications folder after successful logon.

11.7.2 Installing DOS and Windows Public Applications

The steps to be followed to define DOS and Windows public applications are similar to the ones described in 11.7.1, “Installing an OS/2 Public Application” on page 218.

The steps are as follows:

1. Install the DOS or Windows application on the server. Follow the instructions that come with the product.

2. Create a file alias.

3. Create an Access Control Profile for the alias.

4. Create a DOS or Windows Public Application using the DOS Template as shown in Figure 87.

   With previous LAN Server products, the LAN administrator had to create an OS/2 command file that invoked the DOS or Windows application. This command file was then defined as an OS/2 public application. With LAN Server, the LAN administrator can now use the DOS Template to define DOS and Windows public applications. In addition, the LAN administrator defines the real program executable file.

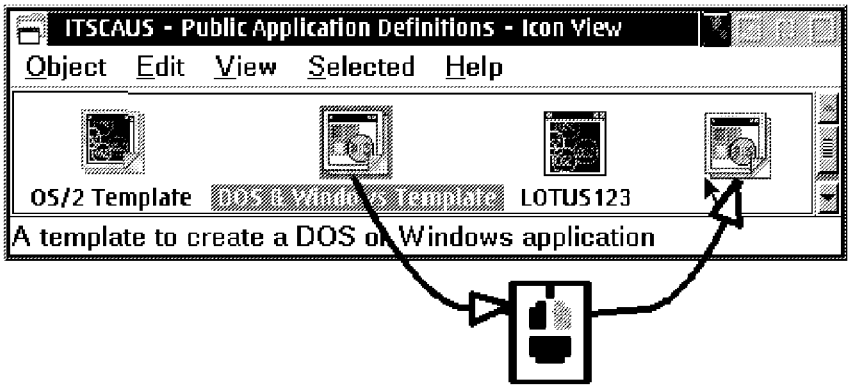

Figure 87. Public Application Definitions Folder

5. In Figure 88 on page 226 complete the Identity, Invocation, Program Location, and Work Directory pages.

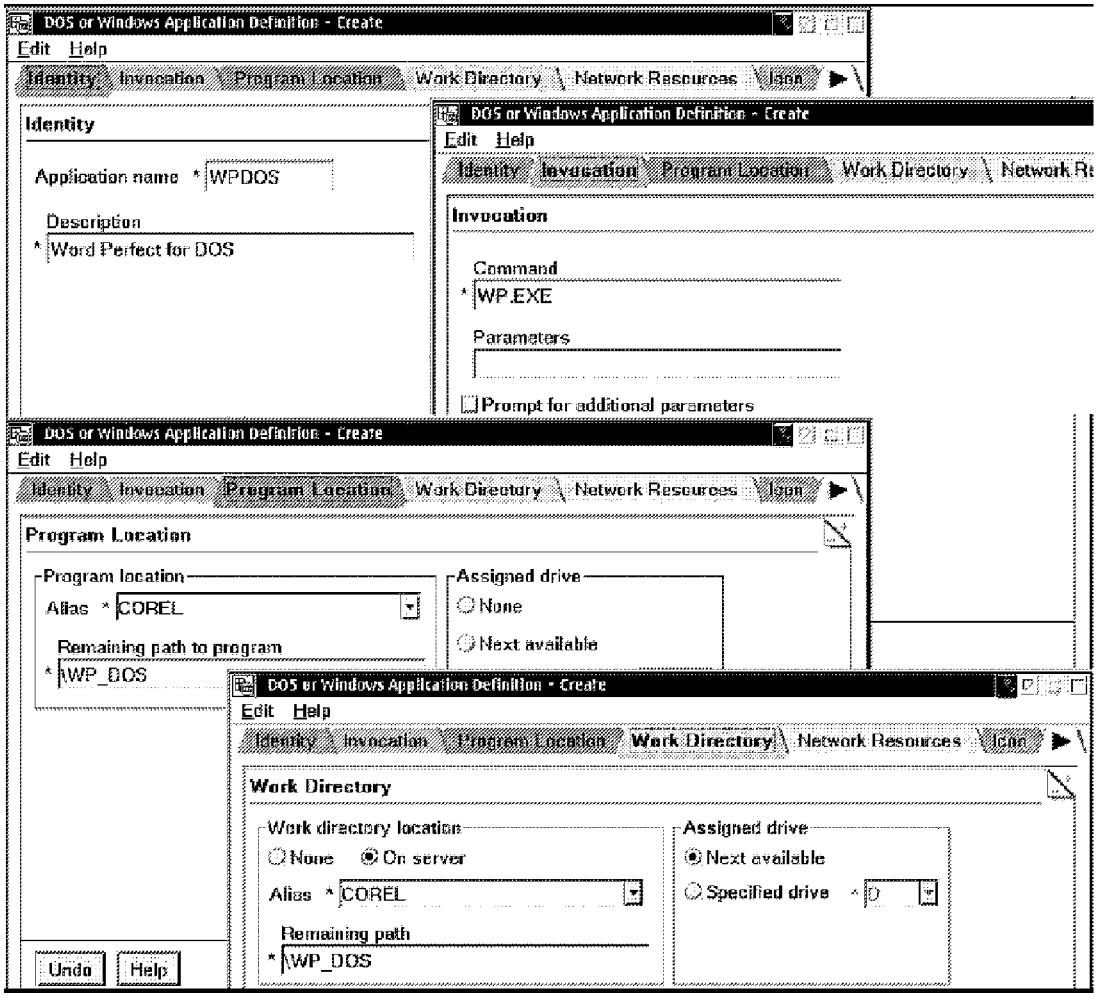

Figure 88. DOS Template - Settings View Notebook

6. Add the DOS or Windows application to the user's Network Application folder.

After you have installed and configured the network applications and assigned them to the appropriate users, these users will receive the Network Applications folder after successful domain logon. Figure 89 on page 227 shows an example of how this folder looks with the Lotus 1-2-3 application:

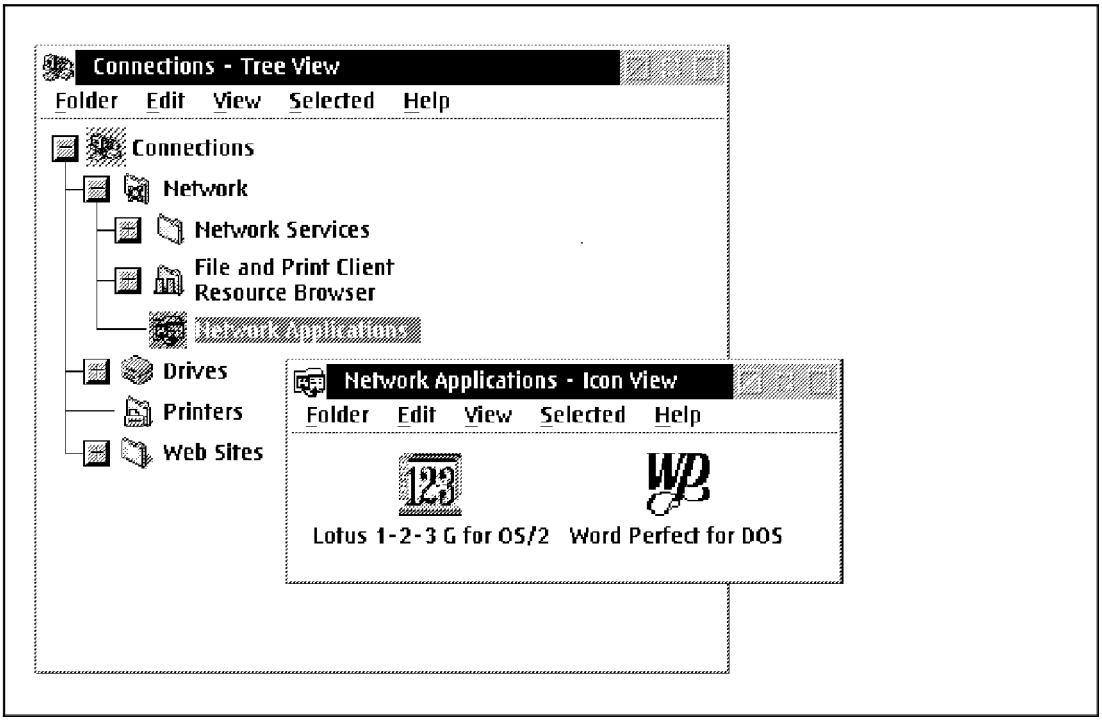

Figure 89. Network Applications Folder

Notes:

1. The location of the Network Application folder has changed in OS/2 Warp Version 4.0. If you are using a previous version of OS/2 the Network Application folder will be placed right onto your desktop after a successful logon. OS/2 Warp Version 4.0 places the Network Application folder directly into the Network folder, located in the Connection folder. As I mentioned earlier (see 11.2, “Where to Find the LAN Server Graphical User Interface” on page 201), the Network folder is not browsable via the Warp Center.

In file and print sharing services of Warp Server, the program properties page of a user's network application contains the actual application that is being invoked, unlike in previous releases as mentioned above. In addition, each program object now has the icon that corresponds to the program rather than only the standard OS/2 icon.

OS/2 and Windows programs normally have an icon associated with the program. DOS programs don't have icons. However, as shown in Figure 89, you can see our DOS program (Word Perfect for DOS) has an icon associated with it. To associate an icon to a DOS program, do the following:

1. Create an icon with the OS/2 Icon editor as shown in Figure 90 on page 228.

2. Copy the icon to the directory in which the DOS executable program file exists.
3. Name the icon the same as the DOS executable program file name. The file extension of an icon file always is .ICO.

In our example, the DOS program is called WP.EXE; the name of the icon file must be WP.ICO and must be placed in the same directory in which the EXE file resides. The same rule applies to batch (.BAT) files.

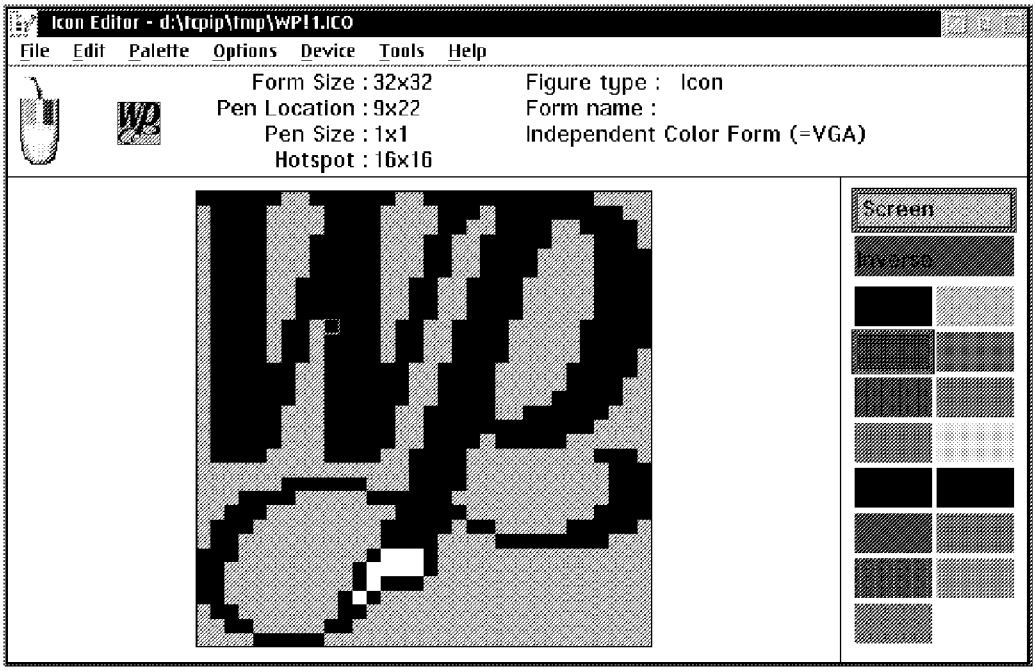

Figure 90. Icon Editor Window

11.7.3 Dynamic Link Library Considerations

Most OS/2 applications require access to Dynamic Link Library (DLL) modules to run. There are two alternatives available to do this:

- Ensure that the requester's machine has a period in the LIBPATH statement in the CONFIG.SYS file, so the current directory is searched for DLLs. This method relies on the user *not* changing the current directory while using the networked application.
- Create another alias for all DLLs, and have users also connect to this as one of the logon assignments (or working directory).

  An extra setup step is then required to copy DLLs to the common DLL subdirectory for the public applications. The Access Control Profile should be created, and the DLL alias should be assigned a drive for the

user at logon time. In addition, the OS/2 requester workstation needs to modify its LIBPATH statement in CONFIG.SYS to access the common DLL on the server.

11.7.4 Defining Network Applications from the OS/2 Desktop

The OS/2 desktop supports seamless interfaces to the application object. The user doesn't need to be aware of where the application is located. If it is a network application, a logon screen will appear with the appropriate logon panel when the object is double-clicked. If the application belongs to a LAN Server, its logon panel will appear. If it belongs to a NetWare server, the NetWare login window will appear. Here, we discuss the steps to define the network applications and how they must be defined to run correctly.

11.7.4.1 Creating a Network Application Using Templates

The following steps show you how to define a network application from a template:

1. Drag and drop a **Program** template from the Templates folder to an open area on the OS/2 desktop.

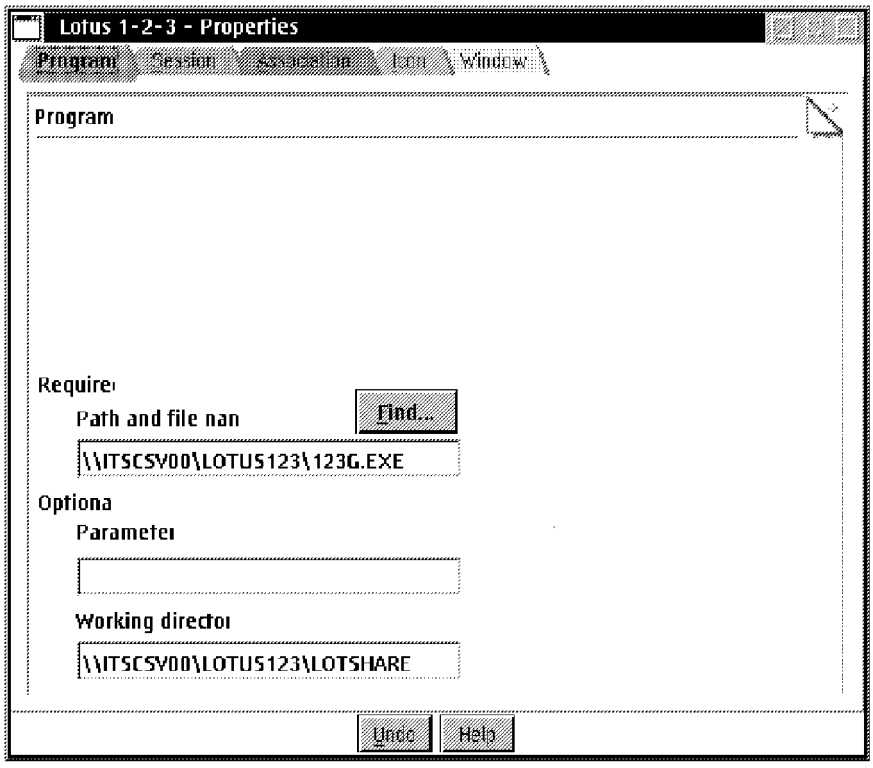

Figure 91. Program Template Notebook

2. In Figure 91, in the Path and file name field, enter the Universal Naming Convention (UNC) of the application. You can use the Find function to help locate the shared resource.

   In the Working directory field, you can enter a working directory.

3. Specify the name of the application in the **General** page of the notebook.

11.8 Multiple Domain Administration

There are two ways to make other domains known for cross-domain administration:

1. Specifying other domain names in the [requester] section of the \IBMLAN\IBMLAN.INI file by using the `othdomains` (other domains) parameter.
2. Using the `NET CONFIG REQUESTER` command from an OS/2 command line. For example, to make the domain called CONSTELLATION known to the LAN Server Graphical User Administration interface, type the following command:

   ```
   NET CONFIG REQUESTER /OTHDOMAINS:CONSTELLATION
   ```

   The next time you start the GUI, the domain called CONSTELLATION is available for administration in the LAN Server Administration folder.

In the following IBMLAN.INI file, you can see that we have specified additional domain names by using the `othdomains` (other domains) parameter. Using this parameter, you can specify up to four domains separated by commas. To have the ability to administer up to the maximum of six domains, you just have to log on to a domain that is neither specified by the `DOMAIN` parameter nor specified as a domain by the `othdomains` parameter. The `DOMAIN` and `othdomains` parameters are shown in Figure 93 on page 232.

To administer different domains from one workstation, you must have set up the same user IDs and passwords on all domains you want to administer. This is an absolute prerequisite!

A useful tool to ensure synchronized passwords on all domains you would like to administer is Network SignON Coordinator (NSC/2), which is also shipped with the Warp Server product.

Note: For more information about the Network SignON Coordinator, please refer to 1.2.2, “Password Coordination” on page 8.

```
; OS/2 LAN Server Advanced initialization file

[networks]
  net1 = NETBEUI$,0,LM10,100,200,14
; This information is read by the redirector at device initialization time.
[requester]
  COMPUTERNAME = ITSCSV00
  DOMAIN = ITSCAUS
; The following parameters generally do not need to be
; changed by the user.
  charcount = 16
  chartime = 250
  charwait = 3600
  keepconn = 600
  keepsearch = 600
  maxcmds = 16
  maxerrorlog = 100
  maxthreads = 10
  maxwrkcache = 64
  numalerts = 12
  numcharbuf = 10
  numservices = 16
  numworkbuf = 15
  numdgrambuf = 14
  othdomains = LS40DOM,ACCTDEPT,FINANCE,PERSONNEL
  printbuftime = 90
  sesstimeout = 45
  sizcharbuf = 512
  sizerror = 1024
  sizworkbuf = 4096
  useallmem = No
```

Figure 92. Extract of the IBMLAN.INI File to Ensure Multiple Domain Administration

In the example shown in Figure 93 on page 232, we are logged on to the ITSCAUS domain. Because of the changes made to the `othdomains` parameter (shown in Figure 92), we are now able to administer the LS40DOM, ACCTDEPT, FINANCE, and PERSONNEL domains as well.

Note: Although cross-domain administration is possible while logged on to one domain, you cannot drag and drop objects from one domain to other domains. To set up cross-domain resources, see 11.9, “Creating Cross-Domain Resource Definitions” on page 232.

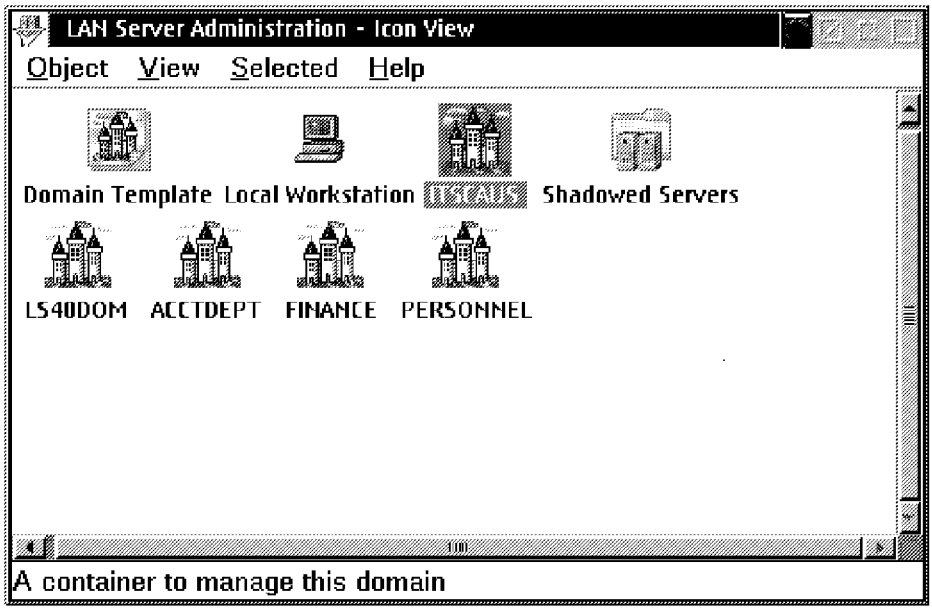

Figure 93. Administering Multiple LAN Server Domains

11.9 Creating Cross-Domain Resource Definitions

The following steps show you how to create an alias for a resource outside of the current domain, called a *cross-domain resource* or external alias, by using the GUI:

1. Within the Domain Controller Content folder, double-click on the **Resource Definitions** object.
2. Drag and drop a **Directory Template** to an open area in the Resource Definitions folder.

   The Directory Alias - Create properties notebook will be opened. It is shown in Figure 94 on page 233.

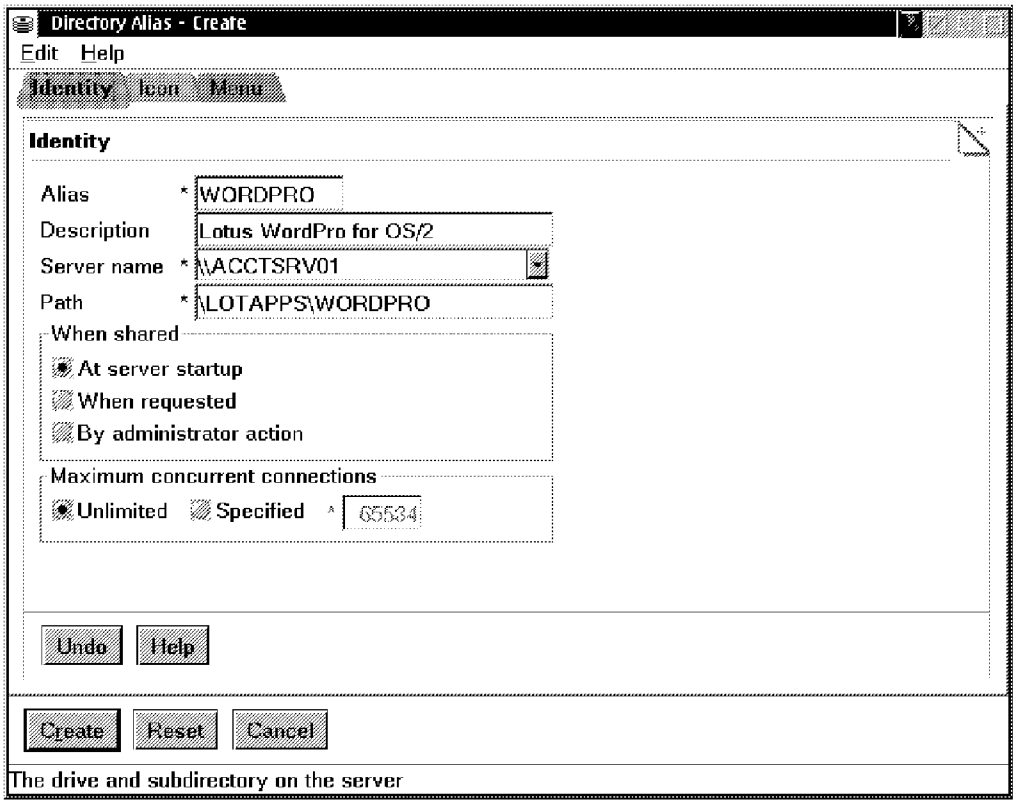

Figure 94. Directory Alias - Create Settings Notebook

3. Enter values as follows:

   - Alias - Type in the netname (sharename) of the resource on the external server.

     Note: It is assumed that the resource is already shared on the external domain, which is a requirement for cross-domain access.

   - Type in a description that you like.
   - Server name - Type in the name of the external server.

     Note: Do not use the pull-down menu here since the external server does not belong to the current domain.

   - Path - Type in the drive and path to the resource on the external server.
   - In the When shared list, select the radio button for either At server startup, or When requested, or By administrator action.

4. Press the **Create** button.

As with any resource, for the user to successfully use the cross-domain resource, access permissions must be set properly. Ensure that either you (if you are the administrator on both domains) or the administrator of the other domain does at least one of the following:

1. Set up the user ID (with the *same* passwords) in both domains, and then grant permissions to the resource in the other domain through the user ID.
2. Grant the desired access permission to the resource through the GUEST user ID on the external domain.

You can do one or both of the above so that users can then access the cross-domain resource transparently; that is, the resource appears to be in the local domain. If the user ID is *not* known to the external domain, then the user will be granted the GUEST user permissions. If the user ID *is* known to the external domain, then the user is granted the permissions for that ID.

Note: If the user ID is known on the external domain, but the passwords are not the same on both domains, then the user can still access the cross-domain resource by specifying the external password when requesting resource use, for example:

```
NET USE X: TEMP password
```

11.10 Managing Machines

This section shows how you can define additional servers and shadowed severs via the GUI.

Note: A Backup Domain Controller (BDC) also is considered as an additional server as far as the Shadowed Servers folder is concerned.

In Figure 95 on page 235 you can see the following four folders:

- LAN Server Administration
- Domain Contents (ITSCAUS)
- Shadowed Servers
- Defined Servers

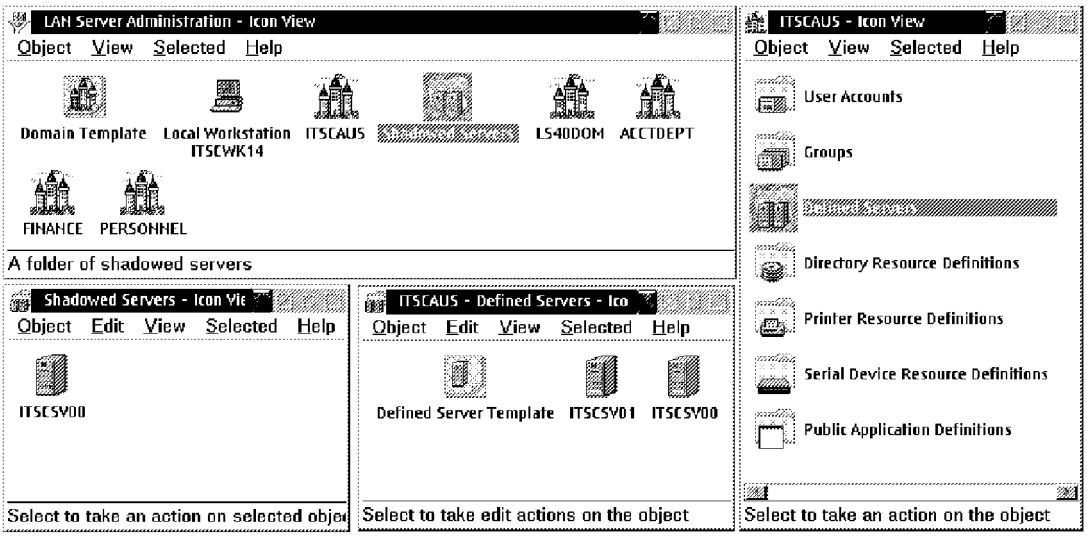

Figure 95. Additional Servers and Shadowed Servers

11.10.1 Defining an Additional Server

The following steps show you how to define an additional server:

1. Double-click on the **ITSCAUS** icon in the LAN Server Administration folder.

   The folder ITSCAUS Domain Contents folder will be opened.

2. In the ITSCAUS folder, double-click on the **Defined Servers** object.

   The Defined Servers folder will be opened.

3. To create an additional server, drag and drop a **Defined Servers Template** to an open area in the Defined Servers folder.

   The Defined Servers - Create properties notebook will be opened. Now you can specify the settings for the additional server. As a result, you have an additional server object in the Defined Servers folder. In our example, the server name is ITSCSV01.

11.10.2 Defining a Shadowed Server

The Shadowed Servers folder just displays shadows of servers that are defined in the Defined Servers folder. Use this folder to have shadows of servers for a quicker access to the servers you administer most often. This may be useful if you are administering multiple domains from the GUI, and would like to group the servers in one folder.

To make a shadow of a server, drag any server from a Defined Servers folder and drop it in the Shadowed Servers folder. In our example, we

dragged the ITSCSV01 object from the Defined Servers folder to the Shadowed Servers folder.

Have a look at the following steps to see what exactly can be done with local and additional servers:

1. Select the server you would like to administer in the Shadowed Servers folder or in the Defined Servers folder. In our example, the local server ITSCSV00 is used.
2. Choose the pull-down menu item **Selected**; click on the right arrow of **Open**.

This shows you a list with the following menu items:

- **Open files**

  Select this menu choice to display a list of files that are open on this server. Files can be closed by using the Open Files window.

- **Active sessions**

  Select this menu choice to display a list of sessions that are active on this server. Sessions can be deleted by using the Active Sessions window.

- **Statistics**

  Select this menu choice to select the type of statistics to display. You can choose server or requester statistics. Statistics can be refreshed, printed, or cleared by using the selected Statistics window.

- **Current shared**

  Select this menu choice and then select **Directories, Printers or Serial Devices** to display or change the devices currently shared by this server.

- **Current assignments**

  Select this menu choice to display a list of the current device assignments that are active for the workstation.

Figure 96 on page 237 shows you that we also can administer the services on local and additional servers by first double-clicking on the server object and then double-clicking on the **Services** object. We then can start, stop, or pause services on local and additional servers.

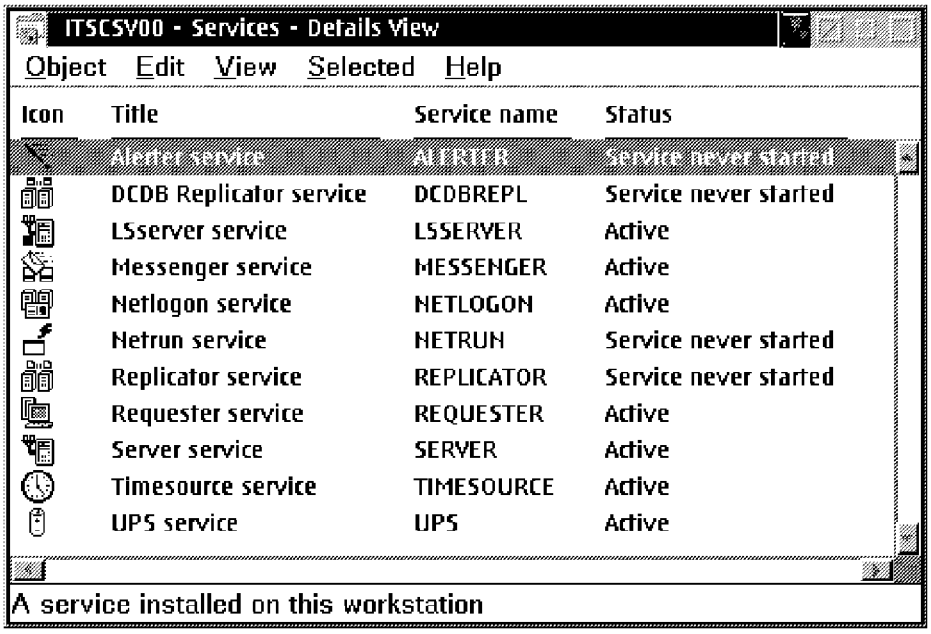

Figure 96. Server Services

11.10.3 Server - Settings View Notebook

In this section, you will find information about server settings and what you can administer within the server properties notebook.

1. In the LAN Server Administration window, select either **Local Workstation** object if your local workstation is a server or go to the Defined Servers folder or Shadowed Servers folder and select one server.
2. Press the right mouse button for object manipulation.
3. Select the right arrow of the **Open** item.
4. Then select **Settings**.

The opened server properties notebook is displayed in Figure 97 on page 238.

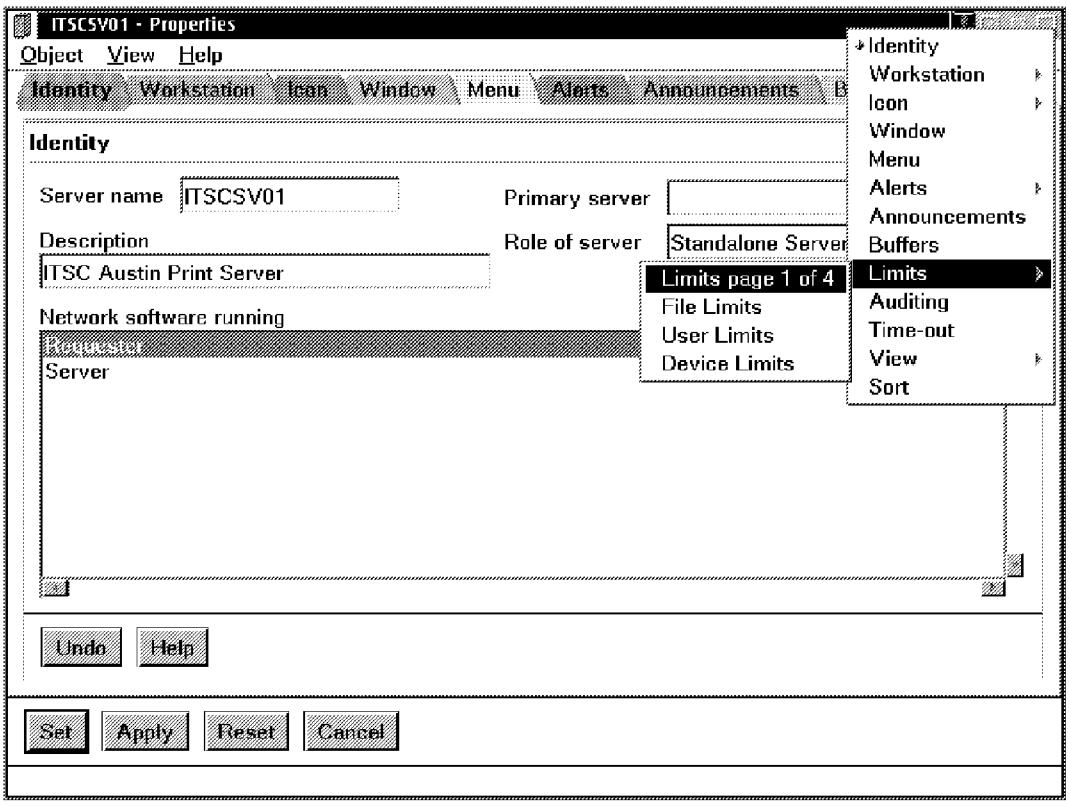

Figure 97. Server - Settings View Notebook

You can use the following notebook pages to change or display different settings for the server:

- Select the **Identity** notebook page to display the server name and other descriptive information about the server. You can also change the comment in the **Description** field.
- Select the **Workstation** notebook page to display the workstation name and other descriptive information about the workstation.
- Select the **Alerts** notebook page to specify alert recipients. An alert is a notification about an event that usually indicates an error, either on the network or in a user operation. Typical events for which alerts are sent include logon violations, access violations, or a decrease in available disk space below some acceptable level.

  Alert recipients may be specified as user IDs or machine IDs.
- Select the **Announcements** notebook page to display or change information about the frequency with which this server announces its presence on the network.
- Select the **Buffers** notebook page to display information about the server buffers provided.

- Select the **Limits** notebook page to display information about the limits the server sets for activity on the network.
- Select the **Auditing** notebook page to display the types of events that can be audited for this server. It also displays which events are audited when auditing is enabled on the server. If the status field of an event is blank, the event is not audited. The following events can be audited on a server:
  - Service state changes
  - Successful session requests
  - Unsuccessful session requests
  - All session requests
  - Successful domain logon requests
  - Unsuccessful domain logon requests
  - All domain logon and logoff requests
  - All domain logon and session requests
  - Successful share requests
  - Unsuccessful share requests
  - All share requests
  - Changes to the user and group account database
  - Changes to the access control database
  - Resource access as defined by per resource auditing options
  - Logon limit violations
- Use the **Time-out** notebook page to display or change information about time-out conditions on the server.

11.11 Fixing a Corrupted NETGUI.INI File

If the NETGUI.INI file is damaged, possibly as the result of a trap or abnormal end of the LAN Server Administration GUI, the icons may have default titles or previously saved information may no longer be available.

To fix this problem, replace the NETGUI.INI file from a backup copy. The two files that store persistent information for the GUI are NETGUI.INI and NETGUI.PDB. Use the following steps to replace these files:

1. Close the LAN Server Administration GUI.
2. Copy C: IBMLAN NETPROG NETINI.BAK to C: IBMLAN NETPROG NETGUI.INI, where C: is the drive on which the LAN Server is installed.
3. Erase NETGUI.PDB.
4. Reopen the LAN Server Administration GUI. The NETGUI.PDB file will be recreated.

11.12 GUI Versus Batch Processing

The graphical user interface of OS/2 LAN Server and OS/2 Warp Server is a very easy, straight-forward way to manage the server environment.

Nevertheless, it might not be an adequate solution if you have to manage a very large environment (100 or more users). OS/2 LAN Server and OS/2 Warp Server provides you with the ability to manage everything not only via the GUI but also using a batch procedure.

To write a batch application for manipulating your LAN environment, you do have several choices:

1. Use an ordinary DOS-like batch file calling NET commands

   Advantage Same oldfashioned way if you already know DOS batch files and like this kind of programming

   Disadvantage DOS batch relies completely on the LAN Server `NET` command. Unfortunately the `NET` command does not support all possible LAN Server APIs. For example, for some features of the GUI there is no corresponding `NET` command. Besides that major disadvantage, the `NET` commands are comparably slow to the solution we suggest below.

2. Use REXX as your primary choice for writing batch files.

   REXX not only gives you the ability of writing batch code in a real structured way (REXX supports IF THEN ELSE, DO WHILE, DO UNTIL and so forth) but also supports external libraries to implement functions into the programming language that are not included there by default. There is one excelent external library (DLL) from Ingolf Lindberg (IBM Denmark), LSRXUTIL.DLL, which provides you with an REXX interface to every OS/2 LAN Server or OS/2 Warp Server file and print sharing API. You can download this DLL and other OS/2 Warp Server related tools from the following Internet address:

   `http://www.software.ibm.com/warp-server/download.htm`

   Advantage You can perform every possible manipulation of your LAN via a batch file. Additionally, it is remarkably faster than any `NET` command.

Disadvantage You need to have basic REXX programming skill. You do not have to be an expert in writing REXX programs, but at least you should know the basic principles.

3. Finally you can use the LAN Server Management Tools (LSMT) which were written by Hermann Pauli (IBM Germany) and Alain Rykaert (IBM Belgium) during a OS/2 Warp Server residency in Austin, Texas in 1995. These tools are written in REXX using the LSRXUTIL.DLL to demonstrate the usage of this very powerful DLL. These tools are included in an IBM Redbook titled *How to Manage PC Server Environments*, SG24-4879. These tools are also available on an IBM internal disk called OS2TOOLS that every IBMer has access to. To obtain a free copy of these tools, just ask your local IBM representative to download it for you.

   Advantage You do not have to write your own REXX programs, you just run these ready-to-use batch files, and if they do not fit your needs, you can adjust these programs by editing the REXX source code. One very nice feature of LSMT is its ability to exclude all information stored in your DCDB or within ACLs to readable and changeable ASCII files. This enables you to use the LAN Server/Warp Server GUI for single tasks and the REXX procedures for multiple tasks. Every single change using the GUI can be tracked by running all the GET programs of LSMT, which will create ASCII files containing the most actual information about your LAN Server/Warp Server domain, including all GUI-performed jobs. These files can also be used for restoring purposes, which is an major advantage and very useful for migration tasks.

   Disadvantage You need basic REXX skill to adjust the programs to special needs. You also should be aware of that due to the fact that these tools are provided at no charge, there is no IBM support available, and if you use them, it is for your own purpose and absolutely at your own risk. LSMT is an excellent tool for performing multiple tasks, but for manipulating single values, you should use the LAN Server/Warp Server GUI.

In Figure 98 on page 242 through Figure 100 on page 243, you will see extracts of batch files that create a user using the described methods above.

```
REM *** Add User Janine to domain   ***

NET ADMIN \\ITSCSERVER /C NET USER JANINE PASSWORD /ADD
         /ACTIVE:YES /PRIVILEGE:USER /PASSWORDREQ:YES
         /USERCOMMENT:"Janine_Rachel"
         /HOMEDIR:H:\ITSCSERVER\D$\HOMEDIRS\JANINE

MD \\ITSCSERVER\D$\HOMEDIRS\JANINE

NET ADMIN \\ITSCSERVER /C NET ACCESS D:\HOMEDIRS\JANINE /ADD JANINE:Y
```

Figure 98. The Old-fashioned DOS-Style to Add a User using Only NET Commands

Note: All numbered statements above have to be executed in one single line each.

```
/* Load the LSRXUTIL.DLL      */
call RxFuncAdd 'LoadLsRxutFuncs', 'LSRXUT', 'LoadLsRxutFuncs'
call LoadLsRxutFuncs

/* Add User Janine to domain   */
NETUSER = 280
SrvName = '\\ITSCSERVER'

 userInfo.name       = 'JANINE'
 userInfo.password    = 'PASSWORD'
 userInfo.priv       = 'User'
 userInfo.home_dir    = 'H:\ITSCSERVER\D$\HOMEDIRS\JANINE'
 userInfo.comment     = 'Janine Rachel'

 myRc = NetAdd(NETUSER, 'userInfo', SrvName)

 if myRc <> '0' then do
  say 'Got error from NetAdd() ' myRc
  call DropLsRxutFuncs
  exit 9
 end
 else do
  say
  say "User created successfully"
 end
```

Figure 99 (Part 1 of 2). Adding a User Using the REXX Programming Language and the LSRXUTIL.DLL

```
/* Drop the LSRXUTIL.DLL      */
 call DropLsRxutFuncs
 call RxFuncDrop 'LoadLsRxutFuncs'

 exit 0
```

Figure 99 (Part 2 of 2). Adding a User Using the REXX Programming Language and the LSRXUTIL.DLL

To add users to a domain, you would type the following command from an OS/2 command line:

```
SETUSERS /SRV:ITSCSERVER /INP:USERS.CSV
```

where `USERS.CSV` is an ASCII file containing the user information needed to create the account. This file can be created automatically by using the `GETUSERS.CMD` command and then edited to type in the new user.

```
OPT;NAME  ;PASSWORD;PRIV ;FLAGS;USR_COMMENT  ;HOME_DIR
A  ;JANINE;PASSWORD;User ;S    ;Janine Rachel;H:\ITSCSERVER\D$\HOMEDIRS\JANINE
```

Figure 100. USERS.CSV (Extract)

Note: The `A` parameter in the `OPT` section stands for Add User.

11.13 Dynamic TCP/IP in OS/2 Warp Server

OS/2 Warp Server comes with with a DHCP Server, a server that offers IP addresses out of a defined IP address pool to requesting IP Clients, such as OS/2 and DOS LAN Services. Microsoft Windows NT and Windows 95 clients are also supported clients for DHCP.

The second server that comes with Warp Server is a Dynamic DNS Server, which dynamically maps IP addresses to host names. It also does reverse mapping: host name to IP address. Supported clients are Warp Server OS/2 clients and OS/2 Warp 4 clients. In the first quarter of 1997, Windows 95 clients will be supported also for dynamic DNS serving, something that is not offered by Microsoft, not even with Windows NT 4.0. For the Windows 95 DHCP/DDNS client, check the Download section of Warp Server's home page on the Worldwide Web at the following URL.:

```
http://www.software.ibm.com/os/warp-server/
```

In the following section we demonstrate how easy it is to set up a dynamic TCP/IP environment under OS/2 Warp Server.

11.13.1 Configuring and Using DHCP Server

1. Within the TCP/IP folder, open the DHCP Server Services window as shown in Figure 101.

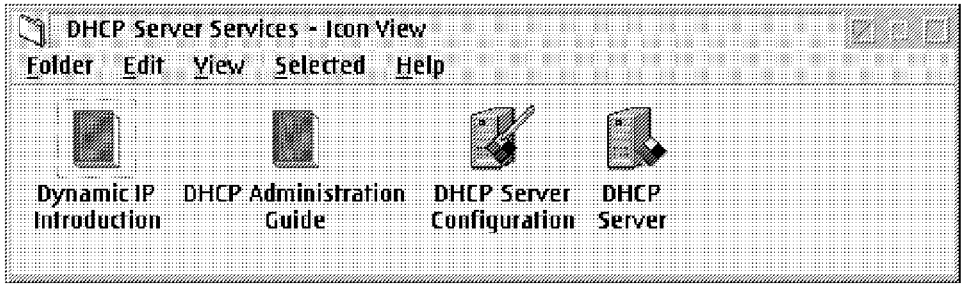

Figure 101. DHCP Server Services Window

2. To begin configuring the DHCP server, open DHCP Server Configuration. You will get the DHCP configuration graphical interface as shown in Figure 102.

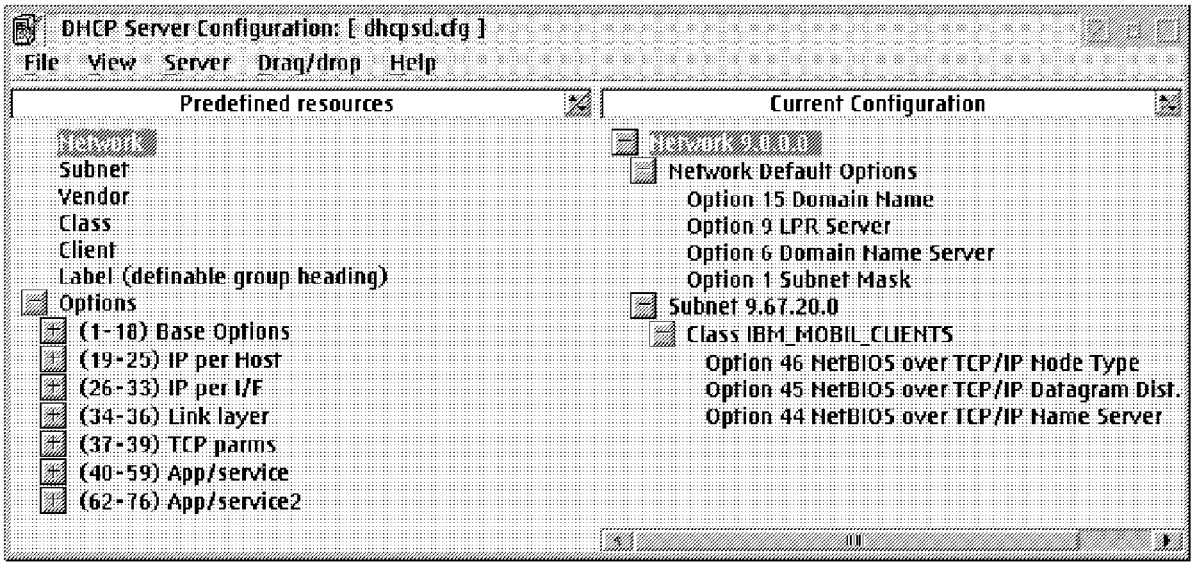

Figure 102. DHCP Server Configuration Window

Configuring is made very easy. Using the drag-and-drop technique, you select one item after another from the list of Predefined Resources and drag it to the list of Current Configuration. Therefore, if you need to delete an item from the list of Current Configuration, you need to drag and drop that item to the Shredder (which usually resides on the Workplace Shell).

This example gives instructions on how to set up a DHCP server for the subnet 9.67.20.0. You can extend the configuration for other subnets by doing the steps shown here for a single subnet.

a. First drag and drop the **Network** item to the list of Current Configuration and double-click on it. Use the spin-button arrows to define the subnet mask. In Figure 103 you find an example for subnetting the 9.0.0.0 network.

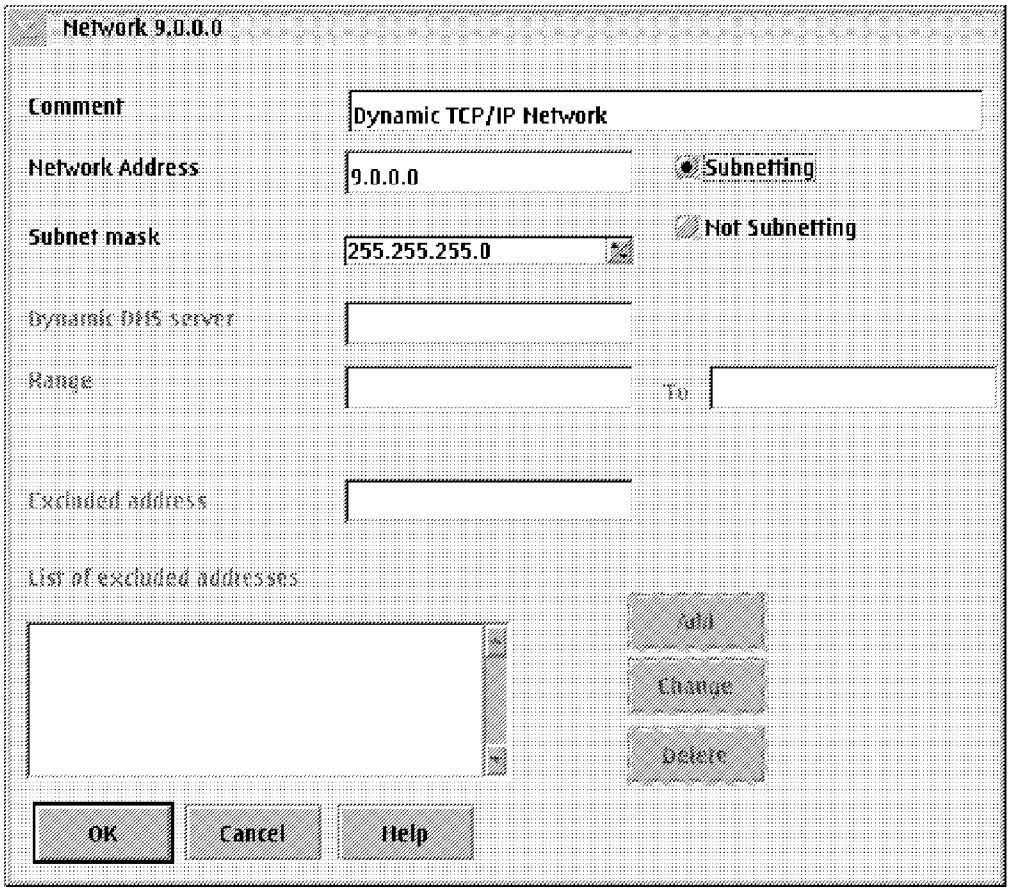

Figure 103. Network 9.0.0.0 Item Window

Note: Since you are subnetting in this example, you cannot type in information for Dynamic DNS server and Range of IP addresses.

b. Drag and drop the Label item on top of the Network 9.0.0.0 item in the list of Current Configuration so that an Expand-Tree sign will appear just next to the Network 9.0.0.0 item. Expand the tree now and double-click on the **Label** item. Provide a name, for example Network Default Options.

Note: You must see the Expand Tree sign left of the Network 9.0.0.0 item; otherwise you did not drop directly on top of the Network item. If this is the case, delete the item by dragging and dropping it to the Shredder, and do this step again.

c. Expand the Options tree and then (1-18) Base Options in the list of Predefined resources. Drag and drop **Option 15 Domain Name** on top of the Network Default Options item in the the list of Current Configuration. Expand the tree and double-click on the new item. Provide information as shown in Figure 104.

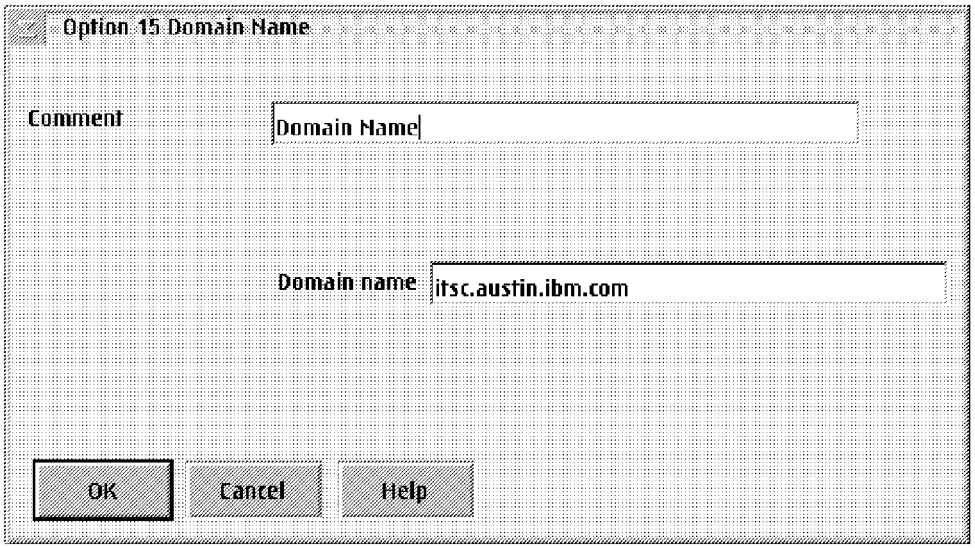

Figure 104. Option 15 Domain Name Item Window

d. If you have an LPR (Line Printer) server in your TCP/IP network, drag and drop **Option 9 LPT Server** to the Network Default Options item and double-click on it. Type in the TCP/IP address of the LPR server (example shown in Figure 105 on page 247).

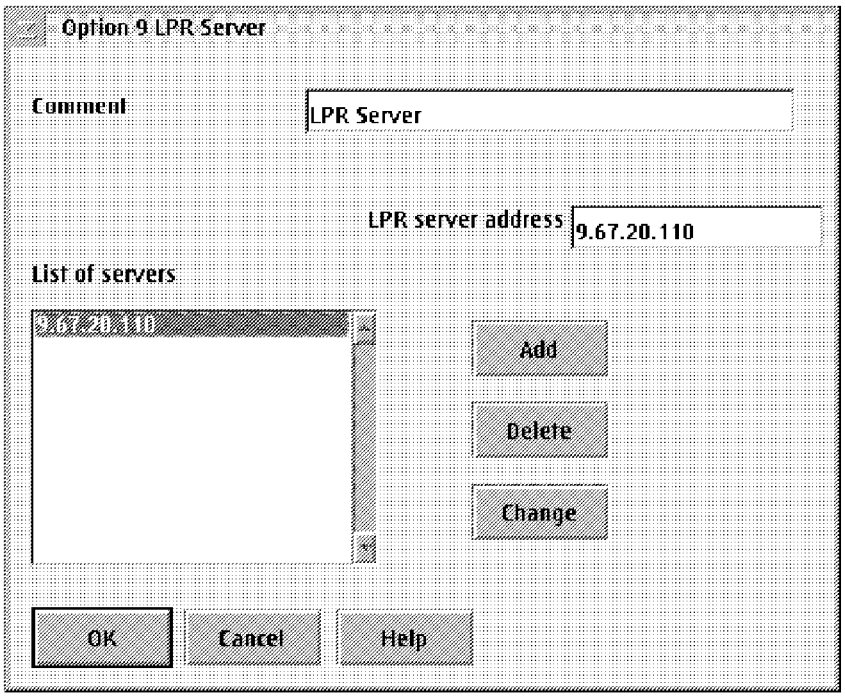

Figure 105. Option 9 LPR Server Item Window

e. Now we need to provide the TCP/IP address of the subnet's DNS server by dragging and dropping **Option 6 Domain Name Server** from the list of Predefined Resources on top of the Network Default Options item in the list of Current Configuration. Double-click on it, and provide information as shown in Figure 106 on page 248.

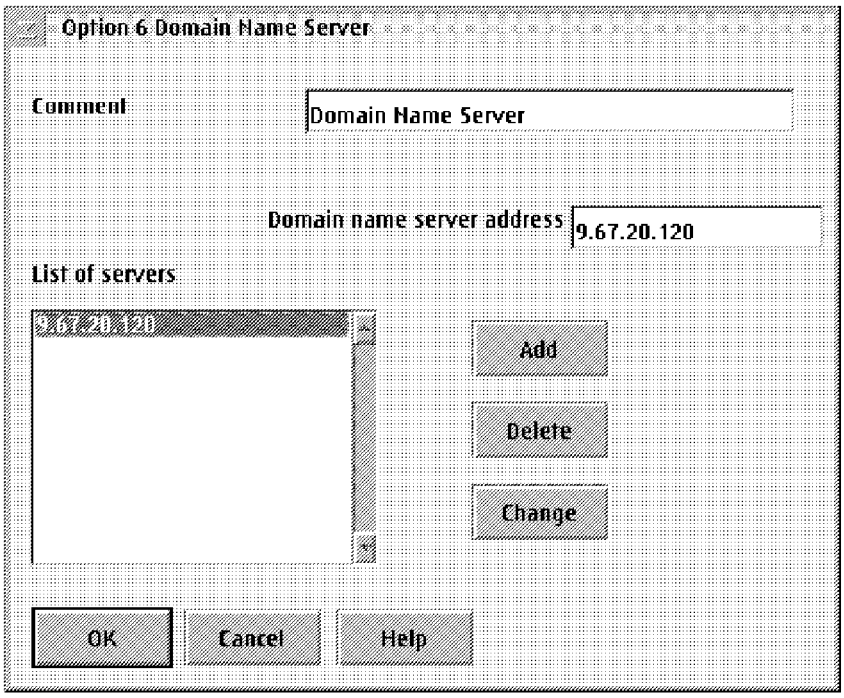

Figure 106. Option 6 Domain Name Server Item Window

f. To provide subnet mask information to the Dynamic IP client, drag and drop **Option 1 Subnet Mask** on top of Default Network Options and double-click on it. Find an example shown in Figure 107.

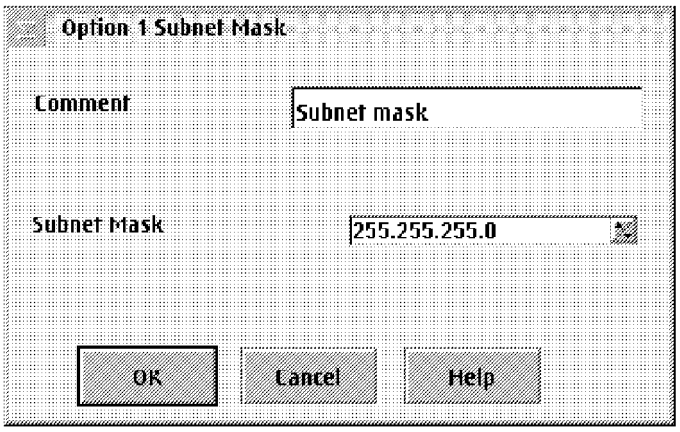

Figure 107. Option 1 Subnet Mask Item Window

g. Now the time has come to define a subnet. To do so, drag and drop the **Subnet** item (second item from the beginning) on top of the Network 9.0.0.0 item in the list of Current Configuration. Provide information as shown in the example of Figure 108 on page 249.

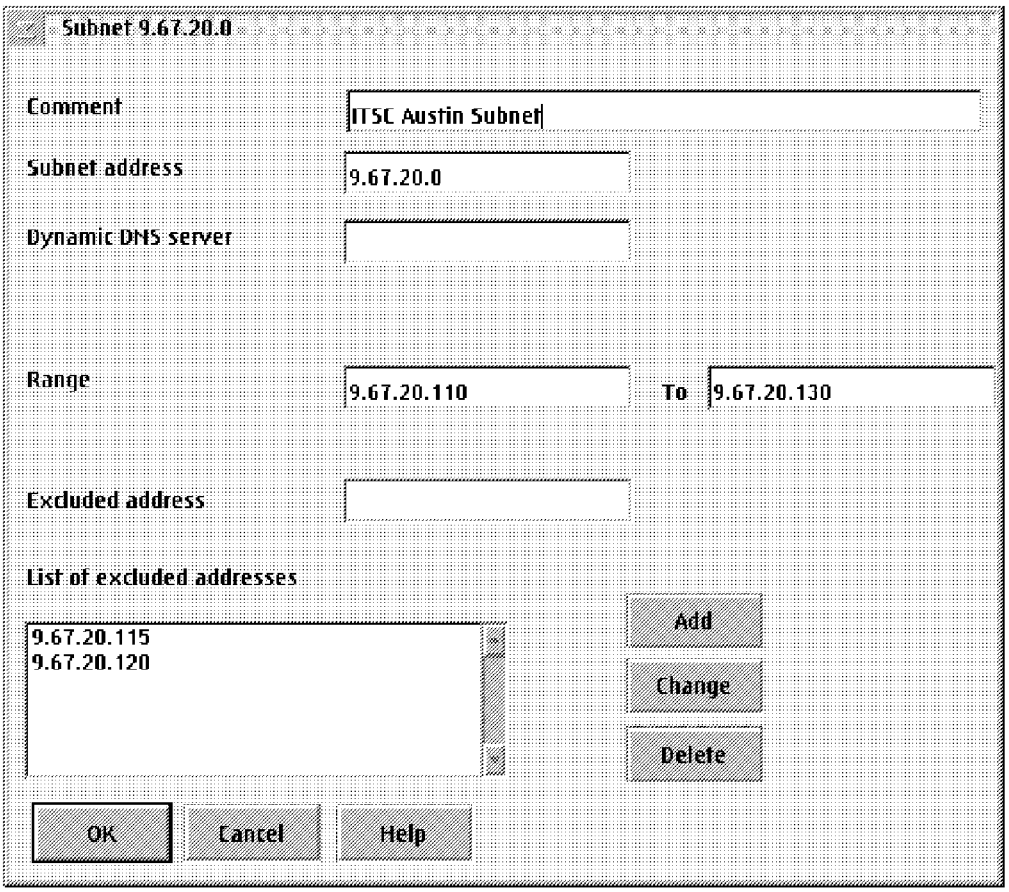

Figure 108. Subnet 9.67.20.0 Window

Specify the range of IP addresses and exclude the ones taken in this range. For example, 9.67.20.115 will be the static IP address for the NetBIOS Name Server, and 9.67.20.120 will be the address of the Dynamic DNS server. Excluded addresses will not be offered to requesting IP clients.

Note: You do not need to provide information about the Dynamic DNS server since this server has been defined already in the Network Default Options section.

The IBM DHCP Server also allows you to choose which configuration information is provided to the client based on who is using the client or what the client is being used for.

Using DHCP classing, you can provide unique configuration information to clients that identify themselves as belonging to a certain group. Further, you can administer parameters to a class either independent of or with respect to the location of the client, or a combination of both. For example, if you wanted all the users of

IBM mobile clients to use the Shadow NetBIOS name server at 9.67.20.115, you would do the following steps:

Note: At the client, you need to configure the client to identify itself as belonging to the Class "IBM_MOBIL_CLIENTS" which is described in 11.14.1, “Configure The Dynamic IP Client for User Class Support” on page 266.

h. Now drag and drop the **Class** item on top of the Subnet 9.67.20.0 item; expand the tree and double-click on the new item. In this example, we defined a class containing dynamic IP configuration for mobile IBM clients (Figure 109).

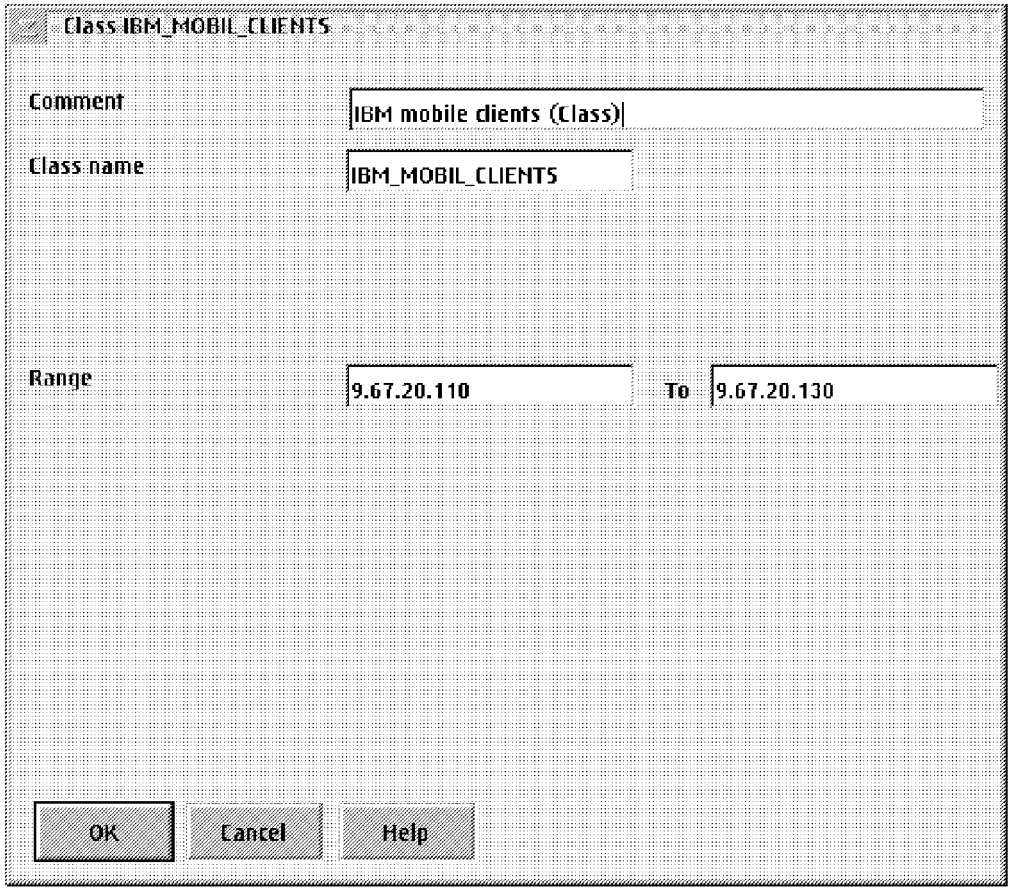

Figure 109. Class IBM_MOBIL_CLIENTS Item Window

For all IBM mobile clients defined in the step before that are configured with the TCPBEUI protocol, we will now provide information for a NetBIOS Name Server which supports datagram distribution (such as the Shadow product as described in 11.13.3, “NetBIOS Name Server Shadow” on page 260 so that these clients get NetBIOS names to IP address resolution from that server.

i. Expand the the (40-59) App/Service Options in the list of Predefined resources. Drag and drop **Option 46 NetBIOS over TCP/IP Node Type** on top of the IBM_MOBIL_CLIENTS item in the the list of Current Configuration and double-click on it. Use the spin-button arrows to the right to provide information as shown in Figure 110.

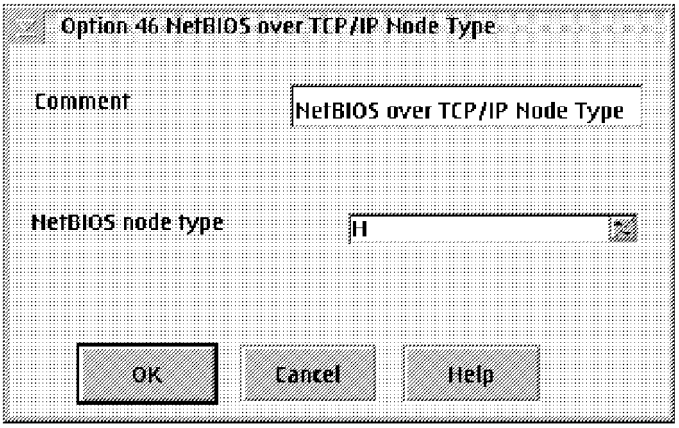

Figure 110. Option 46 NetBIOS over TCP/IP Node Type Item Window

Select **H** as the NetBIOS node type (Hybrid Mode).

j. Drag and drop **Option 45 TCP/IP over TCP/IP Datagram Distribution Server** on top of the IBM_MOBIL_CLIENTS item in the list of Current Configuration; double-click on it and provide the TCP/IP address of the NetBIOS Name Server (in our example it is 9.67.20.115).

Note: WINS as a NetBIOS Name Server does not support datagram distribution.

Figure 111 on page 252 demonstrates an example.

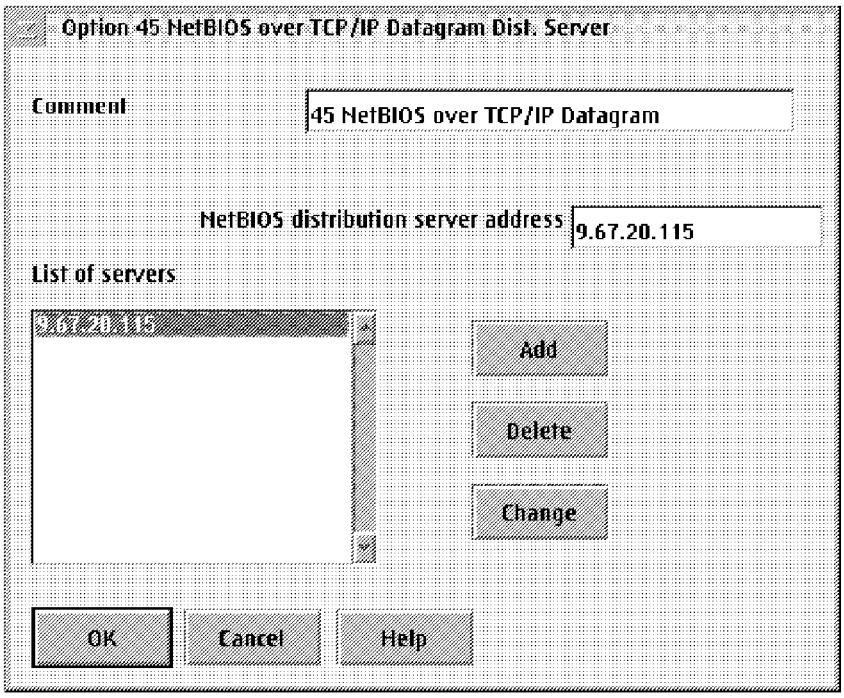

Figure 111. Option 45 NetBIOS over TCP/IP Datagram Distribution Server Item Window

k. Last but not least, drag and drop **Option 44 TCP/IP over TCP/IP Name Server** on top of the IBM_MOBIL_CLIENTS item in the list of Current Configuration; double-click on it and provide the TCP/IP address of the NetBIOS Name Server (in our example it is 9.67.20.115) as shown in Figure 112 on page 253.

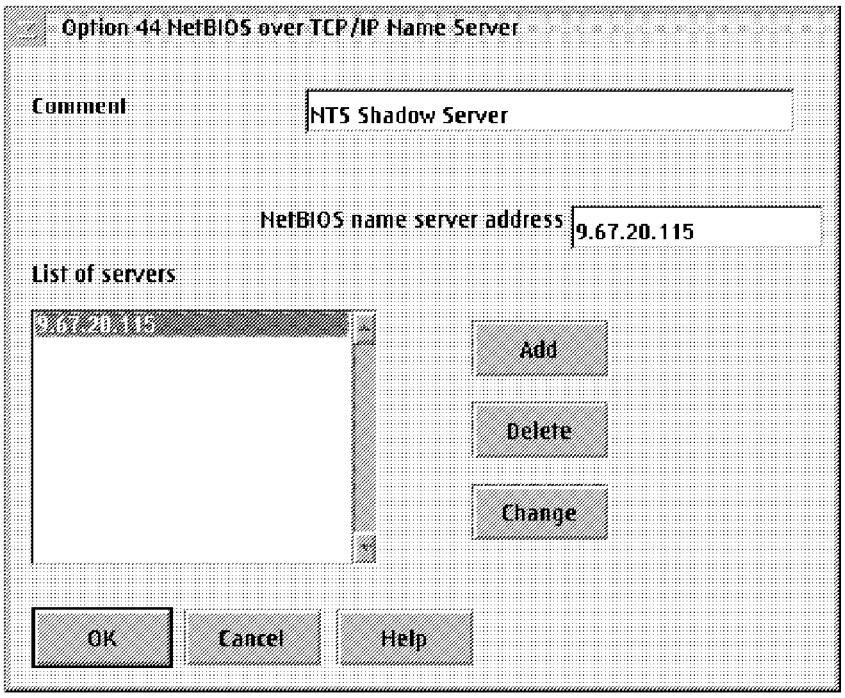

Figure 112. Option 44 NetBIOS over TCP/IP Name Server Item Window

l. Almost done, but you still need to provide server parameters. In the action bar, open the Server's pull-down menu, and select **View/change server parameters**, and provide information as shown in Figure 113 on page 254.

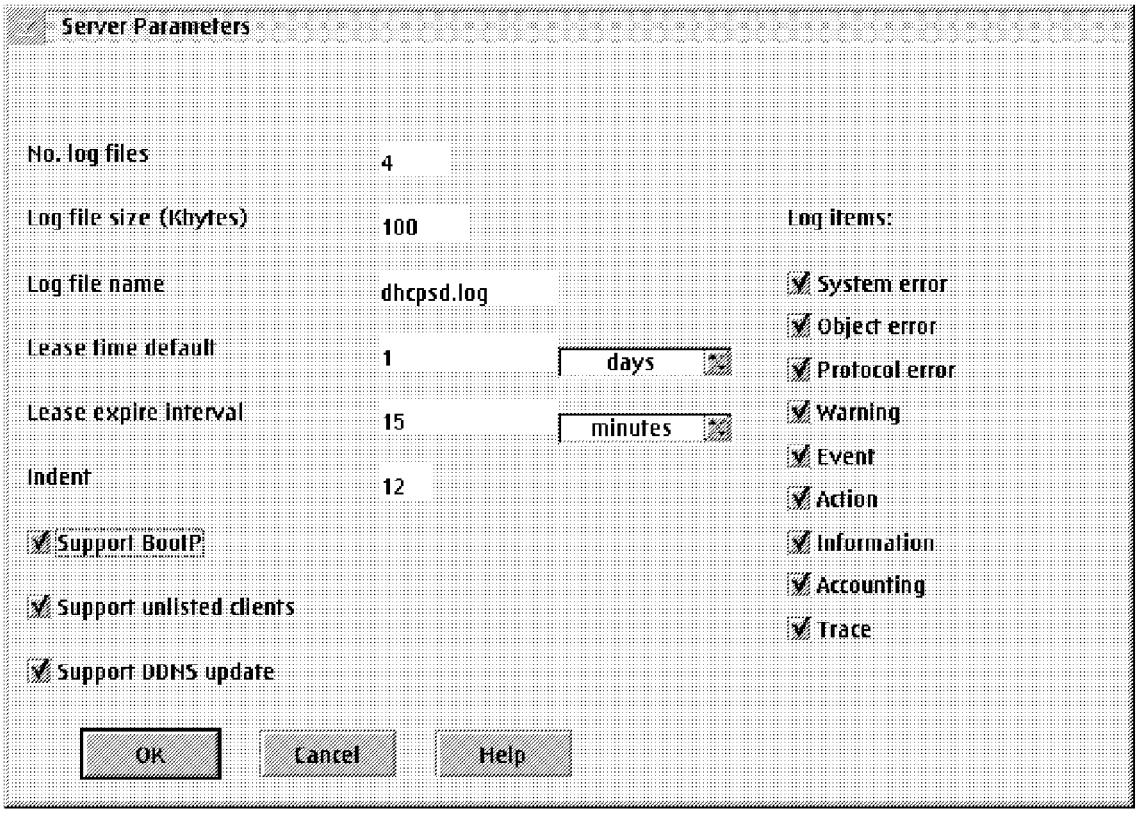

Figure 113. Server Parameters Window

You need to provide Lease time default as well as Lease expire interval information. The Leasing time will apply to all ranges of IP addresses defined in each subnet (in our case, the 9.67.20.0 subnet). Check the box for Support DDNS update so that the DHCP Server is allowed to communicate with the DDNS Server.

m. Save your configuration as DHCPSD.CFG.

11.13.1.1 Running the DHCP Server

Within the TCP/IP folder, open the DHCP Server Services window.

Hint

We recommend you to change the properties (settings) of the DHCP Server object. In the Parameters section, type in `-v`. This will give you information about clients who are trying to obtain an IP address from the DHCP Server once the server is up and running. You have a better control over what is going on in your defined subnet(s).

Start DHCP Server by double-clicking on its object. The following information will be presented to you, as shown in Figure 114 on page 255, when the DHCP Server is up and running, and a dynamic IP client tries to get dynamic IP configuration information.

```
     IBM TCP/IP for WARP Server
  Dynamic Host Configuration Protocol
              Server

  Version:   3.1
  Released:  Nov 14 1995   16:31:23


Server Initialized at Sun Dec  8 12:05:05 1996
Request From:     1-0x08005a2161ae
       Type:     DISCOVER
       Status:   Offering address to the client - REPLY OFFER
       IP Addr:  9.67.20.111
       Options:  1 6 9 15 44 45 46 51 77
Request From:     1-0x08005a2161ae
       Type:     REQUEST
       Status:   Requesting a reserved address - REPLY ACK
       IP Addr:  9.67.20.111
       Options:  1 6 9 15 44 45 46 51 77
```

Figure 114. DHCP Server Window

The requesting dynamic IP client now sort of owns the offered and acknowledged IP address. The next time the dynamic IP boots up, it will request the same address and will get it if still available. The message that would be displayed on your DHCP Server console, is shown in Figure 115.

```
Request From:     1-0x08005a2161ae
       Type:     REQUEST
       Status:   Requesting an existing lease - REPLY ACK
       IP Addr:  9.67.20.111
       Options:  1 6 9 15 44 45 46 51 77
```

Figure 115. DHCP Server Console Window

If the DHCP server returns a `No addresses available` message back to you as shown in Figure 116 on page 256, please refer to 11.14, “Dynamic TCP/IP Client Programs in Warp 4” on page 265 for more information about how to set up dynamic IP clients for certain classes. Remember, we set the DHCP Server up for supporting classes.

```
Request From:     1-0x08005a2161ae
       Type:     DISCOVER
       Status:   No addresses available for the client - NO REPLY
```

Figure 116. DHCP Server No Address Available Message Console Window

To check what IP addresses are available, not available (N/A), or assigned to IP clients, type the `DSTAT` command from an OS/2 command prompt. The following information will be presented to you as shown in Figure 117.

```
Status of DHCP server maverick (9.67.20.120) as of Sun Dec  8 12:19:17 1996
IP Address     Status   Lease Time  Start Time   Last Leased  ClientId
9.67.20.110    N/A
9.67.20.111    Leased    24:00:00  12/09 09:37  12/09 09:37  0x08005a2161a
9.67.20.112    Free
9.67.20.113    Free
9.67.20.114    Free
9.67.20.115    Free
9.67.20.116    Free
9.67.20.117    Free
9.67.20.118    Free
9.67.20.119    Free
9.67.20.120    N/A
9.67.20.121    Free
9.67.20.122    Free
9.67.20.123    Free
9.67.20.124    Free
9.67.20.125    Free
9.67.20.126    Free
9.67.20.127    Free
9.67.20.128    Free
9.67.20.129    Free
9.67.20.130    Free
```

Figure 117. Executed DSTAT Command Window

11.13.2 Configuring and Using DDNS Server

1. Within the TCP/IP folder, open the DDNS Server Services window as shown in Figure 118 on page 257.

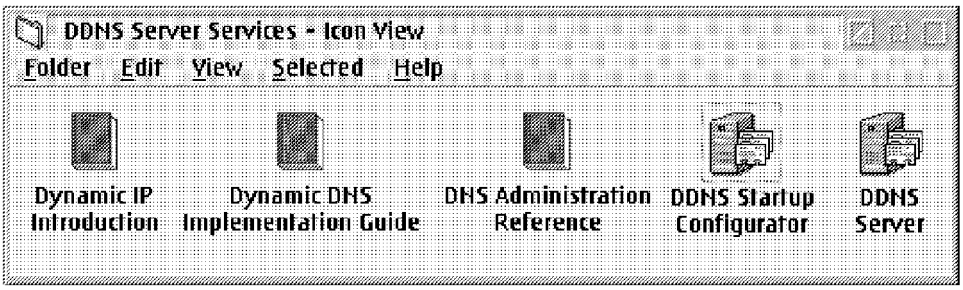

Figure 118. DDNS Server Services Window

If you do not find the DDNS Startup Configurator, you may get it from the Download section of the Warp Server home page at:

`http://www.software.ibm.com/os/warp-server/`

The DDNS Startup Configurator eases the process of configuring a Dynamic DNS server tremendously. If you plan to have both DHCP Server and DDNS Server running on the same machine, you can use that utility. If that is not the case, after creation of the DDNS Server you will be prompted to do manual work on the digital signature files: DHCPSD.DAT and DDNS.DAT

2. Start the configurator by double-clicking on the object. You will be presented with the DDNS Startup Configurator window as shown in Figure 119.

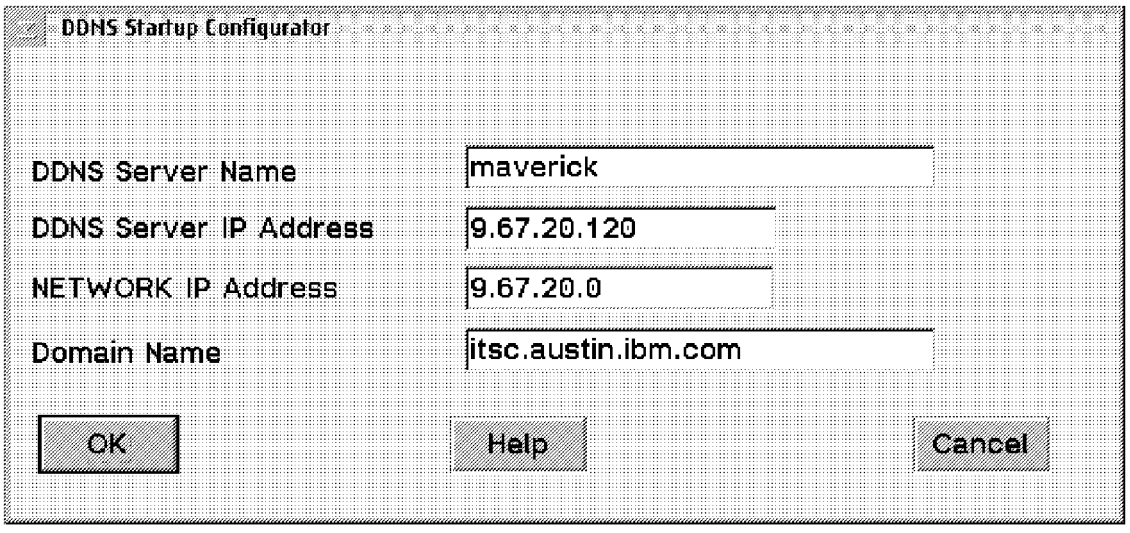

Figure 119. DDNS Startup Configurator Window

3. Press **OK** at the DDNS Startup Configurator window as shown in Figure 120 on page 258 to have the utility create the public keys as well as the NAMED files.

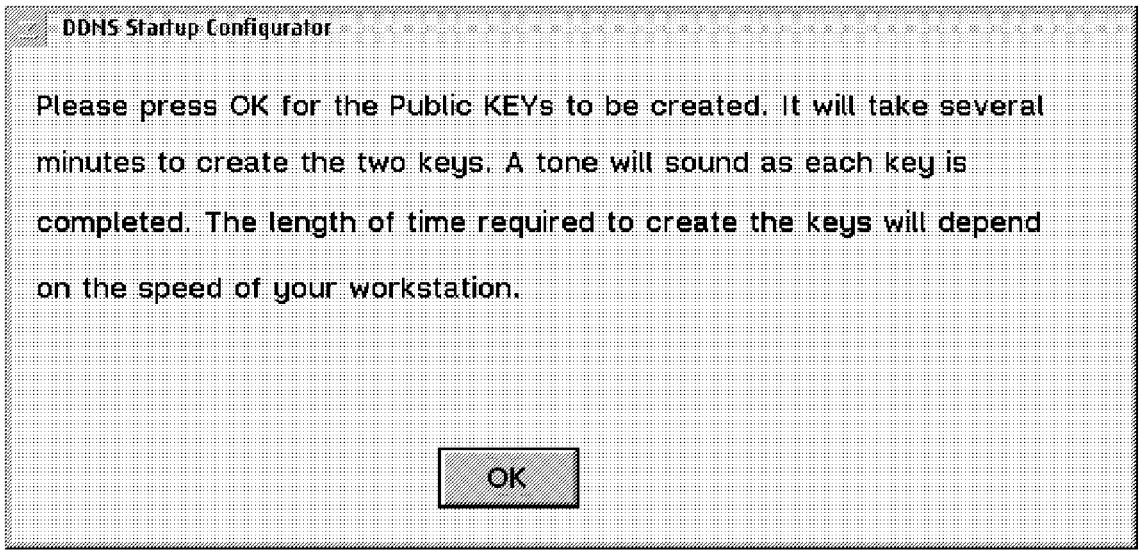

Figure 120. DDNS Startup Configurator Confirmation Window

4. Press **OK** at the Successful Completion window as shown in Figure 121.

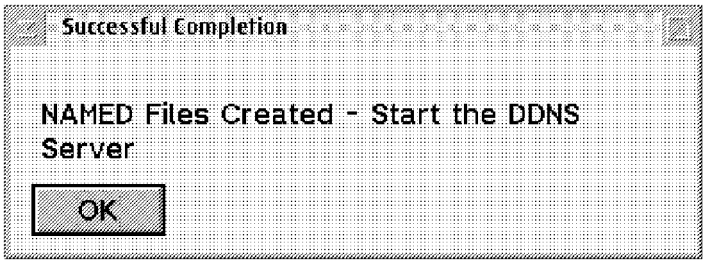

Figure 121. DDNS Startup Configurator Successful Completion Window

Note: If you do not have DDNS Server installed, you will get a pop-up window informing you that you need to have DDNS Server installed. If this is your case, reinstall Warp Server TCP/IP Services with the checked radio button for DDNS Server.

The DDNSSUP utility created three NAMED files for you:

1. NAMED.BT

```
;NAMED.BT file for name server configuration - created by DDNSSP
;
;TYPE         DOMAIN              SOURCE FILE of HOST
primary itsc.austin.ibm.com         c:\\mptn\\etc\\namedb\\named.dom dynar
primary 20.67.9.in-addr.arpa        c:\\mptn\\etc\\namedb\\named.rev dynar
```

Figure 122. NAMED.BT File

2. NAMED.DOM

```
$ORIGIN austin.ibm.com.
itsc          IN     NS     maverick.itsc.austin.ibm.com.   ;Cl=4
              IN     KEY    0x0080   0   1 AQO9VeS1tzum5T1qcDlQ+uASx7nGgybnnyQUl
              IN     SOA    maverick.itsc.austin.ibm.com. maverick.itsc.austin.
              44 86400 300 86400 3600 300 )   ;Cl=4
$ORIGIN itsc.austin.ibm.com.
ns-updates     IN     CNAME   maverick.itsc.austin.ibm.com.   ;Cl=4
maverick       IN     A      9.67.20.120    ;Cl=4
```

Figure 123. NAMED.DOM File (Extract)

3. NAMED.REV

```
;
20.67.9.in-addr.arpa.    IN  SOA  maverick.itsc.austin.ibm.com. maverick.itsc.au
      42    ; Serial number for this data (yymmdd##)
      86400 ; Refresh value for secondary name servers
      300   ; Retry value for secondary name servers
      86400 ; Expire value for secondary servers
      3600  ; Mimimum TTL value
      300 ) ; dynamic update increment time
20.67.9.in-addr.arpa.    IN  NS   maverick.itsc.austin.ibm.com.
20.67.9.in-addr.arpa.    IN  KEY  80 0 1 AQO4vn7wyXTNtsDom2qIJduQjWioIGhU1sBBtH9

; Canonized List
120.20.67.9.in-addr.arpa.    IN  PTR  maverick.itsc.austin.ibm.com.
```

Figure 124. NAMED.REV File (Extract)

Hint

All three NAMED files are very crucial for the DDNS Server to run properly. We recommend that you do backups of these files regularly.

Start DDNS Server by double-clicking on its object. The following information will be presented to you as shown in Figure 125.

```
 ------------------------------------------
| IBM OS/2 Warp Domain Name Server (NAMED) |
|          TCP/IP Version 3.1              |
 ------------------------------------------

bootfile = C:\MPTN\ETC\NAMEDB\NAMED.BT
```

Figure 125. DDNS Server Window

Using DDNS as a Static DNS Server

Although Warp Server's DNS Server is a dynamic one, you also can configure it as a static DNS Server if necessary. You can even migrate an AIX DNS environment to Warp Server's DDNS environment just by renaming the NAMED.Boot file to NAMED.BT.

11.13.3 NetBIOS Name Server Shadow

There is no denying the movement away from other communications protocols to TCP/IP which among other things is, of course, the protocol used on the Internet. Even those customers who may not be moving today are almost certainly planning for the transition tomorrow. Now, there are several challenges in taking NetBIOS applications such as Warp Server onto TCP/IP, not the least of which is how to keep track of names — both IP addresses and NetBIOS names — as the network grows, and as individual users sign on and log off.

Warp Server introduced many innovative features. The dynamic IP approach is a boon to network administrators in automating the control, allocation, and identification of IP addresses. Most people agree that this is a superior solution and the trade press is beginning to sing its praises. On the NetBIOS side Warp Server added support for the mapping of NetBIOS names in the Domain Name Servers for the provision of Local Caches to hold NetBIOS Names, and for "Broadcast Files", which reduce NetBIOS broadcasting. All these additions ease the administration of the NetBIOS names and provide a satisfactory solution in many customer situations.

However, there is another approach to NetBIOS name control that many will agree is more powerful: namely the provision and use of a NetBIOS Name Server (NBNS). The concept is quite simple and was anticipated in the Standards Committee's RFCs 1001/1002. The idea is to have a centrally located and permanently available server on the network where all the NetBIOS names are registered and controlled. New users wishing to register a new name need go to only one place to discover if the name already exists. And any user wishing to locate an application — such as Warp Server — can go immediately to the NBNS and look up its NetBIOS name and corresponding IP address so that it can access the application quickly and easily. Few people will dispute that as networks become larger and more complex, the employment of a NBNS becomes almost mandatory. The criteria for this need will be the subject of another news-sheet.

Warp Server doesn't have a built-in NBNS; Micosoft's NT does — known as the Windows Internet Name Server (or WINS). So Microsoft will be pushing

the NBNS solution even when existing Warp Server approaches may be quite adequate.

However, Warp Server users have another choice, in fact a superior choice. The only other NBNS available commercially today is a product called "Shadow" from Network TeleSystems (NTS) of Sunnyvale, California. NTS' CEO, John Davidson, is a pioneer in this area, having worked on the ALOHA System, ARPANET, and as Chief Technology Officer of Ungermann-Bass Networks. The team he has assembled at NTS is the acknowledged leader for NBNSs. The Shadow product is fundamentally different to WINS in its architecture and capabilities and is designed to serve the needs of the very largest and most demanding networks.

NTS is an IBM partner. They are working with IBM to provide basic education and information about NBNS and Shadow in particular. They are available to visit IBM customers who may be interested in an NBNS, and will discuss configurations, pricing, and other details that the customer needs to know.

For more information (for example a White Paper on NetBIOS Name Server) please visit the NTS Web Site at:

`http://www.nts.com`

or contact:

- Eastern US/Canada

  Regional Sales Manager, Jim Cooney:

  Phone: (617) 944-3220; Fax: (617) 944-8335; e-mail: jcooney@nts.com
- Central US

  Regional Sales Manager, Bill Moore:

  Phone: (972) 663-9383; Fax: (972) 773-1101; e-mail: bmoore@nts.com
- Western US/Can

  Regional Sales Manager, Hal Kroeger:

  Phone: (408) 523-6334; Fax: (408) 523-8118; e-mail: kroeger@nts.com
- International (All other countries):

  Contact NTS Corporate HQ in California, US for nearest distributor

  Phone: (408) 523-8100; Fax: (408) 523-8118; e-mail: sales@nts.com

11.13.3.1 Installing Shadow

This section describes how to install the Shadow server. First install the product by doing the following steps:

1. Install PC DOS 7 (or 6.1 or higher) on an ISA-based machine with an IDE hard disk, set up for one primary partition (the whole disk must be configured as one primary partition) with either an Eagle NE2000 adapter, available from Microdyne, or an IBM 16/4 (or Auto 16/4) Token-Ring adapter, and 8 to 16 MB of memory installed (16 MB of memory is needed to support up to 64000 names).
2. Do not install any networking software. Shadow comes with its own TCP/IP stack.
3. Install the Shadow product (a single diskette). The product replaces the CONFIG.SYS file as well as the AUTOEXEC.BAT file. No High Memory Management drivers are allowed to be loaded since the Shadow server directly communicates with the Network Adapter.

11.13.3.2 Configuring Shadow

Adapt the SHADOWER NTS-NBNS.CFG file. For an Ethernet environment with ThinNet Cabling System, the file may contain the following lines (especially notice the highlighted lines):

```
# NTS-NBNS configuration file
OUTAHERE
NODIAGS
NORESET
NOLOG
NOSAVE

# hardware

#NIU IBMTR
#IO-PORT 0A20
#WINDOW 0D000
```

Figure 126 (Part 1 of 2). Shadow's Configuration File NTS-NBNS.CFG

```
NIU  NE2000
IO-PORT 300                    # NE2000 IO port
WINDOW 0D000                   # NE2000 memory window
#CONNECTOR TPI
CONNECTOR THINNET
#CONNECTOR THICKNET
#CONNECTOR STARLAN

#NIU     GpCniu                # Adapter type
#NIU     PCniuEX               # Adapter type
#IO-PORT  368                  # NIU IO port
#WINDOW  0D000                 # NIU memory window

#NIU  NONE
#MACADDR 0123456789AB
# IP

broadcastaddr host-1s
NETSUBNETMASK   255.255.255.0
arptimeout 30
IPADDR 9.67.20.115
GATEWAYADDR 9.67.20.120
TTL 8
SCOPE NULL
#BACKUPADDR 82
#COSERVER 82
#en chante
```

Figure 126 (Part 2 of 2). Shadow's Configuration File NTS-NBNS.CFG

Notes:

1. Lines starting with a pound sign (#) are read as a comment.
2. GATEWAYADDRESS in this case is the DDNS Server.

11.13.3.3 Configuring Warp Server for Shadow

Usually, servers are not configured as DHCP and DDNS clients. They should have static TCP/IP information. It does not make sense if a server boots up with a different IP address than the one it had before. However, to enable registering at the Shadow server, you need to make changes to Warp Server's IBMLAN.INI file as shown in Figure 127 on page 264.

```
[tcpbeui_nif]

   DriverName = tcpbeui$
   Bindings = ,MACETH_nif
   NODETYPE = "H-Node"
   NBNSADDR = "9.67.20.115"
   NBDDADDR = "9.67.20.115"
   OS2TRACEMASK = 0x0
   SESSIONS = 130
   NCBS = 225
   NAMES = 21
   SELECTORS = 15
   USEMAXDATAGRAM = "NO"
   NETBIOSTIMEOUT = 500
   NETBIOSRETRIES = 2
   NAMECACHE = 1000
   PRELOADCACHE = "NO"
   NAMESFILE = 0
   DATAGRAMPACKETS = 20
   PACKETS = 50
   INTERFACERATE = 300
```

Figure 127. Warp Server's TCPBEUI Section of IBMLAN PROTOCOL.INI

Notes:

1. You also may use the MPTS configuration utility to do the highlighted modifications in PROTOCOL.INI's TCPBEUI section.
2. Warp Server OS/2 clients and Warp 4 clients can be configured the same way unless those machines are set up as dynamic TCP/IP clients, who retrieve NetBIOS Name Server information from Warp Server's DHCP Server.

11.13.3.4 Running Shadow

The Shadow server comes up automatically once the machine is powered on. If TCPBEUI clients are set up correctly with the NetBIOS Name Server information, they will register their NetBIOS names at the Shadow server. The Shadow server maps those NetBIOS names with IP addresses. An example is shown next.

```
 Network TeleSystems Inc. SHADOW Server              3:09:00pm Sat  7Dec
 LOG    NBNS    DHCP  RADIUS  DNS  WEB  Cfg  Inf
 Names:     6, Entries:     6, Name Scroll:     0, Entry Scroll:     0
CHALLENGER-----[00]                    gh  9.67.20.120    6Dec96-15:08
                                     gh  9.67.20.111    6Dec96-15:08:14
SHADOW----------                       sgh  9.67.20.120    Static
SKYWALKER------[00]                     uh  9.67.20.120    6Dec96-15:08
SKYWALKER------[03]                     uh  9.67.20.120    6Dec96-15:08
SKYWALKER-------                       uh  9.67.20.120    6Dec96-15:08:
STARTREK-------[00]                     uh  9.67.20.111    6Dec96-15:08
STARTREK-------[03]                     uh  9.67.20.111    6Dec96-15:08
STARTREK--------                       uh  9.67.20.111    6Dec96-15:08:
SYDNEY---------[03]                     uh  9.67.20.111    6Dec96-15:08
UWE------------[03]                     uh  9.67.20.120    6Dec96-15:08
```

In this example, you will notice that domain names, server/requester names and user Ids are mapped with its unique IP address. For example, CHALLENGER is the domain name and is registered as a group name with the 16th byte of hex 00. SKYWALKER is the server name and is registered twice, both as unique names but each comes with a different 16th byte: hex 00 and hex 03. UWE is an user ID, therefore registered as a unique name with hex 03 as the 16th byte. All those names are mapped with the IP address of 9.67.20.120.

The dynamic IP client with the requester name of STARTREK and the logged-on user SYDNEY also got registered. STARTREK is registered twice, both as unique names but each comes with a different 16th byte: hex 00 and hex 03. SYDNEY is registered as a unique name with hex 03 as the 16th byte.

11.14 Dynamic TCP/IP Client Programs in Warp 4

Besides Warp Server OS/2 clients, Warp 4 clients also support dynamic TCP/IP configuration. When the client requests dynamic TCP/IP configuration information, the server would note that the client belongs to the IBM_MOBIL_CLIENTS class. Therefore, the client would be provided with class-specific configuration information. In our scenario, the DHCP Server as set up in Figure 102 on page 244, the clients who belong to the class IBM_MOBIL_CLIENTS will be provided with the following IP information:

1. An allocated IP address that was assigned to the client, who belongs to the IBM_MOBIL_CLIENTS class, which is 9.67.20.111, with a lease time of 24 hours (retrieved by the server configuration)

.From the definitions made in the Network Default Options section

2. The domain name, which is: itsc.austin.ibm.com
3. The IP address of the LPR Server, which is: 9.67.20.120
4. The IP address of a DNS Server, which is: 9.67.20.120
5. Subnet mask information, which is 255.255.255.0

   From the definitions made in the class IBM_MOBIL_CLIENTS
6. The TCPBEUI Node type, which is: h-node
7. The IP address of the Datagram Distribution Server, which is 9.67.20.115
8. The IP address of the NetBIOS Name Server, which is 9.67.20.115

Since the client was provided with the IP address of the dynamic DNS Server (9.67.20.120), the client will exchange information about its host name with the DNS Server. Information updated in the `NAMED.DOM` file : host name to IP address

The DHCP Server takes care of the reverse mapping: IP address to host name.

11.14.1 Configure The Dynamic IP Client for User Class Support

At the Dynamic IP client, the MPTN ETC DHCPCD.CFG file contains information about the client's class. The default class name is: IBMWARP_V3.1

To configure the client for a different class, you need to modify the `User Class` entry. Edit the file by using an ASCII editor, search for the the line that starts with `option 77` (line 273). Change the entry with the appropriate class name used in our scenario:

```
option 77 "IBM_MOBIL_CLIENTS"          # User Class
```

When done, this particular client can request IP configuration information from the DHCP Server configured in this scenario.

11.14.2 Warp 4 Dynamic IP Utilities

At the Warp 4 client, open the System Setup folder, for example from the context menu you get by clicking on the right mouse button on the Workplace Shell. Besides other utilities, you will be presented with:

- DHCP Client Monitor
- DDNS Configuration

The DDNS Configuration utility will be presented to you automatically once you configured TCP/IP for DHCP and DDNS support as shown in Figure 128 on page 267.

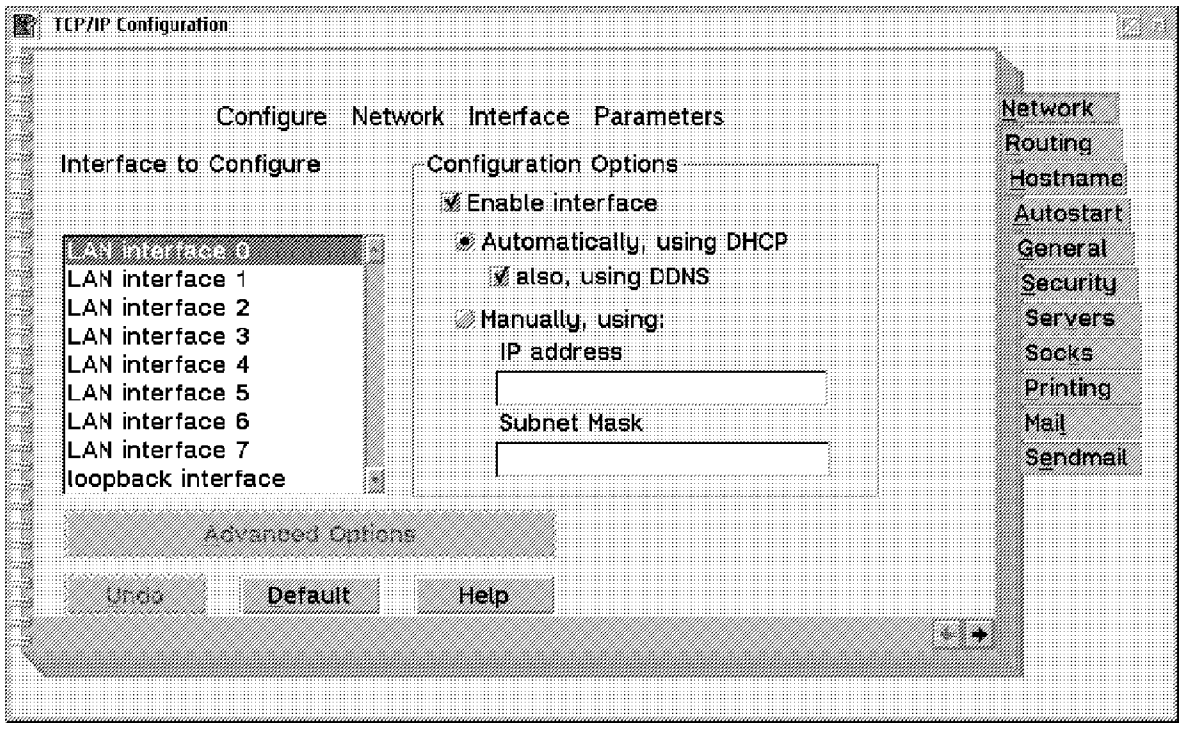

Figure 128. TCP/IP Configuration Window

That basically is everything you need to do to set up a dynamic IP client. All TCP/IP configuration information will be delivered by the DHCP Server.

11.14.3 DHCP Client Monitor

At the Warp 4 client, open the System Setup folder and start the DHCP Client Monitor program which is shown in Figure 129 on page 268.

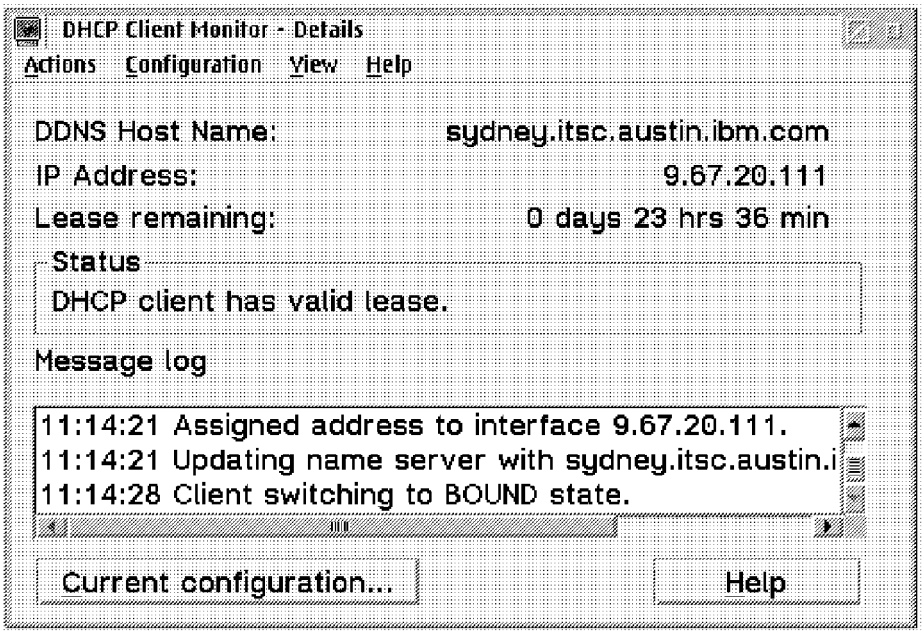

Figure 129. DHCP Client Monitor Window

At the main panel of the DHCP Client Monitor, you will be informed about your host name, IP address, and remaining Lease time of your IP address.

From the Configuration pull-down menu you can get additional configuration information, for example which DNS server is serving you and other things as shown in Figure 130.

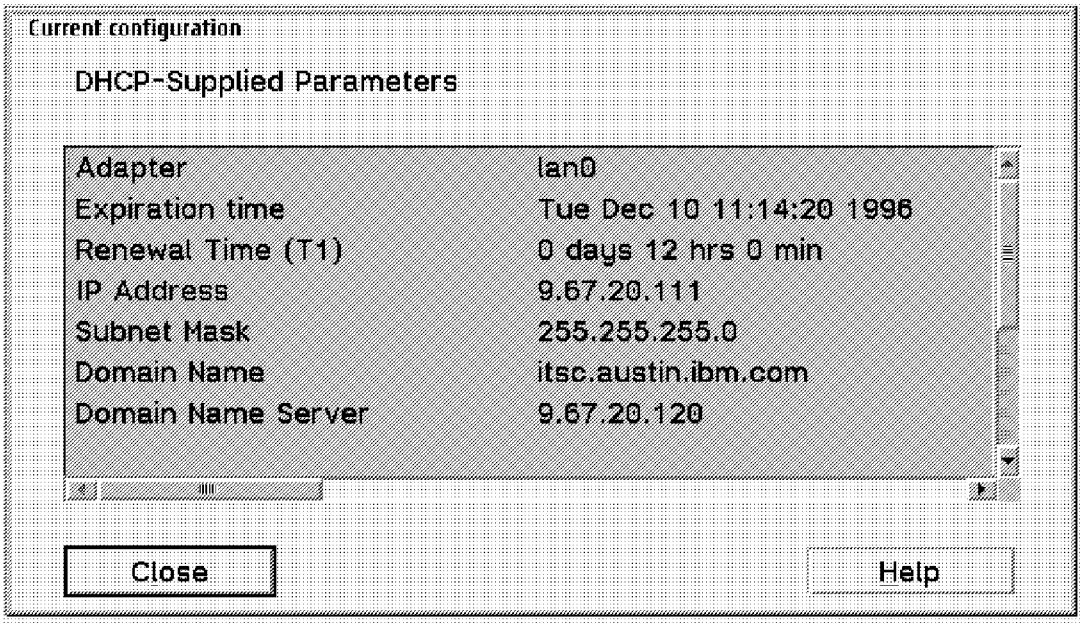

Figure 130. DHCP Client Monitor Current Configuration Window

Note: Renewal time of a leased IP address always is half of the defined leased time. In our scenario, the lease time was set to 24 hours. Therefore, the client will ask for IP address renewal after 12 hours have been expired.

11.14.4 DDNS Client Configuration

At the Warp 4 client, open the System Setup folder and start the DDNS Client Configuration program that is shown in Figure 131.

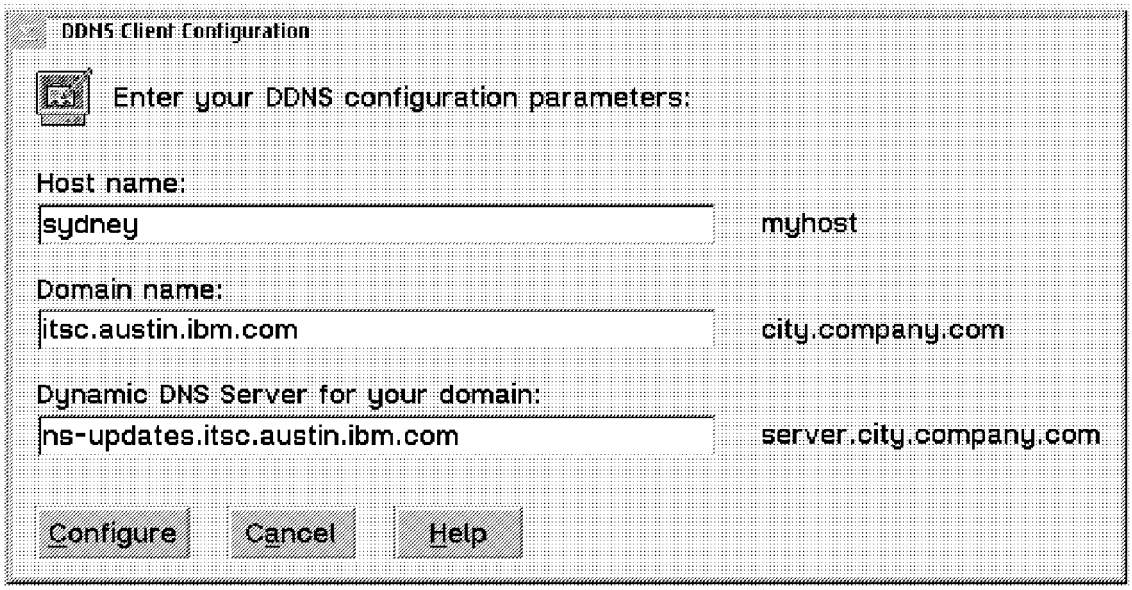

Figure 131. DDNS Client Configuration Window

Note: Domain name and Dynamic DNS Server information is filled out automatically for you.

Again, this is the program that will be presented to you once you have configured your client as a dynamic IP client and you are booting up the machine for the first time after those configuration changes were made.

Once you have typed in a host name and clicked on **Configure**, you are the owner of that host name, assuming successful completion of your host name configuration and that this host name was not taken by somebody else already.

A file named DDNS.DAT will be copied from the DDNS Server to your workstation's MPTN ETC directory. This file contains a digital signature and makes sure that only you will be allowed to change your host name if necessary.

Notes:

1. If you think, this kind of security is not good enough, you can take advantage of having the DDNS.DAT file made available to users who are actually authorized for dynamic DDNS updates. In this case you would supply the DDNS.DAT file on a diskette or on a user's home directory to establish a pre-secured environment for the dynamic DDNS.
2. To make sure that only authorized clients can retrieve dynamic IP addresses, you can establish TCP/IP classes, which is shown in Figure 109 on page 250.

In Figure 132 you can notice, that we changed the host name from sydney to melbourne successfully.

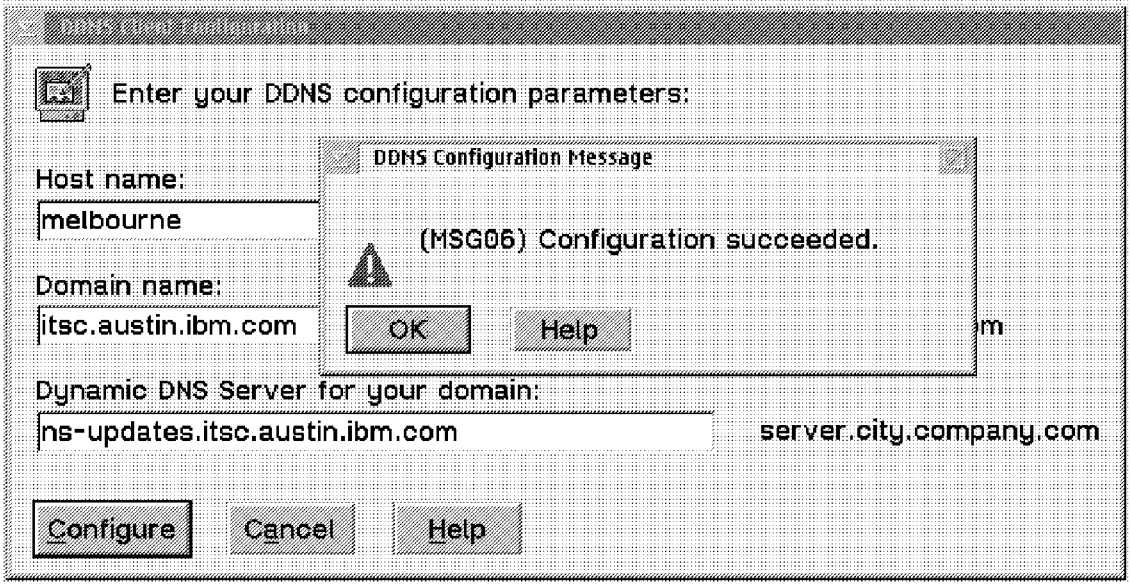

Figure 132. DDNS Configuration Successful Completion Window

From now on, your workstation's host name will be melbourne. However, the former host name, sydney, will be set inactive with the next renewal of the dynamic IP address. This will allow you to inform other people about your host name change and your former host name will not be known to the system anymore.

In case the dynamic IP client does not communicate with the NetBIOS Name Server, although it has gotten all NBNS information from the DHCP Server (see Figure 110 on page 251, Figure 111 on page 252, and Figure 112 on page 253), you may modify the client's PROTCOL.INI file, and insert NBNS information in the TCPBEUI section as shown in Figure 127 on page 264. We experienced no problems with this workaround.

11.15 Introducing TME 10 NetFinity Server

The SystemView for OS/2 Server, originally shipped with Warp Server 4.0 has been completely replaced by TME 10 NetFinity Server 4.0. In 1996 customers were offered to upgrade to TME 10 NetFinity platform at no charge. In 1997 there might be an upgrade charge for that.

TME 10 NetFinity Server provides a systems management solution for a LAN workgroup environment. This type of environment is typical of a small company, or a department within a larger enterprise, that needs workstation interconnection to enable sharing of application data and centralized systems management.

LAN workgroups can also be connected to NetView Distribution Manager for MVS (NVDM/MVS) or to Software Distribution for AIX for enterprise-wide software distribution.

A typical environment is a NetBIOS, TCP/IP, or IPX/SPX LAN with from five to 300 workstations. The workstations on the LAN can run OS/2 Warp, OS/2 Version 2, Windows 3.1 or 3.11, Windows for Workgroups, Windows 95, Windows NT, and NetWare.

Note: Note that Warp Server's systems management component supports all current workstations running on an Intel platform. TME 10 NetFinity Server Clients reside on the CD-ROM's CID CLIENT SYSVIEW2 directory.

To ensure efficient and effective delivery of application services in this type of environment, the NetFinity Server Manager operates as a central point of control that can:

- Monitor and detect problems generated from any workstation and gather relevant data for problem determination
- Track the hardware and software inventory of all workstations
- Keep the software up-to-date on all workstations
- Gain remote access to any workstation to control its processes
- Control usage of licensed software

For managed (NetFinity Server Client) systems (other than NetWare clients), each local workstation can:

- Control its own system resources in order to optimize its application processes
- Request software updates

- Share application software from other workstations
- Prepare software to be distributed by the central point of control

TME 10 NetFinity Client in Warp 4

We need to make certain that you understand that Warp 4 does not ship with TME 10 NetFinity Server Client. The TME 10 NetFinity Client does not support, for example, software distribution, remote workstation control, and IBM's AntiVirus solution. Warp 4 clients need to be upgraded from TME 10 NetFinity Client to TME 10 NetFinity Server Client in order to get those functions mentioned above.

This section covers each topic of TME 10 NetFinity Server Manager Everything that is described in this section comes with the Warp Server 4.0 (Product Refresh), TME 10 NetFinity Server Upgrade for Warp Server users, or Warp Server SMP package.

The NetFinity Server Manager includes the NetFinity Server Web Manager; a special-purpose Web server specifically designed to work with the NetFinity Server services. You can use NetFinity Server Web Manager to remotely access and manage the systems on your network from anywhere in the world by using the Internet and a Worldwide Web (WWW) browser, such as Netscape Navigator.

From the NetFinity Server Manager, as shown in Figure 133 on page 273, you can use the remote system manager service to access and control managed systems. To access the NetFinity Server Manager with a Web browser, load the manager's URL, which is:

```
http://hostname:411/main
```

assuming 411 is set up as the listening port before and the host name is maverick (as defined in the previous Dynamic TCP/IP section).

Software distribution, remote workstation control, antivirus service, and License Use Administration are not available under Web Manager. The screens for NetFinity Server services look different when running under Web Manager.

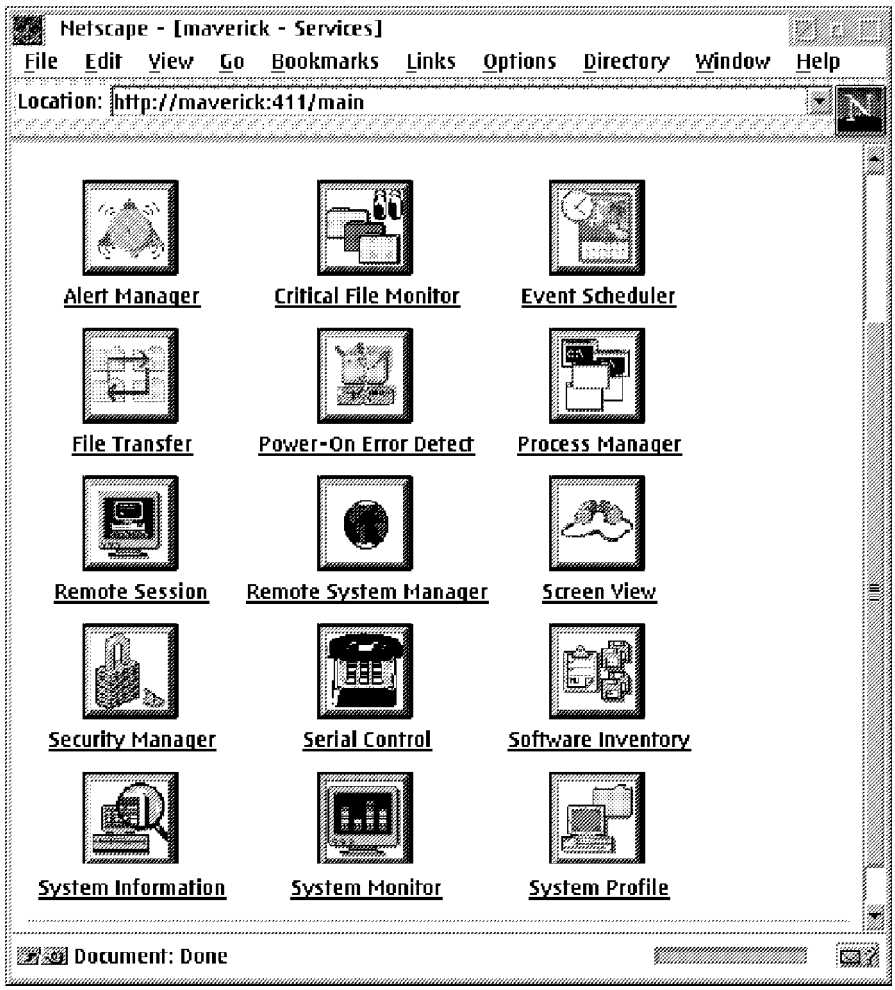

Figure 133. TME 10 NetFinity Server Manager Via Netscape Browser. The loaded URL is: http://maverick:411/main

Many functions, such as Remote Workstation Control or Remote Session require a Web manager that supports Java Script. Therefore, we recommend that you use the Netscape 2.02 for OS/2 browser in order to have those functions supported.

From now on, all major functions of TME 10 NetFinity Server are introduced by from the Netscape Web browser point of view. All basic systems management under Warp Server is covered by another redbook that is titled: *Inside OS/2 Warp Server, Volume 2: System Management, Backup/Recovery and Advanced Print Services*, SG24-4702.

1. From the main menu (`http://maverick:411/main`) select **Remote Systems Manager**. The page retrieved is shown in Figure 134 on page 274.

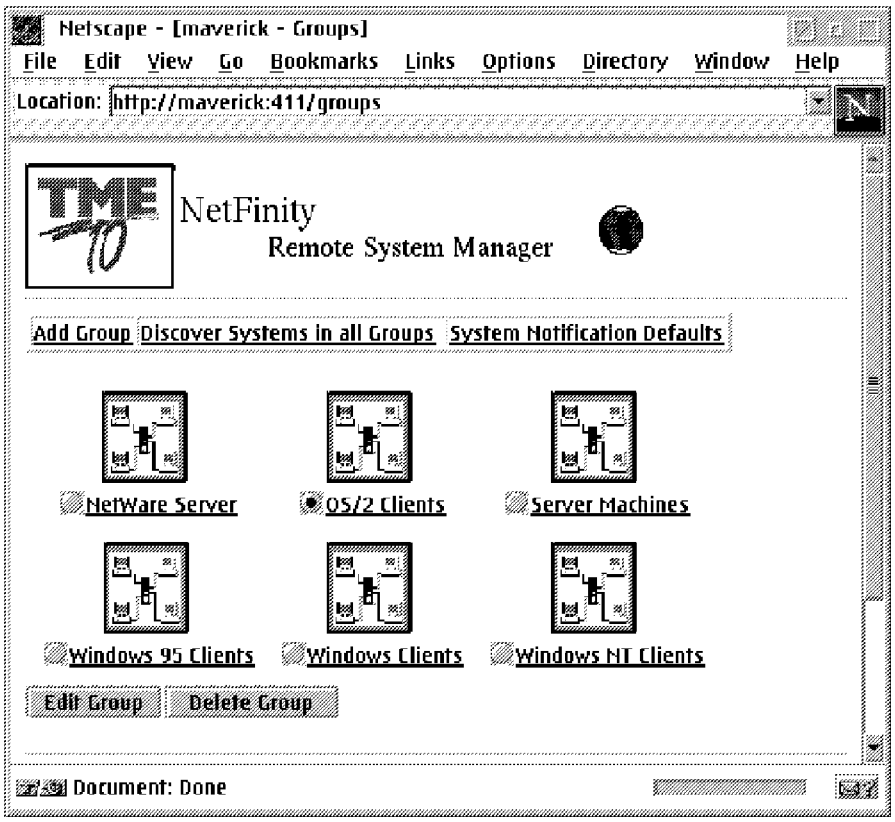

Figure 134. Remote System Manager Via Netscape Browser. The loaded URL is: http://maverick:411/groups

Note: The groups presented in Figure 134 have been created before.

2. From the Groups menu, select **OS/2 Clients**. The OS/2 Clients group window will be presented to you as shown in Figure 135 on page 275.

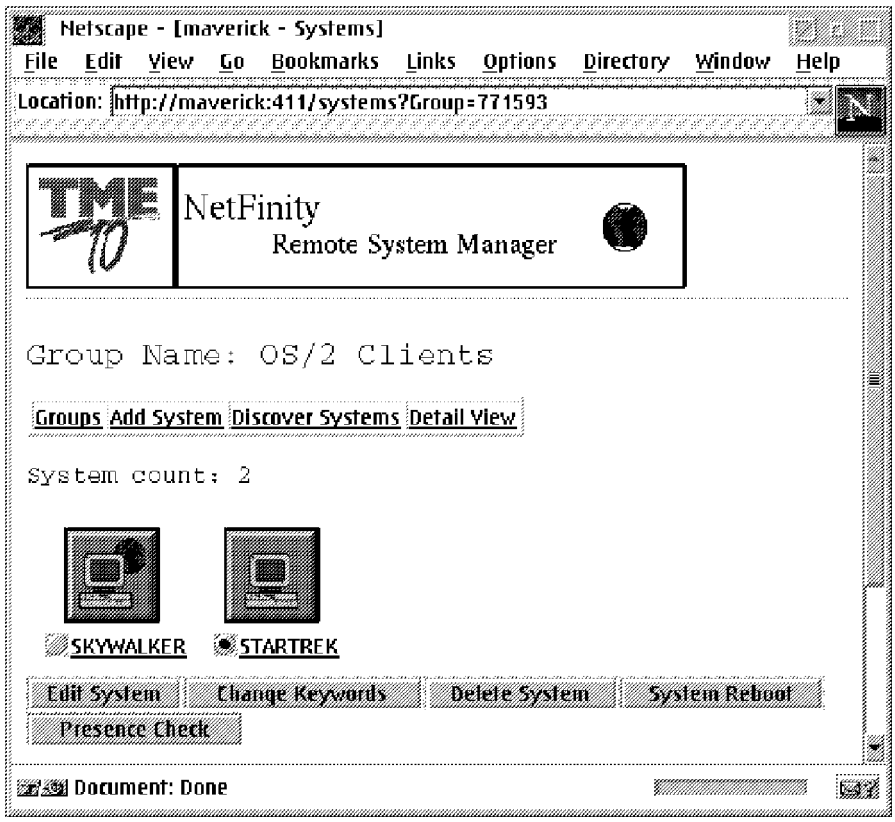

Figure 135. OS/2 Clients Group Via Netscape Browser. The loaded URL is: http://maverick:411/systems?Group=771593

Note: Only two systems will be presented to you: SKYWALKER and STARTREK. Note the two different icons. SKYWALKER has managing functions in comparison to STARTREK, which only has client functions.

3. From the OS/2 Clients Group window, select **STARTREK**. The remote system management functions of STARTREK will be presented to you as shown in Figure 136 on page 276.

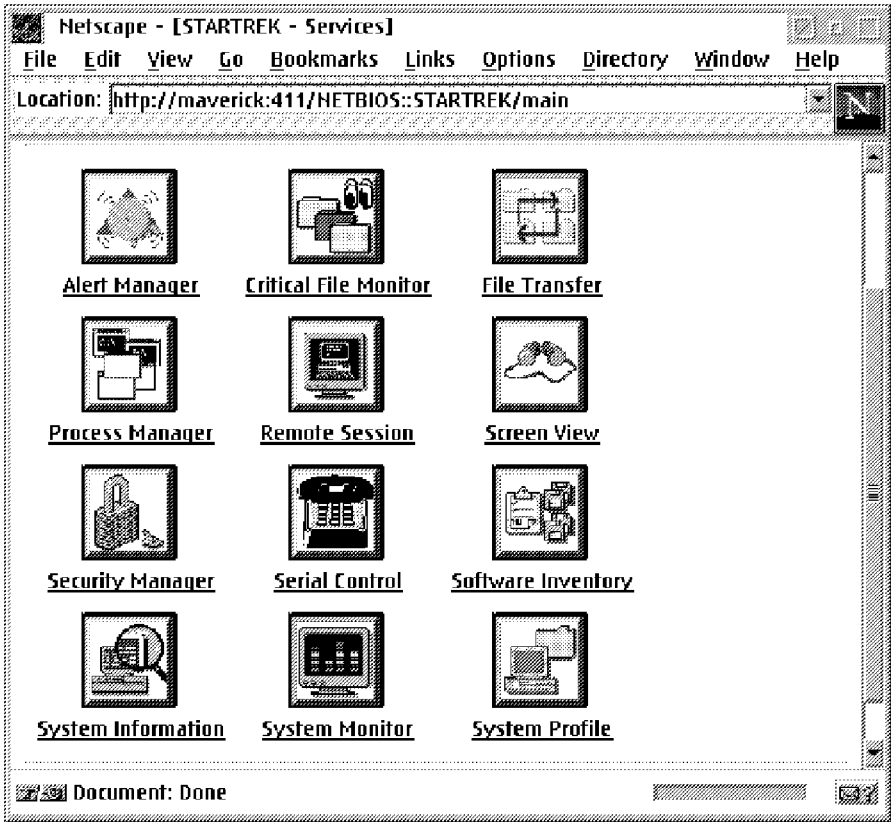

Figure 136. TME 10 NetFinity Server Client STARTREK Via Netscape Browser. The loaded URL is: http://maverick:411/NETBIOS::STARTREK/main

All presented remote system-management functions will be presented to you now.

a. The Alert Manager allows you to activate alerts and set variable alerts and take different actions, for example, generate SNMP alert, page admins, or just generate pop-up warnings. It also offers a console in which you can overview alerts that have been generated since alerts will be logged here.

   Note: Alerts can also be set on logically defined groups of systems. At the group level, alerts can be generated when individual systems go off-line and/or come back on-line. A presence check can be set to provide this function.

b. From the main menu of the client STARTREK remote systems management functions, select **Critical File Monitor**. The client specific critical files will be presented to you as shown in Figure 137 on page 277.

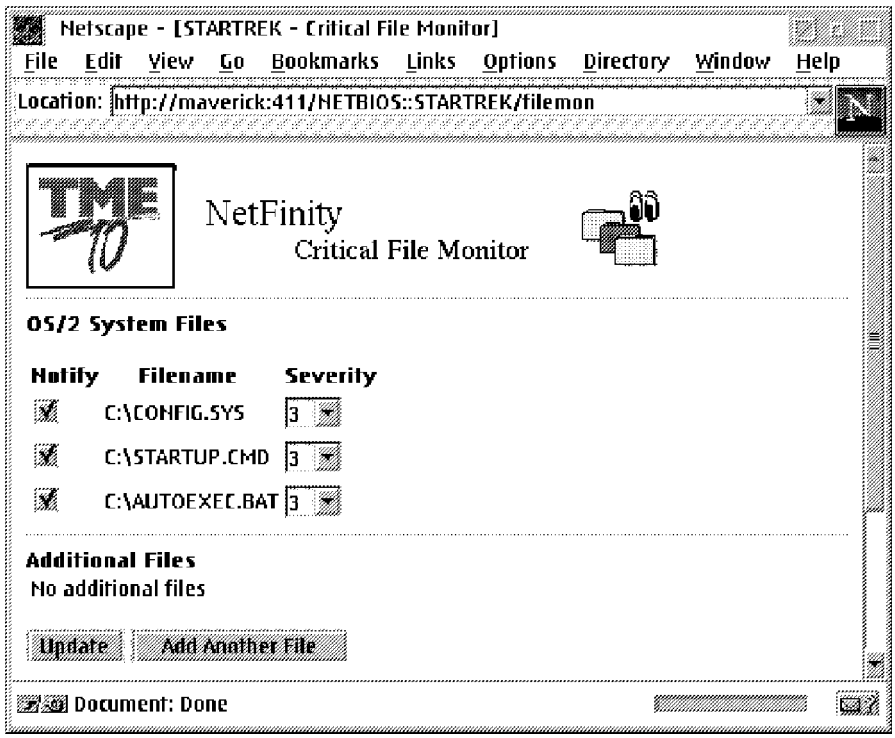

Figure 137. Critical File Monitor Via Netscape Browser. The loaded URL is: http://maverick:411/NETBIOS::STARTREK/filemon

Since the remote client is an OS/2 client, the critical files that can be selected are the CONFIG.SYS file, the AUTOEXEC.BAT file, and the STARTUP.CMD file. If this machine were a NetWare server, the critical files AUTOEXEC.NCF, STARTUP.NCF, SYS$LOG.ERR and VOL$LOG.ERR would be presented to you. Generally the most common files for each monitored operating system are easy check boxes.

If any of those files selected are deleted or modified, a system alert would occur and depending on the severity level, a systems administrator gets an alert on his/her screen or a particular system-administrator gets paged to handle that situation. You also can add additional files to be monitored.

c. The File Transfer function gives you the capability of transferring files from the managing console to the managed client and vice versa.

d. From the main menu of the client STARTREK remote systems management functions shown in Figure 136 on page 276, select **Process Manager**. The client's (which also can be a server of course) processes are presented to you as shown in Figure 138 on page 278.

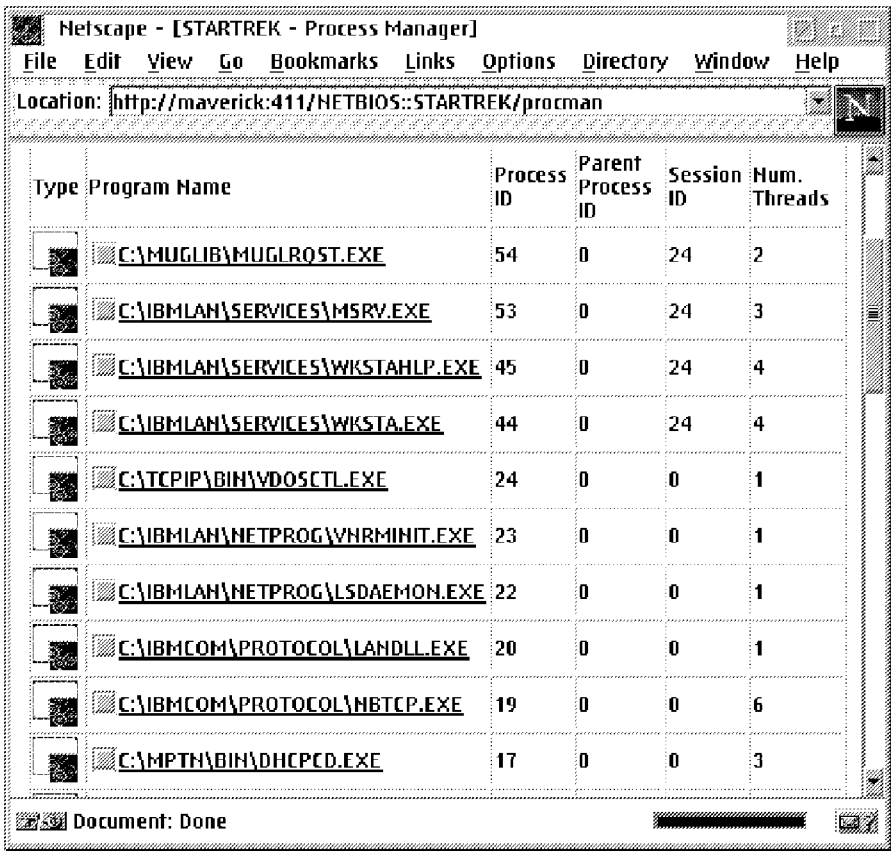

Figure 138. Process Manager Via Netscape Browser. The loaded URL is: http://maverick:411/NETBIOS::STARTREK/procman

If you want to have the client's DHCP program to be monitored, check the box C: MPTN BIN DHCKCD.EXE and link the next page. From that page, you can check the boxes for generating an alert if the program runs, stops, or does not start.

Note: The Kill Process function of the Process Manager is not implemented in the Web version. However, it is part of the native graphical user interface.

e. Remote Session creates a windowed OS/2 command prompt for you in which you can issue text-based commands as if you were local at the remote client. This function requires a Java-enabled browser, such as the Netscape browser.

f. From the main menu of the client STARTREK remote systems management functions shown in Figure 136 on page 276, select **Screen View**. The remote client's desktop is presented to you as shown in Figure 139 on page 279.

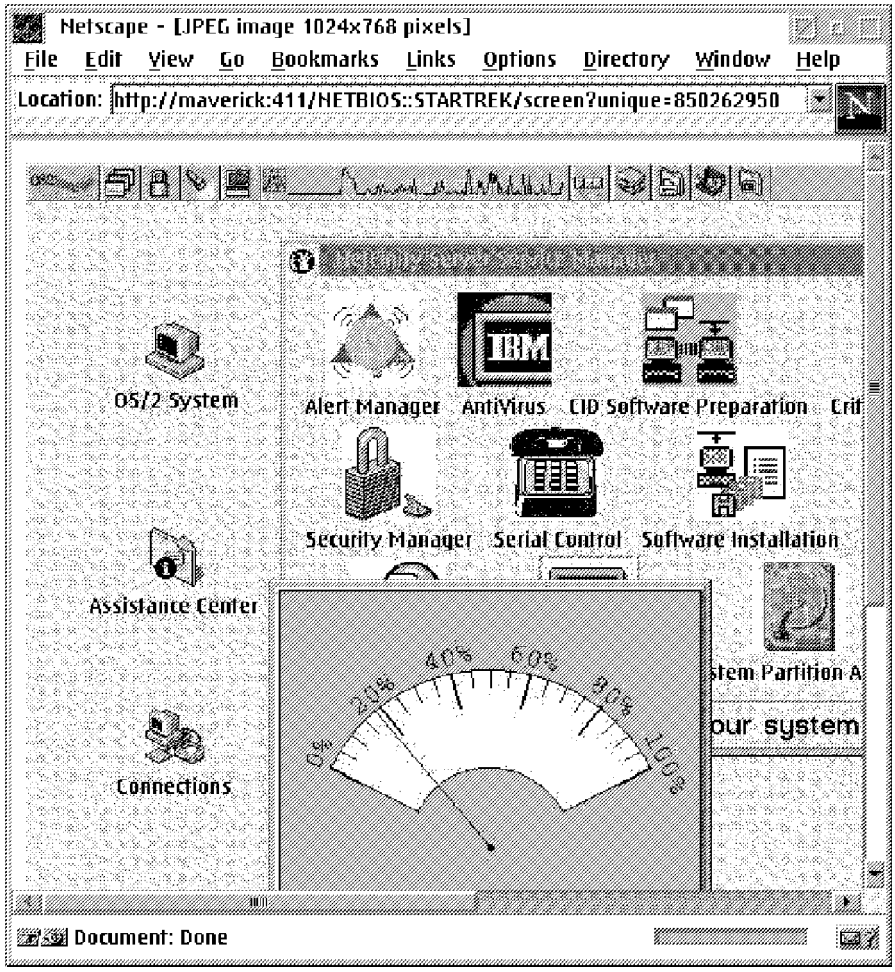

Figure 139. Screen View Via Netscape Browser. The loaded URL is: http://maverick:411/NETBIOS::STARTREK/screen?unique=850262950

Note: The transmission of the remote screen can take a while depending on the speed of the network.

A "snapshot" of the remote session is taken. This can be ideal for getting a common reference point with the user being helped by a help desk function.

g. The Security Manager provides you with the function to set up users and passwords for remote systems management. Note that once TME 10 NetFinity Server Client is installed on a machine, everyone in the network could basically do remote systems management with that machine. It is very recommendable to deselect functions for the public user, which is any user in the network. Once user IDs and passwords are defined, the administrator in charge of doing remote systems management will be prompted to log in to that remote client.

h. The Serial Control provides you with the function to set up the modem installed on the managing or managed workstation. Telephone numbers can be defined here as well in order to page the right people in case of alerts.

i. From the main menu of the client STARTREK remote systems management functions as shown in Figure 136 on page 276, select **Software Inventory**. The remote client's installed software will be presented to you as shown in Figure 140.

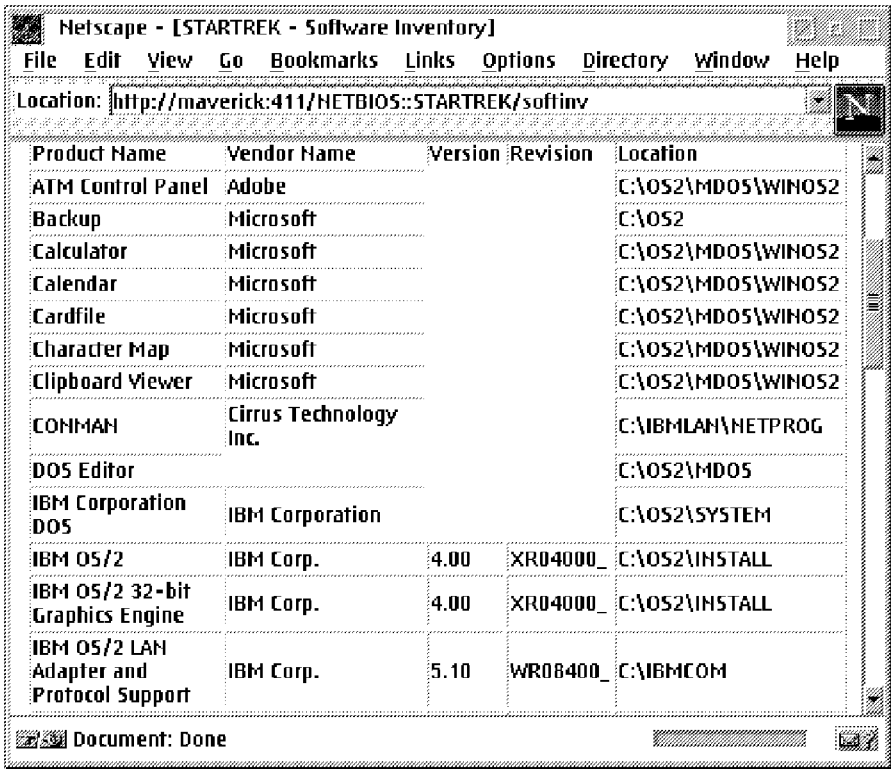

Product Name	Vendor Name	Version	Revision	Location
ATM Control Panel	Adobe			C:\OS2\MDOS\WINOS2
Backup	Microsoft			C:\OS2
Calculator	Microsoft			C:\OS2\MDOS\WINOS2
Calendar	Microsoft			C:\OS2\MDOS\WINOS2
Cardfile	Microsoft			C:\OS2\MDOS\WINOS2
Character Map	Microsoft			C:\OS2\MDOS\WINOS2
Clipboard Viewer	Microsoft			C:\OS2\MDOS\WINOS2
CONMAN	Cirrus Technology Inc.			C:\IBMLAN\NETPROG
DOS Editor				C:\OS2\MDOS
IBM Corporation DOS	IBM Corporation			C:\OS2\SYSTEM
IBM OS/2	IBM Corp.	4.00	XR04000_	C:\OS2\INSTALL
IBM OS/2 32-bit Graphics Engine	IBM Corp.	4.00	XR04000_	C:\OS2\INSTALL
IBM OS/2 LAN Adapter and Protocol Support	IBM Corp.	5.10	WR08400_	C:\IBMCOM

Figure 140. Software Inventory Via Netscape Browser. The loaded URL is: http://maverick:411/NETBIOS::STARTREK/softinv

j. The System Information function gives you information about almost everything you might be interested in: Complete hardware and OS/2 software inventory listed in great detail.

k. From the main menu of the client STARTREK remote systems management functions as shown in Figure 136 on page 276, select **System Monitor**. The remote client's monitored resources will be presented to you as shown in Figure 141 on page 281.

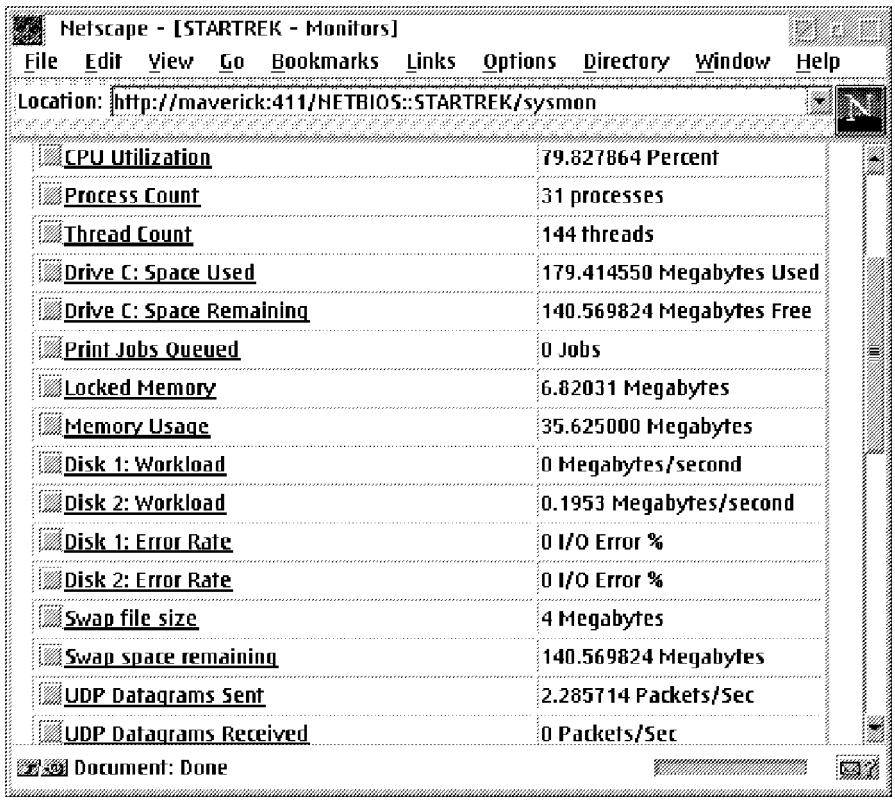

Figure 141. System Monitor Via Netscape Browser. The loaded URL is: http://maverick:411/NETBIOS::STARTREK/sysmon

Using the System Monitor function, you can set thresholds and include monitors to the Alert Manager. For example, when CPU utilization exceeds seventy percent, an alert will occur.

l. Additional information about system, user, location, contacts and miscellaneous can be manually entered in the System Profile editor so that you can set up an information file associated with a networked workstation.

m. As mentioned earlier, Remote Workstation Control is not available under Web manager. However, Figure 142 on page 282 demonstrates how Remote Workstation Control looks like from the native TME 10 graphical user interface (not using the Web browser).

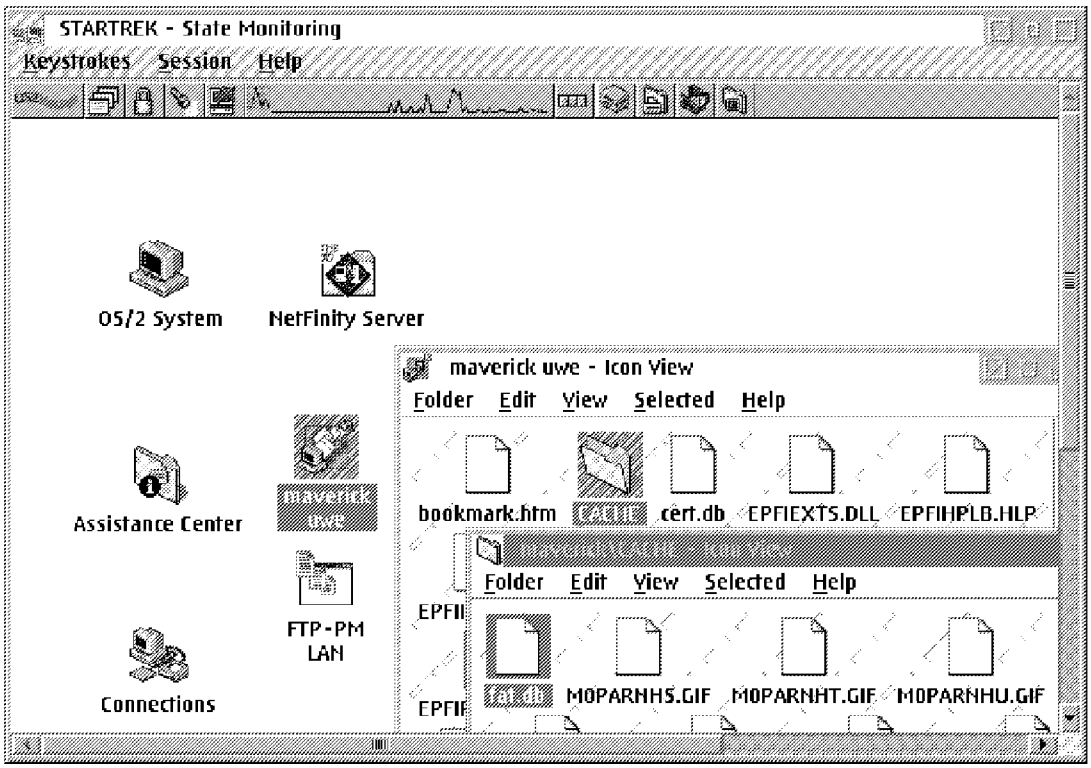

Figure 142. Remote Workstation Control

With the remote workstation control the LAN administrator at the NetFinity Server Manager workstation can control the keyboard and mouse input and monitor the display output of a client workstation without being physically present at the client.

Using remote workstation control, the LAN administrator can:

- Assist the user of the client workstation with running applications
- Do remote problem determination and diagnosis
- Operate unattended workstations remotely
- Have access to data and programs stored on a client workstation
- Remotely monitor work in progress on client workstations (for example, between teachers and students)

Remote workstation control can be useful to the operators of a help desk. For example, a user who is having trouble starting a new accounting program can contact the help desk. The help desk operator takes control of the client workstation and types the commands to run the accounting program. The user observes the

input at the client workstation and learns how to start the program. The help desk operator gives back control of the workstation, monitors the screen to be sure that the user can work independently, and ends the session.

When the remote workstation control session is in the monitoring state, the LAN administrator sees a screen image of the client workstation's display. The end-user has complete control of the operations of the workstation.

When the remote workstation control session is in the active state, the LAN administrator operates and controls the client workstation.

TME 10 NetFinity Server has linkages into Lotus Notes and DB2/2 databases that allows you to create individual views and reports. Via DDCS, information can be uploaded to a host database.

Note: To enable that support, you need to have either Lotus Notes, or DB2/2, or DDCS, or all three of them installed.

11.16 TME 10 Software Distribution

Over the past ten years, the number of workstations in organisations have grown steadily. During this time, operating systems and application software has become larger and more complex. In addition many applications require data or configuration information to be supplied at installation time. All of the above factors make the task of installing and maintaining workstation software within such organizations very difficult. OS/2 and future IBM products have been designed with the above requirements in mind. IBM has designed a method to automate these processes by using redirected input/output on LAN-based client/server systems and named it Configuration, Installation, Distribution (CID).

In this chapter we will briefly discuss the different methods available OS/2 Warp Server of distributing software. This will include software that is CID-enabled as well as the distribution of non-CID-enabled software. A very detailed discussion is put together in an IBM Redbook titled *Inside OS/2 Warp Server, Volume 2: System Management, Backup/Restore, Advanced Print Services*, SG24-4702.

11.16.1 Installation Modes

The installation techniques used to install any kind of software product are classified into three modes:

11.16.1.1 Attended Installation

Attended installation is defined as that requiring a knowledgeable individual to be in attendance at the workstation where the software is being installed. This individual will need to respond to the various prompts that are displayed during the installation and configuration process.

11.16.1.2 Lightly Attended Installation

The phrase lightly attended installation refers to an environment where and individual must be present to initiate the installation process and potentially perform other simple or predefined tasks. However, this individual would require no specialized system knowledge.

11.16.1.3 Unattended Installation

An unattended installation has no requirement for an end-user or administrator to be present at the system being installed. In this instance a Software Distribution Manager handles the initiation of the installation and everything else.

11.16.2 Configuration Installation and Distribution (CID)

The primary goals of CID are to:

- Eliminate human intervention at the target workstation when preparing and executing the configuration, installation, migration, and maintenance processes that are necessary to operate this workstation
- Enable the code executing at the target workstation to perform all required configuration and installation tasks including the integration of previous customizations
- Provide the capability to centralize human intervention to an administrator at a central preparation site

CID conceptually defines six criteria for a software product to be CID enabled:

- Response files
- Command line driven execution
- Redirected drives
- Progress indication and logging facilities
- Standard return codes
- Transfer of product diskettes

With OS/2 Warp Server, there are three methods that can be used to distribute workstation software for CID- and non-CID- enabled applications.

11.16.3 Redirected Installation

When starting a normal software installation, a user inserts a diskette or CD-ROM into a drive and starts the installation program. The product will continue to install from the drive until all the diskettes required by the installation program have been processed or the installation ends.

A Redirected installation defines the capability of the installation program to use a drive other than the diskette or CD-ROM drive. Particularly the ability of the installation program to use a logical drive letter for installation is defined.

Using this method, a workstation can access a server where the contents of the diskettes have been copied and perform the installation.

11.16.4 Product-Specific Response Files

Each product from IBM comes with utilities that create response files for you. A Response File is an ASCII file that supplies the client-specific configuration information required during redirected installation of a product on the client. The Response File contains predefined answers to the configuration questions that users are normally asked during a product installation. A Response File contains pairs of keywords and values that are interpreted during a product installation. Usually, these keyword=value pairs are unique to a particular product. Response Files are stored in their product-specific subdirectories on the connection server

11.16.5 TME 10 Software Distribution Components

TME 10 NetFinity Server Manager provides:

- Easy to use Graphical Interface
- Software preparation processes for CID and non-CID enabled applications
- Support to remove and uninstall applications are included
- Support for installation with deferred activation
- Support for installation with corequisites
- Variable translation on the target workstation at installation time
- User at workstation can initiate software installation via GUI.
- A catalog is maintained of installed software on each machine.

11.16.6 Functions of TME 10 NetFinity Server Manager

1. Software Distribution Server

   The Software Distribution Server maintains a catalog of all software objects and sharable applications. This catalog is used to process all the Software Distribution requests. Both the Software Distribution graphical and command line interfaces are available on the Software Distribution Server. From this server, software can be installed on all managed OS/2 and Windows systems. This component is installed when you select the **Software Distribution Server** when installing TME 10 NetFinity Server Manager.

2. Software Distribution User Interface (Manager)

   This user interface has both the graphical and command line interfaces available. Using this interface, this machine can distribute software to all managed OS/2 and Windows systems.

3. Software Distribution Object Preparation

   This component provides an easy way to prepare CID- and non-CID-enabled software packages. It is a selectable feature on both software Distribution Server and client machines.

4. Software Distribution Agent

   This component actually processes the install requests and is installed by default on all TME 10 NetFinity Server Clients machines. It reports the result to the Software Distribution Server and starts at machine startup.

11.16.7 Functions of TME 10 NetFinity Server Client

1. Software Distribution User Interface (Client)

   This user interface has both the graphical and command line interfaces available. However, this machine can only install software objects from the catalog if it is authorized to that object.

2. Software Distribution Base Client

   This base client does not have the Software Distribution graphical user interface present. It can only be managed by an administrator and cannot issue install requests.

3. Software Distribution Object Preparation

   This component provides an easy way to prepare CID- and non-CID-enabled software packages. It is a selectable feature on both Manager and Server Client machines. Software can be prepared for distribution on the client machines and then cataloged on the server for distribution.

4. Software Distribution Agent

   This component actually processes the install requests and is installed by default on all TME 10 NetFinity Server Clients machines. It reports the result to the Software Distribution Server and starts at machine startup.

11.16.8 Application Definition File (ADF) Considerations

The ADF file contains all the information needed to define and configure a CID-enabled software product. This file could be provided by the company that developed the application, by IBM technical support centers or it can defined by the administrator.

Your Application Definition File must contain four sections:

1. The DEF section contains some basic parameters that must be supplied for every software product.
2. The MCF section contains information that Software Distribution requires to install software.
3. The MRF section is a list of all the Response File keywords for the software product. For each keyword, you can hard-code a value that will be used in every configuration, tell CID Software Preparation to proceed to the variable section to evaluate the value, or set the value after evaluation of a conditional statement.
4. The VAR section contains all the variables specified in other sections of the application description file and specifies how they are to be assigned values in a particular configuration.

11.16.9 How Software Distribution Works

Software distribution works differently for CID-enabled and non-CID-enabled software. Although they are initiated in the same way, different processes occur behind the scenes. These are described in the figure below:

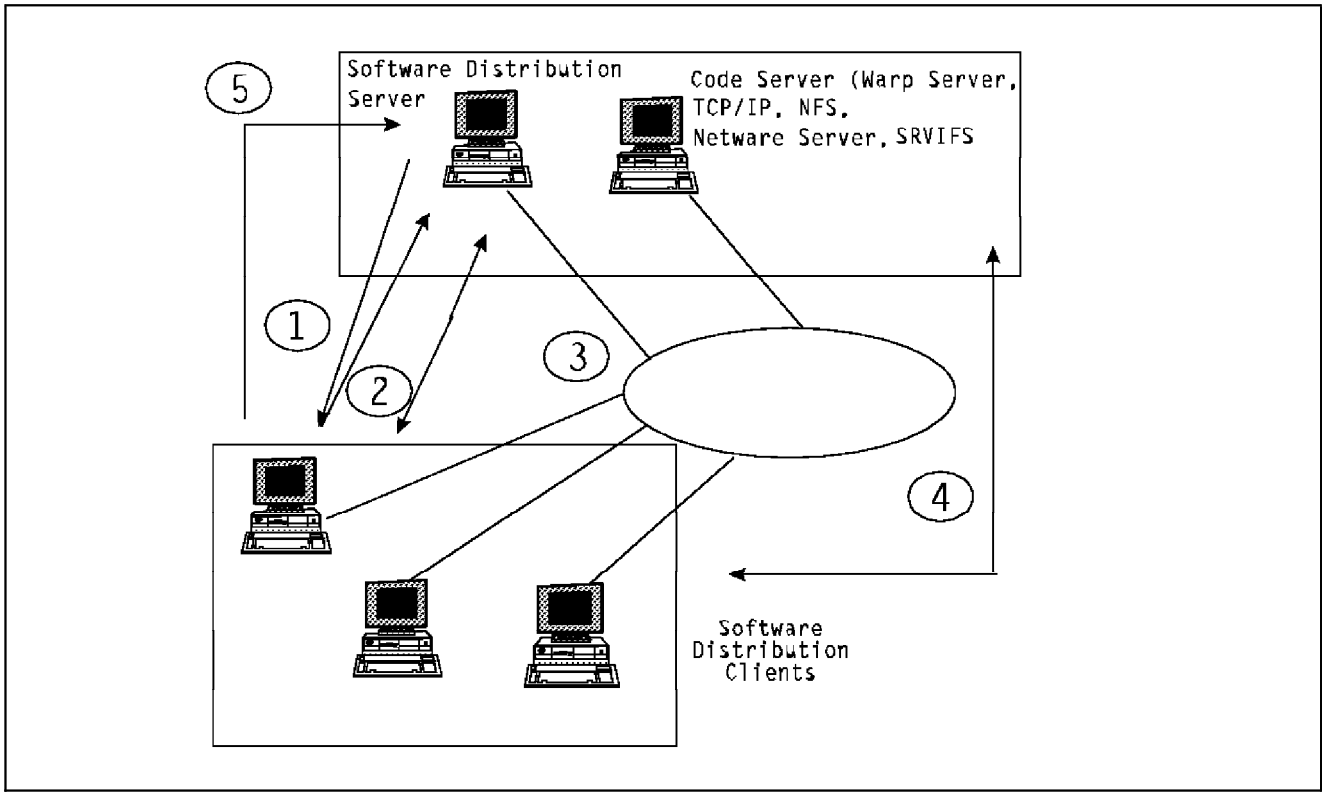

Figure 143. Software Distribution. Processes

1. First, the software distribution administrator decides on which workstation or groups of workstations to distribute to or authorize for the selected software. He/she then starts the distribution or authorization process.
2. If the workstations are authorized to software, then they have to use the software installation option to begin installation.
3. With either 1 or 2, if the software is non-CID software, the communication continues until the software is installed. Once the installation completes successfully or unsuccessfully, the catalog is updated accordingly. The installation ends.

   Should the software be CID-enabled, only control information is passed, and the rest of the installation occurs during step 4.
4. All the information used in step 3 is used to connect to the Code Server, which could physically be a different machine to the Software Distribution Server. Here, the required connections are made, and the installation is initiated. You may need to log on. For example, if the Code Server is LAN Server based, you will need to log on to LAN Server before the process can continue.

5. Once the installation completes successfully or unsuccessfully, the connections to the Code Server are deleted and the client workstations inform the Software Distribution Server as to the outcome of the installation. The catalog is updated accordingly.

11.16.10 Using the Command Line

Many tasks can be performed from the command line, one of these being the installation of software. Using the command line allows you to build batch files that can be processed in whatever sequence you wish.

For further details on the command line options please refer to the online documentation.

11.17 Remote Connection Services in OS/2 Warp Server

We begin the discussion with OS/2 Warp Server and the new functionality that has been added to Remote Access Services (RAS). The Remote Access Services or *Connection Server* is the component of OS/2 Warp Server that provides the remote workstation access to the LAN and consists of the client side and the server side.

The Remote Access Services allows multiple, concurrent, remote OS/2 and Microsoft Windows workstations to connect into a LAN. When connected, the remote workstation has the same abilities and functions as if it were directly connected to the LAN and can directly access any device on the LAN.

11.17.1 Remote Access Protocol Options

OS/2 Warp Server provides you with a wide range of supported networking protocols and communication adapters that you may use in many combinations to suit your requirements for a server system. Adapter and Protocol Services may be called the *communications engine* of OS/2 Warp Server since they provide communication support for any of the other components of this product.

Adapter and Protocol Services is a very complete set of networking protocols which can be used in a LAN environment as well as for wide area networking. Remote Connection Services supports LAN applications that work with NDIS directly or use a LAN protocol that is NDIS-compliant. NDIS is a network driver architecture that allows a workstation to support multiple network adapters and protocols. The following networking protocols are based on the Network Driver Interface Specification (NDIS) standard:

- NetBIOS
- TCP/IP

- IEEE 802.2
- IPX/SPX
- NetBIOS over TCP/IP
- NetBIOS over IPX support

11.17.2 Implementing Security

Adding remote access capabilities to your LAN can make your LAN and its resources vulnerable to unauthorized remote access. The security features provided by the Remote Access Services product control access to the Connection Server and help prevent LAN access by unauthorized users.

The Remote Access Services security subsystem provides two main services:

1. Protects the LAN from casual, unauthorized, external access.

   When an external WAN circuit is established at a Connection Server, the security service ensures that, until the caller is authenticated:

   - No LAN frames are transferred onto the WAN circuit.
   - No WAN frames are transferred onto the LAN wire.

2. Continuous validation of remote requests

   When a Connection Server receives a request for service, it can determine whether the:

   - The request was sent by an authorized user.
   - The request received has not been modified in transmission.
   - The current message is not a copy of a prior message.

   Before a remote workstation sends requests to a secured Connection Server, the user at the remote workstation must first be authenticated by the Connection Server.

11.17.3 Security Features

The Remote Access Services security feature is a configuration option that can be enabled on a remote workstation workstation as well on the Connection Server. This function is not available on Windows workstations (however if it is enabled on the Connection Server, then both the OS/2 and the Windows requester must supply a user ID and password).

If security is disabled, any person can access the configuration interface at the Connection Server and enable its security option. However, once security is enabled, only a user designated as a security administrator can log on to the secured workstation and disable the security subsystem. Understand that the user database used for the Remote Access Services

does not interface to any other user database (such as User Profile Management used by the File and Print Services).

Enabling or disabling security at a remote workstation is a local operation only and cannot be performed remotely. That is, a security administrator must be physically located at the machine when operating the configuration user interface that toggles the state of the security subsystem.

Password Phrases: To minimize the possibility of offline dictionary attacks to discover user passwords, the security database supports *passphrases*. Up to 32-case sensitive characters can be used to build individual tokens that comprise a password phrase. The passphrase is one-way encrypted using a *hash* algorithm. The resulting *password key* is eight bytes in length.

User Permission Types: The user accounts database on each remote workstation is maintained independently. The Connection Servers user database can be configured to operate independently or to use a shared database. This database contains information on each user such as the user ID, password key, and user type.

The three user types are:

- User
- Administrator
- Security administrator

Single Logon: A user is required to log on and be authenticated by each Connection Server before accessing the server's services. For example, a user that has been authenticated can:

- Use Dialer services
- Use Management services
- Access the target LAN wire

However, a user need only be involved in a single logon task (that is supplying a user ID and passphrase) provided the user has the same user ID and passphrase at each of the secured Connection Server workstations that the user subsequently attempts to access. The user ID and password key used during the first logon are saved (in memory only) by the workstation security component and used first for each of the following logon attempts at the other secured Connection Servers. The user is required to participate in a second logon only if the user ID or passphrase is different at the next secured Connection Server .

If a Connection Server has security enabled, then remote workstation users (both OS/2 and Windows) are prompted for the user ID and passphrase after they dial and establish a link with the Connection Server. If an OS/2 remote workstation has security enabled, the user at that workstation must also log on *locally* before accessing local services (such as Settings). In addition, users at remote workstation (both OS/2 and Windows) attempting to access an OS/2 remote workstation where security has been enabled must first log on to that remote workstation, just as they would to a secured connection server. This additional function is not available for Windows remote workstation users.

If security is enabled at an OS/2 remote workstation and if the user ID and the passphrase match between the OS/2 remote workstation and the Connection Server, the user is prompted for only one logon (the first local logon); an implicit logon occurs after a connection is established. If the user ID and the passphrase do not match between the remote workstation and the Connection Server, the user is prompted to log on again to the connection server after the link has been established.

After the remote logon and filtering has completed, it is the responsibility of the LAN-based applications, such as OS/2 LAN Server, to provide security for their own applications. Logons to these applications are separate from the remote logon.

11.17.3.1 User Authentication Protocol

The Remote Access Services security subsystem implements a two-party, two-way entity authentication protocol based upon an IBM patented protocol called 2PP. The Remote Access Services user authentication protocol is based on the use of *Message Authentication Codes* (MACs).

A Message Authentication Code is an 8-byte *cryptographic* checksum attached to the message. It is derived using a secret key and the content of the message. The Message Authentication Code scheme uses Data Encryption Standard (DES) and adheres to the X9.9 standard.

After a successful mutual authentication (client to server and server to client), the client and server both share a session key that is used to build the certificates that authenticate all subsequent workstation service requests sent to the Connection Server. A different session key is used during each separate logon session. The protocol satisfies the following requirements:

- The protocol provides mutual authentication between a client and a server. In the process of authenticating one another, the client and server come to share a random session key.

- The client initiates the protocol. The client has no information about the server except the server's address. The client has a user ID and a user supplied passphrase for authentication.
- The server has no information about the client besides the client's user ID and passphrase-derived key.

User Authentication Protocol's Data Flow: The protocol requires three rounds and is shown in Figure 144.

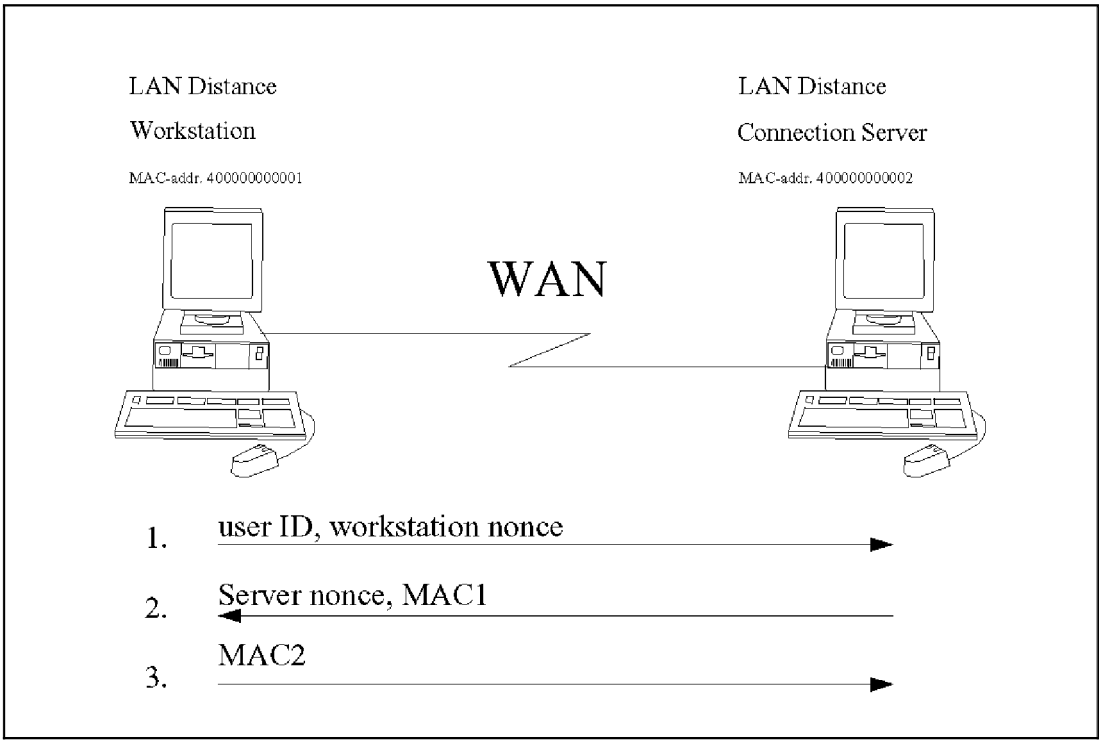

Figure 144. LAN Distance Protocol Data Flow. Although the user types in a user ID and passphrase, note that the passphrase does not go across the link.

Figure 144 shows you the protocol data flow.

1. In this protocol, a user on the remote workstation logs on to the Connection Server. The user submits only the user ID and the passphrase. The remote workstation actually sends the following to the Connection Server:
   - User ID
   - Remote workstation *nonce*. A nonce is a random value generated for this session only and will not be repeated in subsequent sessions.

   Note: The user passphrase, its associated one-way encrypted password key, and the resulting common session key do *not* appear on the

link. A new one-way encrypted password key derived from a new passphrase (that is, when the user changes the passphrase) does appear on the link, but it is encrypted using the logon session key.

2. The Connection Server responds by creating its own server nonce and returns both the Connection Server nonce and a Message Authentication Code (MAC1) based on the following information and encrypted using the one-way encrypted password key from the user's database account:

   - User ID
   - Remote workstation nonce
   - Connection Server nonce
   - Connection Server LAN adapter address (400000000002)

   Since the remote workstation also knows the above information, the remote workstation can generate the same MAC1 using the one-way encrypted password key derived from the passphrase supplied by the user.

   The remote workstation compares the Message Authentication Code (MAC1) received from the Connection Server with its locally generated Message Authentication Code. If they match, the remote workstation accepts the connection server as authentic.

3. The remote workstation then returns a new message authentication code (MAC2) back to the Connection Server based on:

   - Workstation nonce
   - Connection server nonce

   When the Connection Server receives MAC2, it computes its own Message Authentication Code based on the same information and compares its code with the one received from the remote workstation.

   If the two codes match, the Connection Server accepts the remote workstation as authentic.

As a result of the exchange, after each side has authenticated the other side, each party separately generates a common *session key*. It is good for that session only.

Note: The session key never goes across the link.

This session key is then used to verify all Remote Access Services commands. The session key is not used to verify data for other applications going across the link.

If the remote workstation user is authenticated successfully (that is the user provides the correct passphrase), the remote workstation can generate *server certificates* (that is, Message Authentication Codes) that can be added to requests sent to the Connection Server.

When a Connection Server receives a request containing a certificate, it can validate the certificate and verify that the user sending the request is authentic and authorized by the Connection Server to request the service.

Moreover, a valid certificate contains proof that the request itself has not been modified since being sent and is not a copy of a certified request (sent earlier by a valid user) that was introduced by a hacker masquerading as the valid user.

11.17.3.2 Security Policy Options

Several user-authentication *security policy* options can be configured by a security administrator when setting up a Connection Server, such as the following:

- Maximum Age
- Minimum Age
- Minimum Length
- Duplicates Checked
- Maximum Logon

Note: Security policy is a set of rules that can be customized to enable the security requirements of a particular user environment.

11.17.3.3 Additional Security Options

This section covers additional security options that are available. These are:

- Callback
- Workstation Address
- Logon Time intervals

Callback: The Remote Access Services security supports an optional *Callback* feature for remote workstations only. Callback to Remote Access Services requesters and to Connection Servers are not supported. The Callback option configured within the caller's account is not checked unless the call is placed from a remote workstation.

Note: Callback is a feature, active during LAN Distance connection establishment, in which the answering workstation reinitiates the connection by placing a callback to the dialing workstation. The original dialing workstation must be a remote workstation.

Figure 145 on page 297 shows you the general Callback procedure.

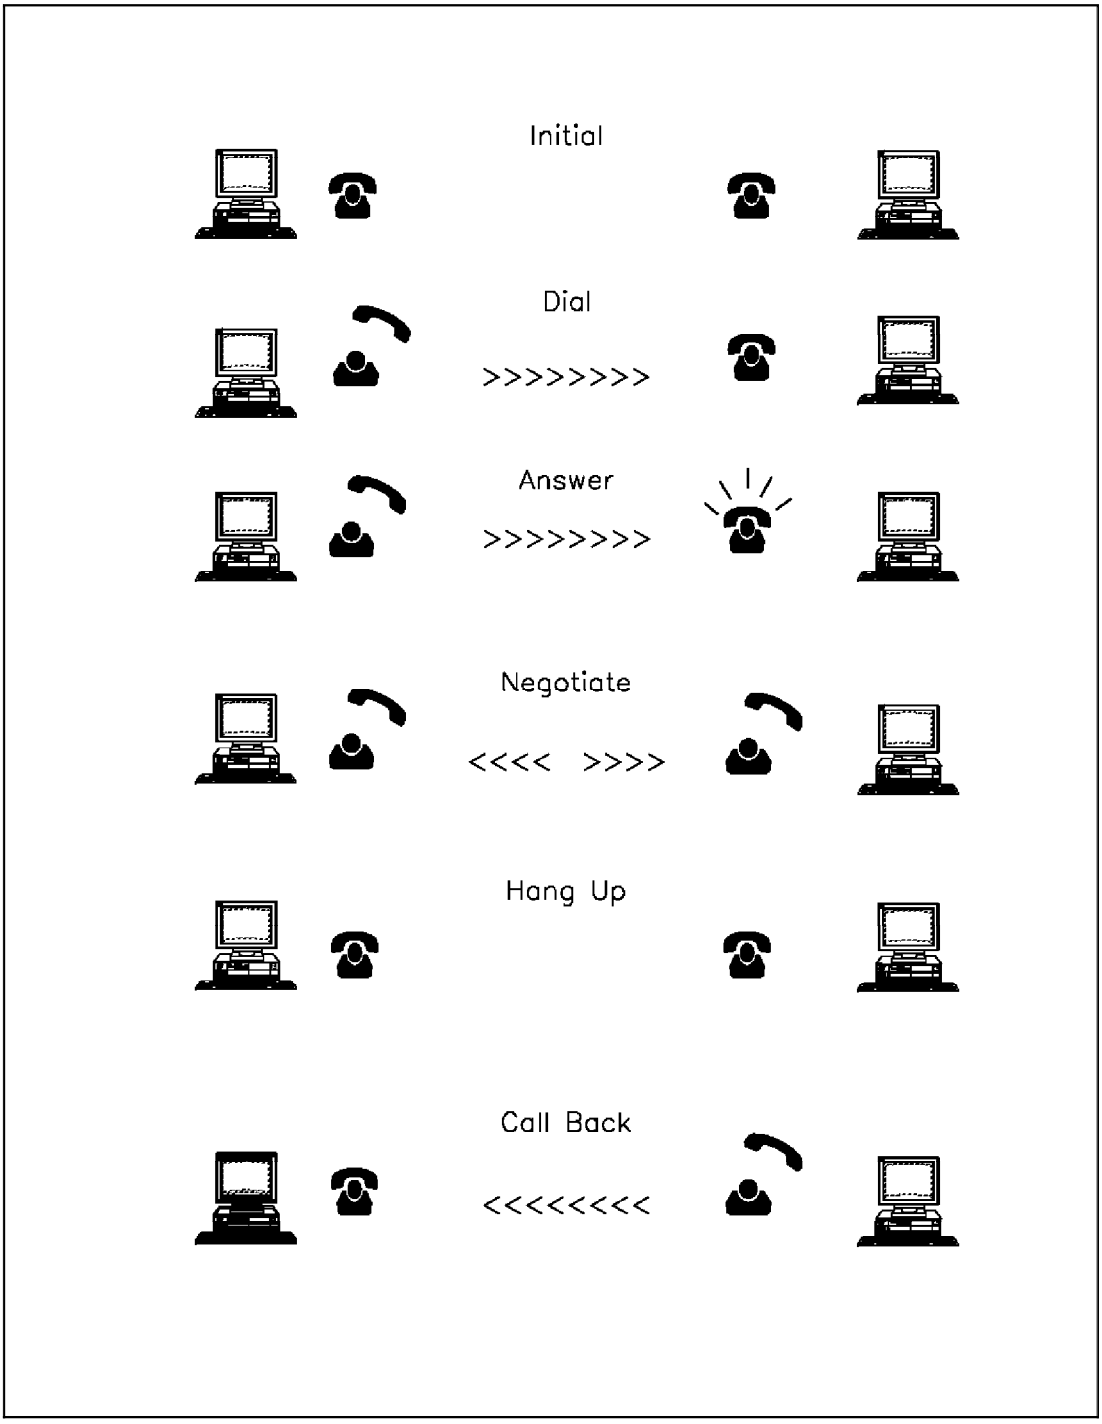

Figure 145. Callback

Callback can be configured in a user account as follows:

- Callback not required

  These users are never called back by the called Connection Server.

- Fixed Callback

  These users are called back at a fixed configured telephone number.

- Mobile Callback

  This is part of the logon protocol. The Connection Server can then use the telephone number submitted to it for the callback.

The caller is authenticated both:

- Prior to the Callback (this prevents harassment calls)
- And also after the Callback is complete (this guards against known hacker techniques that can normally only be avoided using special telephone equipment or service options)

Callback can be useful if reversal of telephone charges is needed. For example, the majority of the charges for a call from a hotel room can be charged to the central site instead of to the traveler at the hotel.

Workstation Address Identification: A security administrator can configure up to eight workstation LAN addresses within a user account. The caller must call from a workstation that has been configured with a Remote Access Services logical adapter network address that matches one of the MAC addresses stored in the caller's account; otherwise, the logon attempt fails.

Valid Logon Time Intervals: A security administrator can configure the days of the week and the time of the day during these weekdays that a user is allowed to log on to his account at the Connection Server.

A logon attempt at a time that is not within the specified time intervals specified in the user's account, fails.

11.17.3.4 Protecting Your Passphrase

The following diagram shows you what to consider if you work with Remote Access Services and need to log on to different systems in the LAN (for example to the OS/2 LAN Server and to a 3270 host).

Important

It is assumed that Connection Servers are physically secured. Access should be limited to a few trusted administrators.

In Figure 146 on page 299, you can see that a user on a remote workstation that would work with the OS/2 LAN Server and a S/370* host has to make a logon (with a user ID and a password) to:

1. Remote Access Services
2. OS/2 LAN Server
3. S/370 host

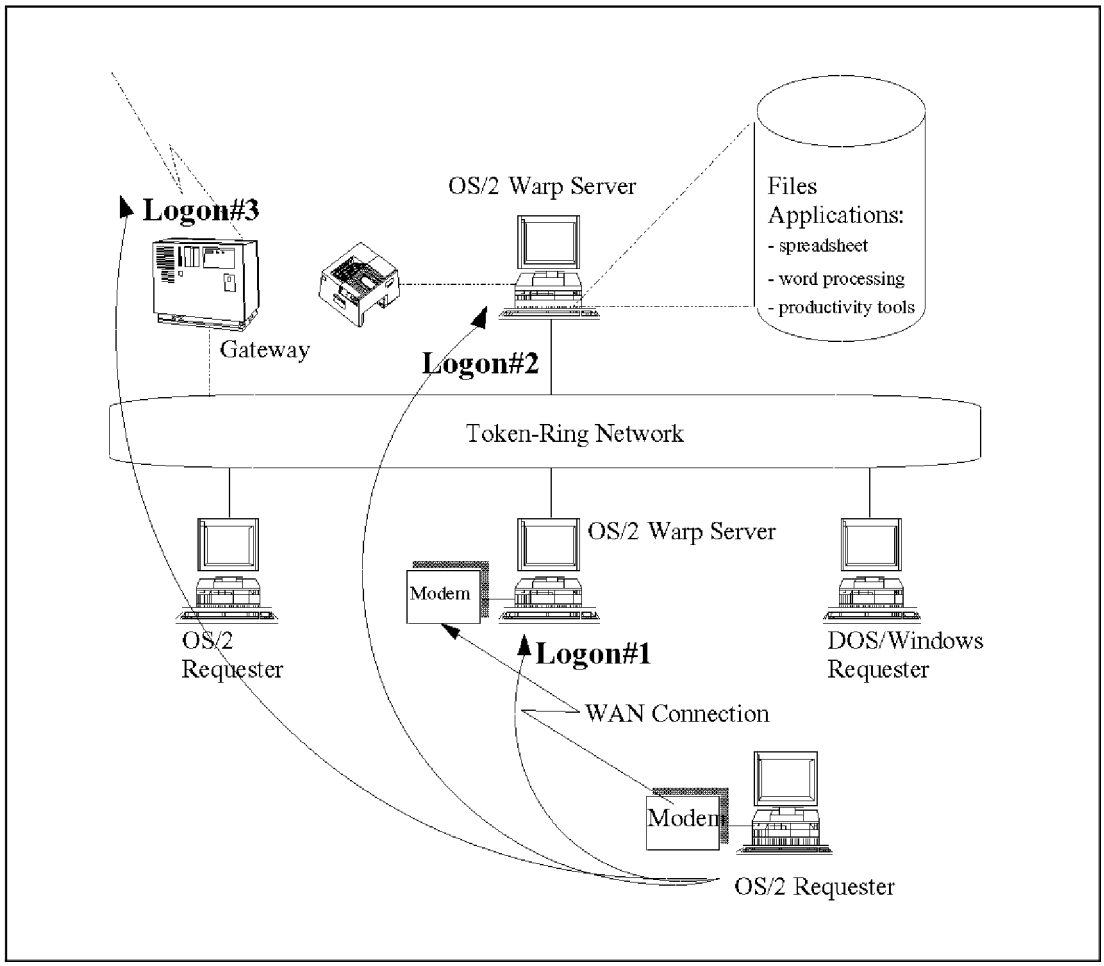

Figure 146. Protecting Your Passphrase

Note: Make certain the passphrase used to log on to the Connection Server and the passwords used for the OS/2 LAN Server and 3270 host sessions are different. This is important because *Remote Access Services does not encrypt application data that appears on its link.* Special equipment (such as a LAN sniffer) cannot see the Remote Access Services passphrase, but can see other passwords from other applications that do not encrypt their passwords. So, use different passwords for LAN Distance and for other applications, and inform the remote workstation users about this rule.

If you have used the same password for Remote Access Services and for your 3270 logon, a hacker may trace the communications link, identify the 3270 password, and try this password for the Remote Access Services passphrase. This may enable a hacker to dial into your LAN.

11.17.3.5 Remote Access Services Security Options.

The Remote Access Services supports two optional types of security for restricted access to the LAN and its resources. The first type of security is included with the Remote Access Services component and is provided by the User Account Management. The second type of security that Remote Access Services supports is a user-provided *User Exit package*. Remote Access Services supports any OEM-provided security User Exit package that is developed in conformance with the LAN Distance Generalized Security User Exit API.

You can also set up the Remote Access Services to use either or both of the security options defined above. By default, security is disabled. The Remote Access Services Notebook is used to enable security. Before security can be used, it must be enabled.

11.17.3.6 The Security User-Exit Package

The security User-Exit package consists of two User-Exit modules: one for the client and one for the server. The client and server User-Exit modules work together to implement the user authentication protocol defined by the security User-Exit package.

A user authentication protocol is a series of User-Exit messages/tokens exchanged between the client and server User-Exit modules when validating the user of a remote workstation that is calling a Connection Server.

One client workstation can use a different security User-Exit package to access each different Connection Server it calls. A Connection Server can use only one security User-Exit package to allow access from all remote workstations that call it.

Security User-Exit packages can be used with or without Remote Access Services security (User Account Management). If LAN Distance security is used with the security User Exit package, the authentication will take place first through the user-exit and second through LAN Distance security.

11.17.3.7 Shared User Database

A large environment that has the requirement for multiple Remote Access Services servers poses some problems, one of them being that of registering users on each server. With previous versions of LAN Distance Connection Server, the user database had to be duplicated. Previous

versions of LAN Distance also allowed only one security administrator to log on at any one time. One of the new features of Remote Access Services is the ability to share the user database between servers and allow multiple administrators to log on simultaneously.

The database sharing is achieved by using a shared file on a redirected drive provided by a file server. In this way multiple Remote Access Services servers can share the Security Database file. The integrity of the Security Database is protected by serializing all modify request to make the Security Database.

Because the Remote Access Services servers rely on the file server to access the database, this file service should be up and running at all times. One of the problems that one may encounter while running in SHARE mode is that of database backup. The Remote Access Services does not provide a database backup mechanism; it relies on network software for the backup.

11.17.4 Mobile File Sync

Mobile File Sync (MFS) is a file system that supports mobile OS/2 Warp 4 clients. Mobile File Sync allows users to physically disconnect from the OS/2 Warp Server or Windows NT Server and still have access to their server files. Warp clients can be either LAN Requester clients or OS/2 Peer clients. MFS is available as part of the OS/2 Warp AttachPak.

Mobile File Sync caches the accessed files and directories to the client machine. When the client is disconnected from the server, the user can continue accessing the files and directories previously cached from the server. The user can read the cached files, update them, or create new files. The user can also list contents of cached directories, create new directories, or delete existing ones. Mobile File Sync keeps track of all updates by recording them in a "Client Modification Log." All updates are propagated to the server when a new connection is established in a process called "Reintegration."

MFS is can be used on its own or together with Remote Access Services for LAN Distance.

11.17.4.1 MFS Functions

Mobile File Sync provides the user with three levels of functions:

- Basic
- Intermediate
- Advanced

Once you have installed MFS, you need to configure the the level of function that you are going to use. No further configuration is required for basic functionality. Further details on using MFS is available on the online documentation.

11.17.5 Inactivity Timeout Feature

This feature is available on both Connection Servers and remote OS/2 workstations. When enabled, this feature subjects every machine that connects to it to a usage test every minute. The usage test checks how many LAN frames have passed across a link and, depending on the values set, decides whether the link should remain up or disconnected.

Figure 147 shows the Shuttle Option window.

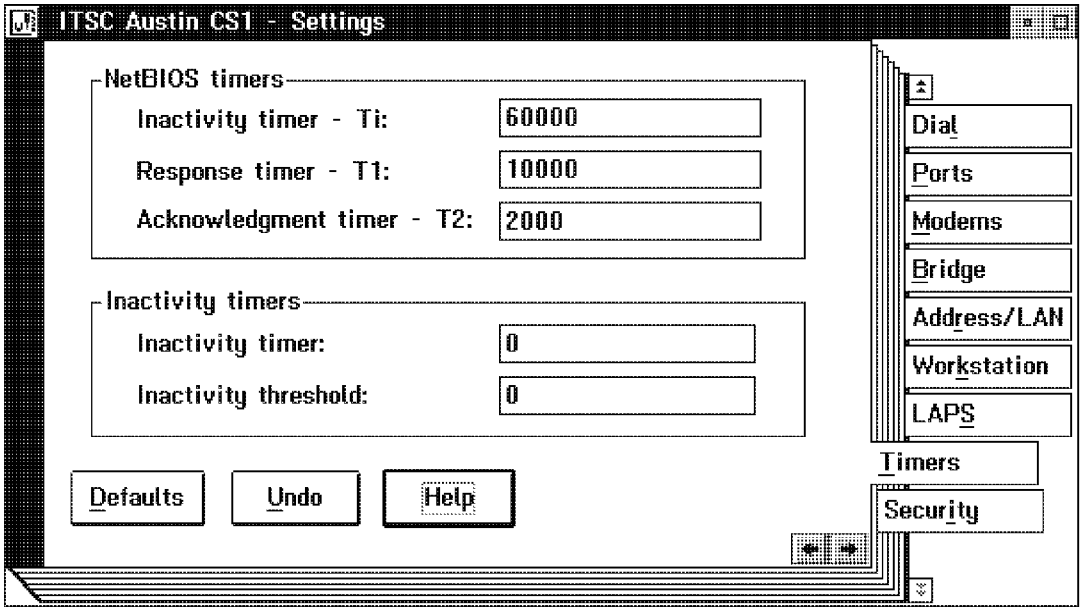

Figure 147. Inactivity Timeout Option

Figure 147 shows the Timers Notebook page where the Inactivity timeout is set. Two values need to be set:

1. **Inactivity timer**

   This timer specifies the maximum amount of time in minutes to allow idle activity before the connection is disconnected. The default value of 0 disables the inactivity timeout function. The range of the value is between 0 and 999.

2. **Inactivity threshold**

   This timer specifies the minimum number of frames required to cross the WAN link on an active connection within one minute. If the number

of frames per minute falls below this timer value for the length of time specified in the *Inactivity timer* field, the connection will be disconnected.

The default value of 0 disables the inactivity timeout function. The range of the value is between 0 and 99999.

11.17.6 PIF Files for Uncertified Modems

A Product Information File (PIF) is used to initialize a modem. The PIF file contains all needed string information and configuration values for your modem.

To set up a modem, initialization strings are needed. A modem initialization string is an `AT` command string passed to the modem when Remote Access Services server or requester is first started. The initialization string is used to configure and optimize the modem for use with Remote Access Services. The PIF file has two parameters that are used to initialize the modem, `Initialization1` and `Initialization2`.

If you have a modem that is not supported, and you cannot get it to work using another supported modem type, it is usually because the initialization string is incompatible. There are a number of parameters that may need to be modified in a new modem PIF file. To help you, the `CFMODEM` utility is shipped with OS/2 Warp Server

The `CFMODEM` utility is a small application to modify and create PIF files. This graphical utility should help you to create the needed PIF files for unlisted modems.

To create and modify modem strings and Remote Access Services PIF files, you need to have some technical knowledge on modems. Also you need to refer to your modem manual to find the correct commands.

11.18 Backup and Recovery Services

OS/2 Warp Server Backup/Restore is available on OS/2 Warp Server systems and integrated with other system services. It provides backup and restore facilities for its system and operational data. These facilities can be extended to other servers that OS/2 Warp Server can access through either LAN Alias or Logical Drive Definitions. As you will see a little later, it provides quick-start guides, hints, default backup and restore routines, and a user-friendly graphical user interface. The GUI is both useful and attractive and provides for quick and easy setup for backup/restore and disaster recovery procedures. It makes backup administratively easy.

OS/2 Warp Server Backup/Restore is based on the *Personally Safe and Sound* (PSnS) product and gives you the possibility to save your server system against lost of data. Warp Server Backup/Restore actually does not back up locked files at the moment. This feature is being reviewed by development and will be made available with the first Warp Server ServicePak. Once this ServicePak is applied to the system, Backup/Restore will be able to handle open files or locked up files in an effective way. Availability of this ServicePak is scheduled for first quarter in 1997.

In order for you to be protected against all kinds of data loss, operational errors, hardware damages and any other kind of disaster, a backup strategy has to be set up. You can choose to install all services at one time or add or delete services at a later date as the need of your environment changes.

Warp Server Backup/Restore supports the OS/2 file system features for:

- Extended Attributes
- Network Access Permissions
- Long File Names

Warp Server Backup/Restore provides support for:

- Locally attached hard disks
- Remotely attached hard disks, for example, accessed with LAN drivers
- FAT formatted drives
- HPFS formatted drives
- HPFS386 ACLs

Warp Server Backup/Restore allows you to determine the importance of your data and group it into different backup/restore priorities. After you have double-clicked on the Warp Server Backup/Recovery folder object, you will see Warp Server Backup/Restore folder contents as shown in Figure 148 on page 305.

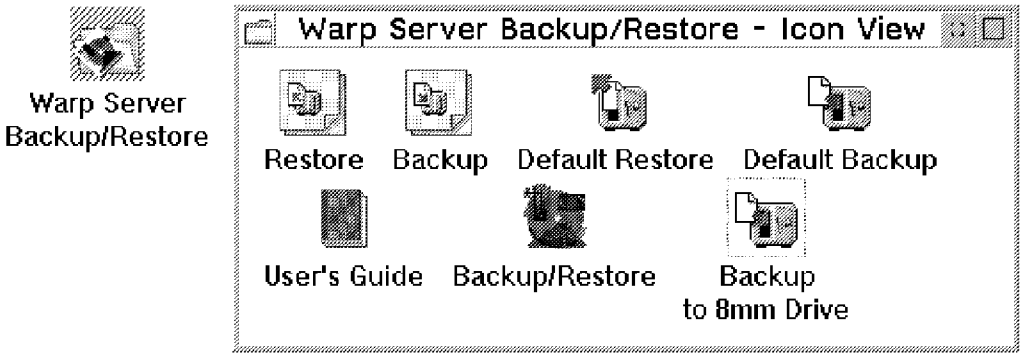

Figure 148. Warp Server Backup/Restore Folder

Double-click on the Backup/Restore object and you are taken to the main menu. By opening the **Tools** pull-down menu you can define all your needs as shown in Figure 149.

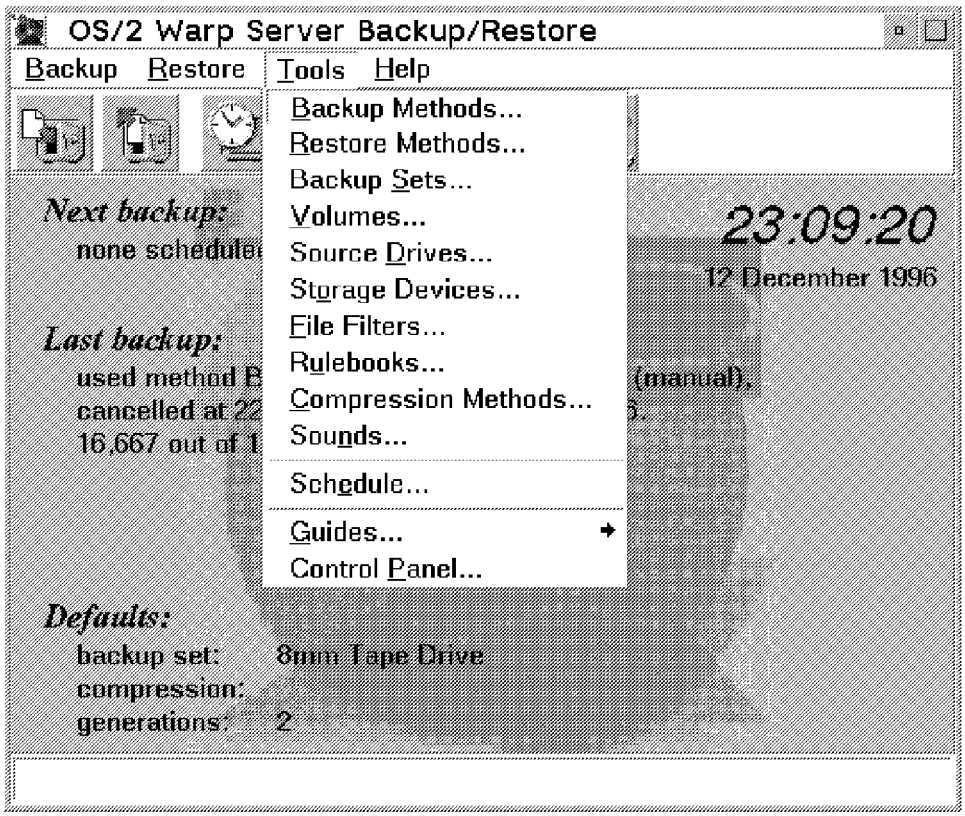

Figure 149. Warp Server Backup/Restore Tools Pull-Down Menu

As you can see, with Backup/Restore you have various solutions to choose from. It could be necessary to require multiple backup strategies to address unique data requirements. OS/2 Warp Server Backup/Restore implements

backup strategies through the definition of a *Backup Set*, a *Backup Method*, and *Index Files*:

11.18.1 Backup Set

A *Backup Set*, as shown in Figure 150, is a logical collection of backed up files. These files are loaded onto a specific storage devices that resides among a group of supported storage devices. When you define a new Backup Set, you must specify a *unique Backup Set name*, *the storage device* on which to store the data, and the *properties* to be used in the backup for the storage devices. Warp Server Backup/Restore manages these backups as a group.

Ask the Backup/Restore feature for the file you need; the program will know exactly where it was backed-up and how to retrieve it. It is a good operative choice to keep a separate Backup Set for each project or activity on your machine. This gives you the possibility to transfer out their Backup Sets when they have finished and keep the backup as a project archive.

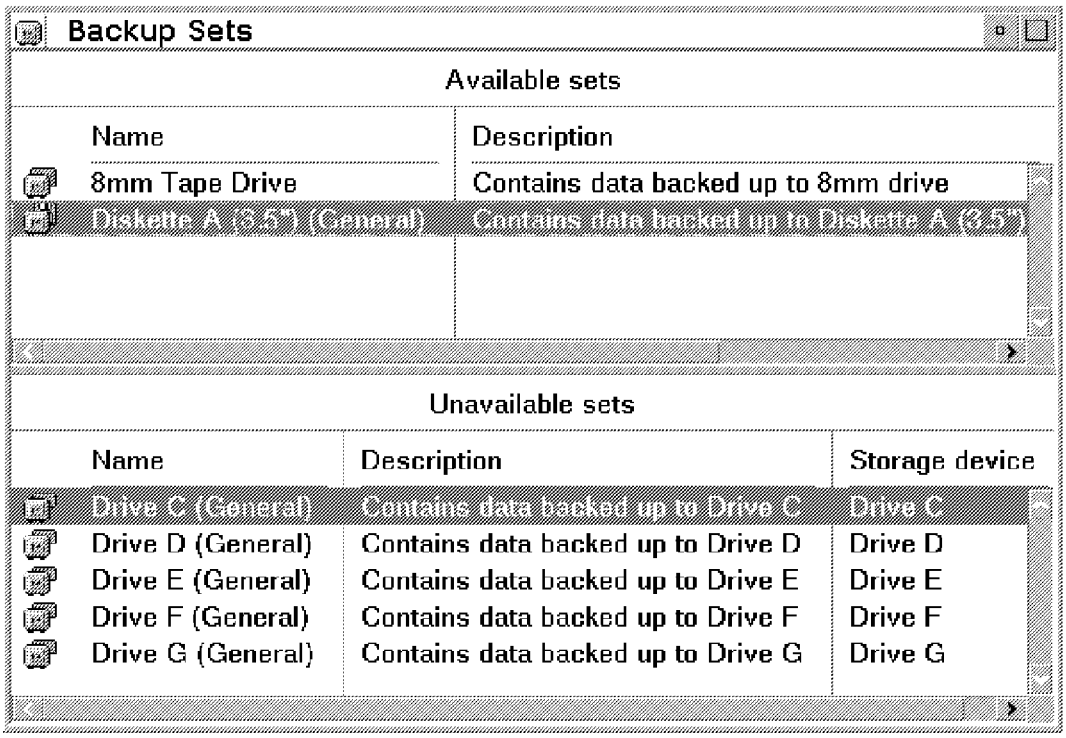

Figure 150. Warp Server Backup/Restore Backup Sets Window

The GUI interface for Backup/Restore is designed so that you can see the logic flow through connecting pipelines. If you follow the pipelines from the top-left of the window down to the bottom-left, you can see that Backup/Restore enables you to create a flexible and easy solution for changing the settings, seeing the flow of your backup, verifying your

backed-up data, verifying that the data actually written to your storage device is correct, and choosing compressing the files.

11.18.2 Backup Method

Every backup is controlled by a backup method that contains all the information needed to execute your backup. It will tell you what is backed up by selecting a source directory and using file filters to exclude unwanted files by specifying whether you want to do an incremental backup, what method of compression you wish to use, and how many generations you wish to keep. You can choose where the data is backed up by specifying a destination Backup Set. It is also possible to see a preview of the data that is about to be backed up. You can choose which files and folders to back up and where to store them. You can restore data simply by selecting the files or folders you want restored. Warp Server Backup/Restore keeps an inventory of all the files that were backed up, where they were stored, and if they had attributes.

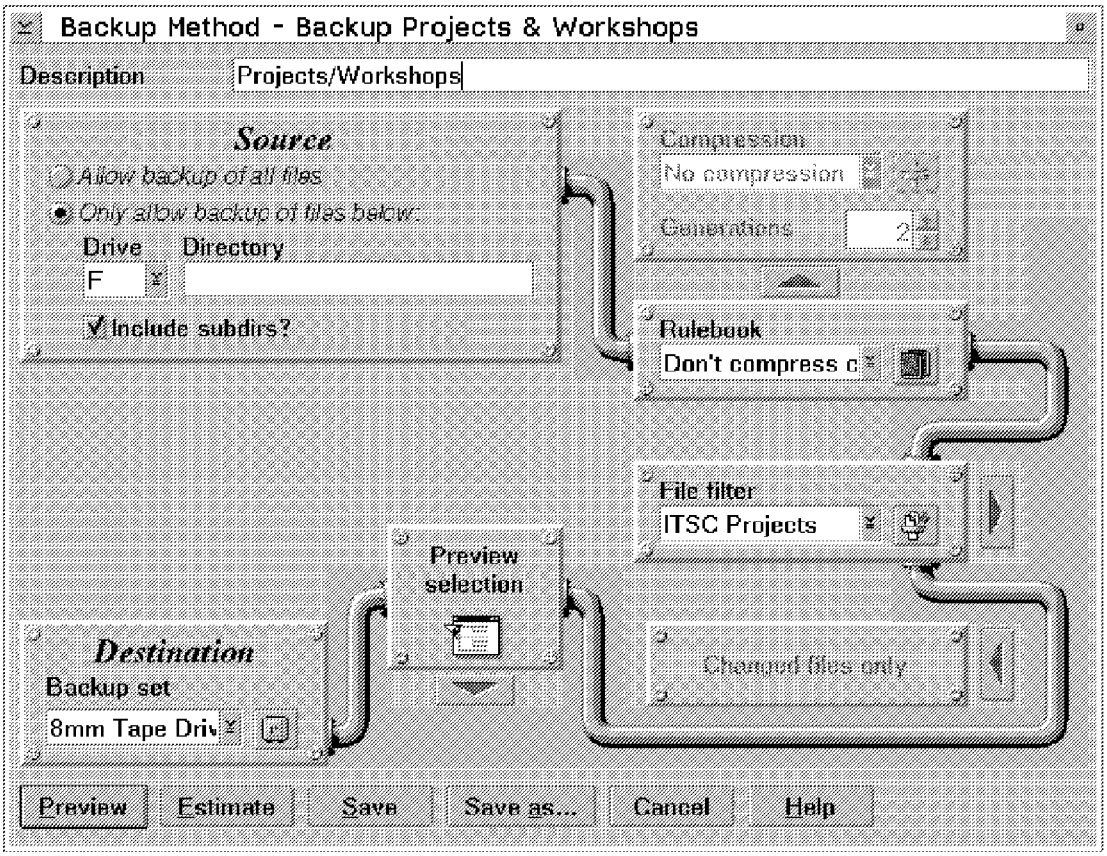

Figure 151. Warp Server Backup/Restore Backup Method

Following the connecting pipelines flow as shown in Figure 151, you can decide to:

- **Allow backup of all files:**
  - This first option will include all the objects to be found on all the source drives that are currently checked for the backup.
- **Only allows backup of files selected below:**
  - This permits you to specify a drive or a subdirectory to backup from.
- **Drive**
  - The letter of the source drive to backup from. If you want to backup all source drives, enter an asterisk *.
- **Subdirectories**
  - This option is to include all objects in subdirectories of the directory specified above.
- **Compression/Pipe Rulebook**
  - The first choice that you have to make is to follow the pipeline flow using the red arrow button passing through the option you need, and then select the specific items. This specifies how files and folders are backed-up. You can specify a *Compression Method* and number of generations for all the files and folders that are backed up or you can select the *Rulebook*, which allows you to specify the Compression Method and generations individually for each file and folder.
- **File Filters**
  - Following the pipeline flow, you can pass through the *file filters*, as shown in Figure 152 on page 309. This allows you to define which files and folders to back up. These can be defined as either "Tree-based", which gives a graphical representation of the drives files and folders on your machine, or "Rule-based", in which case you define some *rules* in the file filter as to whether Insert, Delete, Exclude, or Include files and folders.

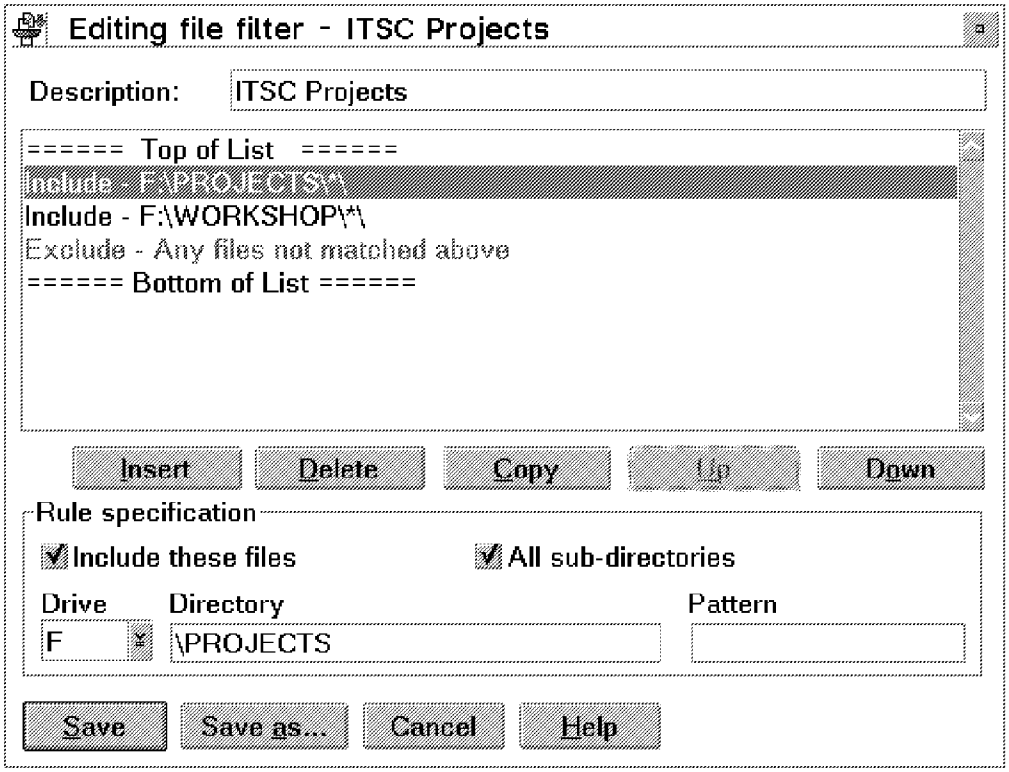

Figure 152. Warp Server Backup/Restore ITSC Projects File Filter

- **Changed files only**
  - In this case the Backup Method will perform an incremental backup.
- **Preview selection**

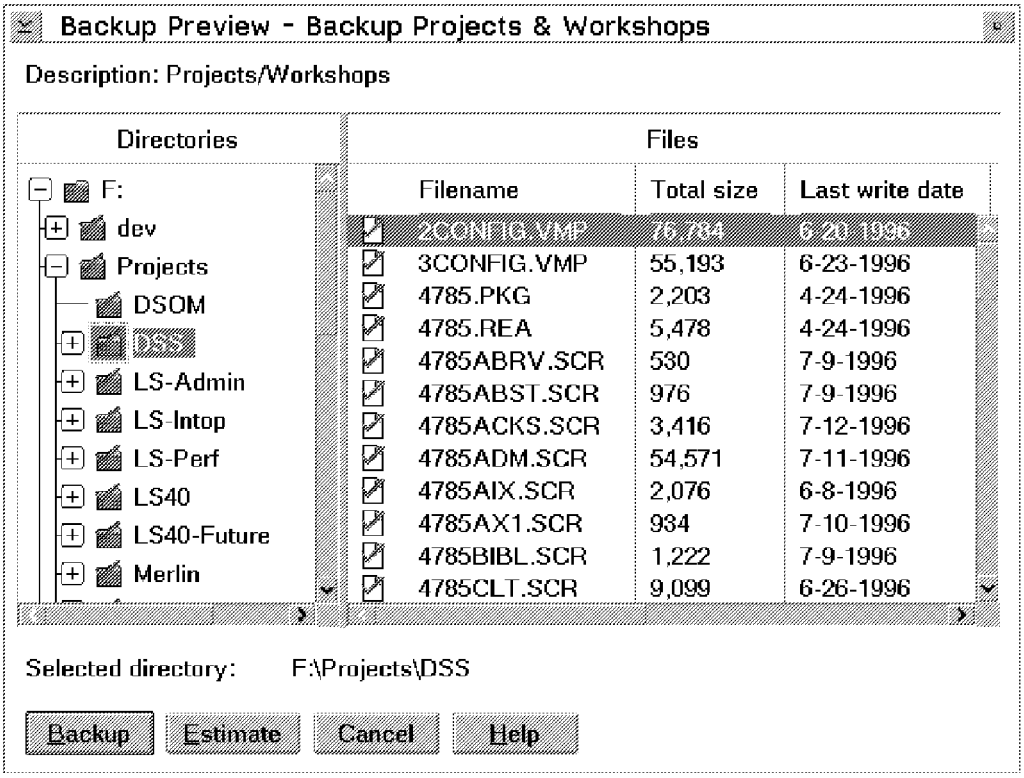

Figure 153. Warp Server Backup/Restore Preview Window

- Before starting your backup, you can preview the objects that will be backed up. This allows you always to check that the right files and folders have been backed up. It is also possible to deselect the files that you do not want to include. Select **Preview Selection** and then select **Preview**. As shown in Figure 153, a tree will appear with a list of files and folders included in the backup. Select the **Backup** button and your backup will start. If you do not use the preview selection, the backup button will be on the Backup Method main menu.

11.18.3 Volumes

As shown in Figure 154 on page 311, Warp Server contains facilities for management of the volumes that contain backed up data. A volume is an item of backup media, for example a hard disk, tape, or optical disk. Removable volumes, such as diskette and/or tape, can only contain data for a single backup. Fixed volumes, such as hard disk, can contain data for multiple Backup Sets. You can have information about the current status of volumes at any time, add volumes to the system, and disable volumes that have become lost or corrupted. When you first install OS/2 Warp Server Backup/Restore, it creates volumes for each hard disk attached to your machine. Whenever you create new storage devices, new Volumes will be

created for you also. If you double-click on a Volume, additional information will be presented to you:

- **Free space**
  - The amount of free space left for backups on the volume.
- **Backup sets**
  - There will be an entry for each Backup Set that the volume uses. This is the amount of space taken by the data backed up to the Backup Set on the volume.
- **Volume type**
  - The type of volume, for example Hard disk, 8 mm drive, or 3.5 Diskette.
- **Removable**
  - Whether the volume is removable or fixed.

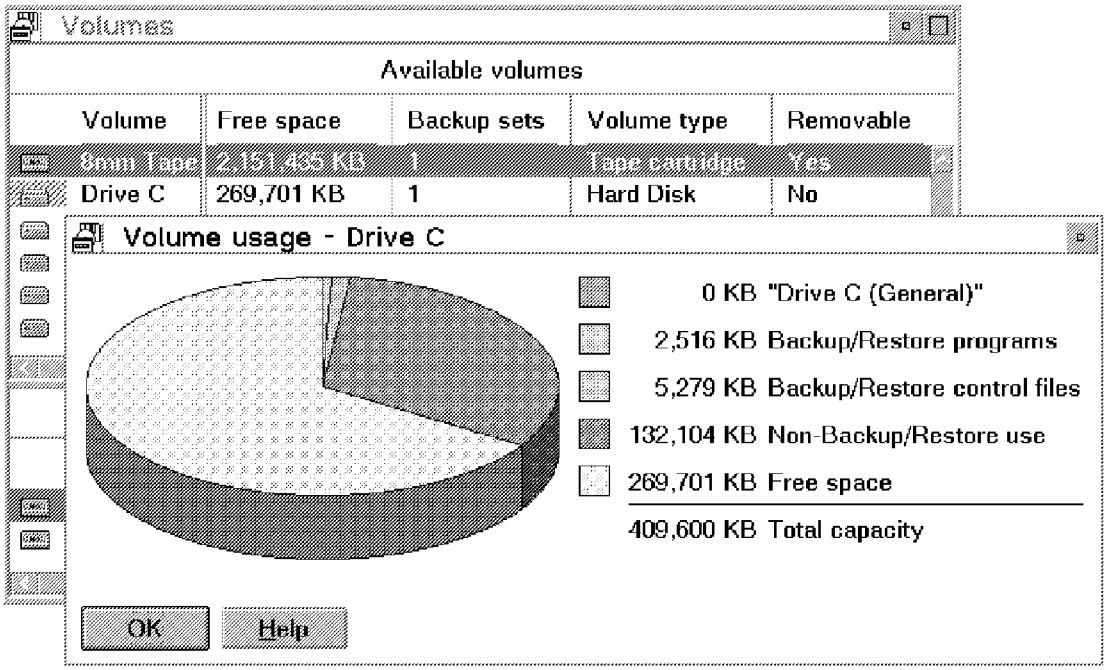

Figure 154. Warp Server Backup/Restore Volumes Window

11.18.4 Source Drives

OS/2 Warp Server detects all drives that have a drive letter, it automatically backs-up data from all the local hard disks attached to your machine. It does not automatically back up data from any other drives that may be attached such as LAN Drives, diskettes or CD-ROM

drives. As shown in Figure 155 on page 312, when you open the Source Drives window, you will see the source drives checked for backup. In the second half of the dialog box there are the ignored source drives. It is possible to drag the drive you are interested into the upper part of the source drive dialog box. You can refresh the list of source drives after you have logged on to a LAN so that LAN drives are listed as available source drives.

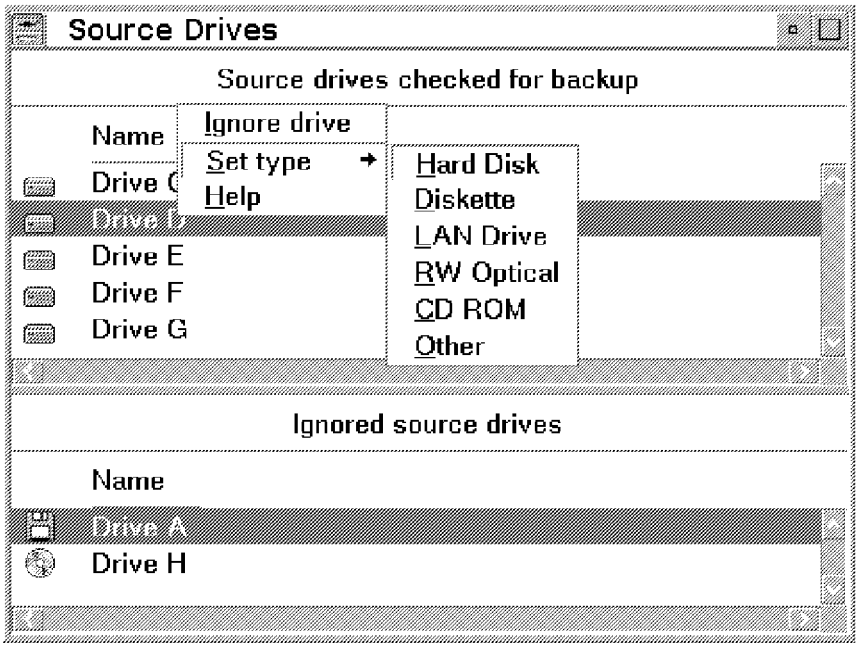

Figure 155. Warp Server Backup/Restore Source Drives Window

11.18.5 Storage Devices

As you can see in Figure 156 on page 313, in OS/2 Warp Server, data is always backed up to Backup Sets. Each Backup Set resides on a storage device. A storage device is a functional unit into which data can be placed, in which it can be retained, and from which it can be retrieved.

Storage devices are divided into two categories:

1. *Removable Volumes - Tape Drives and Diskettes*
2. *Fixed Volumes - Local Hard Disks and LAN Drives*

When you first install Warp Server Backup/Restore, it automatically detects the storage devices that are locally attached to your machine. However, it does detect remote storage devices such as LAN Aliases or ADSTAR Distributed Storage Manager (ADSM) devices. In the Storage

Device dialog box you can see the devices that have been detected and configured on your system, and the name and the class.

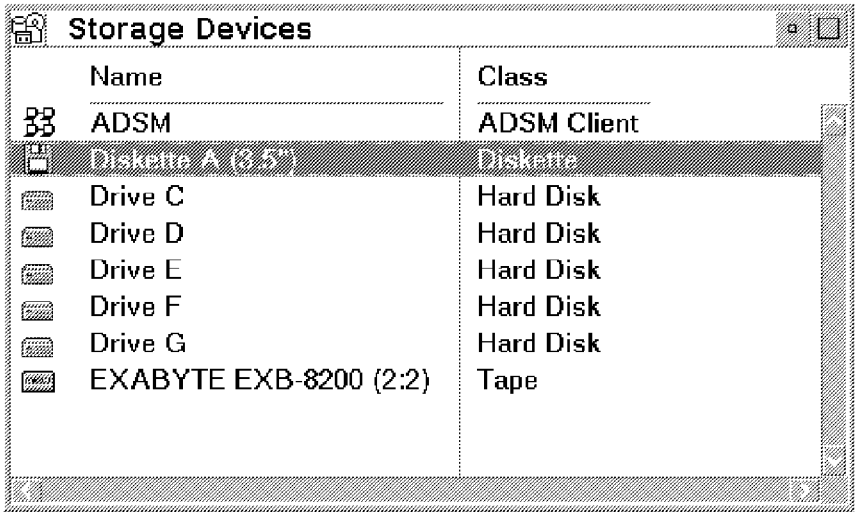

Figure 156. Warp Server Backup/Restore Storage Devices Window

11.18.6 Index Files

The Index Files is a special feature. Backup/Restore remembers where it has backed up data by using the Index Files. The files are automatically backed up by OS/2 Backup/Recovery features, which ensures that your data will always be recoverable.

11.18.7 Backup Invocations

Backups can be done in different ways. You can schedule the event to occur automatically, or make it run manually.

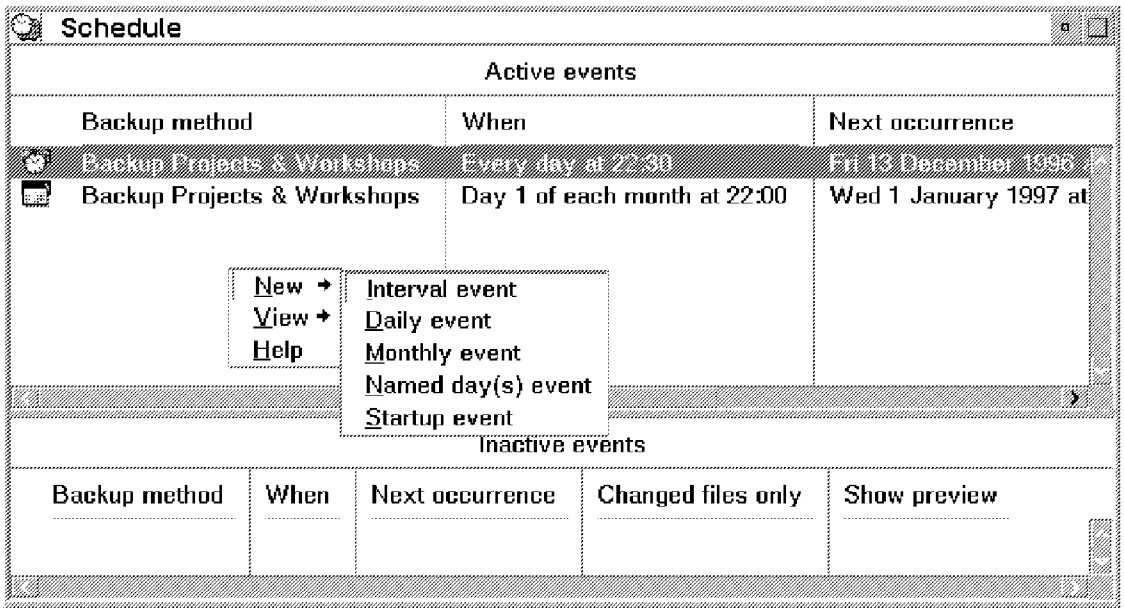

Figure 157. Warp Server Backup/Restore Schedule Window

Schedule events can be defined as in shown in Figure 157:

- **Interval Events**: Interval events occur at fixed intervals of time, such as backing up a transaction every 20 minutes.
- **Daily Events**: Daily events occurs at a specific time every day, such as backing up shared users data at 6 a.m. every morning.
- **Monthly Events**: Monthly events occur at a specific time on a specific day of the month.
- **Startup Events**: Startup events can start a few minutes after Warp Server has finished booting.

Any of these scheduled events can be automatically activated. If you want to delay the start of any of the above-mentioned scheduled events, you must deactivate it. When it is needed again, the selection must be once again activated.

11.18.8 Starting the Backup Process

Whether you select **Backup** from the Backup Method screen or select **Backup** from the **Preview** function of the Backup Method screen to start the backup manually, or the backup is invoked by the scheduler automatically, you will get the Progress Window displayed on your monitor as shown in Figure 158 on page 315.

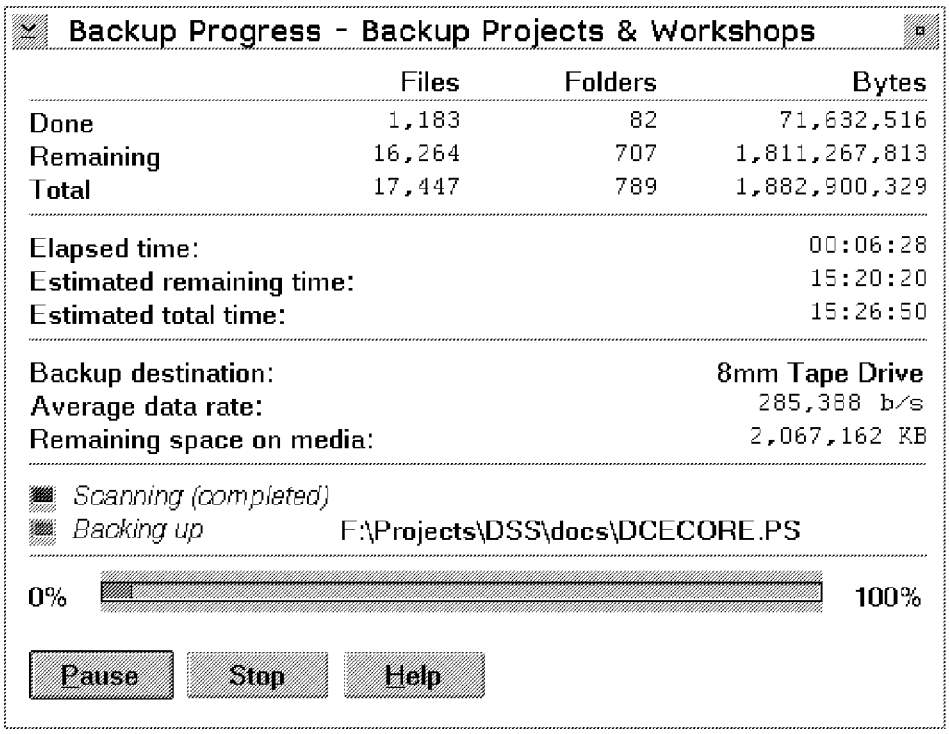

Figure 158. Warp Server Backup/Restore Backup Progress Window

11.18.9 Disaster Recovery Utility

Normally, if a server goes down for any reason and cannot be restarted you would have to first reinstall the operating system then install and second install the backup utility before it is possible to restore your backed-up data.

With Warp Server Backup/Restore, it is possible to restore an Warp Server workstation from locally attached devices, such as tape drive, hard disk, or diskette, without reinstalling your operating system. This feature is achieved by using the Warp Server Backup/Restore Disaster Recovery utility, which creates a set of diskettes from which you can boot your server workstation after the failure occurred. After booting the server workstation from these diskettes, you can restore an entire backup set to recover your system, and you can check, format, and repartition system hard drives and boot sectors.

11.18.9.1 Prepare the Disaster Recovery Utility

To prepare the disaster recovery in advance before your workstation breaks-down, you have to:

1. Create three bootable diskettes. These diskettes will contain:
   - OS/2 Warp systems files

- HPFS386 drivers
- Warp Server Backup/Recovery program files

2. Ensure that the workstation has been backed up to a supported storage device such as:
   - Diskette
   - Backup/Recovery-attached tape device driver support
   - Locally attached read/write optical drive
   - Locally attached hard disk drive

As we have seen up to now, OS/2 Warp Server can easily work in environment where you have a single OS/2 Warp Server server to back up. We even know that OS/2 Warp Server Backup/Restore can be used to back up single machines and small groups of machines connected via a LAN. IBM's ADSM should be considered as an integral part of your backup strategy.

11.18.10 Backup/Restore and ADSM

Warp Server is a business server that addresses many different environments, such as small workgroups, department's LANs, and corporate networks. Each of these environments need different backup/restore requirements, the backup/restore utility provided with OS/2 Warp Server can be extended or integrated to larger and more complex networks by using the powerful backup/recovery features of ADSM, a client/server-based management system. With ADSM, you have one effective storage management plan that encompasses backup archive, Hierarchical Storage Management, and disaster recovery.

When OS/2 Warp Server Backup/Restore is used with ADSM, data is sent to a central server rather than being stored on a locally attached device. Once there, it is centrally managed and protected along with an enterprise's other corporate data, regardless of whether it originated on a personal computer, mid-range computer or mainframe.

ADSM provides backup/restore support for both files and directories. ADSM's cross-user and cross-platform restore provides you with significant flexibility. Cross-user restore enables you to authorize someone else to restore your files. Cross-platform restore is very helpful when you have to migrate to new workstations platforms. For example if a certain day you have to work in a different office, with different workstations, you will still have the access to the data you backed-up. The ADSM server is capable of multitasking; so multiple clients can back up their data concurrently.

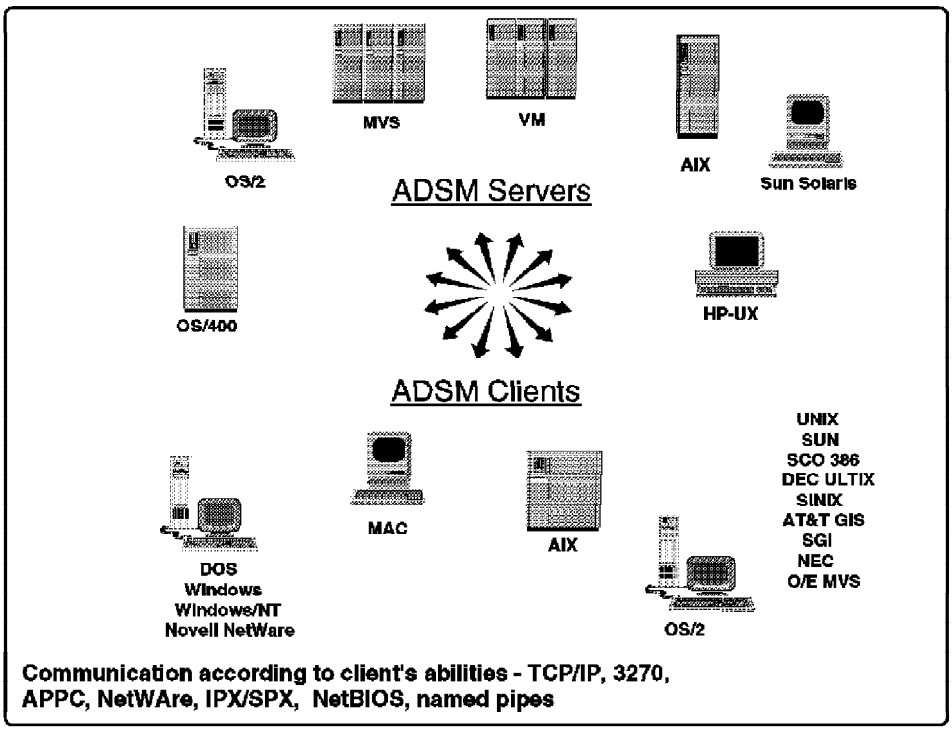

Figure 159. ADSM Platform Support

ADSM supports the direct backup of LAN Server data files and system files that are open and locked during server operation. That is not very common in other backup/restore programs. While installing and configuring OS/2 Warp Server Backup/Restore, you can define the interaction you require with other ADSM systems. It is possible to have:

- OS/2 Warp Server Backup/Restore environment without access to any ADSM systems:
  - This is an environment that uses OS/2 Warp Server Backup/Restore without ADSM. OS/2 Warp Server Backup/Restore is responsible for the backup and recovery of the Warp Server system and data.
- OS/2 Warp Server Backup/Restore coexisting with your ADSM systems and using their storage facilities:

  This is an environment that uses OS/2 Warp Server Backup/Restore as an ADSM API application. OS/2 Warp Server Backup/Restore is responsible for the definition and operation of the backup and recovery of your OS/2 Warp Server system and data. After that data is associated with the ADSM storage media, OS/2 Warp Server Backup/Restore releases control of the data, and ADSM assumes responsibility for managing and protecting it. ADSM can move the

data between its physical storage volumes, depending on the data hierarchy and migration thresholds defined.

- OS/2 Warp Server where ADSM takes full responsibility:
  - This is an environment that uses an ADSM Backup/Archive Client. ADSM is responsible for the backup and restore of your OS/2 Warp Server system and data. It is possible to place backups on the ADSM server.

  Following are some main features of ADSM:

  - Enterprise storage management solution providing automated, unattended backups and long-term data archives
  - Direct integration with applications such as Lotus Notes and DB2/2, for more extensive and customized application backup
  - Very good administrator capabilities to manage the ADSM server from any ADSM client platform
  - Easy-to-use Graphical User Interfaces
  - Robust Server Database
  - Online incremental backup/restore for OS/2 Lotus Notes databases
  - Good security capabilities that ensure that only authorized systems can back up or restore data
  - Open API providing critical online backup services to data-intensive applications. As you can see in Figure 159 on page 317, ADSM clients and servers run on a wide range of the most popular machines and platforms:
    - Microsoft
    - Novell
    - AT&T
    - HP
    - NEC
    - DEC
    - Siemens

ADSM is covered in more detail in an IBM Redbook titled *Using ADSM to Back Up OS/2 LAN Server and Warp Server*, SG24-4682.

11.18.11 Recovery

OS/2 Warp Server Backup/Restore restores data that has been backed-up on your system or transferred from another OS/2 Warp Server system. To restore your data you must define a restore method. A restore method tells OS/2 Warp Server Backup/Restore everything it needs to know in order to perform a restore. The window is the same that we use for backup; it contains information like:

- **The data to be restored:**
  - We select the backup set to use, and the files and/or folders that we want to restore.
- **The versions of the files that are to be restored:**
  - If your backup method specifies multiple kinds of backup, when you restore you must specify which version to restore. You can restore the most recent version, or you can even choose a version closer to a specific date and time.
- **Where the data is to be stored:**
  - You can restore the data to its original location and name or decide to change location and/or name. It is also possible to specify a desire to see a preview of the data that is going to be restored.

To see the Restore Methods that are defined for your system, open the Restore Methods container, as shown in Figure 160 on page 320. This container displays all the available Restore Methods.

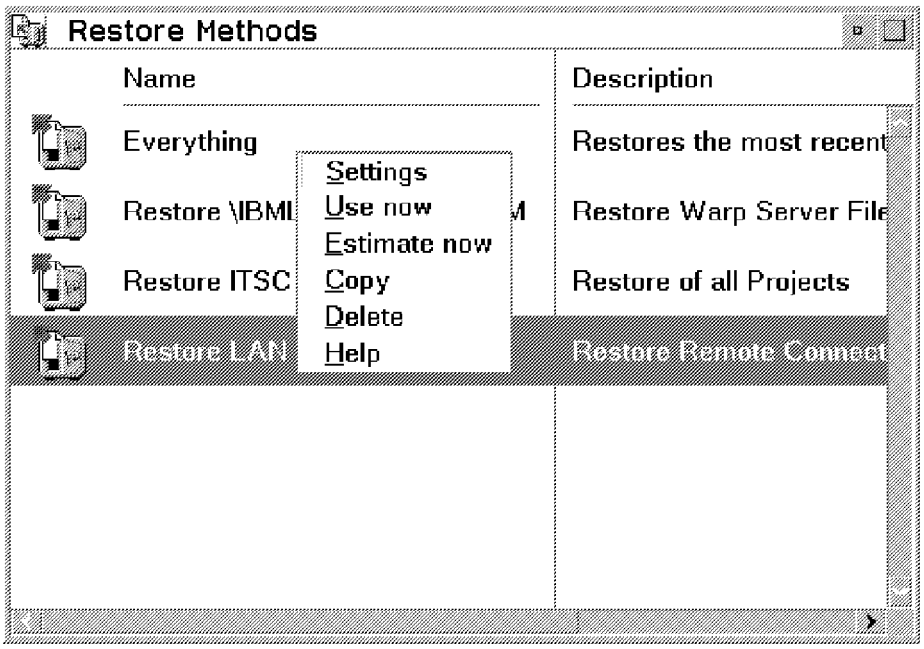

Figure 160. Warp Server Backup/Restore Restore Methods

By pressing the right mouse button, you will get the pop-up menu that gives you the option to do:

- **Settings**
  - Opens the Restore Method screen so that it can be edited.
- **Use now**
  - This will start using the Restore Method, still giving you the possibility to not deselect some files or some subdirectories.
- **Estimate now**
  - The Restore Method screen has the Preview feature, too. This option scans the disk and calculates how many files and folders will be backed up by the Restore Method.
- **Copy**
  - This option gives you the possibility to copy an existing Restore Method to a new Restore Method with a different name.
- **Delete**
  - Deletes your Restore Method.

Another way to get into your Restore Method is to double-click on the Restore Method container:

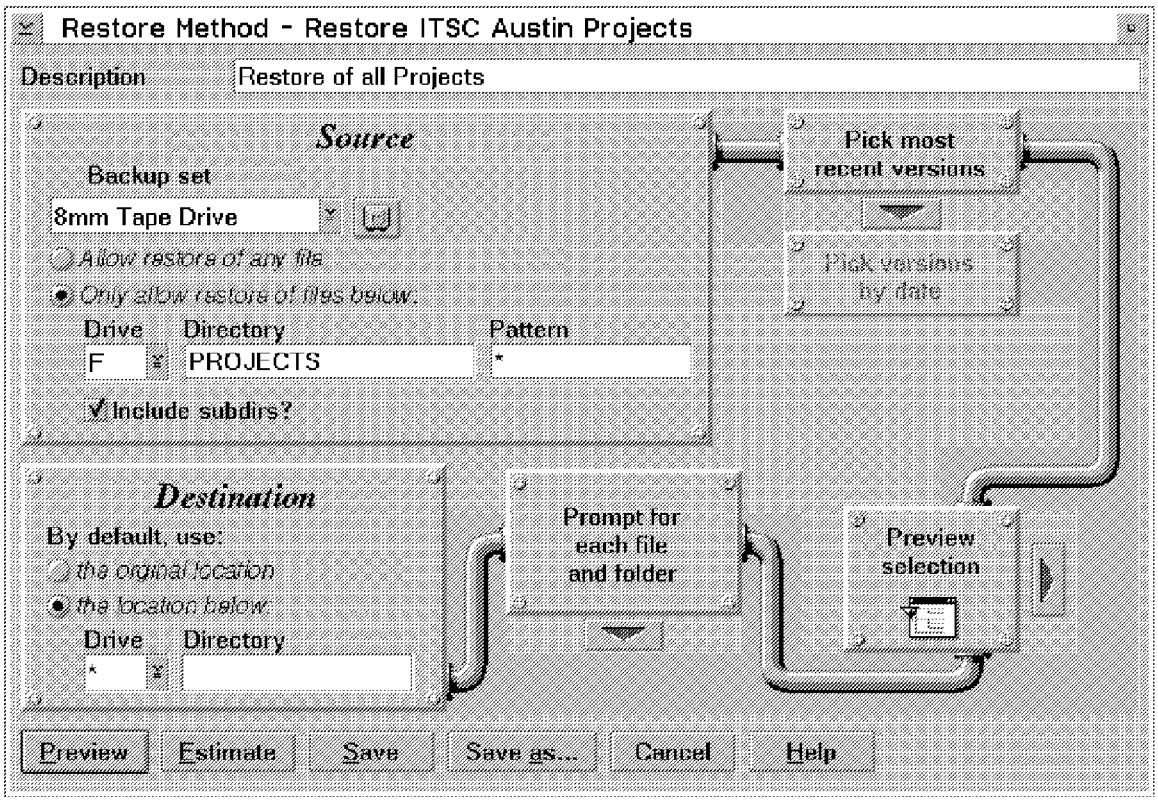

Figure 161. Warp Server Backup/Restore Restore ITSC Austin Projects

As you can see in Figure 161, following the flow of the pipelines, you can select which kind of restore you need. After you have defined a Restore Method for your needs, you can preview if your choices are exact:

1. Preview presents you with a tree view of the objects that are selected. In this selection you have the possibility to deselect the objects that you do not want to restore.
2. Estimate determines the total amount of data to be restored and estimates how long the time for restore will take.
3. Use now runs the Restore Method immediately and captures statistical information on the restore process, such as the amount of data that has been processed, the time taken to perform the restore, and the average data rate.

11.18.12 Backup and Restore Sound Feature

OS/2 Warp Server Backup/Restore can play audio files when certain events occur. To use these facilities your PC should have a sound card. You can now see some of these events and when to set them to start:

- Starting Automatic Backup whenever a scheduled event occurs

- Starting Manual Backup when a manual backup is started
- Starting Restore when a restore is started
- Startup whenever OS/2 Backup/Restore is started
- Successful Backup when all the objects in the backup have successfully been backed-up
- Unsuccessful Backup when a backup finishes its process and could not successfully back up one or more objects
- Successful Restore when a restore terminates and all specifics objects have been successfully restored
- Unsuccessful Restore when a restore process finishes and one or more objects were not able to be restored
- Insert diskette when OS/2 Warp Server Backup/Restore needs another diskette during Backup or Restore
- Insert RW Optical when another Read/Write optical disk is needed during Backup or Restore
- Tape Cartridge Required when another tape cartridge is required during Backup or Restore
- Tape Cartridge Full when a tape cartridge is full and a new one is required
- Tape Drive rewinding when the tape is being rewound at the end of the Backup/Restore

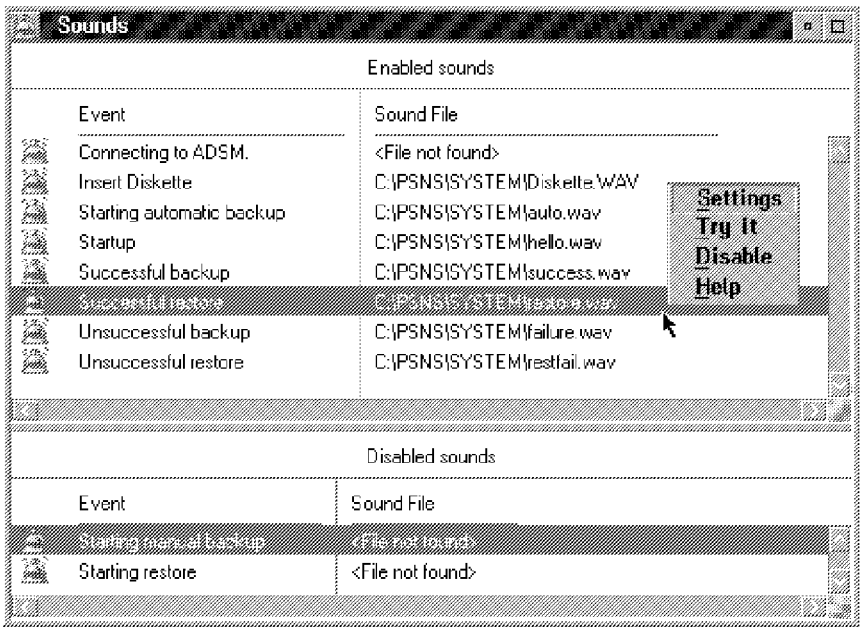

Figure 162. Backup and Restore Sounds

To see the sound events as shown in the figure above you need to open the Sounds window in the Backup/Restore main window. As you can see in Figure 162, the window shows the Sound Events that are available on your system and even the associated Sound File. As you can see, they are .WAV files and can be replaced by your personal .WAV file if you want.

Chapter 12. Windows NT

The administration of a domain or of the trusted domain must be done carefully and reliably. So the first thing you have to think about is how to organize the servers, domains, users, and user-groups. This means you have to make a design that best fits your business needs.

Looking at the Windows NT Server model, you will see that the domain concept is enhanced by extending it to a Trusted Domain model. However, it is not an enterprise-wide solution that could be used to operate a worldwide network (check 1.3, “Microsoft Trusted Domains” on page 10. for more information on why Windows NT lacks enterprise support from an architectural point of view). Also the standard Windows NT Server is not OSF 1.1 and DCE compliant. But the solution fits smaller and midrange networks and could be easily extended for using different domains and trusted domains in big companies.

To administrate a Windows NT Server domain, the administrator has to rely on a few managerial tools:

- The User Manager for Domains, accessible via the following path: [**Start — Programs — Administrative Tools — User Manager for Domains**], creates user accounts and defines what kind of actions a user can do. Also the administrator is aware of and oversees relationships between domains.
- The Server Manager tool lets you control what resources are shared and which services run off the server machines in the network. Also the NT Explorer and the Control Panel can do some of the Server Manager's job. Unlike the Server Manager, Windows NT Server lets you control any accessible server on your network.
- Printer access and print queues are controlled with the Print Manager tool.

In addition to the different managing tools, there are different privileged levels that can use all or only parts of the administration tools. These different rights are:

- Administrator
- Server Operators
- Accounts Administrator
- Backup Operators
- Print Operators

All these administration rights are classified with groups; you will find a detailed overview of the rights in section 12.3, “Group Administration Using the User Manager for Domains” on page 348. These different privilege rights for administration are a good way for an administrator to delegate some of his/her responsibilities. In addition, the user and the department administrator are pleased because they have the power to administrate some or all of the server components, depending on the rights assigned.

12.1 User Administration

A server network without users is useless. So the first step to administrating a server network is to define users. Windows NT Server also allows you to define different groups. The groups can define different workgroups, projects, or other organizational parameters.

The basic security units of an NT Server are domains. All these groups of servers share a common security database. The servers of the domain share all the security policies and all account information. There is also a possibility to back up the common security database by using a Backup Domain Controller (BDC).

The minimum for a domain is one computer running Windows NT Server and acting as a domain controller. This enables the functionality of a secure network environment. Therefore, it is important to have an account and user management. This is handled by the User Manager for Domains, which provides administrators the ability to:

- Manage a domain's security policy
- Manage trust relationships between different domains in the network
- Assign Logon Scripts to user accounts
- Define a user's network connections and desktop environment
- Create, modify, and delete user accounts in the domain
- Manage groups and group membership within the accounts in a domain

As mentioned before, there are different administration rights for different users. If you are the domain administrator, you have all the rights necessary to manage accounts.

User administration is done by using the User Manager for Domains. It is accessible via the path [**Start — Programs — Administrative Tools — User Manager for Domains**]. You will see a list of all accounts and all groups defined in the domain as shown in Figure 163 on page 327.

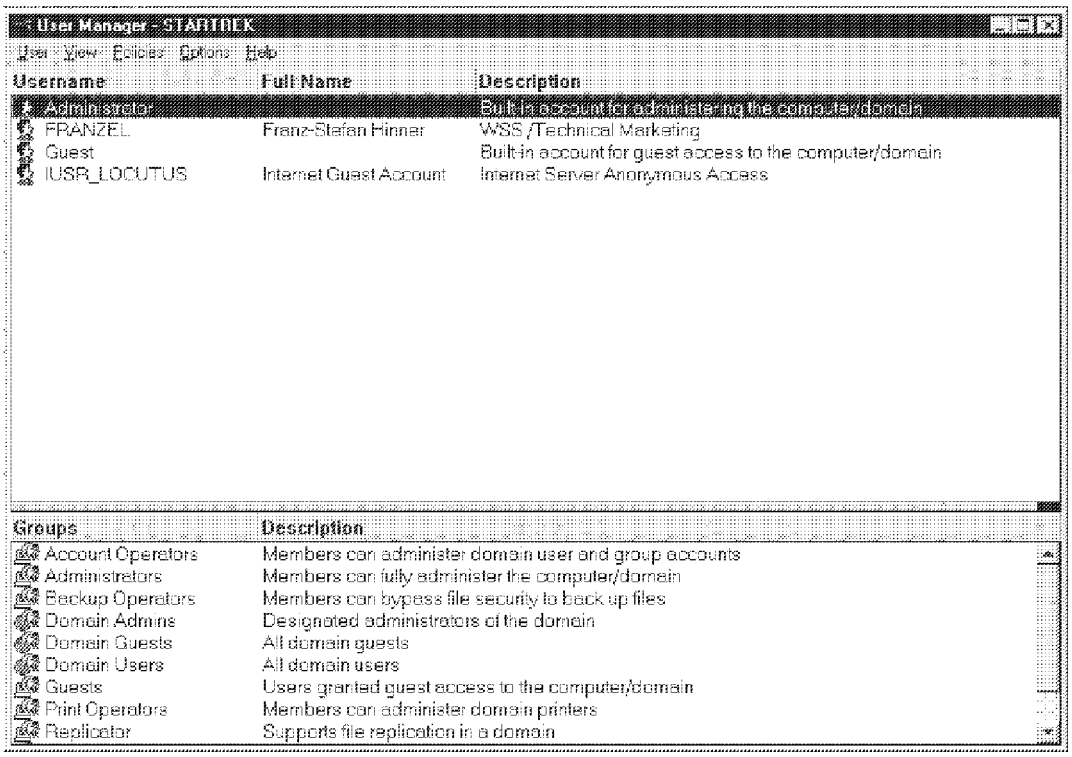

Figure 163. Adding New User Accounts in the User Manager for Domains

All displayed accounts, users, and groups are part of the domain you are administrating.

12.1.1 Creating User Accounts

A user account in Windows NT Server contains information about the user name, privileges, group membership, and password. Table 19 on page 328 describes in detail the different entry fields and push buttons.

When user accounts are created for the first time, a Security Identifier (SID) will be automatically assigned. This unique number identifies the account inside NT Server's security system.

If you are deleting an account, its corresponding SID will be deleted; that means SIDs are always unique. This implies that after deletion and redefinition of a user ID, this ID will not inherit the previous rights and permissions even if the same ID is used.

1. Select **User** from the action bar of the User Manager for Domains window to get the User's pull-down menu.
2. Form the pull-down menu select **New User...** to create an account. The window presented is shown in Figure 164 on page 328.

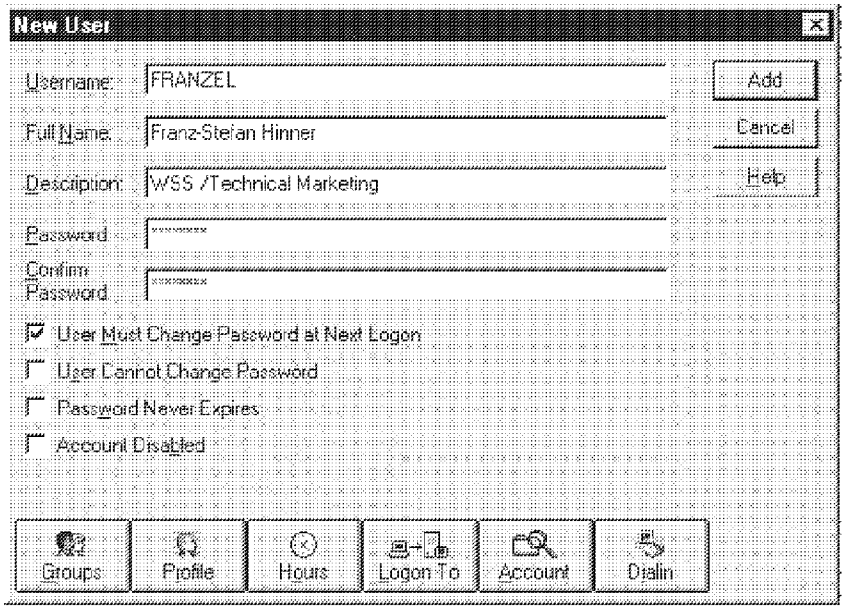

Figure 164. Creating a New User Account in the User Manager for Domains

3. Type in the required user information as described in Table 19 and press **Add**.

Rules for User Names

A user name may contain up to 20 characters, and blanks are allowed; however, using blanks requires you to enclose the user name within quotes when used in commands.

The following characters are not allowed: " / \ [ ] ; : | = , + * ? < >

Be aware that the user name and the user's password are both case-sensitive.

Table 19 (Page 1 of 2). Detailed Description of the User Properties Window

Entry Field / Push Button	Description
User name	Describes the unique user ID typed to log on to the server.
Password	The user's password for logon with his/her account, used in combination with the user name.
Full name	An advanced description that could hold the complete user's name.

Table 19 (Page 2 of 2). Detailed Description of the User Properties Window	
Entry Field / Push Button	**Description**
Hours	Defining the hours the user is allowed to logon with this domain account.
Logon To	Restricts the account to logon only from defined workstations. By default, a user can work with every workstation.
Account	Pressing this push-button, you are able to enter the expiration date of the user ID and the user's account type (local or global account).
Profile	Gives you the opportunity to define the user's home-directory and the user's profiles (Logon Script and the profile path).
Groups	Shows you the actual group assignments of this specific user and allows you to change or add new group assignments (see Figure 165 on page 331).
Dial In (only visible if RAS is installed)	Defines whether the user is allowed to use Remote Access Services (dial in) and if so which callback option to use.

The following table describes what options administrators have when they define accounts and passwords.

Table 20 (Page 1 of 2). Options for Password and Account		
Option	**Default**	**Description**
Account Disabled	OFF	This checkbox disables an account. No one can logon with that account until the checkbox is unchecked. Disabling an account does not remove the account from the account database. It can be reactivated any time. You can use this feature for temporary accounts or for template accounts.
Password Never Expires	OFF	When this box is checked, the password expiration policy is disabled. This is useful for accounts that need access to the network with the same permanent password, for example, replicator accounts, emergency administrator accounts, or guest accounts.

Table 20 (Page 2 of 2). Options for Password and Account		
Option	**Default**	**Description**
User Cannot Change Password	OFF	To prevent the changing of the password for an account, you need to check this box.
User Must Change Password at Next Logon	ON	The user will be forced to change his/her password the next time he/she logs on. Automatically, this parameter is set to OFF after the first logon. It is useful when setting up an account or when resetting an account password to default.

At the bottom of the New User dialog box are five buttons where you can define the properties of the user account. You can do this directly when creating an account or after the creation by doing a double-click on the account inside the User Manager for Domains.

The additional properties of the user account, like Groups, Profile, Hours, Logon To and Account, are discussed in the following sections.

12.1.2 Assigning User Accounts to Groups

To assign the same user rights to different users, you may group these users and assign these rights only once to the group and not to each member of the group itself. Only user-specific rights should be added via selection of a user instead of a group. The groups are also an organization criteria to tie users with the same tasks together. This topic and a discussion of several predefined groups can be found in 12.3, “Group Administration Using the User Manager for Domains” on page 348.

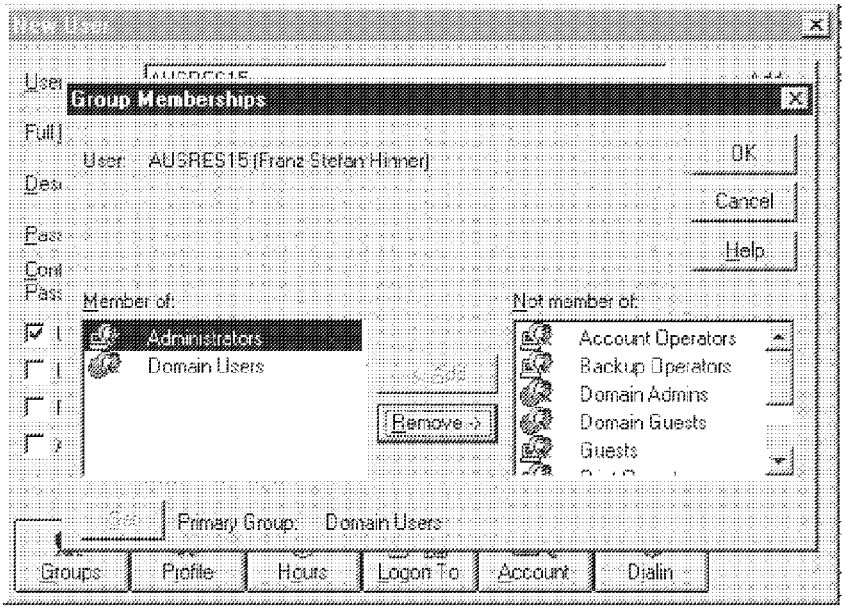

Figure 165. Add Group Membership to User Account

When adding group memberships to a user, you will notice that there are different kinds of groups. You will find a detailed discussion in section 12.3, "Group Administration Using the User Manager for Domains" on page 348. Note the icons next to the group names:

- A global group is indicated by a globe behind two faces.
- A local group is indicated by a computer behind two faces.

The scope of a local group is only within the domain it is defined in. In contrast, global groups are accessible by the entire network.

12.2 Set Up Environment of User Account

Under Windows NT Server, it is possible to define a user-specific environment (desktop, home directories, locales, and environment variables) for each user account that will be activated at logon time.

Most of the environment parameters are helpful in managing your network.

12.2.1 Set Up Logon Scripts

User Logon Scripts are simple batch files or executable programs that run automatically whenever a user logs on. The same script could be assigned to one or many user accounts. This makes it easy to set up profiles for groups.

The Logon Scripts work on computers running:

- DOS
- Windows 3.1/3.11
- Windows 95
- Windows NT

The Logon Scripts are located by default in the \WINNT\SYSTEM32\REPL\IMPORT\SCRIPTS directory.

12.2.1.1 Creating Windows NT Server Logon Scripts

The content of Windows NT Server Logon Scripts is similar to that of OS/2 Warp Server Logon Scripts except that there is no REXX language available today on NT. The Logon Script allows you to run local and network commands.

The Logon Script should be kept short, because of a problem concerning Windows 3.x users. If you assign a Logon Script to a Windows 3.x user, a virtual DOS session will start at logon time, and the script will be executed. This session typically lasts only about 30 to 45 seconds and is not configurable. This means if executing the script takes too long, it will quit, and the user will get a system-integrity violation error. This happens often when the Remote Access Services (RAS) are used over slow links.

12.2.1.2 Disable Windows NT Server Logon Scripts

When using slow modem links or connections through a TCP/IP gateway over asynchronous modem lines, it is useful to know how to disable Logon Scripts in those cases. This can be done by doing the following steps:

1. Select the **Control Panel** and choose the **Network** icon.
2. Choose the **Networks** button.
3. In the box **Other Networks In Use**, select **Microsoft LAN Manager**.
4. Select the **Settings...** option.
5. Unmark the checkbox for **Logon to LAN Manager Domain**.

12.2.1.3 Windows NT Server Logon Script Variables

Sometimes it is very helpful to have some special variables for Logon Scripts. Windows NT Server provides a couple of useful parameters for managing multiple users who are not members of the same domain. In the following table, you find a list of the supported variables:

Table 21. Variables for Windows NT Server Logon Scripts	
Variable	**Function**
%HOMEDRIVE%	Connects a user's local workstation drive letter to the user's home directory.
%HOMEPATH%	Variable for the whole home directory path name.
%HOMESHARE%	The user's home directory share name.
%OS%	The operating system of the user workstation.
%PROCESSOR%	The type of processor of the user workstation (such as 80386).
%PROCESSOR_ARCHITECTURE%	Detects the hardware architecture that is used and is only Windows NT related, because only here are differences found between the processor families like Intel, Alpha, and MIPS.
%PROCESSOR_LEVEL%	Specifies the processor level for the user's workstation. Examples for this are 21064 for Alpha systems or 5 for an Intel Pentium.
%PROCESSOR_IDENTIFIER%	Tells which processor type it is.
%PROCESSOR_REVISION%	Defines the internal version number of a processor. This is relevant for finding out if it is a Pentium version with the divide-by-zero error or not.
%USERDOMAIN%	The domain containing the user account.
%USERNAME%	The user name of the user.

Notice FAT Problems with Variables

Don't use the wildcard %USERNAME% when assigning home directories that are on FAT volumes if at least one of the selected user accounts has a user name longer than eight characters.

Remember the discussion in the 9.4, “File Allocation Table (FAT)” on page 162. FAT is still working with a 8 + 3 naming convention. This limit does not exist on NTFS partitions.

12.2.2 Creating Home Directories

Home directories are assigned to a particular user account. These directories are accessible for the user, and the user has access control over his/her directory. To provide home directory service for a network user, two things are needed:

1. Assign a logical drive letter for that directory to the user.
2. Create a directory with the proper access rights.

The creation of home directories for NT and non-NT users is inconsistent. The user's home directory will be assigned during logon, if accessible.

12.2.2.1 Set Up Home Directories for NT Users

After logon, Windows NT's default directory will be set to the user's home directory. For example, the command prompt, File Open and Save As dialog boxes, and applications that do have not have a predefined working directory will default to the user's home directory.

A home directory can be located on a Windows NT Server or on a Windows NT Workstation. However, it is recommended to put the home directory on a server so that it is accessible through the network.

To define a home directory, do the following steps:

1. Open the **User Manager for Domains** window.
2. Select **New user...**, or select **Properties...**, or double-click on a user entry to get **User Properties**.
3. Select the **Profile** button.
4. Check the radio button for **Connect**, select a drive letter, and type the path for the user, as shown in Figure 166 on page 335.

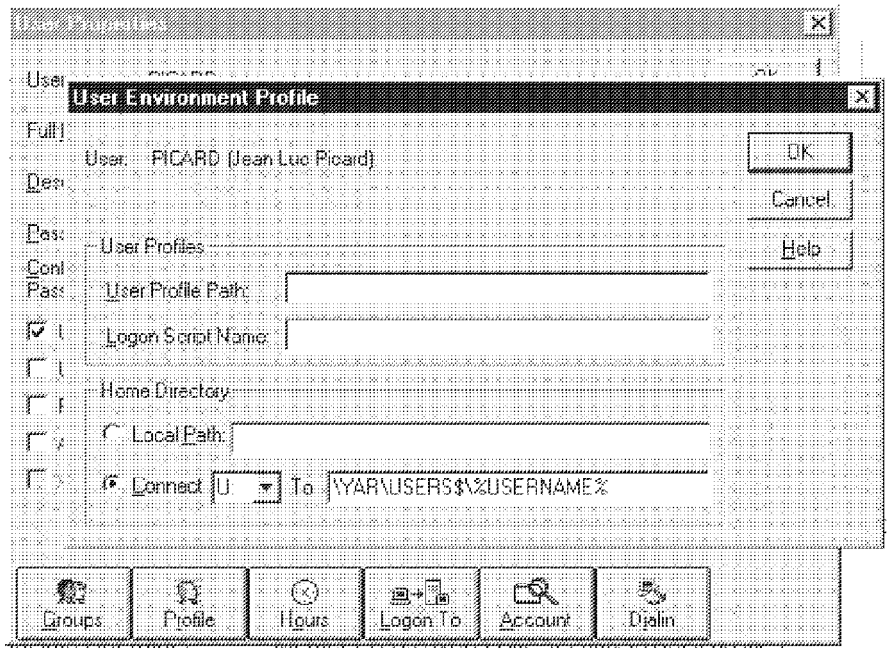

Figure 166. User Environment Profile Setup for Home Directory

It is possible to use variables like `%USERNAME%`. If you look, after setting up the user's home directory, at the user environment profile, you will see that it has written the defined user name, here PICARD, to the path. This option is especially useful when you create a template. After a copy, the user name will again be replaced with the variable.

If a user name is bigger then eight characters and you are setting up the home directory as in the example, a message comes up that explains that this home directory can only be used by a Windows NT workstation but not from a DOS-based desktop (like Windows).

12.2.2.2 Set Up Home Directories for Non-NT Users

The creation steps for the home directories for non-NT users are the same as described in section 12.2.2.1, “Set Up Home Directories for NT Users” on page 334, but additionally you need to run the following command:

```
NET USE U: /HOME
```

assuming the drive letter `U` is used as the user's home directory.

This procedure must run once to connect the home directory under the following clients:

- DOS
- Windows for Workgroups
- Windows 95

12.2.2.3 Conclusions on Home Directories

The home directory functionality of Windows NT Server is very limited. It is very useful for Windows NT users but difficult to implement and manage for others.

Also the Windows NT Server does not support Disk Quotas now. There are offerings from third-party suppliers, but this functionality is not implemented in the standard environment.

However, the home directories are usable without difficulties with Windows NT. Since support for other client platforms is limited, this home directory design makes it a typical Windows NT function. It is not very easy to use in comparison with the other home directory structures that are used in IBM OS/2 Warp Server Advanced and Novell NetWare Version 4.1.

The biggest problem managing home directories is that the size of a home directory cannot be limited. For example, IBM OS/2 Warp Server Advanced provides DASD limits for home directories as well as for other directories.

12.2.3 Managing and Limiting User Accounts

To manage a huge server environment, you need at least a couple of sorting, duplicating, and command-prompt options. Also it is useful to have functions to limit the user's account.

The network management is easier when you have commands you can run using an input file or batch file in addition to using the graphical user interface. This makes it a lot easier to add a large quantity of users to the domain. An important requirement is the possibility to update user information or to make templates for standard users.

On the the other hand, it is useful to limit some functions such as time, logon PC, or password rules. This has an influence on the security.

12.2.4 Copying User Accounts

It is very useful when setting up user accounts that should have the same functionality to copy them or use standard user-templates. This makes it easier for an administrator to set up many user accounts with the User Manager for Domains graphical administration tool. To copy a user account, do the following steps:

1. Start the User Manager for Domains.
2. In the list of users, select the user account, or template user account, to copy from.
3. Select the **Copy...** option from the User pull-down menu.

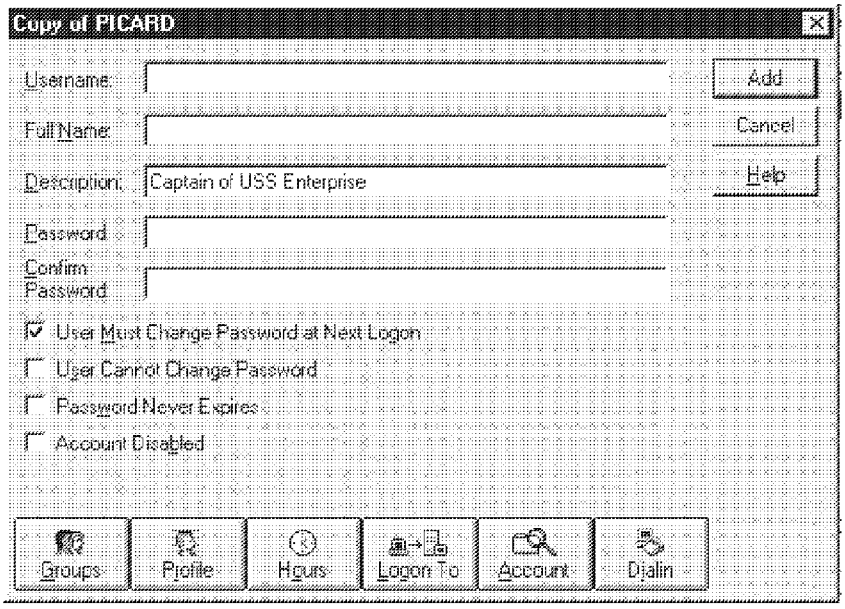

Figure 167. Copy User Account in the User Manager for Domains

4. Modify the properties of the account by typing in information about the new user.

   As you will notice, description information is copied for you as well as previously selected checkboxes.

5. Click the **Add** button to add the new user account.

   The Copy of <username> window will remain open for you so that you can copy as many user as you need to.

6. Once done, select **Close** to exit.

12.2.5 Deleting User Accounts

Sometimes you have to delete a user account when it is not needed in the domain anymore. Keep in mind that you can disable an account when it is not needed for a while, but will be needed again in the future.

To delete user accounts, perform the described steps:

1. From the list of users in the User Manager for Domains, click once on the user account you want to delete.
2. From the User pull-down menu, choose **Delete**.
3. Select **OK** at the User Manager for Domains confirmation window.
4. To finally delete the user account, select **Yes** at the next window.

12.2.6 Adding Many Users at Once

If you have to create a new network environment, it is necessary that you have some special tools or functions to add many users at the same time. The User Manager for Domains is one way, but will not be a practical solution for adding 100 users to a domain. So there must be another way to make these functions available for the administrator.

The code example shown in Figure 168 illustrates a batch file that uses standard `NET USER` commands to add many users at once. This is not as comfortable as the graphical interface, but it is a much faster way to customize a new server and set up the user accounts.

```
NET USER PICARD   password  /ADD  /FULLNAME:"Jean Luc Picard"
NET USER RIKER    password  /ADD  /FULLNAME:"John Riker"
NET USER TYAR     password  /ADD  /FULLNAME:"Tasha Yar"
NET USER DTROI    password  /ADD  /FULLNAME:"Dana Troi"
NET USER DATA     password  /ADD  /FULLNAME:"DATA"
NET USER ROLAREN  password  /ADD  /FULLNAME:"Ro Laren"
NET USER NEELIX   password  /ADD  /FULLNAME:"Neelix"
NET USER KES      password  /ADD  /FULLNAME:"KES Ocampa"
```

Figure 168. NET USER Commands Issued from a Windows NT Command Prompt

This example adds each user with the password "password" and enters the user's full name in the FullName field. The NT Resource Kit, which is separately available for Windows NT, holds another possible way to add user accounts, the program ADDUSERS. We will not discuss this here because we compare only the base functions that come with the network operating system.

12.2.7 Manual Disable/Enable an Account

To disable or enable an account because of security reasons — maybe a user is working with another project or is on a long vacation — perform the following steps:

1. Start **User Manager for Domains**.
2. Double-click the appropriate user ID or select **Properties** in the User pull-down menu.
3. Check or uncheck the box for **Account Disabled**.

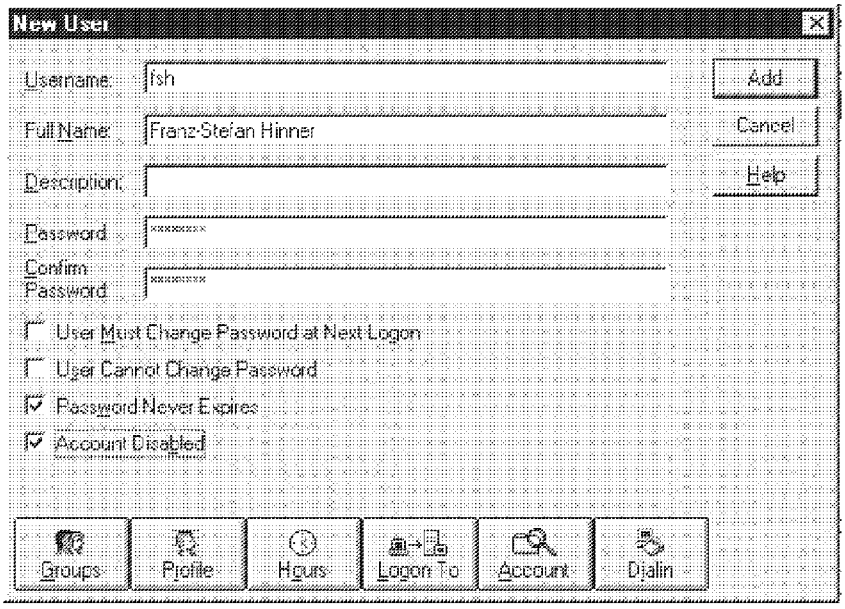

Figure 169. Disable/Enable User Account from the User Manager for Domains

4. Click on **OK** to close the User Properties dialog box.

12.2.8 Limiting Logon Time

The ability to control, due to security or organizational reasons, the time a user account is able to log on to the domain might be important. A user is allowed to log on within the defined time frame and optionally can be forced to log off when the time-slice expires. This is especially helpful if an administrator has only a short maintenance window for the domain. This option to force a logoff can be set in the account policy dialog box, as discussed later in this chapter. Be aware that you have to carefully decide which options are set because every change needs an administrator action. For example, a balance department needs access for the year-end and has to be manually set to have the access rights for this period of time.

Most of your user accounts do not need access to the server environment on weekends. The balancing of different users is another issue. So it is possible that 1000 users have access to the server but only a few during the day, another group of users only in the morning, and the rest from 8 p.m. to 11 p.m.. The advantage is that you can force the logoff for users who have forgotten to log off.

Also the logon control is useful for backup periods so you can be sure that there are no datafiles in use during the logon process. To restrict the logon hours, do the following:

1. Start the **User Manager for Domains**.

2. In the list of users, select the the user you want to change the properties for.
3. As shown in Figure 170, select **Properties...**. from the User pull-down menu.

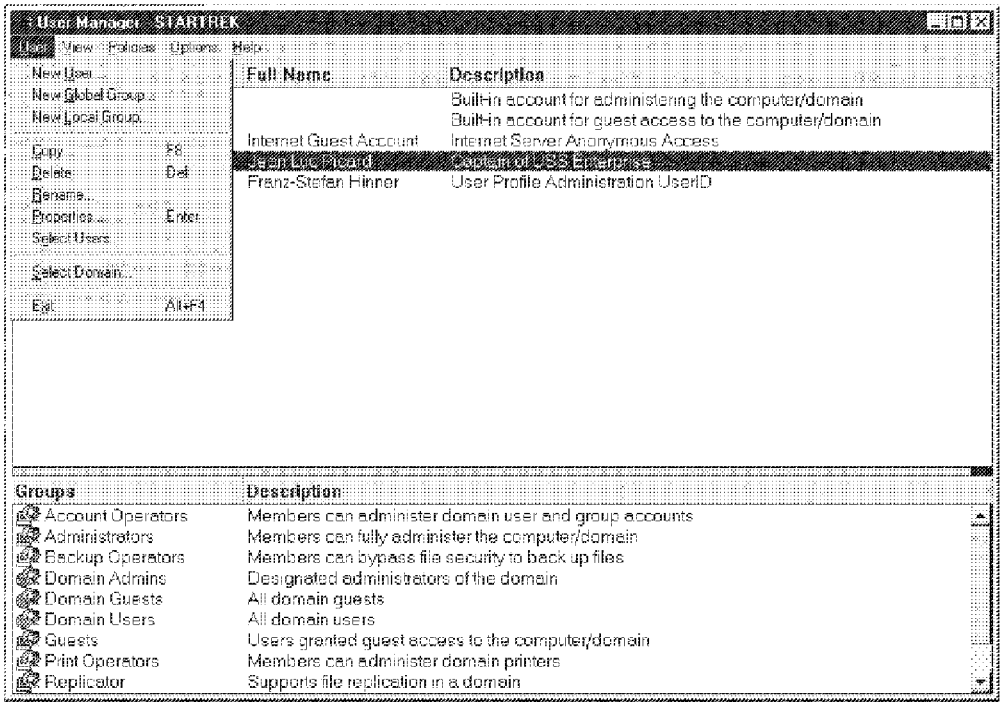

Figure 170. User Menu-Properties

4. Click on the **Hours** button to open the Logon Hours window.
5. Mark the hours the user can log on, and select **Allow** or define hours the user is not allowed to log on; then select **Disallow**.

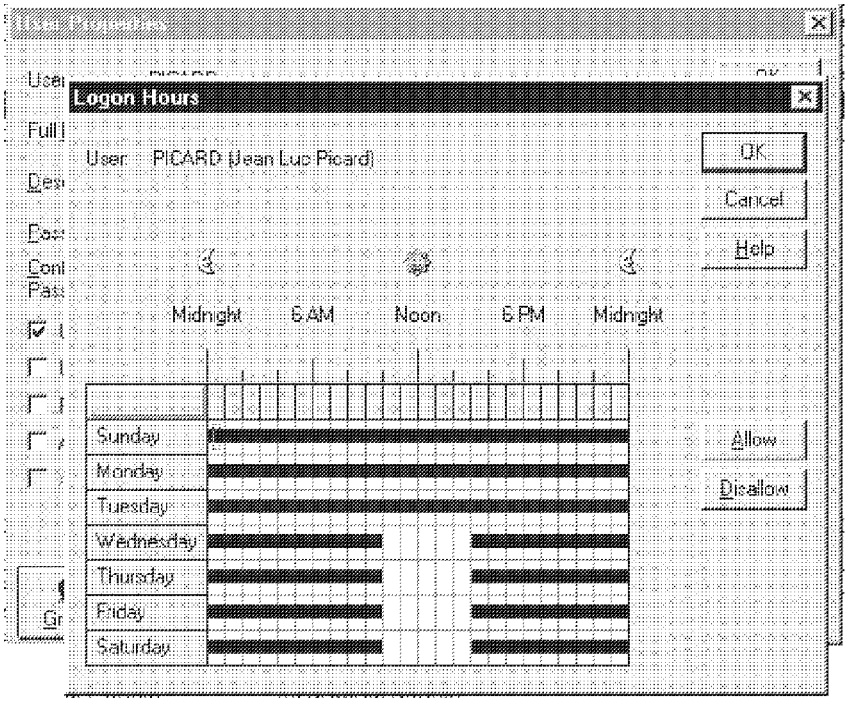

Figure 171. Logon Hours Definition

6. Choose **OK** to close the dialog box.

To set up the forced logout, choose **Account** from the User Manager for Domains' Policies pull-down menu. We will explain the detailed menu in the next step. To complete the Logon Hours limitation, we want to describe how to set the forced logout using the account policy option and also by using the `NET ACCOUNT` command, which forces users to log off. To set up the forced logoff, follow these steps:

1. Start **User Manager for Domains**.
2. Select **Account...** from the Policies pull-down menu.

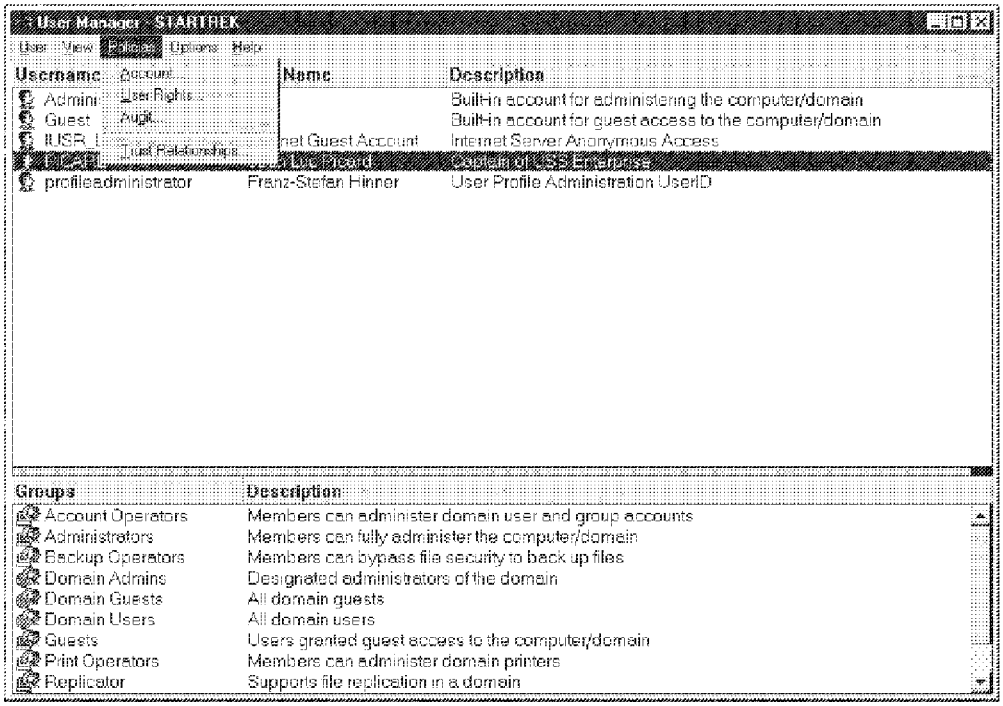

Figure 172. Modify Account Policies

3. In the Account Policy window shown in Figure 173 on page 343, check the box for **Forcibly disconnect remote users from server when logon hours expire**.

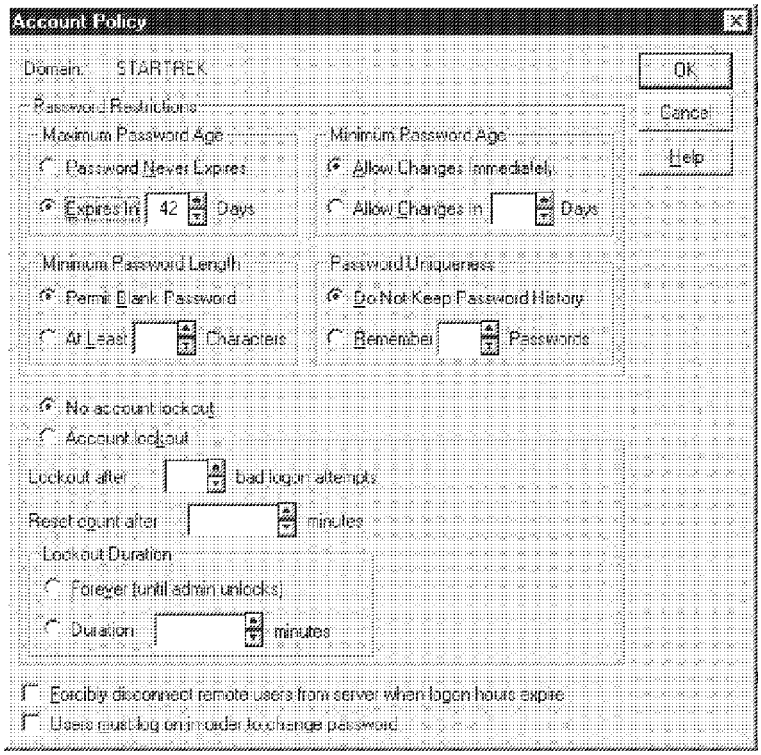

Figure 173. Forced Logoff

4. Click **OK** at the Account Policy window.

To set up the forced disconnection, you can also use the `NET ACCOUNTS` command to change the account policy.

```
NET ACCOUNTS /FORCELOGOFF:<minutes> /DOMAIN
```

The parameter `minutes` describes the number of minutes to notify the user before forced logoff will occur; the `DOMAIN` parameter is only needed when the command is executed on a Windows NT workstation. Only the Windows NT Server automatically updates the domain security database.

To turn off the forced logoff function when the logon time expires, enter the following `NET ACCOUNTS` command:

```
NET ACCOUNTS /FORCELOGOFF:NO /DOMAIN
```

12.2.9 Conclusion on Logon Hours

The idea to manage these logon times and force users to log off when their logon time expires is great. However, the graphical Logon Hours menu has a couple of difficulties. It is not possible in one step to define logon hours like Monday-Wednesday from 8 a.m. until 12 p.m. and Thursday-Friday from 8 a.m. until 11 p.m. with a maintenance phase from 6 p.m. to 8 p.m.. Such

definitions will take three steps. The graphical user interface gives an easy overview of the possible logon times.

You can also look at the defined logon hours by issuing the `NET USER` command from the command prompt.

```
NET USER PICARD
```

assuming `PICARD` is the user name. The output from this command is shown in Figure 174.

```
User name                    PICARD
Full Name                    Jean Luc Picard
Comment                      Captain of USS Enterprise
User's comment
Country code                 000 (System Default)
Account active               Yes
Account expires              Never

Password last set            12/8/96 4:44 PM
Password expires             1/20/97 3:31 PM
Password changeable          12/8/96 4:44 PM
Password required            Yes
User may change password     Yes

Workstations allowed         VULCAN,HUMAN,CARDASSIAN,TALAXIAN,OCAMPA,
FERENGIE,ROMULAN,BAJORAN
Logon script
User profile
Home directory
Last logon                   Never

Logon hours allowed          Sunday 12:00 AM - Wednesday 10:00 AM
                             Wednesday 3:00 PM - Thursday 10:00 AM
                             Thursday 3:00 PM - Friday 10:00 AM
                             Friday 3:00 PM - Saturday 10:00 AM
                             Saturday 3:00 PM - Sunday 12:00 AM

Local Group Memberships
Global Group memberships     *Domain Users
The command completed successfully.
```

Figure 174. Output of the NET USER PICARD Command

12.2.10 Limiting Workstations to Logon

Another possibility to limit logons is to define fix machines. This can extend the security. You can define up to eight workstations.

1. Start **User Manager for Domains**.

2. Double-click on the account you want to modify to get the User Properties window or use the **Properties...** option from the User pull-down menu.
3. Click the button **Logon To**.
4. Check the box for **User May Log On To These Workstations**, and type in the computer names the user is allowed to log on to as shown in Figure 175.

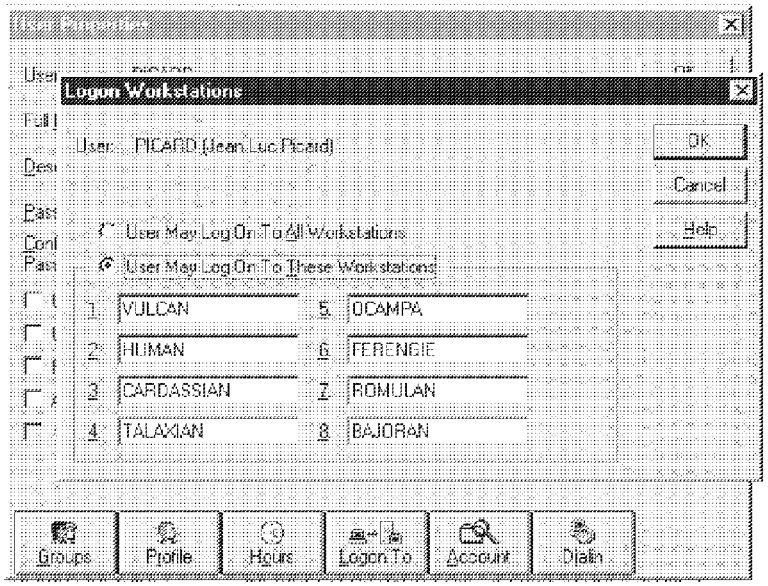

Figure 175. Logon Workstations Window

5. Click **OK** to close the dialog box.
6. Select **OK** to close the User Properties dialog box.

12.2.11 Special Account Information

The last option for limitation and user account control is the Account Information function. Within this dialog, you can set up an expiration date as well as an account type.

The account type defines whether the account is a local account or a global account. The limitation for a local account restricts the account for use on the local NT system only and does not allow the user to participate in a domain. By default, user accounts will be set to global accounts.

The account expiration option setting is a good way to manage temporary accounts or to manage students working only a limited time in the company.

1. Open **User Manager for Domains**.

2. Double-click the account you want to modify or select **Properties** from the User pull-down menu.
3. Select the **Account** button. The window presented to you is shown in Figure 176.

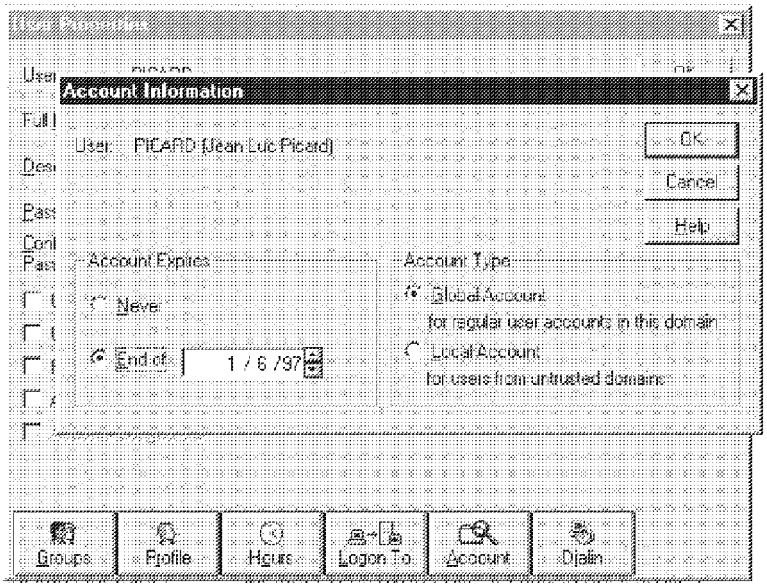

Figure 176. Add Account Information

4. Select the appropriate radio-buttons that suit your needs.
5. Select **OK** to end dialog box.
6. Choose **OK** to close the User Properties window.

12.2.12 Account Policy

To complete the account management tools, you need a tool to control passwords and set expiration times, set password lengths, limit the number of incorrect logon attempts, and establish password uniqueness.

These are all security-related functions. So you, as an administrator, can decide the minimum length of a password to protect your network from user accounts with a password of only two or three letters. Also you can protect your network from having users that are continuously using the same password by keeping a history. To set up the account policy, do the following:

1. Start the **User Manager for Domains**.
2. From the Policies pull-down menu, select **Account...**
3. In the Account Policy window make the necessary changes as shown in Figure 177 on page 347 and select **OK** to close the dialog box.

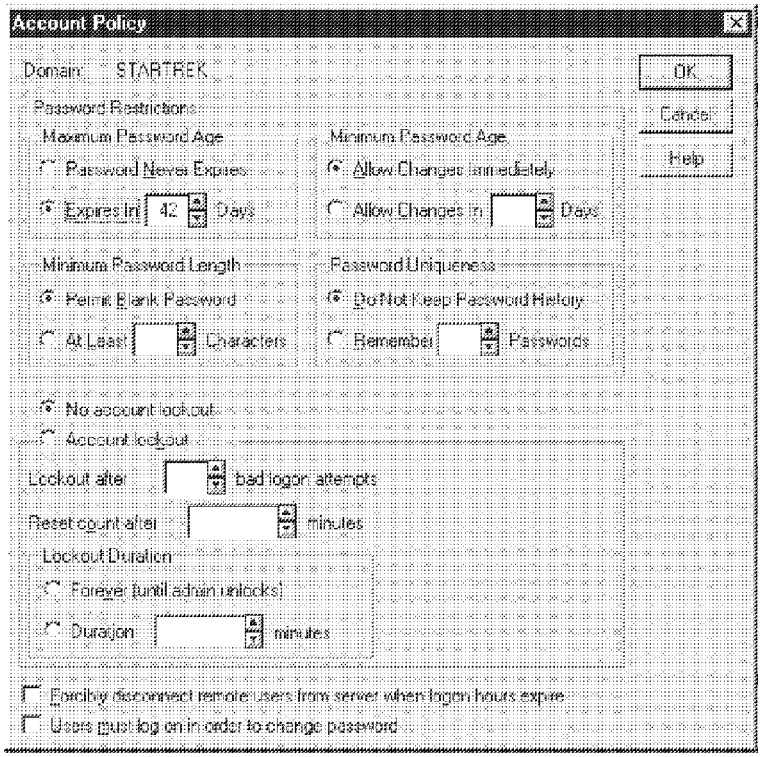

Figure 177. Account Policy Window

In the Account Policy window, you will find different things to change. There is a Password Restrictions section and an Account Lockout section. At the bottom of the dialog box you find a section with special functions for forced logout and password change. With the different fields, you can set up the displayed functions in the following table.

Table 22 (Page 1 of 2). Account Policy Dialog Functions

Field	Description
Maximum Password Age	Sets the time in which a password expires and requires it to be changed by the user.
Minimum Password Age	Sets the time value after which the password could be changed again. Normally, you will allow changes immediately.
Minimum Password Length	Defines how many characters must be used for a password.
Password Uniqueness	Specifies the number of passwords that will be recorded in a history record before it can be reused. This selection requires a minimum password age.
Account Lockout	Will prevent anyone from logging on to the account after a certain number of failed attempts.

Table 22 (Page 2 of 2). Account Policy Dialog Functions	
Field	**Description**
Lockout after	Defines how many bad logon attempts can be performed before the user's account will be revoked.
Reset count after	This defines the time in which the bad logon attempts count starts over.
Lockout Duration	Defines if the administrator must unlock the account manually or the amount of time that must pass before the user can try another logon attempt.
Forced disconnection for remote users from server when logon hours expire	Establishes that the user can't stay connected after his logon hour assignment as assigned in the Logon Hours dialog box. The system will perform an automatic disconnect.
User must log on to in order to change password	Checking this allows password change after legitimization.

12.3 Group Administration Using the User Manager for Domains

The group design allows you to summarize users in a group with identical rights in the network. This option is a great administration facility for managing many users. Also it makes it easier and faster to grant multiple users access to network resources and simplifies security and general account management.

After the user account definition, this is the second most important point for simple network management.

Windows NT also distinguishes between local and global groups. The global groups are represented by an icon with two faces superimposed over a globe, and the local group is represented by two faces superimposed over a computer.

12.3.1 Local Groups

When creating a local group, you should keep in mind that it is only useful on a Windows NT Server or on a Windows NT Workstation. However, local groups can contain global groups from the domain or trusted domains, user accounts, and user accounts from the workstation.

There are predefined local groups with different functionality, as summarized in Table 23 on page 349.

Table 23 (Page 1 of 5). Default Local User Groups

Group Name	Description	User Rights	Special Abilities Granted to Local Groups / Functions
Administrators	This is the most powerful group.	Log on locally	Create, manage user accounts
		Access this computer from the network	Create, manage global groups
		Take ownership of files	Assign user rights
		Manage auditing and security log	Lock the server
		Change the system time	Override the server's lock
		Shutdown the system	Format server's hard disk
		Force shutdown from a remote system	Create common groups
		Back up files and directories	Keep a local profile
			Share, stop sharing directories
			Share, stop sharing printers

Table 23 (Page 2 of 5). Default Local User Groups

Group Name	Description	User Rights	Special Abilities Granted to Local Groups / Functions
Server Operators	The Local group "Server Operators" has all of the required rights for the domain server management. It can create, delete, and manage networks shares as well as printer shares. Also it has the right to backup and restore files, format server fixed disks, unlock and lock the server and change the system time. In addition, server administrators can shutdown servers and logon to the network from the domain server.	Log on locally	Lock Server
		Change System time	Override server lock
		Shutdown system	Format server's hard disk
		Force shutdown from remote System	Create common groups
		Backup files and directories	Keep local profiles
		Restore files and directories	Share, stop sharing directories
			Share, stop sharing printers

Table 23 (Page 3 of 5). Default Local User Groups

Group Name	Description	User Rights	Special Abilities Granted to Local Groups / Functions
Account Operators	Membership in this group allows use of the User Manager to manage user accounts and groups for the domain. The members cannot delete one of the following groups: Administrators, Domain Admins, Account Operators, Backup Operators, Print Operators, and Server Operators. They have no rights for deleting or modifying user accounts of administrators. They use the Server Manager option to add computers to a domain, but cannot administrate the security policies. The members can shut down servers and log on to the servers.	Local Logon	Create and manage user accounts, global groups and local groups
		Shut down the system	Keep local profile

Table 23 (Page 4 of 5). Default Local User Groups

Group Name	Description	User Rights	Special Abilities Granted to Local Groups / Functions
Print Operators	The Print Operators group is designed for users that must have full management for printers. The members can create, delete, and manage printer shares for an Windows NT Server. In addition they can shut down the server.	Logon locally	Keep local profile
		Shut down the system	Share, stop printers
Backup Operators	To backup and restore files on primary and backup domain controllers, you have to be a member of this group. The Backup Operator also has the permission to log on to the server interactively and shut it down.	Local Logon	Keep local profile
		Shut down the system	
		Backup files and directories	
		Restore files and directories	

Table 23 (Page 5 of 5). Default Local User Groups

Group Name	Description	User Rights	Special Abilities Granted to Local Groups / Functions
Users	The Users, members of the User group, have only minimal rights. They have the right to create and manage local groups, but because they have no access to the User Manager tool because they are not allowed to log on locally at the server (domain controller), they are only able to perform this task remotely.	(NONE)	Manage and create local groups (*)
Guests	For occasional access to the network, the local group Guests is designed. Guests usually contains the Domain Guests global group.	(NONE)	(NONE)

Note: (*) In order to actually create and manage local groups, the user must either have the right to log on locally at the server, or must have access to the User Manager tool.

For maintaining security permissions on a local domain, groups are very useful because local groups can contain global groups and users from trusted domains. This means it is possible to create a group on a domain that contains both users on trusted domains and users on the current domain. They can also contain global groups.

Special Groups: The Windows NT Server also includes other default groups that contain no members, by default. These groups are only for definition issues and are used to set up how users can make use of the system. Also there are special groups reserved for special tasks, like replication. These groups are described in the following table.

Table 24. Special Default Groups

Group Name	Description
Creator/Owner	The user who creates or takes ownership of a resource.
Everyone	This group does not appear in the User Manager list, but you can assign rights and permissions to it. Every defined user ID is automatically a member of the Everyone group.
Interactive	Interactively logged on system users
Network Users	Users who access a resource through the network.
Replicator	Only user accounts used to log on to the replicator service should be held in this group. The Replication group is designed to provide support for replicator functions.
System	Operating System

Special groups are useful when Windows NT Server presents a list of users for tasks such as setting permissions. The generic groups also allow the definition of special tasks for users and assignments of permissions by user types. For example, if you want to prevent access to a particular directory or file over the network, simply assign no access permission to the Network users group.

12.3.2 Global Groups

Global groups are only used on Windows NT Server domains. They contain global groups from trusted domains and user accounts from the domain. Local groups cannot pass through trust relationships, but global groups can. A global group can be accessed anywhere you can see users and groups, such as in the permission dialog box in File Manager. At installation time the Windows NT Server creates the following three global groups:

- Domain Admins
- Domain Users
- Domain Guests

These three different global groups have, exactly like the local groups, different rights. In the following table you find a detailed description:

Table 25. Default Windows NT Server Global Groups	
Group Name	**Description**
Domain Admins	This is the classical administrator account group. The users of this group have administrator abilities. The members can administrate other trusted domains that have been added to the Domain Admins global group of their own administrator's local group. They can also administrate the home domain and the workstations of the domain. The Domain Admins global group is member of the domain's Administrators local group and the Administrators local group for every NT workstations in the domain. Automatically the built-in administrator user account for the domain is a member of the Domain Admins global group.
Domain Users	If you are member of the Domain Users global group, you have normal user access permissions to the domain itself and for any Windows NT in the domain. This group contains all domain user accounts and is by default a member of the user's local groups for the domain and for every Windows NT on the domain.

To avoid confusion between local groups and global groups, it is a good idea to add a prefix, like Microsoft has implemented the default groups. Especially on monitors with high resolution, it is hard to notice the difference between the two icons of local and global groups.

When using the same names for global and local groups, the groups will be treated as two different groups, and they will not pass their defined permissions to the other group.

12.3.3 Creating a New Group

The creation of new global and local groups happens in the same manner; the main difference is whether you choose "New Local group" or "New Global group". The definitions are done with the User Manager for Domains by following these steps:

1. Start the **User Manager for Domains**.
2. In the list of users, select the users for the new group by using the mouse and the CTRL or SHIFT key.
3. From the User pull-down menu, select **New Global Group** or **New Local Group**, depending on which type of group you want to create, as shown in Figure 178 on page 356.

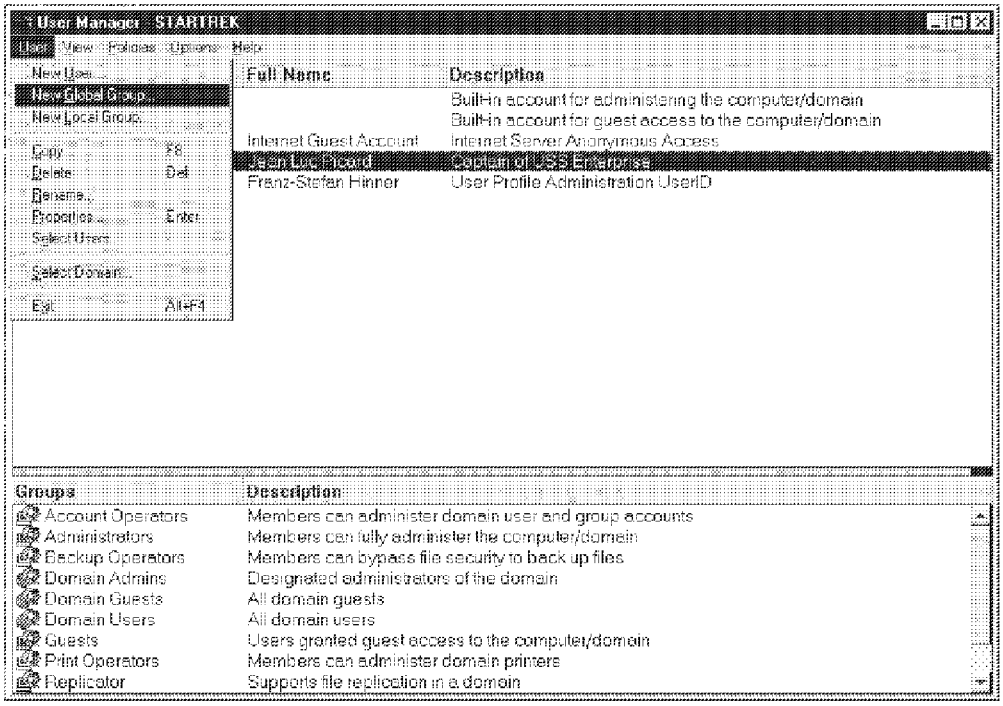

Figure 178. Creating a New Global Group

You will get the New Global Group window shown in Figure 179.

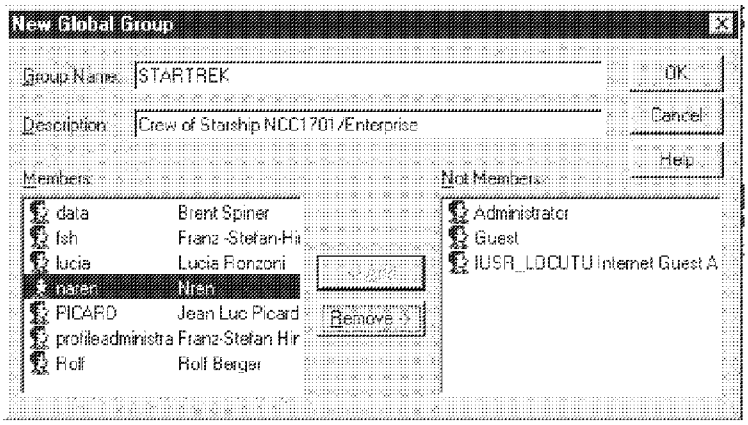

Figure 179. New Group Window

4. Add the group name, in this example **STARTREK**. Enter the optional group description.
5. Choose the **OK** button to create the new group.

12.3.4 Copying Groups

Sometimes it is useful to copy a group; so you use a formerly defined group as a template. To do this, do the following:

1. Start the **User Manager for Domains**.
2. In the list of groups, select the group to copy.
3. As shown in Figure 180, choose the **COPY** item from the User pull-down menu.

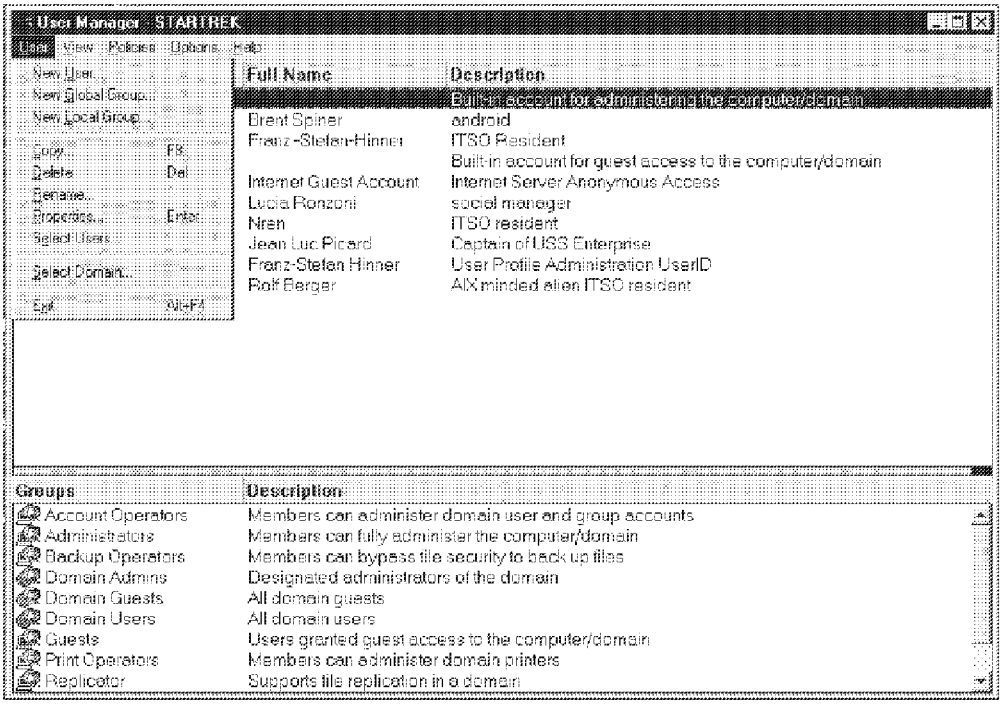

Figure 180. Copy a Predefined Group

4. At the **New Local Group** or **New Global Group** window, type in group name information and optionally overtype the description.
5. If you need to add additional members to the group, select the **Add** button and select user and/or group accounts in the following Add Users and Groups window. When done, select **Add**.
6. Choose the **OK** button to close the dialog box.

12.3.5 Deleting a Group

Sometimes it is necessary to delete complete groups, especially in case of finished projects. To delete a group, perform the following steps:

1. Start **User Manager for Domains**.
2. Select the group you need to select by clicking on it once.

3. Choose **Delete** from the User pull-down menu or press the **Del** key.
4. At the User Manager for Domains pop-up window, confirm the deletion of the group by choosing **OK**.
5. Confirm the deletion again by selecting the **YES** button at the subsequent information window.

12.3.6 Modify Group Properties

For some projects or groups, it is important to change the group properties to add additional rights to a defined user group to meet the project needs. To do this do the following steps:

1. From the User Manager for Domains window, double-click on the group to be modified to open the group's properties window.
2. Modify the desired properties.
3. Click the **OK** button to close the dialog box.

12.4 Managing Users within Groups

As described in the user section, you can add or remove group settings from a user ID. Especially in huge environments, it is not useful to touch every user account to add or remove group settings. It is easier to manage this in the group properties section.

12.4.1 Adding/Removing Users to/from a Local Group

Remember that users of local groups have special abilities on the local machine depending on the membership of the group. To add users, perform the following procedure:

1. Start the **User Manager for Domains**.
2. Open the Local Group Properties by double-clicking on a local group (global and local groups are differentiated by their icons). The Local Group Properties window will be presented to you as shown in Figure 181 on page 359.

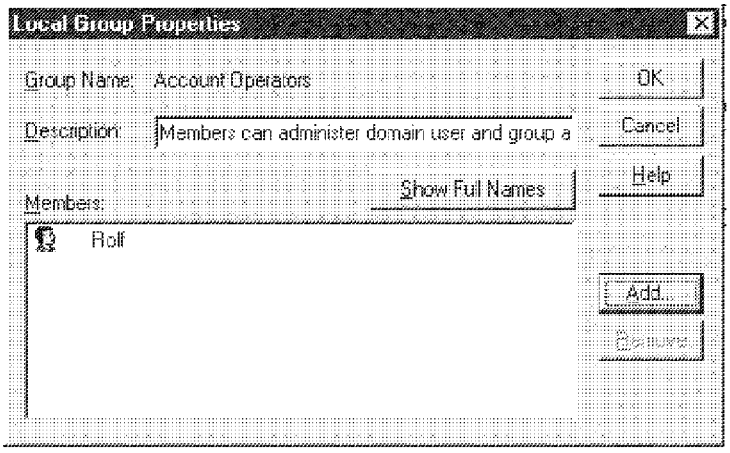

Figure 181. Properties of Local Group

3. Click the **Add** button to display the Add Users and Groups window shown in Figure 182.

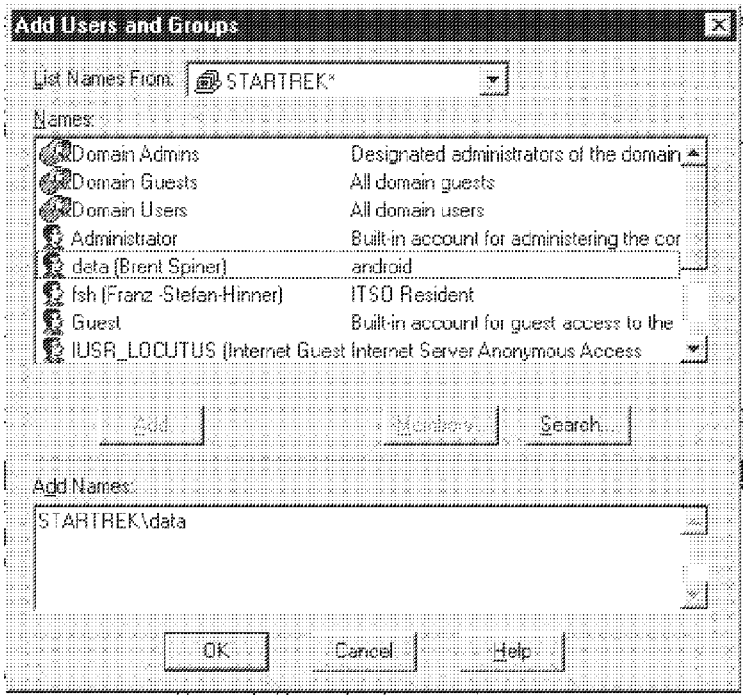

Figure 182. Adding Users to Local Groups

4. Select the users and groups you want to add to the local group and click on the **Add** button.
5. Click on the **OK** button to close the Local Group Properties window.

12.4.2 Adding/Removing Users to/from a Global Group

You can add or remove users and groups to/from a global group and assign appropriate rights to a group of users.

1. Start the **User Manager for Domains**.
2. Double-click on a global group to modify the properties.
3. In the list of Members, select the users you need to remove from the group.
4. In the list of Not Members, select the users you need to add to the group.

   The Global Group Properties window is shown in Figure 183.

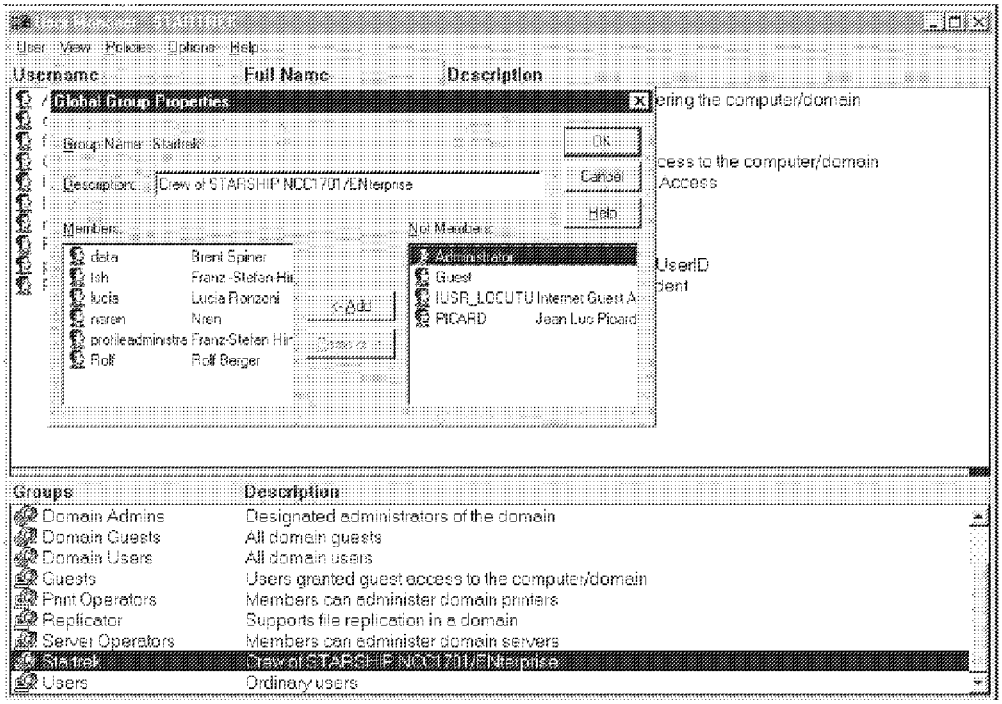

Figure 183. Add Users to Global Groups

5. Click the **OK** button in the Global Group Properties window to make changes active.

12.5 Sharing Resources

After setting up users and groups, the users should have a repository of files and directories that can be accessed. There are certain directories, files, and printers the user needs to have access to. Before the user can access a resource, the resources must first be made known to the network and shared.

12.5.1 Sharing Files and Directories

User-shared directories are treated like other hard disks. Once a directory is shared, the user can connect to it from his/her own workstation by issuing `NET USE` commands or by using the NT Explorer. If you share a directory at a server, everything in the directory is shared, the subdirectories and the data files. Permission restrictions can be applied and they are discussed in the next section.

To share directories or files, you have to be logged on to the server with a user account with administrative rights allowing the possibility to share directories. There are three different tools/ways to share directories.

- Server Manager
- NT Explorer (formerly called File Manager)
- Command Prompt

12.5.2 Using The Server Manager To Share Resources

To share the directory with the Server Manager, perform the following steps:

1. Make sure that the directory you want to share exists.
2. Start the **Server Manager** ([**Start — Programs — Administrative Tools — Server Manager**])
3. From the list of computers, select one computer on which you want to create a new shared directory.
4. From the Computer's pull-down menu, select **Shared Directories...**

   The Shared Directories window will be opened for you.

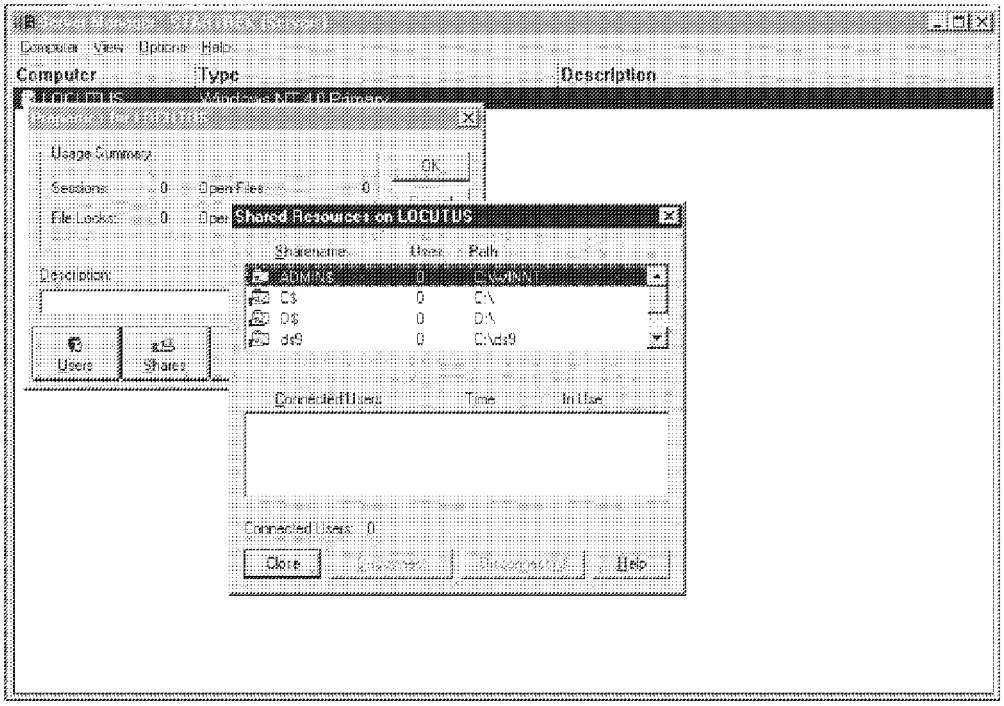

Figure 184. Shared Directories View

All currently shared directories at the selected computer will be displayed in the Shared Directories window.

5. To create a new share, select the **New Share...** button.
6. In the New Share window, provide a Share Name, a Path, and a descriptive comment. Select the appropriate radio button to set user limits.
7. Select the **Permissions** button to get the Access Through Share Permissions window. Select **groups** and/or &hp2.users and specify the type of access.
8. Click on **OK** to close dialog box.
9. Select **OK** on the New Share window.
10. Select **Close** on the Shared Directories window.

Note: Be aware that the new shared directory is generally available to all users because it automatically has given Full Control to the group Everyone if you have not defined otherwise. In addition keep in mind that share names must fit the DOS characteristic naming scheme (8+3) so that DOS users can access them.

12.5.3 Using the NT Explorer To Share Resources

To share another directory that does not yet exist, we want to use the NT Explorer in this scenario.

1. Start the **NT Explorer** ([**Start — Programs — Windows NT Explorer**]).
2. From the action bar, select the **File** pull-down menu.
3. Select the item **New** and **Folder** as shown in Figure 185.

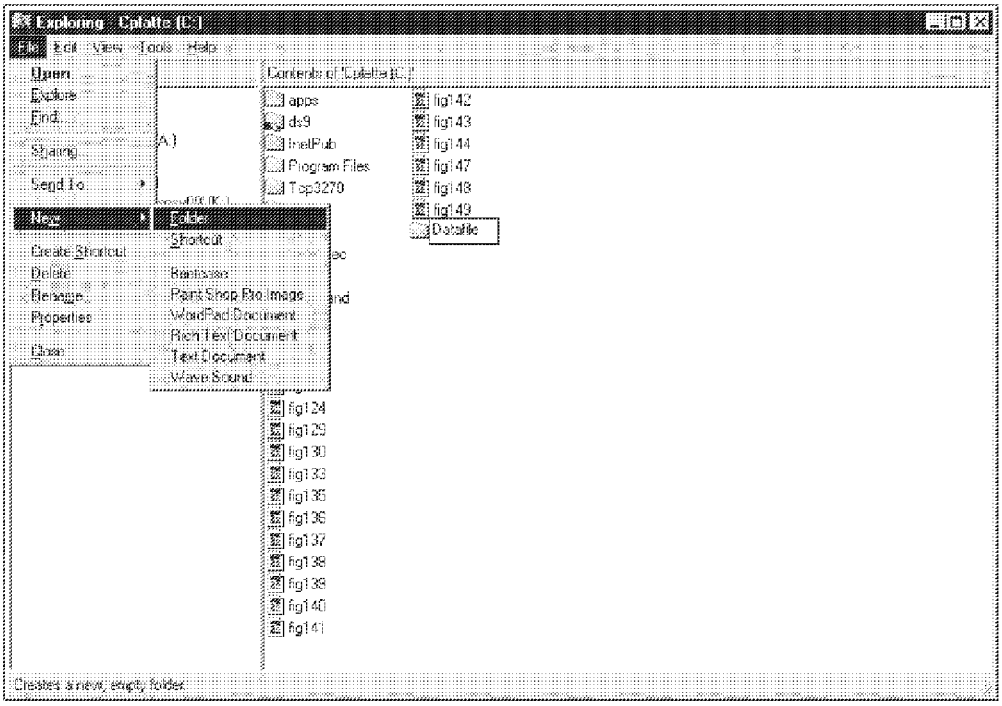

Figure 185. Create Directory for New Share

A new folder with the name "New Folder" has been created for you in the list of Contents.

4. In the list of Contents, overtype the directory name "New Folder" with a new name and click once outside of the folder.
5. Click with the right mouse button on the newly created directory and select **Sharing...**
6. Select the radio button for **Shared As** as shown in Figure 186 on page 364.

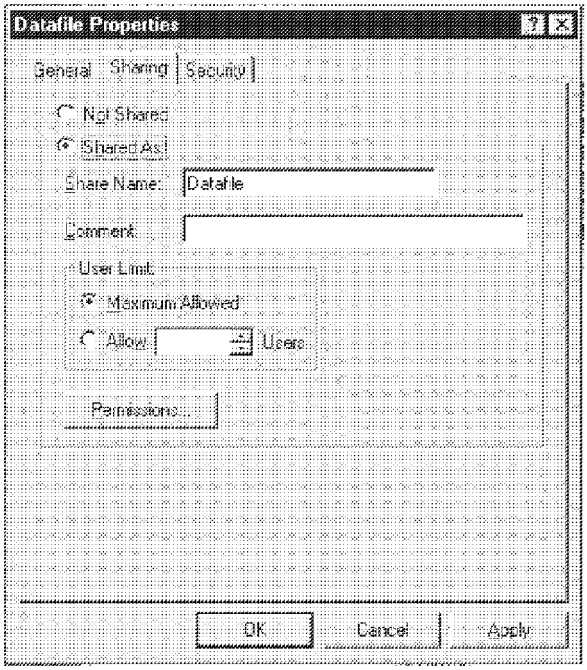

Figure 186. Properties Window of the Newly Created Share

7. Provide a share name, comment, and the User Limit information if needed.
8. Select the **Security** page if you need to provide additional security, auditing, and ownership information.
9. Select **OK** to close dialog box.

12.5.4 Using Commands To Share Resources

Sharing directories by issuing commands at the command prompt can be useful when building up a new server or when you only want to temporarily share a resource. To add a temporary share to the system, issue the following command at a command prompt:

```
NET SHARE [SHARENAME]=[PATH]
```

where the `SHARENAME` parameter stands for the freely selectable name of the share, and the `PATH` parameter identifies the drive and subdirectory to share. The `NET SHARE` command has additional parameters. Table 26 provides more details about parameters of the `NET SHARE` command.

Table 26 (Page 1 of 2). NET SHARE Parameter

Parameter	Description
sharename	Name of the new share

Table 26 (Page 2 of 2). NET SHARE Parameter	
Parameter	**Description**
path	Path and drive of the directory to share
/users	Maximum number of users that can use the share (/user=10 sets the maximum to 10)
/unlimited	Specifies a unlimited number of users for the share
/remark	Specifies a description for the share
/delete	Removes a share definition

The `NET SHARE` command also can be used to change and remove a current share. Also you can, only by using the command without parameters, see a list of the current shares.

12.5.5 Stopping Directory Sharing Using the NT Explorer

Sometimes it is desirable to stop or remove a share. This can be done with the `NET SHARE` command from the command prompt, but it can also be used with the Server Manager and the NT Explorer.

Users who are members of the Administrators and Server Operators group can stop directory shares by doing the following:

1. Start the **NT Explorer**.
2. In the All Folders list, click on the drive where the shared directory is defined.
3. Select the shared directory you want to stop sharing.

   Note: You can recognize a shared directory by its icon. It is an icon that represents a folder with a serving hand.
4. Click with the right mouse button on the shared folder icon.
5. At the directory's Properties window (as shown in Figure 186 on page 364), select **Sharing...** and select the radio button for **Not Shared**.
6. Select **OK** to close the directory's Properties window.

Notice that the icon of the formerly shared directory changes immediately.

12.5.6 Stopping Directory Sharing Using the Server Manager

To stop sharing a directory by using the Server Manager, perform the following steps:

1. Select the **Server Manager**.

2. In the list of Computers, select the server that provides the shared directory resource you want to stop sharing.
3. From the action bar, select **Computer** and then **Shared Directories...** as shown in Figure 187.

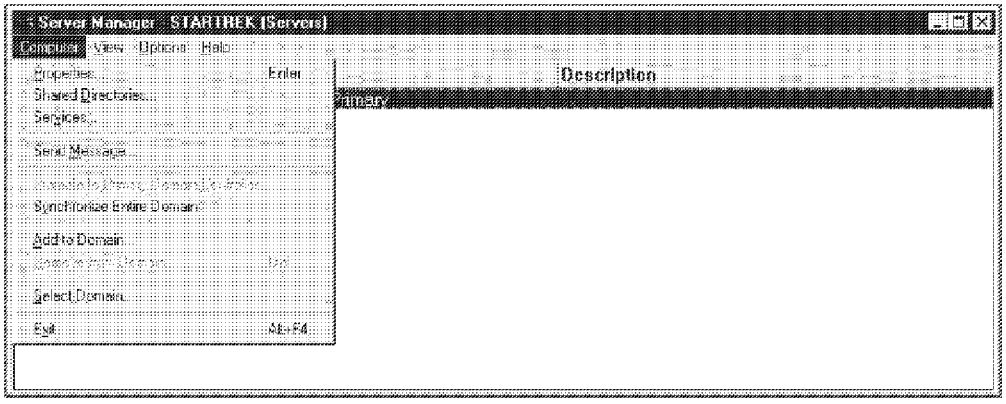

Figure 187. Stop Sharing Directories Using Server Manager

4. At the Shared Directories window, select the directory you want to stop sharing as shown in Figure 188.

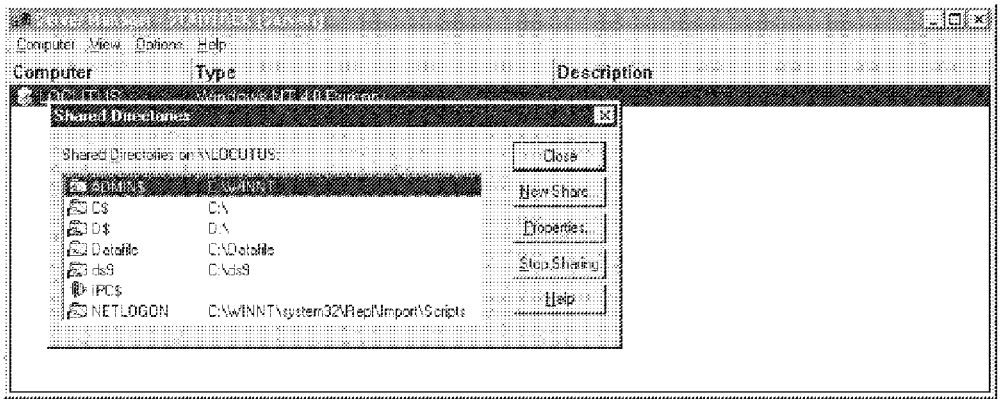

Figure 188. Stop Sharing Directory

5. Select the **Stop Sharing** button.
6. Optionally, you can select additional shares to stop by selecting them one at a time, or close the dialog box by selecting the **Close** button.

12.5.7 Create Printer Shares

After obtaining access to network drives, a user may want to use printers that are connected to Windows NT machines to avoid the necessity of needing another printer. It is important to economize business environments by allowing many users to share expensive printers.

Printer resources are shared with the Printers folder. To share a printer resource, do the following:

1. First, create a printer object.
   a. Open the Printers folder from the Control Panel folder ([**Start — Settings — Printers**]).
   b. Double-click on the **Add Printer** icon. The Add Printer Wizard will be presented to you.
   c. In the Add Printer Wizard window, select the radio button for **My Computer** and select **Next >**.
   d. Check the checkbox of the port you want to configure and select **Next >**.
   e. At the Add Printer Wizard window select the manufacturer of your printer, the type of printer, and then select **Next >**.
   f. Provide a printer name and select **Next >**.
   g. Select the radio button for **Shared** and provide a share name. Also provide information on operating systems that will use the shared printer. When finished, select **Next >**.
   h. At the end you will be prompted whether or not you want to print a test page. Select **Finish** to exit the Printer Wizard.
2. Second, use the printer's properties to specify more options.
   a. In the Printers folder select the printer you need to make ready for printer sharing. Press the right mouse button. Select **Properties**. The printer's Properties window will be presented to you as shown in Figure 189 on page 368.

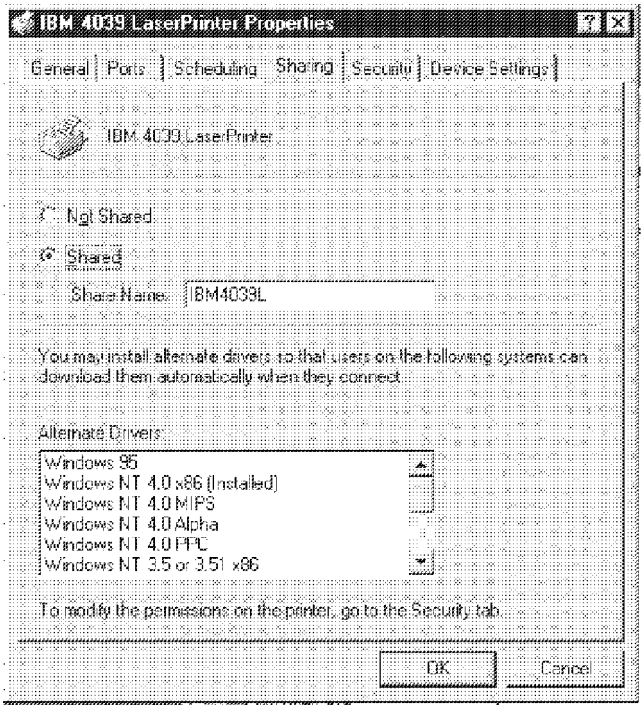

Figure 189. Share a Printer Resource

b. Go to the **Sharing** page to specify whether or not to share the printer.

c. Go to the **Security** page to specify permissions, auditing, and ownership.

d. Once done, select **OK**.

12.5.8 Changing Properties of Directory Shares

After successfully adding shares, it is useful to change the properties of shares. So it is possible to change the previously added items.

The user interfaces for changing the properties are the NT Explorer and Server Manager. The functionality sometimes differs between the two methods.

You have the ability to change the following items for a shared resource:

- Change the shared permissions for the directory.
- Edit the comment box for the share.
- Change the number of users allowed to access the share.
- Change the path of the share.

Only the Server Manager can change all the items, the NT Explorer is only able to change the User Limit, Comment,and Permissions. So we will show how to change Comment and User Limit with the NT Explorer and how to change the Path with the Server Manager.

12.5.8.1 Changing Share Properties with NT Explorer

The NT Explorer cannot change all the described items. So we describe how to change the User Limit and the Comment. We do not want to talk about changing permissions because this is discussed in section 12.6, “Administering / Changing Access Permissions to Resources” on page 371. Keep in mind that you have to be member of the Administrators or Server Operators group to make these changes. To make the required changes perform the following:

1. Start **NT Explorer** ([**Start — Programs — Windows NT Explorer**]).
2. In the list of All Folders select first the drive and then the shared directory, which is symbolized by a folder with a serving hand, and which properties you want to change.
3. Click with the right mouse button and select **Sharing...**.
4. Change the Comment and/or the User Limit as required, and click on **OK** to submit the changes.

12.5.8.2 Changing Share Properties with Server Manager

Sometimes it is necessary to change the path of a share behind the scenes. For example, if you want to put a special function on another disk or you are running out of disk space, you might change the path of the share.

The idea is to change only the path within the share information, but the user still gets his/her defined connections over his/her Logon Script or user profile. Do the following to set up a new drive or path for a shared resource:

1. Start the **Server Manager**.
2. In the list of computers, select the computer where the share resides.
3. From the Computer's pull-down menu, select **Shared Directories...**.
4. Select the shared directory you want to change and select **Properties**.
5. Make necessary changes to User Limit, Comment, and Path. Select **OK** to submit changes.

12.5.9 Change Properties of Printer Shares

You can also change the properties of printer shares, which is what is needed when you move your printer to another building or floor. Also you can change, behind the scene, the printer driver and port. The following changes can be done with the printer's properties:

- Select another Printer Name
- Change the Device Driver for the printer
- Change the description of the printer
- Choose another Printer Port for the printer
- Change the Share Name of the printer
- Change the Location information of the shared printer

To change one or all oft the properties do the following:

1. Open the Printers folder from the Control Panel folder
2. In the Printers folder select the printer you need to make changes to. Press the right mouse button. Select **Properties**. The printer's Properties window will be presented to you.
3. Select the General page as shown in Figure 190.

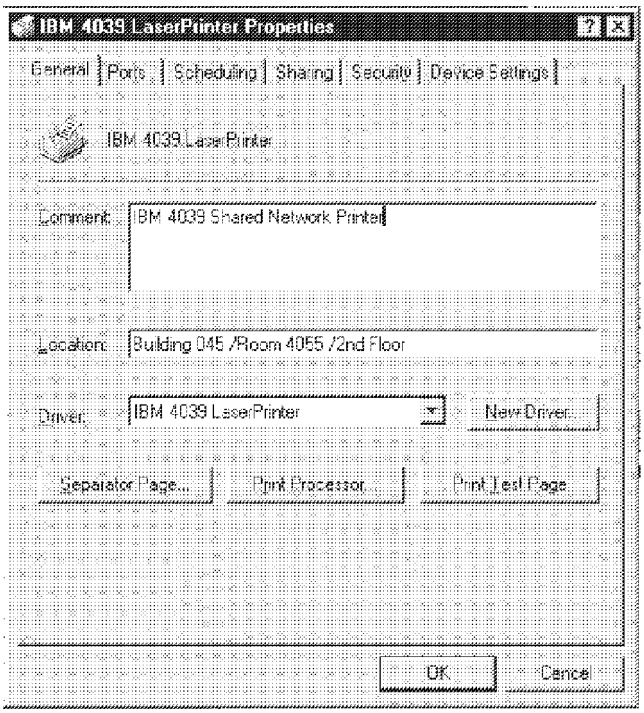

Figure 190. Change required Printer Properties

4. Provide changes as necessary.
5. Select **OK** to update properties.

12.6 Administering / Changing Access Permissions to Resources

The access to resources, like files, directories, and printers, is controlled by permissions to these devices. This means you can give a user different access rights to fit his or her needs.

Windows NT distinguishes between two types of security:

1. Local security, which controls the permissions for all local resources
2. Network security, where the permissions are defined for resources that are accessible through the network

Local security is only available for NTFS; therefore if you need local security for your resources, you have to use NTFS.

However, this restriction does not apply to network resources. For network resources, you will be able to define rights and permissions for files and directories on all NT-supported file systems (NTFS, FAT).

Windows NT uses the Access Control List (ACL) to manage permissions. In contrast to FAT, within NTFS partitions, the Access Control List information is directly attached to the different types of resources (files, directories, and printers). The ACL will be created as soon as the first permission is defined. If this permission information is changed or added, an Access Control Entry (ACE) is added to the Access Control List for the object. ACEs indentify the type of permission as well as user or group ID assigned to the permissions.

When users try to access resources, Windows NT reads the ACL for the object to determine which access rights the user has. The **deny access** permission is the only permission that overrides all other permissions on resource objects. For example, if a user has inherited certain permissions from a group membership and his or her user ID received additional or different rights, the permissions will be accumulated, giving the user the total of his/her individual and group permissions. This does not apply to **deny access**. If this permission is set, it overrides all others regardless of whether they are granted to the group or to the user personally.

12.6.1 Adding Permissions for Network Resources

Use the Windows NT Explorer to add, change, or delete permissions to network resources. Be aware that you can add permission to users and

groups. Especially when setting permissions to groups, you can set local permissions to local groups and global permissions to global groups. To add new permission to a subdirectory, proceed as follows:

1. Start the **Server Manager**.
2. In the Computer list, select the machine you want to control permissions for.
3. From the action bar, select **Computer**. From the pull-down menu, select **Shared Directories...** The Shared Directories window will be opened for you.
4. From the Shared Directories window, select the network resource you want to add permissions to and click on **Properties...**. The Share Properties window will be opened for you.
5. From the Share Properties window select the **Permissions...** button; it will open the Access Through Share Permissions window as shown in Figure 191.

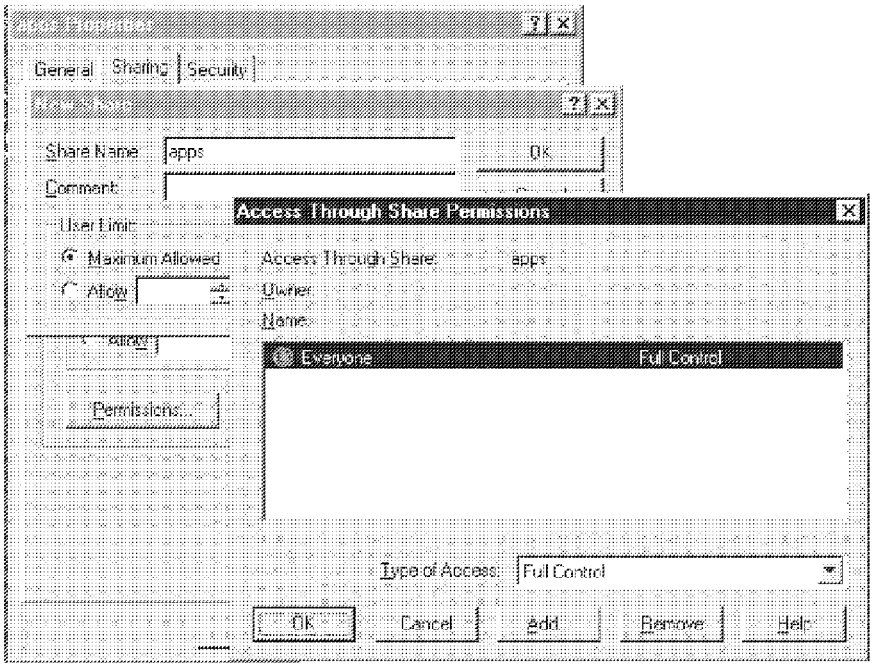

Figure 191. Access Through Share Permissions Window

6. To add a new access permission for a user or a group, click on the **Add...** button. This will open Add Users and Groups window for you. The Names list, by default, will show you only groups. If you also need to display the defined users, you have to click on the **Show Users** button. The displayed window is shown in Figure 192 on page 373.

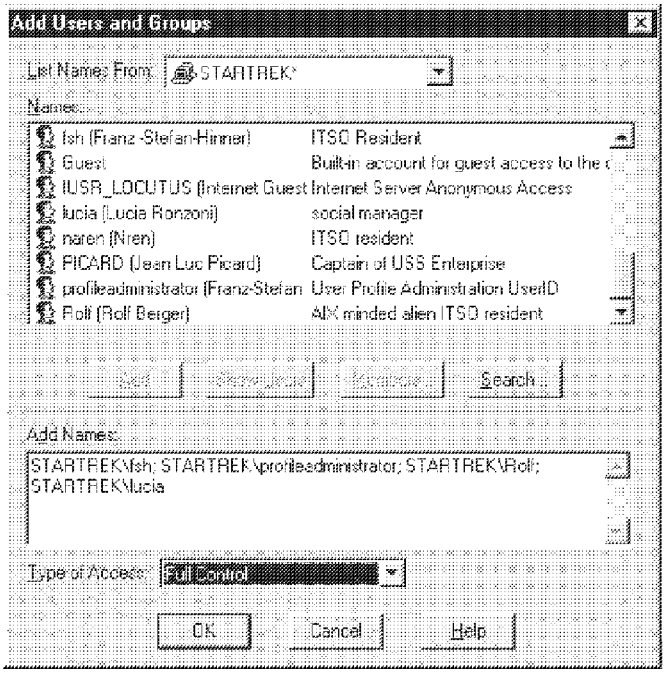

Figure 192. Add Users and Groups Dialog

7. Select the groups and users you want define access permissions to, and press the **Add** button to add them to the list of Add Names.
8. Finally choose the appropriate **Type of Access** and press **OK** (see Table 27 for a description of selectable rules).
9. Exit all windows by pressing **OK**.
10. Select **Close** at the Shared Directories window to exit.

In the following table, we describe the different types of accesses you can select in the Add Users and Groups window.

Table 27. Different Type of Access

Type of Access	**Description**
No Access	The user cannot access a resource at all. This is even the case if he/she is member of a group granted with further access permissions to that specific resource.
Read	This enables the user's read and executes the files if they are executable.
Change	Permits the user to read, write, and delete files.
Full Control	Users with with full control can read, change, delete and own files. Also they have the possibility to give away access rights.

Another enhanced way to add access permissions is handled with the Windows NT Explorer. With the NT Explorer, you can add additional access rights especially designed for directories. This is helpful because you have more details for distinguishing the different access rights of users. Be aware that changing the rights for directories and files in the NT Explorer also effects the local access rights for the user on a NTFS drive.

This means if you set up for the group Everyone No Access to a specific resource, no member of this group, including the administrator, will have access to this resource anymore. To add access rights to a directory or file, do the following:

1. Start the **NT Explorer**.
2. Select the directory you want to add or change permissions on, press the right mouse button, and select **Properties** from the directory folder's pop-up window.
3. At the Properties window, select the **Sharing** page to enter the required information like Share Name and Comment.
4. Click on **Permissions...** and proceed as described in 12.6.1, “Adding Permissions for Network Resources” on page 371. Figure 193 shows you all the windows opened in the preceding steps.

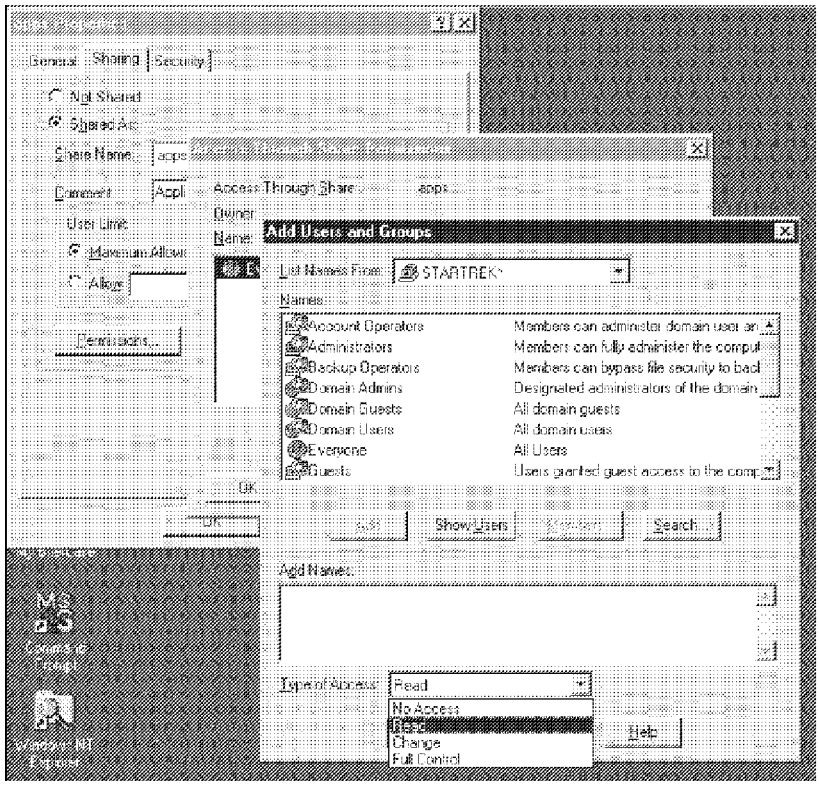

Figure 193. Adding New Access Permissions with the Windows NT Explorer

As mentioned before, an NTFS drive has an additional type of security, local security. In this case the Share Properties window contains an additional page: Security.

Besides ownership and auditing information, you can define the access permissions for the local security. In Table 28 you will see a detailed overview of the additional File/Directory permissions applicable only for local security on NTFS drives:

.

Table 28. Additional Access Permissions for Local Security

Permission	**Description**
List	This allows the user to list files and directories and change to a subdirectory. The user does not have any access to newly created files and directories.
Add & Read	This permission is a combination of the read and the add permission. Users can create new files, read files, and execute program files.
Special Directory Access	With the special directory access permission, the administrator can set up which permissions he/she wants to give groups or users. He/she can selectively give the permissions to read, write, execute, delete, change permissions, and take ownership.
Special File Access	This special function is the same as for directories, but is used only for files. So the administrator can define permissions on a file-by-file basis. The rights are read, write, execute, delete, change permissions, and take the ownership of a file.

Notice on Access Permissions

Please keep in mind that these rights are also used locally. This means if you remove all rights on a shared resource for the administrators, the administrator is no longer able to share the resource.

Adding permissions is also applicable to printer resources. To do this, you have to go through the Printers folder settings to add new printer access permissions, do the following:

1. By clicking on a printer you want to share with the right mouse button, select **Properties**.

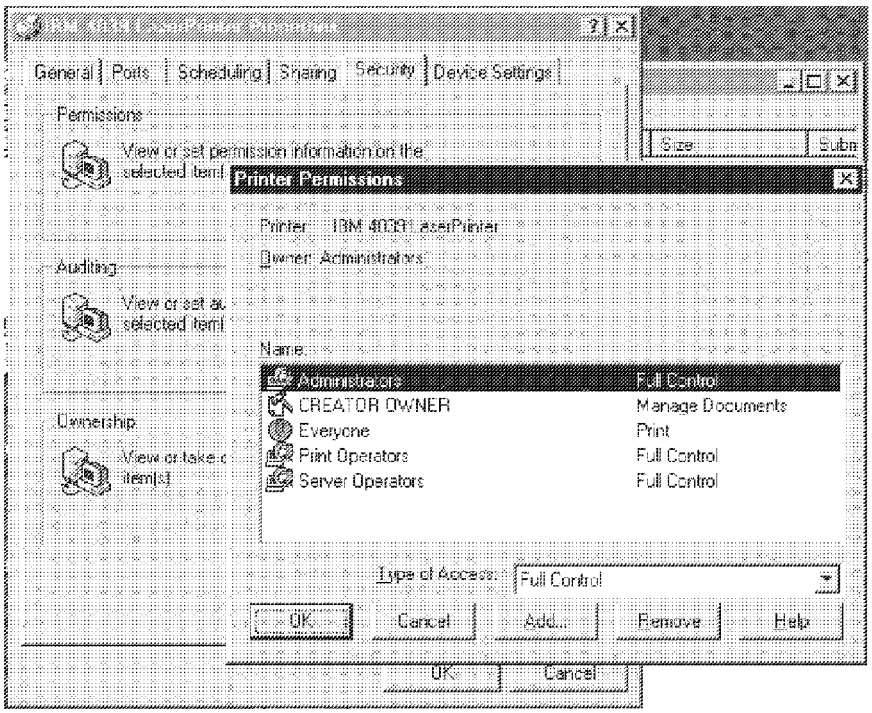

information.

Figure 194. Adding Users and Groups Printer Permissions

2. Go to the Security page and click on **Permissions** and provide information as required (see Table 29 for a description of printer access rights).
3. Select the **Sharing** page, and enter the required information, and proceed as you would when sharing files and directories as described earlier.

.

Table 29 (Page 1 of 2). Access Permissions for a Printer Resource

Permission	**Description**
No Access	Allows no access to the printer.
Print	User can print documents.
Manage Documents	The user can print documents and change print settings for documents. Additionally he/she is able to pause, resume, restart, and delete documents.

Table 29 (Page 2 of 2). Access Permissions for a Printer Resource	
Permission	**Description**
Full Control	User can change the print settings for a document. He/she has the right to print, pause, resume, restart, and delete documents, and he/she is also able to change the order of the documents in the queue. Furthermore, the user has the rights to modify printer settings, delete the printer, change printer permissions, and to pause, resume, and purge the printer.

12.6.2 Changing Permissions of Network Resources

To change access rights, do the following:

1. Start the **Server Manager**.
2. In the Computers list, select the computer where the network resource resides.
3. From the Computer pull-down menu, select **Shared Directories...**.
4. Select the shared directory you want to changes to and select **Properties**.
5. At the Shared Properties window select **Permissions...** Select a name and change the type of access as shown in Figure 195.

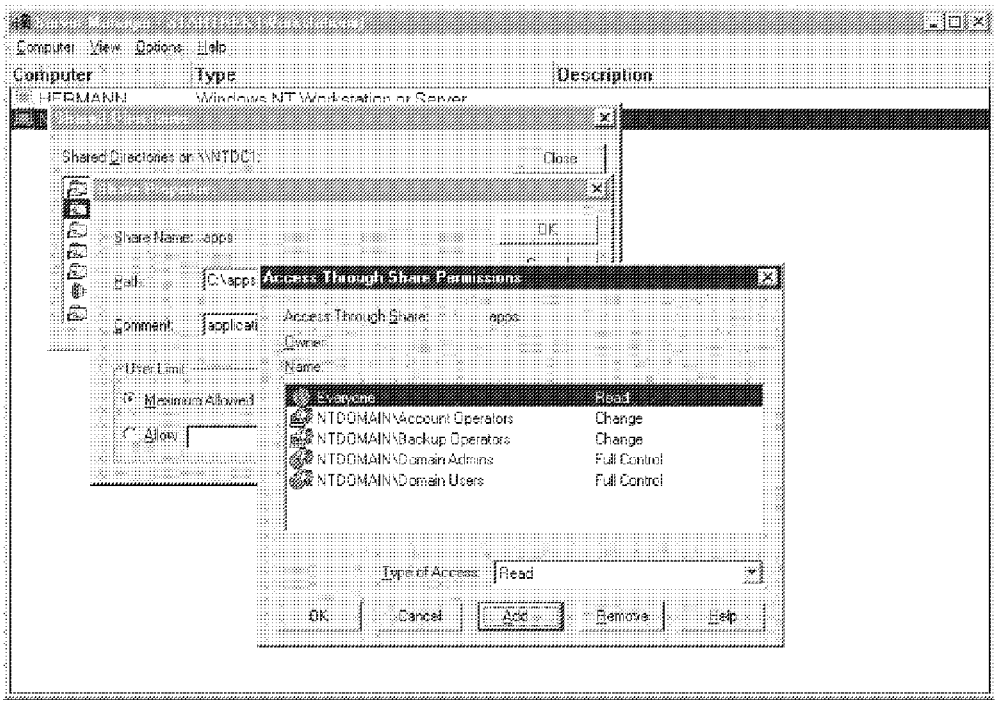

Figure 195. Change the Access Rights for Groups or Users

6. Select **OK** to submit changes and exit all open windows.

12.6.3 Removing Permissions of Network Resources

To remove access permissions proceed with steps one to five as described in 12.6.2, “Changing Permissions of Network Resources” on page 377. In the Access Through Share Permissions window, select a name you want to remove access rights from and select **Remove** to delete his/her access rights to the selected resource.

12.7 Dynamic TCP/IP in Windows NT Server

Windows NT 4.0 Server comes with a DHCP Server, a server that offers IP addresses out of a defined IP address pool to requesting IP clients, such as Windows NT Workstation and Windows 95.

The second server that comes with Windows NT 4.0 Server is a NetBIOS Name Server called WINS (Windows Internet Name Service; by the way, although the name Internet is part of the name of WINS, it has nothing to do with the Internet) with an integrated static DNS server. In comparison to the previous version of Windows NT, NT 4.0 now comes at least with a static DNS server. However, WINS is the "big brother" of Microsoft's DNS server which means, in order to work with DNS, you must work with WINS.

WINS sets pointers to the DNS database, which is still not a good solution in comparison to Warp Server's Dynamic DNS server.

In the following section we demonstrate how to set up DHCP services under Windows NT 4.0 Server. WINS is described in 12.7.2, “Windows Internet Name Service (WINS)” on page 387.

12.7.1 Configuring and Using DHCP Server

1. Using the Windows NT navigation system, follow the path [**Start — Programs — Administrative Tools (Common) — DHCP Manager**] as shown in Figure 196 on page 379.

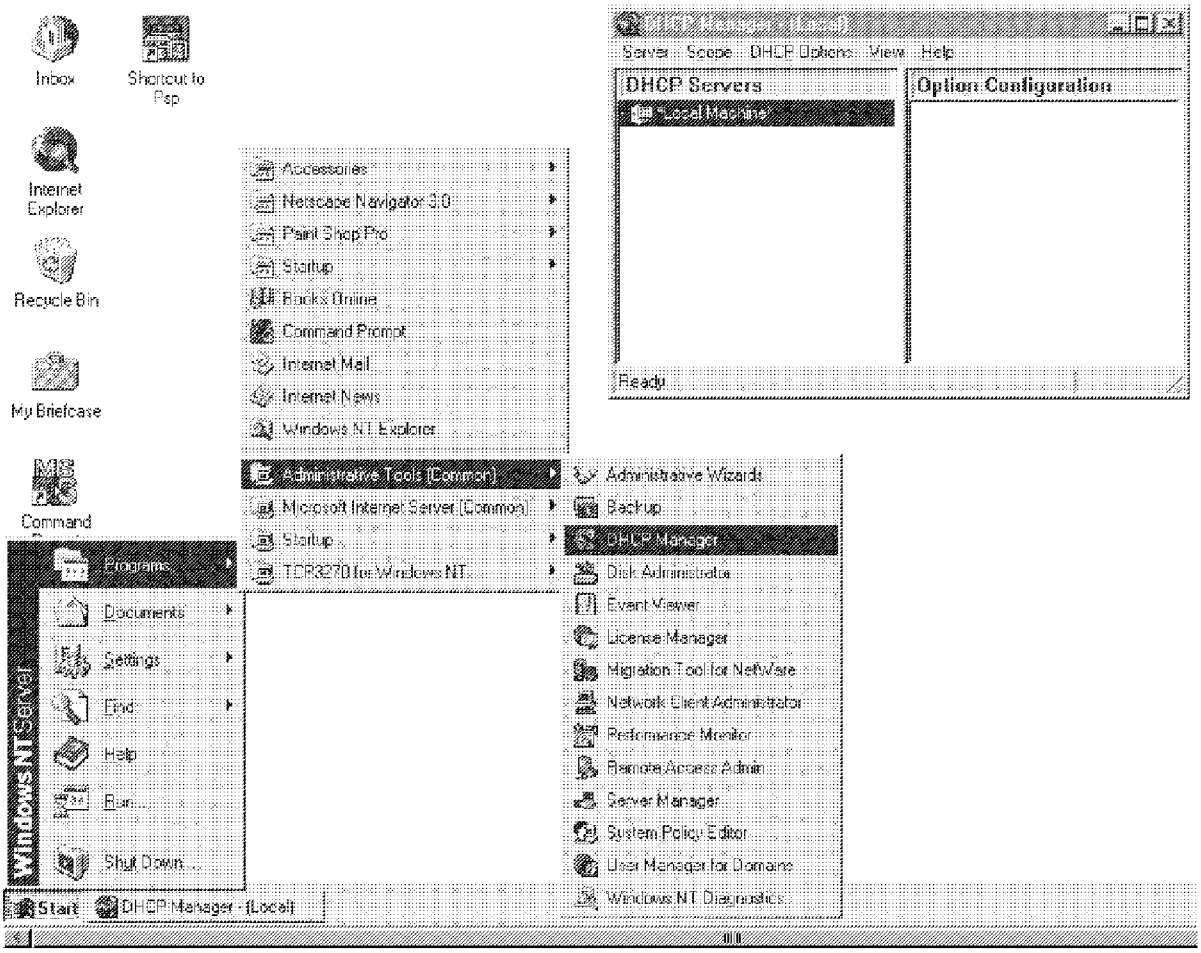

Figure 196. Navigation Path to the DHCP Manager

This example gives you instructions on how to set up a DHCP Server for the subnet 9.67.20.0. You can extend the configuration for other subnets by doing the steps shown here for a single subnet. At the end of this DHCP Server section, we will have the server defined as shown in Figure 197.

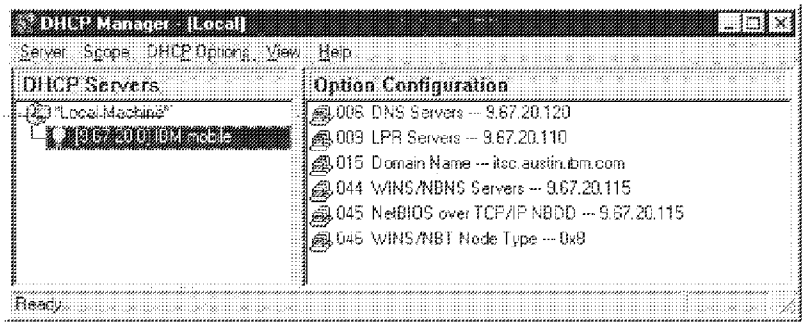

Figure 197. MS DHCP Manager Window

DHCP Manager Configuration to work with Shadow

To make configuring comparable to the method used in Warp Server, we use the same IP addresses. That also means that we configure this DHCP Server to work with the NetBIOS Name Server of NTS, called Shadow, since this is the only NetBIOS Name Server that supports Datagram Distribution. If you want to configure WINS as your NetBIOS Name Server always replace the IP address mentioned in our example by the Windows NT Server's local IP address (9.67.20.120). The same rule applies to DNS server configuration.

2. To begin configuring the DHCP Server, select **Scope** from the DHCP Manager's action bar.
3. Select **Create...** from the Scope pull-down menu. The Create Scope window will be opened for you as shown in Figure 198.

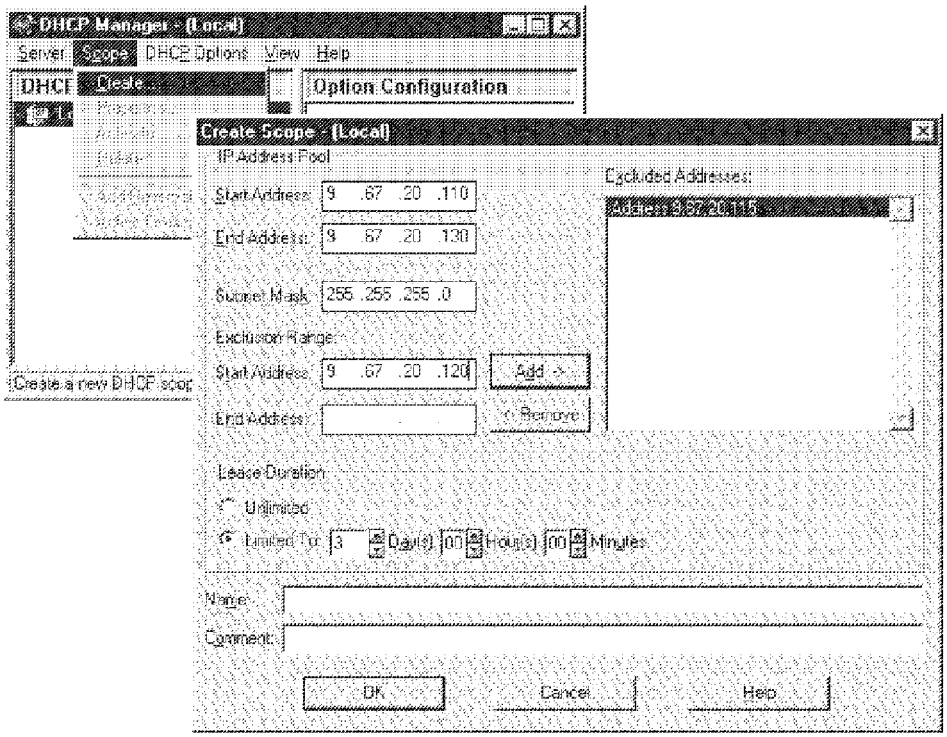

Figure 198. MS DHCP Manager Create Scope Window

Provide the following information:

- IP Address Pool
  - Start Address: 9.67.20.110
  - End Address: 9.67.20.130

- Subnet Mask: 255.255.255.0

• Exclusion Range

- Start Address: Enter 9.67.20.115 and select **Add ->**
- Start Address: Enter 9.67.20.120 and select **Add ->**

• Lease Duration

- Using the up and down arrow buttons, specify a lease duration of 1 day.

• Provide a name, for example, IBM mobile, and then select **OK**.

4. At the DHCP Manager pop-up information window as shown in Figure 199, select **OK** to have the newly created scope activated.

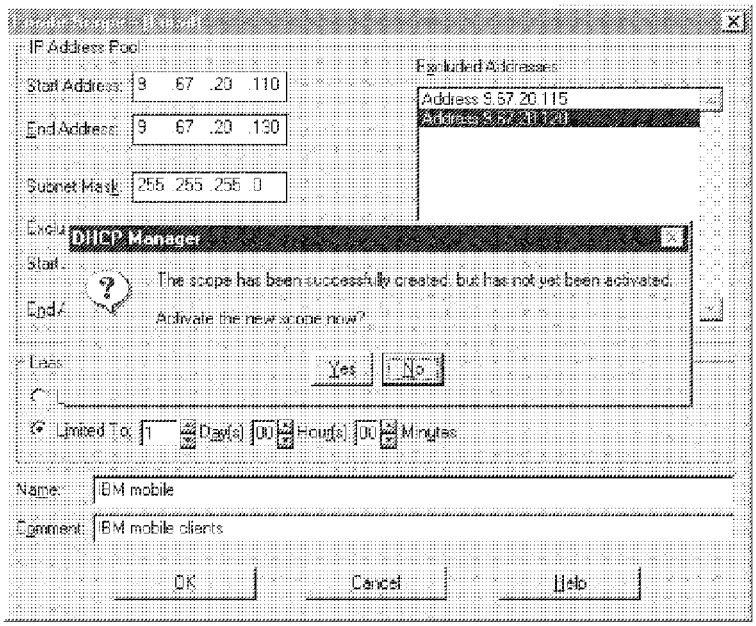

Figure 199. DHCP Manager Pop-Up Information Window

5. In the DHCP Manager's list of DHCP Servers, select the new scope. Then select **Scope...** from the **DHCP Options** pull-down menu. The DHCP Options: Scope window will be opened for you.
6. From the list of Unused Options, select **006 DNS Servers** and click on **Add->**. Then select **Edit Array...** and provide IP address information about the DNS server (in our example it is 9.67.20.120) as shown in Figure 200 on page 382.

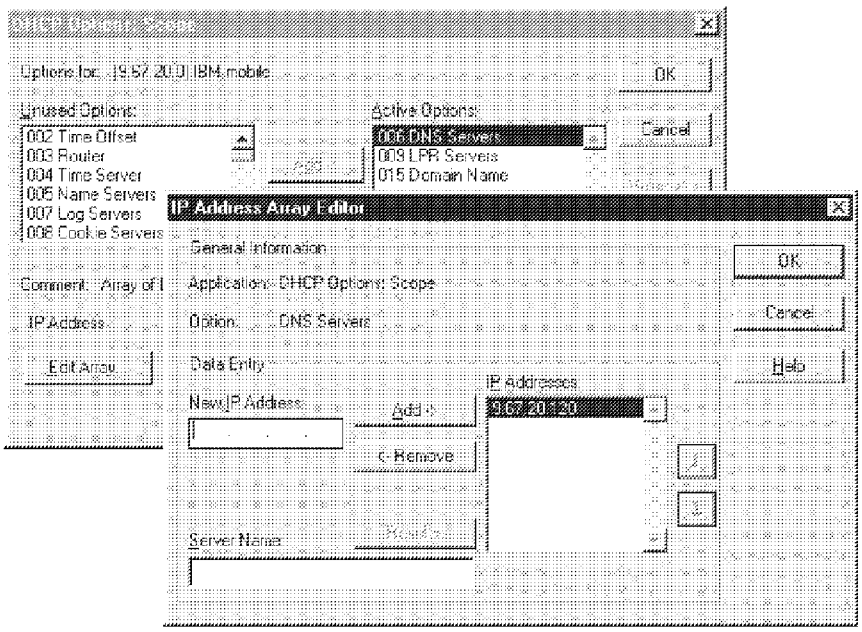

Figure 200. IP Address Array Editor For DNS Servers Window

7. From the list of Unused Options, select **015 Domain Name** and click on **Add->** and provide domain name information (in our example it is itsc.austin.ibm.com) as shown in Figure 201.

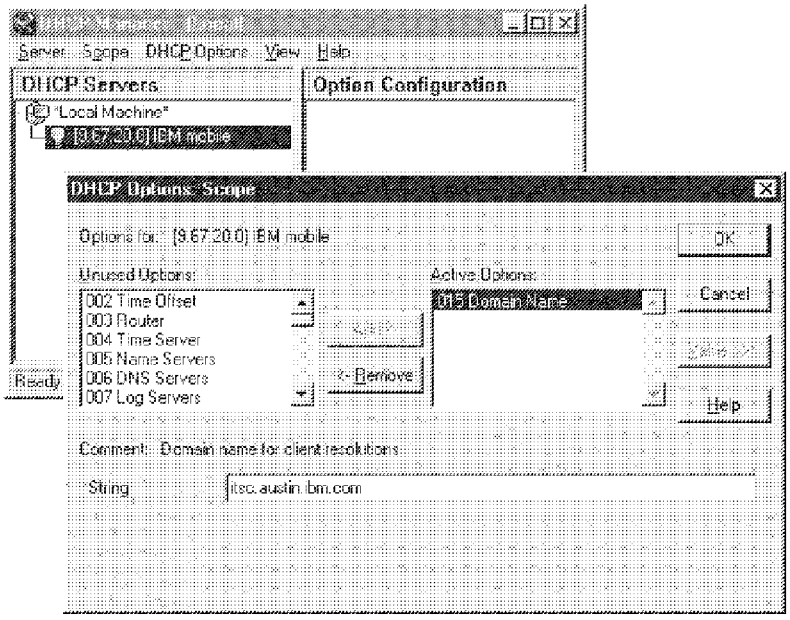

Figure 201. Providing Domain Name Information

8. From the list of Unused Options, select **009 LPR Servers** and click on **Add->**. Then select **Edit Array...** and provide IP address information

about the LPR server (in our example it is 9.67.20.110) as shown in Figure 202 on page 383.

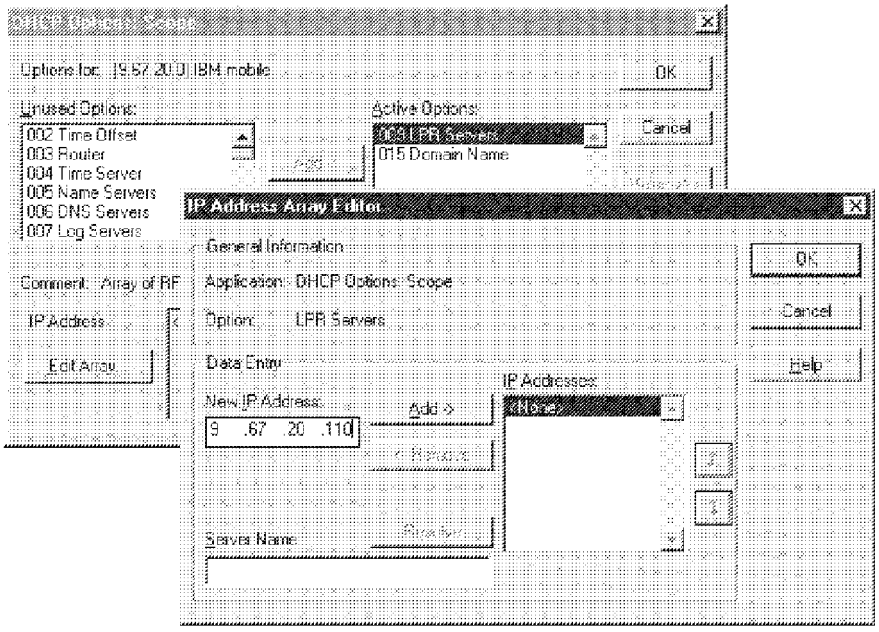

Figure 202. IP Address Array Editor For LPR Servers Window

9. From the list of Unused Options, select **044 WINS/NBNS Servers** and click on **Add->**. Then select **Edit Array...** and provide IP address information about the WINS/NBNS server (in our example it is 9.67.20.115) as shown in Figure 203 on page 384.

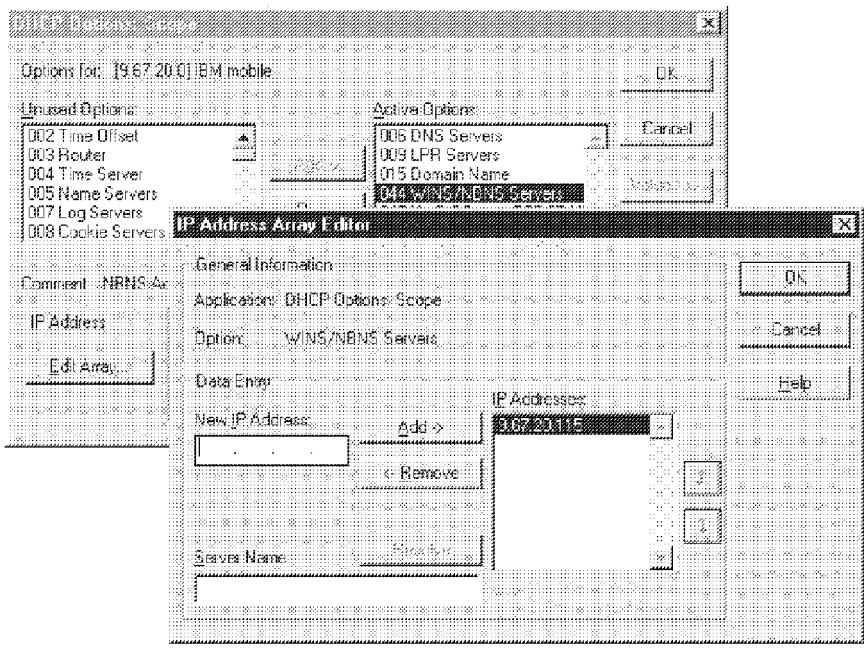

Figure 203. IP Address Array Editor For WINS/NBNS Servers Window

Note: If you want WINS to be your NetBIOS Name Server, choose IP address **9.67.20.120** instead.

10. From the list of Unused Options, select **045 NetBIOS over TCP/IP NBDD** and click on **Add->**. Then select **Edit Array...** and provide IP address information about the Datagram Distribution server (in our example it is 9.67.20.115) as shown in Figure 204 on page 385.

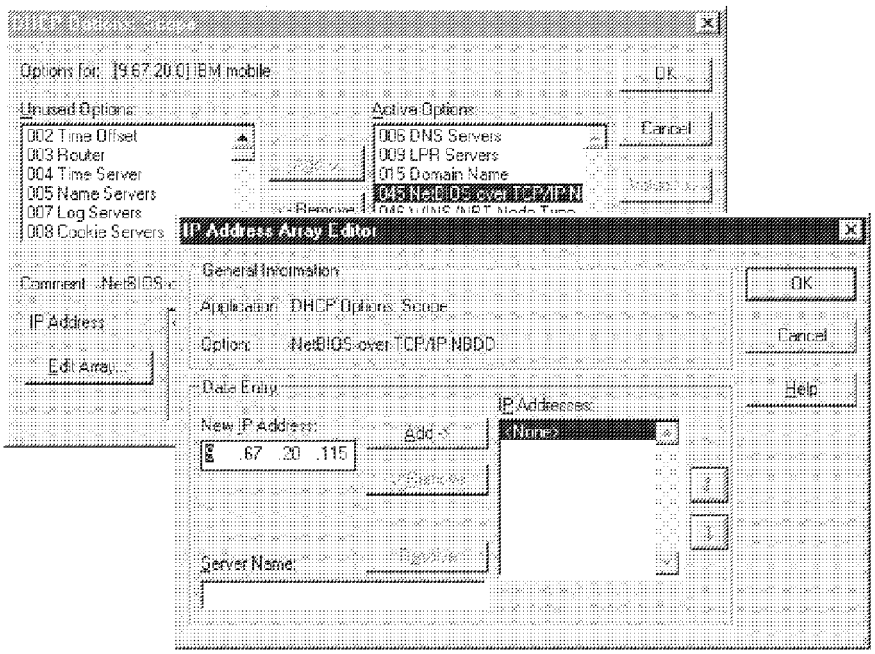

Figure 204. IP Address Array Editor For NetBIOS Over TCP/IP NBDD Servers Window

Note: In case you use WINS, you not not have to select option 045 since WINS does not support Datagram Distribution.

11. From the list of Unused Options, select **046 WINS/NBT Node Type** and click on **Add->** and provide h-node type information (0x8) as shown in Figure 205.

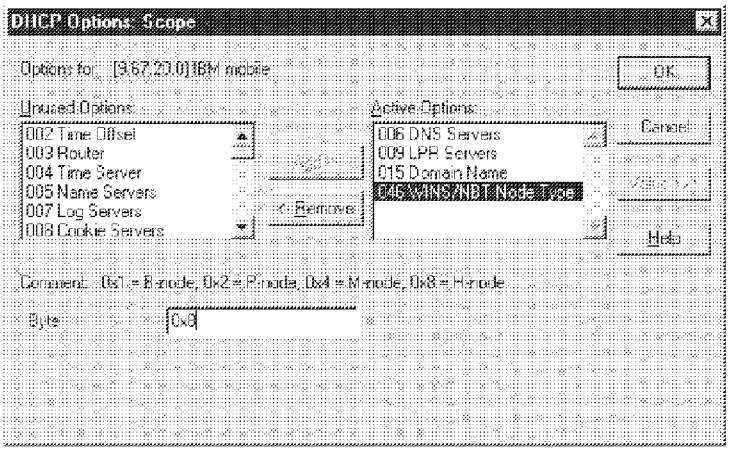

Figure 205. WINS/NBT Node Type Information Window

12. Select **OK** at the DHCP Options window to make the the configuration active.

To check, which dynamic IP clients obtained which IP address, double-click on the previously created scope in the list of DHCP Servers. The Active Leases window will be presented to you as shown in Figure 206 on page 386.

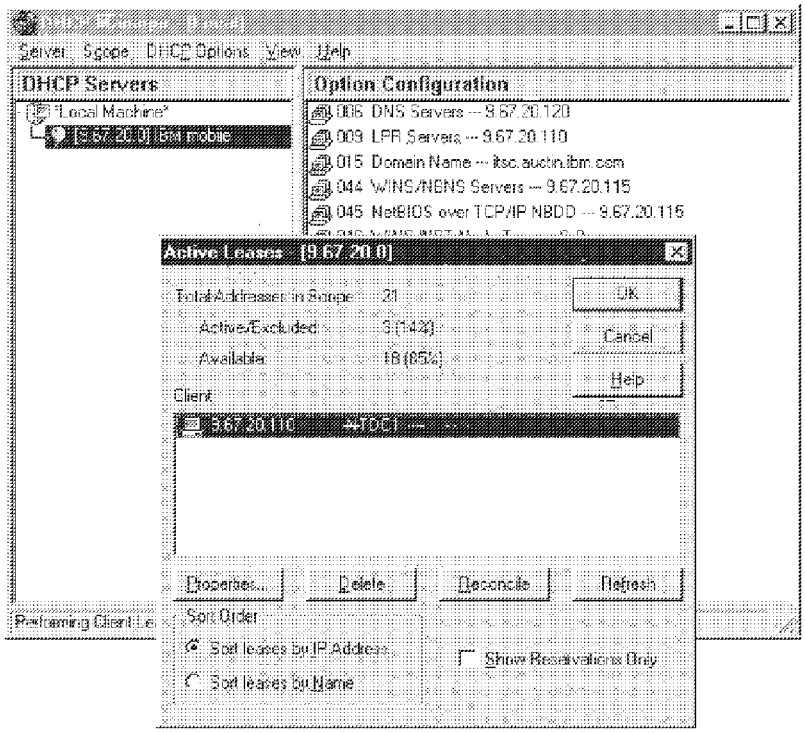

Figure 206. Active Leases Window

In this example you see that a client with the computer name of NTDC1 obtained the IP address of 9.67.20.110. Since we did not configure WINS as the NetBIOS Name Server, the computer name is not made equal to the host name.

To enable a Windows NT Workstation for obtaining dynamic TCP/IP configuration, you need to navigate to the Control Panel window ([**Start — Settings — Control Panel**]). From here, you double-click on the **Network** icon.

1. In the Network window, select the **Protocols** page.
2. Make sure the TCP/IP protocol is in the list of Network Protocols.
3. In the list of Network Protocols, double-click on the **TCP/IP Protocol**.
4. Check the radio button for **Obtain an IP address from a DHCP Server** and select **OK** on the subsequent windows to exit.

12.7.2 Windows Internet Name Service (WINS)

WINS provides a distributed database for registering and querying dynamic NetBIOS names to IP address mapping in a routed network environment. The LMHOSTS file addressed only one disadvantage of broadcast-based systems: It allowed resolution of names across routers.

Since the system itself was still broadcast-based, the problems of broadcast traffic and load on local nodes were not solved. RFCs 1001/1002 address these problems. They define a protocol that allows name registration and resolution through unicast datagrams to NetBIOS Name Servers (NBNS). Because unicast datagrams are used, the system inherently works across routers. This eliminates the need for an LMHOSTS file, restoring the dynamic nature of NetBIOS name resolution.

This, in turn, allows the system to work seamlessly with DHCP. For example, when dynamic addressing through DHCP results in new IP addresses for computers that move between subnets, the changes are automatically updated in the WINS database. Neither the user nor the network administrator needs to make manual accommodations for name resolution in such a case.

The WINS protocol is based on, and is compatible with, the protocols defined for NBNS in RFCs 1001/1002; so it is interoperable with other implementations of these RFCs.

Note: To be honest, WINS only implemented a subset of what is defined in RFC 1001 and 1002 and made proprietary extensions to WINS.

Another RFC-compliant implementation of the client can talk to the WINS server, and similarly, a Microsoft TCP/IP client can talk to other implementations of the NBNS server. However, because the WINS server-to-server replication protocol is not specified in the standard, the WINS server will not interoperate with other implementations of a NetBIOS Name Server. Data will not be replicated between the WINS Server and the non-WINS NBNS. Therefore, the WINS system as a whole will not converge, and name resolution will not be guaranteed.

WINS consists of two main components, the WINS server and WINS client.

12.7.2.1 WINS Server

The WINS Server does the following things:

- Handles name registration/releases requests from WINS clients and registers/releases their names and IP addresses

- Responds to name queries from WINS clients by returning the IP address of the name being queried (assuming the name is registered with the WINS Server)
- Replicates the WINS database with other WINS Servers

12.7.2.2 WINS Client

The WINS client does the following things:

- Registers/releases its name with the WINS Server when it joins/leaves the network
- Queries the WINS Server for remote name resolution

12.7.2.3 Benefits of Using WINS

WINS has the following benefits:

- Dynamic database maintenance to support computer name registration and resolution.
- Centralized management of NetBIOS name database.
- Reduction of IP broadcast traffic in the internetwork, while allowing the clients to locate remote systems easily across local or wide area networks.
- The ability for the clients (Windows NT 3.5 (or newer), Windows for Workgroups 3.11, and Windows 95) on a Windows NT Server-based network to browse remote domains without a local domain controller being present on the other side of the router.
- On a Windows NT network, the ability to browse transparently across routers (for domains that span multiple subnets). To allow browsing without WINS, the network administrator must ensure that the users primary domain has Windows NT Server or Windows NT Workstation computers on both sides of the router to act as master browsers. These computers need correctly configured LMHOSTS files with entries for the domain controllers across the subnet.

12.7.3 WINS/DNS Integration

In Windows NT 4.0, Microsoft's implementation of DNS is tightly integrated with WINS. This allows non-WINS clients to resolve NetBIOS names by querying a DNS server. Administrators can now remove any static entries for Microsoft-based clients in legacy DNS server zone files in favor of the dynamic WINS/DNS integration.

Note: Do not be mistaken. DNS itself is still a static Domain Name System server and has nothing in common with IBM's Dynamic Domain Name System server. There are pointers from the WINS database to

DNS. This trick gives the impression of having a dynamic DNS built into Windows NT.

For example, if a non-Microsoft-based client wants to get to a Web page from an Worldwide Web server that is DHCP/WINS enabled, the client can query the DNS server; the DNS server can query WINS, and the name can be resolved and returned to the client. Prior to the WINS integration, there was no way to reliably resolve the name because of the dynamic IP addressing.

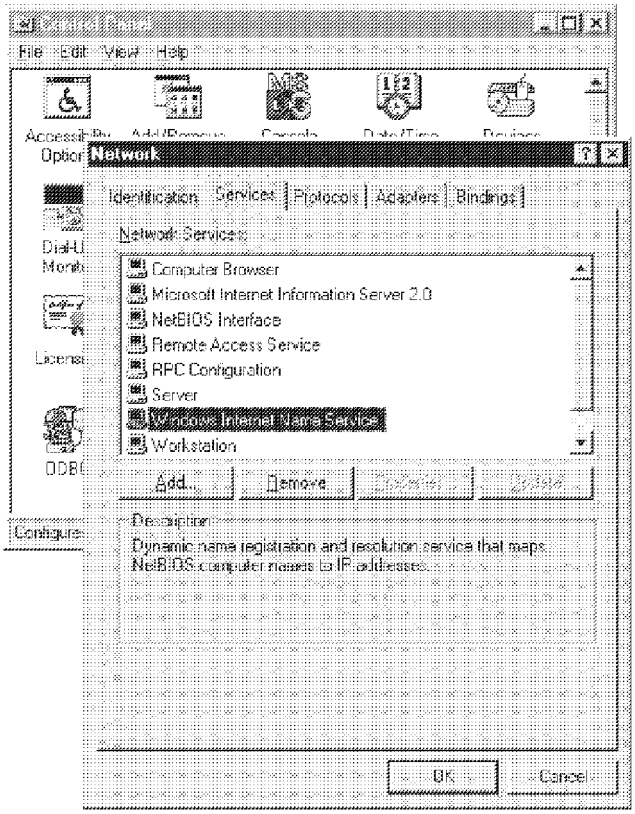

Figure 207. Selection of the WINS Service

12.8 Systems Management with the Systems Management Server (SMS)

The Systems Management Server allows a central control of the client computers in the entire organization. The purpose of SMS is to provide administration and maintenance to the LAN and WAN environments. SMS is a separate product and has to be purchased in addition to Windows NT, and because SQL Server is a prerequisite that has to be loaded, this package has to be purchased as well. Therefore, three different software packages need to be purchased, installed, and configured. After a successful installation of SMS, you have the following functionality:

Software Management: You can distribute, install, and configure software on the network server and clients. When installing software, you can choose to just let the user interact with the installation's Setup program, or you can write a script to automate the installation. You can also exercise control over several versions of the same software installed on different computers on the network.

Inventory Management: You can remotely inventory all hardware and software on your network servers and clients. The inventory is maintained in a SQL Server database in a central location.

Remote Administration: You can remotely administer any client computer, subject to client permissions, on the network. The security is controlled at the client rather than at the administering server. This allows the user to be in complete control of what functions an administrator can perform and when to allow access to a client workstation. To enable remote control, the client must activate the remote control agent located in the folder of the SMS client folder on the desktop. To enable one or more functions, the client must select the **Helpdesk** icon, and from here he can enable or disable helpdesk functions. You can view the display (or capture the display, keyboard, and mouse to take remote control) of MS-DOS, Windows 3.x, Windows 95, and Windows NT computers only.

Alerts: You can create alerts based on a conditional statement for servers and clients. To create a new alert, the administrator first has to create a query. This query is then used to periodically query the SMS database. When the condition is met, SMS can perform the following actions:

- Log an SMS event in the event log
- Execute a command line from the system path where the alerter is installed
- Send a message to a computer or user on a local network

Network Protocol Diagnostics: You can use the SMS Network Monitor to locally (where you execute the SMS Network Monitor) capture network packets to diagnose numerous network-related problems. You can even remotely (where the Network Monitor Agent software resides) capture network packets to reduce the network bandwidth load on a different network segment.

Customization: Microsoft has published application programming interfaces (APIs) and provides a software development kit so that developers can extend the SMS functionality. You also can use custom frontends to access the SMS database for informational purposes or write custom applications to interact with the SMS database.

Server Management: Windows NT Server Performance Monitor provides information for problem detection and performance optimization. NT Server has been designed to be self-tuning. The system monitors several counters and adjusts related parameters to provide better performance. You use the Performance Monitor to monitor a system in real time. This means the event objects you are monitoring are occurring right now, and the value that you see for the event object reflects the actual value with a minimal time lag. The Performance Monitor covers a wide range of system behavior, some of which may be used daily and some of which could be used when trying to isolate problems. NT provides four different views and provides you with the data in these different formats: Chart, Log, Report and Alert. The Performance Monitor tracks many server objects, including:

- Cache

  Monitors the utilization of the disk cache.

- Logical Disk

  Monitors a logical drive to which a letter has been assigned.

- Memory

  Monitors memory usage and virtual memory operation in the system.

- Physical Disk

  Monitors separate physical disks that contain logical disks.

- Process

  Monitors processes within the system.

- Processor

  Monitors activity on each processor on the system.

- System

  Monitors all the processors on the system as a group.

- Thread

  Monitors threads within the system.

- Paging File

  Monitors paging file activity, such as paging usage and page peak activities.

Examples of Windows NT's Performance Monitor are shown in Figure 208 on page 392 and Figure 209 on page 393.

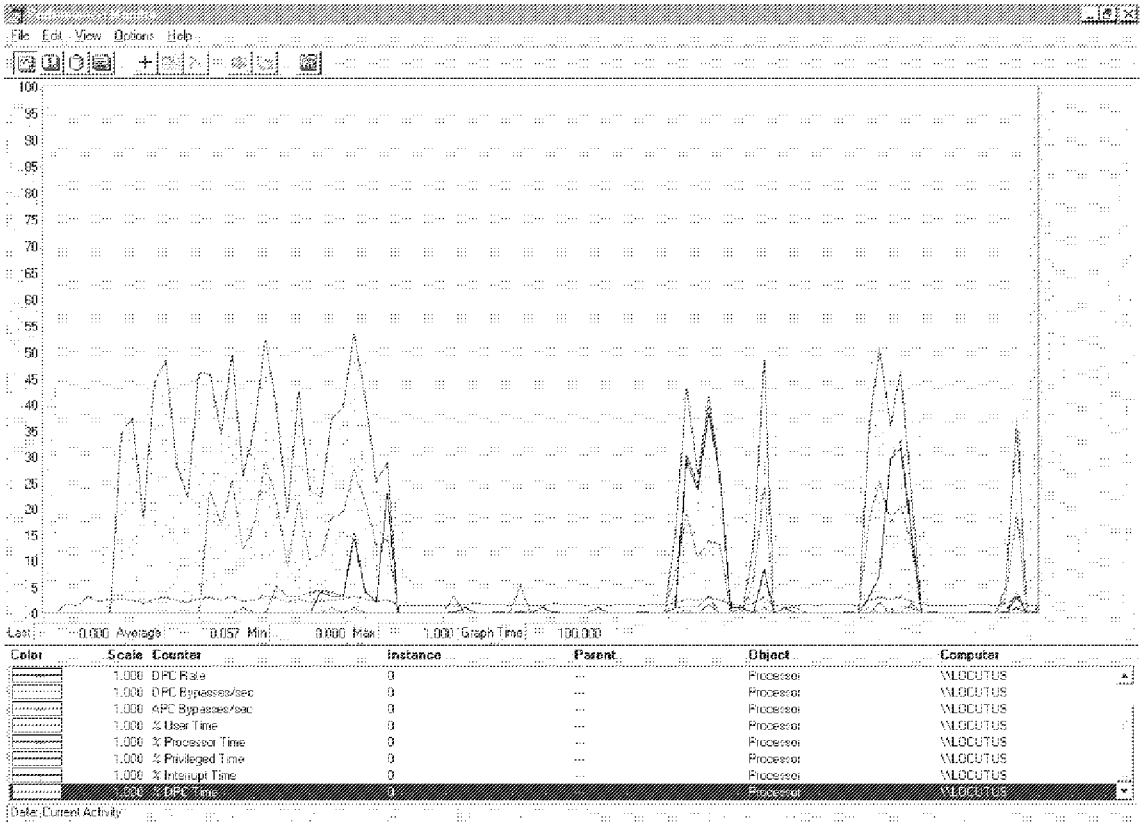

Figure 208. Performance Monitor with a Few Objects Selected

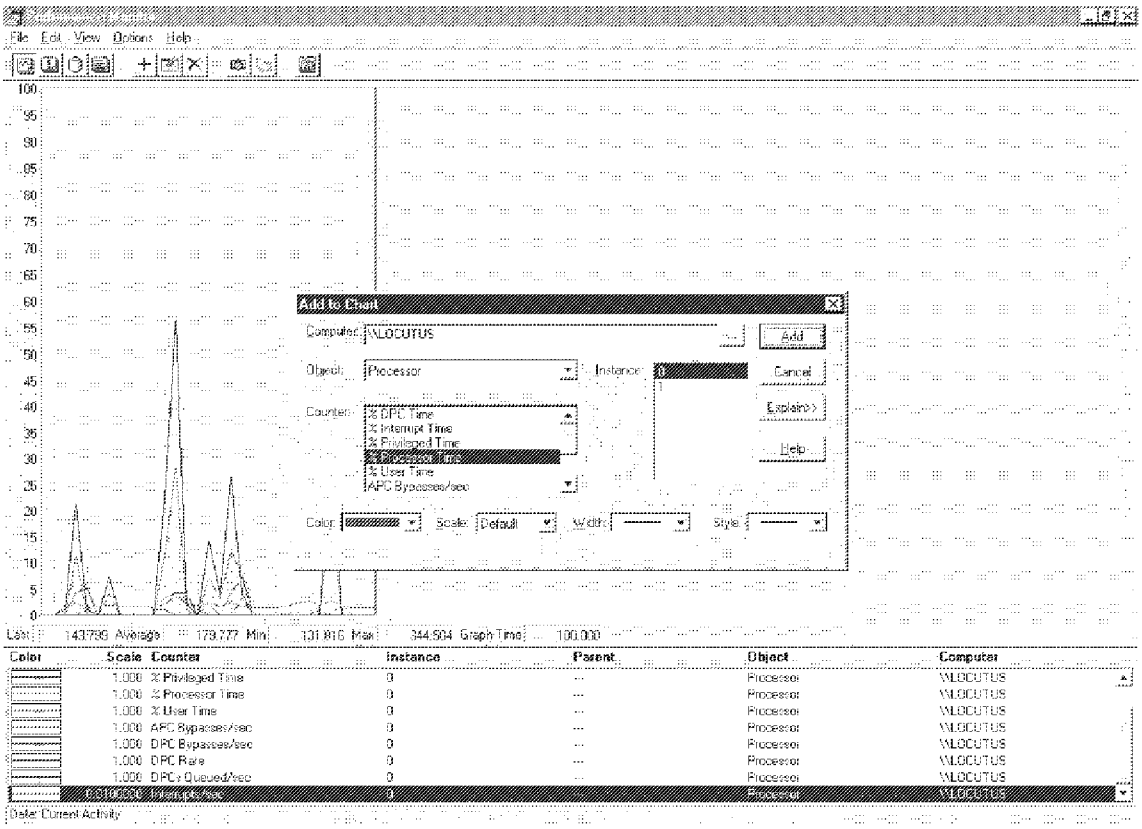

Figure 209. Objects and Counters That Can Be Added to a Chart

12.8.1 What SMS Cannot Do For You

Although SMS can perform many useful functions, some features are missing, such as:

- ***Complete Automation of Inventory:*** Automatic inventory of network routers, bridges or similar devices is currently not possible. Nor is the automatic inventory of various computer peripherals, such as tape drives, printers, scanners, or similar equipment.
- ***Software Licensing:*** There is no built-in means to support license requirements for network applications. There are third-party solutions to accomplish this task.
- ***Real-time Support:*** Depending on the size of the network and the available bandwidth, the installation of packages could take hours or days to be distributed. This also means that the inventory from these remote sites could take just as long to be recorded in the central site's database.
- ***Remote Control:*** Limited to MS-DOS, Windows 3.x and Windows 95. A Windows NT computer can support only remote access, not remote

control, to do basic diagnostics (WINMSD), performance monitoring (Performance Monitor), event log, account database, and Server Manager. Macintosh and OS/2 workstations are even more severely limited.

- ***Installation Scripts:*** Currently, Microsoft only provides Microsoft Test for automating installation scripts, and this tool is not easy to use. There are a number of third-party tools that can make the task much easier.

12.9 Remote Access

The services and features included in Microsoft Windows NT Remote Access Service provides WAN support, protocol support, and security.

Note that the remote clients use standard tools to access resources. For example, the Windows NT Explorer is used to make drive connections, and **Add Printer** within the Printers folder is used to connect printers. Connections made while connected through the LAN via these methods are persistent; so users don't need to reconnect to network resources during their remote sessions. Since drive letters and UNC (Universal Naming Convention) names are fully supported via RAS, most commercial and custom applications work without any modification.

12.9.1 RAS Components

RAS service is made up of three components:

- RAS User Interface

  This component is responsible for all user interactions and includes the Windows GUI and character interfaces.

  Note: No command-line utilities are available for RAS administration.

- RAS Service

  This component is responsible for providing NetBIOS gateway, router management, and authentication services.

- RAS Subsystem

  This component is responsible for providing connection management and integrating drivers for various media and device types.

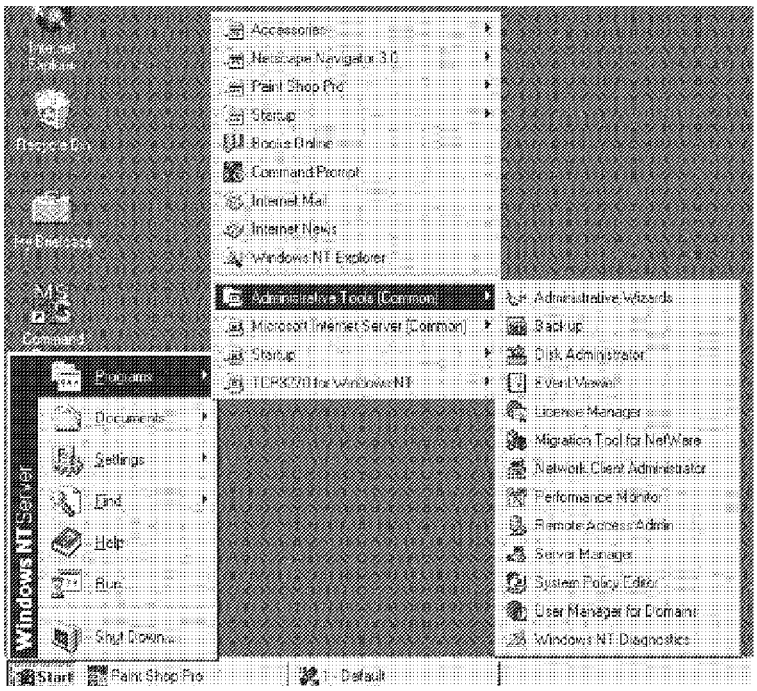

Figure 210. Windows NT Remote Access Admin Pop-Up Window

12.9.2 RAS Protocol Options

The protocols that are used over the physical connections outlined in the previous section are the Serial Line Interface Protocol (SLIP), Point-to-Point Protocol (PPP), and the Microsoft RAS protocol.

12.9.2.1 Serial Line Interface Protocol

SLIP is an extremely basic protocol developed for the UNIX environment. It operates without error checking, flow control, or security. SLIP remains popular because it operates with little overhead and provides good performance. RAS supports SLIP for dial-out purposes, enabling clients to access UNIX computers and Internet providers. SLIP does not support protocol multiplexing and demultiplexing. This means that a SLIP connection cannot be used for transmitting different protocol traffic from multiple sessions between two computers. RAS does not provide a SLIP server in this release of NT Server.

12.9.2.2 Point-to-Point Protocol

PPP overcomes the limitations of SLIP. It can be used for protocol multiplexing and demultiplexing and because RAS 2.0 enables clients to load any layers of NetBEUI, or IPX , and TCP/IP, it performs error checking and recovery and can cope with noisier lines. Like SLIP, it is popularly used for dial-up connections to Internet hosts and is also used to provide

point-to-point connections between routers. Although PPP has a slightly higher overhead, it has become the preferred protocol for remote access, and RAS supports both dial-in and dial-out operations.

Both SLIP and PPP are being widely used, but over time PPP is expected to be more widely used than SLIP due to the benefits of multiprotocol routing.

12.9.2.3 Microsoft RAS Protocol

The Microsoft RAS protocol is a Microsoft proprietary protocol that uses NetBIOS for RAS connections and is supported by all versions of RAS. Microsoft RAS rquires RAS clients to use the NetBEUI protocol. The RAS server acts as a gateway for other protocols such as IPX and TCP/IP.

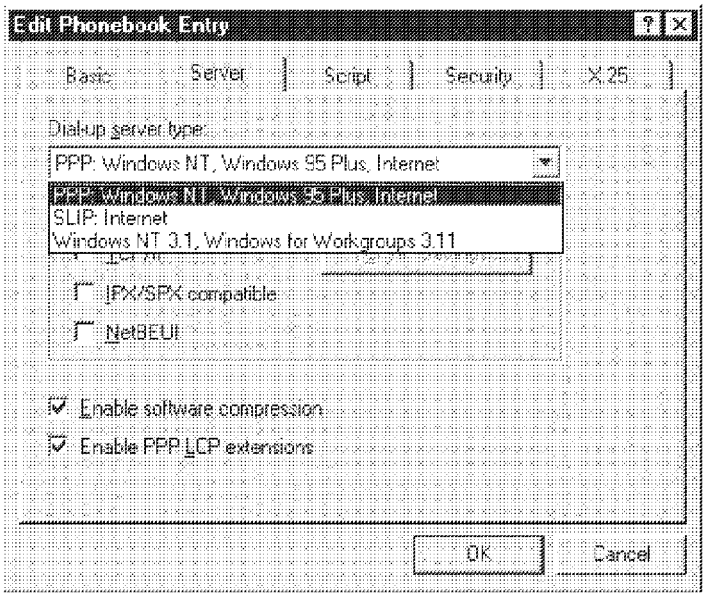

Figure 211. Windows NT Network Protocol Settings

12.9.3 Security Considerations

Windows NT RAS implements a number of security measures to ensure that the remote user is a valid remote access user on the network.

12.9.3.1 Integrated Domain Security

Windows NT Server provides for a enterprise-wide security by using a Trusted Domain, a single-network logon model. This eliminates the need for duplicate user accounts across a multi-server network. The single-network logon model extends to RAS users. Under the single-network logon model, once a user is authenticated, he/she carries with him/her his/her access credentials. If he/she tries to access a resource on the network, NT presents his credentials for him.

The RAS server uses the same user account database as the Windows NT-based computer. RAS access is granted from the pool of all Windows NT user accounts. An administrator grants a single user, group of users, or all users the right to dial into the network. Then, users use their domain login to connect via RAS. Once the user has been authenticated by RAS, he/she can use resources throughout the domain and in any trusted domains. This allows for easier administration because the users log on with the same user accounts that they use in the office. This also ensures that they have the same privileges and permissions they normally have while in the office.

To connect to a RAS server, a user must have a valid Windows NT user account as well as the RAS dial-in permission. Users are authenticated by RAS before they are even allowed to attempt to log on to Windows NT.

Windows NT provides the Event Viewer for auditing. All system, application, and security events are recorded to a central secure database that, with proper privileges, can be viewed from anywhere on the network. Any attempts to violate system security, start or stop services without authorization, or gain access to protected resources, is recorded in the Event Log and can be viewed by the administrator. Microsoft's RAS makes full use of the Event Viewer in Windows NT.

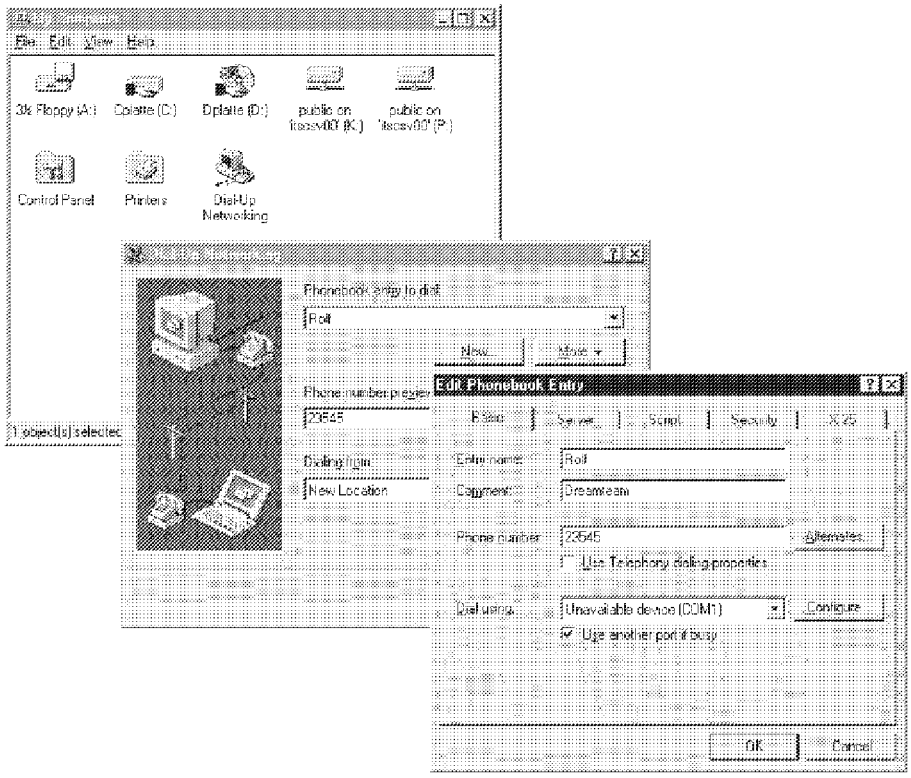

Figure 212. Remote Access Users Defined on the Server

12.9.3.2 How RAS Authenticates User Connections

All authentication and logon information is encrypted when transmitted over the phone line.

To offer a high degree of interoperability, NT provides a number of authentication methods for remote access connections.

When a remote user dials in to a RAS server, the following steps occur:

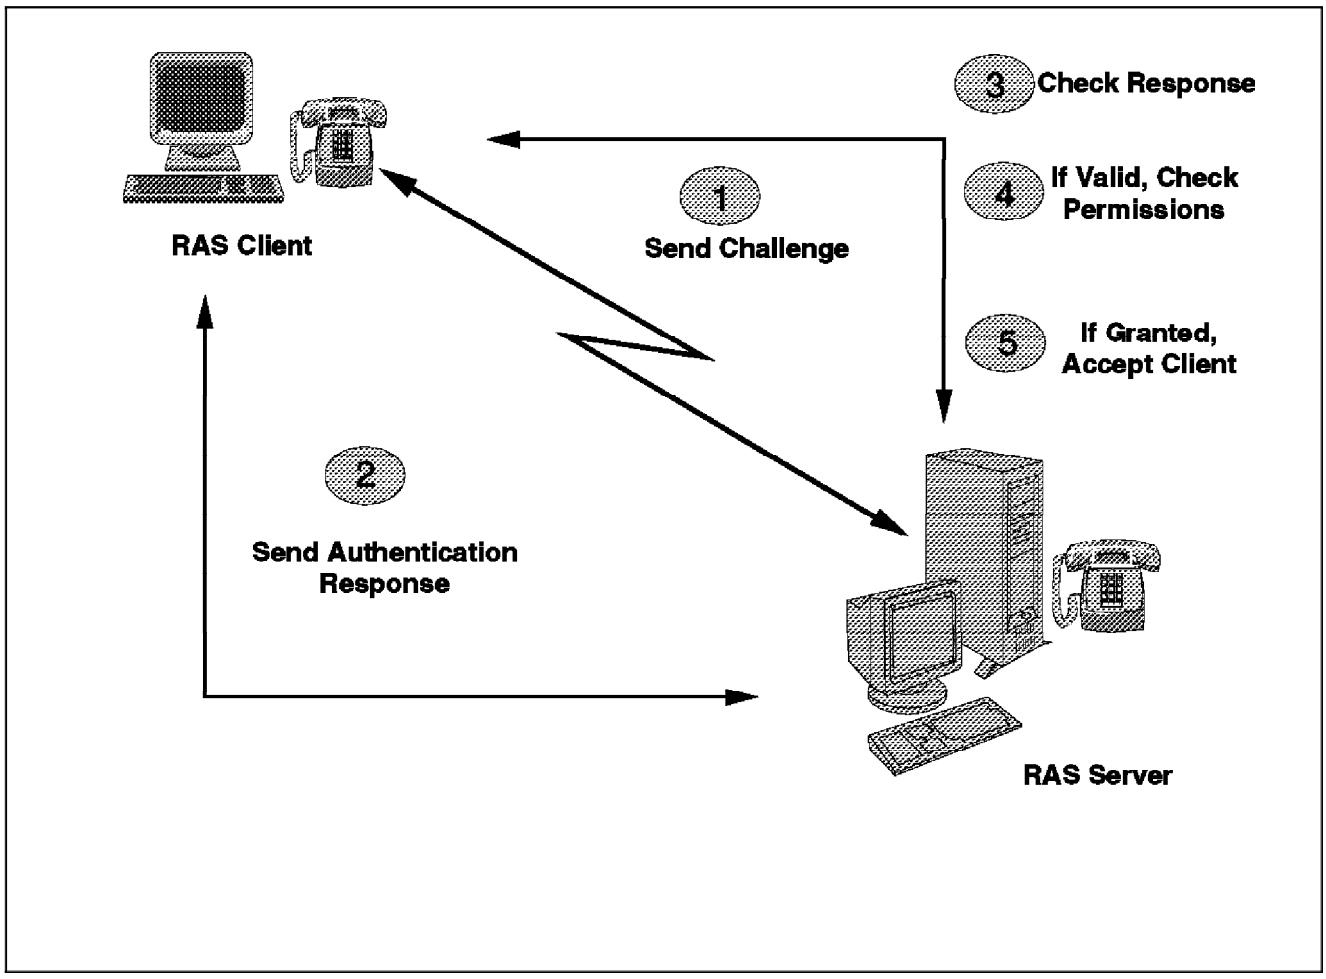

Figure 213. Windows NT RAS Authentication

1. The RAS server sends a challenge to the RAS client.
2. The RAS client sends a response to the RAS server using one of the following authentication methods:
   - RSA Message Digest 5 (MD5) - Challenge Handshake Authentication Protocol (CHAP) (RAS client only)

     This algorithm was designed for speed, simplicity and compactness on a 32-bit architecture. NT supports MD5 for outbound dialing allowing NT clients to connect with virtually all third party Point-To-Point (PPP) servers. Because RSA MD5 requires the availability of the clear text password at the server, NT does not support MD5 for inbound dialing.

     Note: If you use a packet analyzer to watch the traffic, you can read user names and passwords.

The CHAP server sends a random challenge to the client. The client encrypts the challenge with the user's password and sends it back to the server. This prevents someone other than the client from gaining access by recording the authentication and playing it back to the server. Since the challenge is different on each call, a recorded sequence would obviously fail.

- RSA MD4 or MS-CHAP

  A Microsoft version of RSA MD4 is enabled on the NT RAS server by default. It is the most secure encryption algorithm that the NT RAS server supports. Administrators and users can also implement data encryption when using Microsoft Encrypted Authentication. This option uses the RC4 algorithm to encrypt RAS session user data transmitted on the wire if encryption is negotiated between the client and the server at RAS-connection setup time. Either the client or the server can require data encryption to be negotiated.

- Data Encryption Standard (DES)

  Data Encryption Standard (DES) is an algorithm for encrypting data that was designed by the National Bureau of Standards. DES is supported for backward compatibility with LAN Manager-based systems.

- Password Authentication Protocol (PAP) - Clear Text Authentication

  NT RAS supports clear text authentication through PPP PAP. RAS supports both domain\user name and user name formats for the PAP Peer ID. PAP is a plain-text password authentication method and is not very secure. It should only be used when dialing into SLIP servers or PPP servers that do not support encrypted authentication. It is supported for down-level compatibility with third-party applications that work only with PAP. RAS server turns this off by default for security reasons.

- Shiva Password Authentication Protocol (SPAP) - (RAS server only)

  SPAP is a version of PAP implemented by Shiva in their remote client software. NT RAS server supports SPAP to allow interoperability with Shiva clients. SPAP allows you to authenticate with existing Shiva servers. Implementation is a reversible encryption scheme. It is not as secure as CHAP, but it is more secure than PAP. Shiva's SPAP, unlike standard PAP, does not send the clear text password on the wire.

3. The RAS server checks the response against the user accounts database.

4. If the account is valid, the RAS server checks for Remote Access Permission.
5. If the Remote Access Permission has been granted, the RAS server accepts the RAS client. If callback is enabled, the RAS server calls the RAS client back and repeats steps 1 to 5.

After these steps have been completed and the RAS connections have been established, the user is logged on to the network and can access network resources.

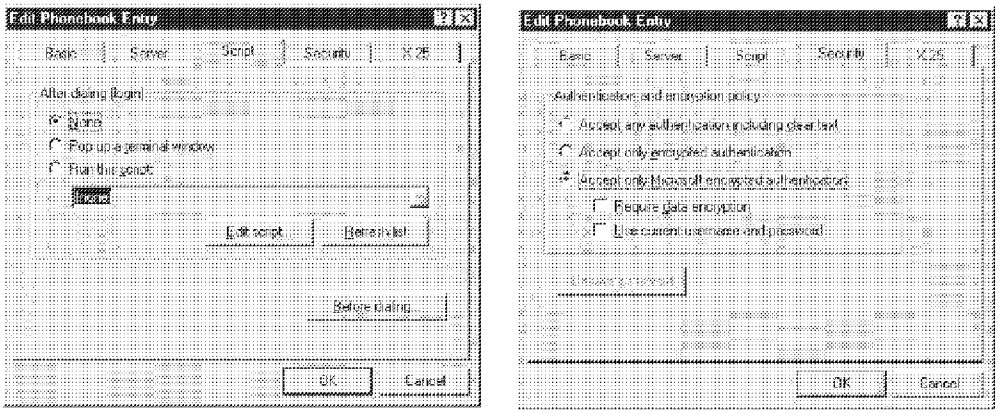

Figure 214. Windows NT Security Settings

12.9.3.3 Callback Security

The RAS server can be configured to operate in dial-back mode as a means of increasing security. This allows the RAS server to call the remote user to verify the connection to the local network. It can be also configured to call back to a preset number; so that way, it is not enough for an intruder to have a valid user name and password, but he/she must be calling from a specific telephone to get in.

All callbacks occur after the user has been authenticated by RAS, but before the user has logged on by Window NT.

12.9.3.4 Network Access Restrictions

Remote access to the network under RAS is under the complete control of the system administrator. In addition to all of the tools provided with Windows NT Server (authentication, trusted domains, event auditing, C2 security design, and so forth), the RAS Administrator tool gives an administrator the ability to grant or revoke remote access privileges on a user-by-user basis. This means that even though RAS is running on a Windows NT Server-based PC, access to the network must be explicitly granted for each user who is to be authorized to enter the network via RAS.

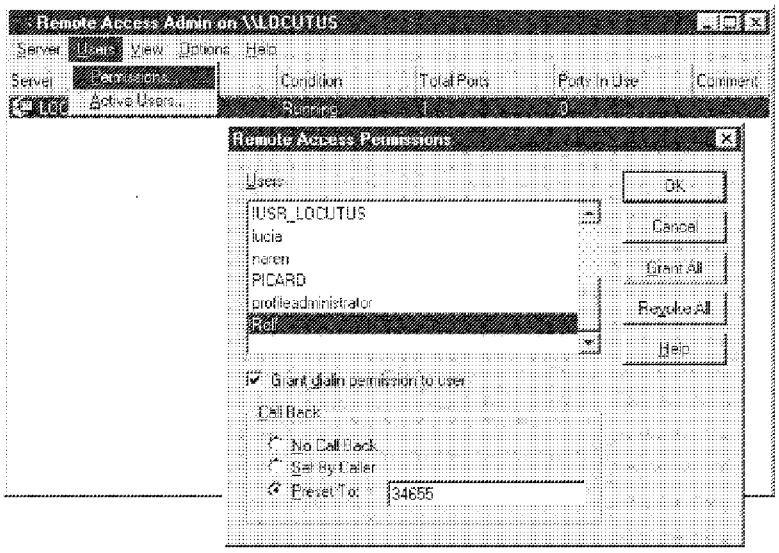

Figure 215. Remote Access Permissions

12.9.4 Third-Party Security

RAS supports third-party security hosts. The security host sits between the remote user and the RAS server. The security host generally provides an extra layer of security by requiring a hardware key of some sort in order to provide authentication. Verification that the remote user is in physical possession of the key takes place before they are given access to the RAS server. This open architecture allows customers to choose from a variety of security hosts to augment the security in RAS.

12.9.5 RAS NetBIOS Gateway and Routers

12.9.5.1 RAS Gateway

The Window NT RAS server has the ability to act as a NetBIOS gateway between NetBEUI and the other NetBIOS-compatible protocols installed on the Window NT RAS server. The NetBIOS gateway allows the remote clients to access NetBIOS resources such as file and print servers on any network to which the Window NT RAS server is attached. It is possible because the NetBIOS gateway receives NetBUEI requests from the RAS clients and passes the requests to any NetBIOS compatible protocol installed on the RAS Server.

12.9.5.2 RAS Routers

Window NT enhances the RAS architecture by adding IP and IPX router capabilities. Window NT remote RAS clients can run TCP/IP and IPX locally and can run Windows Sockets applications over their TCP/IP and IPX protocols across the RAS connection.

A RAS server that is configured with IPX or IP can act as a static IPX or IP router. Static routers do not participate in dynamic routing update exchanges as do normal routers and can route only between networks they are connected to.

For example, a RAS client has the ability to access remote NetWare servers by using the Client Service for NetWare over IPX and RAS or by using FTP over TCP/IP and RAS.

12.9.6 Modems

In order to maximize performance on limited speed-lines, RAS takes advantage of modem data-compression features or compresses data in software when you use modems that do not do their own compression. Software compression is based on the Microsoft DRVSPACE compression algorithm, which can be as much as eight times faster than a connection without compression and it has an average 2:1 compression ratio.

If for some reason you have a modem that is currently not supported by RAS, then you have to tell NT how to communicate with your modem by modifying the MODEM.INF file.

The MODEM.INF file contains information describing each modem supported by RAS. Each entry contains the modem's maximum data rate, the maximum port speed to use connected to the modem, and the character strings to send to the modem to make it dial, enable, and disable data compression and perform other tasks RAS may want it to do.

Adding an entry for a new modem requires the command sequences for the modem. These can usually be found in the documentation that accompanies the modem, although Microsoft recommends not editing this file at all.

12.9.7 Remote to Central Server in NT

Microsoft NT has the same feature as OS/2 Warp Server, which allows you to install RAS without the need for a network card in the machine. During the installation, you have to set the machine up as a server, of course, and NT will automatically detect that there is no network card present. You will then be given a choice of aborting the setup or installing RAS. After RAS has been set up and NT has completed installing and has rebooted, you have to install the MS Loopback Adapter using the Network icon in Control Panel so that the remote machines can connect to the servers' resources. We have tested two simultaneous connections and logons, but Microsoft has claimed support for up to 256 simultaneous connections (provided of course that you can find the hardware).

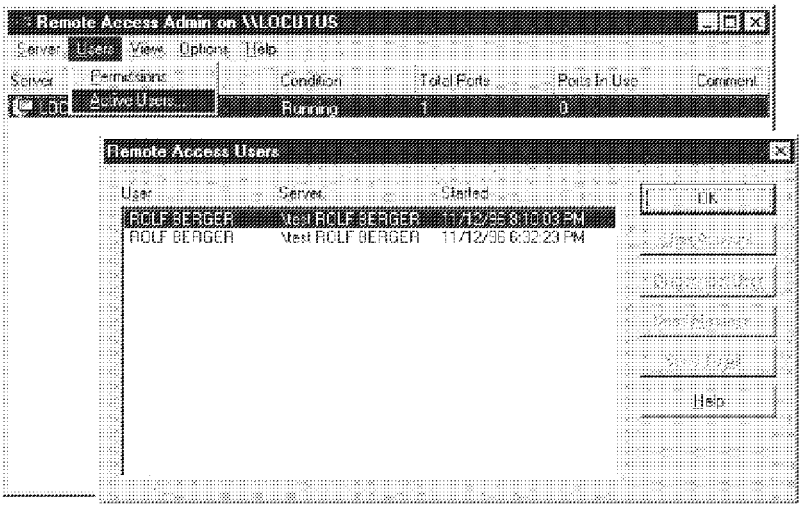

Figure 216. Current RAS Users on the Server

12.10 Backup and Recovery

Windows NT has a simple backup/restore program that comes with the product that lets you back up your data by copying from your hard disk to tape cassette, thereby saving your data from being lost or from any other kind of damage. This backup feature can even be done across the network to the tape device provided that a local disk letter has been assigned.

The limitation of this backup/restore product is that it supports only tape-device with SCSI interfaces and a very limited choice of other tapes. Therefore to use the Windows NT Backup/Restore utility, you will need to use a compatible tape. It is possible to see which tape device is available by consulting the latest Windows NT hardware list. It does not support backup from one server's hard disk to another and provides no support for dynamically connecting and disconnecting network drives during a backup operation.

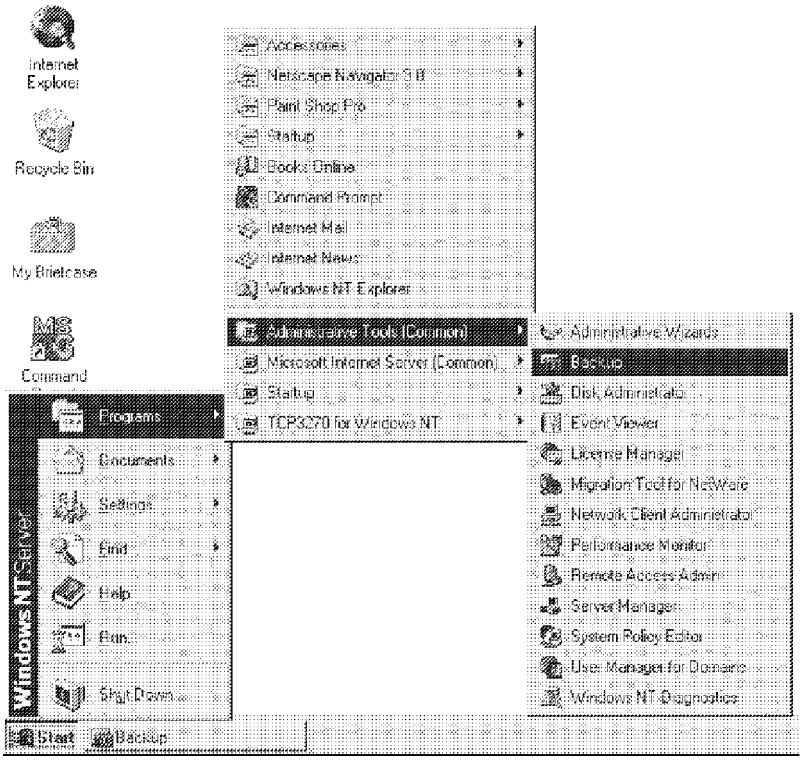

Figure 217. Navigation to Windows NT's Backup Program

Before starting the backup process, you should first check that you have the correct tape device. To do this, you should:

1. Open the Windows NT Control Panel (Start-button — Settings — Control Panel)
2. Open the Tape Devices window and select the **Devices** page.
3. Select the **Detect** button to have Windows NT search for connected tape drives. Once found it will be included in the list of tape devices. If nothing has been found, you will not be able to use Windows NT's backup/restore program.
4. Select the **Drivers** page and click on **Add...**.
5. At the Install Driver window, select first the manufacturer of your tape drive in the list of manufacturers and then select the tape device in the list of Tape Devices.
6. Select **Have Disk...** and provide path information where the drivers can be found.
7. When the install has finished, you might have to reboot your system to make the changes effective.

As you can see in Figure 217, to start the backup program, you must navigate through Windows NT's navigation system until you detect the **Backup** entry in the list of Administrative Tools.

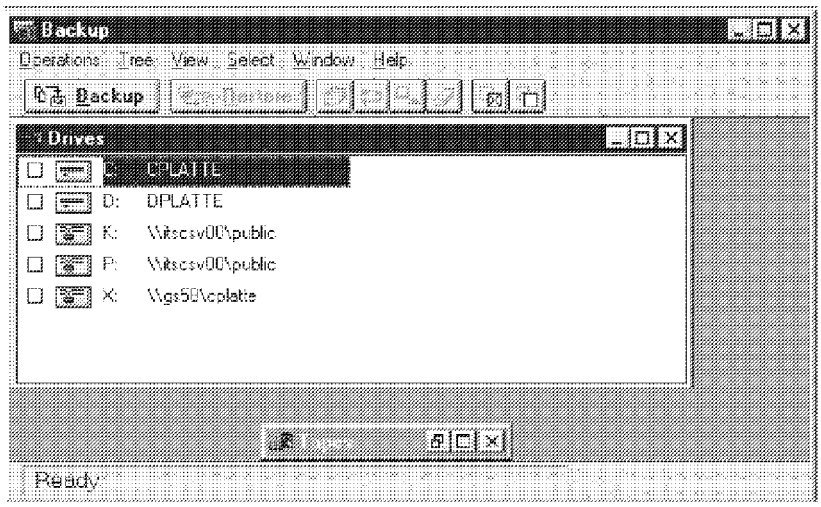

Figure 218. Backup Tree

You can decide to back up the whole drive you have selected or individual files. What you have to do is just click in the checkbox situated to the left of what you want to backup. As you can see in Figure 218, you can give the backup process three different kinds of information: the files you want to backup, the kind of backup you need (full, only changed files, daily, and so forth) and you can even choose a description of the backup, which is helpful when restoring your files. After you have selected the files you need to back up, select the **Backup** push-button. The Backup Information window will come be opened for you as shown in Figure 219 on page 407.

Note: We attached the SCSI tape device to a Windows NT 3.51 workstation; therefore the following screen-captures of an actual backup and restore give you an impression of how it looks in Windows NT 3.51. If you are using Windows NT 4.0, the content of the screens will be the same; only the window controls have changed.

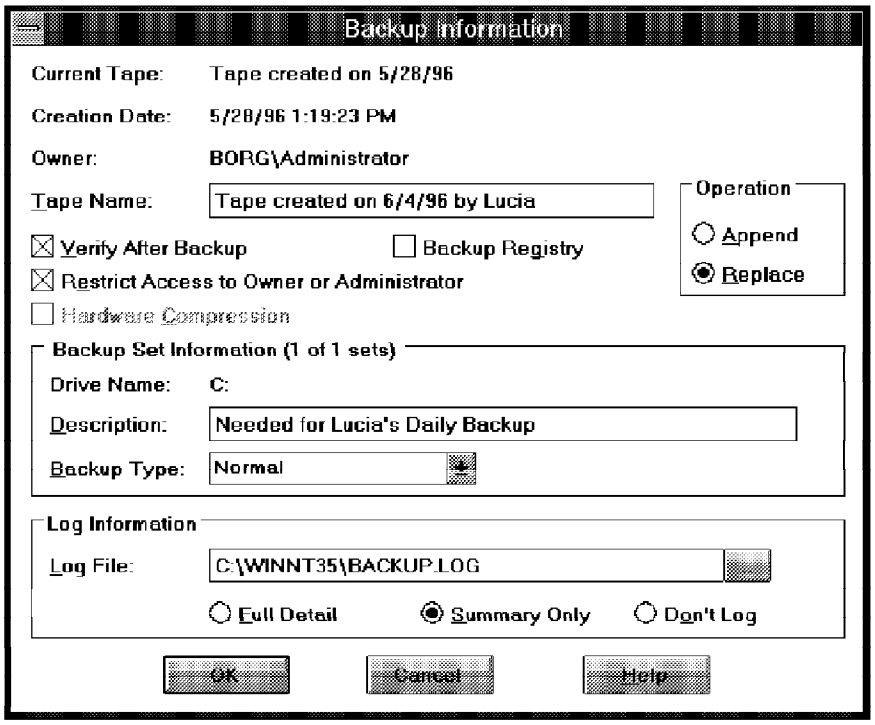

Figure 219. Backup Information Dialog Box

The Backup Information window, as shown in Figure 219, gives you a lot of information on the tape and on what kind of backup you want to make. Here is a description of the fields:

- Current Tape

  This gives you the information on the tape that you put in the drive.

- Creation Date

  It tells you when the current tape was created.

- Owner This is the domain and user name of the person who made the backup.

- Tape Name

  You can write any kind of information that will help you to remember what is on the tape. If you leave the default, it will be the creation date.

- Verify After Backup

  If you select this option, the backup program will compare that the backup matches the original data on the disk; choosing this option, the backup procedure takes longer.

- Backup Registry

This choice will do a copy of the local registry files in the Backup Set. The local registry files are your configuration information.

- Operation

  In this checkbox you have two choices:

  - Append
  - Replace

  If you have data on your tape, you can use one of these two choices to decide what you want to do with the data already on the tape. If you select *Append*, you will add the new data that you have just backed-up to the old backup. If you select *Replace*, you will replace the old backup with a new one. Be sure you will not need the old one anymore because there will be no way to get your overwritten data back.

- Restrict Access to owner or Administrator

  No one except for the owner or administrator can have access to the backed-up files.

- Drive Name

  The drive name is the name of the drive you selected before getting into this dialog box. If it is not correct, you must cancel, go back, change the drive, and start all over again.

- Description

  It is a blank field where you can write in whatever may kelp you to remember the contents of your tape.

- Backup Type

  If you click on the arrow on the side of this box, you will have five different kinds of backup types:

  - Normal

    A full backup of the selected files. Files are marked as backed-up, and the files' archive bit is reset.

  - Copy

    A full backup of selected files, but files will not be marked as backed up; the archive bit is unchanged.

  - Incremental

    Backup selected files modified since the last backup (the archive bit is on); files are marked as backed up (the archive bit is turned off).

  - Differential

It's the same as incremental, but files are not marked as backed-up (the archive bit is unchanged).

- Daily Copy

  Backup files that have been modified on the current day; files are not marked as backed up (the archive bit is unchanged).

• Log Information

  Before starting to back up, you can choose if you need a backup log. To have this option, you must check one of the radio buttons for either Full Detail or Summary Only or Don't Log in case you decide you don't need a backup log.

  Full Detail is a log that will tell you step-by-step what is being done, how many files were backed-up, how many files were skipped (if any), if there were errors, and how long the backup took. In addition, it records the name of every file backed-up.

  Summary Only just has the major events described above.

  - If necessary, the backup procedure will ask you for a second tape.
  - To stop a backup procedure at any time, click on **Abort**.

Once you have determined all your backup decisions for a correct backup strategy, you are ready to select the **OK** button. At this point, a Replace Information window will pop-up as shown in Figure 220.

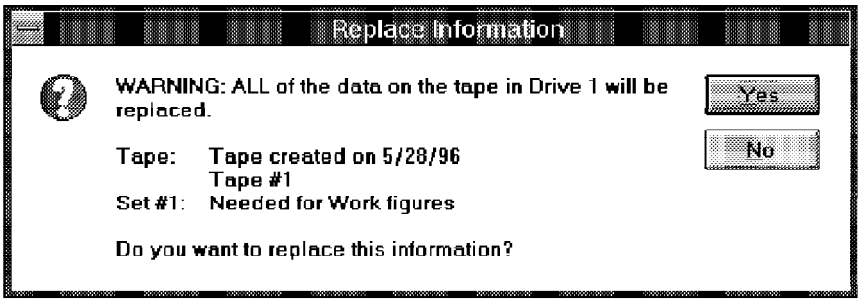

Figure 220. Replace Information Dialog Box

After that you have selected **Yes**, another backup status screen will pop up and shows you what is going on during the backup procedure. It is possible to use more than one tape for the backup procedure, and it is very easy. If you choose **Append** at the Backup Information window, or if you have a very big hard disk drive to back up, when the tape is finished, the backup program will prompt you for a new by presenting you a pop-up window. Insert the new tape and press **OK**. It is also possible to back up to a floppy disk in case you don't have a tape device drive. Due to the small capability

of diskettes, we recommend only to backup a few of your most important files. To do so, open a command prompt and use either the XCOPY command or the BACKUP command.

Once a backup has been completed successfully, you will be presented with a Verify Status window as shown in Figure 221.

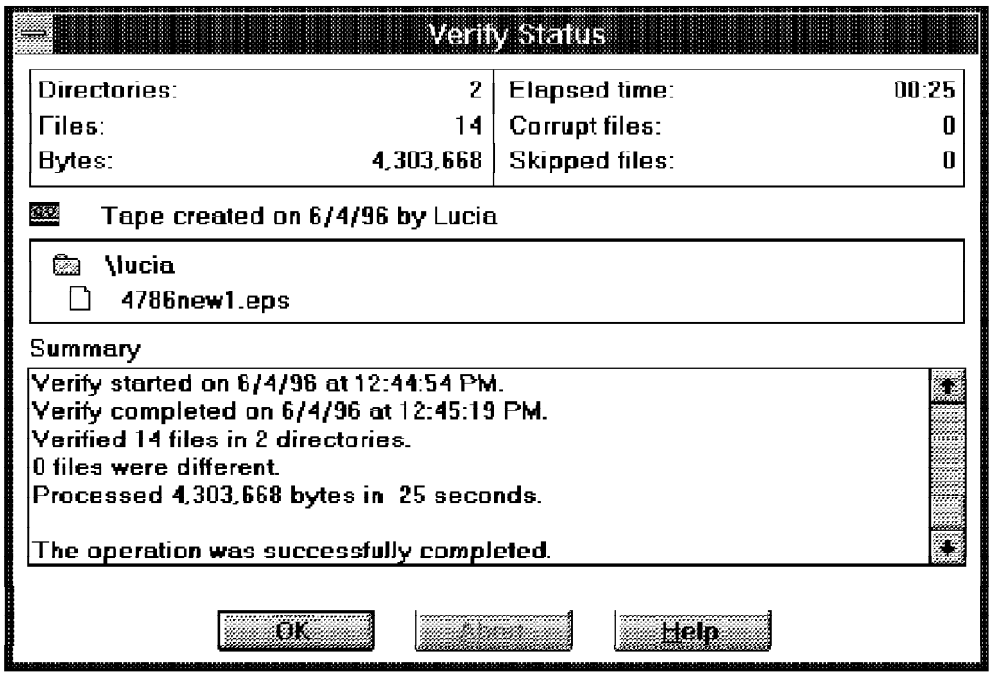

Figure 221. Backup Verify Status Window

12.10.1 Backing Up a Network Drive

Even if you are using an internal tape device driver on your server, you can still use that drive to back up other hard disks at your system. At your system, run the Windows NT Explorer and map a drive or directory you want to back up:

1. Start the **NT Explorer**.
2. From the action bar, select the **Tools** pull-down menu.
3. From the Tools pull-down menu, select **Map Network Drive...** and provide requested information about the networked drive you want to back up.

The networked drive will appear in the backup program's list of drives. It will be represented by new icon that is different to the icon of local drives. In the case of backup, network drives are treated like local drives.

12.10.2 Disaster Recovery Utility

If for any reason you have changed your system configuration so that you cannot boot Windows NT anymore, there are some ways to try to fix it before reinstalling Windows NT. If you watch your machine while it is booting up, you will see a message on a black screen that says:

```
Press spacebar NOW to evoke the Last Known Good Configuration
```

If you press the spacebar as requested, you will see a menu giving you the choice whether you want to:

1. Use the current configuration.
2. Use the Hardware Profile / Last Known Good menu (the configuration that was used the last time the machine booted successfully) option.
3. Restart the computer.

If neither of the three options were successful, you still have another choice before reinstalling the operating system: Using Emergency Disk Repair.

Every time you make a successful change to your system configuration, you should back it up. The backup disk, where your system configuration information is stored, is your Emergency Disk Repair. If you did not create an emergency repair disk, Windows NT creates a WINNT REPAIR directory. Remember that the repair disk contains only a register based on the initial setup.

To update your configuration with all the new settings, you have to:

- Put Disk #1 of the Operating System Setup Disks into drive A and reboot.
- After inserting Disk #2, you will be shown the setup screen, where you will be prompted for:
  - Learn more about setup?
  - Setup now?
  - Repair a damaged installation?
  - Quit setup?

  At this point you must choose **R** to have your system repaired. A message will appear prompting you to insert the Emergency Disk Repair into drive A:

  You will see a message that `Setup is reading REPAIR.INF`. Once it has finished reading, it will prompt you again for the Setup disk.

After you insert Disk #1 again, Setup will ask you if you want to do the following:

- Inspect registry files?
- Inspect startup environment?
- Verify Windows NT system files?
- Inspect boot sector?

By default, all of these options will be checked. To deselect them, use the up and down arrows to select an option, and then press **Enter** to select or deselect the option you want. Select **Continue**, press **Enter** to start adapter detection, and you will be asked to insert Setup Disk #3 to load various device drivers. At this point, you will be asked to insert the Emergency Disk Repair. Setup will start checking your C Disk or the disk where you have installed your operating system. When the Setup has ended, you can press **Enter** to restart your computer. In addition to Emergency Repair Disk, you can save your partition information so that you can replace it at any time by doing the following :

- Start Disk Administrator (Start-button — Programs — Administrative Tools — Disk Administrator).
- From the **Partition** pull-down menu, choose **Configuration**, and then **Save...**.

  An Insert Disk window will be presented to you.
- Insert a formatted disk labeled "Emergency Repair Disk" into drive A: and then press **OK**. An information window will appear confirming that the the configuration was saved.

Disaster Recovery Features

The Hardware Profile / Last Known Good menu will not work if you have made a change more than one successful boot ago. For the Emergency Disk Repair there is a registry based on your initial setup. None of the permissions that you have established are on the Repair Disk.

12.10.3 Windows NT Restore

As you can see in Figure 222 on page 413, to restore the files is much like backing them up. It is possible to restore all the files or just some selected files you need.

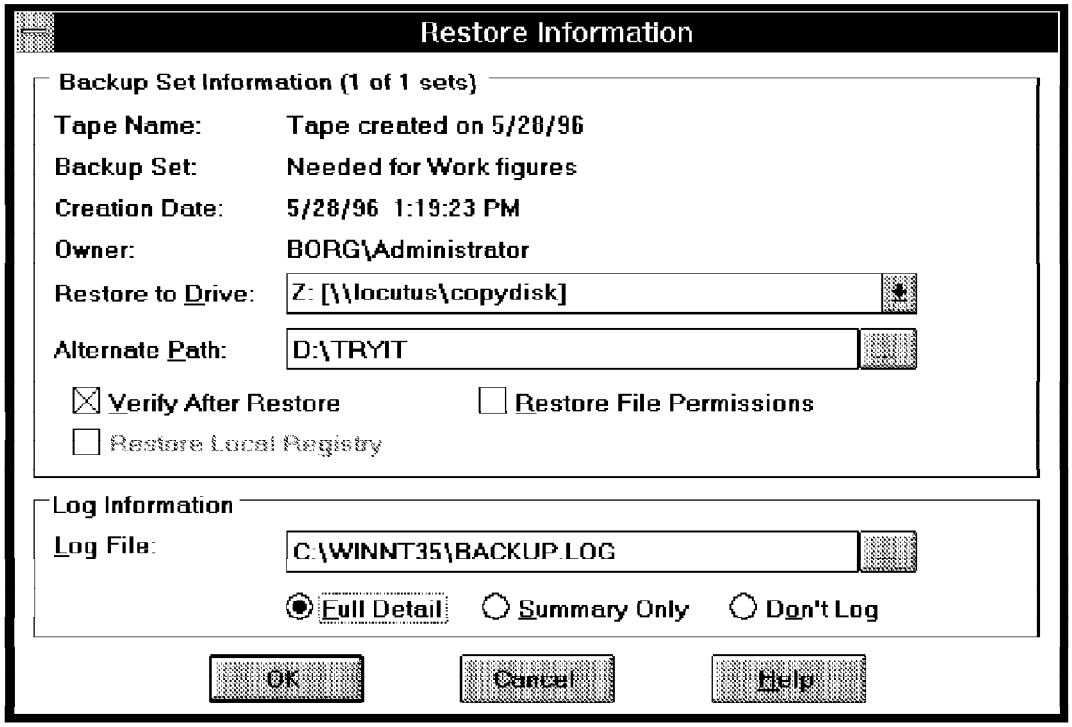

Figure 222. Files Restore Screen

To restore files:

- Start the Backup program (Start-button — Programs — Administrative Tools — Backup).
- From the action bar, select **Operations**.
- From the operations' pull-down menu, select **Restore...**.
- Insert the tape into the tape device and click on the tapes icon.
- Backup displays all the tape information on the left side of the Tapes list. It shows the the drive backed up, the backup tape, and the date and time of the backup. Select the tape containing the file/files you need to restore: double-click the tape's icon, select **Operation/Catalog**, or click on the **Catalog** button in the toolbar.
- Backup displays a list of the backup sets in the backup window. A question mark is displayed with each icon, meaning that the list of files has not yet been read from the tape's directory. Select the backup set you need; double-click on the backup set icon, select **Operation/Catalog**, or click on the **Catalog** button in the toolbar.
- The program displays the list of directories and files in the hierarchy in the Tape File Selection window:

1. To restore all files, select the checkbox for the tape, and choose **Select/Check** or click on the button in the toolbar.
2. To restore an individual file, select the file's checkbox.
3. To restore multiple files that are not contiguously listed, press and hold the **Shift** key and select the first and the last files. Choose **Select/Check** or click on the **Check** button in the toolbar.

- At this point we are ready to restore. Click on the **Restore** button; a dialog box similar to the Backup Information dialog box will open. There are a few decisions to make before you click on the **OK** button:
  - You must decide on which drive you want to restore you tape. It is possible to restore to a different drive than the one from where you backed up.
  - You can decide to verify your data after restore. Even if it takes longer when you select this feature, you will be sure that you restoring process ended properly, although it is not possible to restore registry information to a different drive than from the one you backed it up from.
  - You can Restore File Permissions that were in place for the file when you backed it up.
  - You can also choose what kind of log file you want and where you would like it to be stored.

When you have finished, click on **OK**, and the restore feature will start. While it is in progress, it is possible to see the process on the screen, as shown in Figure 223 on page 415. At the end of this process, your files are successfully restored to your disk.

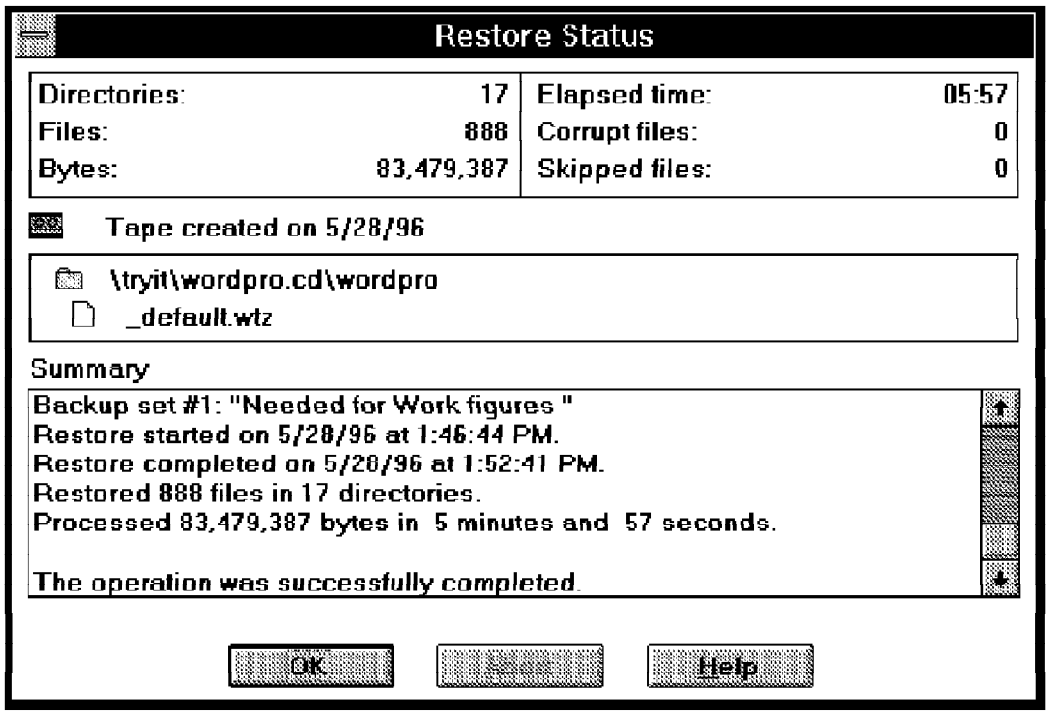

Figure 223. Restore Status

Restoring Data After Reinstalling NT Server

If for some reason you decide to reinstall your Windows NT Server, you must reinstall all the system software that was previously installed, including service pack patches, and only then restore the data. If you do not do it in this way, your backed-up data will never be restored, and you will get a message saying:

```
The system generated an application error.
```

and then the backup program will close. So only when the Windows NT Server is reinstalled with all the software needed, your restore will work. It is also important to remember that Windows NT does not have a standard disaster recovery availability.

Chapter 13. Novell Directory Services

Novell Directory Services technology is a distributed name service that provides global access to all network resources regardless of where they are physically located. Users log on to a multiserver network and view the entire network as a single information system.

This section describes the basic administration of Novell NetWare 4.1 Server. We will discuss user, group, drive, and printer administration. Any additional functions in the NetWare network operating system administration is not part of discussion in this chapter.

13.1 NetWare 4.1 Administration Tools

To administrate a Novell NetWare 4.1 Server you need to use the NetWare administration tools that come with the NetWare Requester package. As the name implies, these administration tools can only be used from a requester workstation. They can not be invoked at a server. When NetWare Version 4.0 was introduced two years ago, you had the choice of a OS/2-based version or a non-OS/2-based version. With the OS/2 version, you could benefit of OS/2's multitasking and multithreading architecture. You could run the NetWare administration tools directly on the server.

Since these server-based NetWare administration tools are not offered with Novell NetWare Version 4.1 anymore, there is no support for running NetWare Server and administrate NetWare on the same machine. However, administration tools run either on DOS, Windows 3.1 or higher, or on WIN-OS/2 of OS/2 Warp Connect or OS/2 Warp 4.

Note on NetWare Client for OS/2

OS/2 Warp Connect, OS/2 Warp Server, and OS/2 Warp 4 include Version 2.11 of the NetWare Client for OS/2. Since October 1996, Novell has offered an update. You can download the NetWare Client for OS/2 Version 2.12 from Novell's support page on the World Wide Web at the following URL.:

`http://support.novell.com/home/client/os2/updates.htm`

Once you are linked to the page of NetWare Client for OS/2 2.12, you can also download OS/2 Utilities for NetWare 3.12 and 4.1. NetWare Client for OS/2 v2.12 also is part of NetWare 4.11 and IntranetWare. The new requester code provides NDS support for the global WIN-OS/2 and DOS sessions. This support enables NDS-aware applications and utilities in the global WIN-OS/2 sessions in OS/2. It also enables NDS-aware DOS utilities in the global DOS sessions.

When administrating Novell NetWare 4.1 Server, Novell makes use of objects within Novell Directory Services. For managing Novell NetWare Version 4.1, you use the NetWare Administrator, which is a Windows-based application. It manages Novell Directory Services as well as parts of the file system.

Novell NetWare Version 4.1 also comes with a DOS version of the administration program. However, due to the lack of functions and features of this version, we concentrate on the Windows-based utility `NWADMIN`.

The following list gives you an overview of the different administration possibilities when using these NetWare Objects:

- Create additional objects, such as user and printer objects
- Change the login restrictions of users
- Change user's access to resources
- Change the trustees of objects
- Grant other users supervisory or subadministrator rights to objects on the network
- Specify groups of users and create login profiles for those users
- Create and edit system-wide and individual user login scripts
- Arrange and organize the structure of the Novell Directory Services tree and its partitions

13.2 Understanding NDS Objects

In NDS we work with directory objects that consist of categories of information known as properties and the data included in those properties. This information is stored in a directory database.

The directory database consists of three different types of objects:

- [Root] object (directory tree name)
- Container objects
- Leaf objects

The NetWare Administrator window shown in Figure 224 is the window where you can define objects.

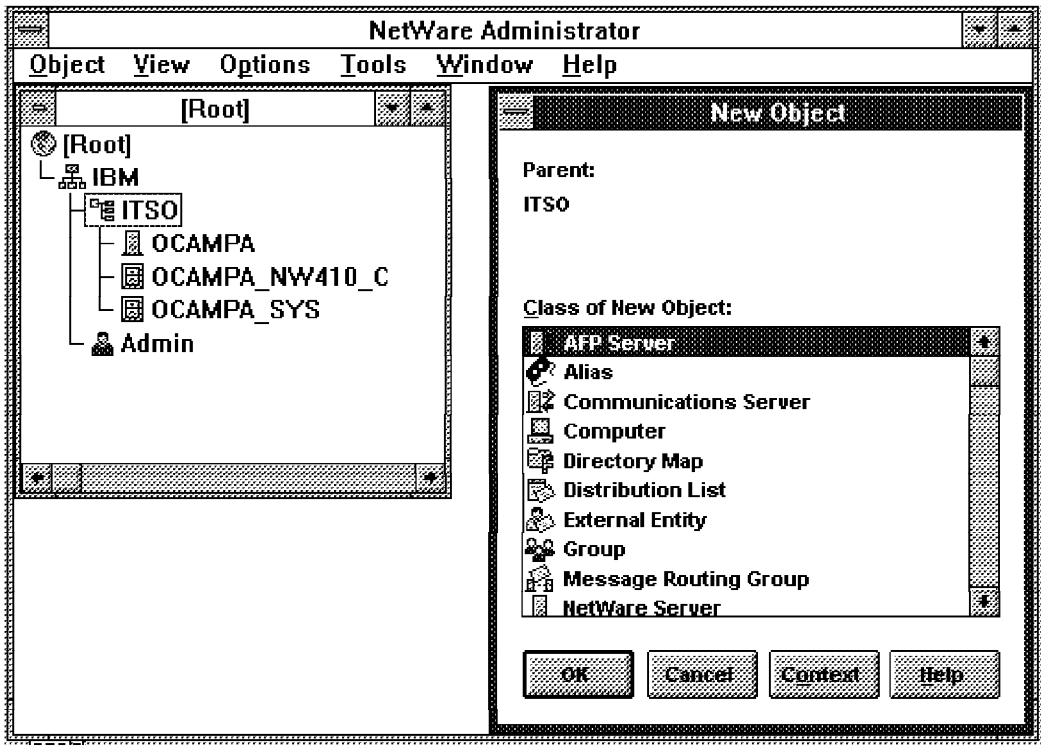

Figure 224. NWADMIN Window with Opened Object Creation Box

In Figure 224, you see an example of Novell Directory Services with the opened Create New Object window. In the Class of New Object list you can get an idea what kind of objects can be defined.

The [Root] window shows that a directory tree has been created with the following objects:

- Container object titled Organization = IBM

- Container object titled Organizational Unit = ITSO

It is very important to understand the whole idea of directory services and how this technology is implemented in Novell Directory Services. Generally, Novell Directory Services has a lot in common with IBM's Directory and Security Server and vice versa. A directory structure gives you an easy way to reflect a big enterprise, large account, or a worldwide operating company or organization in a directory tree. This means that it must be known how a firm is organized. For example, which departments communicate with each other, which ones do not, and where do resources, such as data, reside and how they are shared? These things must be known before you can start to implement a directory structure. Everything is defined by using objects.

Novell Directory Services is very flexible in its architecture. If an organization changes or reorganizes itself — as many firms do more than ever nowadays — you can immediately reflect those changes in the directory just by point-and-click operations.

Figure 225 demonstrates a basic directory tree.

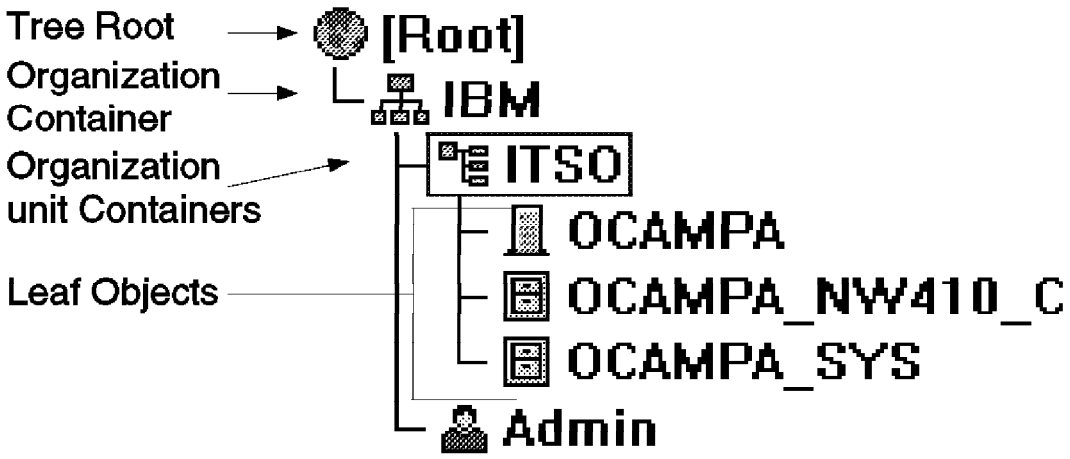

Figure 225. Object in the Directory Tree

In Table 30 you find a description of definable Container Objects.

Table 30 (Page 1 of 2). Organization Container Objects in the Directory Tree

Object	Description
Country	If you run an international network, Novell Directory Services offers you an additional object to split the tree into country-specific containers with suborganizations.
Locality	This object is used to designate regions of a network.

Table 30 (Page 2 of 2). Organization Container Objects in the Directory Tree	
Object	**Description**
Organization	Use this object to define different companies within an organization or different divisions within a company. At least one organization object is required in Novell Directory Services.
Organizational Unit	Organizational units can represent objects like subdivisions or branches in the directory tree. They are used to organize leaf objects. Departments, workgroups, and business units are typically organizational units.

The first time you log in to the tree, it might be necessary that you navigate inside your directory tree, depending on the location of the leaf object that represents your user ID. After the initial installation, the login points to the root directory of Novell Directory Services. This means that you will need to specify the start of your logon. Alternatively, you can do this by using the command:

```
CX.[Organization]
```

Or by entering the user ID with the organization path as shown in Figure 226.

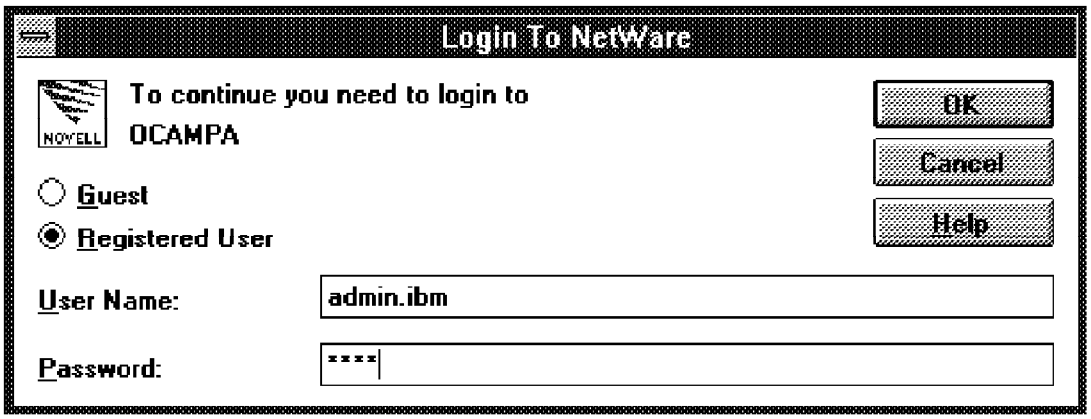

Figure 226. Login with Organization Object Name

When you have created the container structure for organizations, you can add leaf objects. These objects are added to the containers and can represent users and resources. Table 31 on page 422 describes the different leaf objects that can be defined and administered:

Table 31 (Page 1 of 2). Leaf Objects in the Directory Tree

Object	Description
Alias	An alias object points to an object somewhere in the directory tree. You create an alias for an object in one branch of the directory tree that you often need to use to refer to another branch of the directory tree.
AFP Server	An AFP server is an Apple file server running the AppleTalk protocol.
Bindery Object	This is an object that was created by upgrading from former versions of NetWare that used server-specific binderies, like NetWare Version 3.12.
Computer	The computer object holds information like its serial number, node address, user and location of department, and workstation-related information.
Directory Map	The location of applications is specified in the directory map. It also simplifies mapping of directories for a large number of users. Directory maps are used in Login Scripts to map directories for users. If the location of an application changes, you change the directory map rather than each user's Login Script.
Distribution List	A collection of Message Handling System (MHS) mailboxes that can be used to address MHS messages.
External Entity	Is a reference to an object outside of the Novell Directory Services tree.
Group	The group objects are used to group users in different organization groups, like project groups, management groups, mail groups, and groups for other different management purposes. Groups help an administrator to manage his/her network easier.
Message Routing Group	Novell's electronic mail and message delivery system MHS (Message Handling System) is addressable over the Message Routing Group. This is a group of Message Handling System servers that is used to exchange messages.
Messaging Server	Object for a server running Message Handling System services.
NetWare Server	Each single NetWare server on the network, within the Novell Directory Services, is represented by the NetWare Server object.

Table 31 (Page 2 of 2). Leaf Objects in the Directory Tree

Object	Description
Organizational Role	You can assign users as members to the organizational role object. It is similar to the group object. The object has assigned access rights to the system that defines a role on the network. It is a special group, such as department managers or clerks who update user account information. Also you can make users temporary members of the role in case you want them to manage the position. It is recommended to create a role for temporary employees for management purposes.
Printer	This object represents printers attached to a print server or to a workstation and shared in the network.
Print Queue	Network print jobs directed to printers are represented by a print queue object. The users send print jobs to queues, not directly to printers. A print queue can hold one or more printers.
Print Server	Network print servers are represented with a print server object. The network print server may be part of a NetWare server or a stand-alone print server.
Profile Object	To share profiles, Novell NetWare 4.1 Server uses profile objects. In this case the profile is a special login script that is shared by more than one user. The profile script is executed after the script of the user's container but before the user's login script. The profile scripts make it easy to set up a network environment for a group of users. However, the users do not have to belong to the same containers.
User Objects	User accounts are represented by a User object.
Volume Object	A Volume object represents a physical NetWare volume on the hard drive of a file server. It also holds statistics about the volume. However, the assigned volume name does not have to be the same one as the hard disk volume name.

To manage objects, Novell NetWare 4.1 Server includes five major utilities that can be used to manage these objects within Novell Directory Services. These utilities are shown in Table 32 on page 424.

Table 32. Major NetWare Administration Tools	
Administration Tool	**Description**
NWADMIN	The NetWare Administration utility runs on Windows or OS/2's WIN-OS/2 environment. Using this utility, you can manage the Novell Directory Services tree and its objects.
NETADMIN	This is the administration utility to use on text-based systems, like DOS. This also is useful for Windows or OS/2 Warp workstations when administrators want to to issue commands from the command line. The NETADMIN utility does not include file management features. Therefore you have to use the FILER utility or NetWare file commands to work with the file system.
NETUSER	The NETUSER utility is designed to offer the user with the needed functionality to log in, log out, access drives, access printers, and to exchange messages with other users. It is a text-based utility for use with DOS, Windows, or OS/2 from the command prompt.
NetWare User	It is a Windows-based utility that is similar to the NETUSER utility that is text-based.

After you have created the objects that represent your network resources, you can manage those resources in each container. The Graphical User Interface, NetWare Administrator, is used to do this. After installing the NetWare client software, there will be a NetWare Tools group in the Windows Program Manager.

13.2.1 Adding the NetWare Administrator Tool to Program Manager

To simplify the administration you should add the NWADMIN utility to your Program Manager (Windows or WIN-OS/2). Before you do this, you need to log in and mount the SYS:PUBLIC directory.

Note: All users have Read and File Scan trustee directory assignments to the SYS:PUBLIC directory. These directory rights are assigned to the [Public] trustee and therefore to all users. All users have Read and File Scan access rights to the SYS:PUBLIC directory.

The default login script takes care of that when you log in with the LOGIN command from the default-mounted F: drive. We describe the way to do it by using the NetWare User Tools.

To add the NetWare Administrator to your Program Manager you have to mount the SYS:PUBLIC directory. This is done by a login to NetWare. Remember that you have to add the organization, mount the directory using

the NetWare User Tools, and add the NetWare Administrator program to the Windows or WIN-OS/2 Program Manager by following these steps:

1. From NetWare User Tools, select the drive letter you want to use, then select **Map**.

   Note: We recommend to use the Y: drive because this is the standard drive letter added when making the default login from the command prompt.

2. To add this directory/drive for permanent mounting, select **Permanent**.
3. Put the NetWare Tools window into the foreground, and select **New** from the File pull-down menu.
4. Select the **Program Item** radio button, and then select **OK**.
5. Type in Property information as shown in Figure 227. When finished select **OK**.

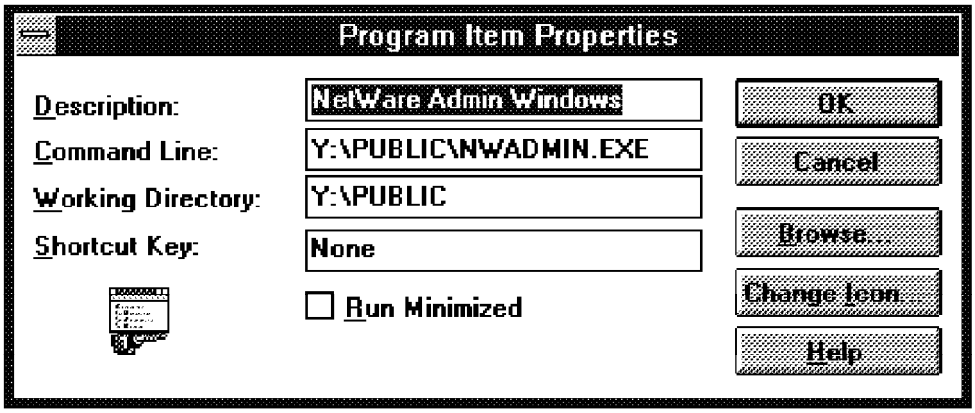

Figure 227. Adding NetWare Administrator to Program Manager

NetWare Administrator is used to browse, create, and manage the directory tree. In Table 33, we summarized the tasks you can perform with NetWare Administrator.

Table 33 (Page 1 of 2). Browsing Tasks

Task	Steps
Expand or collapse a container object view	1. Double-click **Container** object, or open the View menu, and select **Expand** or select **Collapse**.
Change the starting Context of your directory tree display	1. Open the View menu and select **Set Context**. 2. Enter information or select **Browse** for the new context in the dialog box that appears.

Table 33 (Page 2 of 2). Browsing Tasks	
Task	**Steps**
Open a new directory tree browsing screen	1. Open the Tools menu, and select **Browse**.
Select which object types to view	1. Open the View menu, and select **Include**. 2. From the dialog box that appears, select the object types to view.
Search for objects	1. Open the Object menu, and select **Search**. 2. In the dialog box that appears, type in information where to start the search or use the browse function. 3. Select the object type you want to search for. 4. Optionally, select to search based on whether the value for a certain property matches a certain condition or whether the value exists for the property.

There are also steps to follow for managing object tasks. In Table 34 you can find a description how to create, delete, rename, move, and work with the properties of an object.

Table 34 (Page 1 of 2). Management of Objects	
Task	**Steps**
Create an Object	1. Highlight the container object where you want to place the new object. 2. Either press **Insert**, open the Object pull-down menu, and select **Create**, or click on the container object with the right mouse button and select **Create**.
Delete an object	1. Select the object(s) to delete. 2. Press **Delete** or open the object's Context menu by clicking on the object with the right mouse button and select **Delete**.
Rename an object	1. Select the object that you want to rename. 2. Open the object's Context menu and select **Rename**.
Move an object	1. Select the object you want to move. 2. Open the object's context menu and select **Move**. 3. In the dialog box that appears, type in destination information or use the browser function.

Table 34 (Page 2 of 2). Management of Objects	
Task	**Steps**
Working with an object's properties	1. Select the object that you want to work with. 2. Either press Enter or open the Object pull-down menu and select **Details**, or click on the object with the right mouse button and select **Details**.

13.2.2 Managing User Objects

A user object in Novell NetWare 4.1 Server is your link to the network. Unless a user object has been created for the user in the directory tree, the user has no access to the network. To log in, the object name of the defined account is used.

To access resources in the network, the user object must have access rights on these objects. When managing users, it is helpful to create user templates for managing different project groups or departments by using organization or group objects.

In the following section you will learn how to perform the tasks related to configuring a user account and how to create and manage the following Novell Directory Services tree objects.

- Groups
- User templates
- Organizational roles
- Aliases

This advanced structure helps a network administrator to manage his/her network resources and users properly.

13.2.3 Understanding the User and User-Related Object Properties

Objects include unique characteristics called properties. Before you create a object, you will have to type in the required properties information. Properties can be modified and completed after setup also.

The prompting dialog box for creating a object includes the required properties as a minimum. We want to list these properties for the user objects first:

Table 35. Properties for Creating a User Object

Property	Description	Required
Login Name	The name used for login.	Mandatory
Last Name	The new user's last name.	Mandatory
Use User Template	If this box is checked, the user templates properties that could be defined are used.	Optional
Define Additional Properties	By checking this box, you can add additional properties to the user object that are very useful for organization and management. Find in Table 36 on page 434 a detailed description about additional user properties information.	Optional
Create Another User	When this box is checked, you can create many users in one step, without opening the Object pull-down menu again and again.	Optional
Create Home Directory	If you mark this checkbox, a directory with a specified volume and path is created. It automatically assigns the user trustee assignments to this directory.	Optional
Path	This enables the specification of volume and path for the user object's home directory. The select is done by clicking on the browse icon and selecting the volume and directory in which the home directory should be created.	Optional
Home Directory	When setting up a home directory, the NetWare Administrator assumes that you want the same name as the login name and uses the first eight characters because of the DOS limitations.	Optional

13.2.4 Creating a User Object

After installing a server, the next important thing is to create users to use the resources a network server offers. Because NetWare works with the Novell Directory Services concept, it uses container objects as its organizational units. These container objects symbolize the different branches, projects, or organizations.

Before you create user objects, you should first decide on your strategy for choosing user object names. Novell NetWare 4.1 Server gives you the flexibility to choose user object names that can be up to 64 characters long and can include any alphabetic, numeric, or punctuation characters. Object names can even include spaces, but be careful because some utilities for NetWare convert them to underscores. Do not forget that DOS works with the 8-to-3 naming scheme; so be cautious when it comes to the use of user names in addition to user subdirectory/home directory information.

Notice on Object Names

For compatibility with earlier NetWare versions, do not use the following special characters: / \ : , * ?

It is also recommended not to use the equal sign (=) or the plus (+) sign because Novell NetWare 4.1 Server utilities require you to precede them with a backslash whenever you type the name.

To create a new user object in the Novell Directory Services tree, you must have administrator rights. When you create a new server, you use the default user account which is Admin, that was created at installation time. Then follow these steps:

1. In the Program Manager of Windows or WIN-OS/2, open the NetWare Tools window.
2. In the NetWare Tools window, double-click on **NetWare Administrator**.
3. Select the Organization Container in which you want to create a user. In our example we select the container type "Organization unit" titled **ITSO**, by double-clicking on the Organization object, **IBM**.

   Notice on Creating Users in Organization Containers

   You can create users in every organization container; therefore be cautious when you decide on the container object you want to add users to. Also be cautious with the access administration rights you define for user IDs.

4. Click the right mouse button to open the Settings pop-up menu and select **Create**, as shown in Figure 228 on page 430.

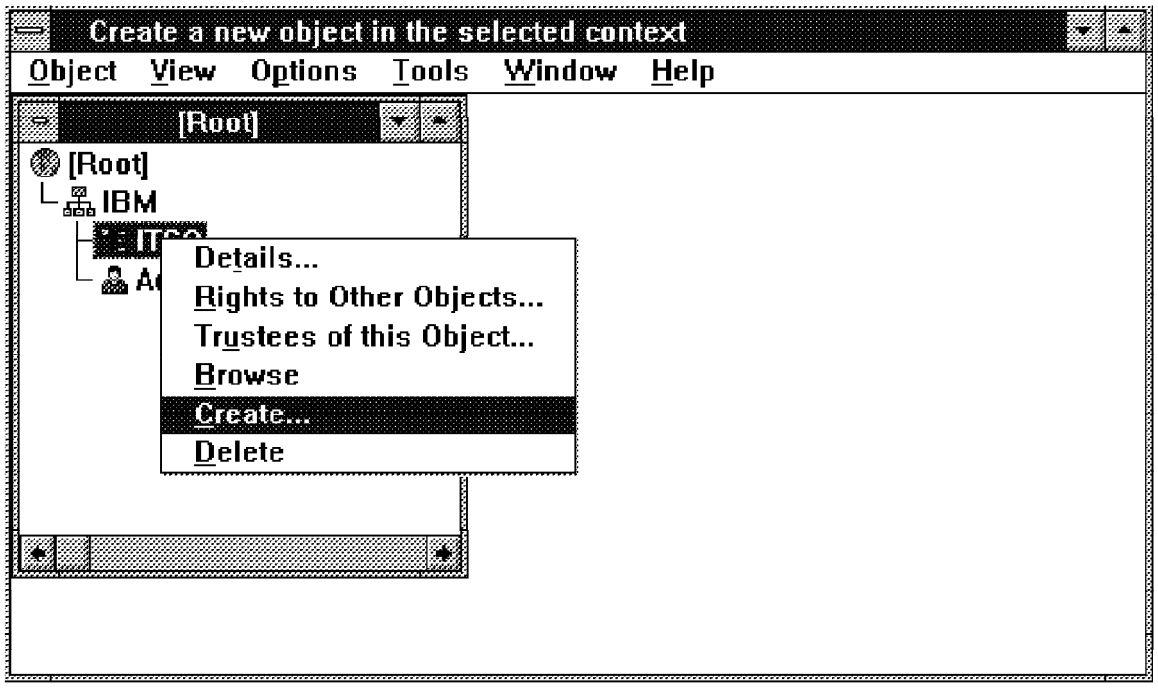

Figure 228. Create New User Object

5. Select the **User** object in the Class of New Objects list in the New Object window and select **OK**.
6. In the next window that pops up, titled Create User, type in the **Login Name**, the **Last Name**, and check the box for **Define Additional Properties** as shown in Figure 229 on page 431.

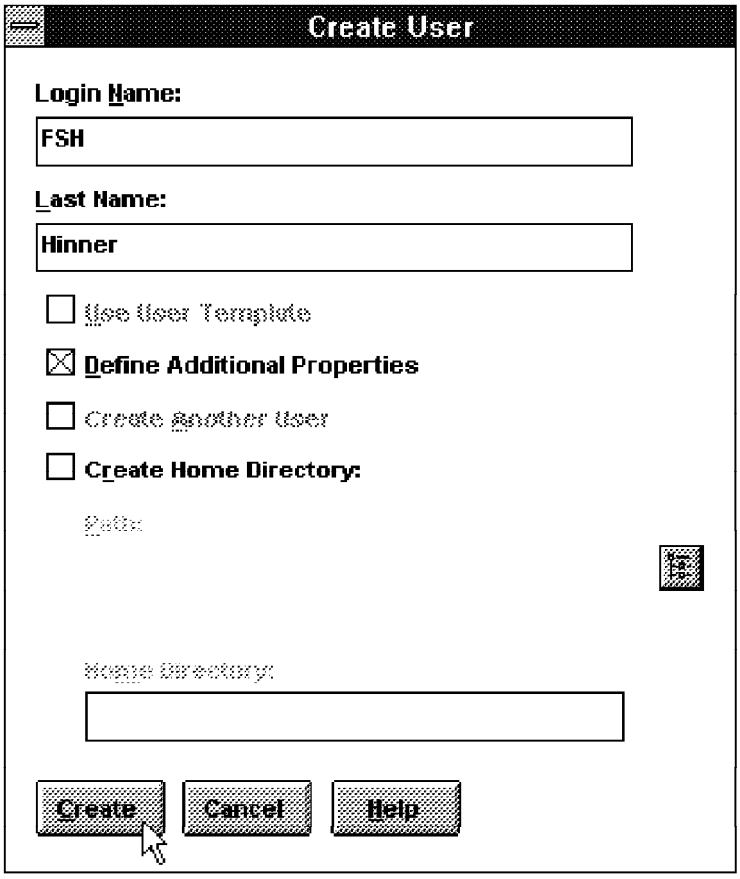

Figure 229. Create User Base Dialog Box

Notice on Creating Users

In this Create User window you can also check the box for Create Home Directory, which would automatically be applied to the selected subdirectory structure in case the check mark is set, which usually is the case. In our scenario, we discuss how to set up home directories separately so that you can get a better idea of the comparison between NDS and Warp Server or Windows NT.

7. Select the **Create** button to create the user, and then select **Additional Properties**.
8. Select the **Identification** page and type in information as needed as shown in Figure 230 on page 432.

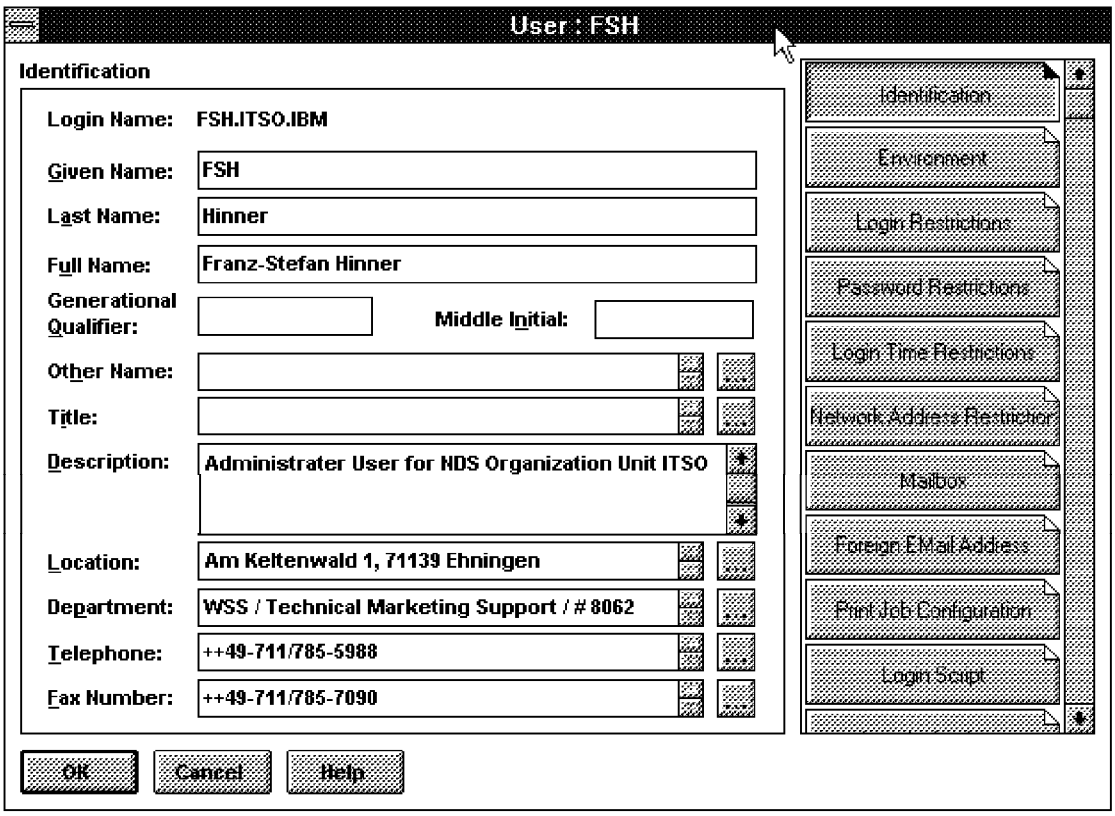

Figure 230. Create User Identification Settings

9. Select the **Password Restrictions** page to provide password information.
10. Check the **Require a Password** checkbox, type in information for the minimum password length, and select the **Change Password** button as shown in Figure 231 on page 433.

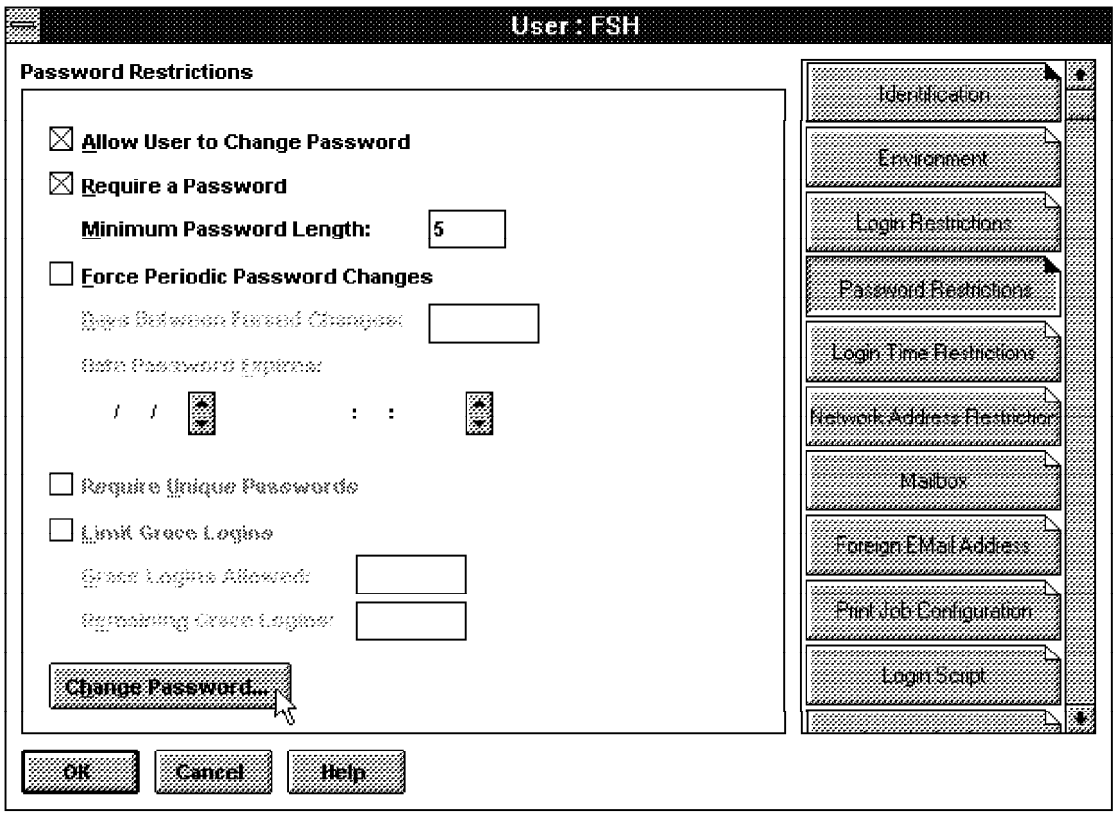

Figure 231. Add Password when Creating a New User

11. In the Change Password window type in the new Password twice and select **OK** as shown in Figure 232.

Figure 232. Type and Retype Password

12. Select **OK** to close the User window.

Notice on User Properties

As you noticed, we did not discuss all pages of the User menu. They are explained in 13.2.5, “Setting Up User Object Properties” on page 434.

After the new user has been created, the user has no rights at the beginning. This means he only can access the drive for login and get the SYS_PUBLIC Volume unless access rights were defined in the User window (see Figure 231 on page 433; one of the pages in the settings notebook contains information of rights to files and directories). So far, the newly created user has the right to only read the default login drive at the preferred server that he/she logs in to. The preferred server can be configured in the requester's NET.CFG file.

13.2.5 Setting Up User Object Properties

So far, we have noticed that Novell NetWare 4.1 Server is working with an object-oriented method of administration. Network objects are ordered alphabetically in the graphical representation of objects.

In this section we want to explain the different pages of user properties in the User Object Settings Notebook. Table 36 gives you a short overview of the different pages and their functions.

Table 36 (Page 1 of 3). User Object Properties

Property Category	Description
Identification	This category is for documentation purposes and allows you to enter information helps an administrator to better identify the object. You can add telephone and fax numbers for easier support also.
Environment	The Environment page is used to see other properties of this user object. So you can view information not included in other pages of this dialog, such as the bindery properties and the network address. Also you see which properties did not migrate when this user was upgraded from NetWare 3.x to NetWare 4.
Login Restrictions	This page allows you to manage basic login restrictions like setting an expiration date for a user account and limiting the workstations that the user can log in to at the same time. Also you can disable the user account.

Table 36 (Page 2 of 3). User Object Properties

Property Category	Description
Password Restrictions	This setting controls how passwords for this user are handled. You can specify if the user can change his password. Also you can set that a password is required, when a password expires, if unique passwords are required, and how many grace logins are accepted. Grace logins allow the user to log in without changing the password, but after the specified number is exceeded, the user account will be locked.
Login Time Restrictions	In this page the administrator can define logon time restrictions for every user, meaning he/she can specify at which time the user is allowed to login. If the user logged in and the logon time expires, he/she would get a warning that his/her login time has expired. This will cause a forced logout after five minutes.
Network Address Restriction	Using these settings, the administrator can specify from which network addresses the user can log in.
Mailbox	The Mailbox location specifies the messaging server where this object's mailbox resides. The Mailbox ID displays an unique name that allows this object's mailbox to be located in the messaging database.
Foreign e-mail Address	Foreign e-mail Address specifies a Novell Directory Services object's mailbox that resides at a foreign e-mail system. For example, Novell Directory Services users can have e-mails delivered to UNIX machines that support SMTP (Simple Mail Transfer Protocol).
Print Job Configuration	In this page, the names of the different print job configurations that can be used are listed. For the administrator, it is an easy way to add additional new job configurations or to modify the parameters of existing configurations.
Login Script	Login Scripts are used to administrate the users. The Administrator can customize standard logon procedures for the user object.
Intruder Lockout	This page displays the user's status for his account after it has been locked.
Rights to Files and Directories	This page gives you the possibility to define trustee assignments to files and directories at a volume. For example, you can see the defined rights of objects to any file or directory, change this object's rights to files and directories, and see trustee assignments that this object has to a directory or file.

Table 36 (Page 3 of 3). User Object Properties	
Property Category	**Description**
Group Membership	The Group Membership page gives you information about which group the user object is a member of. The page is used to add, delete, or change membership settings.
Security Equal To	This page shows which objects this user object is equal to. All rights granted to the objects are listed on this page.
Postal Address	On this page you can add the individual postal address to a user object. This is helpful for administration and network documentation/management. For example, you can add the street address, post office box, city, state, postal-code, and copy this to a label.
Account Balance	To manage the credit on a user account, this page can limit the credit and also shows the credit status of a user. For example, this balance is deleted each time this user uses a network resource for which an accounting charge has been established. This affects the user's account only if the **Allow Unlimited Credit** field is not selected.
See Also	This page gives you a place to list names of objects related to the user's object. Supervisors or managers can use it to record related information about the object.

Although we are not discussing all these additional user object properties in detail, we will take them into consideration for a functional comparison. In the following sections we focus on the most important properties that are useful for administration work.

13.3 Managing NetWare User Objects

To control the network and administer it properly, we show you how to manage the base functions for user objects. Also we describe how to make the creation of user objects easier by using the user template function.

13.3.1 Deleting a User Object

Removing users from the Novell Directory Services should only be done when they are really not needed anymore because all the defined rights and settings will be lost also. Alternatively, you may only want to disable an user object by using the Login Restrictions page of the User Object Settings Notebook. To delete a user object do the following:

1. Start the **NetWare Administrator**, and select the organization from which you want to delete a user object.

2. Select the user object and open the object's pop-up menu by clicking on it with the right mouse button.
3. In the object's pop-up menu, select **Delete** as shown in Figure 233.

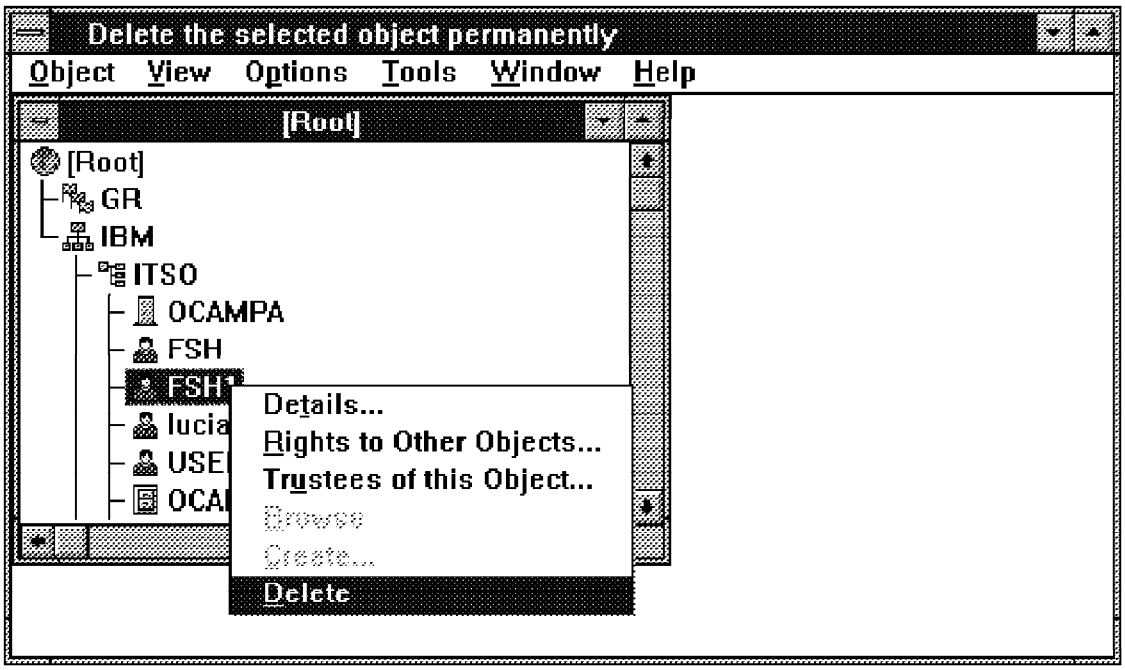

Figure 233. Delete User Objects

4. Click on **Yes** at the deletion window to confirm deletion of the user object.

13.3.2 Disable a User Object

Sometimes it is not useful to completely remove a user object from the Novell Directory Services since there might be user IDs that, for example, can be given to newly employees students when other students go back to school, or the project, that the other students were working on, has been finished.

It is easier to disable a user object before deleting it and setting it up again with all necessary rights and settings when needed again. To disable a user object from the Novell Directory Services you have to do the following:

1. Start the **NetWare Administrator** and select the organization from which you want to change a user's object properties.
2. Select the user object and open the object's pop-up menu by clicking on it with the right mouse button.

3. In the object's pop-up menu select **Details...** as shown in Figure 234 on page 438.

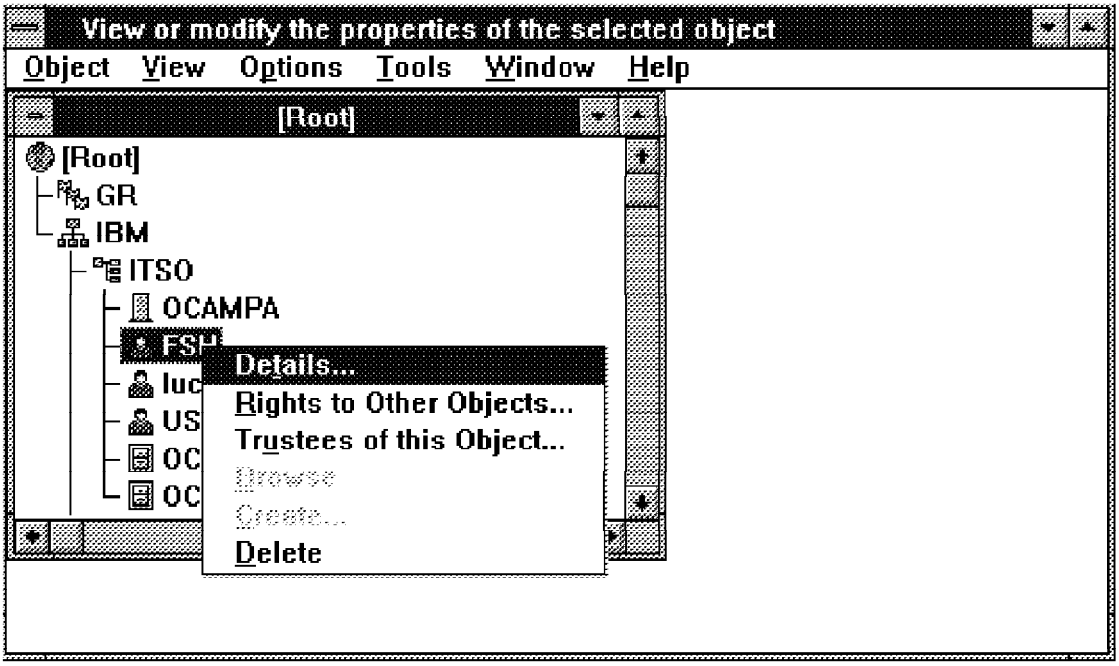

Figure 234. Change Object Details to disable user object from the Novell Directory Services

4. Select the Login Restrictions page and mark the checkbox for **Account Disable**.
5. Select the **OK** button to make the changes effective.

13.3.3 Adding a Home Directory to the User Object

The home directory is an optional directory resource on a server that is typically assigned to one user. In a workgroup environment home directories can even be shared among users whenever there is a necessity. Assigning home directories to users also helps also to increase data protection because the central backup would increase data integrity.

To add a home directory to the user object, do the following steps:

1. Start the **NetWare Administrator** and select the organization from which you want to change a user's object properties.
2. Select the user object and open the object's pop-up menu by clicking on it with the right mouse button.
3. In the object's pop-up menu select **Details...**.
4. Select the Environment page from the Settings Notebook.

5. On the Environment page, click on the **Browse** icon that resides to the right of the Home Directory input fields, as shown in Figure 235 on page 439.

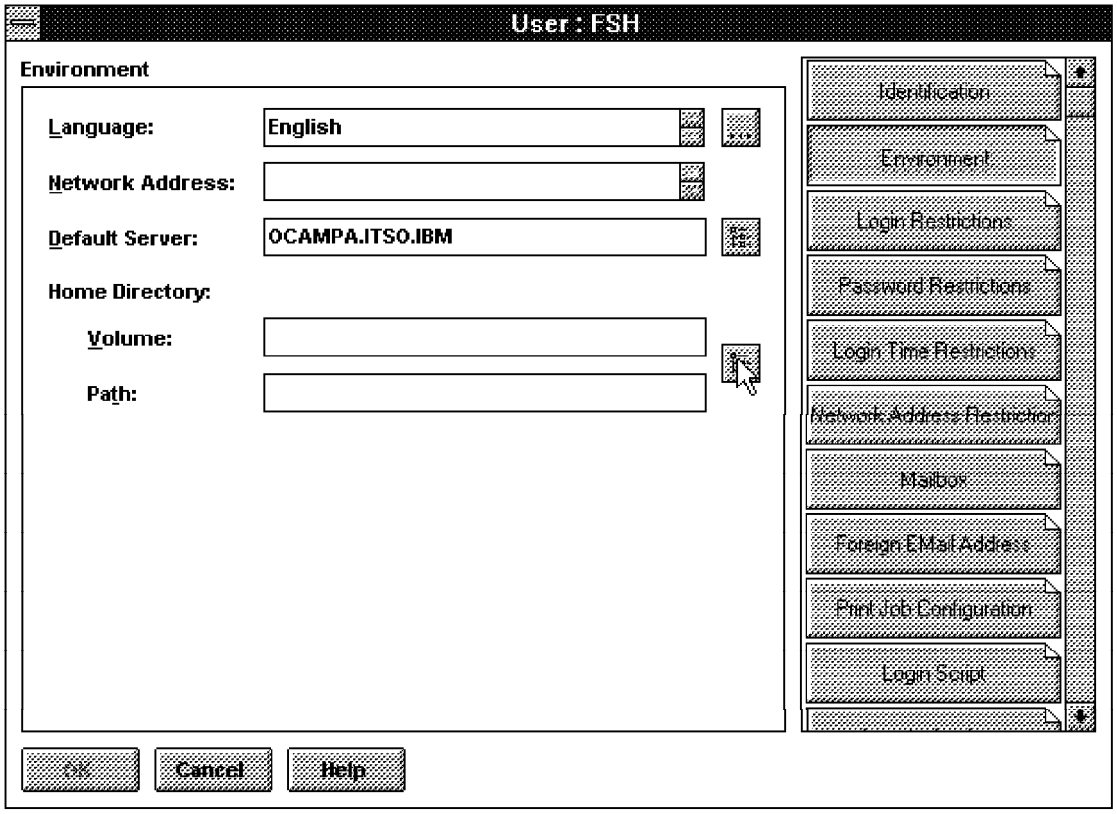

Figure 235. Browse Objects to Select Home Directory Path

6. In the list of Files and Directories, as well as in the list of Directory Context of the Select Object window, select the directory you want to use as the user's home directory, and select **OK** when finished. The Select Object window is shown in Figure 236 on page 440.

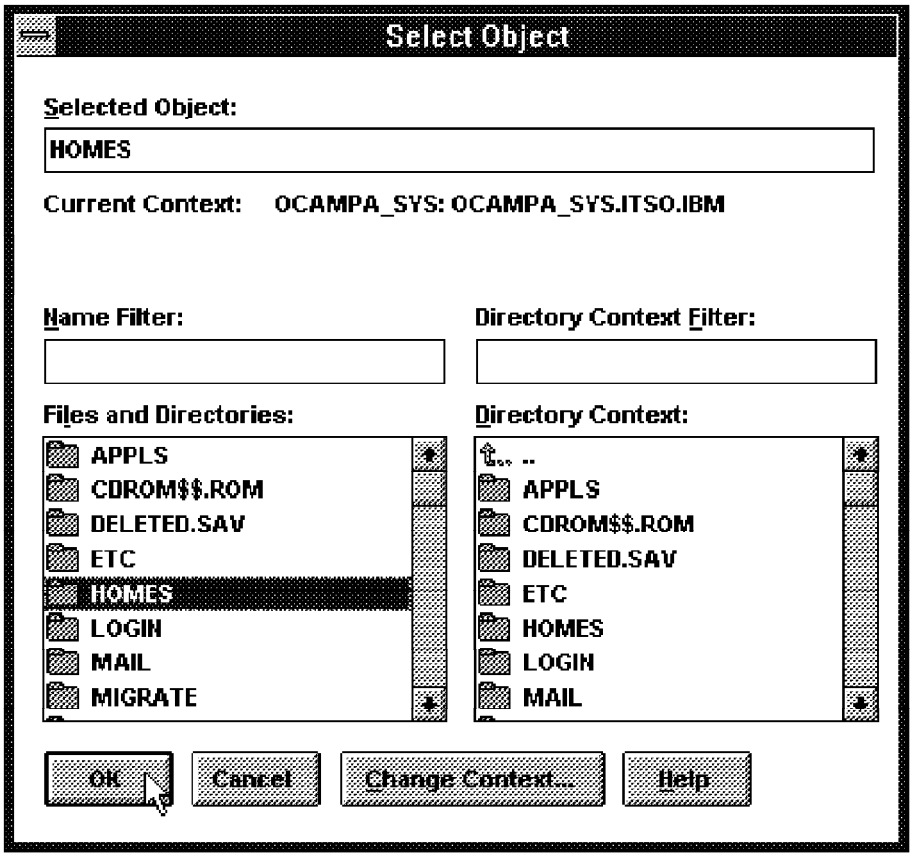

Figure 236. Choose the Directory for Home Directory Path

7. Also select the Default Server in the user's Settings Notebook by clicking on that icon next to the Default Server input field. An example of the filled-out user's Environment page is shown in Figure 237 on page 441.

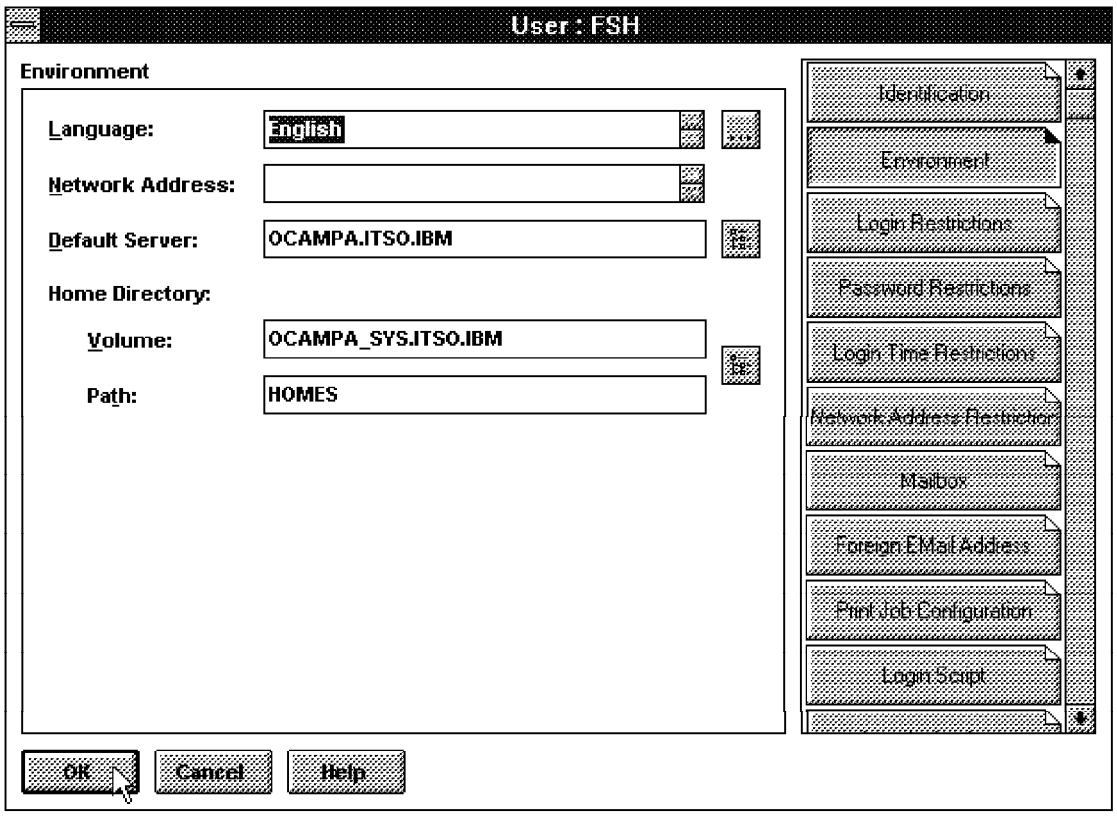

Figure 237. Add Changes to User Object

8. Select **OK** to make the changes active.

13.3.4 Creating User Related Objects

Just having a user ID on the network does not necessarily enable a user to really work with the network. The user also needs files and directories to work with. The administrator needs some functional objects to better organize the NetWare network. Because of this, we want to discuss the following user-related objects in the next section:

- User Templates
- Group
- Organizational Role
- Alias
- Profile
- Computer

Also the user-related objects like alias, computer, group, organizational role, and profile have required properties that are listed in Table 37 on page 442.

Table 37. Required Properties for User-Related Objects

Name of Object	Required Properties
Alias	Name and Aliased Object
Computer	Name
Group	Name
Organizational Role	Name
Profile	Name

13.3.5 Creating User Templates

Because of the advanced architecture of the Novell Directory Services, for large enterprises and internationally operating companies, Novell Directory Services offers the possibility to differentiate between organizational units and organizations by allowing you to set up user templates for each segmented area. This allows you to create multiple users with similar profiles by using user templates.

To do so, follow these steps:

1. Start **NetWare Administrator**.
2. In the [Root] window, select the **Organization** container you want to create the user template for.
3. Select **User Template** from the Object pull-down menu as shown in Figure 238 on page 443.

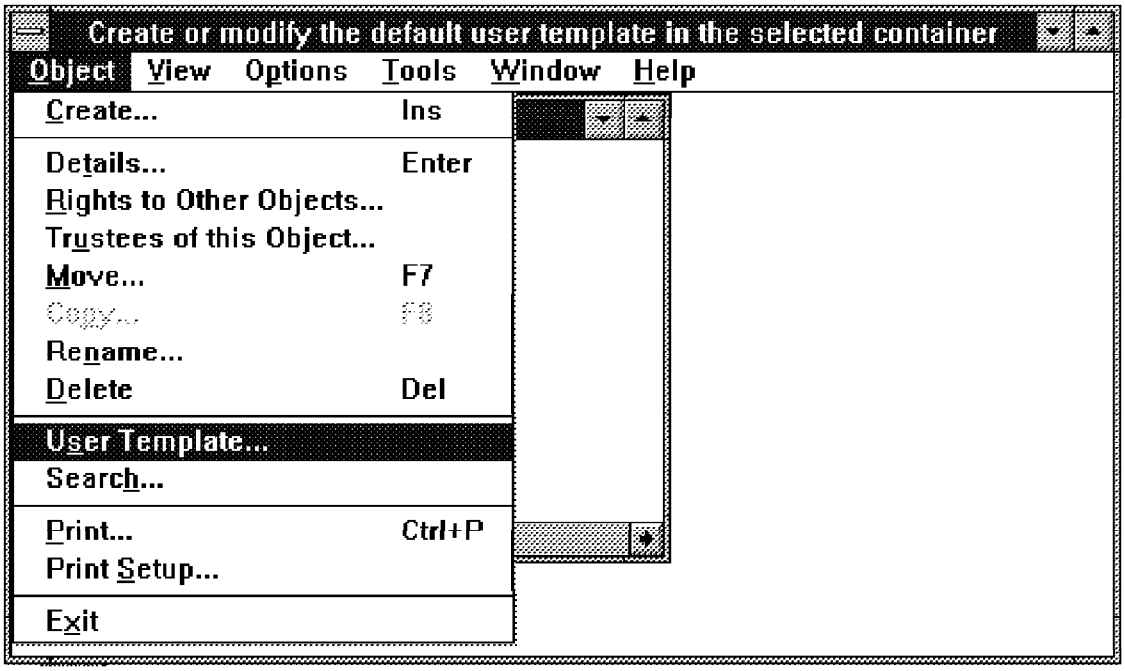

Figure 238. Creating a User Template for an Organization Container

4. Type in general information for the user objects. Look for additional information at Table 36 on page 434 to find a detailed description about the different user menu pages. For example, mark the checkbox for **Require a Password**.
5. Select **OK** when finished to close the dialog box.

To create a user object using the template, you only have to mark the checkbox for **Use User Template** in the Create User window (see Figure 229 on page 431). Then provide a Login Name and additional information as necessary.

13.3.6 Creating a Group

For an administrator, one of the important things is to organize his/her users into groups that represent departments, groups within departmental organizations, or project groups. Groups are perfect for putting people together that need to have the same rights and accesses to resources.

Making user objects members of a specific group provides them access to a common network service. So all members of the group have access rights to a similar set of network resources. The membership of a user in a group is defined in the user's parameters settings on the Group Membership page.

To add users to a group, the administrator can use two paths:

1. By adding user objects to the Member page in the Group properties.
2. By defining a group in the Group Membership page of the user's object properties.

The first step, however, is to create a group. To do so, follow these steps:

1. Start the **NetWare Administrator**.
2. Select **Create** from the Object's pull-down menu.
3. From the Class of New Objects list in the New Object window, select the **Group** object (see Figure 224 on page 419) and double-click on it. Select **Create**. The Create Group window will be opened for you as shown in Figure 239.

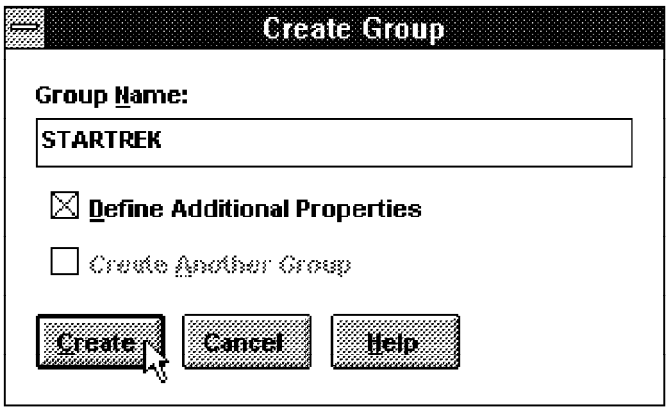

Figure 239. Create a New Group Object

4. Type in the Group Name and mark the **Define Additional Properties** checkbox. If you want to create more than one group, also mark the **Create Another Group** checkbox. The Create Group window will remain opened for you.

   Notice on Creating a Group Object

   As described in Table 37 on page 442, to create a group you only have to define the name, but here we want to go into more detail about showing how to add users to that created group.

5. Click on the **Create** button to create the group. The Additional Properties menu will be opened for you.
6. Type in additional information and add members into the **Member** page.
7. For easier identification, type in information at the **Identification page** and select **OK** to close the dialog box.

To complete the overview we want to give you an overview of additional group object properties and their functions in Table 38 on page 445.

Table 38 (Page 1 of 2). Group Object Properties

Object Page	Property	Descriptions
Identification	Name	The Novell Directory Services name that distinguishes the group object.
	Other Name	Information that helps you to identify or refer to this group object. This is only for your help.
	Owner	Identifies the group owner, if there is any.
	Description	This field could hold a detailed description of the group object.
	Location	Can hold information about the group object's location.
	Department	Holds the department name of the group object in the company or Novell Directory Services organization.
	Organization	Contents the organization to which the group object belongs.
Member	Group Members	The contents of this is a list of the members of this group object. All options that are selected and rights to files or directories are affecting these users.
Mailbox Information	Mailbox location	Provides the messaging server where the mailbox is stored.
	Mailbox ID	Contains the name of the group object mailbox.
	Foreign E-Mail Address	Describes the unique e-mail address for this group object.
	Foreign E-Mail Alias	To communicate with other messaging services, you can add an e-mail alias name.

Table 38 (Page 2 of 2). Group Object Properties		
Object Page	**Property**	**Descriptions**
Rights to Files and Directories	Volume	This field shows the volume object name that this object has a file system trustee assignment to. To add and locate a specific volume object, select the **Find** button. Also use the **Hide** button to clear volume objects from the volumes list.
	File and Directories	After selecting a volume object, this list shows the specific directories and files that have trustee assignments. To include trustee assignments to more directories and files use the **Add** button.
	Rights	To see the specific rights this object has to the file system, select the files or directories you want to see. To add additional trustee assignments to more files and directories use the **Add** button.
See Also		This page includes reference information that is related to this group.

Find in 13.3.22, “Access Rights Administration” on page 479 information about the different rights for the file system and directories.

13.3.7 Creating an Organizational Role

Positions that can be filled with more than one person are called an organizational role. Organizational roles can be filled with container administrators, team leaders, network backup operators, or printing specialists. You can grant specific rights to those who are selected as occupants of these positions, like the group object.

To set up an organizational role object, follow these steps:

1. Start the **NetWare Administrator**.
2. In the [Root] window, select the container object you want to create an organizational role object for, and press the right mouse button to open the object's Context menu.
3. Select **Create** from the object's Context menu and in the Class of New Objects list in the New Object window, double-click on the **Organizational Role** object, or select it and choose **OK**.
4. Enter the object name, which is the only mandatory entry as described in Table 37 on page 442.
5. Mark the **Define Additional Properties** checkbox and select **Create**.

Note on Creating an Organizational Role

The properties for this object are similar to those of the user and group objects. The crucial property information is contained in the **Occupant** field.

6. On the **Identification** page click on the icon next to Occupant as shown in Figure 240.

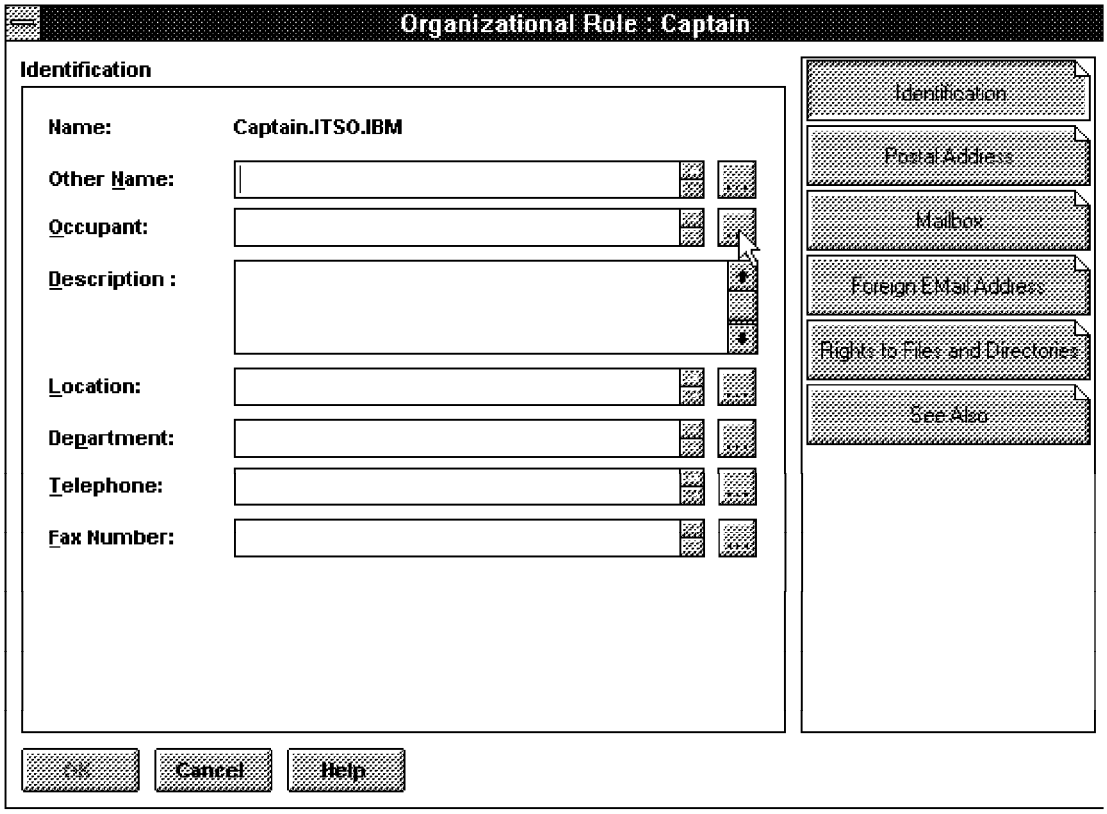

Figure 240. Edit Occupant Field in Organizational Role Object

7. Click on the **Add** button, and select the right location and user object by going through the Objects list as well as the Directory Context list within the Select Object window as shown in Figure 241 on page 448.

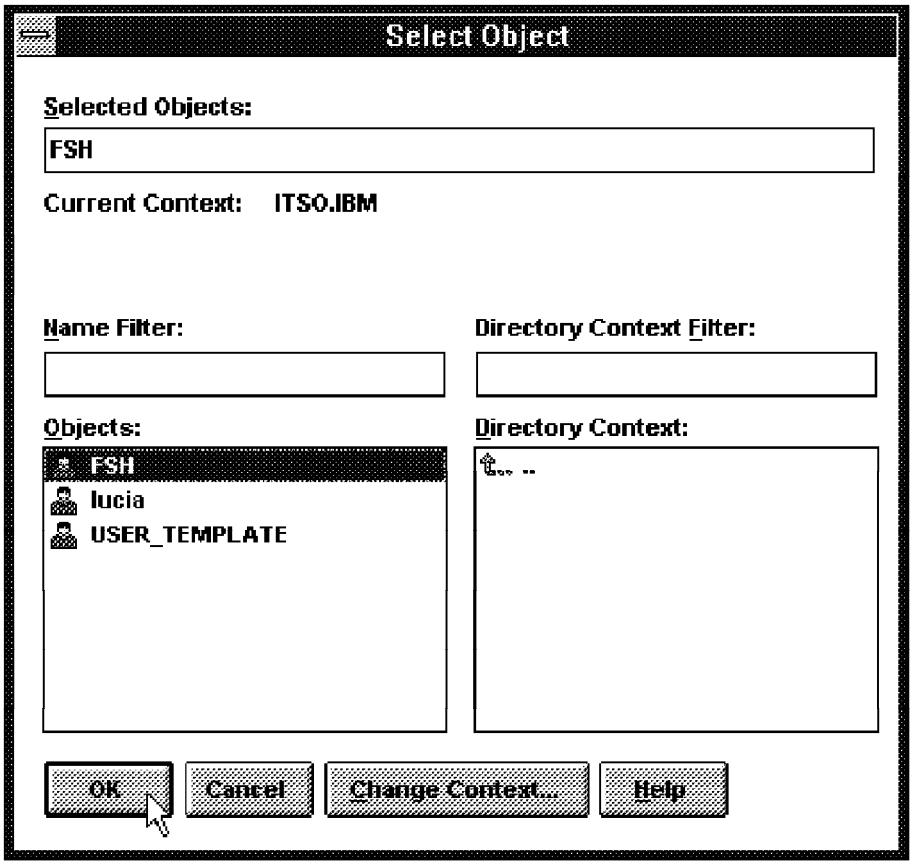

Figure 241. Adding User Objects to the Organizational Role Object

8. After you have selected the User objects you want to add, press on the **OK** button.
9. Select **OK** in the Occupant dialog box to stop adding occupant properties to the organizational role object.
10. Choose **OK** to end the setup of the organizational role or to add additional properties to the object.

13.3.8 Modifying the Organizational Role Object

You not only can add organizational role objects, you can also change, delete, and move them. You have to have NetWare Administrator started to perform these tasks. In the [Root] window, select the **Organizational Role** object and double-click on it. You can make all necessary changes to it, delete it, or move it to a different position in the directory tree.

13.3.9 Creating an Alias Object

An alias object points to another object that is somewhere else in the directory tree and makes it appear as if that object actually exists in the directory tree where the alias object is. Having aliases prevents the user

from typing in the object's full name and context when the object resides in another part of the directory tree. In particular servers, volumes, printers, and print queues are often accessed by users from all parts of the directory tree. Because an alias is just a pointer, it has no properties of its own except its name. An alias object is not required to have the same common name as the object to which it points.

To create an alias object, do the following steps:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree in which you want to place the alias object, and click the right mouse button on the container object to get the object's Context menu.
3. Select **Create**, or press the **Insert** key.
4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Alias**.
5. Press the **OK** button. The Create Alias window will be opened for you.
6. Type in a name that fits your strategy and click on the the browser icon next to the Aliased Object field as shown in Figure 242.

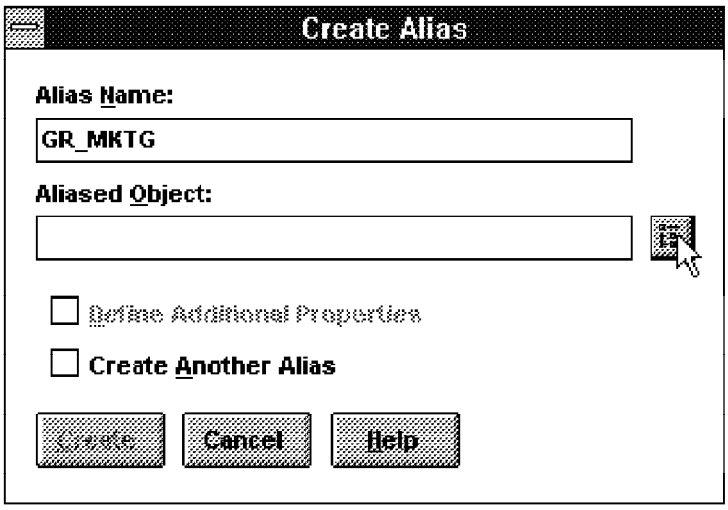

Figure 242. Create Alias Window

The Select Object window will be opened for you.

7. In the Select Object window, by browsing through the directory tree, select the aliased object as shown in Figure 243 on page 450.

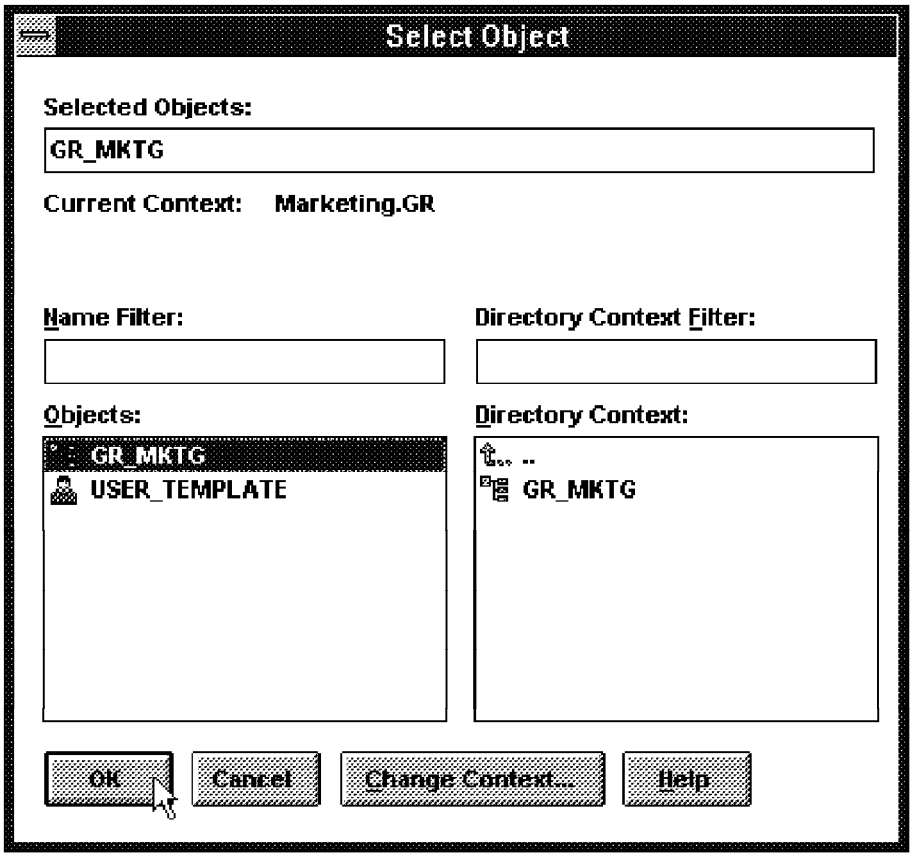

Figure 243. Select Object Window

8. Select **OK** when finished selecting the aliased object.
9. At the Create Alias window, select **Create** to create the alias object.

13.3.10 Creating a Profile Object

A profile object contains a login script that is utilized by users who need to use similar Login Script commands. This helps the network administrator to customize the login procedure and automate the needed steps for the users that he/she has to manage.

If a profile object is listed in the login script page of a user object, it is executed after the global login script of the user's container object but before the users individual login script.

For a user to execute a profile login script, you must select the profile login script in the login script Property page of the user object. Also the users must have Read rights to the profile object's login script property for the profile login script to execute. To create a profile object, perform the following steps:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree for which you want to create the profile object, and click the right mouse button on the container object to get the object's Context menu.
3. Select **Create**, or press the **Insert** key.
4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Profile**.
5. Press the **OK** button. The Create Profile window will be opened for you.
6. Type in a name, mark the **Define Additional Properties** checkbox, click on **Create**.
7. The object's Settings page lets you define different parameters, Type in information as needed, select the **Login Script** page and enter the needed commands. The commands are described in Table 39.
8. Click on **OK** to save the information and end the creation of a profile object.

To understand the syntax of login script commands we show you the commands with the needed parameters and a description of what they are doing. To work effectively with the commands you should pay attention to the following simple rules:

- Each login script command must be written on its own line.
- Blank lines are ignored.

The following list shows commonly used login script commands. There are other commands described in Supervising the Network.

Table 39 (Page 1 of 2). Login Script Commands and Syntax

Command	Parameter	Example	Description
#	Command to execute	#DIR	Executes an external command and returns to the next line of the login script.
ATTACH	Directory tree/server to connect	ATTACH ITSO2/SALES	Connects to a NetWare server in a different Directory tree.

Table 39 (Page 2 of 2). Login Script Commands and Syntax

Command	Parameter	Example	Description
BREAK	ON/OFF	BREAK ON	Allows the use of <Ctrl><Break> or <Ctrl><C> to stop the login script. Default is OFF.
CLS			Clears the display screen.
CONTEXT	Context to change to	CONTEXT USERS.SALES	Changes to a specified context.
DISPLAY	filename	DISPLAY NOTICE.TXT	Displays a text file.
DRIVE	drive letter	DRIVE V:	Specifies the default drive.
EXIT	filename	EXIT MENU ACCOUNT	Terminates execution of a login script.
FDISPLAY	filename	FDISPLAY README.DOC	Filters out control characters and display only ASCII character values 1 to 126.
FIRE PHASERS	number of times	FIRE PHASERS 3	Makes a sound to alert you to certain conditions.
GOTO	label	GOTO LOOP	Skips part of the script and jumps to another label. Usually used with an IF conditional.

To add the now defined login script to a user account, do the following:

1. Start the **NetWare Administrator**.
2. Highlight the user object in the directory tree for which you want to add the Login Script, and click the right mouse button on the user object to get the object's Context menu.

3. Select **Details...**.
4. The User window opens for you. Select the **Logon Script** page and use the browser function to select the profile object an shown in Figure 244.

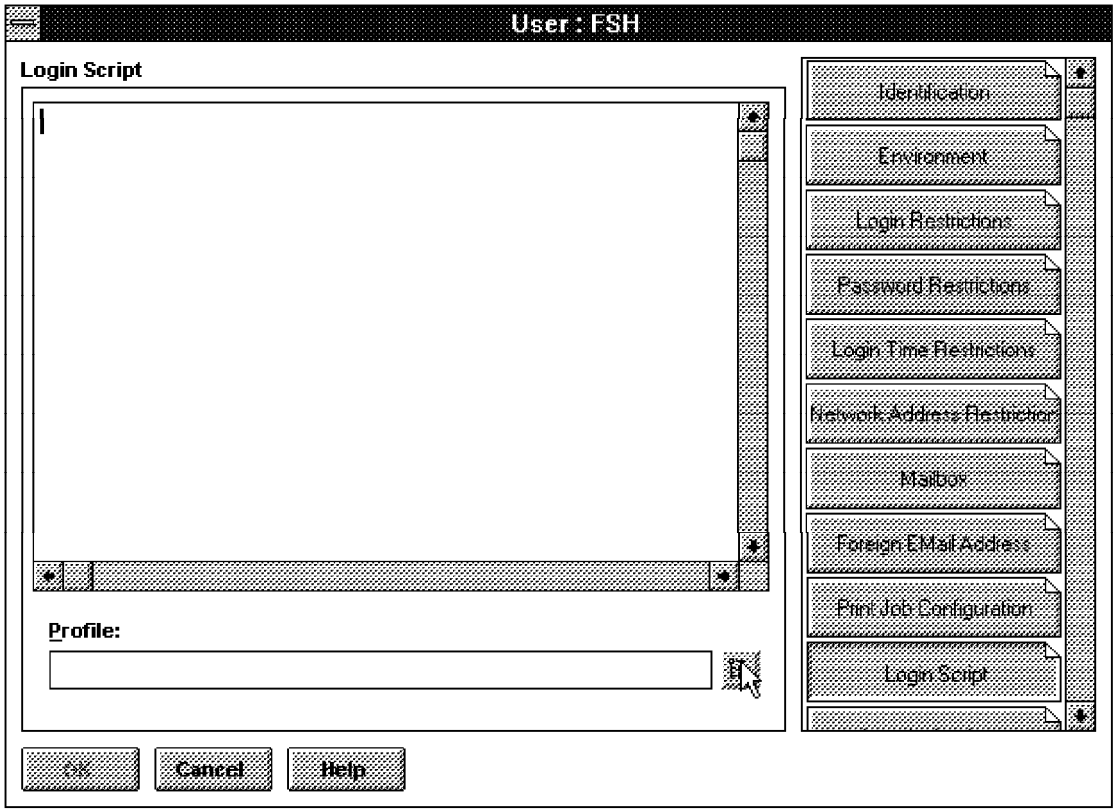

Figure 244. Adding a Login Script to a User Object Using the Browse Icon

The Select Object window will be opened for you.

Browsing through the Objects and Directory Contents lists of the Select Object window, select the profile object as shown in Figure 245 on page 454.

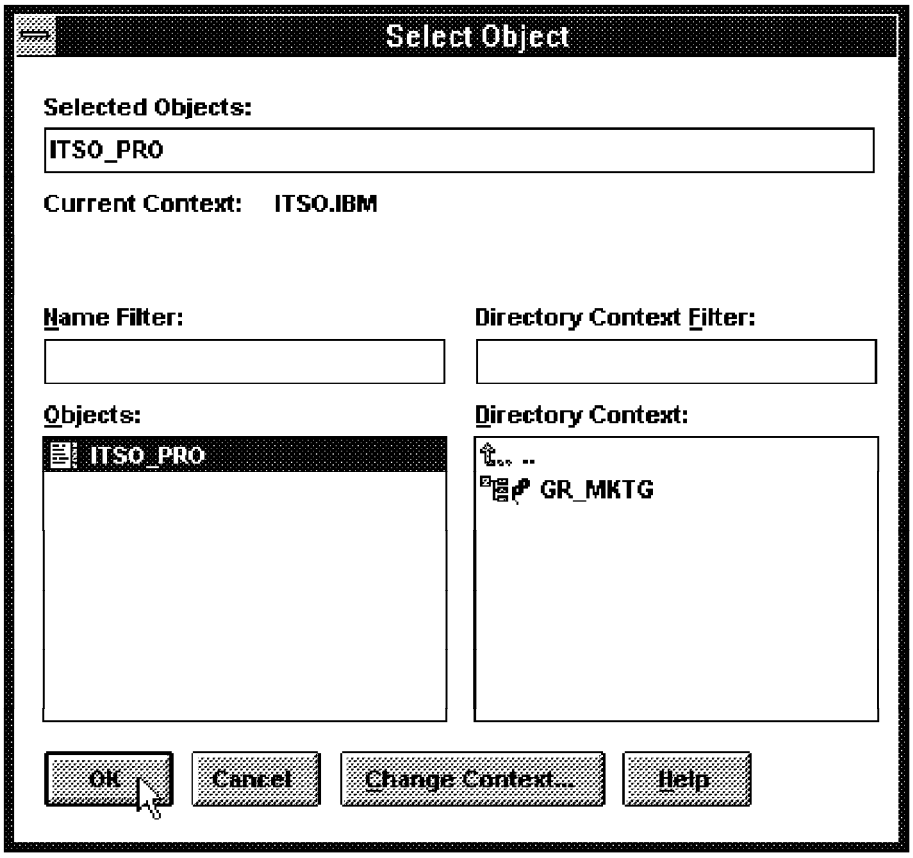

Figure 245. Select the Profile Object from the Objects List.

5. When finished, press the **OK** button to make the changes active.

To make a Login Script active for organizational container, like in our case the ITSO organization unit container, use the additional properties menu and edit the commands to the Login Script Page, as done previously. You can use the same commands as described in the Table 39 on page 451.

13.3.11 Creating a Computer Object

A Computer object represents a computer on the network. This object is for informational purposes only. You can enter description information as well as configuration information, like the network address.

To create a computer object, complete the following steps:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree in which you want to place the computer object, and click the right mouse button on the container object to get the object's Context menu.
3. Select **Create**, or press the **Insert** key.

4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Computer**.
5. Press the **OK** button. The Create Computer window will be opened for you.
6. Type the name for the Computer object in the Computer Name field, and mark the **Define Additional Properties** checkbox.
7. Choose the **Create** button. The Additional Properties window will be opened for you.
8. Type in the information to describe the computer as shown in Figure 246.

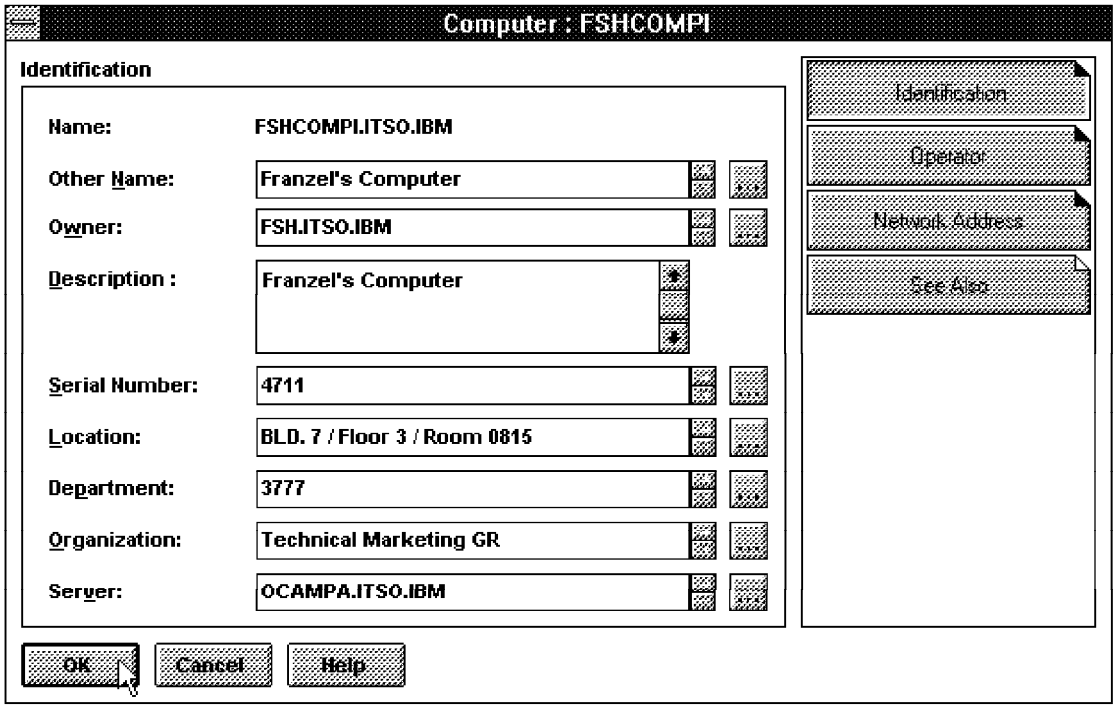

Figure 246. Add Computer Object Properties

9. When finished, choose **OK** to accept entries.

13.3.12 Conclusion on User and User-Related Objects

Because of the good implementing of object and directory technology and the detailed and distinguishable management objects, NetWare is a very good, very manageable network operating system platform.

The detailed properties of the user object and the user-related objects help a network administrator to manage his/her network better. Be aware of the fact that maintaining an up-to-date directory-based network can be

time-consuming. On the other hand, it provides an accurately managed NetWare environment and is very well documented.

Using the NetWare Administrator, the management of the objects is very convenient. All functions included in the NetWare Administrator are object oriented.

13.3.13 Creating and Managing User Objects with UIMPORT

`UIMPORT` is designed to handle two type of situations:

- Adding a large number of user objects at the same time
- Providing a link between your database manager and NetWare 4.1

Here we emphasize adding a large number of user objects at the same time. You need to create two files when you want to use `UIMPORT`:

- Data file — An ASCII file that contains information that is actually imported to the directory tree.
- Control file — A file that lists the fields or property categories contained by the data file and specifies the options that you must invoke when you import the information.

  These options include whether to create a home directory, whether to use the settings of the USER_TEMPLATE user, and which directory context to use as the new object's directory tree location.

13.3.13.1 Constructing the Data File

Following is an example of a data file that lists the user's object name, the user's last name, the user's full name, postal address, telephone number, and fax number.

```
"OSCAR","CEPEDA","Oscar Cepeda","ITSO/2834","678-5634","678-6931"
"UWE","ZIMMERMANN","Uwe Zimmermann","ITSO/2834","678-6007","678-6931"
```

Figure 247. UIMP.DAT File for the UIMPORT Utility

Note: To allow blanks, you need to put values in quotation marks.

13.3.13.2 Constructing the Control File

A control file is divided into two sections. The first section starts with the heading `IMPORT CONTROL`, the second starts with the heading `FIELDS`.

A typical control file looks like the following:

```
IMPORT CONTROL
  SEPARATOR = ,
  USER TEMPLATE = Y
FIELDS
  GIVEN NAME
  LAST NAME
  FULL NAME
  DEPARTMENT
  TELEPHONE
  FAX NUMBER
```

Figure 248. UIMP.CTL File for the UIMPORT Utility

13.3.13.3 Running UIMPORT

After building your data and control files, you are ready to use `UIMPORT` to import the information into the data file. Type `UIMPORT` followed by the name of the control file and the name of data file, and then press Enter.

```
UIMPORT C:\UIMP.CTL C:\UIMP.DAT
```

13.3.14 Understanding Subdirectory Design

One of the first considerations when designing a subdirectory structure for your network is to decide which types of files you plan to use. The next step is to create a directory to hold these files. The following are the basic types of files:

- Application program directories
- Shared data directories
- Individual user directories

In the first part of this section we will shortly explain the three different directories and their contents.

13.3.14.1 Individual User Directories

Individual user directories are discussed a lot in the network community. The types of files stored in the user's individual directory varies. These directories are often referred to as home directories. The individual user directory should give the user a place to store his individual files, data, and programs that should be accessible all over the network for his user account.

The advantage of having individual user directories is that they help to minimize possible data loss which may result from the broad disparity in backup procedures.

13.3.14.2 Shared Data Directories

As the name suggests, shared data directories are designed to exchange and store data that is used within the whole company, within defined departments or within projects.

The access management to shared data directories is done by assigning rights trustee rights to them. These rights can be assigned on a container, group, organizational role or user level. Trustee rights for shared data directories typically include read, write, create, erase, and file scan. Individual users or a group of users are granted with supervisor rights to the directory. The individual trustee rights to files or directories are dependent on the task a user or a group has to fulfill.

13.3.14.3 Application Directories

Application directories contain executable and other application program support files. They are similar to shared data directories, but should be organized differently for reasons of security, integrity, and backup. Application files need to be backed up only when and application has been upgraded or when doing a full system backup. User and shared data directories should be backed up at least once a day.

Trustee rights assignments to application directories often limit users only to read and file scan, which are the minimum rights needed to locate and run applications. Be aware that some applications might require additional rights because of writing profile files to the disk. The rule should be to grant only the rights effectively needed to run the application and to guarantee data integrity.

13.3.15 Creating Directories

To manage directories, you have to select the volume object you want to make changes to. In our example, this is the OCAMPA_SYS volume.

To view directories and files of a volume object, you can double-click on the volume object icon. The tree expands to show directories and files from the root of the volume object. Directories appear as green folders; files appear as white pieces of paper with their upper right edges folded over. To step down the directory structure you can click on the subdirectories.

Figure 249 on page 459 illustrates a volume object and its subdirectories.

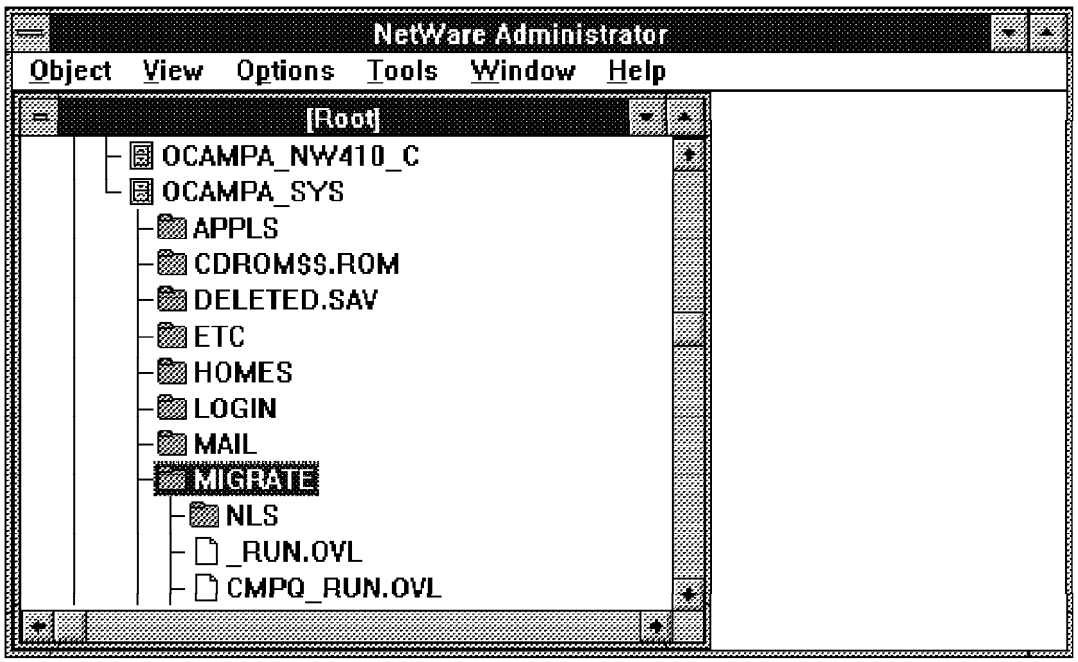

Figure 249. Example of the Subdirectory Structure of a Volume Object

Before you can create a directory using the NetWare Administrator, you must first select the parent to the new directory. This can be accomplished by completing the following steps:

1. From the [Root] window, select the Container that contains the Volume object.
2. Choose the Volume object.
3. Select the Parent Directory.
4. Open the object's Context menu by clicking on the object with the right mouse button.
5. Select **Create** from the object's context menu.
6. Enter the **Name** of the directory you want to create and mark the **Define Additional Properties** checkbox.
7. In the Properties page you can directly add trustee rights to the directory, set a size limit, add ownership of the object, and set attributes as shown in Figure 250 on page 460.

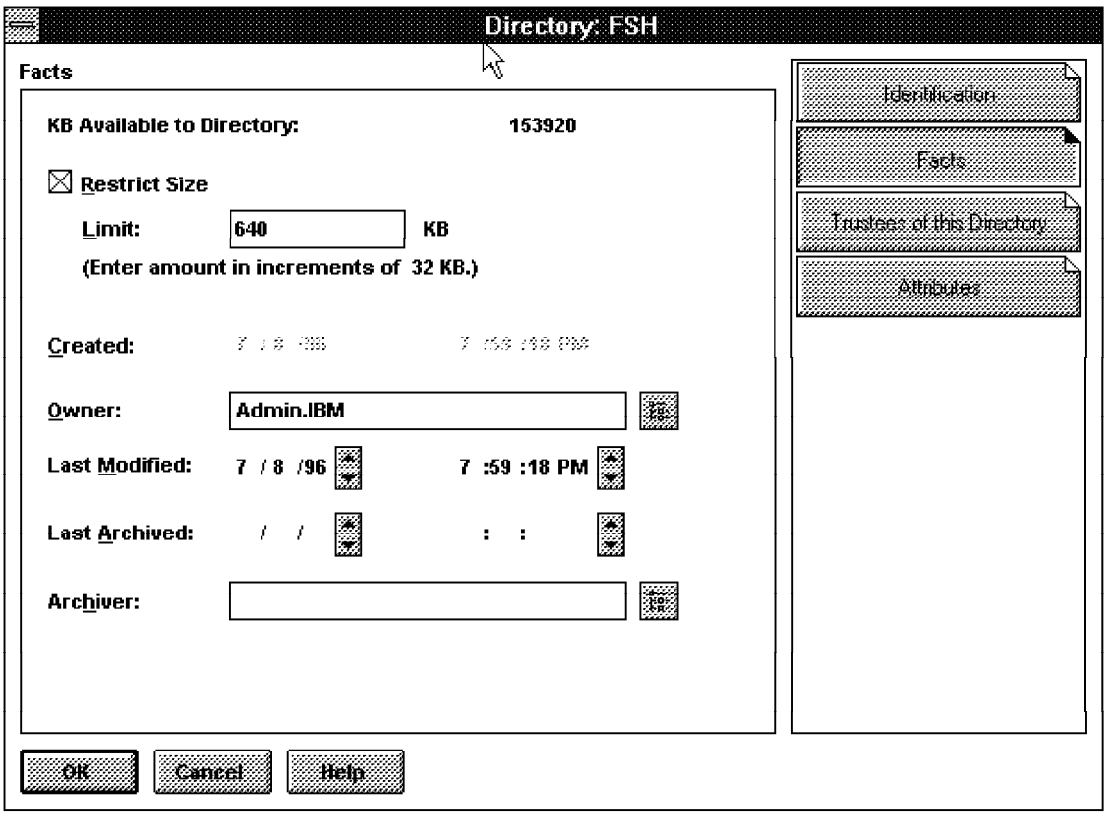

Figure 250. Add Properties in the Create Directory Dialog

8. After you have finished setting up the properties, select **OK**.

Creating and using directories can also be done from the command prompt by using the standard commands, like make directory (`MD`), change directory (`CD`), and remove directory (`RD`). NetWare enhances those commands with an additional utility called `RENDIR` that has the function to rename directories. Especially for creating a large amount of subdirectories these commands can be used with loop commands within a batch file. This is nothing special; these are only operating system commands.

Also there is the DOS interface of NetWare called `FILES`, which we personally find very clumsy to use. The best choice in this case is the NetWare Administrator program because you are using the same interface for all actions and you can also add properties to directories that are described in Table 40 on page 461. That makes it easier to manage large directory and server structures.

Because of its ability to make size restrictions to all subdirectories, it helps a network administrator to manage his/her tasks and space requirements. This means you, as an administrator, can limit individual user directories as well as application or data share directories. The modification and archive

control is very good also and helps to view all the information within the detailed view of directories as well as files. Also there is ownership information for directories and modifier information for files. In addition, all the trustees and attributes can be managed with the NetWare Administrator as shown if Table 40.

Table 40. Properties of Directories

Page Title	Description
Identification	On the Identification page you find a description of the directory you have selected. It holds information about the directory name and which name spaces are available on the volume in which this directory is located. The following name space formats are supported: • DOS, used by IBM PC and compatibles, including Microsoft Windows. • Macintosh, used by all Apple Macintosh computers • FTAM, standard file system used by mainframe and minicomputers. • NFS (Network File System), used by UNIX and RISC workstations. • OS/2, used by OS/2 on IBM PC and compatibles for HPFS. Files and directories can have long names and properties.
Facts	On the Facts page you find statistical information about this directory. So you can change the ownership of directories, view date and time of creation, view date and time of last modification, and date and time when the directory was last archived, and you can also restrict the size that could be used.
Trustees of this Directory	This page shows trustees of the directory or file selected. It shows who is a trustee of this directory or file and allows you to change trustee rights. In the Access rights you can view and change a trustee's access rights. The Effective rights button shows you the rights that any object can use to access this file or directory in a dialog box. The inherited rights filter allows you to change and view the inherited rights of this file or directory.
Attributes	The Attribute information about this directory or volume is shown in the attributes page. These rights could be viewed, deleted, and changed.

After the Directory is created and the proper trustee rights are assigned, every user with the necessary rights can access the subdirectory and do the defined things, like:

- Read
- Write
- Create
- Modify
- Erase
- File Scan
- Access Control

Also there is an additional access right that gives the user supervisor rights to the object. The Supervisor right is the highest right level possible. Detailed information about trustee rights and right levels can be found in 13.3.22, “Access Rights Administration” on page 479.

Administrators can access directories by using the NetWare User Tools or by using the command prompt and issuing the `MAP` command:

```
MAP X:=OCAMPA_SYS.ITSO.IBM:HOMES\FSH
```

13.3.16 Creating a Directory Map Object

To easily assign drive mappings, and for management purposes, it is possible to work with Directory Map Objects. The Directory Map Object can be used in login scripts when assigning drive mappings. Directory Map Objects provide the opportunity to change the location of the application directory without having to change a series of login scripts and drive mapping statements.

For example, if your company uses Lotus WordPro as a word processor, you might install the application in a directory called WordPro, and create a Directory Map Object called WordProcessing and assign it to the WordPro directory. Login scripts or batch files might then be created using the Directory Map Object WordProcessing to create drive mappings. You can then change the directory of the Directory Map Object instead of changing all login scripts and inform users that have their own login procedures that they have to change them.

To give a user access to the created Directory Map object, you have to grant him/her at minimum "Browse object rights" and "Read property rights". To create a Directory Map object, do the following:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree in which you want to place the Directory Map object, and click the right mouse button on the container object to get the object's Context menu.

3. Select **Create** or press the **Insert** key.
4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Directory Map** as shown in Figure 251.

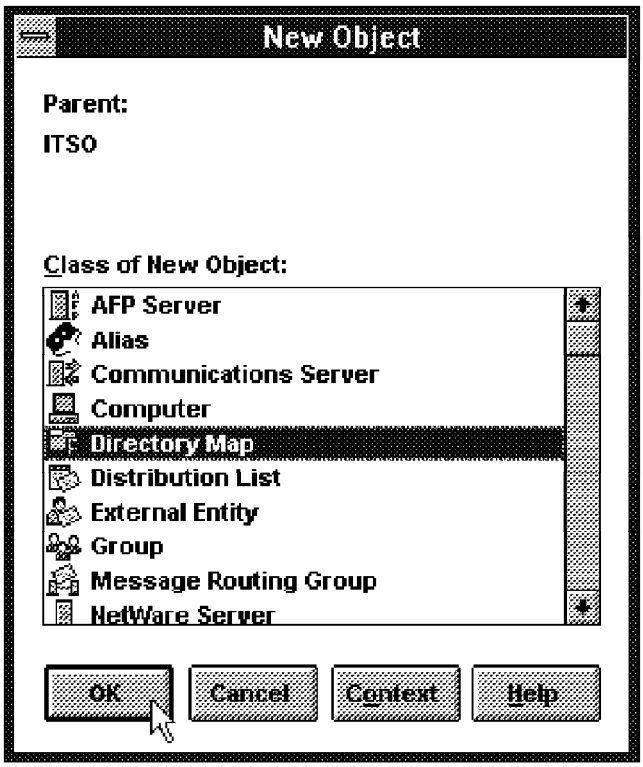

Figure 251. Select the Directory Map Object to Create

5. Press the **OK** button. The Create Directory Map window will be opened for you.
6. Type the name for the Directory Map object, mark the **Define Additional Properties** checkbox, and click on the **Browse** icon to select the volume and directory you want to map. The Select Object window will be opened for you as shown in Figure 252 on page 464.

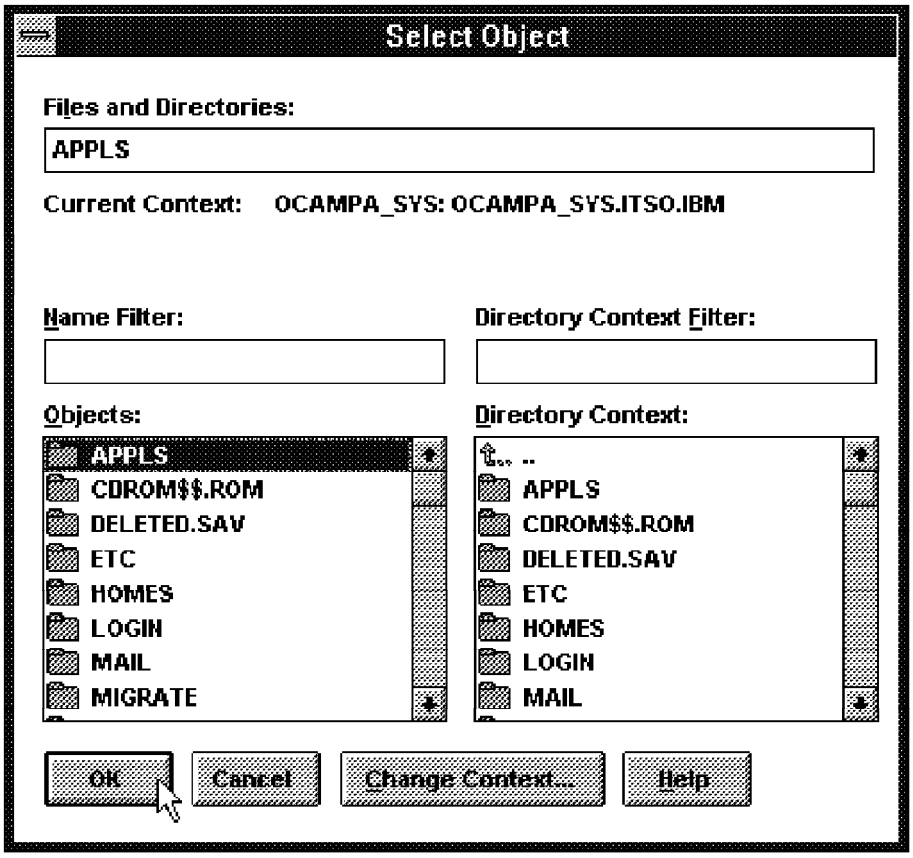

Figure 252. Select Volume and Directory for Directory Map

7. Press on **OK** when finished browsing.
8. Choose the **Create** button at the Create Directory Map window.
9. Type in the additional properties information and select **OK**. For additional information about properties, see Table 41.

Table 41. Properties of Directory Map Objects

Page	Description
Identification	This is the description page of the Directory Map Object. Here you can change and view the description as well as the volume and path of the Directory Map Object.
Rights to Files and Directories	All trustee assignments that this object has to files and directories on a volume are shown here.
See Also	To record information about this object and objects that are related to it, the supervisor or manager can use this page.

The information to the trustee and rights modification and management is discussed in 13.3.22, “Access Rights Administration” on page 479. There are the different trustee rights between objects discussed.

After you have added a Directory Map Object, you can also modify the login scripts with the following statement:

```
MAP INS S2:=.CN=APPLS.OU=ITSO.O=IBM
```

13.3.17 Creating Print Objects

To create a new printer object you have first to create a print queue object. To set up printing, it is very important to remember the links between the three basic elements of printing as shown in the following illustration:

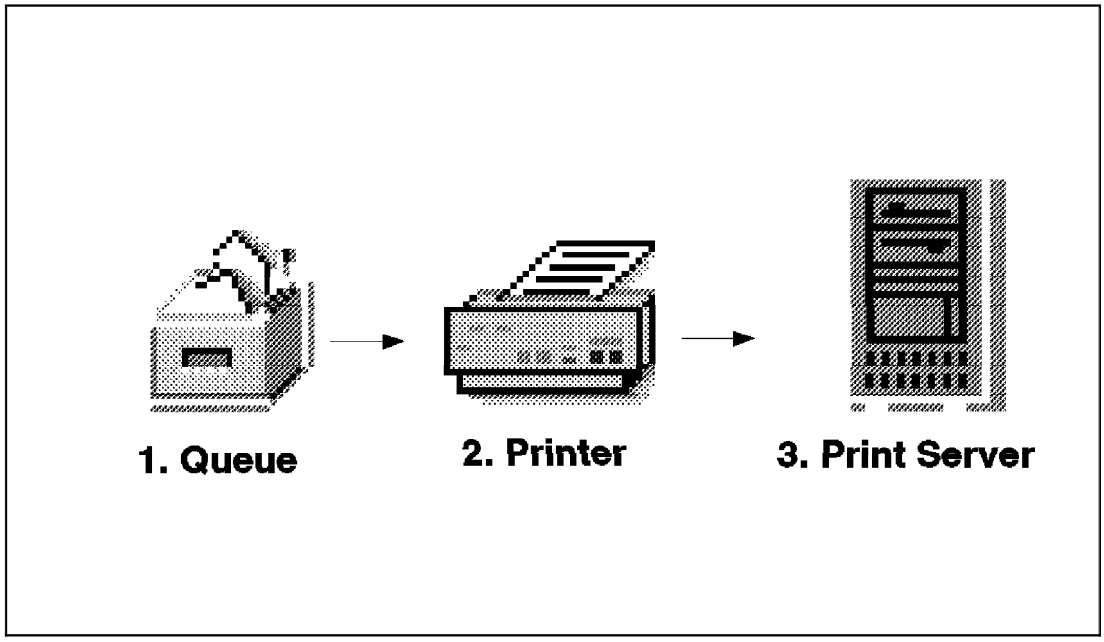

Figure 253. Links Between the Three Basic Elements

As shown in Figure 253, you first have to create a print queue object, then the printer object, and finally the print server object. To make all the associations complete, you have to assign the queue to a printer and the printer to a print server afterwards. In the following section, we show you how to create and connect these three objects together to ensure smooth printing.

13.3.18 Creating Print Queue Object

The print queue is the first step when distributing print services. A queue, where all print jobs are waiting until the printer is ready to receive data, is dedicated to a printer. Also it is easy to control the print service via the queue object.

To create a print queue object by using the NetWare Administrator utility, do the following steps:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree in which you want to place the print queue object, and click the right mouse button on the container object to get the object's Context menu.
3. Select **Create** or press the **Insert** key.
4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Print Queue**.
5. Press the **OK** button. The Create Print Queue window will be opened for you as shown in Figure 254.

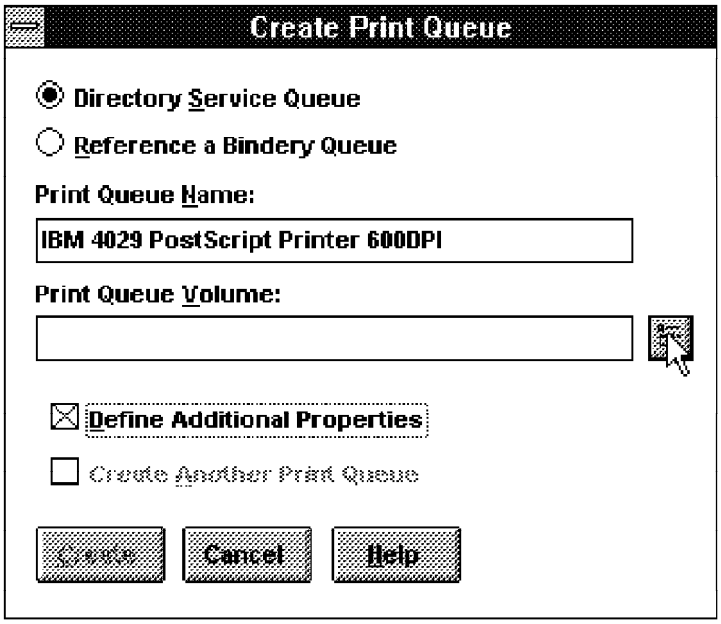

Figure 254. Browse Volume for Print Queue Volume Selection

6. Select the radio button for **Directory Service Queue**, type the name for the print queue object, mark the **Define Additional Properties** checkbox, and click on the **Browse** icon next to the Print Queue Volume field to select the volume that should contain the temporary print jobs until printing. The Select Object window will be opened for you as shown in Figure 255 on page 467.

   Note: At the Create Print Queue window (Figure 254), the **Reference a Bindery Queue** radio button is used if you need to manage a print queue on a NetWare 3.x Server, which is a bindery-based file server. If this is your case, enter the name you plan to use to represent the bindery-based queue in your Novell Directory

Services, then enter the name of the server and queue being referenced in the NetWare Server and Queue dialog box. Select **OK** to exit the NetWare Server and Queue dialog window and **Create** to add the queue. After these steps, you are able to manage the bindery-based queue from NetWare Administrator.

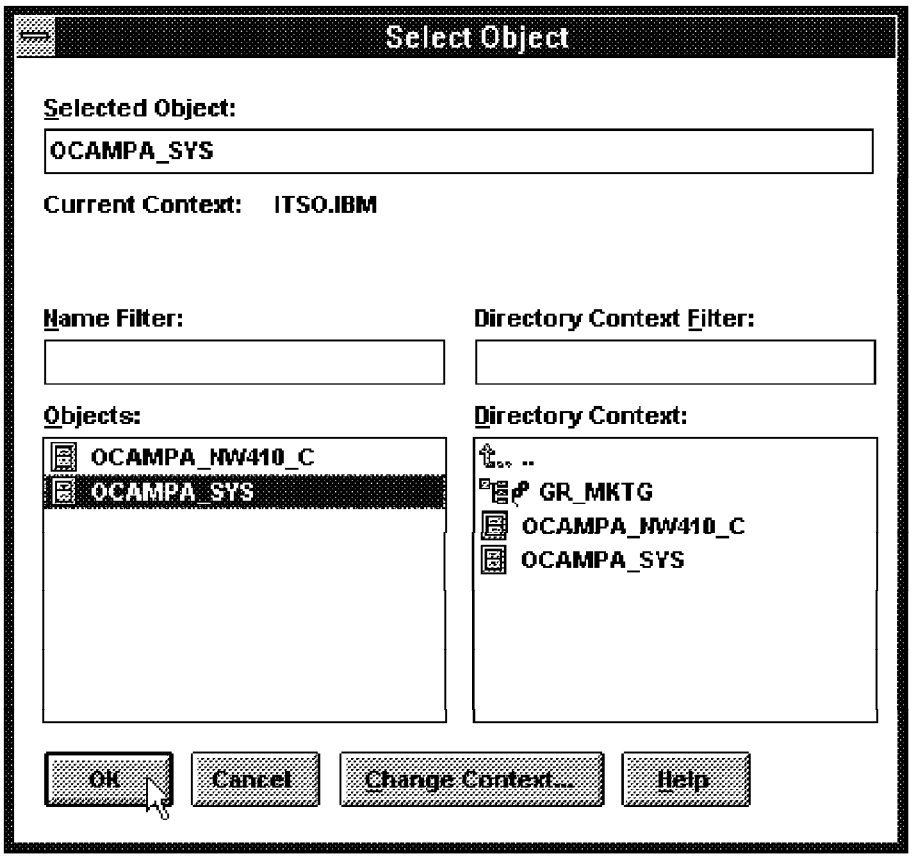

Figure 255. Select Volume for Print Queue

7. Select the Volume that should contain the temporary print jobs and choose the **OK** button. The Create Print Queue window will reappear again.
8. To create the print queue object, select **Create**.
9. When finished adding additional properties, select **OK** to exit the additional property dialog.

By default, the container becomes a print queue user, which means that all objects and all user objects below this print queue object get access to this print queue. In the additional properties you can enhance and restrict the access to the queue by deleting or adding objects to the Users page as shown in Figure 256 on page 468.

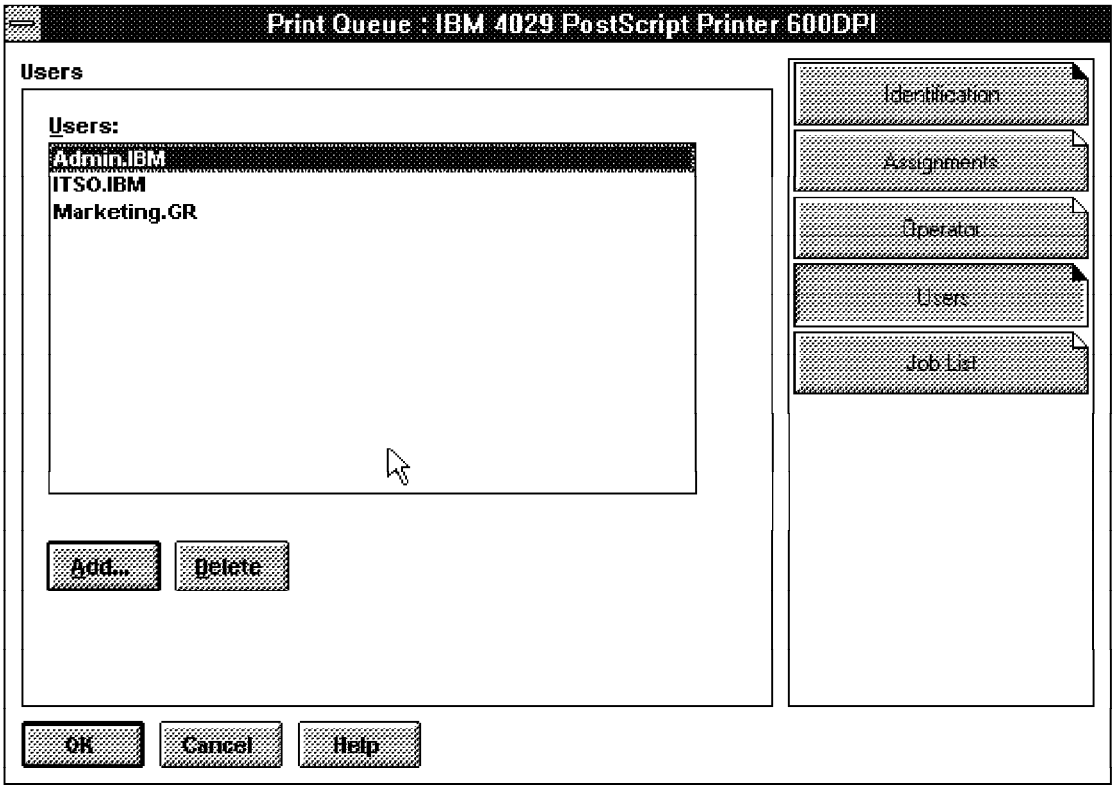

Figure 256. Enhance and Restrict Access to Print Queue Objects

For a detailed overview of all the properties that can be added to the print queue object, see the following table:

Table 42. Print Queue Object Properties

Property Page	Description
Identification	The network administrator can add to the identification page all additional information that describes the object, like other name (which can hold a former name of the object/queue), detailed description, physical location, the department the object belongs to, and the organization. Also you can set the following operator flags: • **Allow Users to Submit Print Jobs** Prevents users from submitting print jobs when the box is not checked. • **Allow Service by Current Print Servers** If this box is not checked, servers are prevented from servicing print jobs in the print queue. • **Allow New Print Server To Attach** Allows servers to attach to the print queue when checked.
Assignments	This page is only an information page where you cannot make changes. It shows you the printers this print queue is assigned to and the print servers assigned to the printers.
Operator	Print jobs in the print queue are managed by operators. This management includes activities like deleting print jobs or changing their order.
Users	This page is used to restrict and enhance access to print queue objects. You can add user objects, group objects and organizational objects to enhance access.
Print Jobs	To get information about each print job you can view the table in the Print Jobs page. It lists the print jobs that are currently in the print queue.

13.3.19 Creating Printer Objects

The next step is to create a printer object to build a relation to the printer queue. In the creation process you directly have the possibility to assign the printer object to a print queue object. To create a new printer object do the following:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree in which you want to place the printer object, and click the right mouse button on the container object to get the object's Context menu.
3. Select **Create**, or press the **Insert** key.

4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Printer**.
5. Press the **OK** button. The Create Printer window will be opened for you as shown in Figure 257.

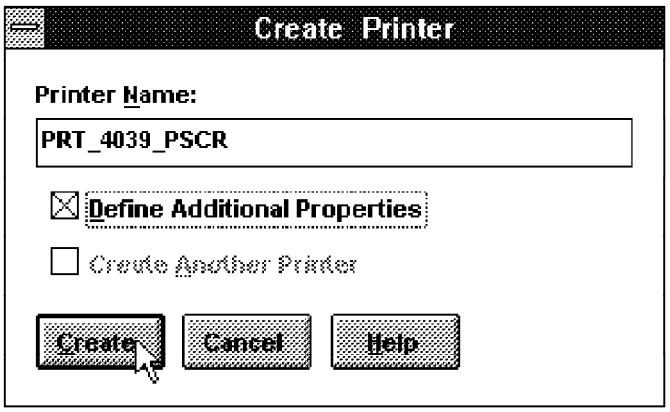

Figure 257. Create Printer Dialog Box

6. In the Create Printer window, type the name for the printer object, mark the **Define Additional Properties** checkbox, and click on **Create**.
7. In the Additional Properties window, select the **Assignments** page.
8. Click on the **Add** button to add the print queue object. The Select Object window will be opened for you as shown in Figure 258 on page 471.

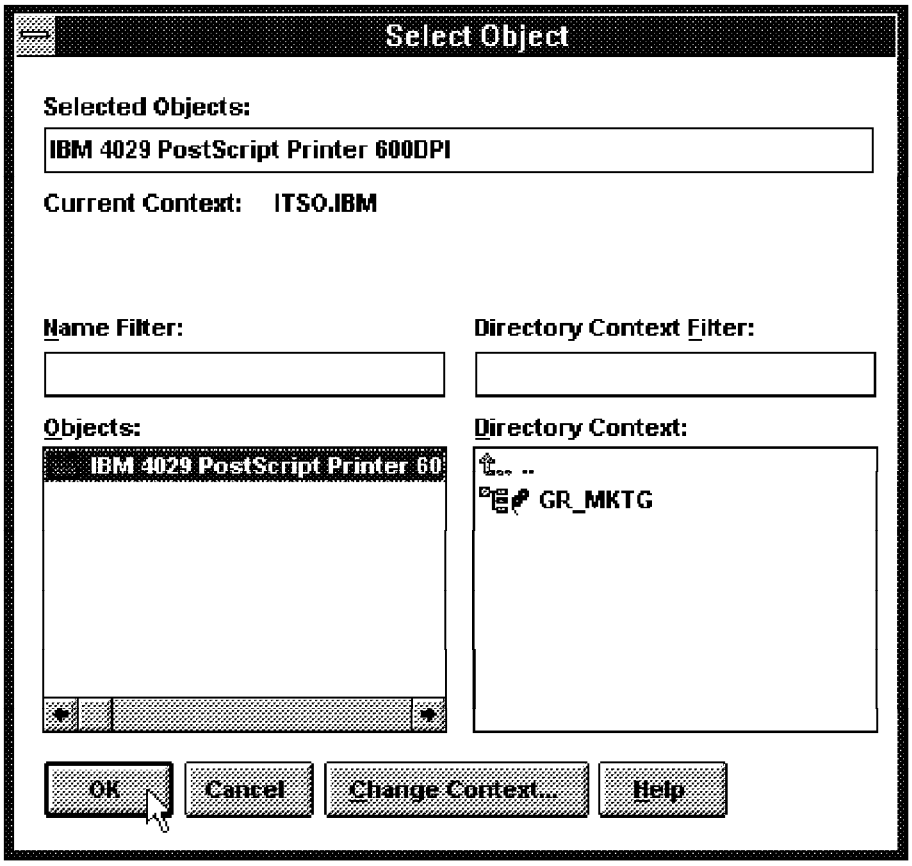

Figure 258. Assign the Print Queue Object to the Printer Object

9. At the Select Object window, select the **Print Queue Object** in the Objects list. When finished, click on the **OK** button.

Notice on Printer Object to Print Queue Assignments

It is possible to assign more than one printer to a queue and more than one queue to a printer. You might want to assign several printer to a queue if you are using identical printers and want to increase your printer throughput because when multiple print jobs are received in the queue, the next available printer gets a print job.

In cases where you have many print jobs but need to prioritize them, you might want to assign more than one queue to a single printer. When this is the case, a priority can be assigned to each queue. Jobs in a queue with a higher priority number will be serviced before jobs in queues with lower priority. The highest priority is 1.

10. Modify the the Configuration page according to the information summarized in Table 44 on page 473 as well as the additional properties.

11. Select **OK** to activate the additional properties.

If you look at the different restriction possibilities, keep in mind that defining too much control costs you performance. Also keep in mind that many of the specific restrictions are for special cases to really restrict sophisticated users. For an overview of all the additional properties for the print object, see the following table:

Table 43. Print Object Properties

Page	Description
Identification	The Identification page could hold all the detailed information to describe the printer. There is information included for • **Other Name**, like the former printer name. • **Description**, can hold a detailed information about the printer and its functions. • **Network Address** • **Location** • **Department** • **Organization**
Assignments	This page controls which printer object is assigned to which print queue object.
Configuration	To use the printer it must be configured in this page. You can define information about the hardware configuration of this printer and how it is used in the network. For details, see Table 44 on page 473.
Notification	On this page the administrator could control which objects are notified when the printer requires service, for example when it is out of paper or jammed.
Features	The Features page stores information about the printer, such as printer language, amount of memory, and supported typefaces. For example, using a search function helps you to find a printer that has at least 8 MB of memory installed and PostScript language support so that you can print out a letter using the Helvetica font.
See Also	This page gives you the possibility to list the names of objects related to the object.

To set up printer configuration you should have detailed information about the features and functions of the printer used. Also you have to add the connection information, like parallel, serial, AppleTalk, UNIX, and so forth. These options are summarized in table Table 44 on page 473.

Table 44. Configuration Options of the Printer Properties	
Options	**Description**
Banner Type	The banner type can be either Text or PostScript.
Service Interval	Represents in seconds the interval the print server will check the status of the printer. The shorter the interval is set, the faster the printer response. Be aware that shorter intervals will take more CPU power.
Buffer Size in KB	Defines the buffer size maintained by the printer server in sending information to the printer. To give a print server more control over sending jobs to the printer, simply raise the buffer's size. The default buffer size is 3 KB.
Starting Form	This represents the form number the printer is assumed to have mounted when first powering up.
Network Address Restriction	To prevent unauthorized users from placing a printer to a remote section of cabling and possibly capturing print jobs from the print queue, you can restrict on which cabling network the printer can be located.
Service Mode for Forms	This represents the way in which the printer reacts from changes. This functions is particularly needed for DOS Clients.

13.3.20 Creating a Print Server Object

The last of the three basic print objects is the print server object. You have to create the print server object to connect the other two objects, print queue and printer, to it. Keep in mind that NetWare 4.1 only allows you to set up one print server per file server.

To create the print server object do the following:

1. Start the **NetWare Administrator**.
2. Highlight the container in the directory tree in which you want to place the printer server object, and click the right mouse button on the container object to get the object's Context menu.
3. Select **Create** or press the **Insert** key.
4. The New Object window opens and prompts you to choose the object type for the object you are creating. Select **Printer Server**.
5. Press the **OK** button. The Create Printer window will be opened for you.
6. Type in the name of your print server and mark the **Define Additional Properties** checkbox.
7. Select **Create**. This opens the additional property pages.

8. To make the assignment to the printer, select the **Assignments** page and click on the **Add** button. The Select Object window will be opened for you as shown in Figure 259 on page 474.

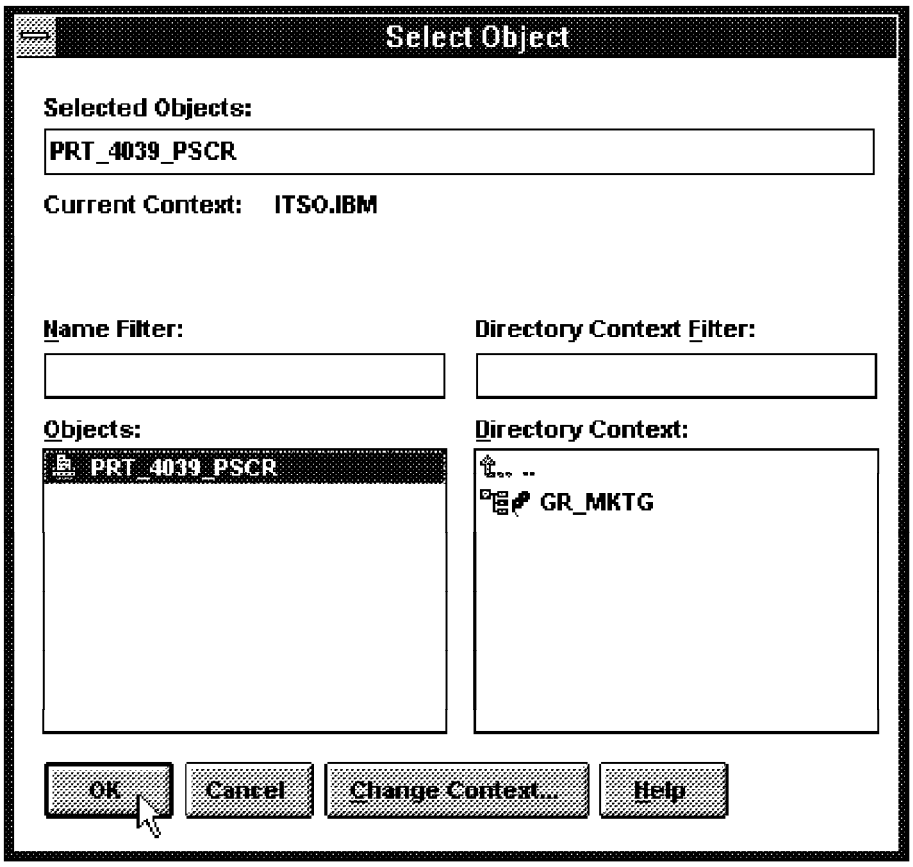

Figure 259. Assign a Printer Object to the Print Server Object

9. Choose the printer object from the Objects and Directory Context list. When finished, click on the **OK** button.
10. Add the other needed additional properties to the print server object and select **OK** when finished.

To give users access rights for the newly created print server object you have to specify those in the Users page. Keep in mind that the container object will be added automatically that grants access to the print server object to all users under the organizational container object because of inheritance. The following additional properties can be set up for the print server object:

Table 45. Properties of the Print Server Object	
Property Page	**Description**
Identification	The additional descriptions to the print server object can be added here. The fields in this page are for your information and reference only. The Change Password option allows you to ensure security for your print server by limiting access with a password.
Assignments	Here the administrator or operator of this print server could add the printer assignments.
Users	The objects that are listed as users of the print server are allowed by Queue Management Services to see the status of this print server.
Operators	To manage the print server and printers serviced by this server you have to add the user object to the operators page. Be aware that you can also add organizations and groups. This means all members of these objects are granted operator rights.
Auditing Log	To track print jobs, you have the possibility to create a auditing log file. The information stored in the log file records information like which printer the job was sent to, how long it took to print, and how large a job was. By default, the auditing is not enabled, and no log file is created unless you enable auditing and reload the print server.
Print Layout	To see an assignment map of your printing setup, you can view the Layout page. Here you find the assigned print object and each print object's status. This information is helpful to determine whether printing is up and running, or which print objects may have problems.

13.3.21 Starting the Print Services

After installing and connecting the base objects, you must load the Print Services at the Novell NetWare 4.1 Server. To do so, you have to load the PSERVER.NLM, which is a NetWare Loadable Module that is used to start up the print server you have defined.

To load the NetWare Loadable Module you must have access to the Novell NetWare 4.1 Server. Use the RCONSOLE.EXE to load the Print Services or load it directly from the Novell NetWare 4.1 Server system console by typing the following command:

```
LOAD PSERVER [Name of print server object in Novell Directory Serv
```

It is essential that you write the correct object name. In our example this is PSRV_PAPERWISH; otherwise the Print Services cannot be started. If

needed, you also have to make this information known to the organizational container, like ITSO.IBM.

After starting the Print Service, you have the possibility to access the following information via the PSERVER.NLM menu:

- Printer Status
- Print Server Information

The text mode window displayed looks like the one shown in Figure 260.

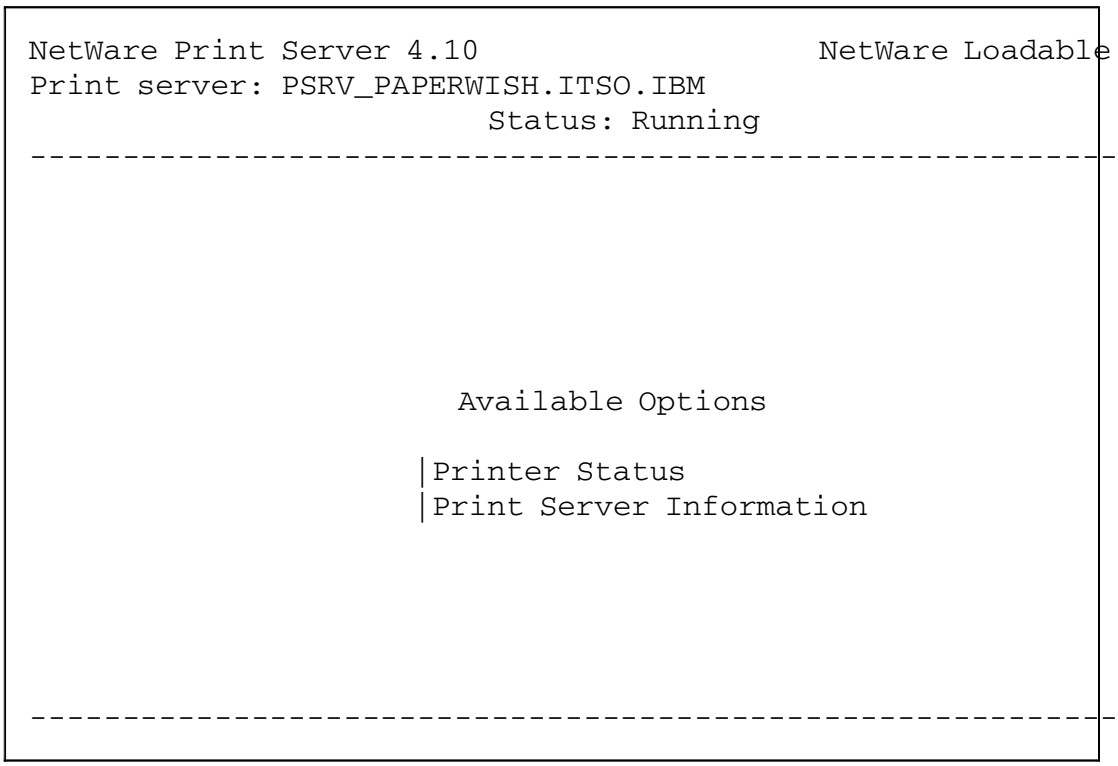

Figure 260. PSERVER.NLM Menu

After successfully starting, the print server will automatically log in to the file server, which you can identify because the License counter has raised by one. Keep that in mind when calculating the licenses.

Via the different menus, you can select information about the printers installed. For example, select **Printer Status** and then select the printer you want information about as shown in Figure 261 on page 477.

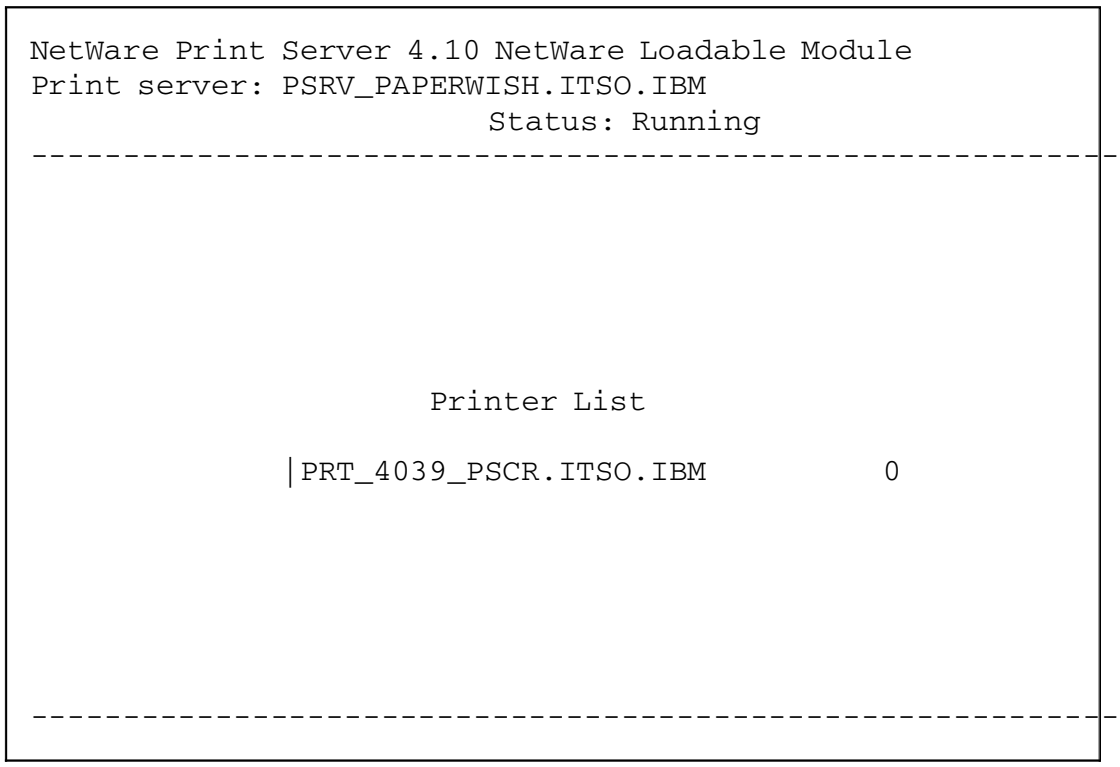

Figure 261. Printer Status with Printer List

Also you can view information about the print server as shown in Figure 262 on page 478.

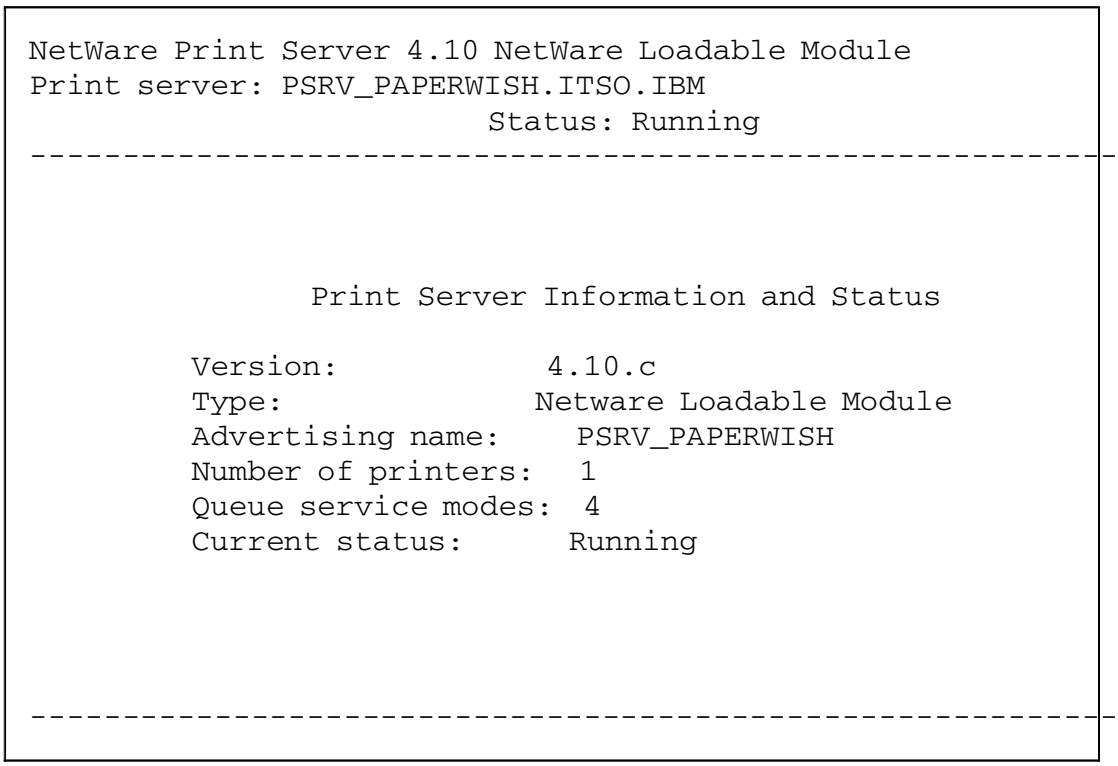

Figure 262. Print Server Information and Status

After you have set up the Novell Directory Services objects and loaded the Print Service (PSERVER.NLM), you also must configure the printer. This means that you have to decide how you want to connect the printer to the network, at a workstation or directly at the server.

To control a printer that is directly connected to a Novell NetWare 4.1 Server, you have to load another NetWare Loadable Module, which is called NPRINTER.NLM. This NetWare Loadable Module must be loaded for every single printer. The syntax of the command that can be run from a remote console or from the server console is:

```
LOAD NPRINTER [Print Server Name] [Printer Number]
```

Viewing the results of the command via the command prompt is:

```
Print server:     PSRV_PAPERWISH
Printer:         3
Printer name:     PRT_4039_PSCR
Printer type:     LPT1
Interrupt:       Polled

Status:
   NPRINTER status: Waiting for Print Job
   Printer status:  Out of Paper
```

To avoid affecting the server performance too much, it is recommended not to connect too many printers on a server that also has the task of being a file server because you have to load the NetWare Loadable Module NPRINTER.NLM for each printer connected to the print server.

To connect a printer to a workstation and make it accessible, you have to load the TSR (Terminate and Stay Resident) program NPRINTER.EXE. The syntax of this program is similar to the one of NPRINTER.NLM. To start the defined printer in our example, you would type:

```
NPRINTER PSRV_PAPERWISH 4
```

Instead of the printer number you can also use the complete name of the printer as shown below:

```
NPRINTER .CN=PRT_4039_PSCR.OU=ITSO.O=IBM
```

If you start NPRINTER.EXE without parameters, you get a selection menu to select the print server as shown in Figure 263.

```
NetWare Print Server 4.10                Friday 12 July 1996 18:4
              User on NetWare Server OCAMPA Connection 2
-------------------------------------------------------------------



                    Active Print Servers

                     | PSRV_PAPERWISH   |




-------------------------------------------------------------------
Select the print server for the printer attached to this workstat

-------------------------------------------------------------------
Enter=Select    F4=DS Printer Object Mode  ESC=Exit           F1=
```

Figure 263. Select the Print Server for a Workstation Printer Attachment

13.3.22 Access Rights Administration

The previously created user account does not have access rights thus far except access rights to the default login drive. In addition, you can set up other user-related things.

With the implementation of Novell Directory Services, there are a lot of changes in comparison to the former NetWare structure. This especially is true for the entire rights management. Novell Directory Services's implementation not only covers servers, it covers an entire network. This means that Novell Directory Services can be used to manage little and small networks as well as huge enterprise-wide networks. The advantage is that Novell Directory Services is designed to describe really large environments and not only one server. However, for small networks, this detailed rights structure is not useful everytime, but it is very helpful when it comes to managing a worldwide enterprise network.

13.3.23 Object Rights

The Novell Directory Services structure is object oriented and based on a non-object-oriented operating system. All defined rights are those of objects. This helps to control and manage the access to the single objects. From the security perspective, only authorized users can access the objects. Table 46 shows you the different rights that can be designed for objects within Novell Directory Services:

Table 46. Different Object Rights

Object Right	Description
Browse	To see the object within the directory tree you need to have at least the browse right. For user objects, it is automatically assigned for the respective user to make sure that he/she can view the related object in the Novell Directory Services
Create	This right is only given for container objects. It enables the user to create new objects within it. This includes the creation of container and leaf objects.
Delete	This right enables you to delete leaf objects and empty container objects. Before you can delete an organizational container object, you have to delete the leaf objects that are included in it.
Rename	The rename right grants you the functionality to change the name of objects. This right is only possible for leaf objects because organizational containers cannot be renamed.
Supervisor	Granting this right to a user makes him to a super user for the whole or only a part of the tree objects. It gives him/her the right to change, create, delete, and rename objects. Also he/she is allowed to modify object properties.

The Supervisor right is different to all the other rights that can be granted to users because it automatically includes all rights to object properties. All other rights represent no function to manipulate object properties.

13.3.24 Property Rights

To give users the right to modify property rights without granting supervisor rights, there is another category of rights that allows you to grant the right to modify object properties. The following rights could be disposed:

Table 47. Table of the Possible Property Rights

Property Right	Description
Compare	Enables you to compare property values. It is needed to make queries in the Novell Directory Services. It does not give you the right to view the contents of properties.
Read	To also see the contents of object properties, the user needs the read right, which also includes the compare capability.
Write	This right gives you the possibility to change and delete the contents of property pages. If the "write" right is given, Add / Delete Self is included also.
Add / Delete Self	Allows an object to add or remove itself as a value of a property. It does not affect other values in the property. This right is only used for properties that contain object names as values, such as lists of group members or mailing lists. This allows users to add or remove themselves from lists stored in an object.
Supervisor	The Supervisor right grants all access privileges. A trustee who has the Supervisor right automatically has all other rights to the property. The only way to block Supervisor rights is by using an Inherited Rights Filter.

The property rights could be varied for the different properties of an object. This means that it is possible to grant the right to a user so that he/she can modify his login script, but only have the right to view the Identification and Environment page.

13.4 Dynamic TCP/IP in NetWare 4.1

Novell Inc.'s NetWare/IP 2.2 is a software option that is free for NetWare 4.1 customers, and it provides TCP/IP networking. The TCP/IP services built into NetWare 4.1 offer simple IP services but no DHCP or DNS capabilities. NetWare/IP 2.2 fills this gap by providing both DHCP and DNS servers.

13.4.1 NetWare/IP 2.2

NetWare/IP 2.2 provides both DHCP and DNS servers. The DHCP Server is configured by a NetWare Loadable Module (NLM) that allows IP address-range creation and management. It also lets you reserve DHCP and BootP leases. The interface is NLM-like and offers little in the way of extended functionality. In some cases, such as exclusion ranges, the configuration tool is quite limited.

The DNS server can be used as a name-resolution server for any client capable of DNS lookups. The `Unicorn` administration utility provides easy configuration and management of the DNS server. However, the DNS server is quite static — it cannot perform any form of dynamic DNS updating. This simply means that using the DNS server for anything other than centralizing your DNS server with your DHCP Server is unrealistic.

NetWare/IP is now part of IntranetWare, Novell's latest version of NetWare that includes Internet and intranet capabilities. The Internet technologies in IntranetWare do not require NetWare/IP, but it allows users to access NetWare-specific services such as the NetWare Core Protocol. NetWare/IP is backward-compatible with NetWare-based applications, including IPX applications.

For those of you Netware 4.1 users who do not have NetWare/IP, it is available for free from Novell's Web site at the following URLs:

```
http://www.novell.com/corp/offices/san_fran.us/download.html
http://support.novell.com/Ftp/Updates/unixconn/nwip22/Date0.html
```

13.4.1.1 Experiences with NetWare/IP 2.2

In comparison to Warp Server's and Windows NT's implementation of installing and configuring a DHCP Server, NetWare/IP's administration capabilities are very poor. The product cannot exclude addresses from a range, and there is no minutes setting for a lease period. The product is server-centric, so multiple servers cannot be managed at the same time.

Although the ranges allowed a subnet mask, they did not allow a range of exclusion addresses — you had to put in every IP address that you did not want included in the range. Also, without some editing the lease terms, only accepted hours and days, not minutes or months. NetWare/IP does not offer DHCP classes support, just like Windows NT DHCP Server. However, DHCP classes are used to administer large TCP/IP networks. Only Warp Server supports DHCP classes (see 11.13.1, “Configuring and Using DHCP Server” on page 244 for more information).

Although the DHCP Server supports BootP for both static and dynamic assignment, it does not offer reverse address resolution.

The DHCP management is done through Novell's DHCP configuration product, a stripped-down IP configuration and management tool. DNS configuration is done through the `Unicorn` utility, which allows easy creation and deletion of all resource records. However, it is a quite static DNS server. It cannot perform any form of dynamic DNS updating. Warp Server is the only network operating system platform, that offers dynamic DNS serving. Windows NT's DNS server is static as well and has no integration with DHCP except in conjunction with WINS.

Because of all these limitations, we saw no need to a NetWare/IP implementation in a large-scale environment, except you may be interested in having a central server for both DHCP and DNS. There is lots of room for improvements especially when compared to Warp Server.

Following clients are supported by NetWare/IP:

- NetWare Client 32 for Windows 95
- NetWare Client 32 for DOS and Windows
- NetWare DOS Requester (16-bit, VLM-based) Client that ships with NetWare/IP

Warp Server clients as well as Warp 4 workstations are not supported by NetWare/IP.

13.5 ManageWise 2.0

In this section we discuss a product that can be purchased separately to NetWare 4.1 in order to get systems management functions in a NetWare environment

Management systems are designed to collect information about the network and its components and to display that information on a central system where managers can manipulate and interpret what they see.

Novell's network management services strategy is to provide an open, standards-based platform for enterprise management. Novell and Intel have combined their management products to create ManageWise for NetWare. ManageWise enhances NetWare server management and unifies the management capabilities included with NetWare into one graphical user interface.

The NetWare Management System (NMS) includes support for more than 600 SNMP alarms and alerts, and this lets NMS function as a subordinate management system for enterprises using HP OpenView or Sun Microsystems' SunNet Manager. NMS also integrates with IBM's NetView, including support for native NetView alerts and messages.

ManageWise is an integrated server, desktop, and network infrastructure management system. ManageWise features include:

- Hardware and Software Inventory
  - ManageWise eliminates the need to take physical inventory. It automatically discovers network devices including routers, hubs, servers and desktops, regardless of protocols. To obtain specific information and focus on your particular group, ManageWise lets you target a specific workgroup or domain. Discovery can also run unattended or be scheduled for continuous or periodic updates. ManageWise displays devices discovered in hierarchical, graphical maps. Internetwork maps provide an overview of the network topology, showing how Ethernet and token-ring segments are logically interconnected with routers. Segment maps display active stations on the network for a specific segment, and you can customize these maps to show the geographical location of sites and devices. ManageWise also tracks hardware and software loaded on networked DOS and Windows PCs, Macintosh desktops and NetWare servers. ManageWise automatically scans each node for component information, such as processor type, memory size and installed software applications. You can then define and search for supplemental information during an asset scan. All asset information is stored in an inventory database, making it easy to find sets of stations with similar configurations.
- Desktop Management and Remote control
  - Using ManageWise's desktop management services, you can view, manage, and control networked PCs across the network from a single location. By taking control of the remote workstation's screen and keyboard, you can quickly solve user support problems. You can also gather detailed configuration and performance information, transfer files, execute programs, and remotely reboot the user's system to activate configuration changes. ManageWise also lets you "talk" in real time to users through their desktop screens. When invoked, chat mode displays a split screen to show both sides of the conversation. ManageWise's file transfer functions make it easy for you to selectively transfer files to users' local hard drives. Users do not need to download files, and you can ensure that they have the latest updates of system and application software.

- Virus Protection
  - The ManageWise virus protection feature incorporates sophisticated pattern-based scanning, which identifies known viruses at the server and desktop. It detects mutations of existing strains, polymorphic viruses, and even stealth viruses designed to hide from scanners. Unknown viruses may elude pattern-based scanning, but the intelligent rules-based detection capabilities of ManageWise identify virus-like behavior before a virus can cause havoc on the system. Virus protection operates continuously in real time to protect servers, desktops, portable and stand-alone PCs from virus infection. As files are transferred to and from the server, it checks in real time to ensure they are not infected. It also performs prescheduled and on-demand scanning on servers. In addition, ManageWise scans users' workstations continuously or at login to identify any viruses on the desktop. Whenever mobile computer users reconnect their portable computers to the network, the ManageWise virus protection feature automatically checks for virus activity and updates the server log as well as the portable computer's virus pattern files. The virus protection in ManageWise includes an "automated download feature", which automatically logs into the Intel Bulletin Board System (BBS) and downloads the most current virus pattern file. ManageWise also shares virus pattern file information among multiple file servers. Both features save considerable amounts of network administration time.
- Distributed Network Analysis
  - ManageWise provides network analysis based on Novell's NetWare LANalyzer software. It monitors the interaction among all network devices. ManageWise provides distributed analysis tools for management of network activity. For example, ManageWise identifies the nodes that are generating the heaviest demand on network resources and shows you exactly what the device is doing to create that demand. ManageWise provides detailed information on packet and data rates for the whole network, each station, and each individual conversation. This capability lets you identify overloaded network devices, segments, or users generating heavy traffic and lets you rebalance the network load. ManageWise monitors network traffic and alerts you to potential problems, such as network errors or duplicate IP or node addresses. After receiving these alerts, you can direct ManageWise to begin capturing packets. These can be stored in server memory and then examined and interpreted to show the state of each station on the network. The network analysis capabilities of ManageWise include full packet capture and decode and support the most common networks

deployed today. These include Ethernet, token-ring, and variants of these standards (for example, 100 Base T and 100 VG-AnyLAN). Along with supporting common hardware standards, ManageWise supports common network protocols, including all versions of NetWare and its Internetwork Packet eXchange, (IPX) protocol, TCP/IP, AppleTalk and SNA. ManageWise facilitates the management of IPX and IP network addresses by displaying a tabular report of network numbers in use, along with the physical media type, system name, and subnet mask. Quick identification of assigned addresses eliminates duplications when adding new nodes to the network.

- Print Queue Management
  - ManageWise monitors and manages print queues on the network. It displays network print queue activity by the number of jobs or the total number of bytes. It shows graphically which queues and servicing printers are the most active and displays queue-related information about servers, queues, and jobs.
- Server Management
  - By installing ManageWise on each NetWare server, you can monitor, maintain, and manage all NetWare servers from a central site. You can also compare multiple servers to optimize server configuration and performance. ManageWise allows unattended monitoring of NetWare servers 24 hours a day, seven days a week. Before problems reach a critical state, ManageWise will immediately alert you to those problems. It monitors 378 server conditions, including directory, disk drive, volume, memory, logged-in user, and NetWare Loadable Module (NLM) software, and covers 378 alarm conditions. You can customize alarm thresholds, such as file activity, memory usage, and logged-in users, to their environment. ManageWise provides complete NetWare management from one administration console. It incorporates familiar NetWare tools, such as NetWare Administrator, to administer servers and graphical `SET` commands to change NetWare server configurations and resources. ManageWise monitors trends on server performance, such as file activity, volume data, disk drive configuration, adapter card, CPU, memory, and print services. You can analyze data collected from servers and take action if necessary, such as reallocating volume space or redirecting jobs in print queues. With ManageWise, you can control problems and perform daily NetWare tasks from a single, intuitive interface.

If you work for a small- to medium-sized organization, you don't need to be an expert to effectively manage your network with ManageWise. Familiar NetWare security features and NetWare administrative tools, such as

NetWare Administrator, and intuitive navigation through maps make ManageWise easy to use. ManageWise guides you to quick resolution on common user problems and lets you plan for growth by using its historical data trending capability.

ManageWise also works over on-demand dial-up links, letting you outsource management tasks to service bureaus easily and cost effectively. These and other capabilities make ManageWise the best solution for small- and medium-sized networks. If you work for a large organization, you can customize ManageWise to your needs. Your site managers, for instance, can get details about the local network and proven technology to assist users, while your enterprise managers get a broad overview of network activity and problems across the organization.

Because ManageWise is built on standards, such as SNMP, enterprise and site managers can cooperatively share management of the network using their console of choice, whether a third-party, UNIX-based console or the included ManageWise console. No matter which console they use, they all get access to the same information and network maps. Many enterprise networks have multivendor components that need to be managed.

ManageWise is based on industry standards, such as SNMP, RMON, IPX and TCP/IP, so it provides interoperability with any network system. Since it is installed on NetWare servers, it scales to the network's processing power, making it easy to expand management capability as the network grows. ManageWise also works well over wide area network (WAN) connections, letting network managers cost-effectively administer geographically dispersed servers.

13.6 Backup/Restore with HSM (Hierarchical Storage Management)

HSM is a Cheyenne Hierarchical Storage Management software feature that is not provided with NetWare 4.1. HSM provides the ability to define and automate the decisions regarding where and when data is stored. The name HSM means that the software has the capability to move files along a hierarchy of storage devices that are ranked in terms of cost per megabyte of storage, speed of storage and retrieval, and overall capacity limits. Files are migrated along the hierarchy to less expensive forms of storage-based on rules tied to the frequency of data access. File migration and retrieval is transparent to end users.

Cheyenne HSM provides a completely automated system for network data and storage management for NetWare 3.11, 3.12, and 4.x servers. HSM provides a flexible set of storage-management parameters that define data migration from the server hard drive through the storage hierarchy. This

migration can be managed by server hard drive utilization, file usage, age of files, or even a combination of both. A number of filters also give you the possibility to include or exclude some specific directories or files based on directory and file name, size, owner, name space, and file attributes.

A typical three-step strategy could be:

1. Hard drives as primary storage on file servers
2. Rewritable optical as the secondary storage type
3. Tape device as the final storage location

It is even possible to implement WORM in place of tape as the final storage destination. HSM does not back up files on the server hard drive, but migrates files to other forms of storage. This option leaves free hard disk space. As files are migrated off the primary server, the Novell filing system continues to show the original state of the file on the server even though the contents of the file have been migrated to another device. When a migrated file is requested, a key is set and NetWare requests that HSM migrates the file back from magnetic or optical devices, thus enabling automatic file retrieval and user access. The backup system must ensure that these keys, as well as other files that remain on the primary server, are backed up. The primary benefits of HSM include:

- Better use of available space on the primary disk drive for current data
- Optimized access time
- Reduction in overall cost of storage
- Simplified disk space management

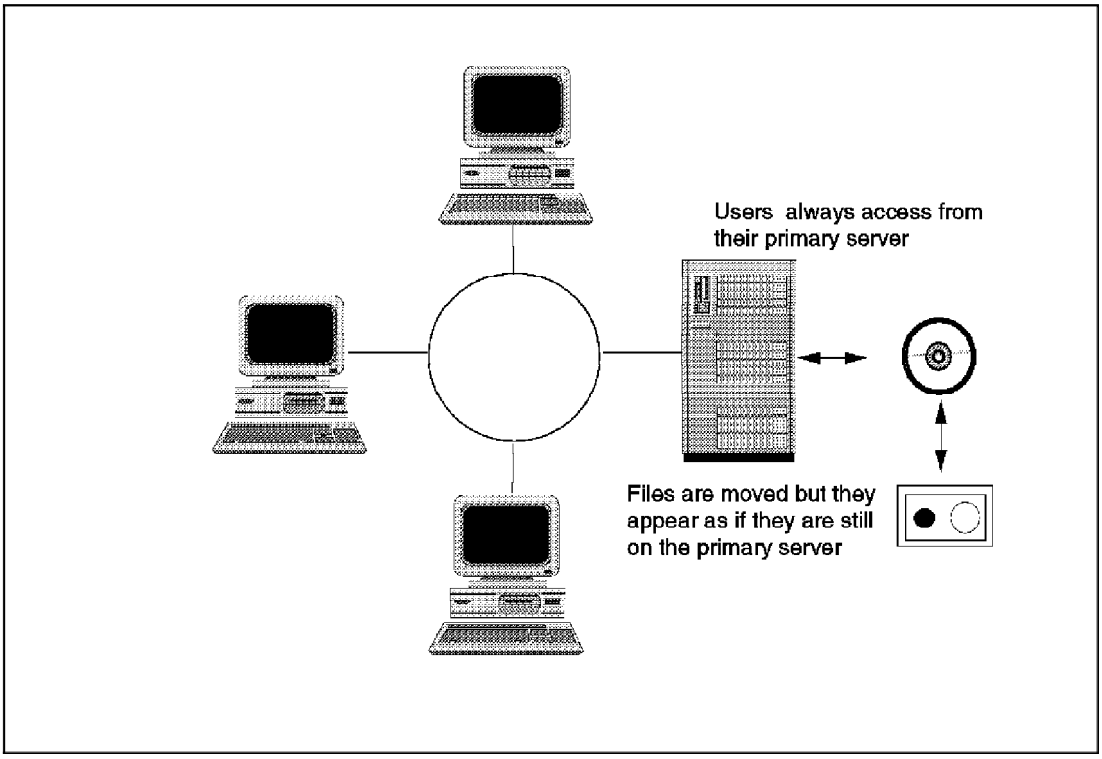

Figure 264. HSM Network View Migration - Demigration

As you can see in Figure 264, HSM automatically transfers infrequently accessed data from file servers or hard disks to storage options such as high-capacity hard drives, optical jukebox, or tape libraries. The primary disk at this point is free to hold only the most frequently accessed data. HSM transparently and automatically manages the distribution data; the network administrator does not have to:

1. Identify inactive data
2. Manually move files to other storage device
3. Maintain catalog and backups for file location

13.6.1 HSM Architecture

During backup, a LAN Administrator archives data from the network volume to a tape drive. During this process, the administrator also attempts to identify the inactive files on the network volume and move them to a less-expensive storage medium. HSM can perform both functions automatically and in a transparent way. The users are not aware that the files have been migrated from the file server's hard disk to another form of storage. In HSM there are several factors used to choose which files to migrate. The primary factor is usually disk capacity:

- The administrator sets critical High and Low thresholds, and the HSM keeps disk capacity between these levels. At this point, when the threshold is crossed, HSM looks for the oldest eligible files to move. HSM identifies which files are the least-frequently accessed by end-users and the files' ages. It is also possible to set HSM to exclude some executable files or DLLs. When a file is migrated off the primary storage medium, HSM leaves an invisible key. This key consists of a pointer to the file's new location. To the user, the file still appears to be stored on the primary stage. When a migrated file is requested, NetWare will request HSM to demigrate the file from magnetic or optical devices by copying the file to the hard disk, thus allowing access by the end-user.

As shown in Figure 264 on page 489, another feature of HSM is that it gives you the maximum flexibility in configuration schemes:

- For example, one server can be shared among the hierarchical storage system and used for other applications, depending upon system resources. Another solution could be to off-load management backup and migrations services to a dedicated server. In this case the primary server would continue as before, and a second server would focus all its processing power on performing secondary and tertiary migration backup. The total distributed strategy would consist of each storage stage being installed on a separate server. This not only provides distributed computing but also offers distributed data storage as well. The configuration chosen depends upon the server CPU and storage utilization, network bandwidth, and physical storage location requirements.

Figure 265 on page 491 gives an idea of how flexible the HSM architecture is for use in hierarchical storage management.

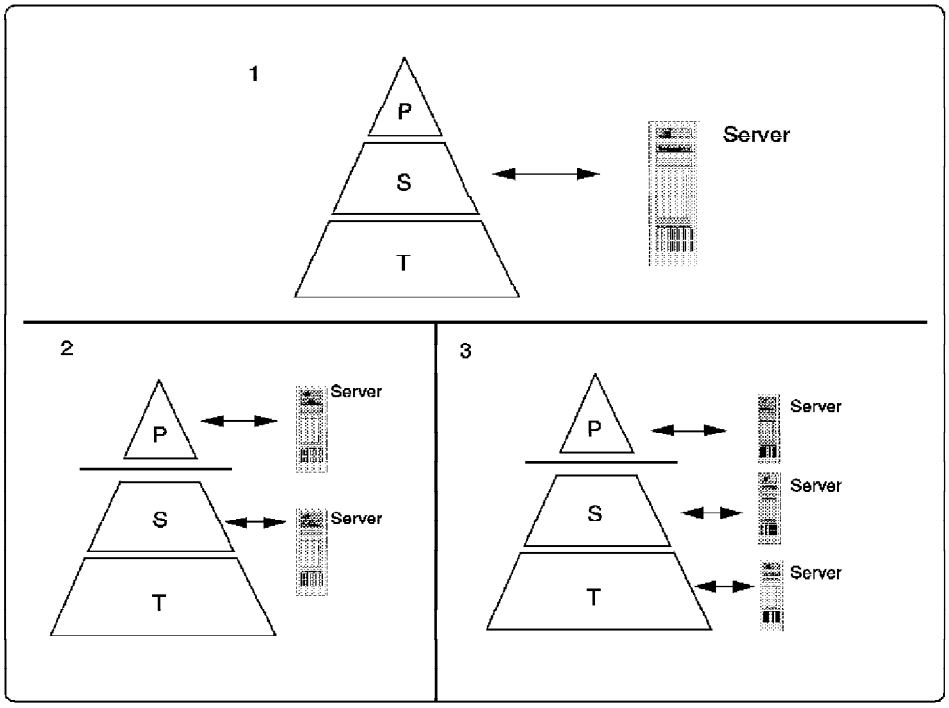

Figure 265. Flexible Architecture

Agenda for Figure 265:

```
P = Primary Server
S = Secondary Server
T = Tertiary Server
```

HSM allows the network administrator to design a hierarchical storage system composed of magnetic, optical, or tape devices at each level of hierarchy. HSM supports many different kinds of optical drives, jukeboxes, tape drives, and autochangers.

Chapter 14. Introduction to Directory and Security Services

The components of the Directory and Security Server are administered using a graphical user interface (GUI), which runs on the DSS client. This administration GUI can be used to administer existing OS/2 LAN Server or OS/2 Warp Server domains, IBM and non-IBM DCE cells, and DSS cells. It is based on the OS/2 LAN Server 4.0 graphical user interface, but it has added features to support the time, security, and directory services. Figure 266 shows the DSS graphical user interface.

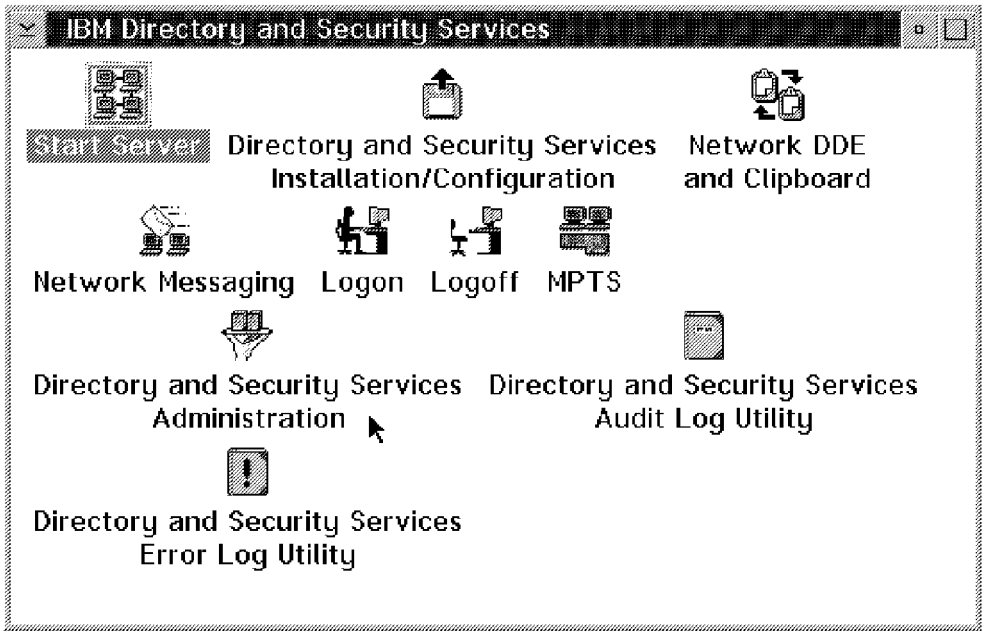

Figure 266. DSS Graphical User Interface

Administration is done from the DSS administration window, as shown in Figure 267 on page 494.

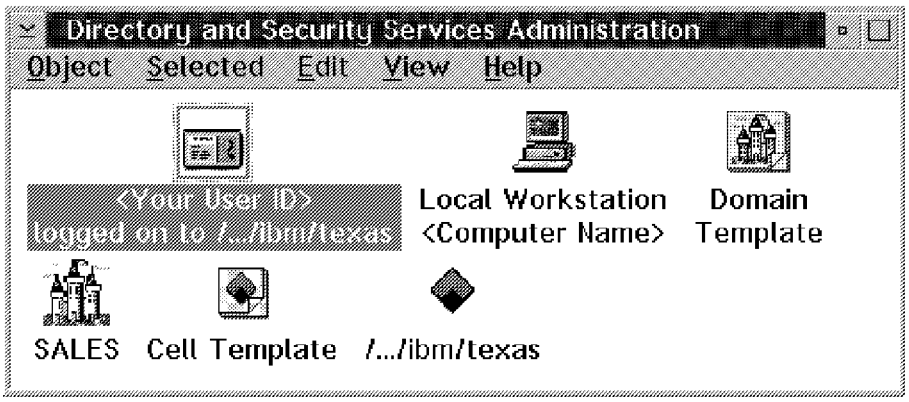

Figure 267. DSS Administration GUI

There is also a DCE-only GUI, which is a subset of the DSS administration GUI and that runs on a DCE client. Most administration functions can be performed remotely. See Figure 268 for the DCE interface.

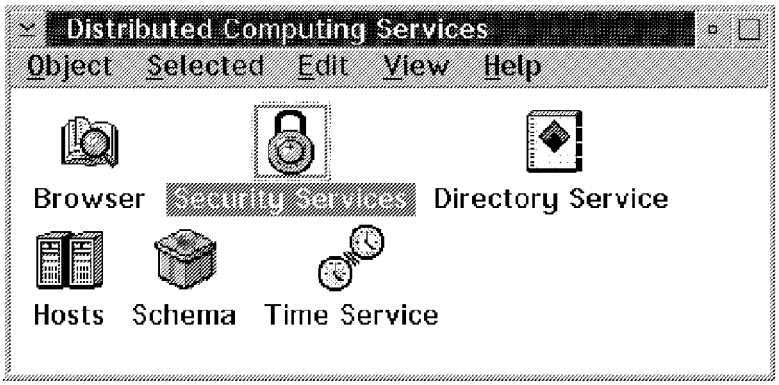

Figure 268. The DCE Interface

The basic scope of administration in the Directory and Security Server is the cell. Users are defined at the cell level rather than at the domain level, as is done in OS/2 LAN Server and OS/2 Warp Server. Cells can include many users and resources and can span many geographical sites. They can consist of existing OS/2 LAN Server and OS/2 Warp Server clients and servers, DSS clients, domain controllers, and additional servers that have been upgraded to DSS. See Figure 269 on page 495 for a cell folder.

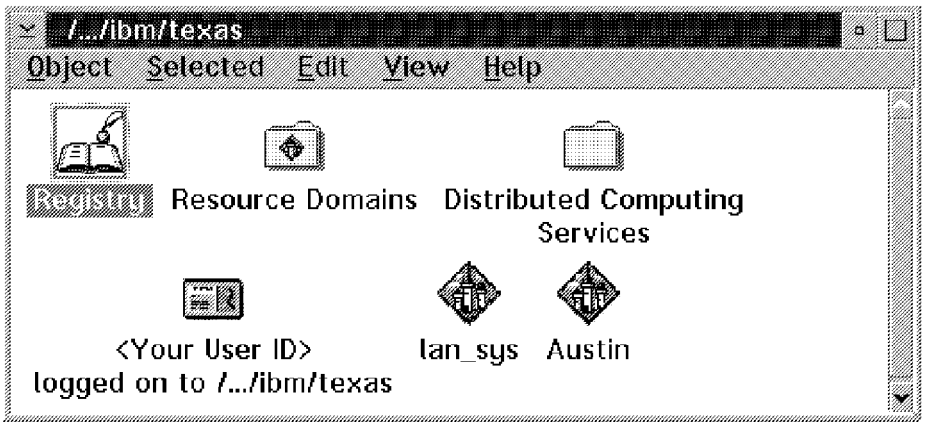

Figure 269. Example for a Cell Folder

In these mixed environments, the existing clients and servers use a protocol that does not recognize the DCE Directory and Security Services. They use the existing OS/2 LAN Server protocols and understand the existing OS/2 LAN Server security and directory structures. For this reason, DSS synchronizes the DCE directory and security databases with the OS/2 LAN Server directory and security databases. This is accomplished with two DSS functions: *registry synchronization* and *directory synchronization*. Registry synchronization is the act of ensuring that all updates to the Directory and Security Server registry database are propagated to the NET.ACC file on the domain controller. Directory synchronization ensures that changes to the Directory and Security Server directory database are reflected in the domain control database.

In order to reduce the traffic on the network, synchronization is done on a resource domain boundary. Synchronization may be enabled or disabled, on an individual basis, for each resource domain in the cell. It is enabled by default when the resource domain is created. When synchronization is enabled for a resource domain, changes to the registry database for users and groups in that resource domain can be reflected in changes to the NET.ACC file on the domain controller for that resource domain and changes to the directory database for resources in that resource domain can be propagated to the DCDB on that resource domain's domain controller.

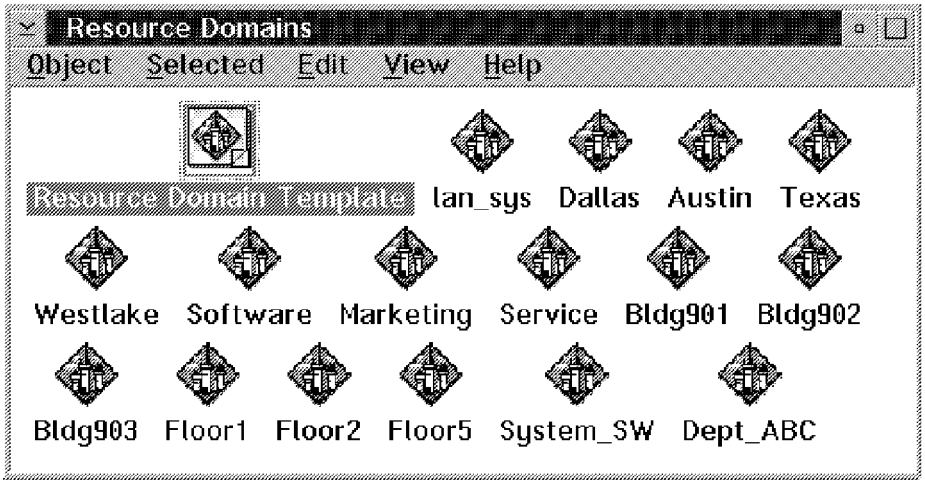

Figure 270. Resource Domain Folder

Resource domains allow a level of granularity in administration that is not possible with OS/2 LAN Server and OS/2 Warp Server domains. The administrative relationships among the resource domains of the cell can be defined by two basic models:

- A *flat* model in which the administrator of each resource domain is independent and autonomous
- A *hierarchical* model in which *ascendant* resource domain administrators have authority over *descendant* resource domains

Figure 271 illustrates the two basic models.

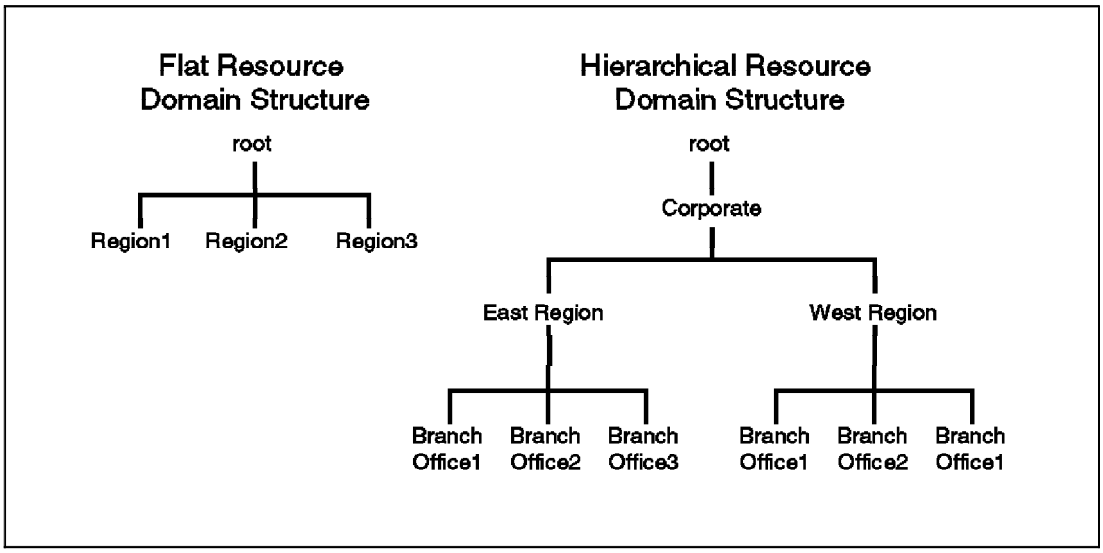

Figure 271. Flat and Hierarchical Resource Domains

For example, a company might have one resource domain for each branch office and each of those resource domains might be children of a resource domain for the entire company. In this case it would be possible for the administrator of the entire company-resource domain to administer the resource domains in any branch office, but the branch office resource domain administrators could only administer their own resource domains. This makes it possible to put boundaries on the resources that a given administrator can control, even though all of the resources are in a single cell.

Figure 272 shows a more sophisticated view of the resource domain structure.

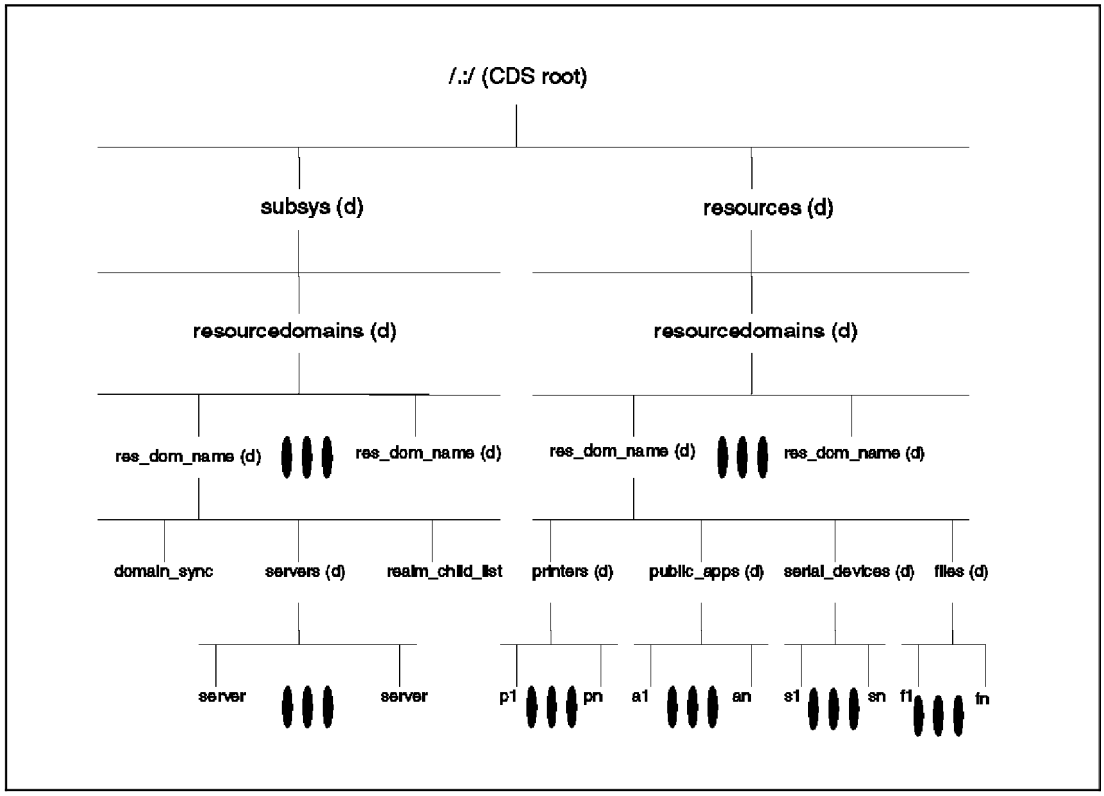

Figure 272. Resource Domain Directory Structure

Some more administrator panels are shown in the following figures. Figure 273 on page 498 shows the contents of the resource domain *Austin* compared to the legacy domain *SALES*, in this case an OS/2 Warp Server domain.

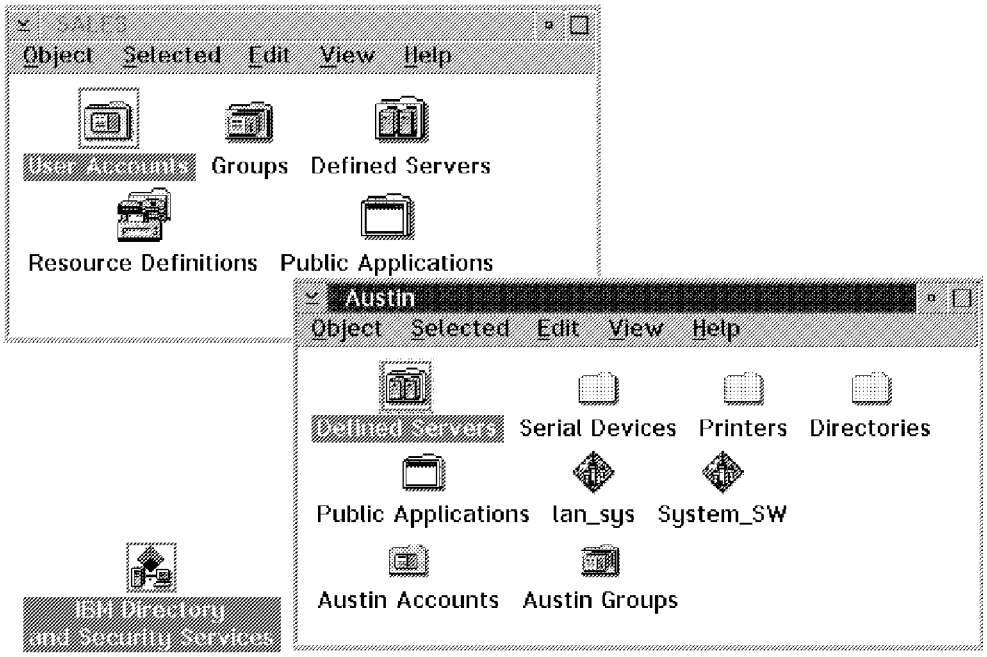

Figure 273. Resource Domain and Legacy Domain Folder

From the DCE folder (see Figure 268 on page 494), the DCE cell can be browsed by using the CDS Browser as shown in Figure 274.

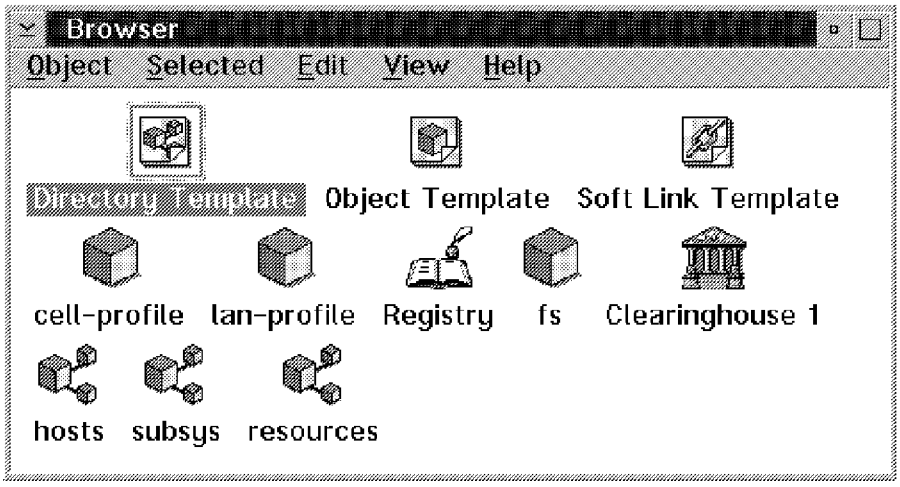

Figure 274. The CDS Browser

From this panel, a hierarchical view of the resources can be opened (Figure 275 on page 499) that shows the integration of LAN Server directories in the resource domain.

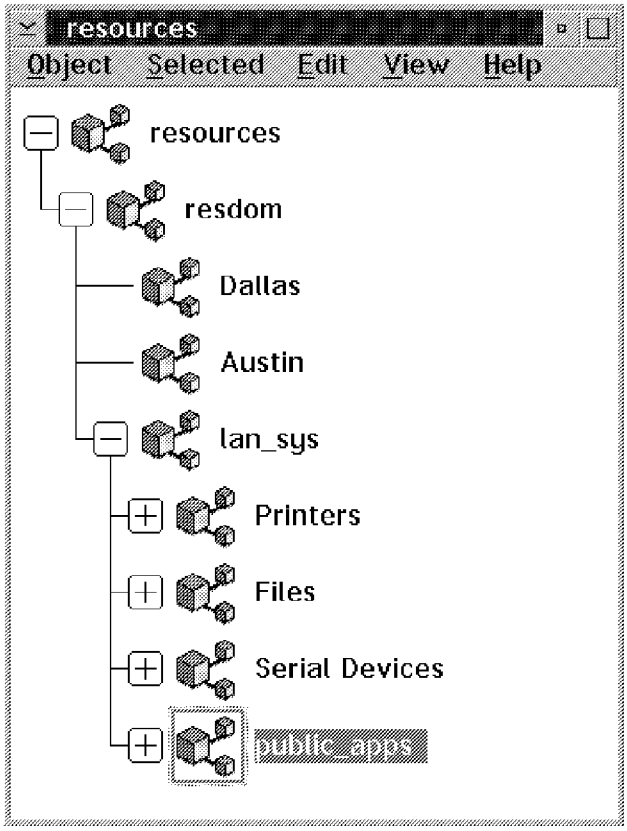

Figure 275. LAN Server Directories Integrated into Resource Domains

14.1.1 Client Capabilities and Interoperability

Several types of clients can be used with the Directory and Security Server. Existing OS/2 LAN Server and OS/2 Warp Server clients can be used unchanged. These clients are capable of accessing resources anywhere in the DSS cell by using a single user ID and password. In addition, they can continue to be used to access resources on OS/2 LAN Server and OS/2 Warp Server domains that are not part of the DSS cell. The Directory and Security Server strengthens the control of the administrator with respect to existing clients. In a DSS cell some of the administration functions that could be performed from existing clients must now be performed from a DSS client. This is done so that DCE security can be used end-to-end on the client that performs the administration.

A user must be a member of a resource domain in order for their user account to get replicated. So if an user account is not a member of an external resource domain, he/she may not be able to access resources in that domain.

DSS clients can access not only DSS domain controllers and additional servers, but also additional servers and domain controllers in existing OS/2 LAN Server and OS/2 WARP Server domains. They can also access unchanged additional servers in the DSS cell. In addition, they can be used to do distributed processing with any OSF DCE-compliant application and seamlessly access DSS resources across cell boundaries.

DSS directory, security, and time servers can also be used to provide services for any OSF DCE-compliant clients.

Again, interoperability between migrated and non-migrated domains are fully supported. Users of a migrated domain can access resources in non-migrated domains. In addition, users of a non-migrated domain can access resources in migrated domains.

The administration of this heterogeneous environment is similar to the multi-domain LAN Server environment. Users must be explicitly defined in each LAN Server domain that they want to access. To access a migrated domain, a user must be defined in the DSS cell registry and also must be associated with the resource domain to which the LAN Server domain was migrated. Users are responsible for keeping their passwords synchronized in each non-migrated domain and in the DSS cell.

The minimum update required to migrate an existing LAN Server domain to a DSS cell is to upgrade the LAN Server domain controller to a DSS domain controller. When you upgrade a certain server domain controller of a domain that you intend to migrate to the DSS cell, the DSS domain controller gains the ability to synchronize user and group definitions in the resource domain to the cell registry. The cell registry supports the concept of a *single user definition* which allows a user to access all resources in the cell with a single ID and password. Users no longer need to synchronize their passwords in multiple domains.

When an existing LAN Server domain is migrated to the DSS cell, its user and group definitions are migrated to the cell registry, and its resources are migrated to a separate resource domain within the cell.

The following figures show the LAN Server integration architecture and the interoperabilities. In these figures some special terms and abbreviations are used. They are explained in Table 48.

Table 48 (Page 1 of 2). Terms Used in DD-related Figures

Term	Description
ACL	Access Control List

Table 48 (Page 2 of 2). Terms Used in DD-related Figures	
Term	**Description**
AS	Additional Server
DC	Domain Controller
DCDB	Domain Controller Database
DSS	Directory and Security Server
Legacy	Legacy client means this is a regular Warp Server client or a DOS client or a Windows95 client with DOS LAN Services, but without DCE or DSS code installed.
PAC	Privilege Attribute Certificate
RPC	Remote Procedure Call
SMB	Server Message Block

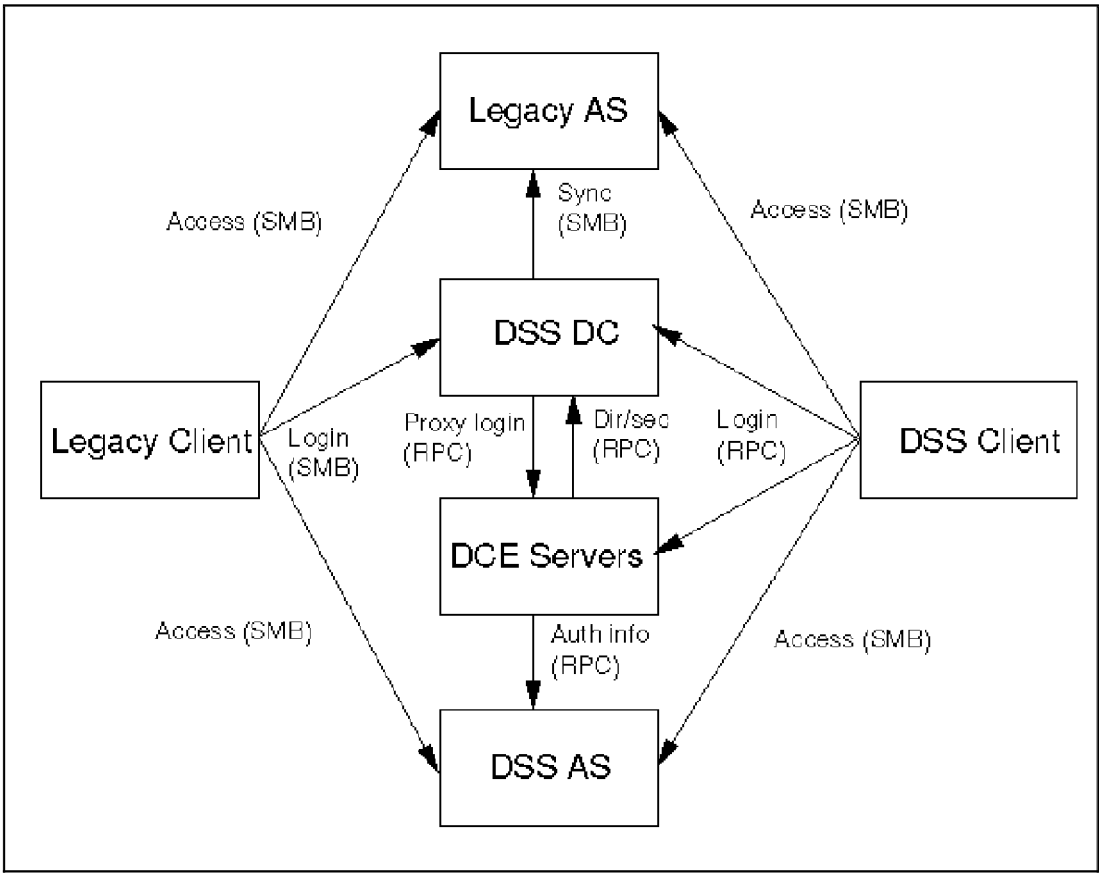

Figure 276. LAN Server Integration Architecture

Figure 276 shows the LAN Server integration architecture in an overview. Every function is explained in detail in the following figures.

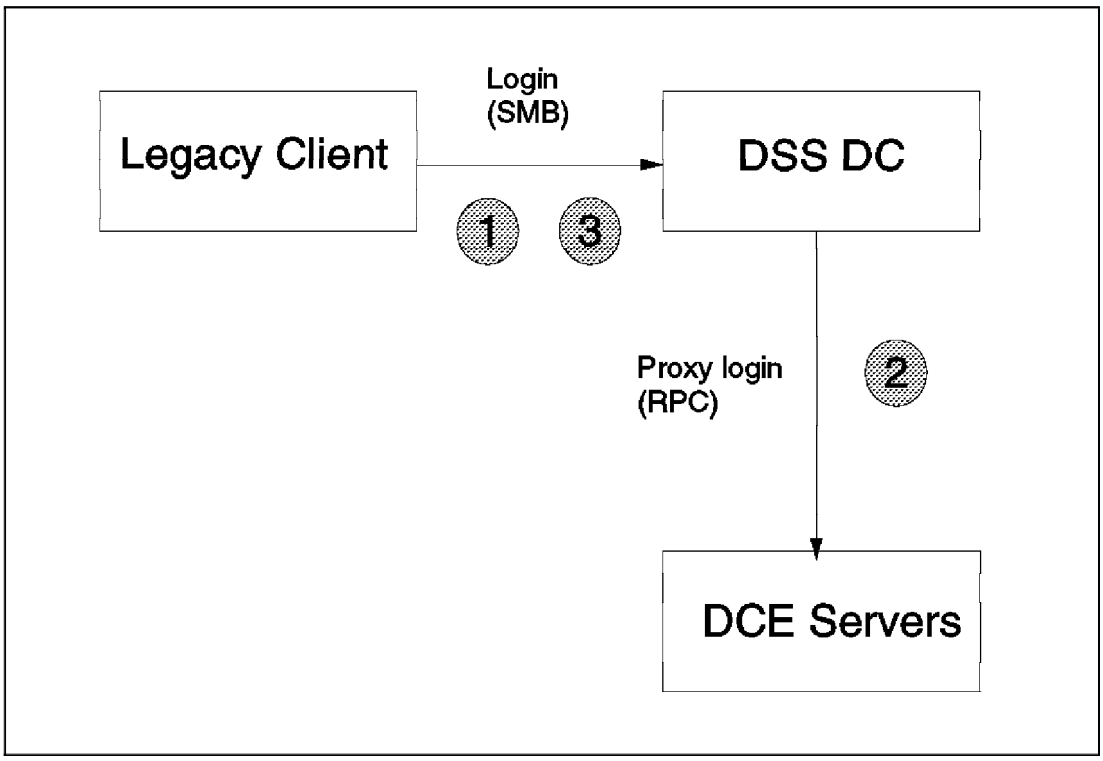

Figure 277. Legacy Client Login Flow

The legacy client login flow, shown in Figure 277 is as follows:

1. The user does a standard UPM login.
2. The domain controller does a proxy login to DCE and saves credentials. The DCE server can be any DCE server, for example an DCE server running on AIX.
3. The domain controller and the client complete LAN Server login.

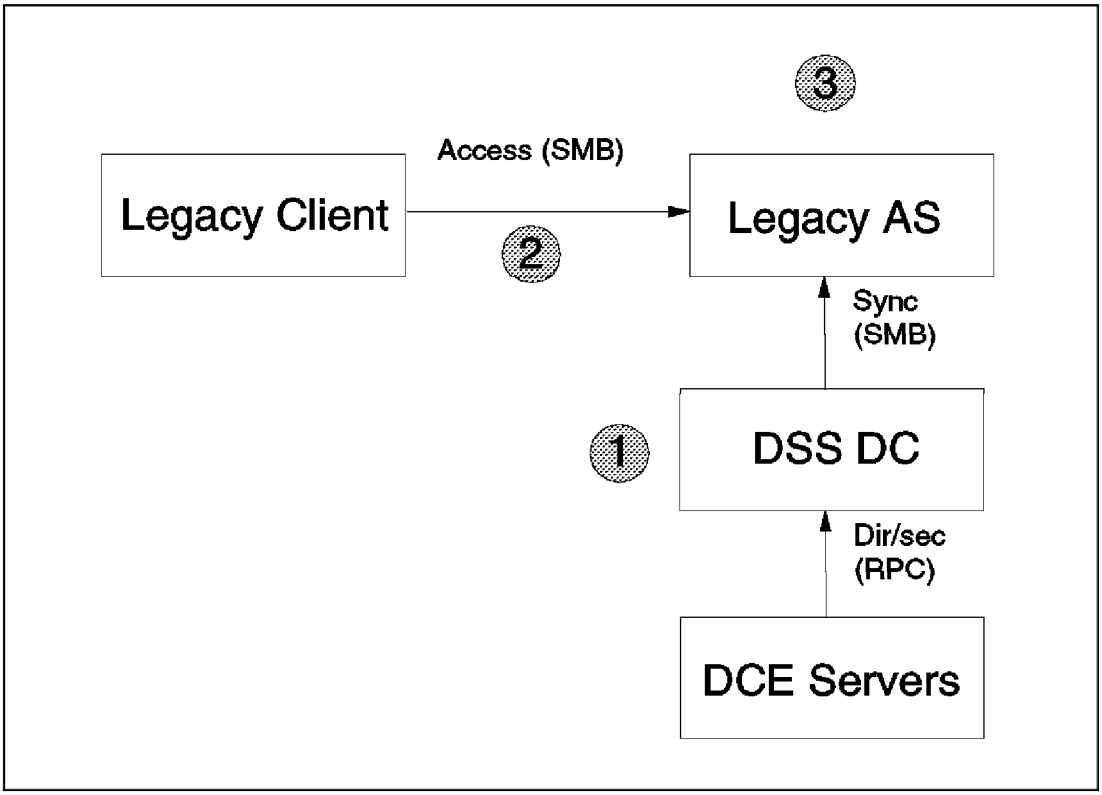

Figure 278. Legacy Client Resource Access to a Legacy Additional Server

The legacy client resource access to a legacy additional server is shown in Figure 278. The legacy additional server in this case can be an OS/2 Warp Server additional server.

1. The domain controller synchronizes DCE directory and security information with NET.ACC and DCDB and flows it to the additional server.
2. The client does a `NET USE` to the additional server which causes authentication.
3. The additional server checks local ACLs and makes a go/no go decision.

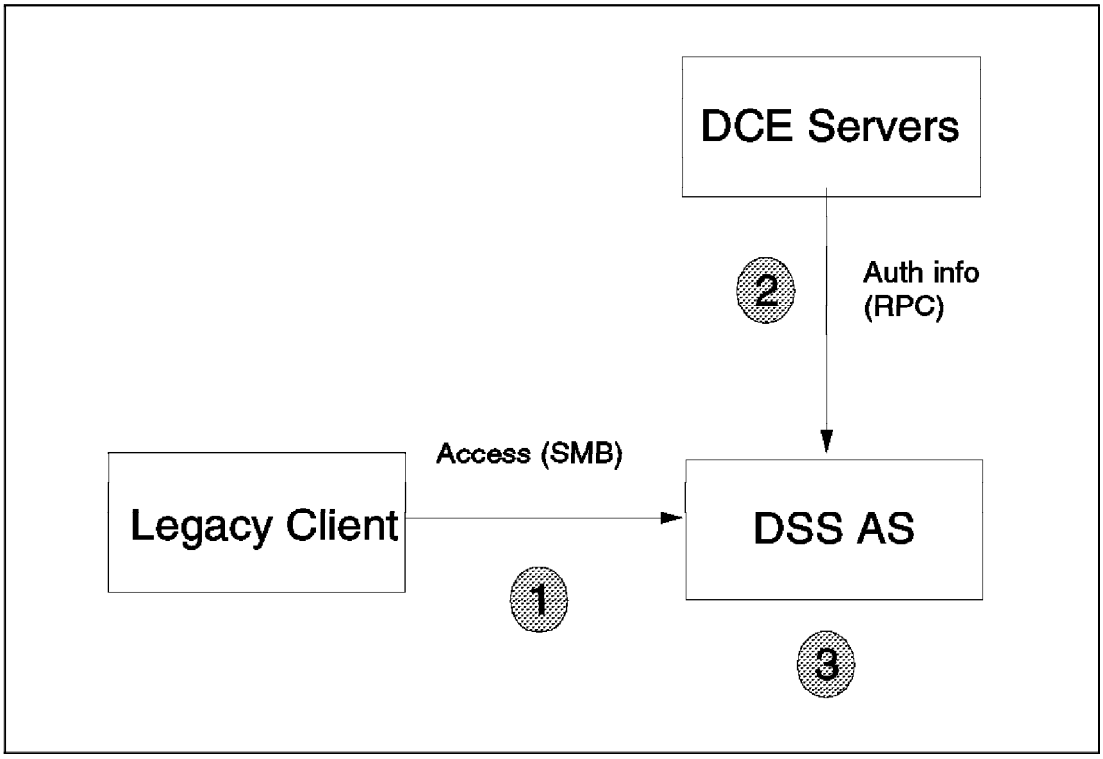

Figure 279. Legacy Client Resource Access on DSS Additional Server

The legacy client resource access to a DSS additional server is shown in Figure 279:

1. The client does a `NET USE` to a DSS additional server.
2. The first `NET USE` causes the additional server to build a pseudo Privilege Attribute Certificate (PAC) and to retrieve client security information from the DCE server.
3. The additional server checks the local access control lists (ACLs) and makes a go/no go decision.

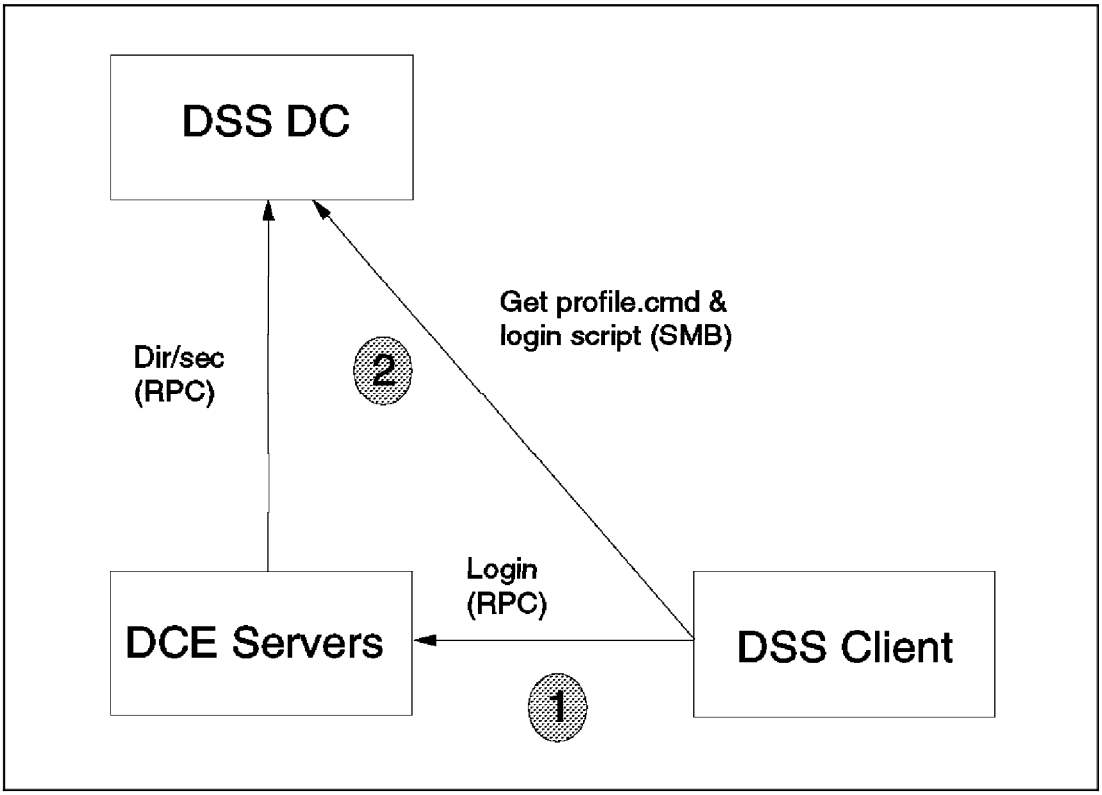

Figure 280. DSS Client Login

Figure 280 shows a DSS client login:

1. The client issues a logon. The login occurs directly to the DCE Security Server using Kerberos third-party authentication. The client gets information to complete the LAN Server login directly from DCE servers.
2. The client issues a new SMB to retrieve the profile.cmd and login script.

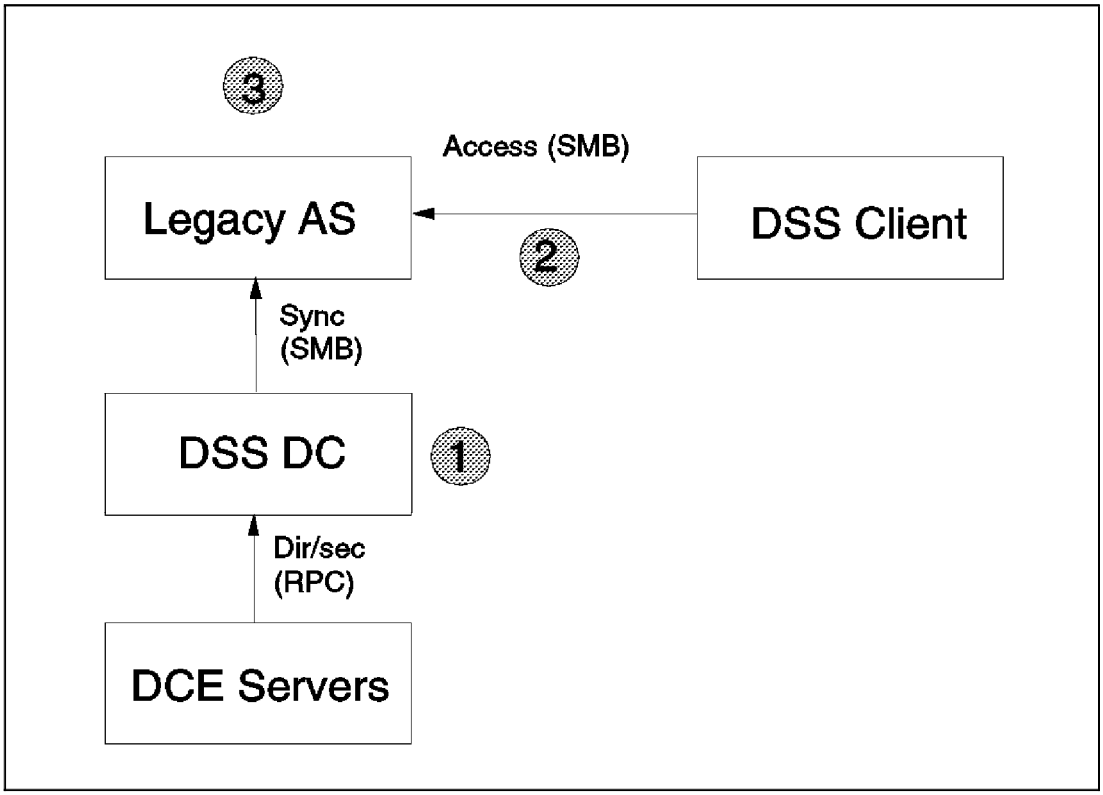

Figure 281. DSS Client Resource Access to a Legacy Additional Server

Figure 281 shows a DSS client resource access to a legacy additional server:

1. The domain controller synchronizes the DCE directory and security information with NET.ACC and DCDB and flows it to the additional server.
2. The client does a `NET USE` to the additional server which causes legacy authentication.
3. The additional server checks local ACLs and makes a go/no go decision.

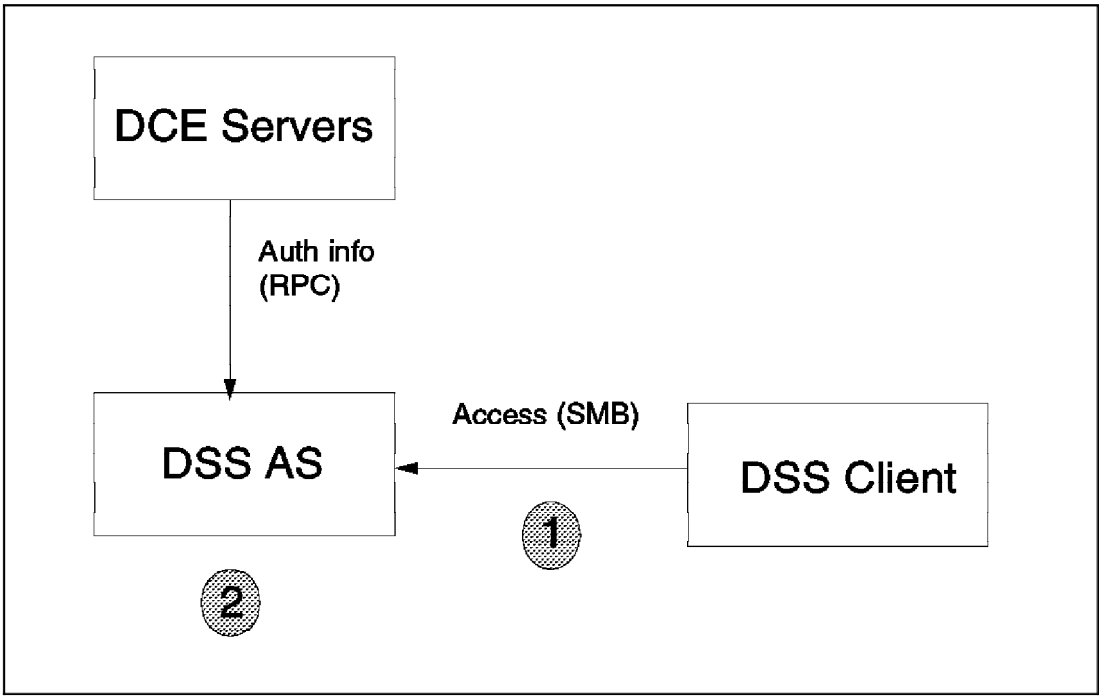

Figure 282. DSS Client Access to a DSS Additional Server

Figure 282 shows a DSS client access to a DSS additional server:

1. The client does a `NET USE` to the additional server which causes DCE authentication using PAC.
2. The additional server checks local ACLs and makes a go/no go decision.

14.1.2 Tuning Options

One of the advantages of DCE, and of using DCE Directory and Security Services with OS/2 LAN Server, is that DCE is designed from the ground up to work in a distributed, heterogeneous environment over both LANs and WANs. This presents many options for tuning the network for a wide variety of requirements. This section is not meant to be a tuning guide, but rather is intended to highlight some of the tunable features of the Directory and Security Server.

Although all of the server components of Directory and Security Server can be installed on a single machine, it is not necessary to do so. Suppose, for example, that we have a campus environment with many departments sharing several geographically close buildings. Today these are served by five OS/2 LAN Server domains averaging 800 users per domain, and we want to migrate them to the Directory and Security Server. We could simply install DSS on each domain controller, installing the same components (OS/2 LAN Server integration server feature, directory server and security

server) on each domain controller. However, a single DSS cell can comfortably support 4000 users with fewer than five directory and security servers. We could save considerable space on the domain controllers by simply installing the directory and/or security server on a single machine (this machine does not have to be one of the domain controllers, but could be), and only installing the OS/2 LAN Server integration server feature on the other domain controllers. If we add users over time and reach a point where performance has degraded, we can simply add one or two replicas of the security and/or directory server as needed. Again, they can be placed on the domain controllers or any other machine in the network that meets the hardware and software prerequisite. Because DCE uses a fair use algorithm, the load is automatically balanced between the master and replicas.

If we have a branch office environment, where the branches are connected using a TCP/IP network with a router in each branch, we can put a directory server in each branch or we can just put one in those branches that are connected via slow lines. Even though DCE uses a fair use algorithm for load balancing, we can take advantage of the fact that it tries to avoid crossing router boundaries when performing directory operations. This causes DSS to look at the directory server(s) on the local subnet first. The directory operation only looks outside the local subnet if it cannot find a server on the local subnet with the desired information.

Directory information is also cached at the client. It is possible for DCE applications to specify whether the local client cache should be used to look for information. In this way, the programmer can make the trade-off between the fast performance obtained by use of the local cache and the increased security obtained by forcing the client to access the server when information is required. Accessing the server every time data is fetched also ensures that the latest copy of the data is used.

In addition to using directory replicas, it is also possible to use registry replicas. This distributes the security operations across several servers, which increases capacity and decreases response time.

In summary, while it is not necessary to tune Directory and Security Server for small installations, when large networks are installed, Directory and Security Server offers many opportunities to tune the network to match the needs of each organization.

14.1.3 Application Development Considerations

The IBM Directory and Security Server for OS/2 Warp provides a powerful distributed application environment. The toolkit, which allows application

developers to take advantage of this environment, is shipped on the DSS product CD.

The toolkit allows the development of applications in either a "pure" DCE environment or in a DSS environment using both DCE APIs and DSS enhancements and modifications to OS/2 LAN Server APIs. The DCE portion of the toolkit contains all the standard DCE development tools, including the Interface Definition Language (IDL) compiler and the Symbols and Message Strings (SAMS) compiler. It also includes the online programming documentation and a wealth of example programs. The header files included with the DCE portion of the toolkit use standard DCE long file names in order to maximize portability. For this reason, DCE application development **must** be done on a machine using the High Performance File System (HPFS). Applications can run on either HPFS or FAT systems. The toolkit does not include a DCE client. In order to develop and test DCE applications programs, the programmer must install a DCE client from the Directory and Security Server package.

Although the DSS toolkit contains all of the OSF DCE development tools, it also contains enhancements that make it easier to create DCE applications that are intended to run only on DSS. The most significant of these is the Managed Object Class Library (MOCL), which abstracts many of the common DCE management API functions, such as creating a server instance into an object-oriented library that can be used by both object-oriented programs and more conventional C programs. Use of the MOCL rather than the analogous DCE APIs can significantly reduce the work effort required to write programs that provide DCE management functions.

The DSS toolkit also includes a full set of OS/2 LAN Server application development tools. The DSS design philosophy has been to change the OS/2 LAN Server APIs as little as possible in order to avoid breakage of existing applications. Most of the changes to the OS/2 LAN Server APIs have been done by mapping existing APIs to the appropriate DCE functions necessary to use DCE Directory and Security Services, thereby making the changes transparent to existing applications. In some cases, however, it is necessary to let DCE directory names, and so forth, show through the OS/2 LAN Server API. Where possible, this has been done by using parameters that were formerly reserved or by adding new information levels to existing APIs.

In general, the API changes for DSS are limited to these sets of APIs: Access, Alias, Application, Group, Requester, Server, User, and User Profile Management. The following is a list of some of the changes:

1. Some existing parameters that formerly accepted OS/2 LAN Server server names or aliases now also accept DCE global names.
2. Some parameters now accept DSS resource domain names.
3. Some administration APIs can no longer be issued from existing OS/2 LAN Server clients. When these APIs are issued from anything other than a DSS client, they usually return the NERR_NotPrimary return code (error number 2226).
4. Some of the NetAccess APIs are not supported for DSS servers. When directed to a DSS server, these APIs usually fail with a return code of NERR_InvalidAPI. This is because Access Control Lists (ACLs) on domain controllers and additional servers, upon which the OS/2 LAN Server integration server feature has been installed, are really DCE ACLs. In this case, it is not possible to map the OS/2 LAN Server ACL-management APIs to the equivalent DCE APIs; so the DCE ACL-management APIs must be used to manage ACLs on servers upon which the OS/2 LAN Server integration server feature is installed.

Like the DCE portion of the toolkit, the OS/2 LAN Server integration application development tools include header files and sample programs. All of the OS/2 LAN Server header files use short file names; so development can be done on FAT systems.

14.1.4 Installation Planning Scenarios

The first step in planning your Directory and Security Server installation is to decide on the number of cells you will need. Cells are logical units of administration that can contain from two to thousands of users, and the servers needed to support them. They can span many network segments and can include both LANs and WANs. Cells are not restricted by geographic boundaries, so they can be defined to include all floors of a building or multiple branch offices spread throughout the world.

There are many factors to consider when deciding how many cells to use and what users and resources to include in them. One set of factors to consider is the set of business needs of your company. If your company has multiple divisions or lines of business that are completely isolated from each other and have no need to share data, then it may be best to create a cell for each of those divisions. On the other hand, if most or many of the people in your company need to share data with each other, then it may be more useful to create a single cell for the entire company. You should also consider the way in which administrators are distributed throughout your company. If you have a great deal of administrative expertise in each major LAN site in your company today, it may make sense for you to use more than one cell. If you want to have administration done at regional sites or

the headquarters site, then it may be more efficient for you to use a single DSS cell. This is especially true if you have a multi-tiered company with varying degrees of administrative expertise at each level of the corporation.

Directory and Security Server cells are compliant with the Open Software Foundation's (OSF's) Distributed Computing Environment (DCE) Version 1.1. DSS security and directory servers can be used to provide security and directory services for any OSF DCE client. Companies that have made the decision to use DCE as the strategic company infrastructure may choose to run other DCE 1.1-compliant applications in their cells in addition to DSS. If this is the case in your company, or if you may encounter this situation in the future, you should take this into consideration when planning the number of DSS cells to use. One of the things to consider with respect to other DCE applications in DSS cells is the fact that DCE is designed to favor operations within a single cell. That means that if a set of users must frequently access both LAN Server resources and other DCE applications or resources, the best performance will probably be achieved by including those DCE applications and resources in the DSS cell. This is the best model as long as the cell servers are geographically separate from each other. If all of the cell servers are in one site (for example, in a single room) and a physical disaster (fire, flood, explosion) should occur that wiped out all of the cell servers, then all applications would be unavailable. The only way to avoid this situation is to distribute the cell servers to geographically separate areas or to use multiple cells that are geographically distributed.

Another point to consider when deciding how many cells to use is that each cell must have at least one directory server and one security server. For this reason, using multiple cells increases the amount of administration that must be done, especially if the cells are small and would normally require only one or two security and directory servers per cell.

If your Directory and Security Server installation includes existing OS/2 LAN Server servers and clients that will not be upgraded to use DSS directly, then there are other factors to consider when deciding how many cells to use. Existing OS/2 LAN Server clients can seamlessly access resources across domains within a single DSS cell. However, they cannot seamlessly access resources across DSS cells. You must use a DSS client to seamlessly access resources across cells. In addition, unchanged existing OS/2 LAN Server additional servers that are part of a DSS cell cannot make their resources available to clients in other cells. If you are migrating your installation to DSS and must intermix unchanged OS/2 LAN Server clients and servers with DSS clients and servers that have been upgraded with DSS, then you should use a single cell for all clients that must access a set of resources and all servers that contain them. If it is necessary to partition

the administration of these resources, this can be accomplished by using Directory and Security Server resource domains.

14.1.5 Examples

The following are a few examples that might make it easier to determine how many Directory and Security Server cells to use in specific situations.

14.1.5.1 Example 1 - A Two-Site Manufacturing Company

It is assumed that the Amalgamated Widget Company has two sites located in cities fifty miles apart. The home office site contains separate LANs for the accounting, accounts payable, accounts receivable, order entry, and inventory control departments as well as additional LANs in the warehouse and shipping areas and in the Widget Finishing manufacturing area. There are 1500 employees at this site.

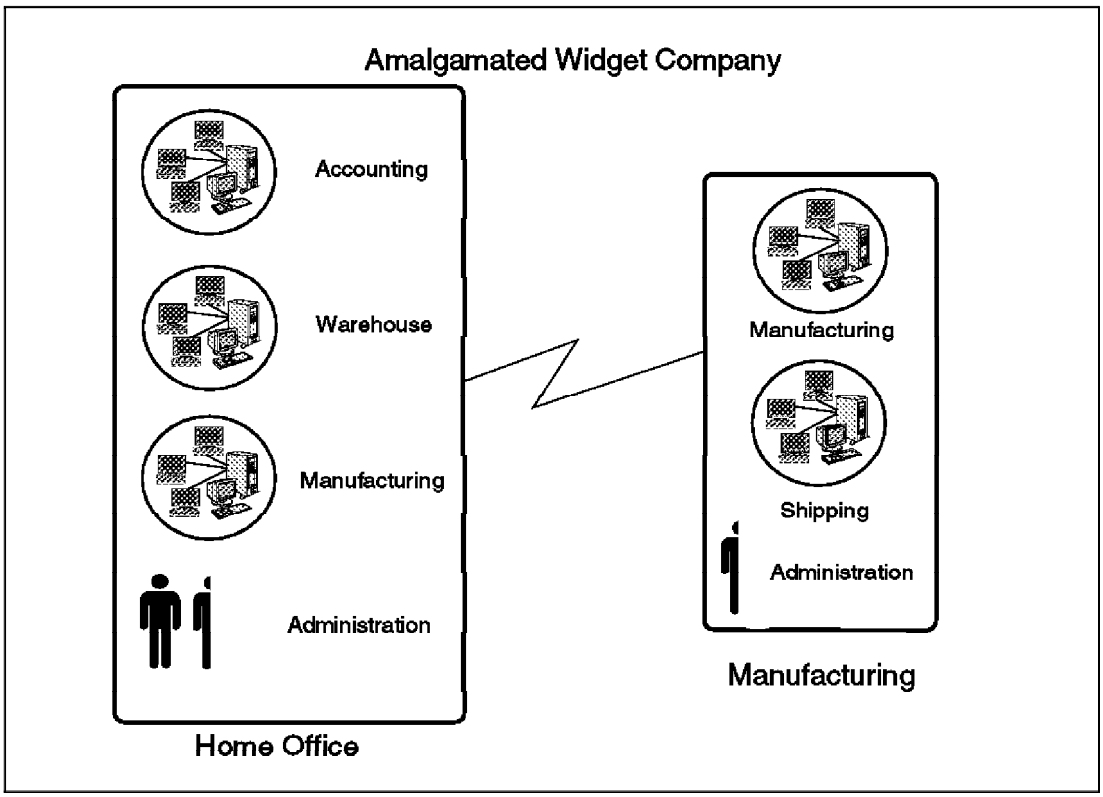

Figure 283. Example 1 - A Two-Side Manufacturing Company

The Widget Stamping manufacturing area in the second city has 500 employees and two LANs, one in the manufacturing area and one in shipping and receiving. The two sites are connected via a TCP/IP connection over a T1 line. Administration of the LANs at the manufacturing only site is done by one person on a part-time basis. There is one full-time and one part-time administrator at the home office site. Directory and

Security Server will be used to tie all of the LANs together for the purposes of consolidating administration and allowing some cross-domain resource access.

Resolution: This situation calls for a single cell. Using one cell allows the administrators at the home office to administer users, groups, and resources at both sites from a single location. In addition, Directory and Security Server resource domains can be used to restrict resource access to only those employees who have a business need to access those resources, but at the same time allow cross-domain and cross-site sharing of data and other resources for employees with a need to do so.

14.1.5.2 Example 2 - A Regional Bank

It is assumed that the New Farmers and Ranch Hands Bank has a central site with five LANs and 1000 employees and forty branches, each with a single LAN and 50 to 100 employees. The total number of employees is 4000. Branches are connected to the central site via TCP/IP using 56 KB or T1 lines, depending upon size. Most branches have a need to access data and other resources within the local branch and at the central site. Central-site employees need to access data and resources at the central site and any of the branches. There is virtually no administration done at most branches. Administration expertise is all located at the central site. Central site employees either drive to the branches for administration or talk the head teller through the activities over the phone. Directory and Security Server is used to tie all of the branches to the central site.

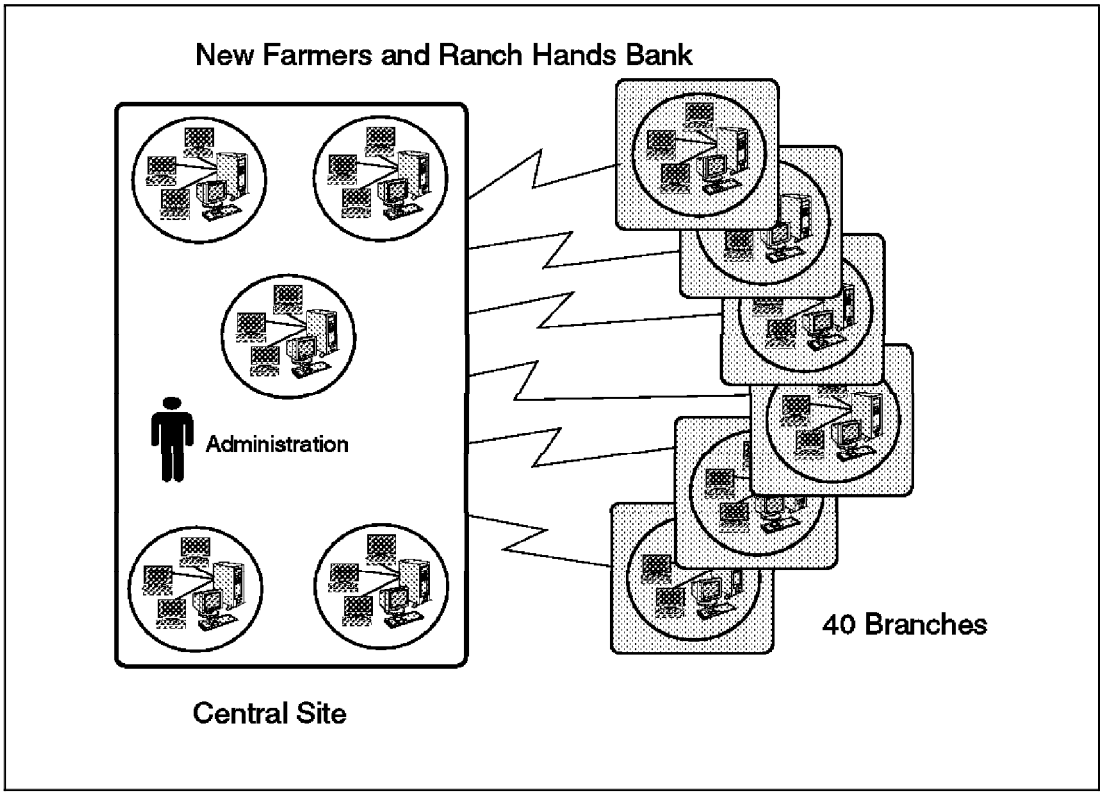

Figure 284. Example 2 - A Regional Bank

Resolution: A single Directory and Security Server cell will be used for this bank, with Directory and Security Server resource domains used to restrict branch employees access to only the data and resources in their branch and selected resources and data at the central site. Administrators at the central site will be able to access and administer the branch resource domains directly from the central site. Those employees with a business need to do so will be able to access data and resources at any branch.

14.1.5.3 Example 3 - A Three-Tiered Insurance Company

The Extremely Safe and Secure Insurance Company has a single corporate office, five regional offices, 100 claims offices, and thousands of agents in branch offices throughout the country. The branches range from five to 150 people with one or more LANs in each branch. There are dedicated administrators at the central site, at the regional sites and at claims offices. Large branch offices have at least part-time administrators on site, but smaller branch offices are administered similarly to the bank example above. In all, there are 30000 users with hundreds of LAN Server domains. Data and resource sharing needs are varied and complex. Branch offices must share data with regional offices. Regional offices and claims offices must share data with the central site. The claims offices, branch offices,

regional offices, and central site are tied together with a variety of lines of varying speed and reliability. Both NetBIOS and TCP/IP are used.

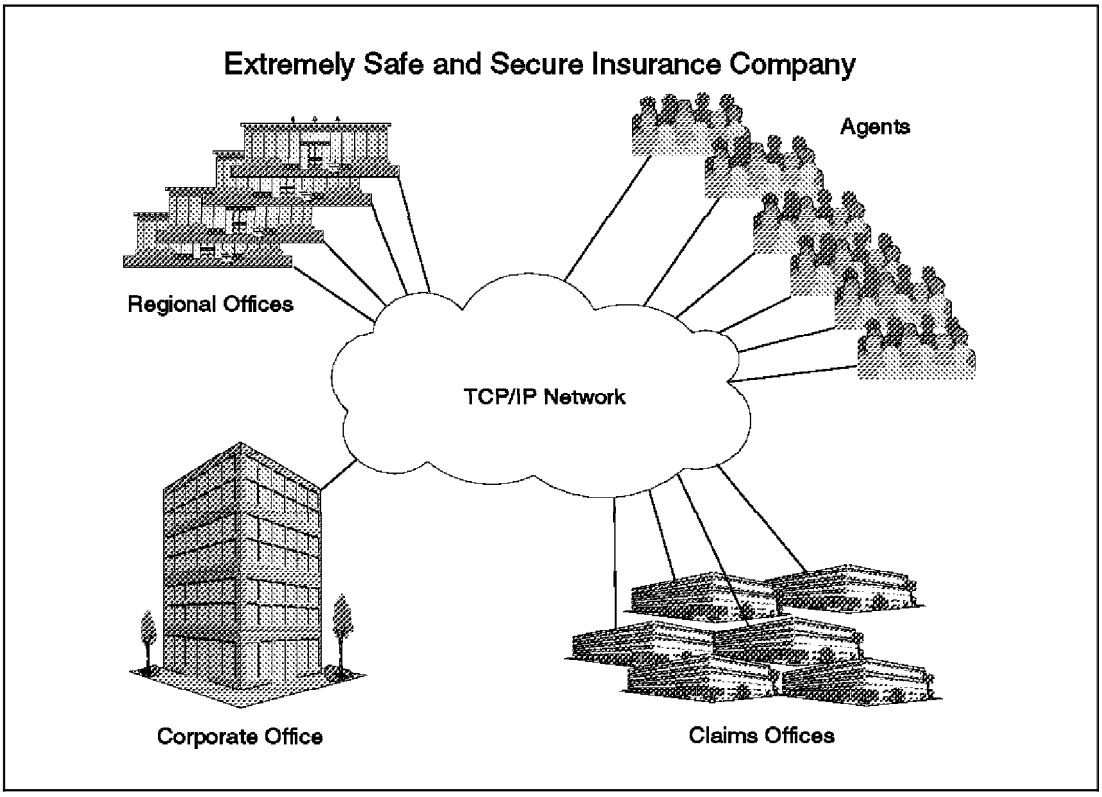

Figure 285. Example 3 - A Three-Tiered Insurance Company

Resolution: This is a case where multiple cells is probably the best answer. Although there are no architectural reasons why 30000 users could not be placed into a single Directory and Security Server cell, hardware limitations and administration complexity combine to make that impractical. Deciding how many cells should be used requires a detailed analysis of the data and resource sharing needs in the corporation. In this instance, it may be most practical to define a cell for each regional office and the branches it serves. This allows Directory and Security Server resource domains to be used to restrict branch office personnel to only those resources that they have a business need to access. At the same time, it allows hierarchical administration of large branch offices and direct administration of smaller branch offices from the regional offices. Depending upon the number of users and resources in the claims offices and their need to share data between offices, it is probably most practical to use a single cell for all of the claims offices. Again, resource domains can be used to restrict access for those employees who have no need to access data or resources across office or domain boundaries. The central site should probably be a separate cell. DSS clients can be used to access resources across cell boundaries

when necessary. These would be needed mostly in the central site, with a few in the regional and claims offices also. Most branches would not need to access resources across cells and could use unchanged OS/2 LAN Server clients.

14.1.6 Directions

IBM currently supports DCE as a strategic infrastructure. IBM is a charter member of OSF and a member of the OSF DCE 1.2 Project Steering Committee. In that role, IBM has worked with the other OSF members to enhance DCE and make it the premier infrastructure for network-centric computing. DCE is a core architecture in the company's Open Blueprint strategy and has been implemented on IBM mainframe, mini-computer, workstation, and PC platforms. The Open Blueprint states IBM's current direction with respect to DCE: to tie together all of IBM's major operating system platforms and resource managers using the common DCE directory and security services. The objectives of this strategy are twofold: to allow an end-user to access all of the resources to which he has been authorized in a single login and to ease the administration burden through the use of a single user definition. The DSS OS/2 LAN Server integration server feature is the first step in realizing this goal. Over time, IBM intends to deliver this level of integration between DCE and all of its major resource managers.

14.1.7 Summary

The IBM Directory and Security Server for OS/2 Warp delivers a powerful set of distributed computing services that help large and small enterprises move to a distributed, network-centric computing environment. These services can be used to install a "pure" DCE network on OS/2 Warp, which can interoperate with OSF DCE-compliant components on a variety of platforms, including mainframes, mini-computers, workstations, and PCs. DSS also extends OS/2 LAN Server from the workgroup to a distributed environment using DCE's powerful directory and security services. This extension is accomplished by additional function that installs on top of OS/2 Warp Server and OS/2 LAN Server 4.0 plus OS/2 Warp.

DSS allows older OS/2 LAN Server clients to take advantage of DCE services with no changes to the clients. These older clients use a protocol that is not aware of DCE. Because of this, and because the Directory and Security Server delivers an environment with a single user definition and single-signon, which strengthens the administrative control of existing OS/2 LAN Server networks, existing OS/2 LAN Server and OS/2 Warp Server clients cannot be used to administer DSS networks. The amount of administration that can be performed by existing clients depends upon the level of the client. OS/2 LAN Server 3.x clients can only change their own password. OS/2 LAN Server 4.x and OS/2 Warp Server clients can also

administer their own logon assignments, private applications, and application selector lists. All other administration must be done by a DSS client.

DSS also changes some of the OS/2 LAN Server APIs. For the most part, these changes involve new information levels and extend the use of existing parameters on existing OS/2 LAN Server APIs. In the case of ACLs on servers upon which DSS has been installed, it is necessary to use DCE APIs to administer the ACLs via a program.

DSS delivers an implementation of OSF DCE 1.1 on OS/2 Warp that is interoperable with any other OSF DCE-compatible implementation of DCE on any platform, using supported transport protocols. All of the OSF DCE 1.1 function is available, including extended registry attributes for easy integration with existing client-server applications. DCE's powerful directory and security services, industry-standard remote procedure call and DFS client are all available in the IBM Directory and Security Server for OS/2 Warp. In addition, DSS adds a graphical user interface administration tool, which allows DCE to be administered with the same ease as that provided by the award-winning OS/2 LAN Server 4.0 administration GUI.

DSS extends the reach of OS/2 LAN Server into the distributed networking environment, allowing OS/2 LAN Server users on unchanged, existing clients to access remote servers anywhere in the DSS cell via the powerful DCE directory services. Using a single identity and a single-signon, OS/2 LAN Server users on existing clients can access resources on any domain in the cell without the need to keep user IDs and passwords in synchronization across multiple domains. Users of DSS clients can seamlessly access resources in *any* cell using end-to-end Kerberos security. The workload of OS/2 LAN Server administrators is reduced because they only have to administer a single identity for each user rather than an identity in each domain. It is not necessary for the administrator to set up a trust relationship between domains in order for a user at an existing OS/2 LAN Server client to access resources in another domain in the cell.

The IBM Directory and Security Server for OS/2 Warp is the first product to implement IBM's Open Blueprint strategy. It is a large step on the way toward network-centric computing.

14.1.8 Where to Get More Information

More information about the IBM OS/2 Directory and Security Server is available online on the World Wide Web at the following URL:

`http://www.software.hosting.ibm.com/is/sw-servers/directory/`

A very detailed description about the principles, security, integration and migration can be found in the soon published Redbook titled *Inside OS/2 Directory and Security Server*, SG24-4785 (in press).

Appendix A. SOCKS (SOCKet Secured) Server

The SOCKS (SOCKet Secure) server is a firewall host that protects a business network from access outside that network. It is similar to, but provides more functions than, a proxy gateway. The SOCKS server verifies that your computer (host name) and optionally your user ID are allowed to access external networks, such as the Internet.

In this appendix, we discuss how to set up your TCP/IP environment to communicate with a SOCKS Server on your LAN.

Warp Server and SOCKS Server

Warp Server's TCP/IP services do not offer enabling a SOCKS server on your LAN. Warp 4 comes with a higher release of TCP/IP services that offer SOCKS enablement. Officially not supported, however technically okay, you can install TCP/IP services of Warp 4 on a Warp Server workstation. Be aware that if you need to have DHCP and DDNS Server, you first have to install Warp Server's TCP/IP services, and then TCP/IP services of Warp 4. All TCP/IP services but DHCP and DDNS Server will be upgraded if you do so. If you plan to do so, make sure you upgrade MPTS before as well.

A.1.1 Enabling SOCKS Server

To enable the SOCKS server on your LAN, do the following things:

1. Start **TCP/IP Configuration (LAN)**.
2. Select the **Socks** page of the TCP/IP notebook.
3. Go to the SOCKS notebook page (1 of 3).
    a. Select the **Enable SOCKS** checkbox to enable the SOCKS (SOCKet Secure) server on your LAN.
    b. Since the Enable SOCKS checkbox is marked, providing a SOCKS Userid is optional.
    c. Specify the SOCKS Domain name that will be used by your local name server to resolve names within your private business network. For example, type in:

        ```
        ibm.com
        ```

    d. Specify the SOCKS name server that will be used to resolve Internet names that are outside of your private business network. For example, type in:

```
129.35.251.68
```

assuming that 129.35.251.68 is your SOCKS name server. You can retrieve IP address information of your SOCKS server by issuing the following command:

```
HOST SOCKS
```

4. Go to the SOCKS notebook page (2 of 3) and add "Direct Routes" that will not go through the firewall.
   a. To add an entry, select the **Add** push-button. When the Direct Connections list box displays, enter the following information:
      - Destination IP Address
      - Mask

      For example, type:

      ```
      9.0.0.0 mask 255.255.255.0
      129.35.0.0 mask 255.255.0.0
      ```

5. Go to the SOCKS notebook page (3 of 3) and add a SOCKS server for traffic through the firewall.
   a. To add an entry, select the **Add** push-button. When the SOCKS Server's list box displays, enter the following information:
      - Host name or the 32-bit dotted decimal notation IP address of the SOCKS server
      - Destination IP Address for each host that can be reached through this server
      - Mask to be applied against the destination IP address

      For example, type:

      ```
      SOCKS Server 129.35.251.68
      Destination IP Address 0.0.0.0
      Mask 0.0.0.0
      ```

6. Close the TCP/IP notebook to save the additions/changes. You do not have to reboot, but you should restart any application you want to use through the firewall. For example, close and reopen the Web Explorer or Netscape (be sure to deselect **Proxy** and select **Socks** in the browser).

Two files that reside in the MPTN ETC directory are created for you:

- SOCKS.CFG

```
direct 9.0.0.0 255.255.255.0
direct 129.35.0.0 255.255.0.0
sockd @=129.35.251.68 0.0.0.0 0.0.0.0
```

Note: Comments are not allowed in the SOCKS.CFG file.

- SOCKS.ENV

```
socks_flag     on
socks_user
socks_domain    ibm.com
socks_ns       129.35.251.68
socks_server
```

Note: Comments are not allowed in the SOCKS.ENV file.

If necessary, you can do manual changes against these files to avoid going through the TCP/IP configuration panels. Use an ASCII editor to modify these files. For example, you can include multiple **direct** and **sockd** statements in the SOCKS.CFG file, but all **direct** statements must precede all **sockd** statements. Statements are processed in the order in which they appear in the file.

A.1.2 Applications that can be used with SOCKS Support

The 32-bit client applications for TCP/IP for OS/2 that can be used with SOCKS support are:

- Telnet
- TelnetPM
- FTP
- FTP-PM
- Sendmail
- NewsReader/2
- Gopher
- Web Explorer
- Netscape Browser for OS/2

The SOCKS support code attempts to determine whether a connection request is internal or external (inside or outside the firewall). If it cannot determine which type of connection is requested, it attempts to establish an external connection.

Appendix B. Understanding Bridging and Filtering

In order to customize the Warp Server's Connection Server for a non-standard configuration, it is important to understand how the bridging and filtering functions work. This section provides an overview of the bridging and filtering functions available. For a more complete discussion on bridging and filtering, please refer to the documentation referenced at the end of the appendix.

B.1.1 Remote Access Services Bridge Considerations

When setting up a Remote Access Services network, there are two points you should consider:

1. Coordinate segment numbers so there are no conflicts in the network.
2. Set the hop count appropriately to:
   - Allow remote workstations to access other systems in the network
   - Minimize the amount of unwanted traffic on the WAN link

B.1.2 Segment Numbers

The Remote Access Services bridges between two LANs. Generally, one LAN segment is a physical LAN segment (token-ring or Ethernet) and the other is the WAN, which is a virtual LAN segment. Segment numbers are used by bridges to route frames from one segment to another. All segments within a network should have an unique segment number.

It is important that the segment numbers used for the LAN and WAN segments are valid when a Remote Access Services server is configured. In a small LAN environment where there are no interconnected segments (only one physical LAN segment exists), the LAN segment number can be any valid value. If the LAN environment is large, with many interconnected segments, then the segment number for the WAN and other remote LANs must be coordinated through a network administrator. Or, you can use `OS2PING` or `CALLBRDG` to be sure you are using a unique number. These are two utilities that ship with OS/2 Warp Server.

Figure 286 on page 524 shows remote LAN segments that are interconnected using Remote Access Services. Each segment within this network must have a unique segment number so that the Remote Access Services bridges can route the frames through the network.

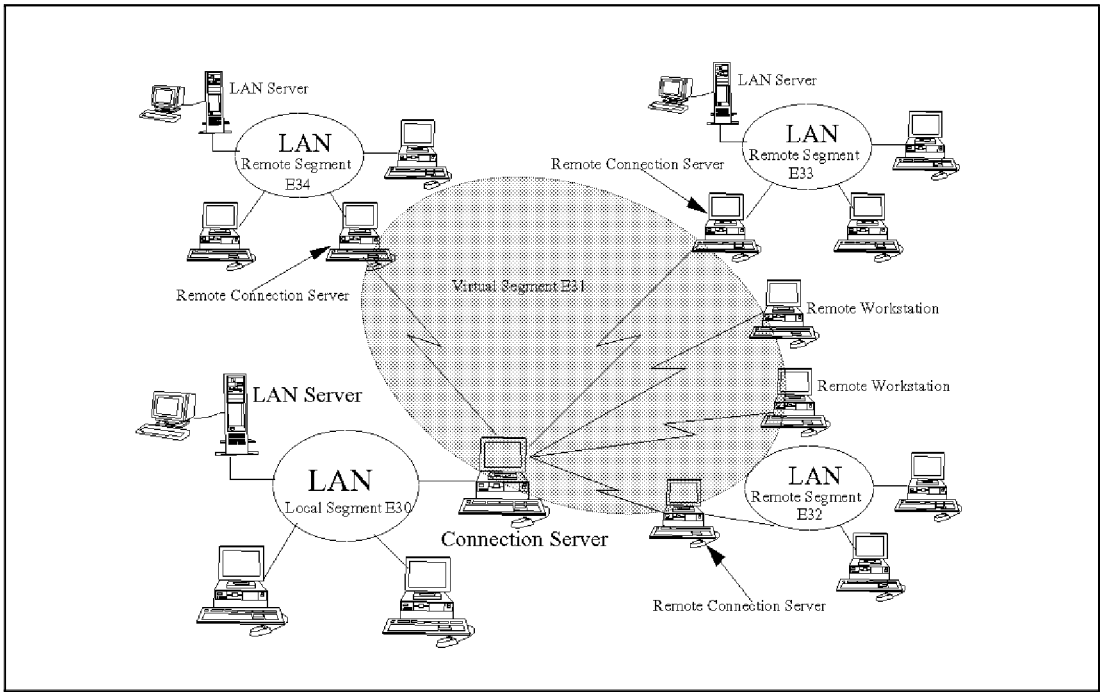

Figure 286. Interconnected LANs Using Remote Access Services

In this example, the network has five segments in total: one local LAN segment (E30), one WAN segment (E31), and three remote LAN segments (E32, E33, and E34). The following table shows how each Remote Access Services would have its segment numbers configured:

Table 49. Remote Access Services Segment Configuration

Connection Server	**LAN Segment Number**	**WAN Segment Number**
Local	E30	E31
Remote 1	E32	E31
Remote 2	E33	E31
Remote 3	E34	E31

B.1.3 Hop Counts

A hop count is used by the Remote Access Services bridge and other LAN bridges to decide whether a frame should be discarded or not. The hop count limit indicates to a bridge the maximum number of bridges a broadcast frame can traverse before it is discarded by the bridge. For the IBM Token-Ring Network Bridge Program 1.x and the Remote Access Services bridge, the hop count affects both Single Route Broadcasts *and* All Routes Broadcasts. For the IBM Token-Ring Network Bridge Program 2.x, the hop count affects *only* All Routes Broadcasts.

The two sides of the bridge, either WAN-to-LAN or LAN-to-LAN, can have different values specified for the hop count limit. The value of (7,7) means that broadcast frames arriving on both sides of the bridge could have already traversed up to six bridges and will still be allowed to traverse this bridge. The number 7 represents the maximum number of bridges that can be traversed.

In a source routing bridge environment, as each bridge in a network is crossed, the bridge adds routing data to a routing-information field. The maximum length of the routing-information field is 16 bytes, which allows a maximum of 7 bridges to add their routing information (2 bytes per bridge, plus 2 bytes of control information).

In a large LAN environment where there are many LAN segments and bridges between a Remote Access Services and perhaps a host gateway, the hop count parameter can become critical to the effective operation of Remote Access Services. Let's look at Figure 287.

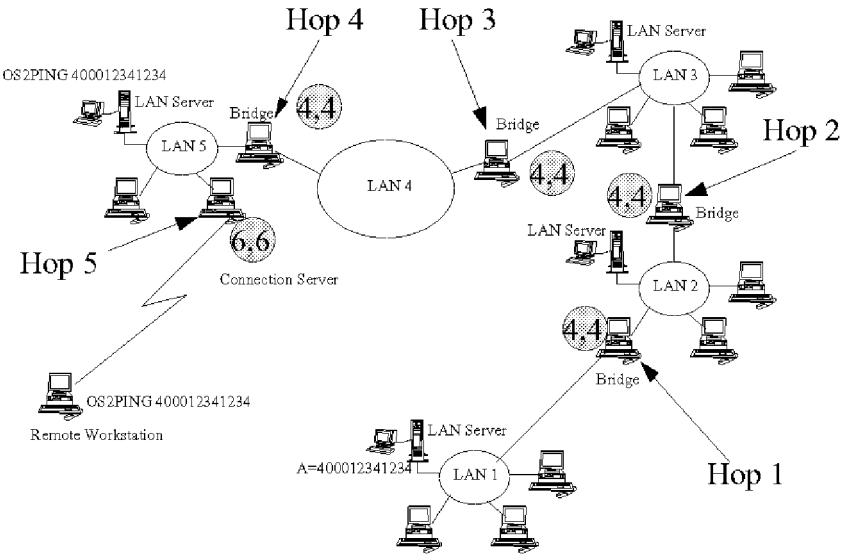

Figure 287. Setting Bridge Hop Counts

In Figure 287, the Remote Access Services on LAN 5 is four bridges away from the LAN Server on LAN 1. Each local bridge in this network has its hop count set to 4,4; thus any frame can traverse between any of the 5 LAN segments shown. When a Remote Access Services remote workstation is introduced and connects with the Remote Access Services, an additional bridge, and thus an additional hop, is introduced. Due to the additional hop

in the network, it is not possible for the remote workstation to communicate with the LAN Server on LAN 1.

With Remote Access Services, if your network is large and contains many bridges and LAN segments, it is important that you understand the network topology.

There are two trade-offs with setting the hop count on Remote Access Services:

- The hop count must be set large enough to enable the remote workstations to communicate with other devices in the network.
- A low hop count can assist with limiting the amount of unnecessary data on the slower WAN communications links.

If you set the hop count high, and there is a lot of broadcast traffic in the network, then the slow WAN link may be too busy transmitting unwanted broadcast traffic to be able to transmit information to and from the remote workstation. In this case, filtering has to be added to permit only certain adapters, NetBIOS names, or protocols to pass data over the WAN link.

If the hop count is set low, then there is less need to filter at the Remote Access Services. In fact setting the hop count low can assist or complement filtering for the purpose of reducing traffic over the WAN link. In some cases, such as with Communications Manager, it is possible to set the hop count on the Remote Access Services to 1 yet still be able to establish a connection with the host gateway.

When Communications Manager at the remote workstation first establishes a link with its gateway, it sends out an All Routes Broadcast to the gateway. As long as the bridges between the Remote Access Services and gateway have their hop count set high enough, the frame will reach the gateway. The gateway responds with a non-broadcast (direct) frame to the remote workstation. Because this frame is a non-broadcast frame, the bridges forward it through the network to the remote workstation. Even if the hop count on the Remote Access Services is set to 1, and the frame has passed over four bridges, the frame is still passed to the remote workstation by the Remote Access Services. The reason for this is that the bridge hop count only restricts broadcast frames.

B.1.4 Filtering

The Remote Access Services implements a bridge between a LAN segment and a WAN segment. As in any bridge, if filtering is not used, the Remote Access Services forwards all LAN traffic that has routing information to the WAN segment and vice-versa. If the Remote Access Services is located on

a busy LAN with large volumes of LAN traffic, the LAN side of the Remote Access Services tries to pass a large number of frames to the WAN side. Because the WAN side is usually much slower, it becomes overloaded, causing performance and connection problems.

In order to prevent these problems, the Remote Access Services allows you to control which frames can be passed between the LAN and the WAN segments. This is provided by the filtering feature. Two types of filtering are available:

- Automatic filtering

  In many cases this may be the only type of filtering needed. It provides an efficient way of preventing traffic from flooding the WAN link without requiring the user to go into the complex filter customization process. It also sets itself up specifically for each port. For example, if your Connection Server has eight ports servicing eight different remote workstations, Remote Access Services determines what type of filtering is required for each individual remote workstation and sets up eight filter criteria.

  Here are some examples of LAN traffic that the automatic filtering function filters:

  - Broadcast traffic
  - Traffic sent to functional addresses
  - Traffic sent to Ethernet multicast addresses or token-ring group addresses
  - Traffic with routing information addressed to stations that are not on the WAN segment

  Automatic filtering works for *most* NDIS-compliant protocols supported by Remote Access Services. Therefore, it can be used for most applications and LAN environments.

- Customized filtering

  The customized filtering feature provides a manual, more advanced way to control traffic flow through the bridge. Like automatic filtering, it allows you to filter frames coming from the LAN side, but adds the capability to filter frames coming from the WAN side as well.

  Customized filtering, however, applies to all ports on the Connection Server. For example, if your Connection Server has eight ports servicing eight different remote workstations, any filtering you set up will be used for all of the remote workstations.

With customized filtering, you must specify what will be filtered using panels in the Remote Access Services notebook. The Connection Server bridge supports the following filter types:

- Source addresses
- Range of source addresses
- Bit mask destination address
- Service Access Point (SAP)
- NetBIOS names

Note: You can combine the two types of filtering. It is important to remember, however, that customized filtering will apply to all ports on the Connection Server (in addition to any automatic filtering).

It is highly recommended that you use filtering to reduce the amount of traffic on your Remote Access Services WAN connections. In most cases, LAN filtering is required in order for Remote Access Services to be effective. Filtering can also be used to obtain a primary level of security by limiting access to resources on the LAN.

The recommendation is to always use some type of filtering (either automatic or customized) in order to improve connection reliability, performance, and security.

Appendix C. Special Notices

This publication is intended to help IBM systems engineers, LAN systems operators and administrators, and other LAN support staff who are responsible for the installation, setup, and customization of network operating systems. The information in this publication is not intended as the specification of any programming interfaces that are provided by products described here. See the PUBLICATIONS section of the IBM Programming Announcement for each product for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and

integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain 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.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

ADSTAR
AFP
AIX
AT
IBM
LAN Distance
NetFinity
NetView
Open Blueprint
OS/2
OS/390
OS/400
SystemView
WebExplorer
WIN-OS/2
Workplace Shell
400

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Tivoli, Tivoli/Courir, TME, and TME 10 are trademarks of Tivoli Systems, an IBM Company.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, Windows 95 and the Windows NT logo are trademarks or registered trademarks of Microsoft Corporation.

DCE, OSF	The Open Software Foundation
DEC	Digital Equipment Corporation
HP, OpenView	Hewlett-Packard Company
Intel,386, 80386, 486, 80486, Pentium	Intel Corporation
IPX, LANalyzer, NE2000, NetWare, Novell, UnixWare	Novell Corporation
Lotus 1-2-3	Lotus Development Corporation
Macintosh	Apple Computer, Inc.
NDIS	3Com Corporation and Microsoft Corporation
Netscape Navigator	Netscape Communications Corporation
NFS, Sun	Sun Microsystems, Inc.
SCO	Santa Cruz Operation, Inc.
Siemens	Siemens Company
Shadow, TCPPRO	Network TeleSystems, Inc.
Windows NT	Microsoft Corporation

Other trademarks are trademarks of their respective companies.

Appendix D. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How To Get ITSO Redbooks” on page 535.

- *How to Manage PC Server Environments*, SG24-4879
- *Inside OS/2 Directory and Security Server for OS/2 Warp*, SG24-4785 (in press)
- *Inside OS/2 Warp Server, Volume 1: Exploring the Core Components*, SG24-4602
- *Inside OS/2 Warp Server, Volume 2: System Management, Backup/Restore, Advanced Print Services*, SG24-4702
- *Using ADSM to Back Up OS/2 LAN Sever and Warp Server*, SG24-4682
- *Workgroup Management Using SystemView for OS/2*, SG24-2596
- *Software Distribution Using SystemView for OS/2 Version 1.2*, SG24-4609

D.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RISC System/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RISC System/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

D.3 Other Publications

These publications are also relevant as further information sources.

- *Understanding OSF DCE 1.1 For OS/2 and AIX*, SG24-4616 or ISBN 0-13-493750-3

How To Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL `http://www.redbooks.ibm.com`.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type `GOPHER.WTSCPOK.ITSO.IBM.COM`
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get lists of redbooks:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

  To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

  For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

  `http://w3.itso.ibm.com/redbooks`

- **IBM Direct Publications Catalog on the World Wide Web**

  `http://www.elink.ibmlink.ibm.com/pbl/pbl`

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to `announce@webster.ibmlink.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Alleroperating systemd
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1) 415 855 43 29 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to `softwareshop@vnet.ibm.com`
- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to `announce@webster.ibmlink.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

O Invoice to customer number

O Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

List of Abbreviations

ACE	Access Control Entries
ACL	Access Control List
ACP	Access Control Profile
ADF	Application Definition Profile
ADSM	Adstar Distributed Storage Manager
API	Application Programming Interface
AS	Additional Server
BBS	Bulletin Board System
BDC	Backup Domain Controller
BootP	Bootstrap Protocol
BRI	Basic Rate Interface
CDS	Cell Directory Service
CDMF	Commercial Data Masking Facility Algorithm
CHAP	Challenge Handshake Authentication Protocol
CID	Configuraion Installation and Distribution
CSD	Corrective Service Disk
DASD	Direct Access Storage Device
DBCS	DataBase Connection Services
DC	Domain Controller
DCAF	Distributed Console Access Facility
DCDB	Domain Controller Database
DCE	Distributed Computing Environment
DDNS	Dynamic Domain Name System (also Dynamic Domain Name Server)
DES	Data Encryption Standard/System
DFS	Distributed File System
DHCP	Dynamic Host Configuration Protocol
DLL	Dynamic Link Library
DLS	DOS LAN Services
DNS	Domain Name System (also Domain Name Server)
DDNS	Dynamic Domain Name System (also Dynamic Domain Name Server)
DSS	Directory and Security Server
FAP	Formats and Protocol
FAT	File Allocation Table
FTP	File Transfer Protocol
GUI	Graphical User Interface
HPFS	High Performance File System
HPFS386	High Performance File System (32-bit)
HSM	Hierarchical Storage Management
HTTP	Hypertext Transmission Protocol
IETF	Internet Engineering Task Force
IFS	Installable File System
IBM	International Business Machines Corporation
IP	Internet Protocol

IPX Internetwork Packet eXchange

ISDN Integrated Dervices Digital Network

ITSO International Technical Support Organization

LAN Local Area Network

LCN Logical Cluster Number

LPR Line Printer Requester

LSL Link Support Layer

LSMT LAN Server Management Tools

MAC Medium Access Control

MFS Mobile File System

MFT Master File Table

MHS Message Handling System

MIT Massachusetts Institute of Technology

MLID Multiple Link Interface Driver

MOCL Managed Object Class Library

MSL Mirrored Server Link

MTU Maximum Transfer Unit

NBDD NetBIOS Datagram Distributor

NBNS NetBIOS Name Server

NC Network Computing

NDIS Network Driver Interface Specification

NDS Novell Directory Services

NFS Network File Systems

NIC Network Interface Card

NLM NetWare Loadable Module

NMS NetWare Management System

NOS Network Operating System

NSC/2 Network SignON Coordinator

NTFS NT File System

NTS Network TeleSystems Corporation

NVDM NetView Distribution Manager

ODI Open Datalink Interface

OSF Open Software Foundation

PAC Privilege Attribute Certificate

PAD Packet Assembler Disassembler

PAP Password Authentication Protocol

PDC Primary Domain Controller

PIF Program Information File

PM Presentation Manager

PPP Point-to-Point Protocol

PRI Primary Rate Interface

PSnS Personally Safe and Sound (Warp Server Backup/Restore)

RAS Remote Access Services

RCD Remote Connection Server

RFC Request for Comments

RMON Remote Network Monitoring (RMONSTER/6000 program product)

RPC Remote Procedure Call

RSA MD5 Rivest-Shamir-Aleman alogorithm Message Digest 5

RWC	Remote Workstation Control
SAMS	Symbols and Message Strings
SAP	Service Access Point
SFT III	System Fault Tolerance Level III
SID	Security Identifier
SLIP	Serial Line Interface Protocol
SMB	Server Message Block
SMP	Symmetric Multiprocessing
SMS	System Management Server
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SPAP	Shiva Password Authentication Protocol
SPX	Sequenced Packet Exchange
SSI	Single System Image
TCP	Transmission Control Protocol
TCPBEUI	TCP Extended User Interface (IBM's implementation of NetBIOS over TCP/IP)
TSR	Terminate and Stay Resident
UDP	User Datagram Protocol
UNC	Universal Naming Convention
UPM	User Profile Management
URL	Universal Resource Locator
VCN	Virtual Cluster Number
VFAT	Virtual File Allocation Table
WAN	Wide Area Network
WFW	Windows for Workgroups
WINS	Windows Internet Name Service
WORM	Write Once Read Multiple/Many
WPS	Workplace Shell
WWW	World Wide Web

Index

E

F

G

H

I

K

L

M

N

O

P

R

S

T

U

V

W

IBML®

Printed in U.S.A.