NetWare Concepts

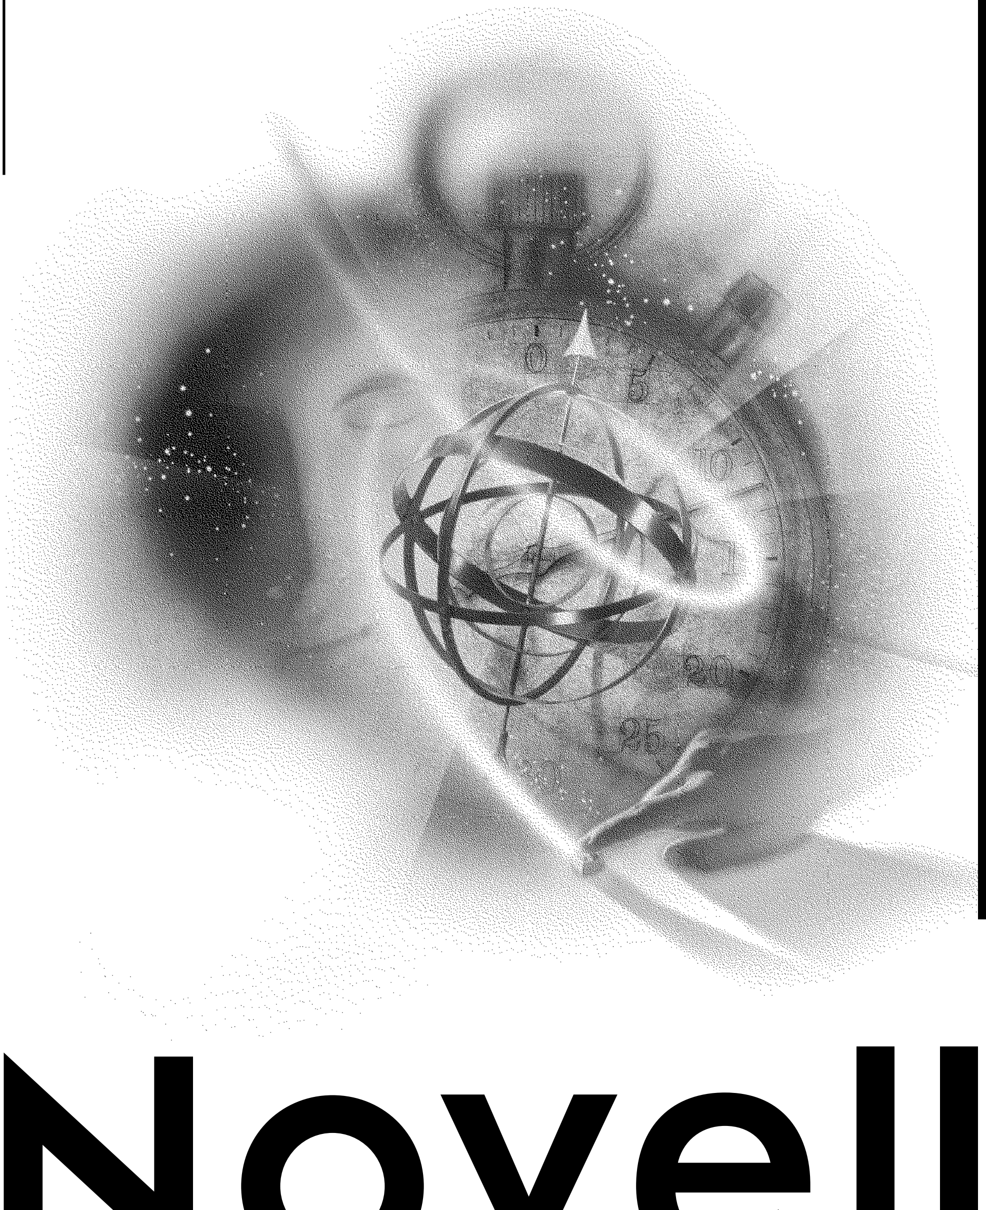

NetWare® 3.12

NETWORKING SOFTWARE

Novell®

Legal Notices

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

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

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

U.S. Patent No. 5,157,663. Patents pending.

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

www.novell.com

NetWare 3.12 Concepts
July 1993
Part Number 100-001715-001

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

Btrieve is a trademark of Novell, Inc.

DR DOS is a registered trademark of Novell, Inc.

Hot Fix is a trademark of Novell, Inc.

Internetwork Packet Exchange and IPX are trademarks of Novell, Inc.

IPX/SPX is a trademark of Novell, Inc.

Link Support Layer and LSL are trademarks of Novell, Inc.

NetWare is a registered trademark of Novell, Inc. in the United States and other countries.

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

NetWare DOS Requester is a trademark of Novell, Inc.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

NetWare Runtime is a trademark of Novell, Inc.

NetWire is a registered service mark of Novell, Inc.

Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Authorized Reseller is a collective mark of Novell, Inc.

Open Data-Link Interface and ODI are trademarks of Novell, Inc.

Packet Burst is a trademark of Novell, Inc.

SFT is a trademark of Novell, Inc.

The N design is a registered trademark of Novell, Inc.

Transaction Tracking System and TTS are trademarks of Novell, Inc.

Virtual Loadable Module and VLM are trademarks of Novell, Inc.

Third-Party Trademarks

All third-party trademarks are the property of their respective owners.

Concepts

1 A

Abend (Abnormal end)

(Abnormal end) A message issued by the operating system when it detects a serious problem, such as a hardware or software failure.

See "Server Abend Messages" in *System Messages*.

Access Control right

See “Security” on page 221.

Access privileges

See “Rights” on page 207.

Account restrictions

See “User account” on page 275.

Accounting

The process of tracking resources used on a network.

Accounting allows network supervisors to

- Charge customers who use file server resources;
- Monitor file server usage.

Using the "Accounting" option in SYSCON, you can

- Install accounting on a file server (you must activate accounting in SYSCON *after* the file server is installed);
- Select methods of charging for file server resource use;
- Select services to be charged for;
- Determine the amount to be charged for each service.

Accounting Options

After accounting is installed, NetWare® tracks user logins and logouts and stores the information in the NET$ACCT.DAT file in the SYS:SYSTEM directory.

You can choose any of the following accounting options that are suitable to your network environment.

- Charging users for resources consumed
- Tracking and charging users for specific services
- Charging users for file server disk space
- Charging users for file server time
- Charging users for file server requests (such as read or write requests)

You can set up accounting to charge for any of these services. If you install accounting but choose not to charge for services, user logins and logouts are still tracked.

You can view system accounting records with PAUDIT.

Accounting Servers

When you set up accounting in SYSCON, you authorize the file server to charge for services. You can authorize other network servers to charge for the services they provide. You can also revoke a server's right to charge.

Add accounting servers. To authorize a server to charge for services, add the server to the "Accounting Servers" list under the "Accounting" option in SYSCON.

Delete accounting servers. If you no longer want a server to charge for its services, delete the server from the "Accounting Servers" list.

Charge Rates

Types of charges. The file server can charge for five types of services:

- *Blocks Read* provides charge rates for the amount of data read from the server. You can specify the amount charged for each block read, in half-hour increments. (One block can be 4 KB, 8 KB, 16 KB, 32 KB, or 64 KB.)
- *Blocks Written* provides charge rates for the amount of data written to the disk. You can change the charge rate in half-hour increments.
- *Connect Time* provides charge rates for the amount of time a user is logged in to a server. You can change the amount charged in half-hour increments. The charge is assigned per minute.
- *Disk Storage* provides charge rates for each block that is stored on the disk for one day. You can specify different charge rates for each half-hour increment. The charge is assigned by blocks/day (the number of blocks stored in a day).

  (We recommend specifying a special time of the day to charge for blocks in disk storage, such as 6 p.m., after most users have completed their day's work.)
- *Service Requests* provides charge rates in half-hour increments for service requests. A service request is any request made of the file server, such as listing the files in a directory. The charge is assigned per request reviewed.

Calculate charge rates. If you plan to charge users for services, you must calculate the amount to charge. The amount charged depends on the network environment.

Before setting charge rates for services, you should do the following:

- Monitor the file server to determine what your costs are and the amount you want to charge over a given period of time.
- Determine what services to charge for and the amount to be made from each service.

  For example, if file server storage capacity is a concern, you should charge for disk storage. If network utilization is high, you should charge for service requests. To prevent users from staying logged in when they are not working, charge for connect time.

- Estimate how much each service is being used by monitoring the file server for two or three weeks.

  For example, if 30% of the file server's charges stem from service requests, you should recoup 30% of the cost through charging for service requests.

- At the end of the monitoring period, use ATOTAL to determine total usage for each service.

  A screen appears listing the total daily and weekly usage of each service.

After determining the amount to be charged for each service and how much each service was used, you can calculate the charge rates.

The charge rate is the charge per unit of the specified service.

Charge rates are specified as multipliers and divisors. This multiplier/divisor ratio is used to convert the amount of service usage to a monetary charge.

The unit of charge is arbitrary, but we suggest you begin with one charge unit equaling one cent. Adjust this ratio if it does not work for your network environment.

Use the formula in "Charge rate formula" on page 18 to calculate a charge rate for services.

Figure 1 Charge rate formula

$$\frac{\begin{array}{c}\textbf{CHARGE}\\ \text{(charge rate multiplier)}\end{array}}{\begin{array}{c}\textbf{ESTIMATED USAGE}\\ \text{(charge rate divider)}\end{array}} = \textbf{CHARGE RATE}$$

For example, if you want to charge $100 a month for blocks read and you find that 250,000 blocks are being read each month, then the charge rate would be 100 dollars divided by 250,000 blocks, or 4 cents per block read.

You must make the necessary conversion to cents (assuming one charge is equal to one cent) per block.

See "SYSCON" in *Utilities Reference* for examples on how to assign charge rates for services.

Account Balances

As the supervisor, you can

- Assign an individual user an account balance that determines how much of a given service the user can use.
- Assign a credit limit indicating how much credit the user can draw upon.
- Assign a system or default account balance. An account balance is assigned automatically to any user created after the default account balance is set up.
- Increase a user's account balance. The user must log out and log in again before any account balance changes are put into effect.

  IMPORTANT: If you install accounting on a file server, you must carefully monitor account balances. Warn users that if they are told to log out because their account balances are too low, they should log out immediately.

  IMPORTANT: If users do not log out, the file server logs them out and they lose any data that has not been saved.

Remove Accounting

To deactivate and completely remove the accounting feature from the file server, you must delete all accounting servers. After deleting the last accounting server, you can remove accounting.

Related utilities: "ATOTAL"; "PAUDIT"; "SYSCON" (*Utilities Reference*).

Active hub

A device used to amplify transmission signals in certain network topologies.

You can use an active hub to add workstations to a network or to lengthen the cable distance between workstations and the file server.

See also “Passive hub” on page 178.

Add-on board

An optional circuit board that modifies or enhances a personal computer's capabilities.

See also “Media Manager” on page 140; “Network board” on page 165.

Address

See “Base I/O address” on page 29; “Base memory address” on page 29; “Controller address” on page 47; “Network numbering” on page 165; “SCSI bus” on page 218.

Application

A software program or program package that makes calls to the operating system and manipulates data files, thus allowing a user to perform a specific task (such as accounting or word processing).

Standalone application. An application that runs from the hard disk or floppy disk in a self-contained, independent computer. Only one user can access the application.

Network application. An application that runs on networked computers and can be shared by users.

Network applications can use network resources, such as printers. Advanced network applications (such as electronic mail) allow communication among network users.

Archive

A transfer of files to long-term storage media, such as optical disks or magnetic tape.

See also “Backup” on page 25.

Archive Needed attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Asynchronous transmission

See “Serial communication” on page 246.

Attach

To establish a connection between workstation and file server.

The server assigns each station a connection number and attaches each station to its LOGIN directory. When you log in, your workstation is attached to the nearest file server.

Attributes

The characteristics of a directory or file that tell NetWare what to do with the directory or file; also called *flags*.

Attributes override trustee, file, and directory rights and prevent tasks that effective rights would allow.

For example, attributes can be used to prevent the following:

- Deleting a file or a directory
- Copying a file
- Viewing a file or a directory
- Writing to a file

Attributes are also used for the following:

- Controlling whether files can be shared, so that only one or many users can access the file at the same time.
- Marking files as modified so that backup utilities can select only the files that have been modified.
- Protecting files from data corruption by ensuring that either all changes are made or no changes are made when a file is being modified.

If users have the Modify right for a directory or a file, they can change the attributes and complete any task allowed with their effective rights.

NetWare uses the following attributes (represented by the indicated letters). "Directory attributes" on page 22 lists the directory attributes; "File attributes" on page 22 lists the file attributes.

Table 1 Directory attributes

Directory Attribute	Syntax	Description
Delete Inhibit	D	The directory can't be deleted.
Hidden	H	The directory can't be seen with a DOS DIR command. It can't be deleted.
Purge	P	Files that are deleted from the directory are purged immediately.
Rename Inhibit	R	The directory can't be renamed.
System	SY	The directory can't be seen with a DOS DIR command; it also can't be copied or deleted.

Table 2 File attributes

File Attribute	Syntax	Description
Archive Needed	A	The file has been modified since the last back up. This is DOS's archive bit.
Copy Inhibit	C	The file can't be copied. (Applies to Macintosh® files only.)
Delete Inhibit	D	The file can't be deleted or copied over.
Execute Only	X	The file cannot be copied or copied over. This attribute can be given only to .EXE or .COM files, and cannot be removed. Some programs do not execute correctly if files are flagged X.
Hidden	H	The file can't be seen with the DOS DIR command. It can't be copied or deleted.
Indexed	I	The file is indexed for faster access. (Assigned automatically to files with over 64 regular FAT entries.)
Purge	P	If the file is deleted, it is purged immediately.
Read Audit	Ra	Has no effect in NetWare v3.12.
Read Only	Ro	The file can only be read; it cannot be written to or deleted. NetWare assigns Delete Inhibit and Rename Inhibit automatically with Read Only.
Read Write	Rw	The file can be read and written to.

File Attribute	Syntax	Description
Rename Inhibit	R	The file can't be renamed.
Shareable	SH	The file can be used by several users simultaneously.
System	SY	The file can't be seen with the DOS DIR command. It can't be copied or deleted.
Transactional	T	The file is protected by the Transaction Tracking System™ (TTS™).
Write Audit	Wa	Has no effect in NetWare v3.12.

The initial letters of these attributes are displayed between brackets by NetWare utilities:

```
[Ro S A X H Sy I T P Ra Wa C D R]
```

By convention, attributes that have *not* been assigned are indicated by blank spaces, as in the following:

```
[Ro S                        D R]
```

Related utilities: "FILER"; "FLAG"; "FLAGDIR"; "NDIR" (*Utilities Reference*).

Automatic rollback

A feature of the Transaction Tracking System (TTS) that returns a database to its original state and abandons the current transactions.

When a network running under TTS fails in the middle of a transaction, the database is "rolled back" to its most recent complete state, preventing corruption from the partially complete transaction.

See also “Backing out” on page 25; “TSA resources” on page 268.

2 B

Backing out

Abandoning an incomplete database transaction because of system failure.

NetWare's Transaction Tracking System (TTS) views a sequence of database changes as a single transaction that must be wholly completed or wholly "backed out" (no changes made at all).

TTS stores all the information necessary to back out of a transaction and return the database to its previous state.

See also “Automatic rollback” on page 23; “TSA resources” on page 268.

Backup

A duplicate of data (file, directory, volume), copied to a storage device (floppy diskette, cartridge tape, hard disk).

A backup can be retrieved and restored if the original is corrupted or destroyed.

The type of backup you perform and the storage media rotation method you use are dictated by

- The number of backup sessions you are willing to restore in the event of data loss;
- The number of duplicate copies of data you want and are willing to store;
- The desired age of the oldest data copy.

Perform backups when the fewest files are likely to be open. (Files in use at the time of the backup aren't backed up.)

How Often to Back Up Files

Files that don't change frequently, such as applications or archived files, don't need to be backed up as often as files that change often.

In deciding which files to back up and how often to back them up, imagine a worst-case scenario: how much time and money would it take to recreate critical information if an unexpected failure were to cause data loss at the worst possible time.

How File Restoration Decisions Determine Backup Methods

Backup methods have different implications for the process of restoring files. Before you decide which backup method works best, make sure you understand the implications.

Rotation Methods

After you decide what kind of backups you will perform, determine how many sets of storage devices you need, and how you will rotate them.

Rotation distributes both current and older data across several storage devices, thereby reducing the risk of all data being lost if one of the devices is corrupted.

Example: The Grandfather Rotation Method

To use the Grandfather rotation method, you need 21 sets of storage media.

Label four "daily" sets Monday, Tuesday, Wednesday, and Thursday. Label four "weekly" sets Friday1, Friday2, Friday3, and Friday4. Label the other twelve "monthly" sets January, February, etc.

NOTE: You may want to add another set of weekly tapes, labeled Friday5, for months in which there are five Fridays.

The Grandfather method of rotation (showing the tape set names) is illustrated in Table 3.

Table 3 The Grandfather method of rotation

Daily	Daily	Daily	Daily	Weekly	Monthly
Monday	Tuesday	Wednesday	Thursday	Friday1	
Monday	Tuesday	Wednesday	Thursday	Friday2	
Monday	Tuesday	Wednesday	Thursday	Friday3	
Monday	Tuesday	Wednesday	Thursday	Friday4	
					January
Monday	Tuesday	Wednesday	Thursday	Friday1	
Monday	Tuesday	Wednesday	Thursday	Friday2	
Monday	Tuesday	Wednesday	Thursday	Friday3	
Monday	Tuesday	Wednesday	Thursday	Friday4	
					February

Example: The 10-tape rotation method

To use the 10-tape rotation method, you need 10 tape sets, each set labeled with a number from 1 through 10.

In this method, a 40-week period is divided into 10 four-week cycles, and each tape set is used an equal number of times during the 40 weeks.

You always have a 12-week-old copy of your data on at least one tape set.

For the first four weeks, use the same tape sets for the Monday (set 1), Tuesday (set 2), Wednesday (set 3), and Thursday (set 4) backups.

On the first four Fridays, use the next sequence of four tapes (5 through 8).

During the second four-week cycle, increment the daily tape set numbers by one, for example, Monday (set 2), Tuesday (set 3), Wednesday (set 4), and Thursday (set 5).

During the second four weeks, increment the Friday tape set number by one also (sets 6 through 9).

The 10-tape rotation method (showing the tape set numbers) is illustrated in Table 4.

Table 4 The 10-tape rotation method

Week 1 M, T, W, Th, F	Week 2 M, T, W, Th, F	Week 3 M, T, W, Th, F	Week 4 M, T, W, Th, F
1, 2, 3, 4, 5	1, 2, 3, 4, 6	1, 2, 3, 4, 7	1, 2, 3, 4, 8
2, 3, 4, 5, 6	2, 3, 4, 5, 7	2, 3, 4, 5, 8	2, 3, 4, 5, 9
3, 4, 5, 6, 7	3, 4, 5, 6, 8	3, 4, 5, 6, 9	3, 4, 5, 6, 10
4, 5, 6, 7, 8	4, 5, 6, 7, 9	4, 5, 6, 7, 10	4, 5, 6, 7, 1
5, 6, 7, 8, 9	5, 6, 7, 8, 10	5, 6, 7, 8, 1	5, 6, 7, 8, 2
6, 7, 8, 9, 10	6, 7, 8, 9, 1	6, 7, 8, 9, 2	6, 7, 8, 9, 3
7, 8, 9, 10, 1	7, 8, 9, 10, 2	7, 8, 9, 10, 3	7, 8, 9, 10, 4
8, 9, 10, 1, 2	8, 9, 10, 1, 3	8, 9, 10, 1, 4	8, 9, 10, 1, 5
9, 10, 1, 2, 3	9, 10, 1, 2, 4	9, 10, 1, 2, 5	9, 10, 1, 2, 6
10, 1, 2, 3, 4	10, 1, 2, 3, 5	10, 1, 2, 3, 6	10, 1, 2, 3, 7

NOTE: To make sure that you have a four-week-old copy of data at the end of the first four-week cycle, back up to tape set 10 as well as to tape set 1 on Monday of the first week.

Keep a Written Backup Log

Keep a written log of all backups performed. The log serves as a record in case electronic log and error files are destroyed.

Record the date, backup type, what was backed up, the media set identification name or number, session log path, data path, and the initials of the person performing the backup.

Recording the data path in a log makes it easier to provide this information if you want to restore the session to a different location than it was backed up from.

See also "Data set" on page 52.

Backup hosts and targets

A backup host is a file server that has a storage device and a storage device controller attached.

A *target* is a server that contains data you back up or a server that you restore data to.

Any server, workstation, or service on the network can be a target, as long as it contains Target Service Agent (TSA) files.

A target can be any server, even the host. If you back up data on the host, the target and host are the same.

See also "Backup" on page 25; "Target Service Agent" on page 259.

Base I/O address

The beginning address of an I/O port.

The base I/O address allows the microprocessor to find the correct port for communicating with a particular device.

See also "Base memory address" on page 29; "Direct memory access" on page 57; "Interrupt" on page 121.

Base memory address

A configuration option available on many network boards.

Network boards often use the base memory address as a buffer. When the network or device attached to the board sends information before the processor is ready, the information is placed in a buffer.

Since the address for each device is unique, the memory address of each buffer should also be unique. A common source for hardware conflicts within a machine is having two devices trying to use the same memory address for a buffer.

See also “Base I/O address” on page 29; “Configuration options” on page 43; “Direct memory access” on page 57; “Interrupt” on page 121.

Baud rate

See also “Printing” on page 183; “Serial communication” on page 246.

Bindery

A database that contains definitions for entities such as users, groups, and workgroups.

The bindery allows the network supervisor to design an organized and secure operating environment based on the individual requirements of each of these network entities.

The bindery has three components:

- *Objects* represent physical or logical entities, including users, user groups, workgroups, file servers, print servers, and any other entity that has been given a name.
- *Properties* are the characteristics of each bindery object. Passwords, account restrictions, account balances, internetwork addresses, lists of authorized clients, and group members are all bindery properties.
- *Property data sets* are the values assigned to an entity's bindery properties.

The NetWare v3.12 bindery consists of three separate files located in the SYS:SYSTEM directory: NET$OBJ.SYS (for objects), NET$PROP.SYS (for properties), and NET$VAL.SYS (for property data sets).

How the Bindery Works

When user STEVE logs in to a file server, the LOGIN program looks in the NET$OBJ.SYS for the object name to determine if he is a valid user.

If an object named STEVE exists, the program looks in the NET$PROP.SYS file for the properties associated with that object (in this case, to see if a password property for user STEVE exists).

If STEVE has a password property, the program prompts him for his password and compares this value with the value in the NET$VAL.SYS file that is assigned to the password property.

If the two values match, STEVE is logged in and allowed to use that network's resources according to the values of other properties (such as account restrictions and trustee assignments) that exist for user STEVE.

See also “Object” on page 173; “Property” on page 196.

Binding and unbinding

The process of assigning and removing communication protocols to and from network boards and LAN drivers.

Each network board needs at least one communication protocol bound to the LAN driver for that board. Without a communication protocol, the LAN driver cannot process packets.

You can

- Bind more than one protocol to the same LAN driver and network board.
- Bind the same protocol stack to multiple LAN drivers.
- Cable workstations with different protocols to the same cabling scheme.

Binding Communication Protocols to Boards and Drivers

To bind communication protocols to network boards and drivers, use BIND. You must enter a BIND command for each network board in the file server.

Until a protocol is bound to the network board, users attached to the cabling scheme from the network board cannot log in to the file server.

When you bind a protocol to a network board, you are prompted for the cabling scheme's network number. This number must be a hexadecimal number different from all other network numbers for cabling schemes or file servers.

See “Network numbering” on page 165 for a complete explanation of how to determine network numbers. See "Bind a Protocol to a Driver" in *System Administration* for an explanation of how to bind a LAN driver to a network board.

Unbinding Communication Protocols from Boards and Drivers

To remove a communication protocol from a network board and driver, use UNBIND. You must enter an UNBIND command for each network board you want to unbind.

If the driver has been entered "re-entrantly" (that is, if you have loaded the driver more than once), you are prompted to select the particular board you want to unbind.

When the protocol is unbound, users attached to the cabling scheme of the network board cannot log in. If they are already logged in, they receive an error message when they try to access the file server for an additional resource.

To connect to the file server, users should do the following:

- Log in again to see if there is an alternate bridge to the file server.
- Wait until the network supervisor binds the network board again. Binding the network board within 15 minutes of unbinding it enables users to retry and re-establish their connections to the file server.

See "Unbind a Protocol from a Driver" in *System Administration* for an explanation of how to remove a communication protocol from a network board.

See also “Communication protocols” on page 41; “Loading and unloading” on page 126; “Network numbering” on page 165.

BIOS

(Basic Input/Output System) A set of programs, usually in firmware, that enables each computer's central processing unit to communicate with printers, disks, keyboards, consoles, and other attached input and output devices.

Block

A unit of stored data on a NetWare volume. A block is the smallest amount of disk space that can be allocated on a volume for a file.

In NetWare, the default block size is 4 KB, or 4,096 bytes, of data. For example, a 40MB hard disk contains roughly 10,000 blocks of data storage area.

You may choose to initialize a volume with a larger block size on larger hard disks if you don't have many small files.

Boot files

Files, like AUTOEXEC.BAT and CONFIG.SYS, that

- Start the operating system and its drivers;
- Set environment variables;
- Load NetWare.

File server boot files include

- *AUTOEXEC.NCF.* Loads modules and sets the NetWare operating system configuration.
- *STARTUP.NCF.* Loads the server's disk driver and name spaces and some SET parameters.

Workstation boot files depend on the workstation type (DOS, Windows, OS/2®, Macintosh, UNIX®). For more information, see your workstation manual.

The Boot Process

To plan your boot files, you need to understand what happens when you turn on a network station or a file server.

POST routine. When the workstation is booted, the power-on self-test (POST) routine built into the ROM-BIOS checks all peripherals (memory, monitor, keyboard, printer, and any hardware installed in the expansion slots).

Boot record. ROM-BIOS determines which device to boot from by checking the boot record. The boot record can be on either a floppy diskette or a local hard disk. ROM-BIOS then loads a short program from the boot record to determine the disk format and the location of system files and directories.

System files. Using the information in the boot record, the ROM-BIOS loads the COMMAND.COM command processor and the system files (including the two hidden files, IBMBIO.COM and IBMDOS.COM).

DOS files. When IBMBIO.COM is loaded, it checks for a CONFIG.SYS file. CONFIG.SYS contains commands to load device drivers and to set parameters for the initial environment. (If you do not create a configuration file, DOS assigns default values.)

When COMMAND.COM is loaded, it checks for an AUTOEXEC.BAT file and executes it. (AUTOEXEC.BAT can contain programs, utilities, and DOS commands.)

Optional DOS files. The AUTOEXEC.BAT file can also load NETBIOS.COM, INT2F.COM (for applications that require NetBIOS from IBM®), and any additional files required by the hardware.

NetWare boot files. Up to this point, a network computer boots like a standalone computer. However, the AUTOEXEC.BAT files for network workstations contain extra commands to log in to the network.

Workstation boot files. Workstation boot files depend on the workstation type (DOS, Windows, OS/2, Macintosh, UNIX). For more information, see your workstation manual.

File server boot files. After the file server's AUTOEXEC.BAT file boots the NetWare operating system (SERVER.EXE), the file server executes the AUTOEXEC.NCF and STARTUP.NCF files.

The files are usually created with INSTALL during installation or with a DOS text editor. The AUTOEXEC.NCF file can also be created in SYSCON.

The AUTOEXEC.NCF file stores the server name and internal network number, loads the LAN drivers and settings for the network boards, and binds the protocols to the installed drivers. The network supervisor may also add executable file server commands (such as LOAD INSTALL or LOAD MONITOR) to this file.

The STARTUP.NCF file loads the server's disk driver.

See also "Load INSTALL" in *System Administration*.

How System Files Are Created

Each boot diskette or workstation hard disk that is formatted with the DOS FORMAT /S parameter contains the following system files.

COMMAND.COM
IBMBIO.COM
IBMDOS.COM

IBMBIO.COM and IBMDOS.COM are hidden files and do not appear when you execute the DIR command.

If the hard disk has already been formatted with the SELECT program, use the DOS SYS command to copy IBMBIO.COM and IBMDOS.COM to the disk.

File Server Boot Files

You can create files that boot your file server automatically and enhance its performance.

AUTOEXEC.BAT. When COMMAND.COM is loaded, it checks the boot diskette or the hard disk DOS partition for an AUTOEXEC.BAT file. The AUTOEXEC.BAT file contains one executable file command: SERVER.

Example

The AUTOEXEC.BAT file below is written to boot a NetWare v3.12 file server:

```
SERVER^Z
```

STARTUP.NCF. SERVER.EXE loads the STARTUP.NCF file. The STARTUP.NCF file contains commands to load the server's disk drivers and any name space support (such as Macintosh).

The disk drivers tell the file server how the network commands move between the main processor and the hard disks.

Example

The following STARTUP.NCF file is created by INSTALL on a v3.12 file server. (The command in this file is generated automatically by INSTALL.)

```
LOAD ISADISK PORT=1F0 INT=E
```

AUTOEXEC.NCF is executed next and contains the following information:

- The server name.
- The server's internal network number.
- The LOAD LAN driver command (sets addresses and interrupts for each installed network board).

  LAN drivers are loadable modules that tell the file server which cabling scheme the network commands will travel on.
- The BIND command (binds the protocol to each LAN driver).
- Other executable file server commands. (See "File Server Utilities" in *System Administration*.)

Example

The AUTOEXEC.NCF file below is created on a file server running on an ARCnet® cabling system. Other file servers on the internetwork are also NetWare v3.12 file servers. Information in this file is generated automatically by INSTALL.

```
FILE SERVER NAME SPEEDYIPX INTERNAL NET 1986ABE1LOAD
  ARCNET PORT=2E0 MEM=D000 INT=2BIND IPX TO ARCNET
  NET=1999EEEE
```

Boot record

Information that ROM-BIOS uses to determine which device to boot from.

The boot record can be on either a floppy diskette or a local hard disk. ROM-BIOS loads a short program from the boot record to determine disk format and the location of system files and directories.

Using this information, ROM-BIOS loads the system files (including two hidden files, IBMBIO.COM and IBMDOS.COM) and the command processor (COMMAND.COM).

Bridges

See “Router” on page 211.

Buffer

See “Cache buffer” on page 39.

Bus

A signal route for transmitting data between various parts of the network.

Several devices can be connected to a single bus, allowing them to share the same data pathway.

Data bus is the primary bus inside a personal computer.

Network bus is the main network cable or line that connects network stations.

3 C

Cabling system

See “Topology” on page 260.

Cache buffer

A block of file server memory (RAM) in which files are stored temporarily.

Workstations can access the data more quickly because reading from and writing to memory is much faster than reading from and writing to disk. Cache buffers greatly increase file server performance.

The default block size is 4 KB. You can change the size of cache buffers to 8 KB or 16 KB.

See also “Cache buffer pool” on page 40.

Cache buffer pool

The amount of memory available for use by the operating system after the SERVER file is loaded into memory.

When the operating system is installed, the supervisor can specify cache buffer size as 4 KB, 8 KB, or 16 KB.

The operating system uses the cache buffers in a variety of ways:

- The operating system loans cache buffers to loadable modules (such as LAN drivers, disk drivers, the INSTALL utility, etc.).

  When a loadable module is removed from file server memory, the module returns the borrowed memory to the cache buffer pool.

- The operating system allocates sufficient cache buffers to cache each volume's entire File Allocation Table (FAT) in memory.
- The operating system allocates enough cache buffers to cache parts of each volume's directory table.
- The operating system uses cache buffers as needed to cache parts of files that users want to access.
- The operating system allocates cache buffers to build a hash table for all directory names.
- The operating system allocates cache buffers to build turbo FAT indexes for all open files that are randomly accessed and have 64 or more regular FAT entries.

Character length

See “Serial communication” on page 246.

Charge rates

See “Accounting” on page 15.

CMOS RAM

(Complimentary Metal Oxide Semiconductor Random Access Memory) RAM used for storing system configuration data, such as number of drives, types of drives, and amount of memory.

The CMOS RAM is battery maintained and is not available to the computer's operating system.

COM1, COM2

See “Serial port” on page 248.

Command format

Instructions that show the correct way to type a command at the computer keyboard.

In NetWare manuals, a command format may include constants, variables, and symbols.

Communication buffer

See “Packet receive buffer” on page 175.

Communication protocols

Conventions or rules used by a program or operating system to communicate between two or more endpoints.

Although many different types of communication protocols are used, they all allow information to be packaged, sent from a source, and delivered to a destination system.

Workstation Protocols

NetWare workstations use network board-specific protocols such as IPX™ (Internetwork Packet Exchange™), SPX (Sequenced Packet Exchange), TCP/IP, NetBIOS, OSI, SMB, and AppleTalk® (for Macintosh).

See "IPX internal network number" on page 122; "SPX.COM" on page 252.

File Server Protocols

NetWare v3.12 has six layers of communication between an application and the hardware in the computer. (See "File server protocols" on page 42) In the file server, the communication protocols allow the Service Protocol Layer to communicate with the Link Support Layer™.

Figure 2 File server protocols

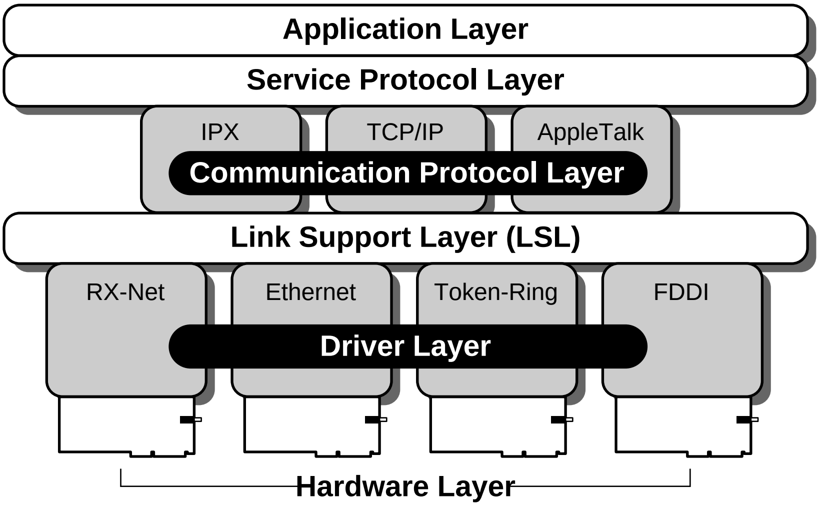

IPX, part of the operating system, is the default communication protocol. You can use more than one protocol on the same cabling scheme because the Link Support Layer allows a driver for a network board to service more than one protocol.

Use the following console commands to view, add, and configure communication protocols in the operating system.

- *PROTOCOL.* Displays the protocols registered with your file server and registers others. See "PROTOCOL" in *System Administration*.

- *BIND*. Binds a protocol to a network board installed in the file server. See "BIND" in *System Administration*.
- *UNBIND*. Unbinds a protocol from a network board installed in the file server. See "UNBIND" in *System Administration*.

See also “Binding and unbinding” on page 31.

Communications

See “Network communication” on page 165; “Serial communication” on page 246.

Configuration (hardware)

The equipment used on a network (such as file servers, workstations, printers, cables, network boards, and routers) and the way the equipment is connected—the physical layout of the network.

Hardware configuration includes

- The specific type of hardware installed in or attached to the computer itself, such as disk subsystems, network boards, memory boards, and printer boards;
- A specific set of parameters selected for a board.

Configuration options

Settings on network boards that allow all boards using the same cabling system to communicate with each other.

Configuration options can include four settings:

- Interrupt
- DMA
- Base memory address
- Base I/O address

The way jumpers or switches are set determines the configuration number for a board. Most network boards are factory set to a default option (0).

When more than one network board is installed in a workstation or file server, the two boards use different cabling systems. The configuration options for those boards must be unique so that conflicts do not occur.

Interrupt

An interrupt setting allows the network board to send an interrupt signal to the file server.

The interrupt signal temporarily suspends the file server's operation. The file server can then perform the task requested by the interrupting device.

Devices such as serial and parallel ports and network boards must have unique interrupts.

DMA

DMA (Direct Memory Access) allows some network boards and the workstation's memory to transfer information back and forth without going through the workstation's microprocessor.

Typically, DMA channels 1 and 3 in a workstation are reserved for network boards. Other channels are dedicated to hard disk drives and floppy disk drives.

Base Memory and Base I/O Addresses

The base memory address and the base I/O address have essentially the same function.

The network boards are assigned a block of I/O address space or memory address space in the workstation where the network board and the operating system can transfer information to one another.

Preventing Conflicts between Settings

The base address location for each configuration option is a hexadecimal number. (The file server operating system shows an extra 0 listed for the number. An address such as D0000 on the file server is the same as address D000 in the documentation.)

The next hexadecimal number address for the next configuration option should be set far enough from the first number that the address space for one setting does not overlap the space reserved for the next setting.

Example

Figure 3 I/O address space

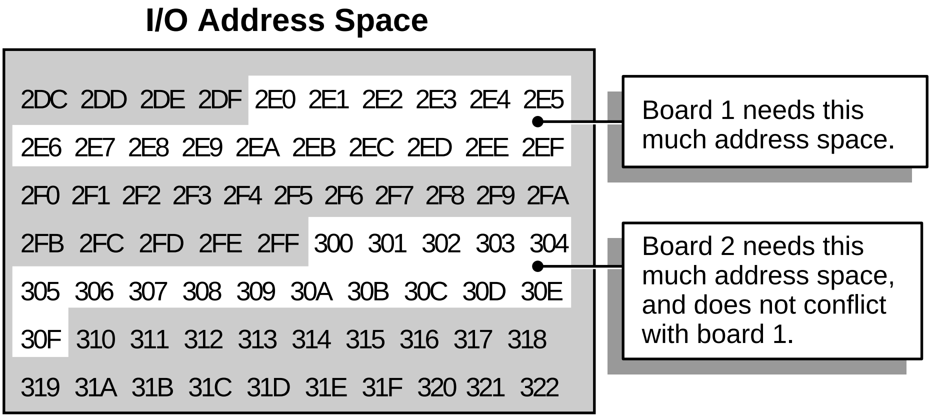
I/O Address Space
2DC 2DD 2DE 2DF 2E0 2E1 2E2 2E3 2E4 2E5
2E6 2E7 2E8 2E9 2EA 2EB 2EC 2ED 2EE 2EF
2F0 2F1 2F2 2F3 2F4 2F5 2F6 2F7 2F8 2F9 2FA
2FB 2FC 2FD 2FE 2FF 300 301 302 303 304
305 306 307 308 309 30A 30B 30C 30D 30E
30F 310 311 312 313 314 315 316 317 318
319 31A 31B 31C 31D 31E 31F 320 321 322

Board 1 needs this much address space.

Board 2 needs this much address space, and does not conflict with board 1.

The NetWare operating system keeps track of hardware settings. As you load disk drivers for the hardware devices in a file server, the system does not allow you to select a hardware option that conflicts with that of another device in the system.

The WSGEN utility keeps track of hardware settings on workstations and prevents you from selecting conflicting hardware options.

See also “Network numbering” on page 165.

Connection number

A number assigned to any station that attaches to a file server; it may be a different number each time a station attaches.

The file server's operating system uses connection numbers to control each station's communication with other stations.

You can find out your connection number by using the WHOAMI or USERLIST command.

Connectivity

The ability to link together different pieces of hardware and software (Macintoshes, PCs, minicomputers, and mainframes) in a network environment where resources (applications, processes, etc.) are shared.

See also "Internetwork" on page 120.

Console

The monitor and keyboard at which you view and control server activity.

At the console, you can enter commands to control printers and disk drives, send messages, set the file server clock, shut down the file server, and view file server information.

Unauthorized use of the console can have disastrous effects on your network.

Preventing Unauthorized Access

You can prevent unauthorized access in two ways:

- Use the lock feature in the MONITOR loadable module to disable keyboard entry until you type the SUPERVISOR password or the console password. (See "Lock the File Server Console" in *Sy*stem Administration.)
- Use the SECURE CONSOLE command to secure your console against breaches of security. (See "SECURE CONSOLE" in *Utilities Reference*.)

Possible Security Breaches

In addition to using SECURE CONSOLE, be aware of the following possible security breaches and of additional software protection available in NetWare v3.12.

- Unauthorized use of SYS:SYSTEM

  In a secure environment, only the system administrator and a backup administrator have rights to SYS:SYSTEM.

- Software tampering

  With access to the CONSOLE, an expert programmer could use the built-in debugger to disable or bypass the security system. Use SECURE CONSOLE to prevent entry into the OS debugger.

- Hardware tampering

  A determined intruder can take the file server apart and either disable the power-on password protection or remove the hard disks to access the data at leisure. Keep your file server in a physically secure location.

- Loadable modules

  Any loadable modules you use should be approved by Novell®. Loading a module that is not approved can cause your server to abend or to corrupt data.

  If you are running third-party loadable modules, check with NetWire® or consult with your Novell Service Representative to make sure they are approved.

Console operator

See also “File Server Console Operator” on page 108.

Controller address

The number that the operating system uses to locate the controller on a disk channel.

The number is physically set (usually with jumpers) on a controller board.

Controller board

A device that enables a computer to communicate with a particular device (such as a hard disk, network board, or tape drive).

The controller board manages input/output and regulates the operation of its associated device.

Coprocessor

See “Disk Coprocessor board” on page 71.

Copy Inhibit attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Create right

See “Rights” on page 207.

Cylinder

Distinct, concentric storage areas on a hard disk (roughly corresponding to tracks on a floppy diskette).

Generally, the more cylinders a hard disk has, the greater its storage capacity.

4 D

Data protection

A means of ensuring that data on the network is safe.

NetWare protects data primarily by maintaining duplicate file directories and redirecting data from bad blocks to reliable blocks on the file server's hard disk.

Protecting Data Location Information

A hard disk's directory table and File Allocation Table (FAT) contain address information that tells the operating system where data can be stored or retrieved from.

If the blocks containing these tables are damaged, the data may be irretrievable.

NetWare greatly reduces the possibility of losing this information by maintaining duplicate copies of the directory table and FAT on separate areas of the hard disk.

If a block in the original tables is damaged, the operating system switches to the duplicate tables to get the location data it needs.

The faulty sector is then listed in the disk's bad block table, and the data it contained is stored elsewhere on the disk.

Every time the server is turned on, the operating system performs a consistency check on both sets of directory tables and FATs to verify that the two copies are identical.

If the tables don't match, you receive a warning and should run VREPAIR.

Protecting Data Against Surface Defects

NetWare hard disks store data in 4, 8, 16, 32, or 64KB blocks. These blocks are specific data storage locations on the disk's magnetic surface. (Block size is the same on all segments of a volume.)

Because of the constant reading and writing of data from and to disk, some storage blocks lose their capacity to store data.

NetWare prevents data from being written to unreliable blocks by employing two complementary features known as read-after-write verification and Hot Fix™.

These features, illustrated in Figure 4 on page 50 and Figure 5 on page 51, enable a hard disk to maintain the same data integrity it had when it was first tested and installed.

Read-after-write verification. When data is written to disk, the data is immediately read back from the disk and compared with the original data still in memory. See Figure 4.

Figure 4 Read-after-write verification

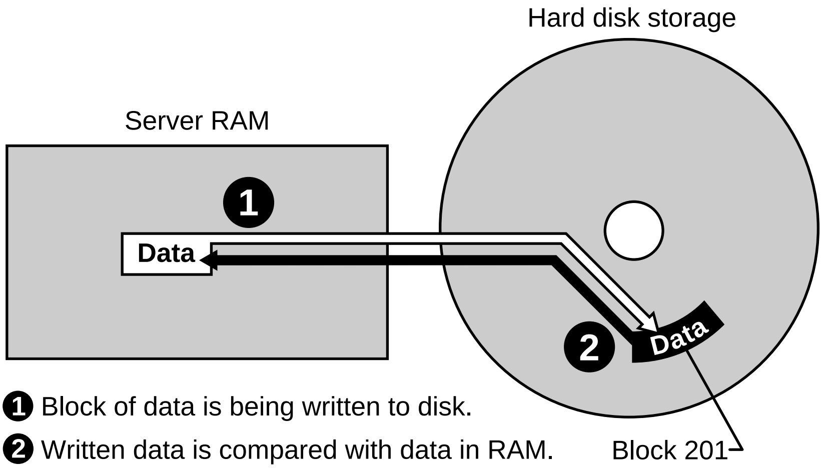

❶ Block of data is being written to disk.

❷ Written data is compared with data in RAM.

If the data on the disk matches the data in memory, the write operation is considered successful, the data in memory is released, and the next disk I/O operation takes place.

If the data on the disk doesn't match the data in memory, the operating system determines (after making appropriate retries) that the disk storage block is defective.

Hot Fix. The Hot Fix feature redirects the original block of data (still in memory) to the Hot Fix redirection area, where the data can be stored correctly.

A small portion of the disk's storage space is set aside as the Hot Fix redirection area. This area holds data blocks that are redirected there from faulty blocks on the disk.

Hot Fix is always activated unless the disk fails or the redirection area is full.

After the operating system records the address of the defective block in a section of the Hot Fix area reserved for that purpose, the server won't attempt to store data in the defective block. See Figure 5.

Figure 5 Hot Fix

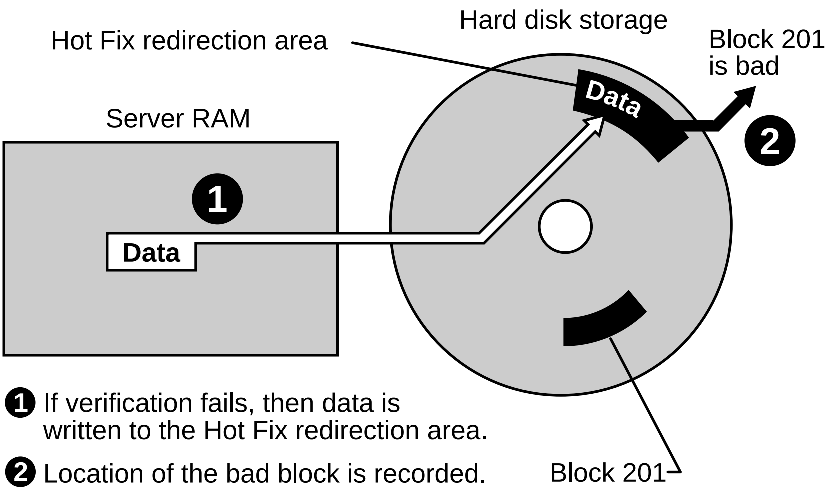

❶ If verification fails, then data is written to the Hot Fix redirection area.

❷ Location of the bad block is recorded.

Read-after-write verification and Hot Fix are transparent to users. The network supervisor or a console operator can view Hot Fix activity in SERVMAN or MONITOR.

Disk mirroring or duplexing. You can also protect your data with disk mirroring or duplexing.

Mirroring stores the same data on separate disks on the same controller channel; duplexing stores the same data on separate disks on separate controller channels.

Duplexing is the preferred method, since two channels rarely fail simultaneously.

See also “Directory table” on page 69; “Disk duplexing” on page 72;“Disk mirroring” on page 75; “File caching” on page 99; “Hot Fix” on page 117.

Related utility: "INSTALL" (*System Administration*).

Data set

A group of data that can be manipulated by SBACKUP.

Data sets can contain different items depending on which TSA they are related to. Figure 6 shows typical NetWare data sets.

Figure 6 Types of information that data sets contain

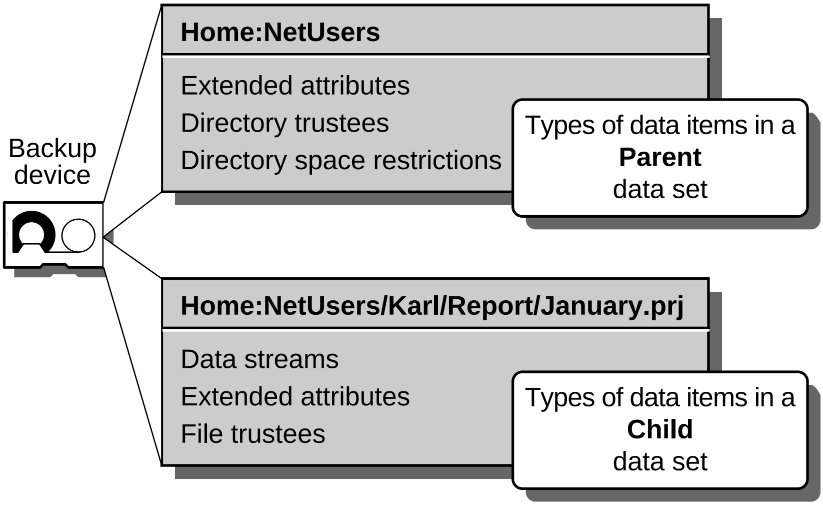
Backup device
Home:NetUsers
Extended attributes
Directory trustees
Directory space restrictions
Types of data items in a **Parent** data set
Home:NetUsers/Karl/Report/January.prj
Data streams
Extended attributes
File trustees
Types of data items in a **Child** data set

See also “Restore” on page 207.

DCB

See “Disk Coprocessor board” on page 71.

Dedicated file server

See “File server” on page 106.

Dedicated IPX drivers

IPX.COM files that are created with WSGEN.

These drivers work only with the IPX protocol. If you want to use other protocols, such as TCP/IP, you need to load ODI™ (Open Data-Link Interface™) drivers.

See also “Open Data-Link Interface” on page 173.

Dedicated mode

See “Router” on page 211.

Default drive

The drive that a workstation is currently using.

The drive prompt (such as A:> or C:>) identifies the default drive letter.

Default server

Usually the first server you log in to.

The LOGIN command lets you change your default server.

Related utility: "LOGIN" (*Ut*ilities Reference).

Delete Inhibit attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Delimiter

A symbol or character that signals the beginning or end of a command or of a parameter within a command.

For example, in the command CHKVOL A: B:, the blank space between A: and B: is a delimiter that marks two distinct parameters.

Other delimiters include the comma (,), the period (.), the slash (/), the backslash (\), the hyphen (-), and the colon (:).

Device driver

See "Disk Coprocessor board" on page 71; "LAN driver" on page 125.

Device numbering

In NetWare, a scheme of numbers that identify hard disks.

There are three different numbers:

- *The physical address* is set with jumpers on the boards, controllers, and hard disks.
- *The device code* is determined by the physical address of the board, controller, and hard disk.
- *The logical number* is determined by the order in which the disk drivers are loaded and by the physical address of the controller and hard disk.

Figure 7 on page 55 shows how the physical numbers correspond to the logical numbers and device code.

Each box in the second column describes the physical address and the device code of a particular hard disk. The "Device #" indicates the disk's logical number.

Figure 7 Physical and logical device numbering

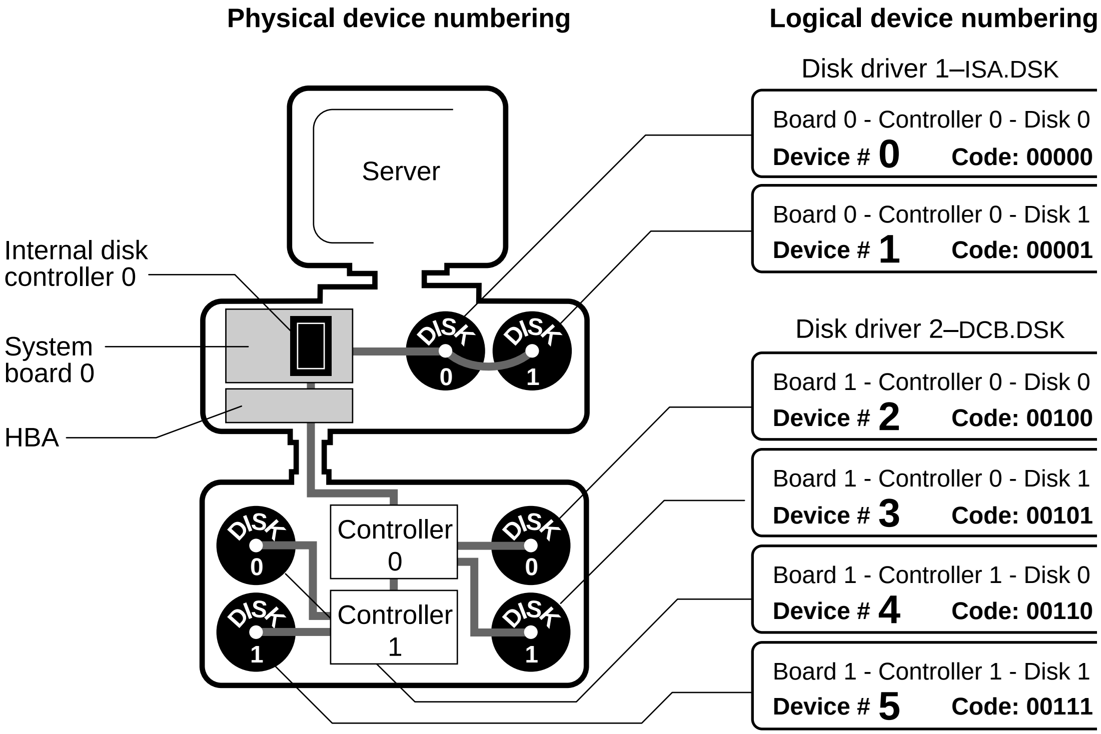

The *physical address* is determined by the driver, based upon jumper settings on the hardware.

The *device code* is determined by the board, controller, and disk. In the device code #00101, the first two digits are reserved for the disk type. The third digit is the board number; the fourth, the controller; and the fifth, the disk.

The *logical number* is determined by the controller that the hard disk is attached to and the order in which the disk drives are loaded. The example in Figure 7 assumes that Disk Driver 1 was loaded first. Hence, Disk 0 on the system board gets the logical number 0. If Disk Driver 2 is loaded first, then the disk that shows the logical device number 2 in the chart becomes 0.

We recommend that you make a chart of your hard disk setup and the order that the disk drivers are loaded in the STARTUP.NCF file.

If you then receive a message stating that Device #4 (00110) has been deactivated due to disk failure, you will know which disk needs to be repaired.

After device numbers are assigned, NetWare assigns physical and logical partition numbers to the partitions created on the hard disks.

Figure 8 on page 56 displays the relationship between the logical devices and the physical partition numbers.

Each logical device is assigned a partition number for each type of partition. NetWare v3.12 recognizes DOS and NetWare partitions. All other partition types, such as UNIX partitions, are listed as non-NetWare partitions.

Figure 8 Logical devices and physical partition numbers

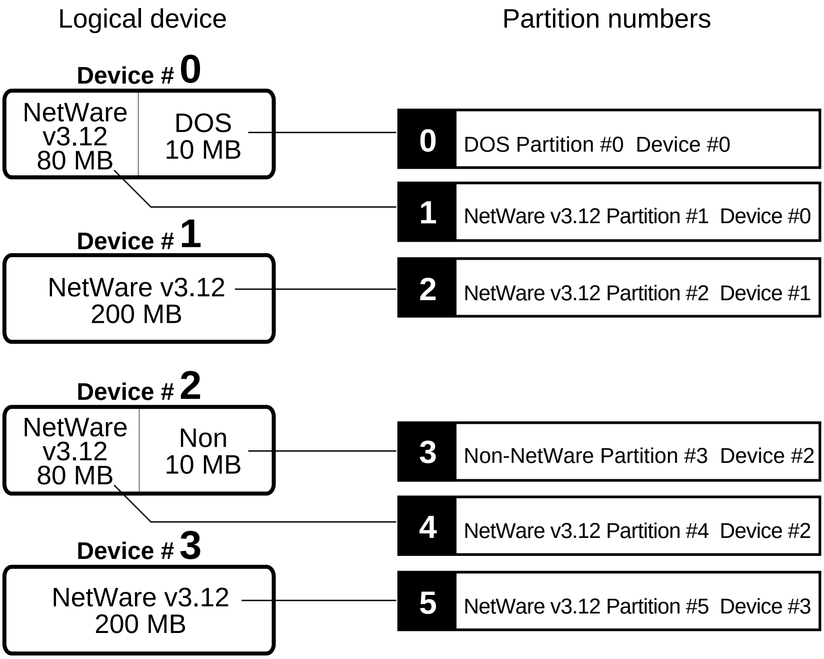
Logical device
Partition numbers
Device # 0
NetWare v3.12 80 MB
DOS 10 MB
Device # 1
NetWare v3.12 200 MB
Device # 2
NetWare v3.12 80 MB
Non 10 MB
Device # 3
NetWare v3.12 200 MB
0 DOS Partition #0 Device #0
1 NetWare v3.12 Partition #1 Device #0
2 NetWare v3.12 Partition #2 Device #1
3 Non-NetWare Partition #3 Device #2
4 NetWare v3.12 Partition #4 Device #2
5 NetWare v3.12 Partition #5 Device #3

Hot Fix messages use the physical partition number when recording which hard disks have blocks of data that need to be redirected.

All physical partitions are assigned logical partition numbers. These numbers are assigned to both the mirrored disks and the DOS and non-NetWare partitions.

Figure 9 on page 57 displays the relationship between the physical partitions and the logical partitions.

Figure 9 assumes that partition 1 is mirrored to partition 4 and that partition 2 is mirrored to partition 5.

Figure 9 Physical and logical partitions

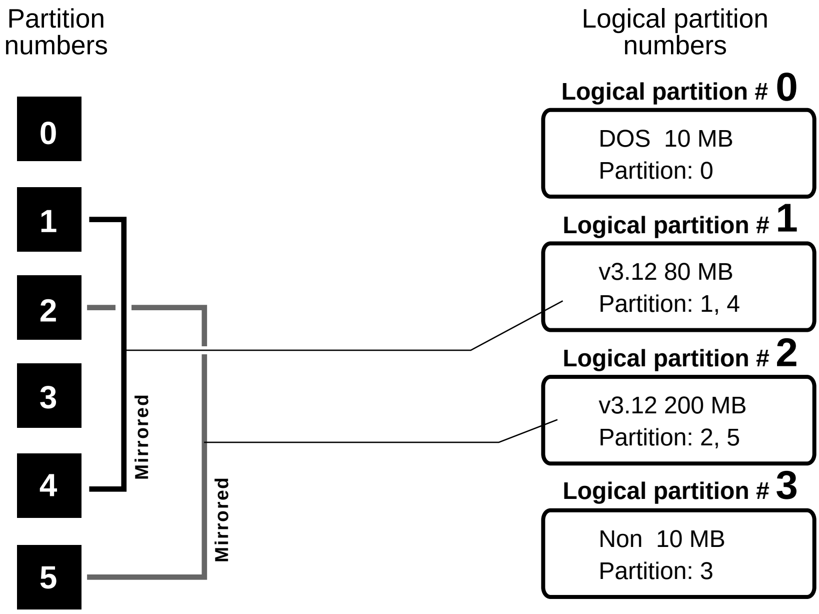

Mirroring messages use the logical partition number to record which hard disks are being remirrored or unmirrored.

Related utility: "SET" (System Administration).

Device sharing

Refers to the shared use of centrally located devices (such as printers, modems, and disk storage space) by a number of users.

By attaching a device to a file server, which in turn serves several workstations, you can use resources more efficiently.

Direct memory access

See “DMA” on page 76.

Directory

See “Directory structure” on page 61.

Directory attributes

See “Attributes” on page 21; “Security” on page 221.

Directory caching

A method of decreasing the time it takes to determine a file's location on a disk.

The File Allocation Table (FAT) and directory tables from the file server's volumes are written into the file server's memory. The area of memory used to hold the most-often-requested directory entries is called directory cache memory.

The file server can find a file's address (from directory cache memory) and the file data (in the file server's cache memory) much faster than if it had to retrieve the information from the hard disk.

As the directory cache fills, the least-used directory entries are eliminated. See Figure 10.

Figure 10 Directory caching

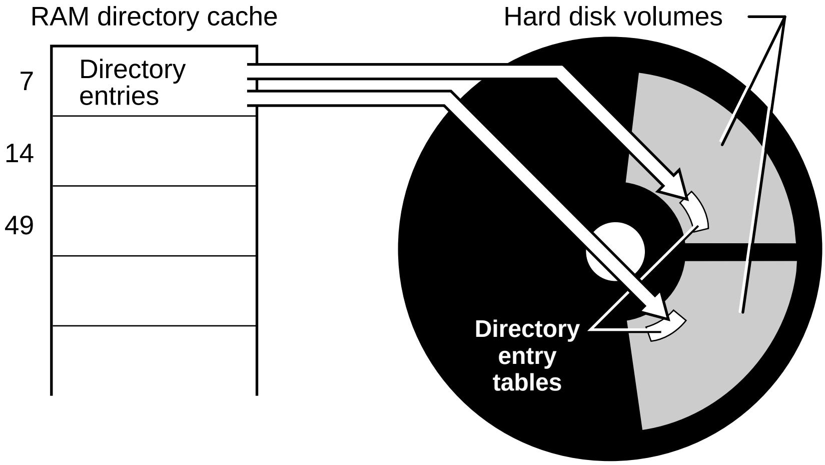

See also “Directory entry” on page 59; “File caching” on page 99; “Hashing” on page 116.

Directory entry

Basic file information located in a directory table on a network hard disk, such as

- Filename
- Owner
- Date and time of the last update
- First six trustee assignments
- First block of the network hard disk in which the file is stored

Directory entries contain information about all files on the volume.

The file server uses directory entries to keep track of a file's location, changes made to the file, and other properties related to the file.

See also “Directory caching” on page 58.

Directory path

The full specification that includes files server name, volume name, and the name of each directory leading to the directory you need to access.

The directory path is the same as the full directory name and begins optionally with the file server name and then specifies volume, directory, and (if necessary) subdirectory and any levels below.

Conventions

Use the conventions illustrated in Figure 11 on page 60 and listed below the figure when specifying a directory path.

Figure 11 Directory path conventions

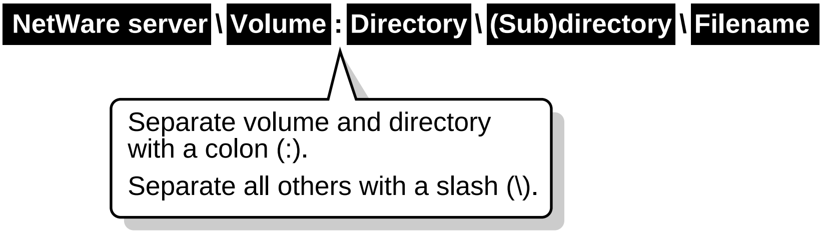
NetWare server \ Volume : Directory \ (Sub)directory \ Filename
Separate volume and directory with a colon (:).
Separate all others with a slash (\).

- Directory names can contain from 1 to 8 characters.
- For separating the levels of directory structure, NetWare recognizes both the slash (/) and the backslash (\), but DOS recognizes only the backslash.
- The volume and the root directory (the directory next to the volume) must be separated with a colon (:).
- A directory path can specify the filename.

Path Restrictions

Restrictions and parameters for different workstation operating environments can vary. This is important if you intend to share files.

For example, NetWare allows 255 characters in a directory path, but DOS permits only 127 characters. Some applications also restrict the path to the number of characters that fit the screen width.

Consult the documentation accompanying third-party applications to determine the maximum path length.

See also "Directory structure" on page 61.

Directory rights

See "Rights" on page 207; "Security" on page 221.

Directory structure

The hierarchical system used to organize network files and directories on a file server's hard disk.

Each file is given a filename and stored at a specific location in a hierarchical filing system so that files can be located quickly.

Directory structure can be compared to a file cabinet, with the drawers in the cabinet, hanging folders in the drawers, manila folders in the hanging files, and documents in manila folders corresponding to different levels of the directory structure.

The levels of this filing system are illustrated in Figure 12.

Figure 12 Directory structure

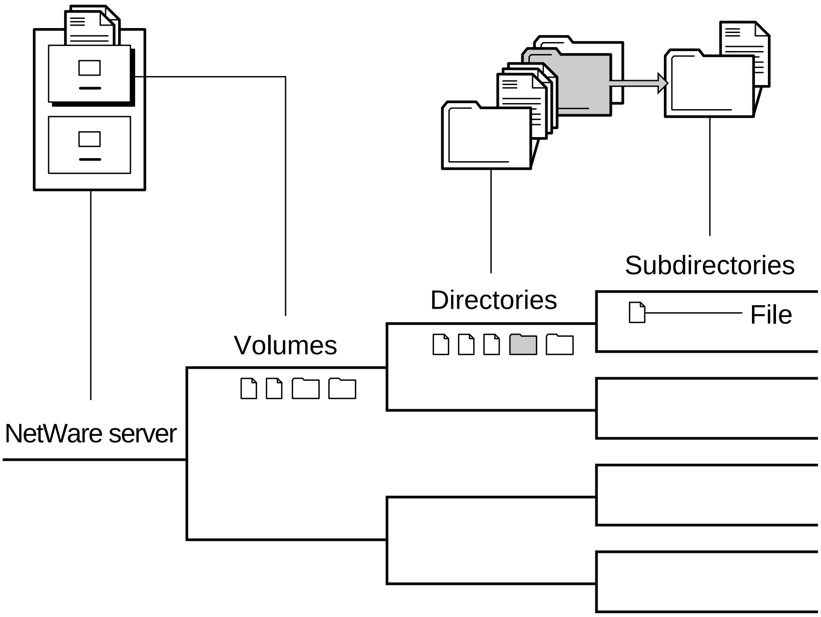

Volumes mark the start of the directory structure. A volume is a logical unit; however, it appears much like a hard disk to a standalone system.

A volume may contain several physical disks. You can store directories at the volume level but, for security reasons, this is not recommended.

Directories are divisions of a volume where you can store files or other directories. These directories-within-directories are sometimes called subdirectories. The term *subdirectory* is relative: a directory is a subdirectory only in relation to the directory above it.

A directory can contain any number of both files and subdirectories, as in the following tree structure. See Figure 13.

Figure 13 Tree structure

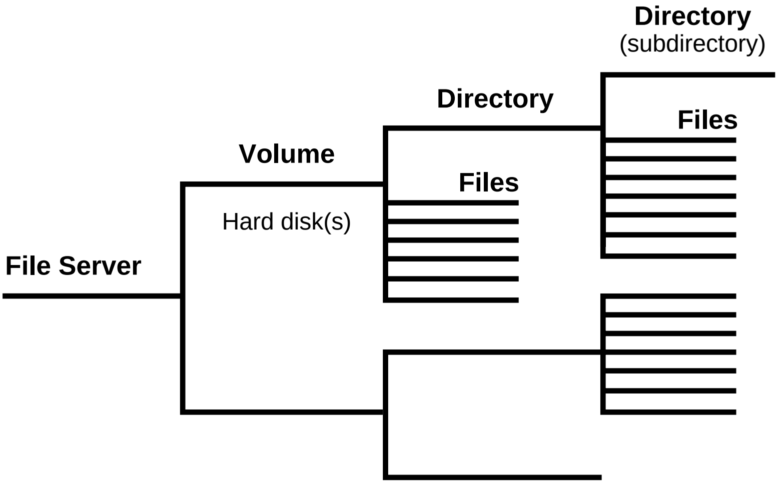

Directory Paths

Each directory is named when it is created. Because each level has a name, the location of any file can be pinpointed by listing the directory name at each level of the directory structure.

The list of levels constitutes both the directory path and the full name of the directory.

See “Directory path” on page 59 for a full explanation of naming conventions for directories.

Basic Directory Structure

When the SYS: volume is created, a directory table is created automatically. The table contains four predefined directories, illustrated in Figure 14 and listed below:

Figure 14 Basic directory structure

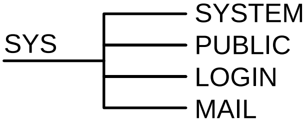

- SYS:SYSTEM is used for system administration and contains operating system files, NetWare utilities, and programs reserved for SUPERVISOR.
- SYS:PUBLIC is used for general access and contains NetWare utilities and programs for regular network users.
- SYS:LOGIN contains the programs necessary for logging in.
- SYS:MAIL is used by NetWare-compatible mail programs. This directory also has an ID number subdirectory for each user that contains the user login script and print job configurations.

Additional Directories

You will create additional directories according to the needs of your organization. However, if your needs are better met by creating separate volumes, the volumes must be created when NetWare is installed.

We suggest you create the following directories or volumes (discussed in detail in “Types of Directories” on page 65).

- One or more DOS directories. We recommend that you create your DOS directory structure in SYS:PUBLIC, since the appropriate security parameters (rights, attributes) are already set up. (See "DOS directories.")
- One or more application directories. You should create a separate directory for each application.
- A "home" or username directory for each user. If you want users to have personal workspace, you can create a separate directory (or volume) for this purpose.

You may also find the following kinds of directories useful.

- Work or project directories. If you want users to have group workspace or if you want to store completed application data files, you can create a directory for each project.

  You may want to create such directories in a separate volume to simplify backup.

- A common directory in the SYS: volume (or a HOME or a WORK volume) to serve as an intermediate point in transferring files.

  The group EVERYONE should be assigned the rights necessary to copy files to and from this directory. If you do not want GUEST to have access to this directory, delete GUEST from the group EVERYONE.

- A directory for storing batch files and utilities.

  If you do not want to store batch files in a public access directory, you can create a directory for your batch files elsewhere and assign rights to groups or individual users.

Accessing Directories

Although you can use commands provided with your workstation operating system to access directories, NetWare provides a more direct way.

Like DOS, NetWare uses the letters of the alphabet to represent various data storage locations, or "drives." You can "map" drives to a particular directory path (name) in the directory structure.

See “Drive mapping” on page 83 for a complete explanation.

Types of Directory Structures

All directory structures are tree structures, but trees may have different proportions. Some are broad and spreading while others are tall and narrow.

A directory structure may be relatively flat with many directories coming off the volume. Or your structure may be deep if you limit the number of directories at the root and create several levels of directories.

The general principle is to keep the directory structure clean and logical. Keeping the structure relatively flat (no more than five levels deep) generally increases its usability. See Figure 15.

Figure 15 Flat vs. deep directory structure

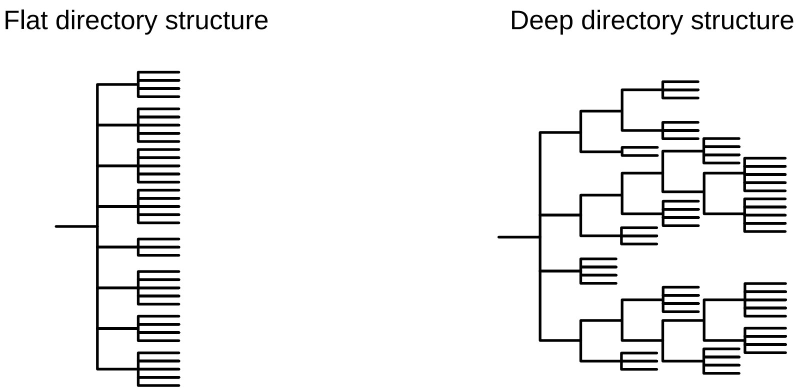

Plan directories by grouping your files logically. Plan subdirectory levels for natural subcategories.

Besides considering the logic of groupings, you may want to limit the number of files in each directory.

Determine which application allows the fewest characters in a directory path. You may need to plan either shorter directory names or a flatter directory structure.

Types of Directories

You can create directories for both executable files and data files. As network administrator, you must determine which types of directories best fit the needs of your network.

DOS directories. Although DOS is the operating system used by individual workstations, it is installed on the network by copying DOS program files into network directories.

You must create one or more DOS directories, copy the files into the directories, and then include a mapped search drive (usually Search2) in the system login script. ("DOS directories" on page 77)

Application directories. Although applications can be accessed from local drives, installing them on the network provides the most convenient access.

To determine the application directory structure that meets your requirements, consult the documentation for the application to see what is recommended and what adaptations you can make. Then plan a logical directory structure similar to one of the following.

- A separate volume for applications with a separate directory for each application off the root. See Figure 16.

  You must make trustee assignments for each application and map search drives in the system login script. This solution hides applications from GUEST and simplifies backup.

Figure 16 Application directories in a separate volume

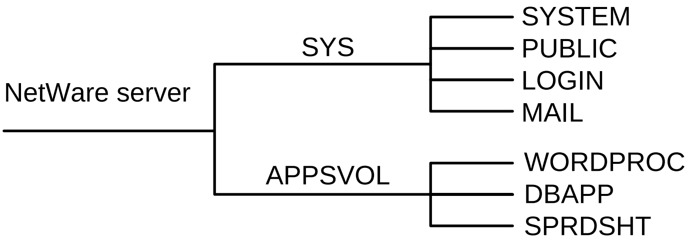

- A separate directory off the SYS: volume for each application. See Figure 17.

  You must make trustee assignments for each application and map a search drive for each in the system login script.

  However, some applications write files to the root. Since you do not want users working at the root level (for security reasons), use MAP ROOT to map a drive to a fake root—in this case, a directory off SYS: in which the user can be assigned rights.

Figure 17 Application subdirectories in volume SYS:

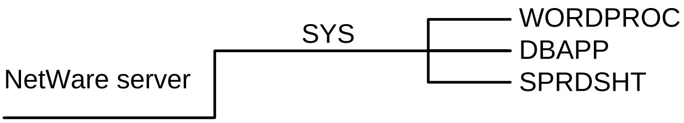

- A parent directory for applications, SYS:APPS, with subdirectories for applications. See Figure 18.

  You must make trustee assignments for each application and map a search drive for each application in the system login script.

Figure 18 Parent directory in volume SYS: for applications

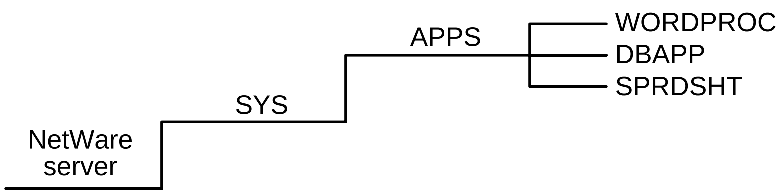

- A parent directory for applications, APPS, in SYS:PUBLIC. See Figure 19.

  Because the group EVERYONE has Read and File Scan rights to SYS:PUBLIC, you do not need to make trustee assignments or map a search drive.

  EVERYONE and GUEST can see and use all applications. Use this directory structure only if you want all users (including GUEST) to have access to all applications.

Figure 19 Parent directory in SYS:PUBLIC for applications

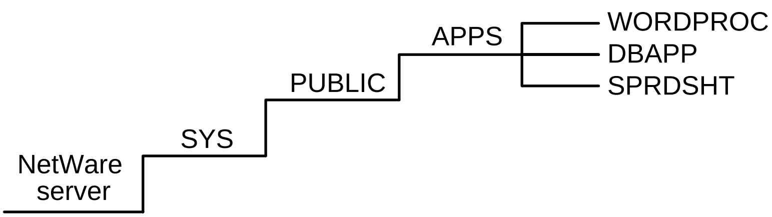

We do not recommend installing applications in the SYS:PUBLIC directory, unless a subdirectory is created for each application.

Upgrading a network is more complicated if you mix NetWare utilities with application program files. An application file might have the same filename as a NetWare utility file or another application's program file. In such a case, one file overwrites the other because two files with the same filename cannot coexist in a directory.

Application data files can be created and stored in personal workspace in the home or username directory or in group workspace in separate work, project, or database record directories (these directories are explained below).

Home or username directories. These directories provide personal workspace for users.

You can create a parent directory in the SYS: volume called HOME or USERS. Then you can create a subdirectory for each user. See Figure 20 for examples.

Figure 20 Home directories

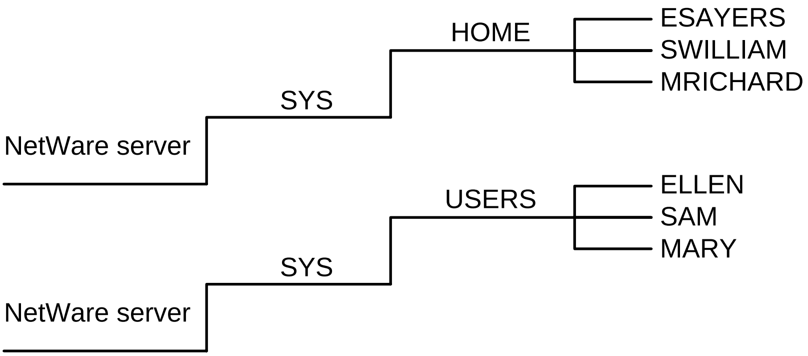

Charting Directory Structure

To plan or to keep track of your directories, make a chart similar to that in Figure 21 on page 69

You do not need to create any directories in SYS:SYSTEM, SYS:LOGIN, or SYS:MAIL. (When you create each user, a subdirectory is automatically created to store the user login script and print job configurations.)

Figure 21 Charting directory structure

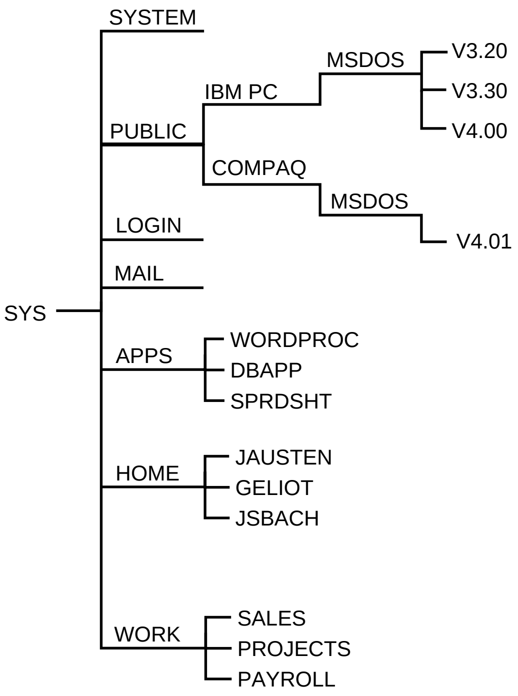

See also “Attributes” on page 21; “Directory” on page 58; “Directory table” on page 69; “DOS directories” on page 77; “Rights” on page 207; “Security” on page 221.

Related utilities: "CHKDIR"; "CHKVOL"; "DSPACE"; "FILER"; "FLAGDIR"; "LISTDIR"; "MAP"; "NDIR"; "RENDIR"; "SYSCON" (Utilities Reference).

Directory table

A table that contains basic information about files, directories, directory trustees, or other entities on the volume.

The directory table occupies one or more directory blocks on the volume. Each block has 4 KB (4,096 bytes) of data. A directory entry is 32 bytes long, so each block can hold 128 directory entries.

Volume SYS: starts out with six blocks for its directory table. When a volume needs to add another block to its directory table, the server allocates another block.

The maximum directory blocks per volume is 65,536. Since each block can accommodate 32 entries, the maximum directory table entries per volume is 2,097,152.

The server doesn't cache entire directory tables; it caches only directory blocks in use.

In NetWare v3.12, a volume can span multiple drives, so each drive can have more than one directory table.

See also "Directory entry" on page 59.

Disable

1. To turn off; to render inactive. For example, the DISABLE LOGIN console command prevents workstations from logging in to the file server.
2. To prevent certain interrupts from occurring in a processing unit (such as a network board) by setting a switch or a jumper, or using some other means.

Disk

A magnetically encoded storage medium in the form of a plate (also called a *platter*).

The following types of disks are used with personal computers:

- *Hard disks* use a metallic base and are usually installed within a computer or disk subsystem. (In some cases, they are removable.)
- *Floppy disks* (also called *diskettes*) use a polyester base and are removable.

- *CD-ROM* (Compact Disc Read Only Memory) is a small plastic optical disk that isn't erasable or writable.
- *Optical disks* are either erasable and writable, or WORM (Write Once, Read Many).

See also “Data protection” on page 49 .

Disk controller

A hardware device that controls how data is written to and retrieved from the disk drive.

The disk controller sends signals to the disk drive's logic board to regulate the movement of the head as it reads data from or writes data to the disk.

Disk Coprocessor board

(DCB) An intelligent board that acts as an interface between the host microprocessor and the disk controller.

The DCB relieves the host microprocessor of data storage and retrieval tasks, thus increasing the computer's performance time.

A DCB and its disk subsystems make up a disk channel.

The NetWare operating system can handle up to four DCB channels. NetWare allows each DCB a maximum of eight SCSI controllers, with each controller supporting up to two disk drives. External SCSI disk drive subsystems can be daisy-chained off the DCB port.

See also “SCSI bus” on page 218.

Disk driver

A loadable module that forms the interface between the NetWare operating system and the hard disks.

The disk driver talks to an adapter that is connected by an internal cable to the disk drives. Depending on the type of disk controller, one or more disk drives can be connected.

The driver is loaded into the operating system at the command line.

Related utility: "LOAD disk driver" in *System Administration*.

Disk duplexing

A means of duplicating data to provide data protection. Disk duplexing consists of copying data onto two hard disks, each on a separate disk channel.

This protects data against the failure of a hard disk or failure of the hard disk channel between the disk and the file server. (The hard disk channel includes the disk controller and interface cable.)

If any component on one channel fails, the other disk can continue to operate without data loss or interruption, because it is on a different channel. See Figure 22 on page 73.

The operating system sends a warning message to indicate when a drive has failed. (You should restore the duplexing protection as soon as possible.)

Figure 22 Disk duplexing

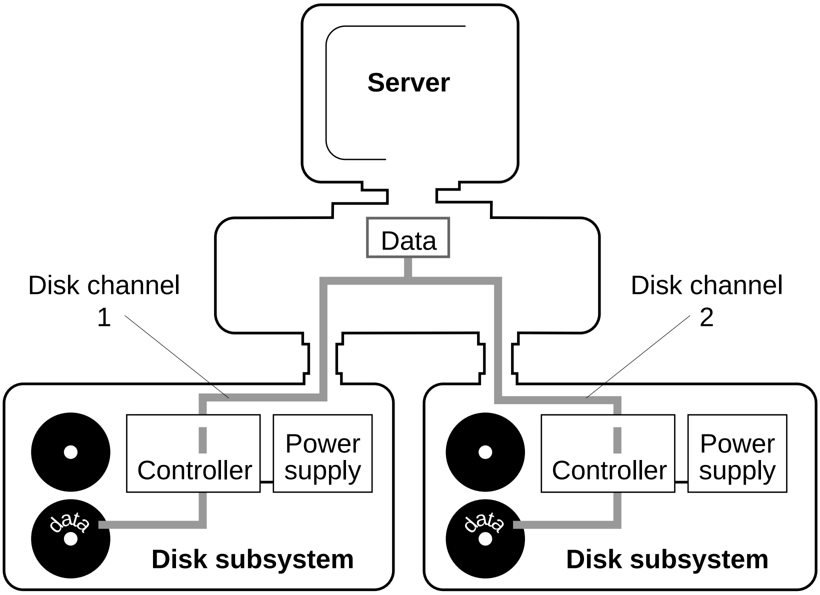

NOTE: Duplexing alone doesn't guarantee data protection. If both disk channels fail at the same time, or if the computer itself fails, you still lose your data. Therefore, you should back up your data regularly. Use SBACKUP or another backup utility.

Disk duplexing allows the same data to be written to all disks simultaneously.

Since the disks are on different channels, data transfer is faster than with disk mirroring, where data is written to the disks sequentially over the same channel.

Disk duplexing also allows *split seeks*: read requests are sent to whichever disk can respond first. Multiple read requests are also split between the duplexed disks for simultaneous processing.

Related utility: "INSTALL" (*System Administration*).

See also "Device numbering" on page 54; "Disk mirroring" on page 75.

Disk formatting

The preparation of a disk by dividing it into sectors so that it can receive data from the computer's operating system.

A NetWare file server works with two kinds of hard disk formatting: DOS and NetWare.

DOS Format

The DOS format allows you to boot the file server from a DOS partition. (This is optional, since a file server running NetWare v3.12 can also be booted from a diskette.)

For more information on the DOS FORMAT command, see your DOS manual.

NetWare Format

NetWare partitions do not need to be formatted initially, unless a large percentage of the disk has bad blocks.

If you need to check for bad blocks, use the NetWare surface test rather than the NetWare format program. Use the nondestructive surface test if you have data saved to the hard disk.

The NetWare format runs in the background on one drive so that you can work on other drives simultaneously. You can format more than one drive at a time, as long as you format one drive per controller.

See also “Data protection” on page 49; “Hot Fix” on page 117; “Surface test” on page 257.

Related utility: "INSTALL" (*System Administration*).

Disk interface board

An add-on board that acts as an interface between the host microprocessor and the disk controller.

See also “Disk Coprocessor board” on page 71.

Disk mirroring

The duplication of data from the NetWare partition on one hard disk to the NetWare partition on another hard disk.

When you mirror disks, two or more hard disks on the *same channel* are paired. Blocks of data written to the original (primary) disk are also written to the duplicate (secondary) disk.

The disks operate in tandem, constantly storing and updating the same files. If one of the disks fails, the other disk can continue to operate without data loss or interruption. See Figure 23.

NOTE: Mirroring alone does not ensure data protection. If both hard disks fail at the same time, you still lose your data. You should back up your data on a regular basis. Use SBACKUP or another backup utility.

Figure 23 Disk mirroring

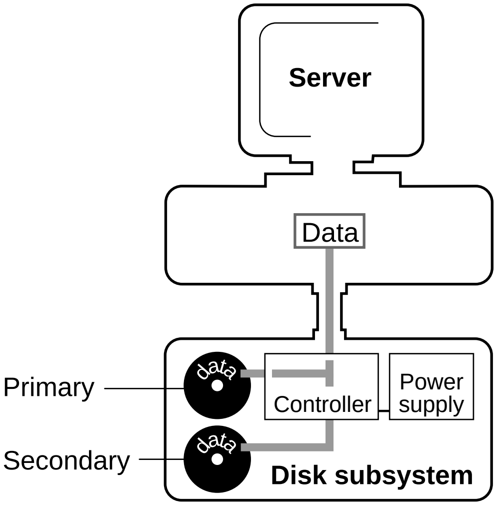

If one disk fails, the operating system sends a warning message to indicate the failure so that the mirroring protection can be restored as soon as possible.

Because disk mirroring duplicates disks on the same channel, it does not protect against failures that may occur along the channel between the disks and the file server. A problem in the channel would cause a failure in both disks.

See also "Device numbering" on page 54; "Disk duplexing" on page 72.

Disk partitions

See "Partitions" on page 177.

Disk subsystem

An external unit that attaches to the file server and may contain hard disk drives, a tape drive, or both.

The disk subsystem gives the file server more storage capacity.

DMA

(Direct Memory Access) A method used to reduce the burden on the processor in sending data to or receiving data from external devices.

The DMA controller chip moves data directly from a device to RAM. When the data transfer is complete, the controller signals the processor that the job is complete.

Since the processor handles many tasks and the DMA handles only data delivery to and from RAM, the DMA chip is usually faster than the processor.

The DMA controller can handle up to four devices. Each device is attached separately to what is called a channel. A common source of hardware conflicts within a machine is two devices trying to use the same DMA channel.

See also "Base I/O address" on page 29; "Base memory address" on page 29; "Configuration options" on page 43; "DOS device" on page 77; "Interrupt" on page 121.

DOS device

A storage unit compatible with the DOS disk format—usually a disk drive or tape backup unit.

The UPGRADE and SBACKUP utilities both write to a DOS device. The DOS device should be a read-write device.

Because the utilities both read and write data, the media the DOS device uses must allow the data to be updated or changed.

The following devices can be used. If the device runs out of storage space, you are prompted to insert another one.

- Workstation floppy disk drives (can't be used with SBACKUP)
- Tape drives with a DOS device driver
- Optical drives that are read-and-write devices and have a DOS device driver

The following devices can also be used as DOS devices; however, be careful not to select more data in one backup session than can fit on the device.

- Workstation hard disk drives
- Network drives
- An optical drive or a WORM drive that has a DOS device driver

NOTE: When you use these DOS devices for UPGRADE, you should ensure that the devices have at least 10% more space than the data occupies on the original file server; if the device runs out of space, the upgrade process abends.

DOS directories

NetWare directories that you create for DOS files.

DOS is an operating system used by individual workstations. You can access it from local drives; however, it is more convenient to copy DOS program files and utilities to the network.

Planning DOS Directories

You must create one or more network directories for DOS and then copy the DOS files to the directories.

To give all users access to DOS, you must make a trustee directory assignment with the Read and File Scan rights [RF] to the group EVERYONE.

If you include a mapped search drive (usually S2) in the system login script, all users whose workstations require DOS can execute DOS commands from any location in the directory structure.

To plan for DOS directories, you must determine the

- DOS version;
- Workstation brand (IBM brand is default).

Planning is easier if you have only one brand of workstation and use one version of DOS.

IBM Brand Workstations and Same DOS Version

Even if all workstations on your network are IBM computers that use the same version of DOS, you still need a separate DOS directory for the DOS program files.

NOTE: Do not copy DOS files to SYS:PUBLIC. Mixing DOS files with NetWare utilities makes upgrading difficult.

You should create a subdirectory in SYS:PUBLIC and name it for the DOS version, using the *Vx.xx* naming convention, as in the following example:

SYS:PUBLIC/V3.30

Load the DOS program files in that directory and map a search drive similar to the following in your system login script.

```
MAP INS S2:=SYS:PUBLIC\V3.30
```

IBM Brand Workstations and Multiple DOS Versions

If you are running more than one version of DOS, you must create a separate DOS directory for each version.

Separate directories prevent files from being overwritten, and also make it possible for workstation requests to be directed to the appropriate DOS directory.

Suppose you want to use three versions of DOS. You need to create directories similar to the following:

SYS:PUBLIC/V3.20
SYS:PUBLIC/V3.30
SYS:PUBLIC/V4.00

The directory structure would be similar to that in Figure 24.

Figure 24 Directory structure for DOS files

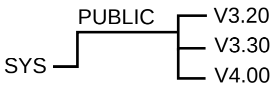

Although you have a different directory for each version of DOS, you need only one search drive mapped for DOS in the login script. Since the directory name follows the convention *Vx.xx*, you can use an identifier variable for the directory named for the DOS version:

```
MAP INS S2:=SYS:PUBLIC\%OS_VERSION
```

IBM-Compatible Workstations

If your network has all IBM-compatible workstations, you can use one of the solutions discussed for one or more DOS directories, as long as you use the format V*x.xx* for the directory name.

One Workstation Brand and Multiple DOS Versions

If all workstations on your network are one brand, and you plan to run more than one version of DOS, we recommend the directory structure illustrated in Figure 25.

Figure 25 Directory structure for DOS files

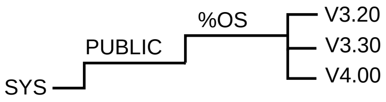

Note that the literal directory name %OS is identical to a generic identifier variable.

When you use identifier variables, you need only one mapped search drive for DOS in the login script:

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

Multiple Workstation Brands

If you have both IBM and non-IBM workstations, and all your machines can run the same type of DOS (for example, IBM PCDOS on COMPAQ® workstations), you can treat them as though they were IBM workstations.

If you have both IBM and non-IBM workstations *and* your non-IBM workstations require a proprietary version of DOS (for example, COMPAQ DOS for COMPAQ workstations and Samsung® DOS for Samsung workstations), you must complete these additional tasks:

- Create a separate DOS directory for the proprietary DOS version.
- Assign the level of directory structure below PUBLIC a six-letter name that specifies the machine name (workstation brand).

  In this case, you would use IBM_PC and COMPAQ.
- Set the LONG MACHINE TYPE command in the NET.CFG file on the boot diskette (discussed below) so the workstation's NetWare shell will know that non-IBM workstation brands will run a proprietary version of DOS.

Name the directories according to the following convention:

SYS:PUBLIC/machine name/DOS version

For example, suppose you have a network with both IBM and COMPAQ workstations. The IBM workstations will run three versions of DOS, and the COMPAQ workstations will run one version of COMPAQ DOS.

In such a case, you need to create four DOS directories. We recommend that you create a directory structure similar to that in Figure 26.

Figure 26 Directory structure for DOS files

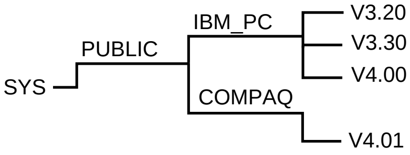

When you follow the naming conventions, you need only one search drive mapping for DOS in the system login script. You can use a generic drive mapping that provides access to specific DOS directories.

The following drive mapping corresponds to the above example of DOS directory structure for multiple workstation brands:

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

To use the %MACHINE variable in the login script to indicate other than the default machine type (IBM_PC), you need to set the following command in the NET.CFG file on each workstation's boot diskette (do not include extra spaces).

```
LONG MACHINE TYPE=name
```

Replace *name* with the six-letter name that specifies the workstation brand.

NOTE: Even if you have only one DOS directory, you can use identifier variables in your system login script when you map a DOS directory to search drives. When you add another DOS version, you do not need to change the login script. Just be sure to follow the same naming conventions.

USERDEF DOS Directories

If you plan to create users with the USERDEF utility, you can let USERDEF create your DOS directories only if

- You plan to create your initial users with the USERDEF utility;
- You boot DOS and run the USERDEF utility from one workstation that represents each unique combination of DOS version and workstation brand.

  (If you have six unique DOS version/workstation brand combinations, you need to run USERDEF from one workstation of each combination.)

In this case, the first time you run USERDEF from each workstation, it creates a DOS directory and then prompts you to load DOS.

USERDEF also provides a user login script that includes a search drive mapping to DOS directories:

```
MAP INS S2:=SYS:PUBLIC/%MACHINE/%OS/%OS_VERSION
```

This mapping contains an additional level of directory structure (indicated by the %OS variable). Do not alter this mapping; it corresponds to the directory path USERDEF creates.

See also "Directory structure" on page 61; "Login script" on page 128.

Related utilities: "FILER"; "USERDEF" (*U*tilities Reference).

DOS Requester

See "NetWare DOS Requester" on page 156.

DOS setup routine

The routine that sets up the system configuration of your workstation or file server.

The setup routine records the system's built-in features (add-on boards, hard drives, disk drives, ports, math-coprocessor) and available system memory. It also lets you set date and time, password, and keyboard speed.

The system configuration is accessed from the reference diskette (for IBM PS/2® systems) or from the setup or user diagnostics diskette (for most other systems).

Instructions for running the DOS setup routine are usually contained in the introduction to your PC's operations guide.

DOS version

The version number and name of the kind of DOS you are using (DR DOS® 6.0, MS-DOS® 3.3, etc.).

Different machine types use different versions of DOS that are generally not compatible.

Since all DOS versions have identically named utilities and command interpreters, you can't place the files of different DOS versions in the same directory.

Create a DOS directory for each workstation type or DOS version you use and load the DOS files into it.

See also “DOS directories” on page 77.

Drive

Physical drive. A storage device that data is written to and read from, such as a disk drive or tape drive. A drive that is physically contained in or attached to a workstation is called a *local drive*.

Logical drive. An identification for a specific directory located on a disk drive. For example, network drives point to a directory on the network, rather than to a local disk.

Drive mapping

A pointer to a location in the directory structure, represented as a letter assigned to a directory path on a volume.

For example, to locate a file, you follow a *path* that includes the volume, directory, and any subdirectories leading to the file.

You can create drive mappings to follow these paths for you. You assign a letter to the path, and then use the letter in place of the complete path name.

Drive mappings can be temporary or permanent.

- *Temporary mappings.* To map a drive so you can use it during your current session, use MAP or SESSION. The mapping is valid only until you log out.
- *Permanent mappings.* To make drive mappings so you can use them every time you log in, place MAP commands in your login script (see “Login script” on page 128).

NetWare recognizes three types of drive mappings:

- Local drive mappings
- Network drive mappings
- Search drive mappings

Figure 27 illustrates which letters you can assign to the different types of drives.

Figure 27 Available drives

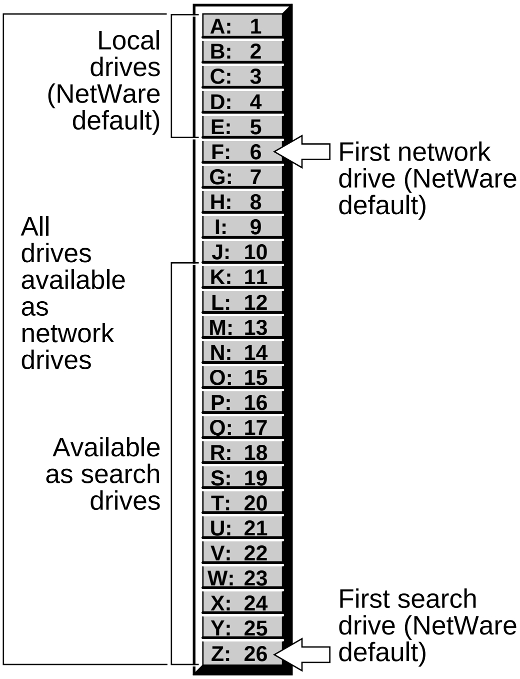
Local drives (NetWare default)
All drives available as network drives
Available as search drives
A: 1
B: 2
C: 3
D: 4
E: 5
F: 6
G: 7
H: 8
I: 9
J: 10
K: 11
L: 12
M: 13
N: 14
O: 15
P: 16
Q: 17
R: 18
S: 19
T: 20
U: 21
V: 22
W: 23
X: 24
Y: 25
Z: 26
First network drive (NetWare default)
First search drive (NetWare default)

Local Drive Mappings

Local drive mappings are paths to local media such as hard disk drives and floppy disk drives.

In DOS 3.0 and above, drives A: through E: are reserved for local mappings. To change this default, use the DOS LASTDRIVE command in your workstation CONFIG.SYS file.

Network Drive Mappings

Network drive mappings point to volumes and directories on the network. Drives F: through Z: can be used for network mappings. Each user can map drive letters to different directories.

To create a network drive mapping, use MAP or SESSION.

Search Drive Mappings

Network search drive mappings are pointers to directories containing applications, DOS files, etc.

Search drives let you execute a program even if it isn't located in the directory you're working in by enabling the system to search for the program.

Search drive mappings are numbered, although they also have drive letters. For example, search drive 1 (or S1) may also be known as network drive Z:.

You can map up to 16 network search drives (letters K: through Z:, starting with Z:). You can't map a search drive and a regular network drive to the same letter.

When you request a file and the system can't find it in your current directory, the system looks in every directory a search drive is mapped to until either the program file is found or can't be located.

Search drive mappings aren't supported on OS/2 workstations. The search functionality is provided with the OS/2 PATH, LIBPATH, and DPATH commands.

Viewing Drive Mappings

When you log in to the file server for the first time (before system and user login scripts are created), a list of default drive mappings appears.

You can view your drive mappings anytime by typing "MAP" at the command line. You can also set your login script to display your mappings every time you log in.

Related utility: "MAP" (Utilities Reference).

Driver

See “Disk Coprocessor board” on page 71; “LAN driver” on page 125.

Duplexing

See “Disk duplexing” on page 72.

Dynamic configuration

A means of allowing the file server to allocate memory according to need and availability.

When the file server boots, all free memory is assigned to file caching. As demand increases for other resources (directory cache buffers, for example), the number of available file cache buffers decreases.

The operating system does not immediately allocate new resources when a request is received. It waits a specified amount of time to see if existing resources become available to service the demand.

If resources become available, no new resources are allocated. If they don't become available within the time limit, new resources are allocated.

The time limit ensures that sudden, infrequent peaks of file server activity do not permanently allocate unneeded resources.

These following parameters are dynamically configured by the operating system:

Directory cache buffers
Disk elevator size
File locks
Kernel processes
Kernel semaphores
Maximum number of open files
Memory for loadable modules
Router/server advertising
Routing buffers
Service processes
TTS transactions
Turbo FAT index tables

Related utilities: "MONITOR"; "SET" (*System Administration*).

Dynamic memory

The most common form of memory, used for RAM.

Dynamic memory requires a continual rewriting of all stored information to preserve data.

A continuous electrical current is necessary to maintain dynamic memory. All data is lost when the power supply is turned off.

5 E

Effective rights

The rights that a user can actually exercise in a given directory or file.

Directory effective rights are determined by trustee assignments, if they exist.

Otherwise the effective rights of the current directory are determined by the intersection of the effective rights of the parent directory and the current directory's Inherited Rights Mask.

No effective rights exist in a volume's root directory until you assign trustee rights.

File effective rights are determined by trustee assignments to the file if they exist.

Otherwise, they are the same as the directory effective rights.

See also “Security” on page 221 (Effective rights).

Related utilities: "FILER"; "RIGHTS"; "WHOAMI" (*Utilities Reference*).

Elevator seeking

A method of organizing the way data is read from hard disk storage devices.

A shared network disk drive and its related channel can quickly become clogged with disk I/O requests. Elevator seeking logically organizes disk operations as they arrive at the server for processing.

A queue is maintained for each disk driver in the server. As read and write requests are queued for a specific drive, the operating system sorts incoming requests into an order of priority based on the drive's current head position.

As the disk driver services the queue, subsequent requests are located either in the vicinity of the last request or in the opposite direction. Thus, the drive heads operate in a sweeping fashion, from the outside to the inside of the disk.

Elevator seeking, therefore, improves disk channel performance by significantly reducing rapid back-and-forth movements of the disk head and by minimizing head seek times.

Embedded SCSI

A hard disk that has a SCSI and a controller board built into the hard disk unit.

Enable

1. To turn on, especially to restore a feature that has been disabled.
2. To place in a state that allows certain interrupts to occur in a processing unit (such as a network board). Interrupts are usually enabled by setting a switch or a jumper.

Encrypted password

A password that is scrambled before it is stored at the file server, to prevent an intruder from viewing or copying it.

Some encryption schemes encrypt the password at the workstation before it is transmitted to the file server. This prevents monitoring of the password over the transmission lines.

See also "Password" on page 179.

Erase right

See "Rights" on page 207; "Security" on page 221 (Rights Security).

Ethernet configuration

The setup that allows communication using an Ethernet environment.

NetWare uses the IEEE 802.3 standard by default. However, you can configure file servers, workstations, and routers to use the Ethernet II standard.

The IEEE 802.3 standard and the Ethernet II (or Ethernet) standard use different frame formats, as shown in Figure 28.

Figure 28 Ethernet frame formats

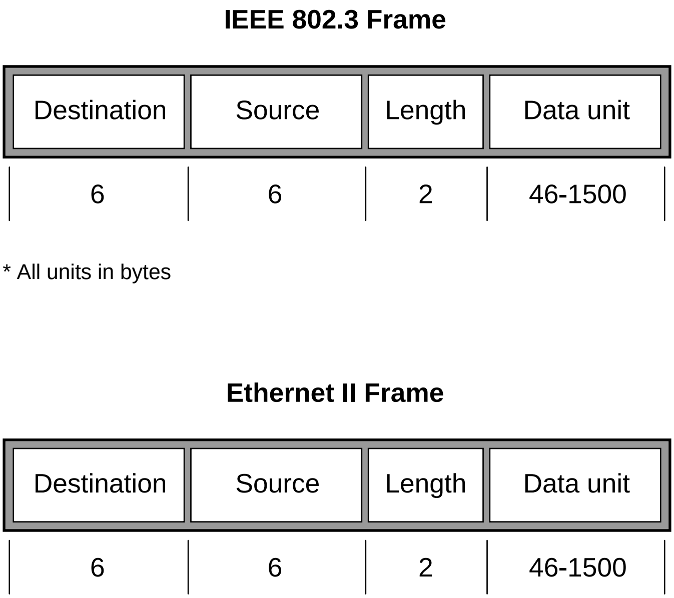

IEEE 802.3 Frame

Destination	Source	Length	Data unit
6	6	2	46-1500

Ethernet II Frame

Destination	Source	Length	Data unit
6	6	2	46-1500

* All units in bytes

* All units in bytes

Ethernet II frames contain a unique protocol ID or PID (represented in the "Type" field of the frame) that IEEE 802.3 frames do not contain.

Stations using the different standards cannot coexist on the same Ethernet cabling system. However, Ethernet II stations using different protocol numbers on an Ethernet II cabling system *can* coexist, but they *cannot* communicate with each other.

NOTE: To configure your workstations and routers to use the IEEE 802.2 or the Ethernet SNAP standard, you need to load the ODI driver. Refer to *Workstation for DOS and WIndows*.

Using Ethernet II

To configure file servers, workstations, and routers to use the Ethernet II standard, you must

- Specify Ethernet II when you load the Ethernet LAN driver during installation;
- Run the PROTOCOL file server utility if you want to specify a protocol other than Novell's IPX;
- Run the ECONFIG utility on the appropriate workstation and router software.

Configuring the file server. You must specify the Ethernet II standard with the FRAME parameter when you load Ethernet drivers. Then NetWare assigns Novell's IPX protocol number (8137) to the Ethernet frames.

To use a different protocol, use the NetWare PROTOCOL command to view a list of protocols registered with your file server. You must bind the protocol you want to the appropriate Ethernet driver using the BIND command.

Configuring workstations and routers. NET.CFG options allow you to configure NetWare workstation and router software so that the Ethernet II standard can be used.

When a station encounters a frame, it reads the protocol number for the frame type and either accepts it or sends it along the cable.

Figure 29 on page 93 illustrates how multiple networks can be configured on an Ethernet cabling system.

Figure 29 Multiple networks on Ethernet

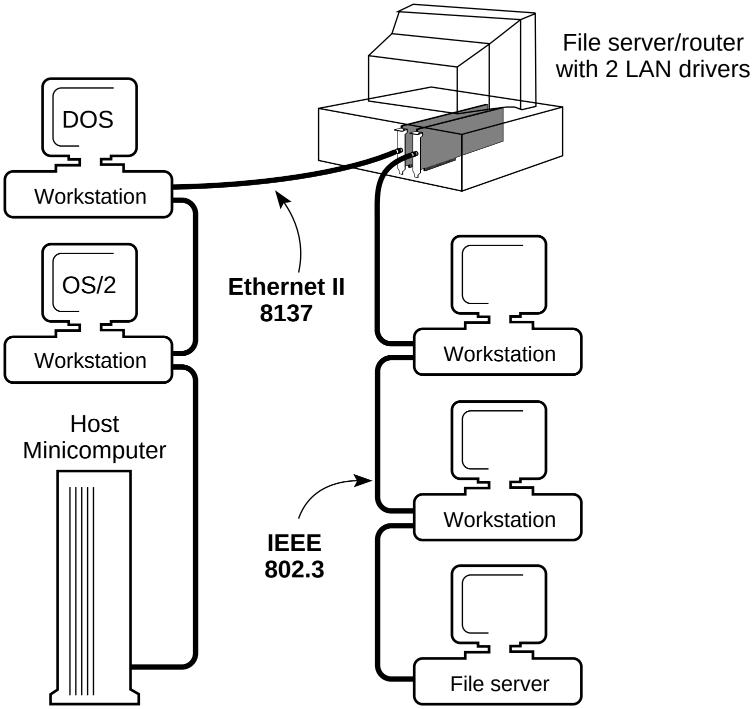

The VAX network was configured with ECONFIG. The utility was run on the OS/2 and DOS workstations *and* on the router LAN driver.

The Novell network was not configured with ECONFIG, since the IEEE 802.3 standard is the default setting.

EVERYONE

A system-created group that includes all users created on the file server.

Users are automatically added as members of EVERYONE when they are created. When all users need the same rights, you can grant those rights to the group EVERYONE.

Trustee Assignments

EVERYONE is assigned Read and File Scan rights to SYS:PUBLIC. These rights allow all users to run NetWare utilities, execute DOS commands, and access application programs residing in that directory.

EVERYONE is also automatically assigned the Create right to SYS:MAIL. This right allows any user to create and send mail to any other user.

The network supervisor can delete any user from the group EVERYONE or change EVERYONE's trustee rights to any directory.

Recommendations

We recommend that you do not remove EVERYONE's trustee assignments in the SYS:PUBLIC and SYS:MAIL directories. Without these rights users cannot access NetWare utilities, use electronic mail, or use print job configurations.

Do not delete the group EVERYONE from a file server. If EVERYONE is deleted and then re-created, you must add all users to the group individually and reassign trustee assignments.

See also “Group” on page 111; “User” on page 274.

Related utility: "SYSCON" (*Utilities Reference*).

Execute Only attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Expanded Memory shell

A version of the NETX shell, called EMSNETX, that moves most of the NETX shell out of conventional DOS memory and puts it in expanded memory (memory in addition to the 640 KB limit on conventional memory).

Virtual Loadable Modules™ (VLMs™) should be used in place of the NETX shell if possible, though both are provided with NetWare 3.12. If you are using Microsoft Windows or Packet Burst™, you must use the VLMs.

See also “NetWare DOS Requester” on page 156.

Extended Memory requirements

See “File server” on page 106.

Extended Memory shell

A version of the NETX shell, called XMSNETX, that moves most of the NETX shell out of conventional DOS memory and puts it in extended memory (memory above 1 MB).

Virtual Loadable Modules (VLMs) should be used in place of the NETX shell if possible, though both are provided with NetWare 3.12. If you are using Microsoft Windows or Packet Burst, you must use the VLMs.

See also “NetWare DOS Requester” on page 156.

External bridge

See “Router” on page 211.

6 F

Fake root

A subdirectory that functions as a root directory.

Some applications cannot be run from subdirectories; they read files from and write files to the root directory. However, for security reasons, you shouldn't assigned rights at the root directory level.

NetWare allows you to map a search drive to a fake root (a directory where rights can be assigned to users).

NOTE: Fake roots work only with NetWare v3.0 and above of NET3.COM and above. If you use older versions of the workstation software, you cannot create fake roots.

To use an application that must be installed at the root, load the files in a subdirectory and designate the subdirectory as a fake root directory, using a command in the login script.

You cannot use the DOS CD (change directory) command at the fake root to return to the original root. To change the fake root back to the original root, remap the drive.

FAT

See "File Allocation Table" on page 98.

Fault tolerance

See "SFT^{TM}" on page 249.

Ferro-resonant isolation transformers

See “Power conditioning” on page 180.

File Allocation Table

An index table that points to the disk areas where a file is located.

Because one file may be in any number of blocks spread over the disk, the FAT links the file together.

In NetWare, the FAT is accessed from the directory table. The FAT is cached in server memory, allowing the server to access the data quickly.

Each volume contains a FAT. NetWare divides each volume into disk allocation blocks that can be configured to 4, 8, 16, 32, or 64 KB. (All blocks on one volume are the same size.)

NetWare stores files on the volume in these blocks. If a file comprises one or more blocks, the file may be stored in blocks that aren't adjacent.

Entries in the FAT correspond to the blocks for that volume. The first entry in the FAT corresponds to the first block on that volume, the second entry corresponds to the second block, etc.

Turbo FAT index. When a file exceeds 64 blocks (and the corresponding number of FAT entries), NetWare creates a turbo FAT index to group together all FAT entries for that file.

A turbo FAT index enables a large file to be accessed quickly.

The first entry in a turbo FAT index table is the first FAT number of the file. The second entry is the second FAT number, etc.

File attributes

See “Attributes” on page 21; “Security” on page 221.

File caching

The use of file server RAM to improve file access time.

Cache memory allocates space for the hash table, the FAT, the turbo FAT, the directory cache, a temporary data storage area for files and NLMs™, and an open space for other functions. See Figure 30 on page 99.

If the cache memory uses the default block size (4 KB) and a file takes more than one block, the file is placed in a second 4KB non-contiguous block both in cache memory and on the volume (on the fixed disks).

NOTE: If you prefer to save data in larger blocks, you can use SERVER with the -C parameter to change the size of the blocks to 8 KB or 16 KB.

Figure 30 Cache memory in RAM

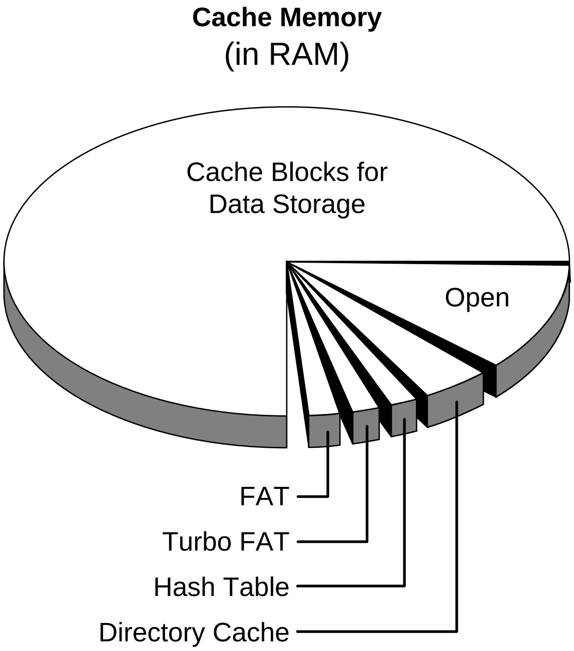

Reading Files from Cache

Workstations can access the file server's cache memory up to 100 times faster than they can access the file server's fixed disks.

When a workstation makes a read request from the file server, the file server executes a hash algorithm to predict a file address from a hash table.

Hashing is a quick way of predicting the file's address in the directory table. See Figure 31.

Figure 31 Reading files from cache

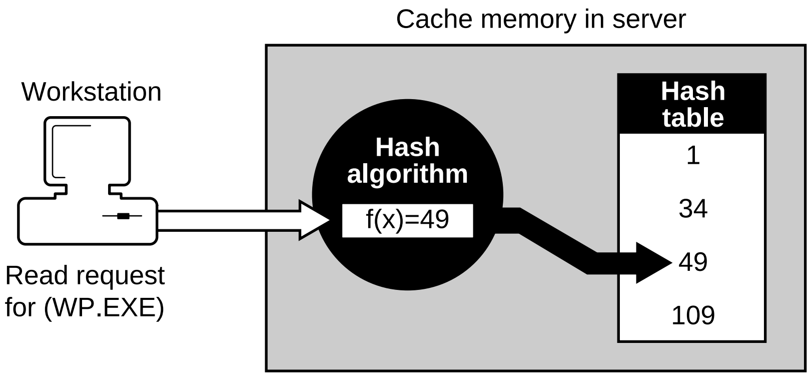

For example, the address on the hash table contains pointers to the first and second probable locations of the WP.EXE file's directory entry in the directory cache.

If the first search is not successful, the file server uses the second pointer to find the directory entry. See Figure 32.

Figure 32 Reading files from cache

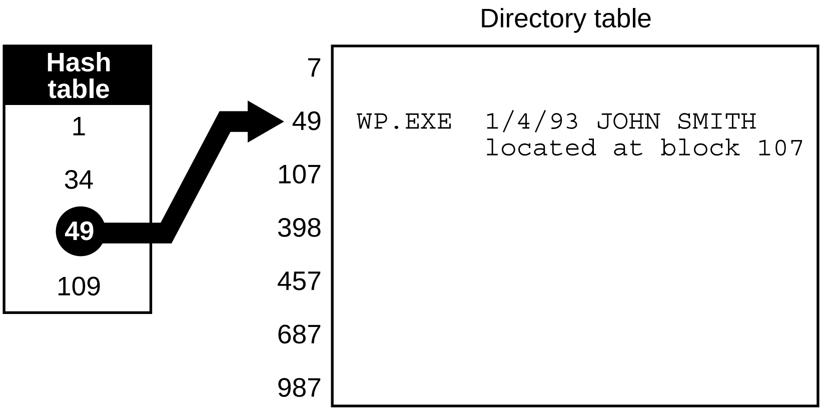

When the directory entry is located, the file server checks its cache memory to see if it has a copy of the WP.EXE file. See Figure 33.

Figure 33 Reading files from cache

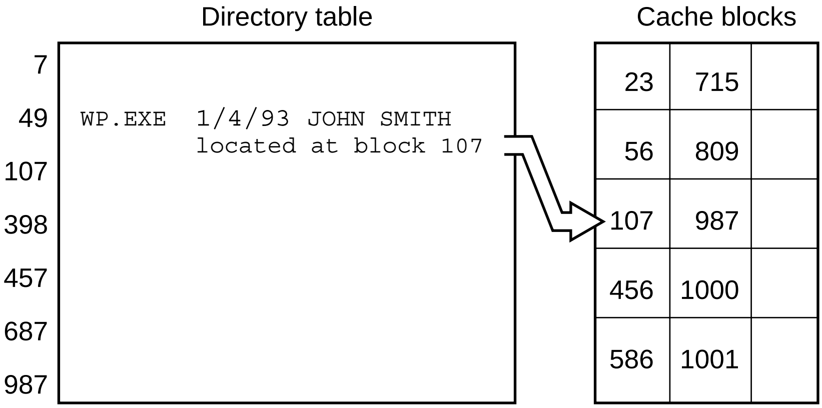

If the file is there, the server sends the file to the workstation directly from cache memory. If the file is not there, the file server retrieves the file from the hard disk and sends it to the workstation.

Writing Files to Cache

When a workstation writes a file to the file server, the file server performs the hash algorithm to find the file's cache buffer.

It writes the file to the designated location and updates the directory table in the directory cache. Since the file has changed, its cache buffer becomes "dirty."

"Dirty" cache buffers indicate that the file in cache memory is different from the file on disk. Since writes to disk take longer to perform than writes to cache, the file server keeps the "dirty buffer" designation on the file in cache until the disk receives the file. See Figure 34.

Figure 34 Writing files to cache

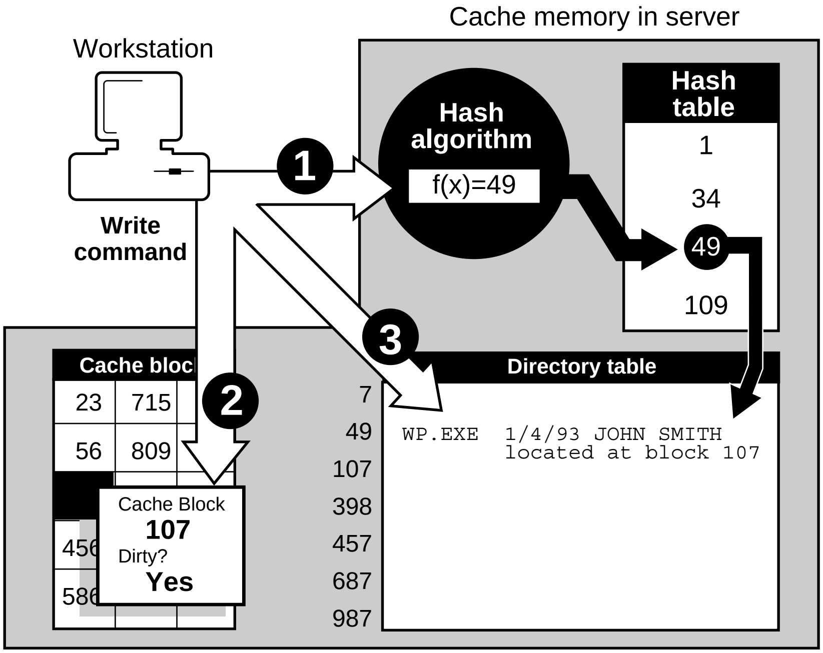

1. Performs hash algorithm
2. Writes file to cache blocks
3. Updates directory table

The file server sends a message to the workstation that the server has received the file, and the workstation is free to complete other processes.

Once the file is written to disk, the file server checks the data in memory against the data on disk. If the data matches, the buffer is no longer dirty.

As the cache memory fills up, buffers containing least-used files and directories are eliminated. See Figure 35.

Figure 35 Writing files to cache

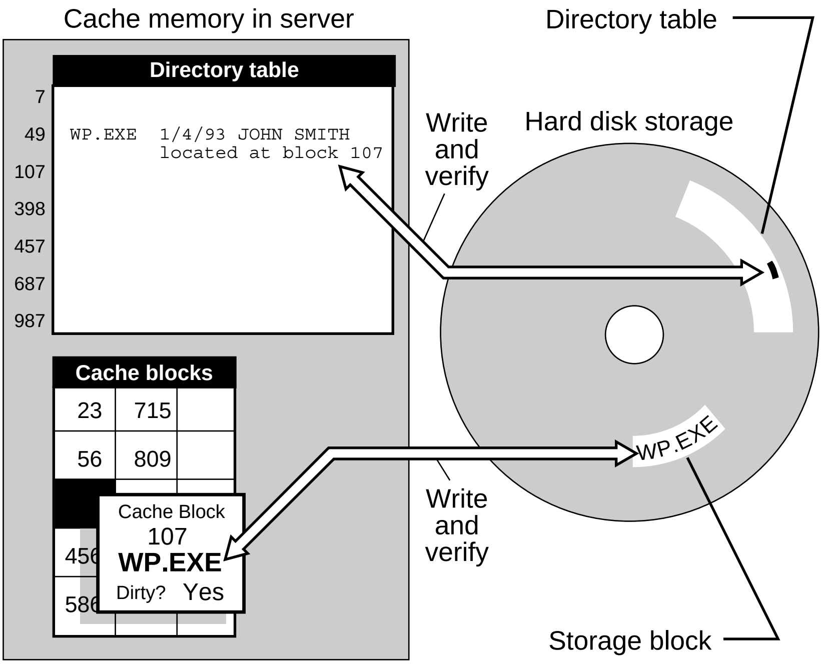

See also “Directory caching” on page 58; “Hashing” on page 116.

Related utility: "SET" (*System Administration*).

File extension

A 3-character addition to a DOS filename. Most extensions are assigned arbitrarily, but some are reserved for special purposes:

- An extension such as .COM, .EXE, or .BAT enables a file to execute a command, program, or batch file.
- Some extensions indicate that the file has a unique function, for example, .ERR, .DAT, and .NLM.
- Some extensions identify the file format, such as .ASC and .BIN.
- Some extensions are assigned by certain companies.
- Some extensions indicate specific programming languages, like .C, .ASM, and .BAS.

DOS and NetWare Extensions

Table 5 shows extensions that DOS and NetWare files share.

Table 5 File extensions shared by DOS and NetWare

Extension	Meaning
BAT	DOS executable batch file
COM	DOS executable command file
DAT	ASCII text file
ERR	Error log file
EXE	DOS executable file
HLP	Help screens in a menu utility
OVL	Overlay file used as part of NetWare files created from the C-Worthy program to share code
SYS	Operating system file

NetWare-Specific Extensions

Table 6 shows extensions that are NetWare-specific.

Table 6 NetWare-specific extensions

Extension	Meaning
DKD	NetWare disk drive configuration file
DSK	NetWare disk driver
LAN	NetWare LAN driver
NAM	NetWare name support
NFC	NetWare executable batch file used to load modules and set the NetWare operation system configuration
NLM	NetWare Loadable Module™
PDF	NetWare printer definition file
Q	NetWare print job file
QDR	NetWare print queue definition directory
SRV	NetWare operating system file for the print queue

File indexing

The method of indexing FAT entries for faster access to large files.

With file indexing, for example, you could go directly to block 128 of a file instead of scanning through the 127 previous blocks.

NetWare v3.12 supports automatic file indexing above 64 blocks. Two levels of file indexing in NetWare v3.12 refer to the size of the table it uses to index the FAT.

The first level indexes 64 to 1,023 blocks; the second level, 1,024 or more blocks.

File locking

The means of ensuring that a file is updated correctly before another user, application, or process can access the file.

For example, without file locking, if two network users were to try to update the same word-processing file simultaneously, one user could overwrite the file update of the other user.

See also "Record locking" on page 202.

File rights

See "Rights" on page 207; "Security" on page 221 (Rights Security).

File Scan right

See "Rights" on page 207; "Security" on page 221 (Rights Security).

File server

A computer that runs NetWare operating system software.

The NetWare operating system enables the file server to regulate communications among the personal computers attached to it and to manage any shared resources (such as printers).

File Server Requirements

A v3.12 file server must

- Have at least one hard disk, either internal or external;
- Have 4 MB of RAM;
- Contain at least one network board;
- Be used only as a dedicated file server.

Booting the File Server

You can set up the file server to boot automatically by placing an AUTOEXEC.BAT file on either the server's boot diskette or a DOS partition on the file server's DOS-addressable internal hard disk.

The AUTOEXEC.BAT file executes the SERVER command.

To complete the boot process, you also need NetWare Loadable Modules (.NLM) for the cabling systems attached to the server and one disk driver (.DSK) for the architecture handling the server's hard disks (AT or compatible, PS/2 or compatible, or disk coprocessor board).

The NetWare boot files, AUTOEXEC.NCF and STARTUP.NCF, are created as part of the installation process and must be included on the file server's boot diskette or DOS partition.

The AUTOEXEC.NCF file

- Stores the file server's name and internal network number;
- Loads the settings for the network boards and binds the communication protocol (such as IPX) to the installed drivers;
- May contain executable file server commands (such as TRACK ON or MONITOR).

The STARTUP.NCF file loads the file server's disk driver.

Related utility: "INSTALL" (*System Administration*).

Bringing Down the File Server

Changes made to data files are often held in the file server's cache buffers until the file server processor has idle time. When you execute the DOWN command, the file server writes the data in the cache buffers to disk, thus preventing data from being lost.

Users on the file server should be notified that the server is going down so that they can save and exit their files.

Related utilities: "BROADCAST"; "DOWN" (*System Administration*).

File Server Console Operator

A user or a member of a group to whom SUPERVISOR delegates certain rights to manage the file server.

A Console Operator has rights to use the FCONSOLE utility. This utility allows the Console Operator to do the following:

- Broadcast messages to users
- Change file servers
- Access connection information
- View NetWare version information
- Change the system date and time
- Enable or disable login for additional users
- Enable or disable TTS

See also “SUPERVISOR” on page 255.

Related utilities: "FCONSOLE"; "SYSCON" (*Utilities Reference*).

File sharing

A feature of networking that allows more than one user to access the same file at the same time.

See “Attributes” on page 21.

Flag

See “Attributes” on page 21.

Form

In a NetWare printer command, the name and size of the paper used for a print job.

Related utility: "PRINTDEF" (*Print Server*).

Format

See “Disk formatting” on page 74; “Surface test” on page 257.

Frame

1. A variation of a protocol, such as Ethernet II, Ethernet 802.3, Ethernet 802.2, Ethernet SNAP, Token-Ring, or Token-Ring_SNAP.

   See “Ethernet configuration” on page 91.

2. A packet.

   See “Message packet” on page 142.

7 G

Gateway

A link between two networks.

A gateway allows communication between dissimilar protocols (for example, NetWare and non-NetWare networks) using industry standard protocols such as TCP/IP, X.25, or SNA.

Group

A network convention that allows you to deal with users collectively rather than individually.

When users are created, they automatically become members of the group EVERYONE and have the rights assigned to that group.

You can use SYSCON to create other groups based on who uses the same applications, printers, or print queues, who performs similar tasks, or who has similar needs for information. Users can belong to a maximum of 32 groups.

You can use groups to simplify trustee assignments and login scripts. With SYSCON and FILER, you can make group trustee assignments for directories and files.

By using groups (formed for common application use) and the conditional IF...THEN, you can map a search drive to that application directory.

You can also use IF...THEN with a group name to exit to a menu created for that group. Or you can prepare a login message to be displayed for group members when they log in.

See also “EVERYONE” on page 93; “Login script” on page 128; “User” on page 274.

Related utility: "SYSCON" (*Utilities Reference*).

GUEST

A username for anyone who needs temporary and restricted access to the file server.

NetWare creates the user GUEST automatically.

You should evaluate the security needs of the network and determine whether to retain GUEST. You must also determine what rights temporary users can exercise and what information they can access.

If someone needs to access the file server for a definite period of time, you might prefer to create a username and user account with an expiration date.

GUEST's Rights

GUEST is automatically a member of the group EVERYONE, and GUEST's rights flow from membership in that group. Any trustee assignments you make to the group EVERYONE apply to GUEST as well.

As a member of EVERYONE, GUEST is granted rights to the SYS:PUBLIC directories. These rights allow all users to run NetWare utilities, execute DOS commands, and access application programs residing in that directory.

GUEST is also assigned a mailbox and the Create right to the SYS:MAIL directory. This right allows GUEST to create and send electronic mail.

Managing GUEST

The GUEST account has no initial password, but you can require a password on GUEST's account. We suggest you assign GUEST a password and change it frequently. Do not allow GUEST to change the password.

For maximum security, you can also take any of the following measures.

- Delete GUEST from the group EVERYONE and make a specific trustee assignment to GUEST, such as EVERYONE's rights to SYS:PUBLIC, the DOS directory, and whichever applications you permit GUEST to use.
- Delete GUEST from the file server if you have no temporary users.

If you want GUEST to have more privileges, you can

- Assign GUEST temporary rights to specific directories and files;
- Create a GUEST subdirectory in the HOME or USERS directory to provide personal work space for GUEST.

See also “User” on page 274.

8 H

Handshaking

The initial exchange between two data communication systems before and during data transmission to ensure proper data transmission.

A handshake method (such as XON/XOFF) is part of the complete transmission protocol.

A serial (asynchronous) transmission protocol might include the handshake method (XON/XOFF), baud rate, parity setting, number of data bits, and number of stop bits.

See also “Serial communication” on page 246.

Hard disk

A high-capacity magnetic storage device that allows a user to write, read, and erase data.

Hard disks can be network disks or local workstation disks.

See also “Data protection” on page 49; “Device numbering” on page 54; “Disk Coprocessor board” on page 71; “Partitions” on page 177.

Hardware interrupt

See “Interrupt” on page 121.

Hashing

A process that facilitates access to a file in a large volume by calculating the file's address both in cache memory and on the hard disk.

When a workstation wants to read a file from the file server, the server performs a hash algorithm that predicts an address on a hash table.

In NetWare v3.12, it is common to find the file 95% of the time on the first try. This method is much more efficient than searching for a file sequentially.

See also "Directory caching" on page 58; "File caching" on page 99.

HBA

See "Disk Coprocessor board" on page 71.

Hexadecimal

A base-16 numeric notation system used to specify addresses in computer memory.

In hexadecimal notation, the decimal numbers 0 through 15 are represented by the decimal digits 0 through 9 and the alphabetic digits A through F (A = decimal 10, B = decimal 11, etc.).

Hidden attribute

See "Attributes" on page 21; "Security" on page 221 (Attribute Security).

Home directory

A network directory that the network supervisor creates specifically for a user.

The user's login script should contain a drive mapping to his or her home directory.

Host bus adapter

See "Disk Coprocessor board" on page 71.

Hot Fix

A method NetWare uses to ensure that data is stored safely.

Data blocks are redirected from faulty blocks on the server's disk to a small portion of disk space set aside as the *Hot Fix redirection area*.

Once the operating system records the address of the defective block in a section of the Hot Fix redirection area reserved for that purpose, the server won't attempt to store data in defective blocks.

By default, 2% of a disk partition's space is set aside as the Hot Fix redirection area.

Hot Fix is always activated unless the disk fails or the redirection area is full. Hot Fix, together with read-after-write verification, enables a hard disk to maintain data integrity.

See also “Read Write attribute” on page 202.

Hub

A device that modifies transmission signals, allowing a network to be lengthened or expanded with additional workstations.

See also “Active hub” on page 19; “Passive hub” on page 178.

9 I

Identifier

See “Login script” on page 128.

Indexed attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Inherited Rights Mask

A list of rights given to each file and directory when it is created that controls which rights users can inherit.

The IRM for any given file or directory is modified by revoking rights.

The *directory's* IRM controls which parent directory effective rights can be exercised in the current directory.

The *file's* IRM controls which of the current directory's rights can be exercised in the file.

See also “Security” on page 221 (Rights Security).

Related utilities: "ALLOW"; "FILER" (*Utilities* Reference).

Initialize hard disk

See “Surface test” on page 257.

Interleave factor

A method of adjusting the speed of the controller to match the speed of the hard disk.

Typically, a hard disk spins faster than a controller can perform a read/write. If the controller is not fast enough to read or write consecutive sectors on a hard disk track, then the controller can be programmed to skip one or more sectors of the hard disk before the next read/write is performed.

If the controller reads or writes one sector and then skips a sector, the interleave factor is 2 (one out of every two sectors is used). The interleave factor can be written as 2:1.

If the controller reads or writes one sector, then skips two sectors (one out of three sectors is used), the interleave factor is 3 or 3:1.

The interleave factor is usually established by the manufacturer or reseller of the hard disk/controller combination.

Someone other than the manufacturer or reseller who puts together the hard disk/controller combination may need to experiment to determine the correct interleave factor.

Internal bridge

See “Router” on page 211.

Internal network number

See “IPX internal network number” on page 122.

Internetwork

Two or more networks connected by an internal or external router.

Users on an internetwork can use the resources (such as files, printers, or disk drives) of all connected networks.

Interrupt

A signal or call to a specific routine.

The microprocessor suspends the current program until the routine is completed. The processor continues with the original program after the routine is completed.

Interrupts are divided into two general types: hardware and software.

A *hardware interrupt* is caused by a signal from a hardware device, such as a printer.

A *software interrupt* is created by instructions from within a software program.

I/O address

See "Base I/O address" on page 29.

IPX

(Internetwork Packet Exchange) A Novell communication protocol that sends data packets to requested destinations (workstations, servers, etc.).

IPX addresses and routes outgoing data packets across a network. It reads the assigned addresses of returning data and directs the data to the proper area within the workstation's or file server's operating system.

IPX is closely linked with other programs and routines that help in the network data-transmission process.

The NetWare DOS Requester prepares data packets in a form understandable to the intended destination before handing them to IPX.

The IPXODI.COM file then uses the services of a LAN driver routine to control the station's network board for data delivery.

Therefore, IPX can route and accept data packets through physically different networks and workstations.

See also "Communication protocols" on page 41; "LAN driver" on page 125; "Open Data-Link Interface" on page 173.

IPX internal network number

A "logical" network number that identifies the individual file server.

The internal network number for each file server must be unique, hexadecimal, and from one to eight digits long.

You are prompted to enter internal network number when you bring up the file server. Then when you want view the file server's internal configuration by using the CONFIG utility, the number appears along with the file server's name and other information.

See also “Network numbering” on page 165.

Related utility: "CONFIG" (System Administration".

IRQ

See “Interrupt” on page 121.

Isolation transformer

See “Power conditioning” on page 180.

10 J

Jumper block

A group of jumper pins used to make hardware configuration settings on a printed circuit board.

11 L

LAN driver

A software routine that understand and controls the physical structure of a network board.

A LAN driver serves as the connection between a station's operating system and the physical network parts.

LAN drivers are specific to a particular network board, but when you run WSGEN and adapt the driver to the IPX protocol, you allow the operating system to communicate on a network regardless of the type of network board.

IPX simply passes information to the LAN driver and lets the LAN driver direct the board on transmission procedures.

After you install the network board in the file server, you load the LAN driver into the operating system and then bind it to the communications protocol.

See also “Binding and unbinding” on page 31; “IPX internal network number” on page 122; “NetWare DOS Requester” on page 156; “NetWare Runtime” on page 161.

Related utilities: "BIND"; ""LOAD LAN driver"; "UNBIND" (System Administration).

Line-surge suppressor

See “Power conditioning” on page 180.

Link Support Layer

(LSL™) An implementation of the Open Data-Link Interface specification that serves as an intermediary between the file server's LAN drivers and communication protocol, such as IPX, AFP, or TCP/IP.

The LSL allows one network board to service several communication protocol stacks.

The LSL also allows several network boards to service the same protocol stack.

See also “Open Data-Link Interface” on page 173.

Loadable module

See “NetWare Loadable Module” on page 158.

Loading and unloading

The process of linking and unlinking NLMs to and from the NetWare operating system.

NLMs are linked to the operating system with the LOAD command; they are unlinked with the UNLOAD command.

When a module is linked with the operating system, the operating system allocates a portion of the memory to the module. The amount of memory the module uses can vary.

When the task is completed, the module returns some of the memory to the operating system. When a module is unloaded, all allocated resources are returned to the operating system.

All modules can be unloaded while the file server is running. See "UNLOAD" (System Administration) for precautions you should take before unloading a disk or LAN driver.

See also “Loading and unloading” on page 126.

Related utilities: "LOAD"; "LOAD disk driver"; "LOAD LAN driver"; "LOAD name space;" "LOAD NLM utility;" "UNLOAD" (System Administration).

Log in

To initialize the user's security rights and the user's environment by using the LOGIN command.

When a user initiates a login request, the operating system scans the bindery and reads the user's bindery information into memory. The user is then asked for a password.

All security information is then placed in the file server's connection list and the user is said to be "logged in."

At this point, the login program executes the login script to initialize environmental variables, map network drives, and control the user's program execution.

Related utility: "LOGIN" (*Utilities Reference*).

Log out

To break the file server/workstation connection and delete any drives mapped to that server by using the LOGOUT utility.

Related utility: "LOGOUT" (Utilities Reference).

LOGIN directory

A directory (SYS:LOGIN) that is created by the system during network installation and cannot be deleted.

The LOGIN and SLIST utilities are copied into the LOGIN directory. Users can log in to the network and view a list of available network servers from this directory.

Login restrictions

Limitations on a user account that control access to the network, including the following:

- *Requiring a password*. If you require a password, you can specify its minimum length, whether it must be changed (and how often), whether it must be unique, and whether the user can change it.

  You can also specify the number of times a user can log in using an expired password and the number of incorrect login attempts allowed.

- *Setting account limits*. If you install accounting, you can assign account limits, such as an account balance or expiration date.
- *Setting an expiration date.* You can specify an expiration date for a user account. The account expires at 12:01 a.m. the next day.
- *Limiting disk space*. You can limit the amount of disk space for users by specifying the maximum blocks available for each user on a volume.
- *Specifying the number of connections*. You can limit the number of times a user can log in simultaneously. You can also specify, by node address, which workstations users can log in on.
- *Setting time restrictions*. You can restrict the hours during which users can log in. You can assign all users the same hours, or you can restrict users individually.

When a user violates login restrictions, NetWare disables the account and no one can log in using that username. This prevents unauthorized users from logging in.

Login script

A file containing commands that set up your users' workstation environments whenever they log in.

Login scripts are similar to configurable batch files and are executed by the LOGIN utility.

After a user initiates a login request and supplies the correct username and password, the login program executes the login script.

Purpose of Login Scripts

You can use login scripts to

- Map drives and search drives to directories.
- Display messages.
- Set environment variables.
- Execute programs or menus.

Login scripts work the same way for DOS, Windows, and OS/2 workstations. (Some commands that apply to DOS and Windows workstations may not apply to OS/2 workstations.)

Types of Login Scripts

NetWare uses two kinds of login scripts.

System login scripts allow the network supervisor to set network drive mappings and search drive mappings for all users. The system login script includes commands that should be executed for every user.

The system login script is created in SYSCON and saved in the SYS:PUBLIC directory (as NET$LOG.DAT).

User login scripts specify the users' individual drive mappings and environment variables. A user login script is created for each user in SYSCON and saved in an ID subdirectory for that user in the SYS:MAIL directory.

If you do not create a user login script, a default user login script executes.

Guidelines for Login Scripts

Use the following guidelines when you plan login scripts:

- Do not include passwords or proprietary information in the system login script.
- Create a system login script before you create user login scripts.

- Include a login script for each user, however minimal.

  Since the login script is kept in the user's mail directory, it is possible for an intruder to create a login script file for any user who does not already have one.

  If you access the box for a user login script in SYSCON and put in even a blank space, the system considers a user login script to exist and prevents the user from accessing the default login script when logging in.

Login Script Conventions

When you create a login script, follow the command format described in "Login Script Commands and Variables" in *Installation.*

The command format specifies the syntax for keywords, variables, parameters, spacing, delimiters, and other characters and punctuation.

The following conventions apply when you create login scripts:

- Login script commands are not case-sensitive.

  However, any identifier variables enclosed in quotation marks must be preceded by a percent sign (%) and typed in upper-case letters.

- Only one command can be entered on each line.
- Command lines cannot exceed 150 characters.

  To increase readability we recommend that you use only 78 characters per line—the width of your screen.

  Consider leaving a line space between groups of commands; for example, put all mappings in a block set off from commands that come before and after.

- Press the Enter key at the end (and only at the end) of each command.

  If a long command continues beyond the width of the screen, the words that wrap automatically onto the next line are still considered part of the command.

- Use REMARK to include explanatory comments so that you will have a record of the purpose of each command or block of commands.

  The comments and notes included in a login script with REMARK are not displayed when the login script is executed.

Commands Used in Login Scripts

Some login script commands have the same name as corresponding commands with similar functions from other systems. Login script commands are based on the following.

- NetWare utilities

  For example, the login script commands MAP and ATTACH have functions analogous to the NetWare utilities.

- DOS commands

  For example, you can set the DOS environment from within the login script because DOS BREAK, SET, and VERIFY have the same function in a login script.

- Programming languages

  For example, the login script command IF...THEN...ELSE is similar to conditional commands in most programming languages (and can also be nested up to 10 levels in versions of NetWare above 3.0).

You can also use the login script to pass a command to the command line (EXIT or # command) or to display text (WRITE or FDISPLAY command).

Login script commands are described individually in Appendix A of *Installation.*

What Should a System Login Script Provide?

The more a system login script accomplishes, the shorter user login scripts can be.

Some commands are essential in a login script, some are recommended, and others are optional or dependent on the needs of your network.

Essential commands provide access to NetWare and DOS.

- NetWare utilities

  Use MAP INSERT to map the first search drive to SYS:PUBLIC. This provides access to NetWare utilities from any directory.

- DOS

  Use MAP INSERT to map the second search drive to the DOS directories. This mapping provides access to DOS system files and commands from any directory.

  For guidelines, see "DOS directories" on page 77.

- Command interpreter

  Use COMSPEC to ensure that the transient portion of COMMAND.COM reloads properly into each workstation when an application is exited. COMSPEC should specify the search drive mapped to the DOS directories.

Recommended commands provide access to frequently used directories.

- Use [DOS] SET to set the prompt to display the current directory path, so that users know where they are in the directory structure:

  ```
  SET PROMPT="$P$G"
  ```

- Application directories

  For example, if all or most users access a word processing program, you could include a search mapping similar to the following:

  ```
  MAP S3:=SYS:APPS/WORDPROC
  ```

  If you have formed groups on the basis of application use, you can assign drive mappings only to the users who need them by using the conditional IF...THEN.

- Home or username directories

  To map the first network drive to each username directory stored in the SYS:HOME directory (or other parent directory), you can use an identifier variable.

  The drive mapping would be similar to the following:

  ```
  MAP F:=SYS:HOME/%LOGIN_NAME
  ```

- Work or database record directories

  To map the second network drive to a directory providing group workspace, you can use a mapping similar to either of the following:

  ```
  MAP G:=SYS:OFFICE/TRAININGMAP *2:=SYS:PROJECTS/PLANS
  ```

Optional commands can include or provide for the following.

- Messages

  Use WRITE to display a brief message. For example, you could display a message by using WRITE commands (including identifier variables) similar to the following:

  ```
  IF DAY_OF_WEEK = "SUNDAY" THENWRITE "WHY ARE YOU
    WORKING TODAY?"END
  ```

  You can also use FDISPLAY to display longer messages created as text files.

- Settings for environmental variables

  Use [DOS] SET to set application environmental variables.

  To determine which variables need to be set, see the documentation for your application program. Generally, variables that can be set in an AUTOEXEC.BAT file can be set in a login script.

- Login script batch files

  Use INCLUDE to call up batch files you have created to shorten the login script, to display messages created as text files with WRITE, or to customize the system login script for groups.

- Additional file servers

  Use ATTACH to access other file servers on an internetwork. For example, to provide access to the file server ORACLE, you could include a command similar to the following:

  ```
  ATTACH ORACLE
  ```

- Comments and notes for the network supervisor

  Use REMARK or its aliases (REM, asterisk, and semicolon) to include explanatory comments so that you have a record of each command's purpose.

What Should a User Login Script Contain?

If you have accomplished your main purposes in the system login script, user login scripts need to contain only commands that apply to the individual user.

Users can change their own login scripts if they are allowed to change their passwords.

A user login script can contain the following:

- Drive mappings
- Environmental variables for applications
- An exit to a menu

  Users can exit to a menu from the user login script. The EXIT command can call up a menu created for an individual user.

Example of a System Login Script

This section provides an example of a system login script.

Assume the following for the example:

- File server ELIOT has three volumes: SYS:, BOOM:, and BAH:.
- The users work at the command line.
- The workstations run more than one version of DOS.
- Two workstations must be able to transfer files to UNIX workstations.
- Word processing, office productivity, and electronic mail programs are installed.
- BOOM: is specified as the master volume for mail.
- The network supervisor has created daily greeting messages and a monthly reminder to users (introduced with a beep).
- The network supervisor also maintains a log file containing login and logout times for users (the accounting feature is not installed).

The network supervisor has created the following login script.

```
map display offdos set mv="eliot/boom:"

map *1:=eliot/sys:map ins s1:=eliot/sys:publicmap ins s2:=eliot/sys:public/
%machine %os_versioncomspec = s2:command.commap ins s3:=eliot/sys:public/wordp-
42rem user-mapped search drives 4-9

rem mappings for mail and office appsmap ins s10:=eliot/boom:mhs/exemap ins
s11:=eliot/boom:atc/exemap ins s12:=eliot/boom:officemap ins s13:=eliot/
boom:wpmail

rem electronic publishing transmission control protocolif
p_station="0000000000F8" then map ins s4:=c:\pctcpendif
p_station="0000000000F7" then map ins s4:=c:\pctcpend
```

```
map display on
```

```
rem daily greetingif NDAY_OF_WEEK = "1" and hour24 < "09" thendisplay
sys:public/hello1.msgendif NDAY_OF_WEEK = "2" and hour24 < "09" thendisplay
sys:public/hello2.msgendif NDAY_OF_WEEK = "3" and hour24 < "09" thendisplay
sys:public/hello3.msgendif NDAY_OF_WEEK = "4" and hour24 < "09" thendisplay
sys:public/hello4.msgendif NDAY_OF_WEEK = "5" and hour24 < "09" thendisplay
sys:public/hello5.msgendif NDAY_OF_WEEK = "6" and hour24 < "09" thendisplay
sys:public/hello6.msgendif NDAY_OF_WEEK = "7" and hour24 < "09" thendisplay
sys:public/hello7.msgend
```

```
rem monthly reminderif DAY < "07" and NDAY_OF_WEEK = "2" and hour24 < "10" then
write "\n\7\7 A Monthly Reminder:"     write "Please delete any unnecessary
files you own on"      write "Eliot (including old e-mail).\n"      pauseend
```

```
rem user logins#command /c z:logfil %LOGIN_NAME %P_STATION %DAY_OF_WEEK
%MONTH_NAME %DAY %YEAR %HOUR %MINUTE %AM_PM
```

Examples of User Login Scripts

This section provides two examples of user login scripts.

Example 1

User JSBACH on the file server ELIOT has customized this user login script by adding a message and firing phasers during login:

```
set wp = "/b-10/u-jsb/"set usr = "jsbach"set pwd = ""attach ENTERPRISE
```

```
MAP INS S5:=ENTERPRISE\SYS:HISTORY\1988MAP INS S6:=SYS:HOME\JSBACH\MACROSMAP
INS S7:=SYS:PUBLIC\UTIL
```

```
#newmail eliot/boom: jsbach
```

```
WRITE ""IF DAY_OF_WEEK = "MONDAY" THENWRITE "Sic Transit Gloria Mundi."WRITE
""FIRE PHASERS 2 TIMESEND
```

Example 2

User JAUSTEN on the file server ELIOT has customized the following user login script by displaying a random quote at login.

```
write ""write ""
```

```
map *4:=boom:print\utilmap *5:=bah:humbug
```

```
set usr "jausten"#newmail eliot/boom: jaustenpause#basica quotepause
```

Model Login Script for an Internetwork

An MIS manager developed the following model to help network supervisors standardize system login scripts for a large internetwork.

Assume the following for the model:

- The model provides for word processing, accounting programs, and electronic mail.
- All file servers in each division begin with the same letter. The directory structure is standardized.
- DOS directories are named in conformity with the pattern in the Search2 mapping.
- Username directories are in a separate volume named HOME:.
- Accounting work directories are in VOL1:.
- Each database is assigned one or more volumes.
- Printer 0 is hard wired and users can print from applications. Other printers require the CAPTURE command.
- Power users are given access to additional batch files and utilities. A third-party auditing utility is used on most file servers.

```
MAP DISPLAY OFFWRITE "Good %GREETING_TIME, %FULL_NAME."WRITE "You are logged
onto connection %STATION."

; environment mappingsMAP INS S1: = SYS:PUBLICMAP INS S2: = SYS:PUBLIC/%MACHINE
%OS_VERSIONCOMSPEC = S2:COMMAND.COMSET PROMPT = "$P$G"DOS SET MV = server/SYS:

; personal directory mappingsIF MEMBER OF "HOME" THEN      MAP P: =
HOME:\%LOGIN_NAME      DRIVE P:END

IF MEMBER OF "HOME2" THEN      MAP P: = HOME2:\%LOGIN_NAME      DRIVE P:END

; default printer mappings by groupIF MEMBER OF "PGROUP1" THEN #CAPTURE
Q=PRINTER1 nb nff ti=3 END

IF MEMBER OF "PGROUP2" THEN#CAPTURE Q=PRINTER2 nb nff ti=3END

IF MEMBER OF "PGROUP3" THEN#CAPTURE Q=PRINTER3 nb nff ti=3END

; mapping required for network applications

; word processing programIF MEMBER OF "WP42" THENMAP S3: = SYS:PUBLIC\WPEND

IF MEMBER OF "WP50" THENMAP S3: = SYS:PUBLIC\WP50END

; accounting programIF MEMBER OF "LOTUS_VI" THEN      MAP INS S16: =
SYS:PUBLIC\NET123      MAP INS S16: =      VOL1:USERS\%LOGIN_NAME\123      MAP
```

```
O: = SYS:PUBLIC\NET123      MAP L: = VOL!:USERS\%LOGIN_NAME\123END

; map miscellaneous search drivesIF MEMBER OF "POWER_USERS" THEN      MAP INS
S16: = SYS:PUBLIC\BATCH      MAP INS S16: = SYS:PUBLIC\UTILEND

; supervisor mappingsIF "%LOGIN_NAME" = "supervisor" THEN      MAP P: =
SYS:SYSTEM      MAP *1: = SYS:      MAP *2: = HOME:      MAP *3: = HOME2:
;etc.      DRIVE P:END

; display login messages as requiredFDISPLAY
SYS:PUBLIC\NEWS\message.txtPAUSEWRITE "any short message not in
message.txt"PAUSE; run miscellaneous programs#SYS:PUBLIC\lantrail

; display all current drive settingsMAP DISPLAY ONMAP
```

Long machine type

A six-letter name representing a DOS workstation brand. (This doesn't apply to OS/2 workstations.)

Use the long machine type in system login scripts (using the MACHINE identifier variable) to automatically map a drive to the correct version of DOS assigned to the station.

IBM computers use the long machine type IBM_PC. If the station is not an IBM computer, create a long machine type for the station in a NET.CFG file.

Use the six-letter name for the long machine type as the subdirectory name when you use more than one brand of workstation. Example: COMPAQ.

Use the same six-letter name for DOS directories that you use for the long machine type.

If you use more than one version of DOS, you must have a separate subdirectory for each DOS version used on each machine type.

See also “DOS version” on page 82.

LPT1

The primary parallel printer port of a personal computer.

See also “Parallel port” on page 176.

LSL

See “Link Support Layer” on page 126.

12 M

MAIL directory

A directory, SYS:MAIL, created during network installation and used by mail programs that are compatible with NetWare.

When network users are created, they are assigned a User ID number. Users are also assigned a subdirectory, or mailbox, in the MAIL directory. The User ID number is used as the mailbox name.

Users have all rights, except Supervisory, in their mailboxes, but are have only the Create right in the MAIL directory.

Each user's login script is stored in his or her mailbox, which allows the login script to be accessed each time the user logs in.

Major resource

A category of data defined by the Target Service Agent and recognized by SBACKUP.

A major resource contains data that can be backed up as a group, for example, server, volume, etc.

See also “Backup” on page 25; “Minor resource” on page 144; “Target Service Agent” on page 259; “TSA resources” on page 268.

Mapping

See “Drive mapping” on page 83.

Media Manager

Routines within NetWare v3.12 and v4.0 that provide applications with a generic view of the different types of backup storage devices (disk, tape, autochanger, etc.).

These routines enable applications to communicate with the different storage devices without having device-specific drivers.

Memory board

An add-on board designed to increase the amount of RAM within a personal computer.

Memory pool

A defined area of the file server's memory that can be allocated for use by server processes, or for recording the status of server resources.

NetWare v3.12 uses memory for many functions. A minimal system (80386 CPU with one 80MB hard disk) requires at least 2 MB of RAM, but 4 MB of RAM are recommended for NetWare v3.12. NetWare can address up to 4,096 MB of RAM.

To initialize the NetWare operating system, the file server must be booted by loading DOS into low memory. Then the NetWare operating system is loaded into high memory (memory above 1 MB).

The memory that is not used for the NetWare operating system or for DOS is given to three main memory pools:

- File cache buffers
- Permanent
- Alloc memory

NetWare allocates as little memory as possible to the permanent and alloc pools, allowing these pools to grow dynamically on demand.

Figure 36 illustrates the various pools and their subdivisions.

Figure 36 Memory pools

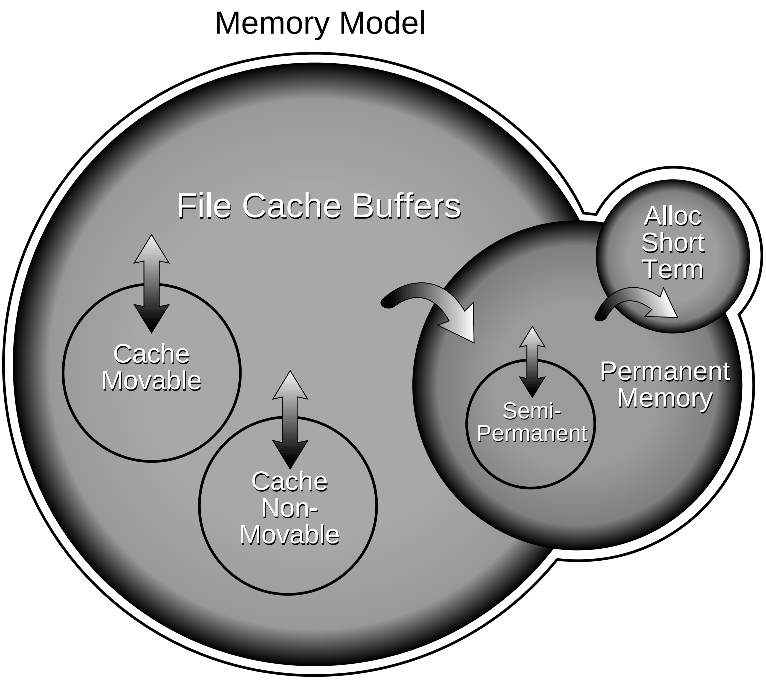

File Cache Buffer Pool

The file cache buffer pool stores the most frequently used files.

It is the pool from which all other pools obtain additional memory. Two pools, movable and nonmovable, obtain and return memory from the file cache buffer pool.

- The movable pool is used for system tables that change size—for example, FATs, and hash tables.
- The nonmovable pool is used for loadable modules, which it stores in memory.

As Figure 36 indicates, both the movable and nonmovable pools return memory directly to the cache buffer pool when the memory is no longer needed.

Permanent Pool

The permanent pool is used for long-term memory needs, such as directory cache buffers and packet receive buffers.

Alloc Memory Pool

The alloc memory pool stores the following information:

- Drive mappings
- Service request buffers
- Open and locked files
- Server advertising
- User connection information
- Messages waiting to be broadcast
- Loadable module tables
- Queue manager tables

Message packet

A unit of information used in network communication.

How Messages Are Sent

Messages sent between network devices (workstations, file servers, etc.) are formed into packets at the source device. The packets are reassembled, if necessary, into complete messages when they reach their destination.

A message packet might contain a request for service, information on how to handle the request, and the data that will be serviced.

An individual packet consists of headers and a data portion. The different headers are appended to the data portion as the packet travels through the communication layers.

Any message that exceeds the maximum size is partitioned and carried as several packets. When the packet arrives at its destination, the headers are stripped off in reverse order and the request is serviced.

Example

For example, when a message is routed,

1. The NetWare Core Protocol™ (NCP™) attaches a write request header and an IPX header to a piece of data to be written.
2. The workstation's IPX communication protocol fills in the IPX header designating the source of the request and the packet length.
3. The device driver adds a hardware or MAC (Media Access Control) header.

Figure 37 shows a message packet traveling from a source device to a destination device.

Figure 37 Message packet routing

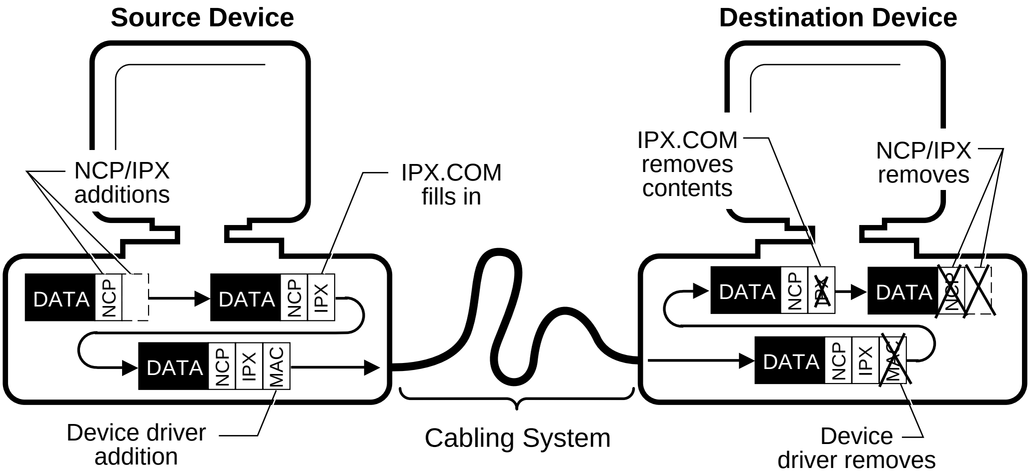

See also “Communication protocols” on page 41; “NetBIOS” on page 155.

Message system

A communications protocol that runs on top of IPX.

It provides an engine that allows a node on the network to send messages to other nodes. A set of Application Programmer Interfaces (APIs) gives programs access to the message system.

NetWare v3.12 supports file server broadcasts and alerts to Supervisors and Workgroup Managers.

Minor resource

A category of data defined by the Target Service Agent and recognized by SBACKUP.

A minor resource might be located in the directory structure below the selected major resource, for example, directories, subdirectories, or files.

See also "Backup" on page 25; "Major resource" on page 139; "Target Service Agent" on page 259; "TSA resources" on page 268.

Mirroring

See "Disk mirroring" on page 75.

Modify bit

A bit set by the operating system when a file is changed to indicate that data has been modified.

When a backup is performed, SBACKUP can check to see whether modify bits are set, and can back up only those files that have their modify bit set.

Different types of backups have different actions on the modify bit and on files that aren't marked with a modify bit (unmodified files).

Table 7 shows how different types of backups handle the modify bit and process the files that aren't marked with the modify bit.

Table 7 Types of backups and their modify bit settings

Type of backup	Modify bit setting	Treatment of unmodified files at time of next backup of this type
Full	Clear after backup	Include
Incremental	Clear after backup	Exclude
Differential	Don't clear after backup	Exclude
Custom	As desired	As desired

See also “Backup” on page 25.

Modify right

See “Rights” on page 207; “Security” on page 221 (Rights Security).

Multiple file server network

See “Multiserver network” on page 146.

Multiple name space support

The method that allows various workstations running different operating systems to create their own familiar naming conventions. In other words, the file system can present multiple workstation views for any given file.

Each file stored on a given volume has a DOS name that any DOS workstation can recognize. This DOS name is stored in a file entry in the volume's directory table.

Different operating systems (for example, DOS, OS/2, Macintosh, UNIX) may have different conventions for naming files, such as these:

- Name length
- Legal characters

- Case-sensitivity or insensitivity
- Data and resource forks
- Length of extensions
- Multiple extensions

For example, a file server configured to support DOS and Macintosh file names would generate two 128-byte file entries for every file.

Volumes that support multiple name spaces use one file/directory entry for each name space supported. The same applies to directory names.

See also “Name space” on page 147.

Related utilities: "ADD NAME SPACE"; "LOAD" (System Administration).

Multiple-byte characters

Single characters made up of more than one byte.

One byte allows 256 different characters. Since the number of ASCII characters equals 256, a computer can handle each ASCII character with one byte.

Asian character sets, however, include more than 256 characters; in this case, a computer must use two bytes for each character.

Although the VLMs and the NetWare 3.12 operating system support multiple-byte characters, the NetWare 3.12 utilities do not incorporate this feature.

Multiserver network

A single network that has two or more file servers operating.

On a multiserver network, users can access files from any file server to which they are attached (if they have access rights).

A multiserver network is not the same as an internetwork, where two or more networks are linked through a router.

See also “Network numbering” on page 165.

13 N

Name space

A loadable module that allows you to store non-DOS files on a NetWare v3.12 file server.

DOS name space is always provided by the operating system.

How Name Space Works

Any file types other than DOS, such as Macintosh or OS/2, must have a name space loadable module linked with the operating system before the file server can store such files.

These modules have an .NAM extension. Once the module is loaded, you must use ADD NAME SPACE to configure the volumes so that you can store the other types of files.

When the name space support is added to the volume, the volume creates another entry in the directory table for the directory- and file-naming conventions of that file system.

For example, a volume that supports Macintosh files has the following for each Macintosh file:

- A DOS filename in the DOS directory space
- A Macintosh filename in the Macintosh directory space

See Figure 38.

Figure 38 Name space support for DOS and Macintosh

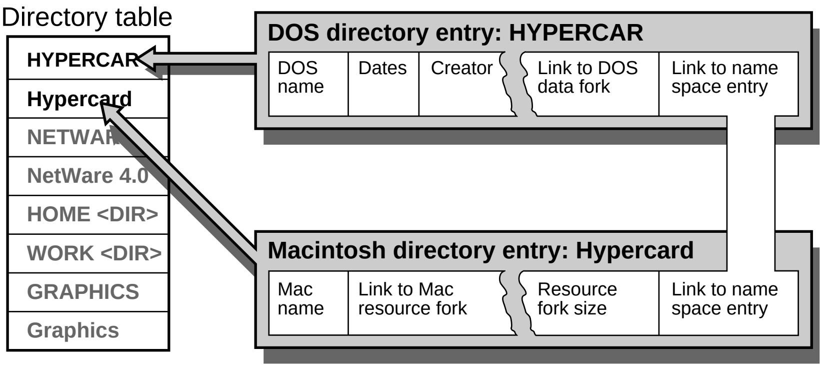

Using Name Space Support

Consider the following when you use name space support:

- Each volume that is assigned to support an additional name space type requires about twice as much memory as a volume without additional name space types.
- Once an additional name space has been added to a volume, it cannot be removed.
- A volume that is assigned to support an additional name space type cannot be mounted unless the loadable module that adds the name space to the operating system has been loaded.

  This LOAD command must be issued before the MOUNT command.
- You should store name space modules with the file server boot files so that you can add name space to any volume.
- You should add the command to load name space support to the STARTUP.NCF file.

See also "NetWare DOS Requester" on page 156.

Related utilities: "ADD NAME SPACE"; "INSTALL"; "LOAD name space" (*System Administration*).

If you are using ODI LAN drivers, refer to Workstation for DOS and Windows.

NCP

See “NetWare Core Protocol” on page 156.

NCP packet signature

An enhanced security feature that protects servers and workstations using NCP (NetWare Core Protocol) by preventing packet forgery.

Without NCP packet signature installed, a workstation can pose as a more privileged workstation to send a forged NCP request to a file server.

By forging the proper NCP request packet, an intruder could gain SUPERVISOR rights and access to all network resources.

NCP packet signature prevents packet forgery by requiring the server and the workstation to "sign" each NCP packet. The packet signature changes with every packet.

NCP packets with incorrect signatures are discarded without breaking the workstation's connection to the server.

However, an alert message about the invalid packet is sent to the error log, the affected workstation, and the server console. The alert message contains the login name and the station address of the affected workstation.

If NCP packet signature is installed on the server and all its workstations, it is virtually impossible to forge a valid NCP packet.

Packet Signature Options

Because the packet signature process consumes CPU resources and slows performance, both for the workstation and the file server, NCP packet signature is optional.

Several signature options are available, ranging from never signing NCP packets to always signing NCP packets. Four signature levels are available for servers, and four signature levels are available for workstations.

The options for servers and workstations combine to determine the level of NCP packet signature on the network.

NOTE: Some combinations of server and workstation packet signature levels may slow performance. However, low CPU-demand systems may not show any performance degradation.

Network supervisors can choose the packet signature level that meets both their performance needs and their security requirements.

Server Levels

Server packet signature levels are assigned by a new SET parameter:

```
SET NCP PACKET SIGNATURE OPTION = number
```

Replace *number* with 0, 1, 2, or 3. The default is 1.

Number	Explanation
0	Server does not sign packets (regardless of the workstation level)
1	Server signs packets *only* if the workstation requests it (workstation level is 2 or higher)
2	Server signs packets if the workstation is capable of signing (workstation level is 1 or higher)
3	Server signs packets and requires all workstations to sign packets (or logging in will fail)

Workstation Levels

Workstation signature levels are assigned by a new NET.CFG parameter:

```
SIGNATURE LEVEL = number
```

Replace *number* with 0, 1, 2, or 3. The default is 1.

Number	Explanation
0	Workstation does not sign packets
1	Workstation signs packets *only* if the server requests it (server option is 2 or higher)
2	Workstation signs packets if the server is capable of signing (server option is 1 or higher)
3	Workstation signs packets and requires the server to sign packets (or logging in will fail)

Effective Packet Signature

The packet signature levels for the server and the workstation interact to create the "effective" packet signature. Some combinations of server and workstation levels do not allow logging in.

Table 8 shows the interactive relationship between the server packet signature levels and the workstation signature levels.

Table 8 Effective packet signature of server and workstation

IF	Server = 0	Server = 1	Server = 2	Server = 3
Workstation = 0	No packet signature	No packet signature	No packet signature	*No logging in*
Workstation = 1	No packet signature	No packet signature	Packet signature	Packet signature
Workstation = 2	No packet signature	Packet signature	Packet signature	Packet signature
Workstation = 3	*No logging in*	Packet signature	Packet signature	Packet signature

When to Use NCP Packet Signature

NCP packet signature is not required for every installation. Some network supervisors may choose not to use NCP packet signature because they can tolerate certain security risks.

Security Risks

The following situations are examples of tolerable risks that may not need NCP packet signature:

- Only executable programs reside on the server.
- All workstation users on the network are known and trusted by the supervisor.
- Data on the file server is not sensitive; loss or corruption of this data will not affect operations.

NCP packet signature is recommended for security risks such as these:

- An untrustworthy user at a workstation on the network.
- Easy physical access to the network cabling system.
- An unattended, publicly accessible workstation.

Signature Level Examples

The default NCP packet signature level is 1 for workstations and servers. In most installations, this setting provides the most flexibility while still offering protection from forged packets. Below are some examples of using different signature levels.

All Information on the Server Is Sensitive

If an intruder gained access to any information on the file server, it could damage the company.

The network supervisor sets the server to level 3 and all workstations to level 3 for maximum protection.

Sensitive and Nonsensitive Information Reside on the Same Server

The file server has a directory for executable programs and a separate directory for corporate finances (such as accounts receivable).

The network supervisor sets the server to level 2, and the workstations that need access to accounts receivable to level 3. All other workstations remain at the default, level 1.

Users Often Change Locations and Workstations

The network supervisor is uncertain which employees will be using which workstations, and the file server contains some sensitive data.

The network supervisor sets the server to level 3. Workstations remain at the default, level 1.

Workstation is Publicly Accessible

An unattended workstation is set up for public access to nonsensitive information, but another server on the network contains sensitive information.

The network supervisor sets the sensitive server to level 3 and the unattended workstation to level 0.

Packet Signature Considerations for Job Servers

Network supervisors should be aware that some job servers do not support NCP packet signature. A job server may produce unsigned sessions if

- It does not operate on top of DOS;
- It does not use standard NetWare shells;
- It is not an NLM;
- It uses its own implementation of the NCP engine (such as embedded print servers in printers).

Minimizing Risks

To minimize security risks associated with job servers, do the following:

- Install queues only on servers with signature level 3.
- Do not allow privileged users to put jobs in queues on servers with signature levels below 3.
- Make sure the job server's account is unprivileged.
- Disable the job server's ability to change to workstation rights.

Disabling Change to Workstation Rights

To prevent a job server from assuming the rights of a workstation, put the following new SET command in the server's AUTOEXEC.NCF file:

```
SET ALLOW CHANGE TO CLIENT RIGHTS = OFF
```

NOTE: *Client* is an alternate term for *workstation*.

The default is ON, because certain job servers and third-party applications cannot function without changing to workstation rights.

Troubleshooting Tips

This section describes some solutions to problems that may be associated with using NCP packet signatures.

Workstations Are Not Signing Packets

Make sure the old *.COM shells are renamed or removed from the directory where the new *.EXE shells reside.

Workstations Cannot Log In

Make sure the packet signature levels on the server and the workstation are correct.

The following situations do not allow logging in:

- Server packet signature = 3, workstation signature = 0
- Server packet signature = 0, workstation signature = 3
- Utilities are old and do not support packet signature
- Shells or requesters are old and do not support packet signature

Third-party NLMs Do Not Work

If the SET parameter "Allow Change to Client Rights" is set to OFF, some third-party NLMs may not function. Set this parameter to ON.

Unsecure Workstations Log In to Secure Server

The workstations are using an old LOGIN.EXE file that does not include NCP packet signature.

Make sure the LOGIN.EXE file is dated 10-16-92.

Add the "Preferred Server" option to the NET.CFG file for all workstations that have access to secure servers (level 3).

Enabling Packet Burst (optional)

The Packet Burst loadable module, PBURST.NLM, must be loaded on NetWare v3.11 servers in order for NCP packet signature to work. However, using the Packet Burst protocol to transfer data between servers and workstations is optional.

Packet burst is a protocol built on top of IPX that speeds the transfer of multiple-packet NCP reads and writes. The Packet Burst protocol eliminates the need to sequence and acknowledge each packet. With Packet Burst, the server or workstation sends a whole set (or burst) of packets before it requires an acknowledgment.

By allowing multiple packets to be acknowledged, the Packet Burst protocol reduces network traffic. The Packet Burst protocol also monitors dropped packets and retransmits only the missing packets.

The file server requires the PBURST.NLM to be loaded in order to transfer data in packet bursts. For a workstation to send and receive packet burst data, it requires the BNETX.EXE file and a new parameter in its NET.CFG file.

NOTE: The Packet Burst protocol is not supported by expanded memory or extended memory workstation shells, or by OS/2 workstations.

NetBIOS

An emulator program provided with NetWare that allows workstations to run applications that support IBM's NetBIOS calls.

NetWare routers

See “Router” on page 211.

NetWare Core Protocol

(NCP) Procedures that a server's NetWare operating system follows to accept and respond to workstation requests.

The process of requesting service from a file server begins in the workstation's RAM where the NetWare DOS Requester or NetWare Requester for OS/2 forms requests according to the definitions of the server's NCP.

The Requester then hands the requests to the station's IPX communication protocol. IPX transmits the request to the server after attaching a header designating the source and destination.

Upon receiving the request, the server removes the IPX header and reads the request.

Because the NetWare Requester formed the request using the exact guidelines of a specific service protocol, the server handles the request according to the protocol rules, resulting in a proper response.

NetWare Core Protocols exist for every service a station might request from a server.

Common requests handled by NCP include creating or destroying a service connection, manipulating directories and files, opening semaphores, and printing.

See also “Communication protocols” on page 41; “IPX” on page 121.

NetWare DOS Requester

The DOS workstation software portion of NetWare v3.12.

The NetWare DOS Requester™ replaces the NetWare shell under v3.12, while maintaining backward compatibility with previous NetWare shell versions.

The NetWare DOS Requester can be called by applications or utilities in one of three ways:

- Before DOS, in the same way as the old NETX shell.
- By DOS, through the INT 2Fh redirector.
- Bypassing DOS, through a pipeline between the shell and post-DOS portions.

VLMs

The NetWare DOS Requester is composed of a number of modules called *Virtual Loadable Modules (VLMs)*. These VLMs are the key to the NetWare DOS Requester's modularity. See "Virtual Loadable Module" on page 292.

The Requester contains categories of services and platforms in the following three layers:

- DOS Redirection Layer
- Service Protocol Layer
- Transport Protocol Layer

Figure 39 shows how these layers and modules fit together.

Figure 39 NetWare DOS Requester layers and modules

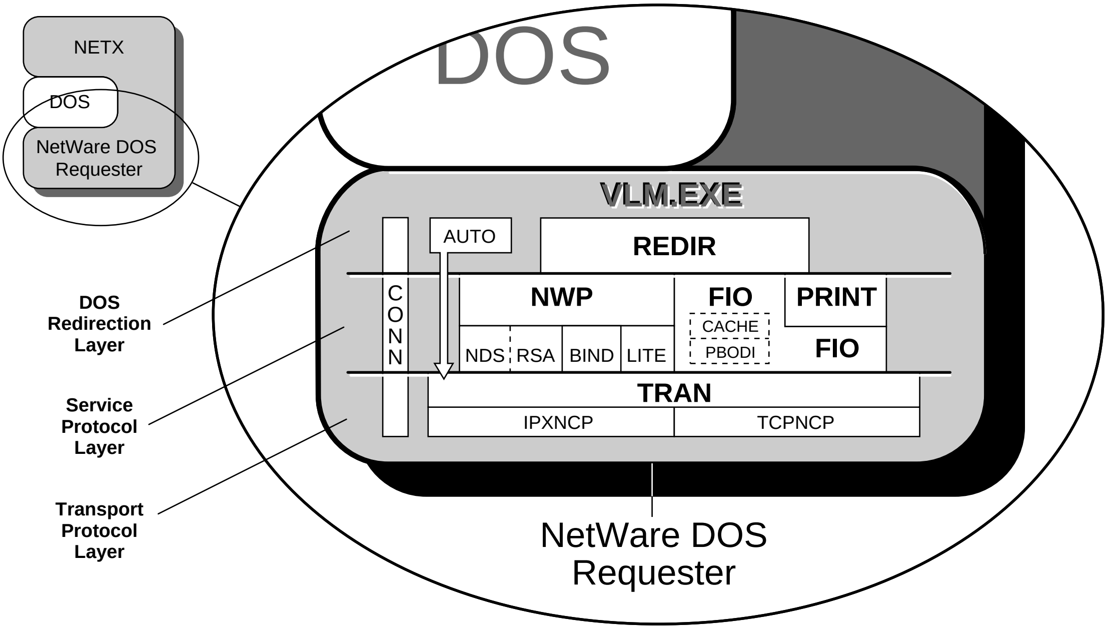

DOS Redirection

The NetWare DOS Requester includes a redirector that, in contrast to the NetWare shell, is called by DOS.

Under the redirector, DOS makes specific requests for services from the redirector (such as for file and print services from the server) that DOS can't provide. (These function were previously performed by the NETX shell, without involving DOS.)

The NetWare DOS Requester also continues to provide network services for file and print redirection, as well as for connection maintenance and other NetWare-specific support.

NetWare Loadable Module

(NLM™) A program you can load and unload from server memory while the server is running.

How NLMs work

NLMs link disk drivers, LAN drivers, name space, and other file server management and enhancement utilities to the operating system.

The file server allocates a portion of memory to the module when it is loaded. The module uses the memory to perform a task and then returns control of the memory back to the operating system when the module is unloaded.

The amount of memory the module uses during run time can vary, depending on the task. Some tasks make calls that cause the operating system to allocate more memory.

When the task is completed, the module returns some of the memory to the operating system. When a module is unloaded, all allocated resources are returned to the operating system.

Types of NLMs

NetWare v3.12 has four types of loadable modules:

- *Disk drivers* control communication between the operating system and the hard disks.

  These loadable modules have a .DSK extension. You can load and unload drivers while the file server is running and users are logged in.

- *LAN drivers* control communication between the operating system and the network boards.

  These loadable modules have an .LAN extension.

  You can load and unload drivers while the file server is running and users are logged in.

- *Management utilities and server applications modules* allow you to monitor and change configuration options.

  These loadable modules have an .NLM extension. You can run VREPAIR on a dismounted volume, add disk space to a mounted volume, or surface test a disk drive while the file server is running.

  Once you are finished with your tasks, you can unload the utility and free the memory for other file server functions.

- *Name space modules* allow non-DOS naming conventions to be stored in the directory- and file-naming system.

  These loadable modules have an .NAM extension.

Some modules, such as utilities, can be loaded, used, and then unloaded. Other modules, such as LAN driver modules and disk driver modules, need to be loaded every time the file server is booted.

NCF files allow you to store loadable module commands that you want loaded every time the file server is booted.

Location of NLMs

The loadable modules released with NetWare v3.12 (except BTRIEVE®) are copied automatically to the SYS:SYSTEM directory during installation.

As you acquire additional modules, you must decide where you want to copy them. The operating system must be able to find the modules when a LOAD command is issued.

The modules can be copied to any of the following areas.

- SYS:SYSTEM directory.
- Any network directory on the file server.

  If you copy modules to a directory other than SYS:SYSTEM, you need to either include the complete directory path, starting with volume name, in the command or use SEARCH to establish directories other than SYS:SYSTEM in which the operating system should look for loadable modules.
- A DOS drive of the file server.

  You can store the modules on a diskette in the floppy disk drive or on a DOS partition on the file server's hard drive (if you created a DOS partition during installation).

  Make sure you include the drive letter for the DOS drive in the command. For example, to load the INSTALL NLM from a diskette in drive A:, you would type

  ```
  LOAD A:INSTALL <Enter>
  ```

See also “Loading and unloading” on page 126.

Related utilities: "LOAD"; "LOAD disk driver"; "LOAD LAN driver"; "LOAD name space"; "LOAD NLM utility"; "UNLOAD" (System Administration).

NetWare operating system

The operating system developed by Novell, Inc.

The NetWare operating system runs in the file server and controls system resources and information processing on the entire network or internetwork.

NetWare partition

See “Partitions” on page 177.

NetWare Runtime

A single-user version of the NetWare v3.12 operating system that provides NetWare services to clients of NetWare Loadable Module (NLM) applications.

Benefits of a Runtime Server

NetWare Runtime™ is a network server platform supporting front-end or back-end applications as well as basic NLM services such as communication services, database servers, electronic mail, and other third-party applications.

NLM application developers have the flexibility to determine which client services will be available in their product.

NLMs, loaded on a NetWare Runtime server, provide client connection services (using IPX, SPX, AppleTalk, or TCP/IP).

How Runtime Works

Network clients attach to the NLM application that runs on top of NetWare Runtime. The NLM provides all services required by the client using NetWare application program interfaces (APIs).

Specific NetWare services are available to clients only through the NLM applications that serve those clients. The NLM application implements and manages any required client services.

Figure 40 shows how NetWare Runtime functions.

Figure 40 NetWare Runtime

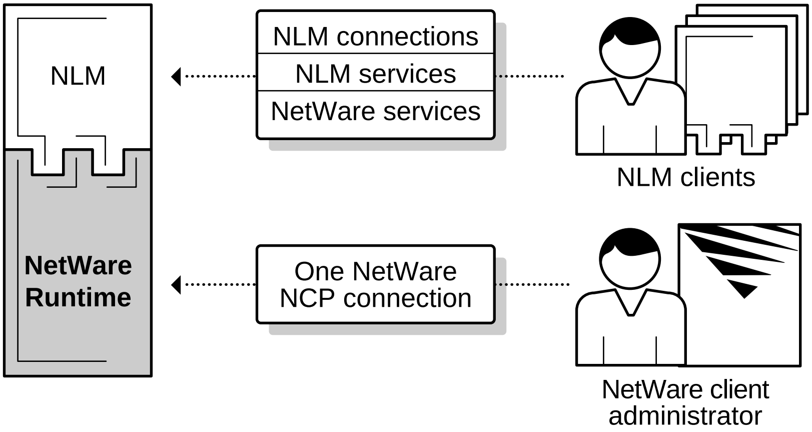

For example, a database NLM application may provide a client with login connections (including authentication), file services, data access, and disconnect services.

Other NetWare services can be built into NLM applications to provide additional functionality.

Limitations of a Runtime Server

Only one user can log in to the NetWare Runtime server at a time. Normally, the network supervisor logs in for administrative purposes using the single allowed NetWare Core Protocol (NCP) connection.

However, this does not mean that only one person can use the application because users access the application through an NLM connection and not the single NCP connection.

Utilities that Don't Apply to a Runtime Server

Though all NetWare v3.12 utilities and commands are available on a Runtime server, some don't apply or wouldn't be used because of the server's single-user environment.

The following utilities aren't likely to be used:

CAPTURE
DCONFIG
NMENU
NPRINT
PCONSOLE
PRINTCON
PRINTDEF
PSC
SEND
SETTTS
SYSTIME
WSUPDATE

In addition to utilities, there are several network supervisor commands you don't need to use, such as commands associated with creating or modifying login scripts or menus, as well as those related to workstation management.

NetWare Runtime Installation

The procedures for installing a NetWare Runtime v3.12 server are the same as for any NetWare v3.12 server.

For information on how to install or upgrade a Runtime server, see Installation and Upgrade.

NetWire

Novell's online information service.

NetWire provides access to Novell product information, Novell services information, and time-sensitive technical information for all NetWare users.

NetWire allows you to

- Remotely access information 24 hours a day;
- Submit questions to a Novell Technician or System Operator;
- Download files and technical information dealing with product updates and modifications.

NetWire is accessed through the CompuServe® Information Service. It requires a PC or compatible workstation, a modem, and a communications program.

For more information on NetWire subscriptions, contact CompuServe at 1 (800) 848-8199 (in the USA); ask for representative number 58 to receive a free introductory membership to NetWire. Outside of the United States or Canada, call 1 (614) 457-0802.

Network

A group of computers that can communicate with each other, share peripherals (such as hard disks and printers), and access remote hosts or other networks.

A NetWare network consists of workstations, peripherals, and one or more file servers.

NetWare network users can share the same files (both data and program files), send messages directly between individual workstations, and protect files with an extensive security system.

Network address

See “Network numbering” on page 165.

Network board

A circuit board installed in each network station to allow stations to communicate with each other and with the file server.

Network communication

Data transmission between network stations.

Requests for services and data pass from one network station to another through a communication medium such as cabling.

Network hard disk

See "Disk" on page 70.

Network numbering

The system of numbers that identifies servers, network boards, and cable segments.

Network Addresses

A network number (address) is an eight-digit hexadecimal number that uniquely identifies a network.

A network address identifies a network cabling scheme, and node numbers identify individual stations along the network cable. Each packet on a network is stamped with a source and a destination address, which consist of a network address and a node number.

A network, therefore, is a single cabling scheme identified by a unique address to which one or more stations are attached, each of which are identified by numbers that are unique along the network.

The figures on the following pages use simple network configurations to illustrate the concepts of network numbering.

The network numbers used are arbitrary; any combination of hexadecimal numbers in the range of 1 to FFFFFFFE are valid network addresses. The illustrations use a bus topology; however, the concepts apply to any network topology.

Single Server Network

Figure 41 shows a simple network configuration—one server attached to one network. The network number AAA1 identifies the network to which the server ADMIN is attached.

Figure 41 Single server network

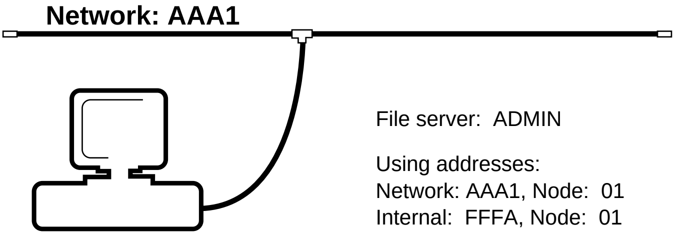

Internal network. The server ADMIN uses the network number FFFA to identify the internal network. An internal network is a "logical" network that routes packets to the physical networks to which a file server is attached.

The node number 1 distinguishes the server ADMIN from any other station on the network AAA1. The operating system automatically assigns the node number of 1 to the internal network.

Multiserver Network

Figure 42 illustrates the first configuration slightly expanded by the addition of a second server, SALES, to the same network.

Figure 42 Multiserver network

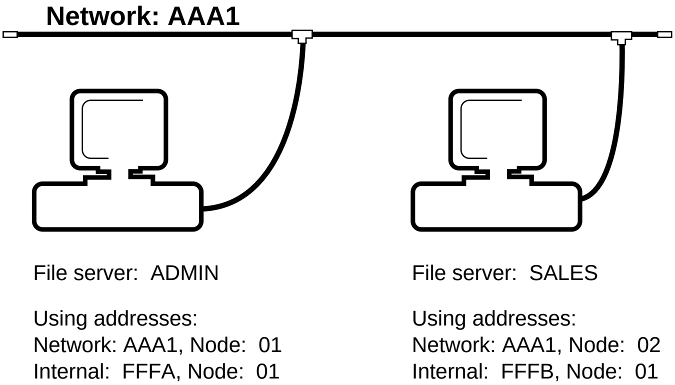

SALES and ADMIN share the same physical network number AAA1. However, because node numbers must be unique, each server uses a different node number on network AAA1.

Because both servers see each internal network as a separate network, the internal networks must use unique network numbers. In this configuration, the server SALES uses the internal network number FFFB.

Multiserver Internetwork

Figure 43 expands the configuration with a second, multiserver network.

Figure 43 Multiserver internetwork with internal router

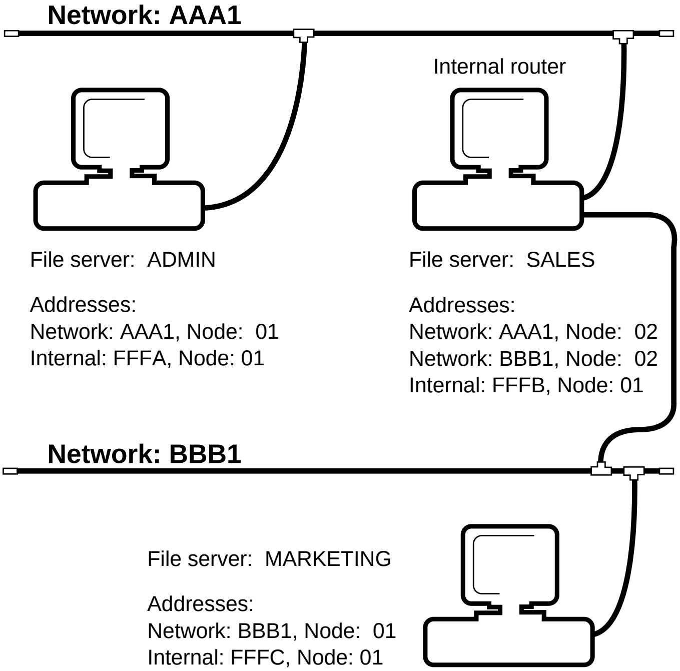

On an internetwork, stations on one network can communicate with stations on other networks.

In this internetwork, the server SALES performs the function of an *internal router*. This means that server SALES routes data frames between networks AAA1 and BBB1.

The configuration includes a third server, MARKETING. Because server SALES is attached to both networks AAA1 and BBB1, all three servers can communicate with each other.

Figure 44 illustrates an internetwork that uses an external router between the networks AAA1 and BBB1.

Figure 44 Multiserver internetwork with external router

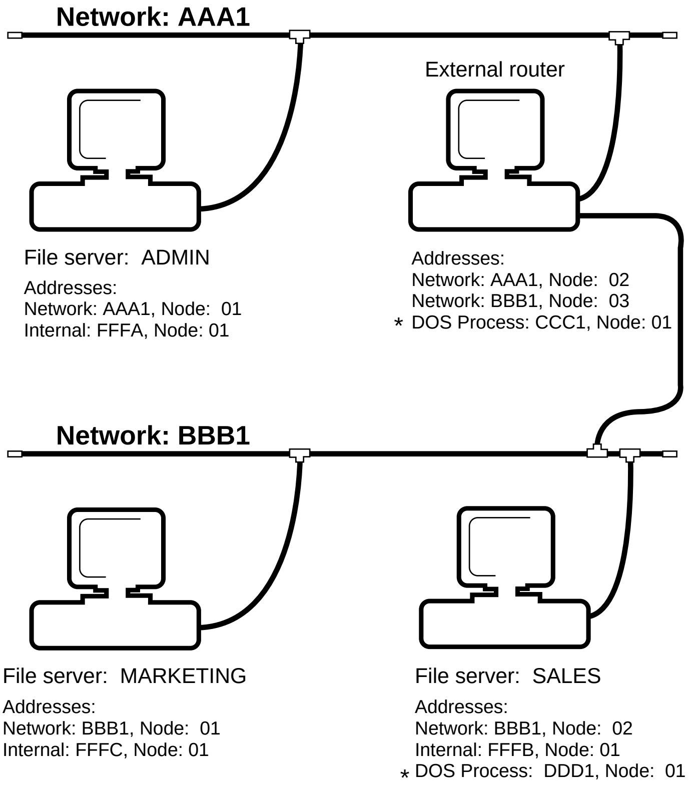

* Used with nondedicated file servers and routers

The external router in this figure performs the same functions as server SALES in Figure 43 on page 168. However, routing data packets between networks is the main function of the external routers.

Therefore, server SALES in Figure 44 provides no routing services, because the external router routes the data packets between the networks AAA1 and BBB1.

The external router in Figure 44 on page 169 is configured as "nondedicated," so it can also function as a workstation. The router's operating system uses a nondedicated DOS process.

A nondedicated DOS process is a special-purpose router used in the operating systems of nondedicated routers and file servers.

Like the internal network of a file server, a nondedicated DOS process is a "logical network." However, this logical network has only one station attached: the workstation function of the router.

The router operating system treats this workstation function of the router as a "logical workstation."

This nondedicated DOS process routes data packets from the logical workstation of the router to the router's operating system.

Because a nondedicated DOS process is a logical network, it must be assigned a network number that is unique to the number of the network to which the router is attached.

This number must also be unique on the internetwork. In a nondedicated file server or router, the logical workstation is assigned node number 1 and the router station number is assigned node number 2.

The configuration in Figure 44 on page 169 also includes server MARKETING. Because server MARKETING is attached to the network BBB1, the router of server MARKETING must use the network number BBB1 and a node number unique to network BBB1.

Network operator

See “User” on page 274.

Network station

Any personal computer connected to a network by means of a network board and a communication medium.

A network station can be either a workstation or a router.

Network supervisor

See “SUPERVISOR” on page 255; “User” on page 274.

NLM

See “NetWare Loadable Module” on page 158.

Node address

See “Network numbering” on page 165.

Node number

A number that identifies a network board (in a server, workstation, or router).

Every station on a network has a unique node number to distinguish it from other stations.

See “Network numbering” on page 165.

14 O

Object

An entity that is defined on the network and thus given access to the file server.

Object types are defined in the file server's bindery, and include users, groups, file servers, print servers, and archive servers.

See also “Bindery” on page 30.

ODI

See “Open Data-Link Interface” on page 173.

Open Data-Link Interface

(ODI) An architecture that allows multiple LAN drivers and protocols to coexist on network systems.

ODI supports media- and protocol-independent communications by providing a standard interface that allows transport protocols to share a single network board without conflict.

See also “Dedicated IPX drivers” on page 53.

Operating system

See “NetWare operating system” on page 160.

15 P

Packet

See “Message packet” on page 142.

Packet receive buffer

An area in the file server's memory set aside to temporarily hold data packets arriving from the various workstations. (The number of packet receive buffers is set during server installation.)

The packets remain in this buffer until the server is ready to process them and send them to their destination. This ensures the smooth flow of data into the server, even during times of particularly heavy input/output operations.

The operating system increases the number of buffers as needed for heavy buffer activity, within the following parameters (also set during server installation):

- *Maximum packet receive buffers*. The *maximum* packet receive buffers that the operating system can use.
- *Minimum packet receive buffers*. The *minimum* packet receive buffers that the operating system can use.
- *New packet receive buffer wait time*. The amount of time the operating system waits before allocating a new buffer in response to a packet request.

  Waiting ensures that new packet receive buffers aren't allocated needlessly during sporadic peak activity.

> If the maximum number of packet receive buffers is reached and a waiting packet isn't processed within the specified wait time, the operating system discards the packet and the station must resend it.

The default range of packet receive buffers should be satisfactory for most server installations, even with many users performing many read/write operations.

Because more system overhead is required to manage large numbers of packet receive buffers, we recommend that you increase the minimum/maximum range *only* if you are running out of buffers.

Packet signature

See “NCP packet signature” on page 149.

Parallel port

A printer interface that allows data to be transmitted a byte at a time, all eight bits moving in parallel.

See also “LPT1” on page 137.

Parent directory

The directory immediately above any subdirectory.

For example, SYS:ACCTS is the parent directory of the subdirectory SYS:ACCTS/RECEIVE.

See also “Directory structure” on page 61.

Parity

See “Serial communication” on page 246.

Partitions

Logical units into which hard disks can be divided.

One of the file server's internal hard disks can contain both an active, primary DOS partition and a NetWare partition.

When the file server boot files are copied to the DOS partition and included in the AUTOEXEC.BAT file, the NetWare operating system boots automatically. You need only one DOS partition; the other hard disks need to contain only a NetWare partition.

Physically, a NetWare partition consists of a Hot Fix redirection area plus a large data area. The term *NetWare partition* refers only to the data area. Consequently, the logical sector 0 of a NetWare partition is the first sector of the data area. See Figure 45.

Figure 45 NetWare partition

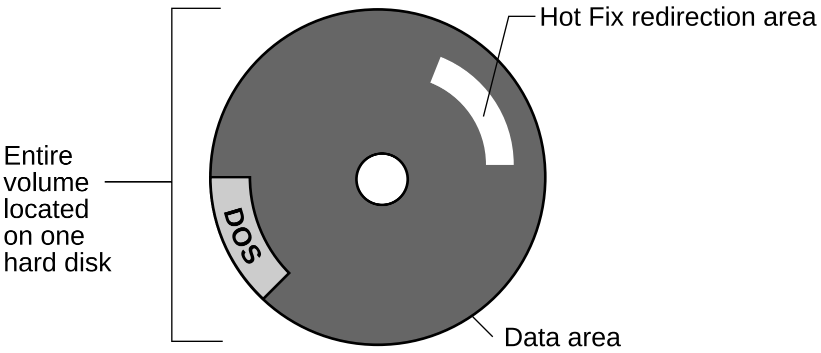

A data area contains four copies of the volume definition table. Each table contains a list of all the volume segments in that NetWare partition.

Four copies are maintained for fault tolerance. If a disk error occurs and one of the tables is corrupted, the error can be detected and corrected.

The rest of the data area can contain one to eight volume segments. Each segment can belong to a different volume. See Figure 46.

Figure 46 NetWare partition

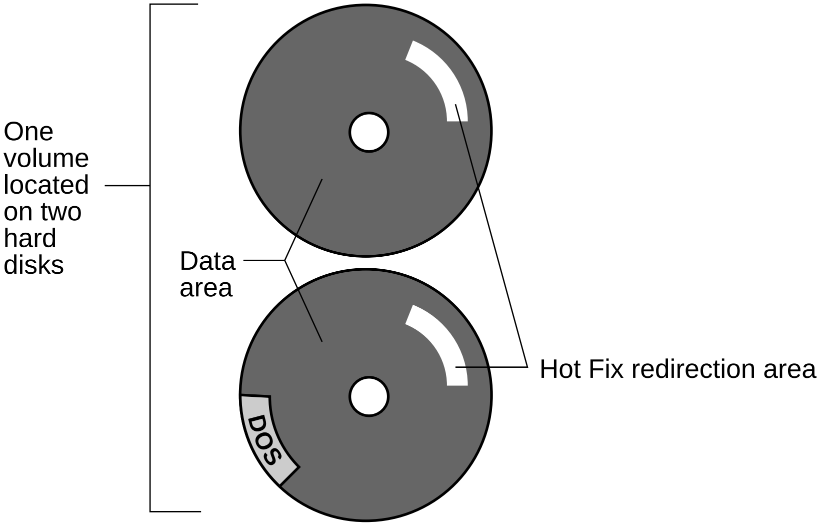

See also "Data protection" on page 49; "Hot Fix" on page 117; "VLM" on page 292.

Related utility: "INSTALL" (*System Administration*).

Passive hub

A device used in certain network topologies to split a transmission signal, allowing additional workstations to be added.

A passive hub cannot amplify the signal, so it must be connected directly to a workstation or to an active hub.

See also "Active hub" on page 19.

Password

The characters a user must type to log in (if a password is required).

NetWare allows you to assign passwords to users on the network.

In NetWare v3.x, login passwords are encrypted at the workstation and put into a format that only the file server can decode. This format prevents intruders from accessing your files.

Passwords in NetWare versions below 2.15 Rev. C are sent across the network in a "clear text format." A packet analyzer can pull in a login packet that uses the clear text format, and an intruder can read a user's password.

See also “Security” on page 221 (Login Security); “User” on page 274; “User account” on page 275.

Path

A variable that appears in command formats.

Path represents a NetWare directory path that includes the file server, volume, directory, and (if necessary) subdirectories and file you need in your command.

When you type commands, replace *path* with a drive letter or the complete directory path.

See also “Directory structure” on page 61.

Port, hardware

A connecting component that allows a microprocessor to communicate with a compatible peripheral.

See also “Parallel port” on page 176; “Serial port” on page 248.

Port, software

A memory address that identifies the physical circuit used to transfer information between a microprocessor and a peripheral.

Power conditioning

Methods of protecting sensitive network hardware components against power disturbances.

Power disturbances can be categorized in several ways:

- A *transient* (sometimes called a *spike* or *surge*)—a very short, but extreme, burst of voltage.
- Noise or static—a smaller change in voltage.
- Blackouts and brownouts—the temporary drop in or loss of electrical power.

Protection against Power Disturbances

Three types of protection are available:

- *Suppression*—protects against transients. The most common suppression devices are surge protectors that usually include circuitry to prevent excess voltage.
- *Isolation*—protects against noise, using ferro-resonant isolation transformers to control voltage irregularities.
- *Regulation*—protects against brownouts and blackouts. The uninterruptible power supply (UPS) is the most commonly used form of regulation.

Proper use of power conditioning devices greatly reduces network maintenance costs.

Make sure the proper amperage is available for each system; dedicated power lines should provide ample amperage.

Also, make sure all outlets are grounded. Power conditioning devices connected to poorly grounded outlets offer very little protection.

See also “Uninterruptible power supply” on page 271.

Print device

A printer, plotter, or other peripheral used to produce hard copy.

See also “Printing” on page 183.

Print function

A printer command that determines the characteristics of a print job.

For example, a print function can specify the style of typeface.

See also “Printing” on page 183.

Print job configuration

A group of characteristics that determine how a job is printed.

The characteristics may include the mode, the form, the number of copies, and the particular printer used.

Users can create print job configurations using the PRINTCON utility.

See also “Print mode” on page 181.

Print mode

A sequence of print functions that determines the appearance of the printed output.

A print mode can define the style, size, boldness, and orientation of the typeface.

The SUPERVISOR uses PRINTDEF to designate print modes, allowing users to quickly select a combination of print functions.

See also “Printing” on page 183.

Print queue

See “Printing” on page 183.

Print queue operator

A printing supervisor with rights to create, manage, disable, and enable print queues.

A print queue operator can also authorize a print server to service a queue.

User SUPERVISOR is a print queue operator by default; however, SUPERVISOR can delegate this responsibility to another user with the PCONSOLE utility (*System Administration*).

See also "Print server operator" on page 182.

Print server

See "Printing" on page 183.

Print server operator

A printing supervisor with rights to do the following, with PCONSOLE:

- Attach to other file servers
- Specify who to notify if the printer needs service
- Issue commands to the printer
- Change forms for the printer
- Change the queues serviced by a print server
- Change queue priority
- Bring down the print server

A print server operator cannot create new print servers or assign other users as print server operators.

User SUPERVISOR is a print server operator by default; however, SUPERVISOR may assign another user as print server operator by using the PCONSOLE utility (*Print Server*).

To establish the proper connection between the print server and the print queue, a user must be both print server operator and print queue operator, or enlist the help of the print queue operator of each queue serviced by the print server.

See also “Printing” on page 183.

Related utility: "PCONSOLE" (*Print Server*).

Printing

The ability to transfer data from computer files to paper.

Network and Standalone Printing

To network users, printing from network workstations seems the same as printing from standalone workstations, but there are differences.

When a user at a standalone workstation sends a print job to a local printer, the job is sent directly to the printer for processing. See Figure 47.

Figure 47 Standalone printing environment

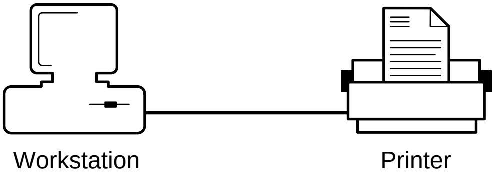

However, when a user on a network workstation sends a print job to a network printer, the job is routed first through the file server and then delivered to the printer by the print server.

When a print job leaves a network workstation, it is stored temporarily in a print queue on the file server before being sent to a printer.

Print queues, which are basically subdirectories, are created when the print server is installed. See Figure 48.

Figure 48 Network printing environment

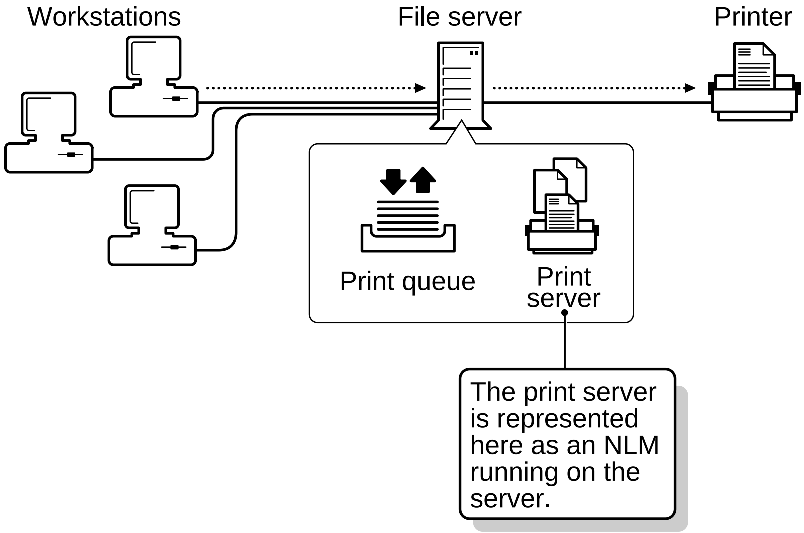
Workstations
File server
Printer
Print queue
Print server
The print server is represented here as an NLM running on the server.

NetWare Printing Services

The NetWare printing services control the network printing process at different stages.

When the print job is ready to be sent from the workstation to the queue, users can use NetWare printing utilities to control and monitor their print jobs.

Once the print job is in the queue, the NetWare print server takes the job and sends it to the printer that has been mapped to that particular queue.

NetWare printing services consist of the following:

- The NetWare print server
- The remote printer software
- The NetWare printing utilities

NetWare Print Server

The NetWare print server allows you to

- Increase the number of printers on your network;
- Locate printers where you need them in the workplace, not just next to the file server.

For example, you can print to a network printer connected to a dedicated workstation print server, or you can print to a remote network printer connected to a workstation attached to the network.

Each print server supports 16 printers and can service queues on up to eight file servers.

Types of print servers. There are three types of print servers available in NetWare: the loadable module (PSERVER.NLM); the dedicated workstation (PSERVER.EXE); and the VAP (PSERVER.VAP).

- *PSERVER.NLM*, a NetWare Loadable Module, is loaded directly onto the file server. It links to the operating system and sends print jobs from file server queues to the printers mapped to the queues.

  PSERVER.NLM can be loaded and unloaded while the file server is running. If you need to add a printer or change your printing setup, you unload the NLM, make the necessary changes in PCONSOLE, and then reload the NLM.

  PSERVER.NLM is copied to the SYS:SYSTEM directory during installation.

- *PSERVER.EXE* is run on a dedicated DOS workstation. Like the loadable module, it sends print jobs from file server queues to the printers mapped to the queues.

  A primary advantage to using PSERVER.EXE is that it reduces the load on the file server.

  If you need to add a printer or modify your printing setup, you must bring down the print server, make the changes in PCONSOLE, and reboot the dedicated workstation so that the changes take effect.

  PSERVER.EXE is copied into the SYS:PUBLIC directory during installation.

- *PSERVER.VAP* is run on a NetWare v2.1x or above file server or external router.

  Like other VAPs, the print server VAP runs "on top" of the NetWare 286 operating system. It ties into the operating system and sends print jobs from file server queues to printers mapped to the queues.

  PSERVER.VAP can be stopped and restarted at the file server or router console while the file server is running.

  If you need to add a printer or change your printing setup, you stop the VAP, make the necessary changes in PCONSOLE, and then restart the VAP on the file server or router.

  PSERVER.VAP is copied into the SYS:SYSTEM directory during installation.

Remote Printer Software

RPRINTER.EXE allows the print server to use a printer attached directly to your workstation as a remote network printer.

RPRINTER is a terminate-and-stay-resident (TSR) program that is loaded into the workstation's memory. It runs in the background until it is removed or until the workstation is rebooted.

If you want to use the printer as a standalone printer, you can remove the program from your workstation's memory or switch to Private mode with the PSC command.

Before you can connect a remote printer, you must use PCONSOLE to

- Define the remote printer in the print server account;
- Assign a queue to the remote printer.

The remote printer files are copied into SYS:PUBLIC during installation.

NetWare Printing Utilities

These utilities allow you to set up, control, and monitor the network print process.

You can use the utilities to

- Print from applications that are not designed for network printing;
- Print from disks;
- Print screen displays.

If you usually print from a network-compatible application, you will seldom need to use the printing utilities except to set up and monitor your print process.

“Printing tasks” on page 187 gives you an idea of what each command is used for. For a complete description of the utilities, see Print Server.

Table 9 Printing tasks

Task	Utility
Set up printing	PCONSOLE, SPOOL
Run software	PSERVER, RPRINTER
Print files	CAPTURE, ENDCAP, NPRINT, PCONSOLE
Control printer/print server	PSC, PCONSOLE
Customize printing	PRINTDEF, PRINTCON

Setting Up Network Printing

Before you can run the print server, you must run PCONSOLE to set up the following:

- Print queues
- Print servers
- Printers

Print Queues

Each print job must be sent to a print queue on the file server. The queue stores the job until the printer is ready. Then the print server takes the job from the queue and sends it to the printer assigned to that queue.

Each print job is stored as a file in the directory created for that queue.

When you create a print queue, the following happens:

- A corresponding directory is created in the SYS:SYSTEM directory.
- User SUPERVISOR is assigned as a print queue operator.
- Group EVERYONE is assigned as a print queue user.

If these default assignments are not sufficient, you can change them using PCONSOLE.

You must create print queues using PCONSOLE before you set up the print server account. We suggest you create one queue for each printer to simplify your printing setup.

Queue operators. The supervisor can assign other users to be queue operators as necessary.

Queue operators can

- Edit any other user's print queue entry request.
- Delete any entry from the queue.
- Modify the queue status by changing the operator flags.
- Change the order in which print jobs are serviced.

Queue users. To send print jobs to a queue, a user must be designated as a print queue user.

The group EVERYONE is assigned as a print queue user of each print queue when it is created. However, you can change this assignment and restrict the users for a print queue.

Queue servers. When you create your print server account, you must assign print queues to the printers you have defined. This authorizes the print server to take jobs out of the queue and move them to the appropriate printer.

The print server can service print queues from as many as eight file servers.

If you want the print server to service queues on another file server, you must create a print server account on each file server and attach to each file server.

You must either have SUPERVISOR privileges on that file server or enlist the help of the supervisor of that file server. For more information, see Print Server.

Print Servers

A print server takes print jobs out of a print queue and sends them to the appropriate network printer. After you create your print queues, you must create a print server account using PCONSOLE.

Print server configuration. You must create the print server configuration in PCONSOLE *before* running the print server. This configuration tells the print server

- The file servers to which the print server should attach;
- The printers and queues the print server supports;
- The queues that are serviced by each printer;
- The users that should be notified if the printer needs service.

Each time you load the print server, it uses this configuration to set up the printing environment.

We recommend that you assign the print server a password.

If you want the print server to service queues on multiple file servers, you must create a print server account on each file server.

When you install a print server, user SUPERVISOR is assigned as a print server operator, and group EVERYONE is assigned as a print server user. If these default assignments are not sufficient, you can change them in PCONSOLE.

Print server operators. These users have special privileges to control the print server. They can do the following using PCONSOLE or PSC:

- Attach the print server to other file servers
- Determine who should be notified if the printer needs service
- Issue commands to the printer
- Change forms for printer

- Change the queues serviced by a printer
- Change queue priority
- Bring down the print server

They cannot assign print server users or other print server operators.

Print server users. These users can monitor the status of the print server.

The group EVERYONE is assigned as a print server user. However, you can change this assignment and restrict the users for a print queue.

You do not need to be a print server user to have the print server send your print jobs to a printer. You do need to be a queue user to put your job in a queue.

Printers

As part of creating the print server configuration in PCONSOLE, you must define printers by specifying whether the printer is local or remote and which port (LPT or COM) the printer is attached to.

A print server can support up to 16 printers. You can usually attach up to five local network printers.

Local network printers are attached to one of the following:

- NetWare v3.12 file servers running PSERVER.NLM
- NetWare v2.1x or above file servers
- Routers running PSERVER.VAP, or dedicated workstation print servers running PSERVER.EXE

Remote network printers are attached to network workstations. They can be used as standalone printers or, if you run RPRINTER, they can be used by other network users.

Since you can attach 16 printers, the number of remote printers you can use depends on how many printers are attached directly to the print server.

Printing ports. Most workstations can support three parallel ports (LPT1, LPT2, LPT3) and two serial ports (COM1, COM2).

If you are using a serial printer, you must specify the following parameters when you define your printer configuration.

- Baud rate
- Word length (data bits)
- Stop bits
- Parity
- XON/XOFF protocol

Printer mappings. When you define printers in PCONSOLE, you also specify which queues are serviced by each printer. We suggest you map one queue to each printer.

One printer can service multiple queues. For example, you could set up a special high-priority queue that is serviced before any other queue. When you specify which queue a printer is serviced by, you are prompted for the priority (from 1 to 10) of that queue.

One queue can be serviced by multiple printers. For example, you could have three identical printers printing jobs from a single queue, enabling you to maintain a first-come, first-served approach to printing large numbers of print jobs.

Printer mappings are stored in the print server configuration files on the file server. Each time the print server comes up, the print server checks for these mappings.

Bringing Up the Print Server

Once the print queues, print servers, and printers are set up, you can bring up the print server.

PSERVER.NLM. You must load PSERVER.NLM on the v3.12 file server using the NetWare LOAD command.

PSERVER.EXE. You must modify the SPX connections on the dedicated workstation's NET.CFG file. Then you boot the workstation, log in to the network, and run PSERVER.

PSERVER.VAP. For a NetWare v2.1x file server, you boot the file server and answer "Yes" to the prompt that asks if you want to load Value-Added Processes.

For an external bridge, you copy PSERVER.VAP to the same directory that contains the bridge's executable file. Then you boot the bridge and answer "Yes" to the prompt that asks if you want to load Value-Added Processes.

Printing Files and Screen Displays

Once the print server is running, you can print from any application, or you can print with the NetWare printing utilities.

To print screen displays, use the CAPTURE utility.

Printing from network applications. Many applications are designed for network printing. The application documentation will include specific directions for printing files.

If you use applications written for earlier versions of NetWare, they may send print jobs to "printer numbers" rather than "print queues." If so, you must set up spooler mappings at the file server console. Spooler mappings translate the printer number used by the application into a queue that NetWare can recognize.

Save the spooler mappings by including them in the system AUTOEXEC. BAT file (using "Supervisor Options" in SYSCON).

NOTE: The printer numbers that you specify in your application are not the same printer numbers used in PCONSOLE to define printers on the print server. The printer numbers used by applications are logical printer numbers (0-4 are used by earlier versions of NetWare). They do not correspond to the 16 printer numbers used by the print server.

Printing from non-network applications. If you are using applications that are *not* designed for network printing, NetWare provides four utilities that allow you to print files:

- CAPTURE
- ENDCAP
- NPRINT
- PCONSOLE

You should specify a default print queue for CAPTURE and NPRINT by setting up a spooler mapping with the NetWare SPOOL command.

If you do not specify a queue, the print server sends the print job to whichever queue Spooler0 is mapped to. If it has not been mapped to a queue, you receive an error message.

CAPTURE/ENDCAP. Use CAPTURE to

- Print screen displays;
- Print from applications that are not set up to print on networks.

  CAPTURE redirects any print command issued to a local printing port and sends it to a print queue to be printed.

  IMPORTANT: Printer numbers are no longer supported with CAPTURE. If you are upgrading from NetWare v2.x and have included the printer option (CAPTURE p=*x*) in your login script, you now need to use the queue option (CAPTURE q=*x*).

- Send data to a file rather than to a print queue.

Enter CAPTURE at the command line or include it in your login script or an AUTOEXEC.BAT file to set up network printing each time you log in to the file server.

Use ENDCAP to end the CAPTURE command and direct your print commands back to the local workstation printing ports.

For complete information, see "CAPTURE" (*Print Server*).

NPRINT. Use NPRINT to print files that were created in applications or as ASCII text files.

NPRINT routes the files to a specified print queue. You can specify printing parameters at the command line or set up a print job configuration as a shortcut.

To print a file with NPRINT, you must specify the directory path and filename of the specific file you want sent to the queue. NPRINT sends your files to the print queue immediately.

For complete information, see "NPRINT" (*Print Server*).

PCONSOLE. Use PCONSOLE to print files that were created in applications or as ASCII text files.

Like NPRINT, PCONSOLE routes the files to a specified print queue. However, PCONSOLE is a menu utility and allows you to select your file and specify printing parameters from a menu rather than at the command line.

For complete information, see "PCONSOLE" (*Print Server*).

Customizing the Printing Environment

If you will print frequently from applications that are not designed for network printing, you can customize your printing environment with PRINTDEF and PRINTCON.

PRINTDEF and PRINTCON are used to simplify printing with CAPTURE, NPRINT, and PCONSOLE.

We suggest you use PRINTDEF and PRINTCON in two cases only:

- If you have problems resetting your printer to default fonts after using different printing setups.

  For example, if you are printing spreadsheets and letters on the same printer, you may have some problems with the printer resetting after each print job.

- If you want to specify different forms or print configurations when you are using CAPTURE, NPRINT, and PCONSOLE.

PRINTDEF creates a database that stores print forms and print device definitions. (Print devices are equipment such as printers and plotters.)

Print device definition files (.PDF). These files are used to issue specific commands to the printer at the beginning of a print request. These files can be set up in one of three ways:

- NetWare includes print device definitions for commonly used printers. The print device definitions are automatically placed in the SYS:PUBLIC directory when you install NetWare.
- You can also import print definitions set up on another file server.
- You can set up your own printer definitions using PRINTDEF if your printer is *not* found in NetWare's device definitions.

Print job configurations. Print job configurations control the way a job is printed.

PRINTCON uses the information set up in PRINTDEF to create customized print job configurations. PRINTCON creates individual databases of print job configurations for each user.

A primary advantage of setting up print job configurations is that users printing a file using PCONSOLE, CAPTURE, or NPRINT can choose a print job configuration instead of manually entering all the print job specifications each time they want to print.

The printer returns to its defaults after the print job is completed.

As you set up printing on the network, you can

- Create print job configurations for all users, using the forms, devices, and modes defined in PRINTDEF;
- Select the default print job configuration;
- Copy print job configurations from one user to another.

You can set up print job configurations for any user. In addition, users can set up their own print job configurations.

Related utility: "PRINTCON" (*Print Server*).

Sending Commands to the Print Server or Printer

You can send commands to the print server or printer with PCONSOLE or PSC. You can perform the following tasks:

- View the status of one or all printers
- Pause the printer temporarily
- Stop printing the current job
- Start the printer
- Mark top of form
- Rewind a print job (only in PCONSOLE)
- Advance the printer to the top of the next page
- Mount a new form
- Bring down the print server (only in PCONSOLE)

With PSC, you can also remove a remote printer from the list of network printers so that it can be used as a standalone printer.

You can use the DOS SET command to set a default print server and printer number for PSC so that you do not have to specify the print server and printer at the command line each time you execute PSC.

Related utility: "PSC" (*Print Server*).

Prompt

A character or message (generated by the software) that appears on the display screen and requires a response (such as a command or a utility name) from the user.

Standard types of prompts include the following:

- The DOS prompt, which displays one of the local drives (A: to E:) followed by a greater than symbol: A:>
- The network prompt, which displays one of the network drives (F: to Z:) followed by a greater than symbol: F:>
- The file server console prompt, which displays a colon (:)
- The OS/2 prompt, which displays a drive letter in brackets: [B]

Property

A descriptive feature of a bindery object such as a password, account restriction, account balance, internetwork address, or list of authorized clients.

See also “Bindery” on page 30; “Object” on page 173.

Protected mode

The mode that 80286, 80386, and 80486 processors run in by default. When running in protected mode, these processors aren't subject to the same memory constraints as 8086 processors.

The 80286 processor uses a 24-bit address bus, and can address up to 16 MB of memory. The 80386 and 80486 processors use a 32-bit address bus, and can address up to 4 GB of memory.

Protected mode allows for multitasking (running more than one application or process at a time).

Protected mode allocates memory to various processes running concurrently so that memory used by one process doesn't overlap memory used by another process.

By contrast, 8086 processors can address only 1 MB of memory, and can run only one application or process at a time.

The 80286, 80386, and 80486 processors can be set to run in *real* mode, in which case they emulate an 8086 processor (and are subject to its memory constraints).

See also "Real mode" on page 202.

Protocol

See "NetWare Core Protocol" on page 156.

Protocol, NetWare Core

See "NetWare Core Protocol" on page 156.

Public access

A security condition that gives all NetWare users access rights to a particular directory.

For example, all NetWare users must be able to access NetWare utilities. Therefore, NetWare utilities are usually placed in a directory (named SYS:PUBLIC) that has public access rights; in other words, all users have rights to Open, Read, and Search for files in that directory.

PUBLIC directory

The SYS:PUBLIC directory is created automatically on the SYS: volume during network installation and cannot be deleted.

The NetWare utilities, as well as the .OVL and .DAT files necessary to run the menu utilities, are copied into the PUBLIC directory during installation.

All NetWare users have a search drive mapped to the PUBLIC directory via the system login script and are assigned Read and File Scan rights to this directory.

Public files

Files that need to be accessed by all NetWare users.

By convention, they are located in the SYS:PUBLIC directory. This general-access directory is created automatically during network installation and cannot be deleted.

NetWare utilities, help files, and some message and data files (such as the system login script file, NET$LOG.DAT) are public files that all NetWare users have Read, Open, and Search rights to.

Purge attribute

See "Attributes" on page 21; "Security" on page 221 (Attribute Security).

16 Q

Queue

See “Printing” on page 183.

17 R

RAM buffer address

See “Base memory address” on page 29.

RCONSOLE

See “Remote Console” on page 204.

Read-after-write verification

A means of assuring that data written to the hard disk matches the original data still in memory.

If the data from the disk matches the data in memory, the data in memory is released. If the data does not match, the block location is recognized as "bad," and Hot Fix redirects the data to the Hot Fix redirection area.

See also “Data protection” on page 49; “Hot Fix” on page 117; “Redirection area” on page 203.

Read Audit attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Read Only attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Read right

See “Rights” on page 207; “Security” on page 221 (Rights Security).

Read Write attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Real mode

The mode that allows an 80286 or 80386 processor to emulate an 8086 processor and run as though it actually were an 8086 processor.

The 8086 processor uses a 20-bit address bus, and can address up to 1 MB of memory. The 8086 processor is also limited to running one application or process at a time.

When running in protected mode, the 80286, 80386, and 80486 processors are capable of multitasking and addressing much more than 1 MB of memory.

When running in real mode, these processors are subject to the same 1MB memory constraint as the 8086 processor, and they can run only one application or process at a time.

However, the 80286 and 80386 processors running in real mode are more efficient than the 8086 processor, because they operate at a faster clock rate.

See also “Protected mode” on page 196.

Record locking

A feature of the NetWare operating system that prevents different users from gaining simultaneous access to the same record in a shared file, thus preventing overlapping disk writes and ensuring data integrity.

Recursive copying

The process of copying a specified source directory to a destination directory until all the files and subdirectories in and below the specified source directory are copied.

Recursive copying copies all directories and files of a logical drive to the destination, keeping them exactly as they were in the source directory.

Whether a trustee's rights are copied with the files and directories depends on what rights are assigned in the destination directory.

The DOS XCOPY and NetWare backup utilities use recursive copying.

Redirection area

A small portion of a network hard disk that is set up as a table to hold data blocks that are "redirected" from bad block locations on the disk.

The Hot Fix redirection area is set aside during installation.

See also“Hot Fix” on page 117.

Remote boot

A method that allows a user to boot a workstation from remote boot image files on a file server rather than from a boot diskette in the workstation's local drive.

When a workstation is booted, the Remote Reset PROM (installed on the workstation's network board) directs the nearest (default) file server to run the remote boot image file commands contained in that server's SYS:LOGIN directory.

See also “Remote Reset” on page 205.

Remote connection

A connection between a LAN and a workstation or network, often using telephone lines.

A remote connection allows data to be sent and received across greater distances than those allowed by normal cabling.

Remote Console

Software that allows network supervisors to manage NetWare 386-based servers from a DOS workstation or from a PC using a modem.

Network supervisors can manage all NetWare v3.x servers on the internetwork from one centralized location, reducing the need for multiple network administrators.

Remote Console allows network supervisors to execute console commands from a workstation over either SPX (direct) or asynchronous connections. Supervisors can perform the following tasks remotely:

- Monitoring the file server
- Installing a file server
- Upgrading a file server
- Accessing the following file server console utilities from the workstation:

  ADD NAME SPACE
  BIND, UNBIND (communication protocols)
  DISKSET
  MOUNT, DISMOUNT (volumes)
  LOAD, UNLOAD (loadable modules)
  REMOVE DOS
  SECURE CONSOLE
  UPS STATUS

Remote Console is designed primarily for large, multiserver networks.

For more information about Remote Console, refer to "REMOTE" in *System Administration.*

Remote Reset

Software that allows you to boot a DOS workstation (including a diskless workstation) from a remote boot image file on a file server, rather than from a boot diskette in the workstation's local drive.

To use Remote Reset to boot a workstation, you must do the following:

1. Install a Remote Reset PROM on the station's network board and run the DOSGEN utility.

   DOSGEN uploads the station's boot files into a *remote boot image file*, NET$DOS.SYS, in the server's LOGIN directory.

   The remote boot image file includes the station's AUTOEXEC.BAT file, used by the station as if the file were present on a local boot diskette.

2. Copy the workstation's AUTOEXEC.BAT file to the remote boot image file, to the LOGIN directory, and to any default directory named in the workstation's login script.

Using Remote Reset with Multiple Servers

If you have multiple file servers on your network, you must copy the remote boot image files onto each server that may come up as the remote boot workstation's default server.

Then, if the first default server isn't available, the station can boot from the next available server.

Using Multiple Remote Boot Image Files

If you want more than one workstation to use Remote Reset, you must do the following:

1. Upload multiple remote boot image files for each station into SYS:LOGIN.

   Upload a separate remote boot image file for each workstation. Name the image files for each user (FRED.SYS for user FRED, JANE.SYS for user JANE, etc.).

2. In the LOGIN directory, create a BOOTCONF.SYS file, which is a DOS text file that identifies the following for each station's network board:

3. IPX external network number

4. Node number
5. Remote boot image filename (FRED.SYS, JANE.SYS, etc.)

Remote workstation

A terminal or personal computer connected to the network by a router or through a remote asynchronous connection.

A remote workstation can be either a standalone computer or a workstation on another network.

Rename Inhibit attribute

See "Attributes" on page 21; "Security" on page 221 (Attribute Security).

Resource

A manageable component of a network.

Resources include the following:

- *Networking components*—cabling, hubs, concentrators, adapters, and network boards.
- *Hardware components*—servers, workstations, hard disks, printers, etc.
- *Major software components*—the NetWare operating system and resulting network services such as file, mail, queue, communication, etc.
- *Minor software components* that are controlled by the operating system of its subsystems—protocols, gateways, LAN and disk drivers, etc.
- *Data structures and other network resources* that don't easily fit into one of the above categories, or are created by a combination of network components—volumes, queues, users, processes, security, etc.

Resource tags

Operating system tags that keep track of file server resources such as screens and allocated memory.

NLMs request a resource from the file server for each kind of resource they use and give it a resource tag name.

NLMs return resources when they no longer need them. When an NLM is unloaded, its resources are returned to the file server.

Resource tags ensure that allocated resources are returned properly.

Related utility: "MONITOR" (*System Administration*).

Restore

A retrieval of data previously copied and backed up to a storage media. Perform a restore if data has been lost or corrupted since the backup.

See also “Backup” on page 25; “Data set” on page 52.

Ribbon cable

A cable in which the wires are placed side by side in the insulation material instead of being bunched together in a circle inside the insulation material.

Rights

Qualities assigned to users and groups that control what tasks they can do with directories or files.

Eight rights can be granted at either the directory or file level: Access Control, Create, Erase, File Scan, Modify, Read, Supervisory, and Write.

How Rights Are Granted

- Rights granted at the directory level can be redefined for a file by making new trustee assignments or revoking rights from the Inherited Rights Masks.
- SUPERVISOR has all rights and grants trustee assignments to users and groups.
- When users are created, they are made a member of the group EVERYONE and obtain any rights granted to that group. All other rights must be granted individually to users or groups.
- Users are granted the right to search to the root of any directory where they have rights to a directory or file. They cannot see any subdirectories unless they are granted rights to the subdirectories or to files in the subdirectories.

Description of Rights

Each right is represented by its initial. NetWare utilities display the initial letters of these rights between brackets as shown below:

```
[S R W C E M F A]
```

By convention, if some rights have either been revoked or have not been assigned, the absence is indicated by a blank space:

```
[  R          F  ]
```

The meaning and effect of each right depends on whether it is assigned to a file or a directory.

Each right is defined twice below, once in "Directory Rights" and again in "File Rights."

Directory rights control general access to a directory, its files, and its subdirectories. When granted at the directory level, the rights apply to all the files and subdirectories in that directory unless the rights redefined at the file or subdirectory level.

Directory rights are defined in Table 12.

Table 10 Directory trustee rights

Right	Syntax	Permissions Granted
Supervisory	S	All rights to the directory, its files, and its subdirectories. Grant other users the Supervisory right. The Supervisory right overrides the Inherited Rights Masks and can be revoked only from the directory where it was granted.
Read	R	Open files in a directory and read their contents or run the programs.
Write	W	Open and modify files in the directory.
Create	C	Create files and subdirectories in the directory. Granting only Create at the directory level and no rights below the directory creates a "drop box" directory, where users can copy files but have no other rights in the directory.
Erase	E	Delete a directory, its files, its subdirectories, and its subdirectory files.
Modify	M	Change directory and file attributes. Rename the directory, its files, and its subdirectories.
File Scan	F	See directory files in a directory listing.
Access Control	A	Modify a directory's or a file's trustee assignments and Inherited Rights Mask. Grant any right (except Supervisory) to other users.

File rights control access to specific files in a directory. They are used to redefine the rights that users inherit from directory rights.

File rights are defined in Table 13.

Table 11 **File trustee rights**

Right	Syntax	Permissions Granted
Supervisory	S	All rights to the file. Modify the Inherited Rights Mask. Grant other users the Supervisory right.
Read	R	Open and read the file.
Create	C	Salvage the file after it has been deleted.
Write	W	Open and write to the file.
Erase	E	Delete the file.
Modify	M	Change the file's attributes and rename the file.
File Scan	F	See the filename when viewing the directory. See the directory structure, from the file to the root of the directory.
Access Control	A	Modify the file's trustee assignments and Inherited Rights Mask. Grant all file rights, except Supervisory, to other users.

See also “Inherited Rights Mask” on page 119; “Security” on page 221 (Rights Security).

Related utilities: "ALLOW"; "FILER"; "GRANT"; "REVOKE"; "RIGHTS"; "SYSCON" (*Utilities Reference*).

Root directory

The highest directory level in a hierarchical directory structure.

With NetWare, the root directory is the volume; all other directories are subdirectories of the volume.

See also “Directory structure” on page 61.

Router

A workstation or file server running software that manages the exchange of information (in the form of data packets) between network cabling systems.

NetWare Routers vs. Traditional Bridges

NetWare routers, unlike traditional bridges, do more than transfer data packets between networks that use the same communications protocol.

NetWare routers are "intelligent." They not only pass packets between different cabling systems, but they also route the packets through the most efficient path.

NetWare routers can also connect cabling systems that use different kinds of transmission media and different addressing systems.

For example, a NetWare router can connect a network using the Ethernet addressing structure and RG/58 coaxial cable to another network using the ARCnet addressing structure and RG/62 coaxial cable.

Local vs. Remote Routers

A *local router* is used within the cable length limitations for its LAN drivers.

A *remote router* is connected beyond its driver limitations or through a modem.

Internal vs. External Routers

An *internal router* runs as part of a file server. It connects separate network cabling topologies or separate networks by way of the file server's NetWare operating system. See Figure 49.

Figure 49 Internal router

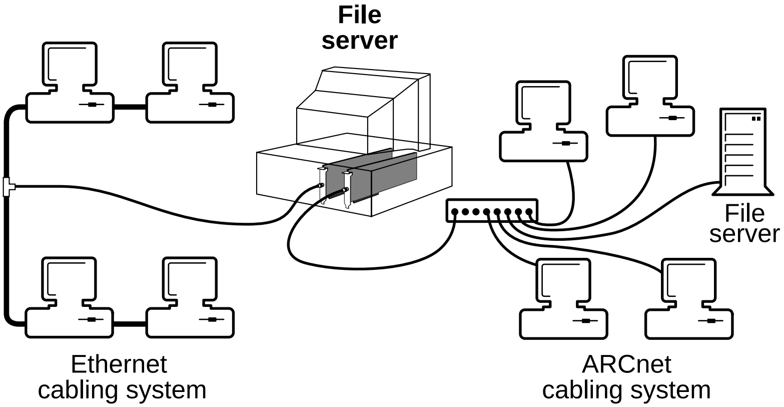

One solution to heavy utilization of the network is to create a "backbone" for the multiple file servers. The central cable handles all the network traffic, thus increasing packet transmission speed and decreasing traffic on the network.

If you have three or more file servers, a backbone is an efficient way to implement internal routing. See Figure 50.

Figure 50 Network backbone

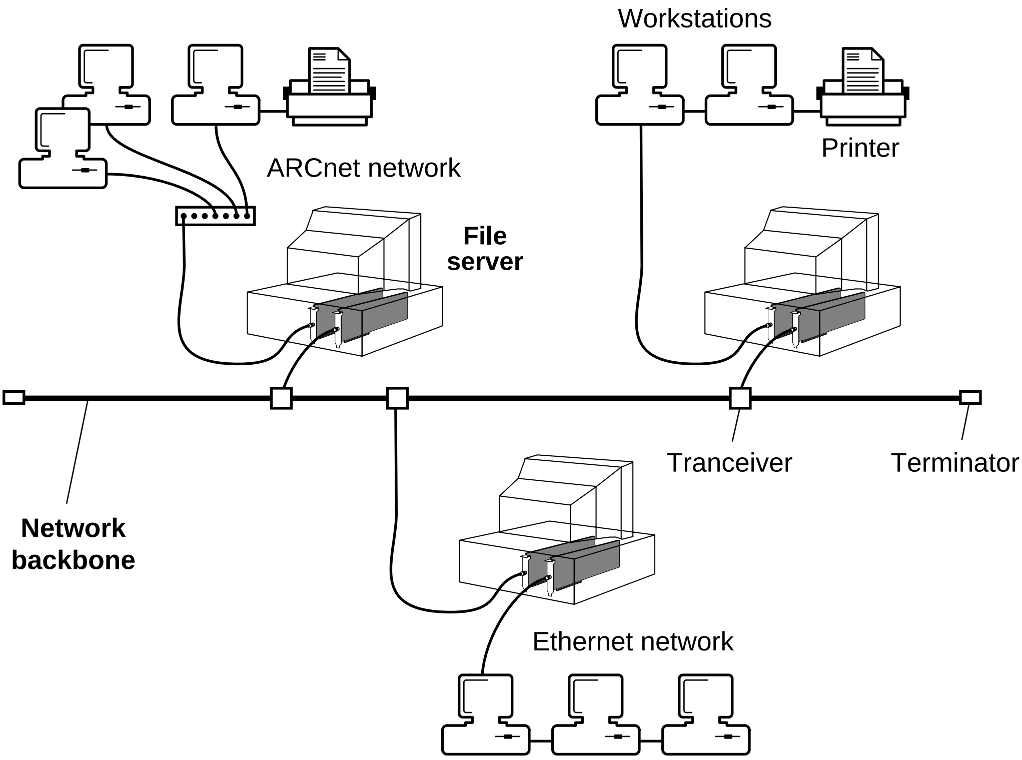

An *external router* runs in a networked computer that is not a file server. It manages packet routing with its ROUTER.EXE file. See Figure 51.

Figure 51 External router

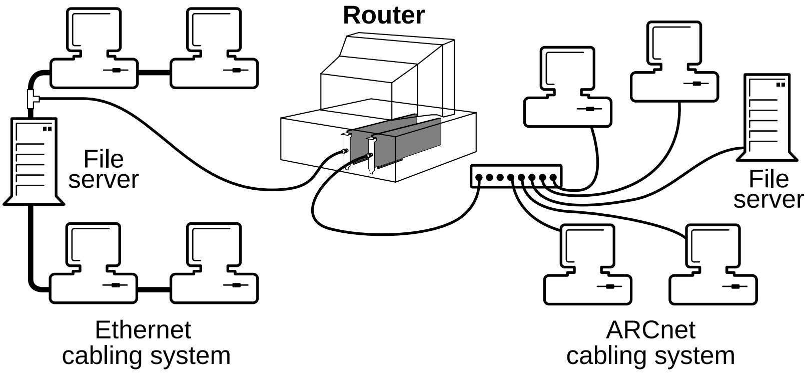

Dedicated vs. Nondedicated External Routers

A *dedicated router* works only as a router; it cannot function simultaneously as a workstation

Dedicated routers are more reliable. Workstation applications cannot hang and cause the router to stop operating.

A *nondedicated router* can function simultaneously as a router and as a workstation. In a nondedicated router, the workstation's NetWare workstation files runs "on top of" the router software.

The principal advantage of using a nondedicated router is that you eliminate the expense of buying an additional computer.

Real Mode vs. Protected Mode

The type of microprocessor your router has and the amount of memory installed determine whether you should run the router in real or protected mode.

Protected-mode router. NetWare routers that contain an 80286 or 80386 microprocessor can have up to 12 MB of RAM (12 MB is a NetWare router limitation.) Anything above 1 MB of memory is called "extended memory," which allows the router and Value-Added Processes (VAPS) to run on the router.

Programs that run in extended memory run in protected mode, meaning they are protected them against interference from other programs. See Figure 52.

Figure 52 Protected mode

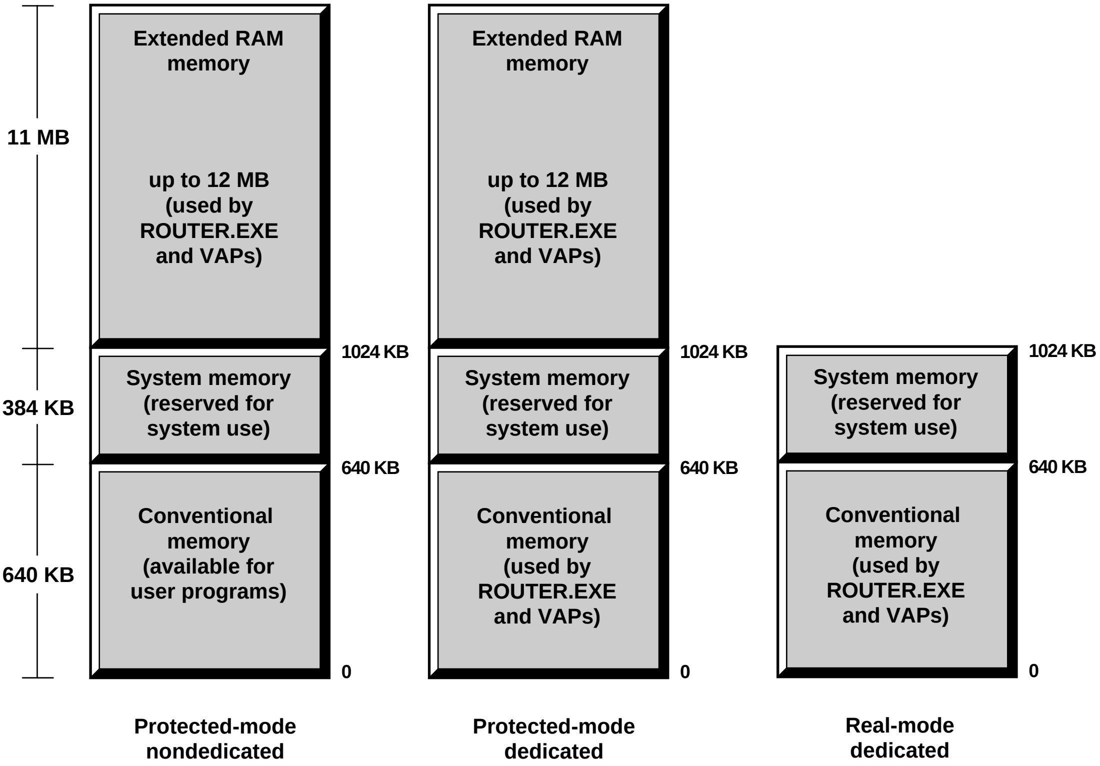
11 MB
384 KB
640 KB
Extended RAM memory
up to 12 MB (used by ROUTER.EXE and VAPs)
System memory (reserved for system use)
Conventional memory (available for user programs)
1024 KB
640 KB
0
Extended RAM memory
up to 12 MB (used by ROUTER.EXE and VAPs)
System memory (reserved for system use)
Conventional memory (used by ROUTER.EXE and VAPs)
1024 KB
640 KB
0
System memory (reserved for system use)
Conventional memory (used by ROUTER.EXE and VAPs)
1024 KB
640 KB
0
Protected-mode nondedicated
Protected-mode dedicated
Real-mode dedicated

Real-mode router. NetWare routers that contain an 8086 or 8088 microprocessor have the standard 640 KB of base memory. They are called real-mode routers.

You can run either a dedicated or a nondedicated router in the real mode. However, the limited memory allows only one or two VAPs to run on the router.

For more information on routers, see *Installation*.

Routing buffers

Portions of memory reserved in a router's RAM.

Routing buffers are used to temporarily store and queue the message packets sent between communicating stations when the network bus is busy.

18 S

Salvageable files

Files that can be recovered, or *salvaged*, after being deleted by users.

NetWare usually stores salvageable files in the directory they were deleted from. If the user deletes that directory, the files are saved in a DELETED.SAV directory located in the volume's root directory.

The user can view a list of deleted files in a directory and recover files using the SALVAGE utility. Recovered files contain information about who deleted the files and when they were deleted.

Deleted files are saved until the user deliberately purges them or until the file server runs out of disk allocation blocks on the volume.

Files deleted from a local drive cannot be salvaged.

When the file server runs out of blocks, it purges deleted files on a first-deleted, first-purged basis.

Files and directories can also be purged as they are deleted. There are two ways to do this:

- Use the SET command at the file server to disable the salvageable file feature.

  This increases performance, but at the cost of losing the salvageable file feature.

- Set the Purge attribute on a directory or file.

  When a file is flagged with the Purge attribute, the file is purged when it is deleted.

When a directory is flagged with the Purge attribute, any file in that directory is purged when the file is deleted. Such files can't be recovered with the SALVAGE utility.

Related utility: "SALVAGE" (*Utilities Reference*)

SCSI

(Small Computer Systems Interface, commonly pronounced *scuzzy*) An industry standard that sets guidelines for connecting peripheral devices and their controllers to a microprocessor.

The SCSI interface defines both hardware and software standards for communication between a host computer and a peripheral.

Computers and peripheral devices designed to meet SCSI specifications have a large degree of compatibility.

SCSI bus

The SCSI bus (interface) connects disk controller boards to controllers and hard disks. You need to properly terminate and address the connected peripherals.

Termination

By terminating a signal, you ensure that it does not "echo" or become corrupted.

You must terminate the end components on a SCSI bus—that is, the first and the last controller and disk in a chain. Otherwise they may not be initialized properly. The default setting is "terminated."

Figure 53 shows which controllers and disks in a subsystem need to be terminated when the DCB controls only a subsystem.

Figure 53 SCSI bus termination

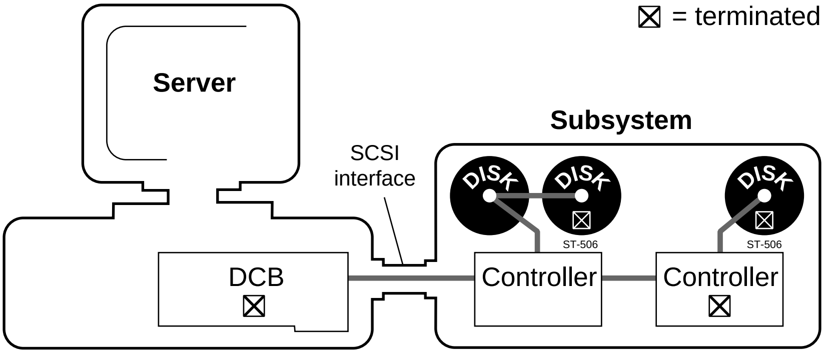

If the computer has no built-in controller on its system board and the DCB controls both internal hard disks and a subsystem, you must terminate the DCB, the controller, and the hard drives, as illustrated in Figure 54.

Figure 54 SCSI bus termination

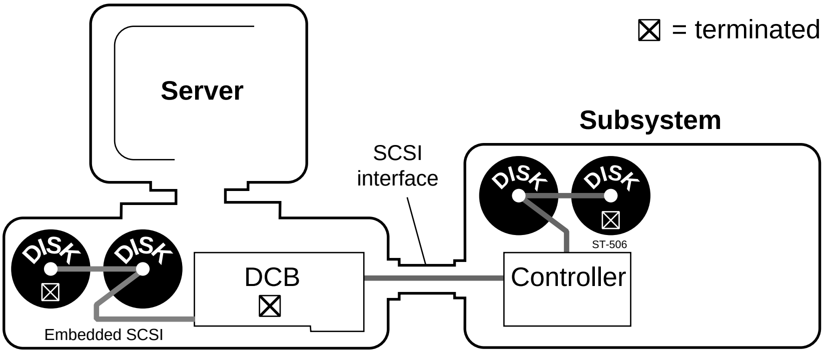

Addressing

Each controller must have an address unique to its disk channel.

You can find the physical address settings in the documentation shipped with the controller. Depending on where you place the controller in your system, you need to change the default setting (usually controller 0).

Within one channel, the addresses can range from 0 to 7. For reference purposes, the controller nearest to the DCB should always be controller 0.

For more information on physical addresses, see "Device numbering" on page 54.

SCSI disconnect

A feature in NetWare v3.12 that allows the SCSI driver to inform a disk that the disk should prepare for I/O.

While that disk is preparing for I/O, the SCSI driver can either send messages to other hard disks or actually perform I/O to other hard disks.

When the disk is ready to receive or send data, the disk sends a message back to the SCSI driver, informing the driver that the disk is ready for I/O. The SCSI driver then performs the I/O.

The SCSI driver does not wait while hard disks prepare for I/O, as it did in the past.

Formerly, a SCSI driver shook hands with a hard disk over the computer's bus, asked the hard disk to prepare for I/O, waited for the hard disk to prepare for I/O, performed I/O to the disk, and then let go.

This disk I/O method was inefficient because the SCSI driver was idle during the time that the disk was preparing for I/O.

The IBM SCSI board supports SCSI disconnect.

Search drive

A drive that is searched by the operating system when a requested file isn't found in the current directory.

Search drives are supported only from DOS workstations.

A search drive allows a user working in one directory to access an application or data file located in another directory.

See also "Drive mapping" on page 83.

Security

Elements that control access to the network or to specific information on the network.

Network security has four levels:

- Login security (see "Login Security" on page 221)
- Rights security (see "Rights Security" on page 222)
- Attribute security (see "Attribute Security" on page 242)
- File server security (see "Console" on page 46)

NetWare security controls

- Who can access the networks;
- What directories and files users can access;
- What tasks users can perform with those directories and files (for example, read or modify a file);
- Who can perform tasks at the file server console.

This entry explains NetWare security, provides guidelines for planning security, and lists the utilities available for establishing and monitoring security.

Login Security

The network supervisor establishes login security by assigning usernames, requiring passwords, and setting up restrictions.

Login security controls access to the network. It determines

- Which users can work on the file server;
- When users can work;
- What workstations users can work from;
- Which resources they can use.

Usernames provide the first level of security. Only network supervisors and workgroup managers can create usernames.

To access the file server, users must know a username (and its corresponding password if one is required).

Passwords are optional. However, if you don't require passwords, any person who knows a valid username can access the file server.

If you require passwords, you can assign them or allow users to choose and change their own passwords.

If you allow users to change their own passwords, you can increase password security by requiring the following:

- A minimum password length.
- A periodic change in the password.
- A unique password.

  This option prevents users from alternating between two favorite passwords.
- A limited number of logins permitted after a password has expired.

Restrictions control when and where a user can log in and protects your network from intruders. You can use three types of login restrictions:

- *Station restrictions* limit where a user can log in.

  You can specify which workstations a user can log in from and the number of workstations a user can be logged in from concurrently.
- *Time restrictions* limit when users can log in. You can restrict users to specific days and hours.
- *Account restrictions* lock a user's account when certain limits are exceeded.

  When an account is locked, no one can log in using that username. You can have an account locked automatically when it expires, when its balance is depleted, and when a set number of incorrect password uses is exceeded.

See also "Plan the Network Security" in *Installation*.

Related utility: "SYSCON" (*Utilities Reference*).

Rights Security

Rights security controls which directories, subdirectories, and files a user can access and what the user is allowed to do with them.

Rights security is controlled by trustee assignments and by the Inherited Rights Mask.

Trustee assignments grant rights to specific users (or groups) that specify how they can use a file or directory (for example, only for reading).

A trustee assignment grants users the right to see to the root of a directory. However, the users can't see any subdirectories unless they have rights to the subdirectories.

Inherited Rights Masks are given to files and directories when they are created.

The only rights the user can "inherit" for a file or subdirectory are rights that are allowed by the Inherited Rights Mask.

The default Inherited Rights Mask includes all rights. But this does not mean that users have all rights; users have only the rights granted in their trustee assignments.

For example, if a user is granted the Read right with a directory trustee assignment, the right to read files in a subdirectory could be revoked by having the Read right removed from the subdirectory's Inherited Rights Mask.

Trustee Rights

Trustee assignments and Inherited Rights Masks use the same eight trustee rights to control access.

Each right is represented by its initial. NetWare utilities display the initial letters of these rights between brackets, as shown below.

```
[S R W C E M F A]
```

By convention, if some rights have either been revoked or have not been assigned, the absence is indicated by a blank space:

```
[  R         F ]
```

The meaning and effect of each right depends on whether the right is assigned to a file or a directory.

Each right is defined twice below, once in "Directory Rights" and again in "File Rights."

Directory rights control general access to a directory, its files, and its subdirectories. When granted at the directory level, the rights apply to all the files and subdirectories in that directory unless redefined at the file or subdirectory level.

Directory rights are defined in Table 12.

Table 12 Directory trustee rights

Right	Syntax	Permissions Granted
Supervisory	S	All rights to the directory, its files, and its subdirectories. Grant other users the Supervisory right. The Supervisory right overrides the Inherited Rights Masks and can be revoked only from the directory where it was granted.
Read	R	Open files in a directory and read their contents or run the programs.
Write	W	Open and modify files in the directory.
Create	C	Create files and subdirectories in the directory. Granting only Create at the directory level and no rights below the directory creates a "drop box" directory, where users can copy files but have no other rights in the directory.
Erase	E	Delete a directory, its files, its subdirectories, and its subdirectory files.
Modify	M	Change directory and file attributes. Rename the directory, its files, and its subdirectories.
File Scan	F	See directory files in a directory listing.
Access Control	A	Modify a directory's or a file's trustee assignments and Inherited Rights Mask. Grant any right (except Supervisory) to other users.

To grant or modify a directory trustee assignment, use "FILER"; "GRANT"; "REMOVE"; "REVOKE"; "SYSCON" (*Utilities Reference*).

To modify a directory's Inherited Rights Mask, use "ALLOW"; "FILER" (*Utilities Reference*).

File rights control access to specific files in a directory. They are used to redefine the rights that users inherit from directory rights.

File rights are defined in Table 13.

Table 13 File trustee rights

Right	Syntax	Permissions Granted
Supervisory	S	All rights to the file. Modify the Inherited Rights Mask. Grant other users the Supervisory right.
Read	R	Open and read the file.
Create	C	Salvage the file after it has been deleted.
Write	W	Open and write to the file.
Erase	E	Delete the file.
Modify	M	Change the file's attributes and rename the file.
File Scan	F	See the filename when viewing the directory. See the directory structure, from the file to the root of the directory.
Access Control	A	Modify the file's trustee assignments and Inherited Rights Mask. Grant all file rights, except Supervisory, to other users.

To grant or modify file trustee assignments, use "FILER"; "GRANT"; "REMOVE"; "REVOKE"; "SYSCON" (*Utilities Reference*).

To modify the file's Inherited Rights Mask, use "ALLOW"; "FILER" (*Utilities Reference*).

Effective Rights

Effective rights are the rights a user can exercise in a given directory or file. To determine a user's effective rights, you must know

- What rights were granted in the trustee assignments for the user;
- What rights were granted in the trustee assignments for any groups the user belongs to;
- What rights were revoked with Inherited Rights Masks.

To view effective rights, use "FILER"; "RIGHTS"; "WHOAMI" (*Utilities Reference*).

The following examples explain how effective rights are determined for a directory and a file. In the diagrams, Inherited Rights Mask is abbreviated to IRM and trustee assignment to TA.

Directory effective rights. In directories, effective rights are determined by one of two methods:

- Calculating the effective rights to the parent directory and then determining which rights the Inherited Rights Mask allows to filter through.
- Granting a user a trustee assignment to a directory.

In a directory, trustee assignments override the directory's Inherited Rights Mask. If trustee assignments are granted, a user's effective rights are determined solely by the trustee assignments (user plus group trustee assignments).

Directory Effective Rights: Example 1

See Figure 55 for an example of determining the effective rights for JAN in the directory WORK\PROJECT.1.

Figure 55 Directory effective rights example 1

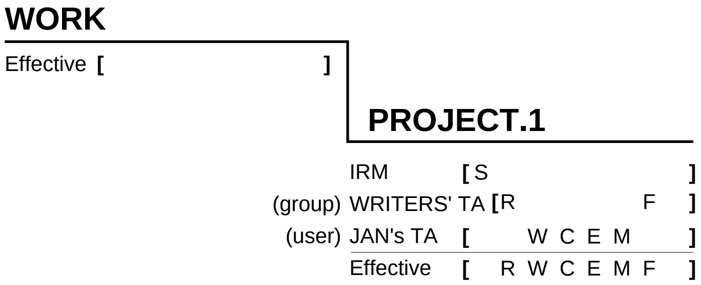
WORK
Effective [ ]
PROJECT.1
IRM [S]
(group) WRITERS' TA [R F]
(user) JAN's TA [W C E M]
Effective [R W C E M F]

Since JAN's user trustee assignment is W, C, E, and M and the group WRITERS is assigned R and F, the sum of these trustee assignments a re her effective rights to the PROJECT.1 directory.

Since she has a trustee assignment to PROJECT.1, JAN has limited File Scan rights in WORK. She can see PROJECT.1 when scanning the WORK directory.

In the examples on the following pages, Example 2 explains how to determine rights if no new trustee assignments are granted to the user. Example 3 explains how to determine rights when a new trustee assignment is granted to the user. Example 4 explains how to determine rights when the Supervisory right has been granted.

Directory Effective Rights: Example 2

If no new trustee assignments are granted, the Inherited Rights Mask determines the effective rights for the directory.

Step 1: To calculate the rights, determine the user's effective rights to the parent directory. From the example in Figure 56, calculate JAN's effective rights for the WORK directory.

Figure 56 Directory effective rights example 2

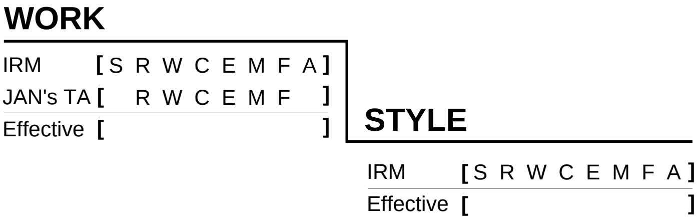

In this example, JAN's effective rights to the WORK directory are R, W, C, E, M, and F.

Step 2: View the Inherited Rights Mask of the STYLE directory.

- If the mask has all rights, JAN's effective rights to the STYLE directory are the effective rights of the WORK directory.
- If the STYLE directory's Inherited Rights Mask has rights revoked, the parent effective rights must be matched with the remaining rights in the mask. Only those that match are effective rights in the STYLE directory.

In the previous figure, JAN's effective rights to the STYLE directory are R, W, C, E, M, and F because the mask allows all rights to be inherited. What are JAN's effective rights in 1988.STY in Figure 57?

Figure 57 Directory effective rights example 2

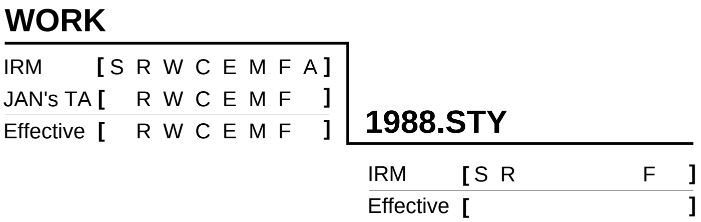

JAN's effective rights to the 1988.STY directory are R and F. The Supervisory right has not been revoked from the Inherited Rights Mask. (This right cannot be revoked from an Inherited Rights Mask.) However, the S right in the mask has no effect if it isn't granted to the user.

Directory Effective Rights: Example 3

If a trustee assignment is made to a directory, the trustee assignment overrides the directory's Inherited Rights Mask and the effective rights from the parent directory.

The user's effective rights are those granted in the trustee assignment.

In Figure 58, JAN's effective rights to the 1988.STY directory are R and F. What are JAN's effective rights to the REVISED directory?

Figure 58 Directory effective rights example 3

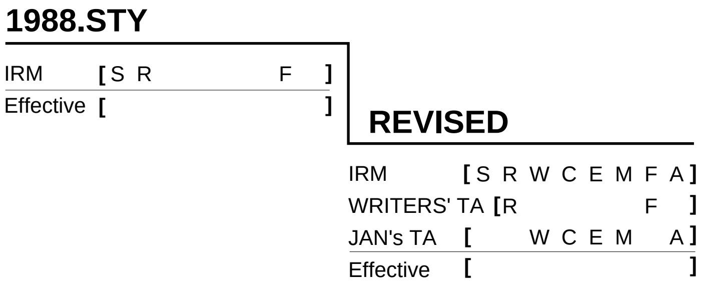

1988.STY

IRM	S	R					F	
Effective								

REVISED

IRM	S	R	W	C	E	M	F	A
WRITERS' TA		R					F	
JAN's TA			W	C	E	M		A
Effective								

JAN's effective rights to REVISED are R, W, C, E, M, F, and A.

Directory Effective Rights: Example 4

If a user has the Supervisory right in the parent directory, the user has all rights to all directories below it regardless of trustee assignments or Inherited Rights Masks.

For example, suppose that JAN is made the supervisor of a new project with the code name "Unicorn." The network supervisor grants her the Supervisory right to the UNICORN directory containing the project.

Refer to Figure 59. What are JAN's effective rights to the UNICORN directory? What are her rights to the PHASE.1 directory?

Figure 59 Directory effective rights example 4

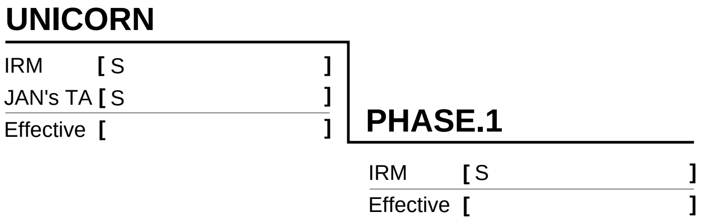

UNICORN

IRM	S							
JAN's TA	S							
Effective								

PHASE.1

IRM	S							
Effective								

JAN's effective rights to the UNICORN directory are all rights: S, R, W, C, E, M, F, and A. She has the same rights to the PHASE.1 directory.

Directory Effective Rights: Example 5

In Figure 60, JAN's rights are redefined in the PHASE.2 directory with a new trustee assignment. What are JAN's effective rights to the PHASE.2 directory?

Figure 60 Directory effective rights example 5

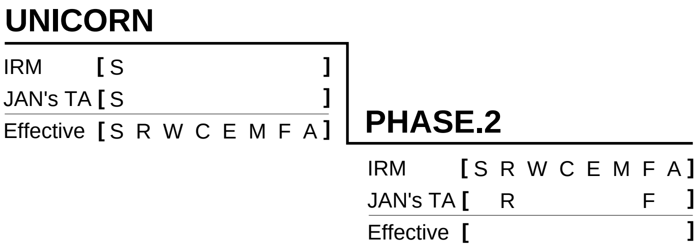

UNICORN

IRM [S]

JAN's TA [S]

Effective [S R W C E M F A]

PHASE.2

IRM [S R W C E M F A]

JAN's TA [R F]

Effective [ ]

JAN's effective rights to the PHASE.2 directory are still all rights. Her rights cannot be redefined below the UNICORN directory because she has the Supervisory right in UNICORN.

Summary for determining directory effective rights. The following figure summarizes the principles governing effective rights to directories.

Figure 61 Effective rights to directories

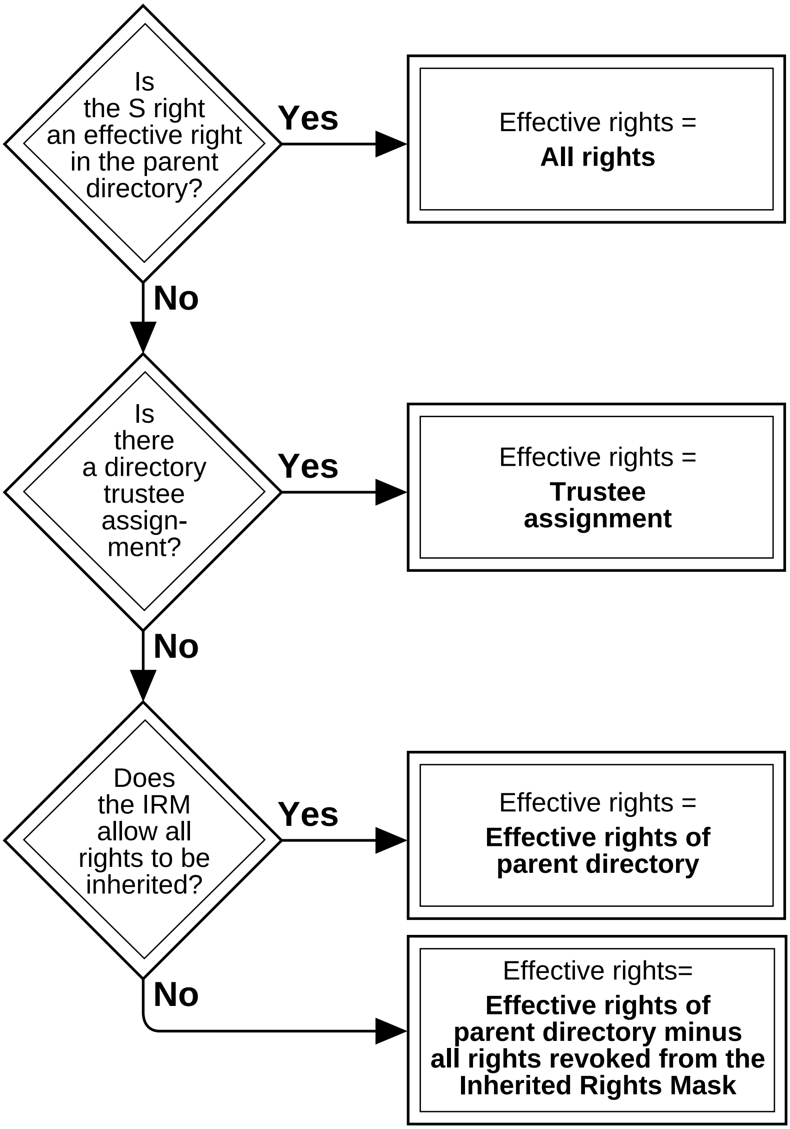
Is the S right an effective right in the parent directory?
Yes
Effective rights = All rights
No
Is there a directory trustee assign- ment?
Yes
Effective rights = Trustee assignment
No
Does the IRM allow all rights to be inherited?
Yes
Effective rights = Effective rights of parent directory
No
Effective rights= Effective rights of parent directory minus all rights revoked from the Inherited Rights Mask

File effective rights are determined in nearly the same way as subdirectory effective rights:

- By calculating the effective rights to the directory and then determining which rights the file's Inherited Rights Mask allows to filter through
- By granting a user a trustee assignment to the file
- By determining the user's effective rights to the parent directory

In the examples that follow,

- Example 1 explains how to determine rights if no new trustee assignments are granted at the file level.
- Example 2 explains how to determine rights when a new trustee assignment is granted at the file level.
- Example 3 explains how to determine rights when the Supervisory right has been granted.

  NOTE: If you are using an application that creates extra files (such as backup files), you need to assign the users of that application the Create right in the directory where the application places those files.

File Effective Rights: Example 1

If no new trustee assignments are granted for the file, the file's Inherited Rights Mask determines which effective rights a user inherits from the directory.

Step 1: To calculate the rights, first determine the user's effective rights to the parent directory. From Figure 62, calculate JAN's effective rights for the WORK directory.

Figure 62 File effective rights example 1

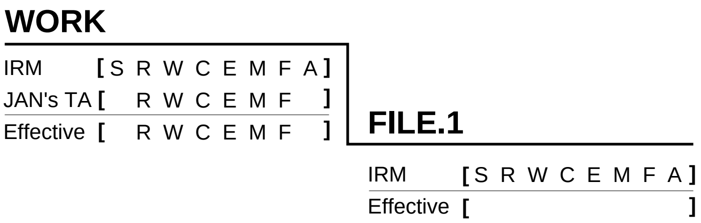

WORK

IRM	[S R W C E M F A]
JAN's TA	[R W C E M F]
Effective	[R W C E M F]

FILE.1

IRM	[S R W C E M F A]
Effective	[ ]

JAN's effective rights to the WORK directory are R, W, C, E, M, and F.

Step 2: View the Inherited Rights Mask of the file.

- If the mask has all rights, the user's effective rights to the file are the effective rights of the parent directory.
- If the file's Inherited Rights Mask revokes some of the rights, the directory effective rights must be compared with the remaining rights in the mask. Only those that match are effective in the file.

In the previous figure, JAN's effective rights to the FILE.1 file are R, W, C, E, M, and F because the mask allows all rights to be inherited. What are they in Figure 63?

Figure 63 File effective rights example 1

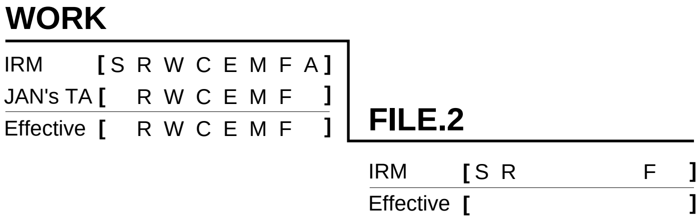

WORK

IRM	[S	R	W	C	E	M	F	A]
JAN's TA	[	R	W	C	E	M	F	]
Effective	[	R	W	C	E	M	F	]

FILE.2

IRM	[S	R					F	]
Effective	[							]

JAN's effective rights to FILE.2 are R and F. Notice that the Supervisory right has not been revoked from the Inherited Rights Mask. (This right cannot be revoked from an Inherited Rights Mask.) However, the S right in the mask has no effect if it isn't granted to the user.

File Effective Rights: Example 2

If a trustee assignment is granted to a file, effective rights are determined the same as they are for a directory.

The trustee assignment overrides the effective rights from the parent directory and the file's Inherited Rights Mask. The user's effective rights are those granted in the user and group trustee assignments.

In Figure 64, JAN's effective rights to the PROJECTS directory are R, W, C, E, M and F. What are JAN's effective rights to the FILE.3 file?

Figure 64   File effective rights example 2

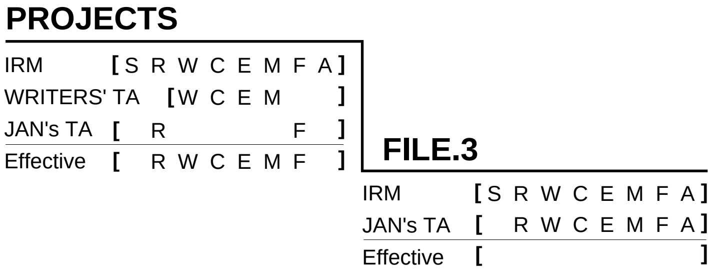

```
PROJECTS
IRM          [S R W C E M F A]
WRITERS' TA    [W C E M      ]
JAN's TA     [  R          F ]
Effective    [  R W C E M F  ]

                  FILE.3
                  IRM        [S R W C E M F A]
                  JAN's TA   [  R W C E M F A]
                  Effective  [               ]
```

JAN's effective rights to FILE.3 are R, W, C, E, M, F, and A.

File Effective Rights: Example 3

If a user has the Supervisory right in the parent directory, the user has all rights to the files regardless of other trustee assignments or the Inherited Rights Masks.

In Figure 65, JAN has the Supervisory right to directory UNICORN. What are her rights to FILE.4?

Figure 65   File effective rights example 3

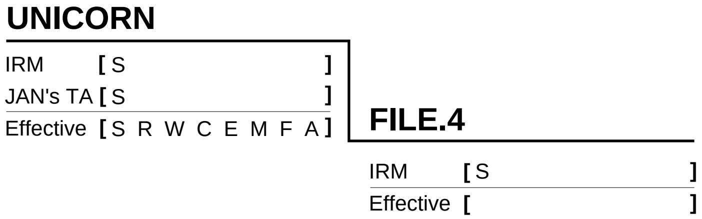

```
UNICORN
IRM        [S                 ]
JAN's TA   [S                 ]
Effective  [S R W C E M F A]

                  FILE.4
                  IRM        [S              ]
                  Effective  [               ]
```

JAN's effective rights to FILE.4 are all rights.

In Figure 66, JAN's rights are redefined in the FILE.5 file with a new trustee assignment. What are JAN's effective rights to the FILE.5 file?

Figure 66 File effective rights example 3

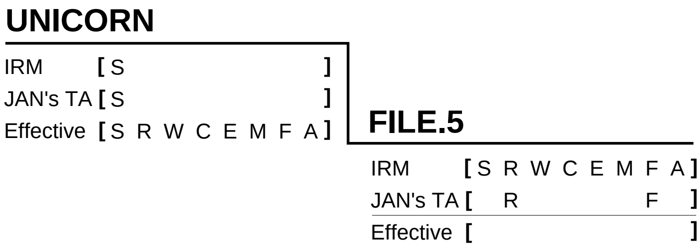

UNICORN	
IRM	[S]
JAN's TA	[S]
Effective	[S R W C E M F A]

FILE.5	
IRM	[S R W C E M F A]
JAN's TA	[R F]
Effective	[ ]

JAN's effective rights to the FILE.5 file are still all rights. Her rights cannot be redefined below the UNICORN directory because she has the Supervisory right in the UNICORN directory.

Summary for determining file effective rights. The following figure summarizes the principles governing effective rights for files.

Figure 67 Determining effective rights to files

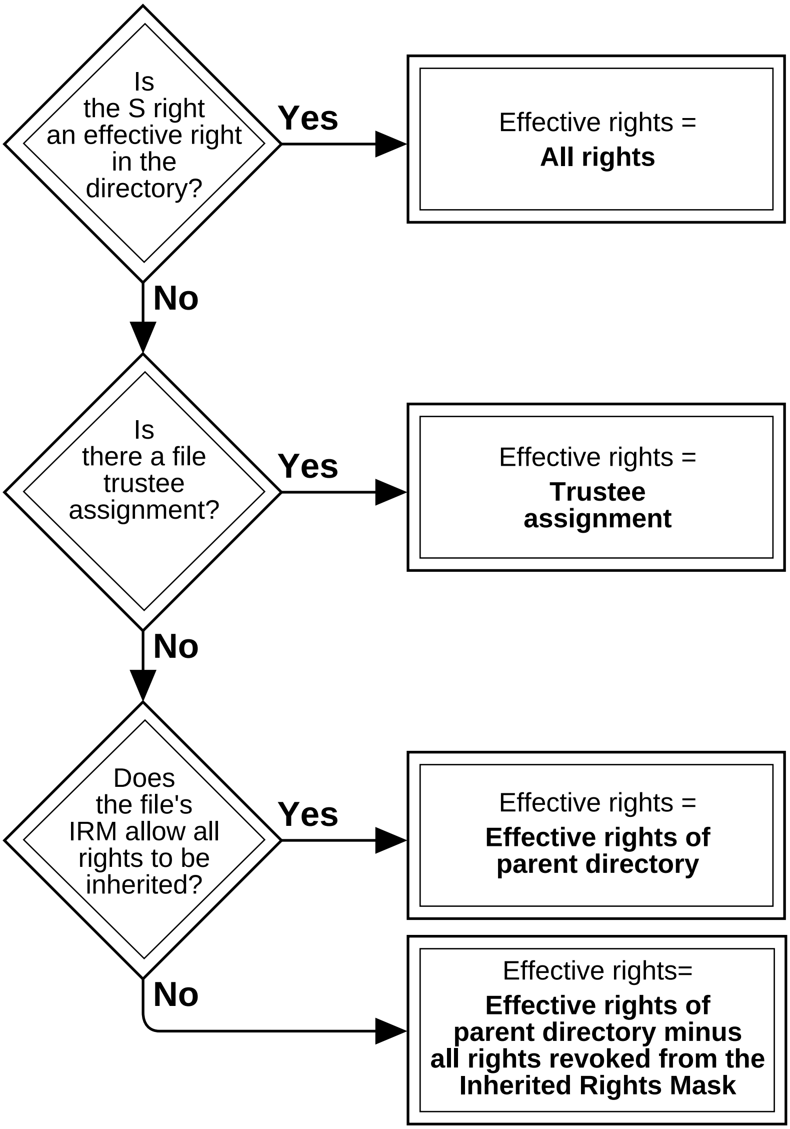

Is the S right an effective right in the directory?
Yes
Effective rights = **All rights**
No
Is there a file trustee assignment?
Yes
Effective rights = **Trustee assignment**
No
Does the file's IRM allow all rights to be inherited?
Yes
Effective rights = **Effective rights of parent directory**
No
Effective rights= **Effective rights of parent directory minus all rights revoked from the Inherited Rights Mask**

Assigning Rights

You can grant any combination of rights, although some combinations are not useful. For example,

- You can grant a user the Supervisory right in a directory and revoke all other rights. However, that user still has all rights to the directory.
- You can grant a user only Read and File Scan rights in a directory for an application that creates temporary files when accessed. This user cannot work because the Create, Erase, and Modify rights have not been granted.

You need to grant rights carefully. If you grant users more rights than they need, they can delete, corrupt, or steal data. If you grant them too few rights, they cannot do the work assigned to them.

When you assign rights to application directories, check the application documentation. If possible, separate the files that users create when they use the application from the executable files for the application.

Normally, users need only R and F rights to directories with executable files, while they need R, F, C, E, and M rights to directories in which they create files.

If a search drive is mapped to applications directories, the users can access the applications from another directory in which they have the additional rights they need to create files.

Sample Scenario for Assigning Rights

The scenario described below clarifies the interaction between trustee assignments and Inherited Rights Masks and demonstrates one method of assigning rights.

The scenario in Figure 68 on page 239 is followed by some questions that clarify the assignment of rights.

Suppose you have the following groups with the supervisors indicated.

MAC	PC	EDITORS
Mel (supervisor)	Pam (supervisor)	Jan (supervisor)
Mike	Pat	Janene
Mary	Paul	Jean
Mona	Peter	Judy

In Figure 68 on page 239 these users and groups have been granted trustee assignments to various directories in the SYS: volume. The rights allowed by the Inherited Rights Mask (IRM) are listed under the directory's name.

For example,

IRM [S C]

indicates that the Inherited Rights Mask has been modified to allow only the Supervisory and Create rights to be inherited.

The trustee assignments are listed below the Inherited Rights Mask and are preceded by either a group's or user's name. For example, [SRWCEMFA] is Pam's directory trustee assignment.

Following the diagram are some questions that ask you to determine the effective rights of users in some of the directories and the tasks that they can perform in certain directories. Following the questions, the correct answers are given with explanations.

Figure 68 Sample scenario

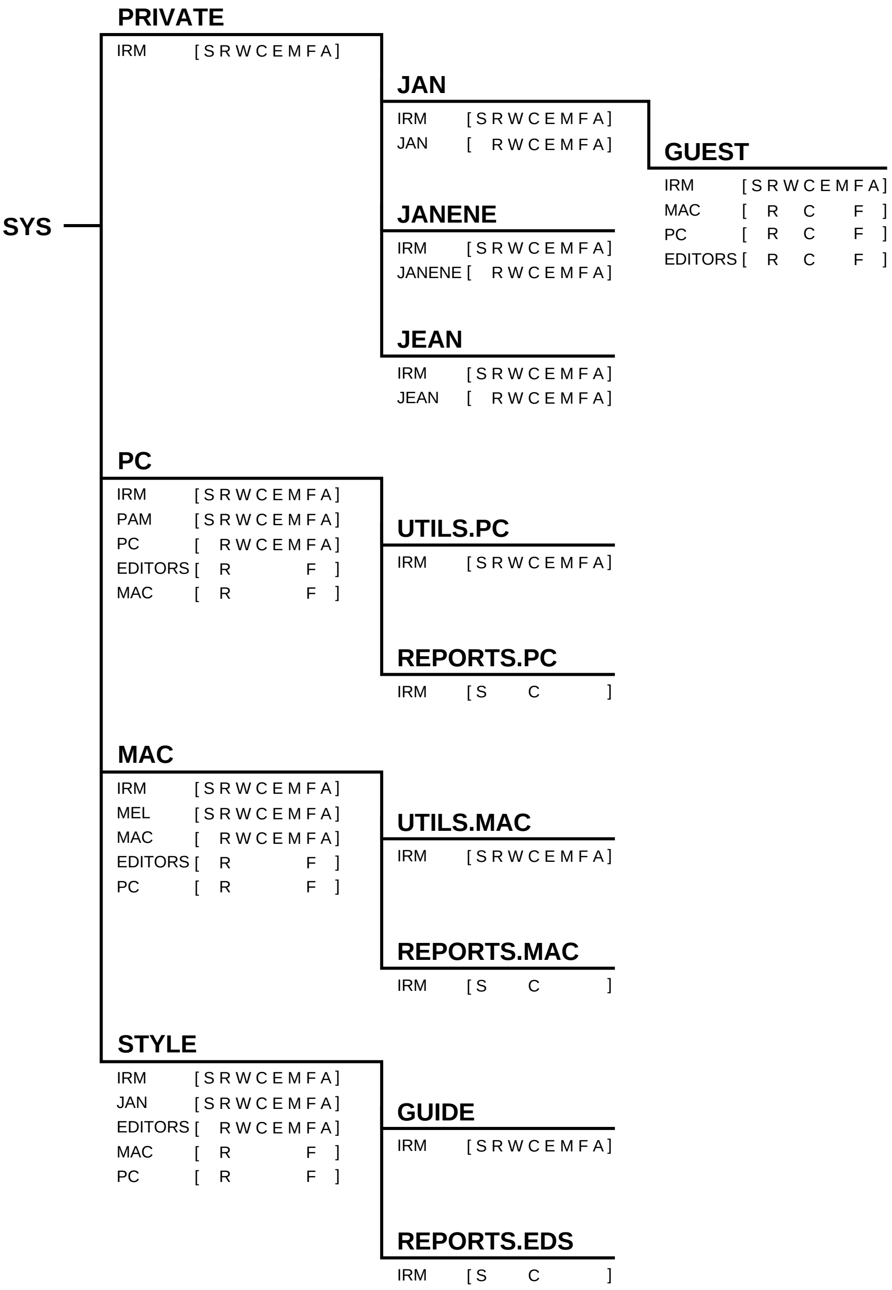

Questions. Answer these questions. Answers and explanations follow the questions.

1. What effective rights can the supervisor of the editing group, JAN, exercise in the following directories? Circle the rights that JAN has.

   SYS:PRIVATE/JEAN   S R W C E M F A

   SYS:STYLE/REPORTS.EDS   S R W C E M F A

   SYS:PC/REPORTS.PC   S R W C E M F A

   SYS:MAC/UTILS.MAC   S R W C E M F A

2. What effective rights can PETER exercise in the following directories?

   SYS:PRIVATE/JAN/GUEST   S R W C E M F A

   SYS:PC/REPORTS.PC   S R W C E M F A

   SYS:PC/UTILS.PC   S R W C E M F A

3. Can MARY, who is a member of the group MAC, copy a file from the directories listed below?

   SYS:PC/UTILS.PC   Yes   No

   SYS:MAC/REPORTS.MAC   Yes   No

4. Can MEL, who is the supervisor and a member of the group MAC, see and read files in any of the directories listed below?

   SYS:MAC/REPORTS.MAC   Yes   No

   SYS:STYLE/GUIDE   Yes   No

5. Can JEAN, who is a member of the group EDITORS, create or copy a file into the directories listed below?

   SYS:STYLE/REPORTS.EDS   Yes   No

   SYS:PRIVATE/JEAN   Yes   No

   SYS:PRIVATE/JAN/GUEST   Yes   No

Answers. The answers to the questions appear below with brief explanations.

1. What effective rights can the supervisor of the EDITORS group, JAN, exercise in the following directories?

   SYS:PRIVATE/JEAN None. JAN cannot even see the directory, because she has no rights there.

   SYS:STYLE/REPORTS.EDS All rights: S R W C E M F A. Since JAN has the Supervisory right to the STYLE directory, she has all rights to its subdirectories, regardless of the IRM.

SYS:PC/REPORTS.PC None. All JAN's rights (RF) are revoked with the IRM; therefore, JAN cannot even see this directory.

SYS:MAC/UTILS.MAC R and F. JAN has these through her membership in the EDITORS group.

2. What effective rights can PETER exercise in the following directories?.

   SYS:PRIVATE/JAN/GUEST R, C, and F. PETER is a member of the PC group, and the PC group was granted RCF rights to the directory.

   (If JAN wants all users on the file server to have the R, C and F, a more effective way is to grant the R, C, and F rights to the EVERYONE group. Since JAN has the A right, she can make this assignment herself.)

   SYS:PC/REPORTS.PC C. All other rights granted to PETER in the PC directory (through membership in the PC group) are revoked in the IRM of the REPORTS.PC directory

   SYS:PC/UTILS.PC R W C E M F. The PC group was granted these rights to the PC directory, PETER is a member of the PC group, and the IRM allows all these rights to be inherited.

3. Can MARY, who is a member of the group MAC, copy a file from either of the directories listed below?

   SYS:PC/UTILS.PC Yes. MARY is a member of the MAC group, and MAC has been granted R F rights to the directory.

   SYS:MAC/REPORTS.MAC No. Even though MARY has R F rights, those rights are revoked in the IRM of REPORTS.MAC.

4. Can MEL, who is the supervisor and a member of the group MAC, see and read files in any of the directories listed below?

   SYS:MAC/REPORTS.MAC Yes. Even though the IRM revokes the R F rights, MEL has the Supervisory right to the MAC directory and, therefore, can exercise all rights in all MAC's subdirectories.

   SYS:STYLE/GUIDE Yes. The IRM does not revoke the R F rights, MEL is a member of the MAC group, and the MAC group has R F rights to the STYLE directory.

5. Can JEAN, who is a member of the group EDITORS, create or copy a file into any of the directories listed below?

   SYS:STYLE/REPORTS.EDS Yes. The IRM does not revoke the Create right, JEAN is a member of the EDITORS group, and the EDITORS group has the Create right in the STYLE directory.

SYS:PRIVATE/JEAN Yes. This directory is JEAN's private directory, and she has all rights except Supervisory in the directory.

SYS:PRIVATE/JAN/GUEST Yes. The EDITORS group has the Create right in the directory, and JEAN is a member of the EDITORS group

Attribute Security

Attribute security assigns special properties to individual directories or files. Attributes override rights granted with trustee assignments and can prevent tasks that effective rights would allow.

For example, attributes can be used to prevent the following:

- Deleting a file or a directory
- Copying a file
- Viewing a file or a directory
- Writing to a file

Attributes are also used for the following:

- Controlling whether files can be shared
- Marking files as modified so that backup utilities can select only the files that have been modified
- Protecting files from data corruption by ensuring that either all changes are made or no changes are made when a file is being modified

If users have the Modify right for a directory or a file, they can change attributes and complete any task allowed with their effective rights.

NetWare uses the directory attributes listed in Table 14 and the directory attributes listed in Table 15.

Table 14 Directory attributes

Directory Attribute	Syntax	Description
Delete Inhibit	D	The directory can't be deleted.
Hidden	H	The directory can't be seen with a DOS DIR command. It can't be deleted.
Purge	P	Files that are deleted from the directory are purged immediately.
Rename Inhibit	R	The directory can't be renamed.
System	SY	The directory can't be seen with a DOS DIR command; it also can't be copied or deleted.

Table 15 File attributes

File Attribute	Syntax	Description
Archive Needed	A	The file has been modified since the last back up. This is DOS's archive bit.
Copy Inhibit	C	The file can't be copied. (Applies to Macintosh files only.)
Delete Inhibit	D	The file can't be deleted or copied over.
Execute Only	X	The file can't be copied or copied over. This attribute can be given only to .EXE or .COM files and can't be removed. Some programs do not execute correctly if files are flagged X.
Hidden	H	The file can't be seen with the DOS DIR command. It can't be copied or deleted.
Indexed	I	The file is indexed for faster access. (NetWare assigns I automatically to files with over 64 regular FAT entries.)
Purge	P	If the file is deleted, it is purged immediately.

File Attribute	Syntax	Description
Read Audit	Ra	Has no effect in NetWare v3.12.
Read Only	Ro	The file can only be read; it can't be written to or deleted. NetWare assigns Delete Inhibit and Rename Inhibit automatically with Read Only.
Read Write	Rw	The file can be read and written to.
Rename Inhibit	R	The file can't be renamed.
Shareable	SH	The file can be used by several users simultaneously.
System	SY	The file can't be seen with the DOS DIR command. It can't be copied or deleted.
Transactional	T	The file is protected by the Transaction Tracking System (TTS).
Write Audit	Wa	Has no effect in NetWare v3.12.

The initial letters of these attributes are displayed between brackets by NetWare utilities:

```
[Ro S A X H Sy I T P Ra Wa C D R]
```

By convention, attributes that have *not* been assigned are indicated by a blank space, as in the following:

```
[Ro S                       D R]
```

Using Attributes

Use directory and file attributes to increase security in areas where many users have access to files.

For example, NetWare utilities are flagged so that even the SUPERVISOR cannot accidentally delete them without removing flags:

- All NetWare files in SYS:SYSTEM, SYS:PUBLIC, and SYS:LOGIN are automatically flagged Ro, S, D, and R.
- The bindery files are automatically flagged Sy, H, and T.

To change or display file attributes, see "FILER"; "FLAG" in *Utilities Reference*.

To change or display directory attributes, see "FILER"; "FLAGDIR" *in Utilities Reference*.

Related utilities: "ALLOW"; "FILER"; "FLAG"; "FLAGDIR"; "GRANT"; "REMOVE"; "REVOKE"; "RIGHTS"; "SYSCON" "WHOAMI" (*Utilities Reference*).

Security equivalence

An assignment that allows one user or group to have the same rights as another.

Use security equivalence when you need to give a user access to the same information as another user. You thus avoid having to review the directory structure and determine which rights need to be assigned in which directories and to which files.

SUPERVISOR can make any user security equivalent to another user or group. Workgroup Managers and User Account Managers can assign security equivalences only within the pool of users and groups they manage.

Maintaining Security

In networks containing confidential data that only selected users have access to, take care that you do not inadvertently give a user access to restricted information through a security equivalence.

Be aware that a security equivalence gives a user who is security equivalent to another user all rights to that user's home directory. Making a user security equivalent to a group does not have that disadvantage.

Be especially cautious in making any user security equivalent to SUPERVISOR. Instead, delegate responsibilities to managers and operators.

See also “User” on page 274; “User Account Manager” on page 282; “Workgroup Manager” on page 296.

Related utility: "SYSCON" (*Utilities Reference*).

Semaphore

A flag that coordinates activities of both programs and processes to prevent data corruption in multiprocess environments.

One use of semaphores is to provide a system of file sharing and file locking.

See also "File locking" on page 106.

Serial communication

The transmission of data between devices over a single line, one bit at a time.

NetWare uses the RS-232 serial communication standard to send information to serial printers, remote workstations, remote routers, and asynchronous communication servers.

The RS-232 standard, developed by the Electronic Industries Association (EIA), enhances the delivery of information from one system to another.

A *system* can be any device or group of devices that can handle and process the data received.

For example, a printer can be thought of as a system that transforms the binary data it receives from the computer into printed text.

Parameters. The RS-232 standard uses several parameters that must match on both systems for valid information to be transferred. These parameters include baud rate, character length, parity, stop bit, and XON/XOFF:

- *Baud rate.* The signal modulation rate, or the speed at which a signal changes.

  Since most modems or serial printers attached to personal computers send only one bit per signaling event, baud can be thought of as bits per second; however, higher-speed modems may transfer several bits per signal change.

  Typical baud rates are 1200, 2400, 4800, and 9600. The higher the number, the greater the number of signal changes and, therefore, the faster the transmission.

- *Character length*. The number of data bits used to form a character.

  The standard ASCII character set (including letters, numbers, and punctuation) consists of 128 characters and requires a character length of 7 bits for transmissions.

  Extended character sets (containing line drawings or the foreign characters used in IBM's extended character set) contain an additional 128 characters and require a character length of 8 bits.

- *Parity*. A method of checking for errors in transmitted data. You can set parity to even or odd, or not use parity at all.

  Serial communication sends information in a stream of bits called a *frame*. Each frame consists of start bits, data bits, an optional parity bit, and stop bits.

  The parity bit is set to 0 or 1 so that the sum of the data bits is even or odd. Upon reception, each transmitted frame is checked to ensure that the parity is still even or odd.

  If parity is incorrect (because a bit was changed during transmission), the communications software determines that a transmission error occurred and can request that the data be retransmitted.

  Table 16 shows examples of parity checking.

Table 16 Examples of parity checking

Character length	Sample bits	Even parity	Odd parity
7 data bits	0010110	00101101	00101100
7 data bits	1110111	11101110	11101111
8 data bits	10001000	100010000	100010001
8 data bits	11011111	110111111	110111110

- *Stop bit*. A special signal that indicates the end of that character. Today's modems are fast enough that the stop bit is always set to 1. Slower modems required two stop bits.
- *XON/XOFF*. One of many methods that prevent the sending system from transmitting data faster than the receiving system can accept it.

Serial port

A port that allows data to be transmitted asynchronously, one bit at a time.

On IBM PC-compatible computers, COM1 and COM2 are asynchronous serial ports.

Serialization

The process of serializing software to prevent unlawful software duplication.

Each NetWare operating system has a unique serial number.

If two NetWare operating systems with the same serial number exist on the same internetwork, each file server displays a copyright violation warning at the server console and at each logged-in workstation.

If you suspect you have received the copyright violation warning in error, contact your Novell Authorized ResellerCM. Or, call (800) NETWARE (638-9273) (in the USA). You may be charged for the technical support to correct the errors.

If you suspect criminal copyright violation, please contact Novell's anti-piracy hotline, (800) PIRATES (747-2837) (in the USA).

Server

See “File server” on page 106; “Print server” on page 182.

Server protocol

See “NetWare Core Protocol” on page 156.

Service protocol

See “NetWare Core Protocol” on page 156.

SFT™

(System Fault Tolerance) Data duplication on multiple storage devices; if one storage device fails, the data is available from another device.

There are several levels of hardware and software system fault tolerance; each level of redundancy (duplication) decreases the possibility of data loss.

See also "Disk duplexing" on page 72; "Disk mirroring" on page 75; "Hot Fix" on page 117.

Shareable attribute

See "Attributes" on page 21; "Security" on page 221 (Attribute Security).

Shell

See "NetWare DOS Requester" on page 156.

Small Computer Systems Interface

See "SCSI" on page 218.

SMS

See "Storage Management Services" on page 253.

Software interrupt

See "Interrupt" on page 121.

Source routing

IBM's method of routing data across source-routing bridges. NetWare source routing programs allow an IBM Token-Ring™ network bridge to forward NetWare packets (or frames).

IBM bridges can be configured as either single-route broadcast or all-routes broadcast. The default is single-route broadcast.

- *Single-route broadcasting*. Only designated single-route bridges pass the packet and only one copy of the packet arrives on each ring in the network.

  Single-route bridges can transmit single-route, all-routes, and specifically routed packets.

- *All-routes broadcasting*. Sends the packet across every possible route in the network, resulting in as many copies of the frame at the destination as there are bridges in the network.

  All-routes bridges pass both all-routes broadcasts and specifically routed packets.

To support IBM hardware and applications, Novell provides ROUTE.NLM for the file server and ROUTE.COM for DOS workstations (or ROUTE.SYS for OS/2 workstations).

These drivers allow users running NetWare v3.12 to communicate across IBM Token-Ring network bridges. They also allow IBM applications that require source routing support to run unmodified on NetWare networks.

Parameters for ROUTE.NLM determine which packets are broadcast as all-routes packets and which as single-route packets.

At the workstation, the ROUTE.COM file determines the type of packets the workstation broadcasts. It executes after IPXODI.COM and before the NetWare DOS Requester.

For more information on the ROUTE.NLM driver and the workstation driver, refer to *System Administration*.

NOTE: NetWare provides two types of drivers for source routing: a dedicated IPX driver generated with WSGEN and an ODI driver. If you want to use a protocol other than IPX (such as TCP/IP), you need to load the Token Ring ODI driver. For information on this driver, refer to *Workstation for DOS and Windows*.

Sparse file

A file with at least one empty block. (NetWare won't write any block that is completely empty.)

Databases often create sparse files. For example, suppose the disk allocation block size for volume VOL1: is 4 KB. Also suppose that a database opens a new file, seeks out the 1,048,576th byte, writes 5 bytes, and closes the file.

An inefficient operating system would save the entire file to disk.

The file would be composed of 256 zero-filled disk allocation blocks (the first 1 MB) and one more disk allocation block with 5 bytes of data and 4,091 zeros. This method would waste 1 MB of disk space.

However, NetWare writes only the last block to disk, saving time and disk space.

Sparse files aren't limited to large files. If a two-block file is created and the first block is empty, the operating system treats the file as a sparse file.

If a program tries to read from a file's empty blocks, the operating system generates and returns zeros.

The NetWare NCOPY command doesn't write to sparse files automatically. The /F option in NCOPY forces the operating system to write to sparse files.

Spool

To transfer data that was intended for a peripheral device (such as a printer) into temporary storage.

The data can be transferred to the peripheral later without affecting or delaying the operating system as it performs other operations.

NetWare v3.12 uses CAPTURE to spool data.

SPX.COM

(Sequenced Packet Exchange) A Novell communication protocol that monitors network transmissions to ensure successful delivery.

SPX verifies and acknowledges successful packet delivery to any network destination by requesting a verification from the destination that the data was received.

The SPX verification must include a value that matches the value calculated from the data before transmission. By comparing these values, SPX ensures that the data packet arrived at its destination intact.

SPX can track data transmissions consisting of a series of separate packets. If an acknowledgment request brings no response within a specified time, SPX retransmits the series.

After a reasonable number of retransmissions fail to return a positive acknowledgment, SPX assumes the connection has failed and warns the operator.

SPX is derived from Novell's IPX using the Xerox® Sequenced Packet Protocol.

See also “Communication protocols” on page 41; “IPX internal network number” on page 122; “NetWare Runtime” on page 161.

Station

See “Network station” on page 170.

Station address

See “Network numbering” on page 165.

Stop bit

See “Serial communication” on page 246.

Storage Management Services

(SMS) Services that allow data to be stored and retrieved. SMS is independent of backup/restore hardware and file systems (such as DOS, OS/2, Macintosh, Windows, or UNIX).

SMS Architecture

SMS provides NLMs and other software modules that run on file servers. Modules used in SMS are

- *SBACKUP.* Provides backup and restore capabilities.
- *SMDR (Storage Management Data Requester).* Passes commands and information between SBACKUP and Target Service Agents (TSAs).
- *Device drivers.* Control the mechanical operation of storage devices and media.
- *NetWare Server TSAs*. Pass requests for data (generated within SBACKUP) to the target file server where the data resides, then return requested data through the SMDR to SBACKUP.
- *Database TSAs.* Pass commands and data between the host server (where SBACKUP resides) and the target database where the data to be backed up resides, then return the requested data through the SMDR to SBACKUP.
- *Workstation TSAs.* Pass commands and data between the host server (where SBACKUP resides) and the target station where the data to be backed up resides, then return the requested data through the SMDR to SBACKUP.
- *Workstation Manager.* Receives "I am here" messages from stations available to be backed up. It keeps the names of these stations in an internal list.

Although all SMS software modules are loaded on the host server, all modules aren't used in all backups.

For information about loading SMS NLMs and drivers, see *Server Backup*.

See also “Backup hosts and targets” on page 29; “NetWare Loadable Module” on page 158; “Target Service Agent” on page 259.

STREAMS

NLMs that provide a common interface between NetWare and transport protocols (such as IPX/SPX™, TCP/IP, SNA, and OSI) that need to deliver data and requests to NetWare for processing.

You need to load STREAMS if you are using an application that requires the C Library (CLIB) interface. Check the application's documentation.

By making the transport protocol transparent to the network operating system, NetWare STREAMS allows the same set of services to be provided across the network, no matter what transport protocols are used.

Network managers can install the protocols of their choice or change the protocols used without affecting the level of services delivered to the user.

NetWare v3.12 implements STREAMS as the following NLMs.

- STREAMS.NLM includes the STREAMS application interface routines, the utility routines for STREAMS modules, the log device, and a driver for the Open Data-Link Interface.
- SPXS.NLM provides access to the SPX protocol from STREAMS.
- IPXS.NLM provides access to the IPX protocol from STREAMS.
- CLIB.NLM is a library of functions that some loadable modules use.
- TLI.NLM (Transport Level Interface) is an application programming interface that sits between STREAMS and user applications, allowing interface with transport level protocols such as IPX/SPX or TCP/IP.

For applications that communicate via IPX/SPX through STREAMS, you must load the modules in the following order.

1. STREAMS.NLM
2. IPXS.NLM and SPXS.NLM
3. CLIB.NLM
4. TLI.NLM (Load only if the application you are using needs it. Check the application's documentation.)

Figure 69 shows how STREAMS operates.

Figure 69 How STREAMS functions

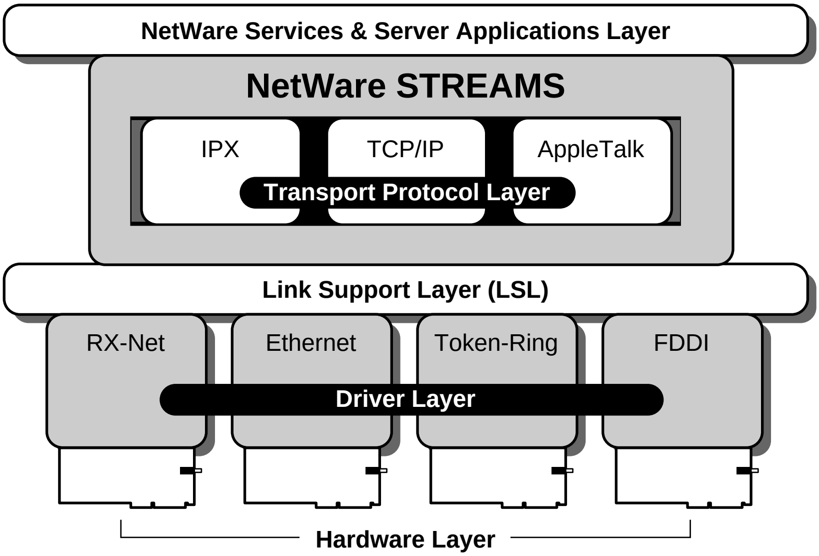

Subdirectory

Any directory below another in the directory structure. For example, in SYS:ACCTS/RECEIVE, RECEIVE is a subdirectory of SYS:ACCTS.

See also "Directory structure" on page 61.

SUPERVISOR

The username for the network supervisor or system administrator present in the bindery as a bindery object when the file server is first brought up.

User SUPERVISOR

- Cannot be deleted or renamed and is assigned ID number 1.
- Has all rights to all directories; these rights cannot be revoked.
- Has no initial password so that the network supervisor can log in to the file server to set up the network environment.

Delegating Responsibility

As necessary, the network supervisor can delegate system administration responsibilities to Management Information System (MIS) staff or organization managers.

After users are created on the file server, they can be designated as one of the following:

- *Workgroup Manager.* An assistant supervisor with rights to create bindery objects (such as users and groups) and manage the user accounts.
- *User Account Manager.* A user with rights to manage user accounts, but no rights to create new bindery objects.

  Workgroup Managers are User Account Managers for users they create. Existing users can be assigned to a Workgroup Manager (or any other user) for account management.
- *File server console operator.* A file server supervisor with rights to use the FCONSOLE utility.
- *Print server operator.* A printing supervisor with rights to manage the print server.
- *Print queue operator.* A printing supervisor with rights to create, manage, disable, and enable print queues.

Security

Because SUPERVISOR has absolute authority over users and information, use caution in making anyone security equivalent to SUPERVISOR.

Instead, delegate supervisor responsibilities by designating operators and managers.

For *maximum security*, we suggest the following.

- Set a password in SYSCON the *first* time you log in as SUPERVISOR. Change this password regularly and do not select an obvious password such as a spouse's name, etc.
- Create a second username to use when you work on the network as a regular user. Log in as SUPERVISOR only when you have supervisor tasks to perform.
- Set up a "back door" to the system in two ways:

- Create a fictitious user with an unpredictable username and make that user security equivalent to SUPERVISOR.
- Create a fictitious user with an unpredictable username and make that user the User Account Manager for both the fictitious user and SUPERVISOR.
- Use caution in making anyone security equivalent to SUPERVISOR. Instead, delegate supervisor jobs by designating operators and managers.

See also “File Server Console Operator” on page 108; “Print queue operator” on page 182; “Print server operator” on page 182; “User” on page 274; “User Account Manager” on page 282; “Workgroup Manager” on page 296.

Related utility: "SYSCON" (*Utilities Reference*).

Supervisory right

See “Rights” on page 207; “Security” on page 221 (Rights Security).

Surface test

A test in the NetWare INSTALL program that lets you test the NetWare partition on a hard disk for bad blocks.

The surface test can run in the background on one or more dismounted hard disks so that you (or other users) can work on mounted volumes on other hard disks.

You can choose either a *destructive* or a *nondestructive* surface test:

- *Destructive test.* Acts like a disk format—it destroys data as it makes several passes over the disk surface, reading and writing test patterns.
- *Nondestructive test.* Prereads and saves existing data while it reads and writes test patterns to the hard disk. Then the program writes the data back to the disk.

Neither surface test can be executed unless the volumes on the hard disk are dismounted.

Surge protectors/ suppressors

See “Power conditioning” on page 180.

Switch block

A set of switches mounted together to form a single component.

In some file servers, a switch block is used to control system configuration data, such as type of monitor, amount of memory, and number of drives.

Network boards often use switch blocks to set system addresses (such as station, base I/O, and base memory addresses).

Synchronous transmission

See “Serial communication” on page 246.

System attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

SYSTEM directory

A directory (SYS:SYSTEM) created automatically on the SYS: volume during network installation.

Copy the NetWare supervisor-only utilities into this directory during installation. This directory cannot be deleted.

System Fault Tolerance

See “SFT^{TM}” on page 249.

System supervisor

See “SUPERVISOR” on page 255; “User” on page 274.

19 T

Tape backup unit

Typically, an external tape drive that backs up data from hard disks.

Target Service Agent

(TSA) A program that processes data moving between a specific target and SBACKUP. SBACKUP, running on the host, sends requests to the TSA, which

- Receives the commands from SBACKUP and processes them so that the target operating system can handle the request for data.
- Passes the data request from SBACKUP to the target.
- Receives the requested data from the target and returns it to SBACKUP in standard SMS format.

Servers and workstations running different software releases, or having different operating systems, require NetWare-compatible TSAs to communicate with SBACKUP.

See also “Backup hosts and targets” on page 29; “Storage Management Services” on page 253.

Terminating resistor

A grounding resistor placed at the end of a bus, line, chain, or cable to prevent signals from being reflected or echoed. Sometimes shortened to "terminator."

Termination

See "SCSI bus" on page 218.

Topology

The physical layout of network components (cables, stations, gateways, hubs, etc.).

There are three basic topologies:

- *Star network*. Workstations are connected directly to a file server but not to each other.
- *Ring network*. The file server and workstations are cabled in a ring; a workstation's messages may have to pass through several other workstations before reaching the file server.
- *Bus network*. All workstations and the file server are connected to a central cable (called a *trunk* or *bus*).

Transaction Tracking System

(TTS) A system that protects database applications from corruption by backing out incomplete transactions that result from a failure in a network component.

When a transaction is backed out, data and index information in the database are returned to the state they were in before the transaction began.

TTS is an integral part of NetWare v3.12; it is not an optional feature as in earlier versions. However, it can be turned on and off.

NOTE: Even if you don't plan to add a multiuser database to your file server, TTS is valuable because it protects the bindery and the queuing database files from corruption.

Benefits of TTS

Mainframe, minicomputer, and network database systems have offered transaction backout capability for some time.

Usually, this capability is implemented as part of the database application software and not as part of the operating system.

NetWare TTS is implemented at the operating system level on the file server. This method provides two major advantages over application-level implementation.

Transaction tracking performance improves with TTS in the file server. Performance increases dramatically because transactions are tracked in the file server, where file writes are made; less data is transferred across the network; and all transactions benefit from the speed of NetWare's disk caching system.

Database applications not designed for transaction backout can have backout capability. When a database application without backout capability makes a physical record lock to a database file or a logical record lock to an open database, the application "implies" that it is making a transaction.

At this point, TTS begins tracking this implicit transaction so that the transaction can be backed out if a failure occurs.

When a database application without backout capability releases physical or logical record locks, TTS can "infer" that the application has completed a transaction. At this point, TTS stops tracking the transaction.

Three kinds of database applications benefit from TTS:

- Applications designed with no transaction backout capability (implicit transactions)
- Applications that have built-in transaction backout capability (such as Btrieve®)
- Applications that use explicit NetWare TTS calls to provide transaction backout capability (such as BEGIN, ABORT, and END)

TTS Protection

A transaction on a network can be saved improperly in any of the following situations:

- Power to a server or a workstation is interrupted during a transaction.
- Server or workstation hardware fails (for example, a memory parity error or a network board error) during a transaction.
- A server or a workstation hangs (a software failure) during a transaction.
- A network transmission component (such as a hub, a repeater, or a cable) fails during a transaction.

TTS protects data from failure in all these cases by making a copy of the original data in the transaction before it is overwritten by new data.

If a failure occurs during the transaction, TTS can back out the transaction and restore the original data.

- If the file server fails, TTS backs out the transaction when the file server comes up again.
- If a workstation or network transmission component fails, TTS backs out the transaction immediately.

TTS can protect against these types of failures for any type of application that issues record-locking calls and stores information in records, including traditional databases, some electronic mail applications, and some workgroup appointment schedulers.

Files such as word-processing files that are not organized into discrete records are not protected by TTS.

How TTS Operates

The purpose of TTS is to guarantee that all transaction changes to a file are either wholly completed or not made at all.

To track transactions on a given file with TTS, you must flag the file as Transactional.

NOTE: A file flagged Transactional cannot be deleted or renamed.

When a workstation begins a transaction in a database file, TTS follows four basic steps to maintain the integrity of the file:

1. TTS makes a copy of the original data so that if the transaction fails, the original data can be restored to the database file.

   The copy of the original data goes in a file external to the database file. This external file contains all transaction backout information; only the operating system uses it.

2. TTS writes the changed data to the database file after the copy of the original data is written to the backout file.

3. TTS repeats Steps 1 and 2 for additional changes within the transaction (a single transaction can consist of a sequence of changes).

4. TTS waits until all changed data in the transaction is written to disk. Then TTS writes a record to the backout file and indicates that the transaction is complete.

   If TTS marks a transaction complete, the transaction is not backed out if a failure occurs.

NetWare TTS supports all types of transactions: single transactions; multiple, concurrently active transactions from either different workstations or different tasks within a single workstation; or both.

Record-locking thresholds. Since some database applications leave one or more records locked at all times (for copy protection), TTS allows you to set a locking threshold in the workstation so that an implicit transaction is not unnecessarily tracked when the application is started and the first record lock occurs.

Setting the locking threshold prevents having an entire database session tracked as a single transaction or having too many transactions per file update.

Related utility: "SETTTS" in *Utilities Reference*.

Two States of TTS

A NetWare v3.12 file server goes through a TTS enabling routine each time it boots. But TTS can also be disabled.

Enabled. When volume SYS: is first mounted, TTS is enabled if enough disk space and memory are available to allow transaction tracking.

If TTS is ever disabled, you can enable it with the ENABLE TTS command. (If a problem occurs that disables TTS automatically, you must fix the problem before you can enable TTS.)

Disabled. If volume SYS: is full or the file server does not have enough free memory, TTS is disabled immediately after the file server boots or after any server process requests more memory or disk space than the file server has.

You can disable TTS indirectly by dismounting SYS: (the TTS backout volume) or directly by issuing the DISABLE TTS command. When TTS is disabled, workstations can make transactions in a database, but TTS does not protect the database.

TTS Handles Special Backout Cases

In addition to handling routine backout tasks, TTS can back out file truncations or extensions and multiple changes to the same data area during a single transaction.

TTS can even back out interrupted transaction backouts. (An interrupted transaction backout can occur if the file server fails while it is in the middle of backing out transactions from a previous file server failure.)

TTS has been enhanced to allow for the increased flexibility in volume mounting and dismounting that NetWare v3.12 provides:

- If SYS: is mounted *after* other volumes, TTS scans the backout file on SYS: at mounting time to determine whether to back out transactions on the volumes mounted before SYS:.

  NOTE: Since TTS is not enabled until SYS: is mounted, you should never back out transactions made on volumes where transactions were made following a file server failure but before SYS: was mounted. Backing out older transactions over the more recent transactions can corrupt your database.

- If SYS: is mounted *before* other volumes, TTS scans the backout file on SYS: as each additional volume is mounted (whether immediately after SYS: or later). If TTS finds transactions that need to be backed out, it initiates the regular backout process.

Example. TTS holds all workstation record locks until a transaction is completed. This prevents the disaster that would result from the following situation:

- The application in Workstation #1 releases a lock on a record before the transaction is completed (written to disk).
- Workstation #2 locks and changes the same record in file server cache memory (also before the first transaction is completed).
- Workstation #1 fails and Workstation #2 completes its transaction (it is written to disk).
- Because Workstation #1 failed, its transaction is backed out—in this case, over the transaction that Workstation #2 completed.

If NetWare TTS did not hold workstation records locks until transaction completion, the database in this situation would not contain the correct information, since the latest transaction (from Workstation #2) would have been incorrectly overwritten with data that existed before the failed transaction (from Workstation #1).

Related utilities: "DISABLE TTS"; "ENABLE TTS"; "MONITOR"; "SET" (*System Administration*); "FILER"; "FLAG"; "SETTTS" (*Utilities Reference*).

Routine Maintenance for TTS: the TTS$LOG.ERR file

The TTS$LOG.ERR file, located in the root directory of the SYS: volume, contains TTS status information that the system administrator can use in troubleshooting. (Refer to "Troubleshooting" in *System Administration*.)

TTS writes the status information to this file when TTS is initialized or disabled and when transactions are backed out. This file is a valuable troubleshooting tool for the system administrator, and TTS never deletes it or decreases its size.

The system administrator should monitor the size of this file and delete it completely or partially when necessary. TTS recreates this file when it needs to.

TTS$LOG.ERR grows more quickly if the "TTS Abort Dump Flag" is set to ON and transaction backouts occur. (If the "TTS Abort Dump Flag" is set to ON, a copy of the original data from the transaction and a copy of the data before the backout is written to TTS$LOG.ERR when a backout occurs.)

The file grows in proportion to the size and number of transactions backed out.

See also "SET" in *System Administration*.

Transactional attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security); “TSA resources” on page 268.

Trustee rights

Permissions that control which directories and files a user or group can access and what tasks the user or group can do with them.

A trustee assignment consists of the rights assigned to a user or group. A user or group that has been assigned rights to work in a directory or file is known as a "trustee" of that directory or file.

Automatic Assignment

If you make a trustee assignment in a directory, the trustee has access to the directory, its files, and its subdirectories (unless the rights are redefined at the file or subdirectory level).

In other words, trustee rights "flow down" through the structure unless

- Other trustee assignments are granted at a lower level of the directory structure;
- The Inherited Rights Mask of a subdirectory or file revokes rights assigned in a trustee assignment.

Default Rights

When you make a trustee directory assignment, the default rights (Read and File Scan) allow a trustee to read the files in the directory and to see the subdirectories and files in the directory.

Any trustee assignment, whether for a directory or a file, also includes the right to see the path leading from the root to that directory or file.

A new assignment of trustee rights at the file level can revoke rights assigned at the directory level or allow additional rights.

SUPERVISOR's Trustee Rights

The user SUPERVISOR has all rights to all directories and files and can assign any of these rights to users and groups.

A trustee must have the Access Control right [A] to make trustee assignments in a directory or file.

Rights Security

Rights security is controlled by both trustee assignments and the Inherited Rights Mask (IRM). Both the trustee assignment and the IRM use the same rights.

When you grant a trustee assignment, it takes precedence over the IRM in the current directory. However, in the subdirectories, the Inherited Rights Mask takes precedence, unless new trustee assignments are granted.

Some common tasks and the rights required to do them are listed in Table 17.

Table 17 Rights required for tasks

Task	Rights required
Read from a closed file	Read
See a filename	File Scan
Search a directory	File Scan
Write to a closed file	Write, Create, Erase, Modify
Create and write to a file	Create

Task	Rights required
Copy files into a directory	Create
Remove an empty subdirectory	Erase
Delete a file	Erase
Change directory or file attributes	Modify
Rename a file	Modify
Change the Inherited Rights Mask	Access Control
Change trustee assignments	Access Control
Modify a directory's disk space assignment for users	Access Control

To grant or modify trustee assignments for either directories or files, use "ALLOW"; "FILER"; "GRANT"; "REMOVE"; "REVOKE"; "SYSCON" (*Utilities Reference*).

See also “Rights” on page 207; “Security” on page 221.

TSA

See “Target Service Agent” on page 259.

TSA resources

Categories of data, referred to as *major resources* and *minor resources*, created by each TSA.

Because these resources vary with each TSA, SBACKUP processes these resources in different ways.

See also “Backup” on page 25; “Major resource” on page 139; “Minor resource” on page 144; “Target Service Agent” on page 259.

TTS

See “Transaction Tracking System” on page 260.

Turbo FAT index table

See “File Allocation Table” on page 98.

20 U

Unbinding

See "Binding and unbinding" on page 31.

Uninterruptible power supply

(UPS) A backup power unit that supplies uninterrupted power if a commercial power outage occurs.

Types of UPS are online and offline:

- *Online UPS*. Actively modifies the power as it moves through the unit. If a power outage occurs, the unit is already active and continues to provide power.

  An online UPS is usually more expensive than an offline UPS, but provides a nearly constant source of energy during power outages.

- *Offline UPS*. Monitors the power line. When power drops, the UPS is activated.

  The drawback to this method is the slight lag before the offline UPS becomes active. However, most offline UPS systems are fast enough to offset this lag.

Because UPS systems can be expensive, most companies attach them only to the most critical devices, such as file servers, routers, and hard disk subsystems.

Attaching a UPS to a server enables the server to close files and rewrite the system directory to disk.

Unfortunately, most programs run on the workstation and data stored in RAM is not saved during a power outage unless each station has its own UPS.

If the UPS doesn't have its own form of surge protection, install a surge protector.

See also "Power conditioning" on page 180; "UPS monitoring" on page 272.

Unloading

See "Loading and unloading" on page 126.

UPS

See "Uninterruptible power supply" on page 271.

UPS monitoring

The process a file server uses to ensure that an attached UPS (uninterruptible power supply) is functioning properly.

A Novell-approved UPS should be attached to each file server in a NetWare network to provide backup power. (You can also attach a UPS to workstations on the network without installing UPS monitoring hardware on the workstations.)

When a power failure occurs, NetWare notifies all current users. After a timeout specified with UPS TIME (*System Administration*), the file server logs out remaining users, closes open files, and shuts itself down.

If you install a Novell-approved UPS, you must also install or set a board in the file server to monitor the UPS. (If you have a microchannel file server, the UPS is monitored through the mouse port and does not require a board.)

The following boards can be used as a UPS monitor board (however, the disk coprocessor board is normally used for UPS monitoring):

- Software-serialized (SS) keycard
- Disk coprocessor board (DCB, 37-pin version)
- Disk coprocessor board with floppy disk controller (DCB/FDC)

- Enhanced disk coprocessor board (EDCB)
- Standalone UPS monitor board

To determine if your disk coprocessor board functions as a UPS monitor board, check the board's mounting bracket. If the board has a stereo phone jack on the mounting bracket, it has UPS monitoring capability.

Use the chart in Figure 70 to help you determine what boards can be set for UPS monitoring.

Figure 70 Boards for UPS monitoring

File server	Disk Coprocessor board	NetWare version	Serialization device (keycard)	UPS monitoring
micro channel	DCB/2	As drivers are available	None	Mouse port
	No DCB/2	v2.12 and above	None	Mouse port
ISA (AT bus)	Enhanced DCB	v3.0 and above	None	EDCB
	AT - compatible DCB (and DCB/FDC for Novell 386AE file servers)	v2.12 and above	None	DCB or DCB/FDC
	AT - compatible DCB (no UPS monitoring)	v2.12 and above	None	Standalone UPS monitor board
	No EDCB, DCB, or DCB/FDC	v2.12 and above	None	Standalone UPS monitor board

See also “Power conditioning” on page 180; “Uninterruptible power supply” on page 271.

Related utilities: "UPS STATUS"; "UPS TIME" (*System Administration*).

User

An identify created on the network that represents a person with access to the network.

Once a username exists as an object in the file server bindery, the user can log in to the file server with that username and access the network.

Although any number of users can be created, only 250 users can be logged in simultaneously to a file server running NetWare v3.12.

You can simplify user access to file server resources (such as applications, printers and print queues, or directories) by creating groups on the file server and giving groups access to resources.

Once a group is created, you can assign users; as members of the group, users inherit the group's privileges.

System-Created Users and Groups

Two usernames and a group exist in the file server bindery when the file server is first installed.

SUPERVISOR is the username for the network supervisor or system administrator.

SUPERVISOR has all rights to all directories. These rights cannot be revoked, and SUPERVISOR cannot be deleted or renamed.

For more information and security tips, see "SUPERVISOR" on page 255.

GUEST is the username for anyone who needs temporary and restricted access to the file server.

GUEST is automatically a member of the group EVERYONE, and GUEST's rights flow from membership in that group. You can change GUEST's rights or delete GUEST.

For more information and security tips, see "GUEST" on page 112.

EVERYONE is the group that includes all users.

EVERYONE is automatically assigned Read and File Scan rights to the SYS:PUBLIC directory and the Create right in SYS:MAIL. You can delete EVERYONE or users from EVERYONE or change EVERYONE's trustee rights.

For more information and security tips, see "EVERYONE" on page 93.

Other Users and Groups

When you set up the network environment, you can define

- *Regular network users.* The persons for whom you create usernames and user accounts.
- *Groups.* Collections of network users who share applications, perform similar tasks, or have similar needs for information.
- *Managers and operators.* Users to whom you assign certain supervisor rights and responsibilities for system administration.

See "User account" on page 275 for detailed information about these types of users.

User account

A set of restrictions, privileges, and rights that, with the username, comprise a user's identity on the network.

Each user has a user account. User accounts are part of network security and also control the user environment.

Some features of user accounts are *automatically* assigned to each user, some *must* be created or assigned, and some are *optional*.

The following features associated with user accounts are described on the following pages.

"Usernames" on page 276
"Group Membership" on page 276
"Home or Username Directories" on page 277
"Trustee Assignments" on page 277
"Security Equivalences" on page 278
"ID Numbers" on page 278
"Mailboxes" on page 278

Usernames

About usernames. The username is also the login name and must be supplied when logging in.

Usernames can be up to 47 characters long and usernames usually follow one of these formats.

- Given name (for example, user JANE and user RICHARD)
- Surname (for example, user DOE and user SMITH)
- Initials and surname (for example, user JDOE and user RDSMITH)

Viewing users. You can view a list of users with SYSCON's "User Information" option. You can also view the user's account features and restrictions. You can view logged-in users with SESSION.

Group Membership

About groups. Users are automatically assigned to the group EVERYONE and inherit the rights assigned to EVERYONE.

Other groups are created in SYSCON as empty sets and then users are assigned or added. Group members inherit the rights assigned to the group.

Viewing group information. By selecting a particular username in SYSCON, you can view a list of "Groups Belonged To." You can also view a list of existing groups with SYSCON's "Group Information" option.

When you select the name of a particular group from the "Group Information" list, you can view the group's "Member List."

Home or Username Directories

The home or username directory serves as personal "workspace." These directories are optional.

If you want each user to have a home or username directory, you should plan a parent directory (such as SYS:HOME or SYS:USERS) for these directories. You may want to set aside a separate volume for home directories.

For login script purposes, each user's home directory name should be the same as that user's username (for example, SYS:USER/JANE or SYS:HOME/RDSMITH).

If you grant all trustee rights to users in their home directories, each user can control access to files in that directory. (But users cannot restrict SUPERVISOR.)

Users can use their home directories to create files and work on projects without allowing other users to have access to their work. Once the work is completed, the files can be copied to a work or project directory where other users have rights to access the information.

The USERDEF utility creates home directories in volume SYS: when users are created unless you create and specify a parent directory for username directories.

When username directories are created by the system, trustee assignments are also made automatically.

Trustee Assignments

Trustee assignments allow users to have access to specific directories and files. You can assign users trustee rights to specified directories and files.

You can view trustee directory assignments or trustee file assignments for any user by selecting that username in SYSCON under the "User Information" option.

Security Equivalences

About security equivalences. With a security equivalence, a user can exercise rights equivalent to those of another user.

Assigning a security equivalence is convenient when you need to give a user access to the same information another user has access to.

In networks containing confidential data that only selected users have access to, take care that you do not inadvertently give a user access to restricted information through security equivalences.

Use caution in assigning SUPERVISOR equivalence. Delegate responsibilities to managers and operators instead.

Viewing security equivalences. You can view security equivalences by selecting a username in SYSCON under the "User Information" option.

ID Numbers

ID numbers are random, hexadecimal numbers assigned by the file server to each bindery object (including users).

As user SUPERVISOR, you can view an ID number for a user by selecting the user in SYSCON and then selecting the "Other Information" option.

Mailboxes

Mailboxes in the SYS:MAIL directory contain user login scripts and print job configurations.

Each user automatically gets a mailbox directory named with the user's ID number. Users receive all but the Supervisory right in this directory.

Because users' login scripts are stored in their mailbox directories, do not delete SYS:MAIL or any of its subdirectories, even if you are not using electronic mail.

User Login Scripts

Login scripts are configurable batch files that customize the network environment for individual users by initializing environment variables, mapping drives, and executing other commands.

You can view a user's login script by selecting the user in SYSCON and then selecting the "Login Script" option.

If you allow users to change their own passwords, they are also automatically allowed to change their own login scripts.

Print Job Configurations

Print job configurations define how a print job is printed.

Each user can use printing defaults, or you can create print job configurations in PRINTCON and copy them from one user to others.

Print job configurations are stored in each user's mailbox directory in SYS:MAIL.

Account Management

SUPERVISOR manages the accounts of all users.

If users are created by a Workgroup Manager, then the Workgroup Manager can manage these user accounts.

You can assign existing users to a User Account Manager, who can be a Workgroup Manager or any other user. You can also have more than one manager for a user account.

You can determine which user or group manages an existing user's account by selecting the user in SYSCON and then choosing the "Managers" option.

If a user (or group) manages other user accounts, you can also view a list of managed users and groups in SYSCON's "User Information" by selecting the manager's username and then selecting "Managed Users and Groups."

User Account Restrictions

Login restrictions are assigned at the account level to make it difficult for unauthorized users to access the file server. When certain limits are exceeded, an account is disabled.

When an account is disabled, no one can log in to the file server under that username.

You can view information about a user's account restrictions by selecting the user in SYSCON and then selecting the type of restriction you want to view information about.

You can restrict logins in the following ways.

- *Account balance.* If you have installed accounting, you can assign initial account balances for users and specify credit limits. When the account balance is depleted, the account is disabled.
- *Expiration.* You can specify an expiration date for a user account. The account expires at 12:01 a.m. the following day. Any attempt to log in after the expiration date disables the account.
- *Password.* You can require a password. You can also specify the following:
- Minimum length of passwords (default: five characters)
- How often the password must be changed (default: 40 days)
- Whether the password must be unique (default: No)

  Unique passwords are different from the previous 10 passwords for the account.
- Whether the user can change the password (default: Yes)
- The number of times a user can log in with an expired password
- The number of incorrect login attempts allowed (default: six times)

  When that number is exceeded, the account is disabled.

Disk space restrictions. You can limit the amount of disk space for each user by specifying the maximum number of blocks available for each user per volume.

Connection restrictions. You can limit the number of workstations a user can be logged in from concurrently by specifying the maximum number of concurrent connections permitted.

Time restrictions. You can restrict the hours during which users can log in. Hours are specified in half-hour blocks. You can assign all users the same hours, or you can restrict users individually.

Station restrictions. You can restrict the physical locations that a user can log in from by specifying the network and node addresses of workstations the user is permitted to log in from.

Station restrictions cannot be set with system default restrictions; they must be assigned individually.

Utilities for Creating Users

You can create network users with the SYSCON, MAKEUSER, and USERDEF utilities. (Groups can be created only with SYSCON.)

Use SYSCON to modify accounts for existing users.

SYSCON. You can create users individually with the "User Information" option.

By setting the "System Default Restrictions," you can automatically assign all login restrictions except station restrictions to new users.

Any features of user accounts not automatically created or created by defaults must be individually created or assigned.

MAKEUSER. You can create multiple users by creating and processing a .USR script with the MAKEUSER utility.

A .USR script is a list of commands in a text file that specify how to create users and user account features (such as group memberships, passwords, and time restrictions).

You can create a .USR script by using the cut-and-paste features of a text editor or by typing the script in the MAKEUSER script box. All commands in the .USR script must be entered in a specific order.

USERDEF. You can create multiple users with similar characteristics with the USERDEF utility. Instead of creating a script, you use a template that allows you to specify parameters.

When you have entered all the usernames, USERDEF creates a temporary .USR script and the MAKEUSER utility processes it.

When users are created with the defaults provided in USERDEF, each new user is provided a basic login script with these essential drive mappings:

- The first search drive is mapped to SYS:PUBLIC.
- The second search drive is mapped to the appropriate DOS directory.
- The first network drive is mapped to the user's default directory.

Although these mappings make it possible to work on the network right away, they are more appropriate for a system login script.

USERDEF also supplies new users with an initial password and SUPERVISOR's print job configurations so that new users can begin work on the network.

You can also customize the way USERDEF creates new users by creating a custom template. To do this, set new default parameters or edit the basic login script.

We recommend that you use USERDEF only in two instances:

- If you must get large numbers of new users on the network quickly, use the default template for users.
- If you want to create users in batches, but do not want to plan MAKEUSER scripts, create a custom template for users.

See also “Account restrictions” on page 15; “Group” on page 111; “User Account Manager” on page 282; “Workgroup Manager” on page 296.

Related utility: "MAKEUSER"; "SESSION"; "SYSCON"; "USERDEF" (*Utilities Reference*).

User Account Manager

A user or group that has rights to manage some users and groups.

Existing users can be assigned to a Workgroup Manager or to any other user or group for account management. Workgroup Managers automatically become User Account Managers for users they create.

Even though account management is delegated, SUPERVISOR still has all rights to manage user accounts.

User Account Manager Tasks

User Account Managers can do the following:

- Delete managed users and groups.
- Assign a managed user to a managed group.
- Assign a managed user as User Account Manager.
- Modify all options in SYSCON's "User Information" submenu not requiring file rights, including the following:

  Account Balance
  Account Restrictions
  Change Password
  Full Name
  Groups Belonged To
  Login Script
  Managed Users and Groups
  Managers
  Security Equivalences
  Station Restrictions
- With assigned file rights, make or modify
- Trustee directory assignments;
- Trustee file assignments;
- Volume restrictions;
- Disk space restrictions.

User Account Managers cannot

- Create users or groups;
- Assign users to a group the manager does not manage;
- Modify the login restrictions of their own accounts (unless they are also assigned management of their own accounts).

Why Have User Account Managers?

If you assign experienced users as User Account Managers, especially in a large system, they can handle routine tasks such as assigning users to new groups or changing drive mappings (in individual login scripts), account restrictions, account balances, etc.

A user can be managed by more than one user account manager, so you can have a backup if the User Account Manager is not available.

See also “User” on page 274; “Workgroup Manager” on page 296.

Related utility: "SYSCON" (*Utilities Reference*).

Utilities

Programs that add functionality to the NetWare operating system.

NetWare utilities are divided into groups according to where the commands for the utilities are executed: file server utilities, workstation utilities, and router utilities.

File Server Utilities

File server utilities consist of two types: console commands and NLMs. All file server utilities are documented in the "File Server Utilities" section of *System Administration*.

Console commands are part of the SERVER.EXE file server program that you run to install a NetWare v3.12 file server.

After you run this file server program, you can perform various installation and maintenance tasks by typing these commands at the file server console. You can also issue screen commands and see file server configuration information.

The console commands are listed in Table 18.

Table 18 Console commands

Installation	Screen commands	Maintenance	Configuration information
ADD NAME SPACE BIND ENABLE LOGIN ENABLE TTS LOAD MOUNT REGISTER MEMORY SEARCH	BROADCAST CLS EXIT OFF SEND	CLEAR STATION DISABLE LOGIN DISABLE TTS DISMOUNT DOWN REMOVE DOS RESET ROUTER TRACK ON TRACK OFF SECURE CONSOLE SET SET TIME UNBIND UNLOAD UPS TIME	CONFIG DISPLAY NETWORKS DISPLAY SERVERS MODULES NAME PROTOCOL SPEED SPOOL VERSION VOLUMES UPS STATUS UPS TIME

NetWare Loadable Modules (NLMs) are code modules that can be linked to and unlinked from the file server while it is still running.

Installing a server consists of running the SERVER.EXE file and then loading additional modules. NLMs are loaded with the LOAD command.

NLMs reside in file server memory with the NetWare operating system and can access a large number of file server functions directly. They can request allocations for file server memory, use the memory to perform some task, and then return control of the allocated resources to the file server.

Third-party developers can provide file server enhancement modules that can be linked to the operating system.

The four types of loadable modules, each with its own filename extension, are listed in Table 19.

Table 19 NLMs

Disk drivers (.DSK)	Name space (.NAM)	LAN drivers (.LAN)	Utilities/file server enhancements (.NLM)
DCB.DSK ESDI.DSK IBMSCSI.DSK ISADISK.DSK MFM.DSK	MAC.NAM	NE1000.LAN NE2000.LAN NE2.LAN NE232.LAN RXNET.LAN TOKEN.LAN 3C503.LAN 3C505.LAN 3C523.LAN	DISKSET.NLM INSTALL.NLM MONITOR.NLM PSERVER.NLM UPS.NLM VREPAIR.NLM

Disk drivers enable the operating system to communicate with the hard disks. When you install NetWare the first time, you must run the SERVER.EXE program and then load a disk driver. Subsequently, you can create a STARTUP.NCF file that loads your disk drivers for you.

LAN drivers allow the operating system to communicate with different types of network boards. For each network board you install in the file server, you must also load a LAN driver.

When you load your LAN driver, you are asked to enter the hardware configuration information.

Because NetWare v3.12 allows multiple communication protocols to run on top of each LAN driver, you must also bind the driver to the communication protocol with the BIND console command.

You can create an AUTOEXEC.NCF file that loads your LAN drivers and binds them to the protocol.

Name space allows the file server to store files created by non-DOS operating systems. However, it does not allow you to connect non-DOS workstations to the file server except through a router.

Name space must be loaded before you mount a volume. You should load name space after you load your disk drivers. You can include code to support loading name in your STARTUP.NCF file.

Utilities/File server enhancements provide added functionality.

DISKSET	Use when you install external drives connected by a DCB or its third-party equivalent.
INSTALL	Use to create, delete, or modify Netware partitions on fixed disks; mirror or unmirror fixed disks; create or modify volumes; create, delete, or modify the NetWare boot files, AUTOEXEC.NCF, or STARTUP.NCF; format a fixed disk; execute a surface test of the fixed disk for bad blocks; or recopy the NetWare v3.12 diskettes.
MONITOR	Use to lock the file server console, see how efficiently your network is operating, and monitor server statistics.
PSERVER	Use to set up NetWare v3.12 print services.
UPS	Use when you attach a UPS to your file server.
VREPAIR	Use to correct minor hard disk problems on a volume without destroying its data.

Workstation Utilities

NetWare workstation utilities are designed to change the network after initial installation or to perform network tasks.

The two types of workstation utilities are command line utilities and menu utilities. These utilities are documented alphabetically in *Utilities Reference.*

All the printing utilities are documented in *Print Server.*

Command line utilities are executed at the DOS prompt.

You can use these utilities to manipulate rights and attributes, copy and print files, log in and out of file servers, view file server information, and map network drives.

Once you are familiar with the command line utilities, they are faster to use than the menu utilities.

The command line utilities are listed by category of task in Table 20.

Table 20 **Command line utilities**

Rights/ Attributes	Directories/ Volumes	Server	Users	Files
ALLOW	CHKDIR	NVER	CASTOFF	NBACKUP
FLAG	CHKVOL	PAUDIT	CASTON	NCOPY
FLAGDIR	DSPACE	SECURITY	HELP	NDIR
GRANT	LISTDIR	SETPASS	SEND	PURGE
REMOVE	MAP	SLIST	USERLIST	
REVOKE	RENDIR	SMODE	WHOAMI	
RIGHTS	VOLINFO	SYSTIME		
TLIST		VERSION		

Menu utilities allow you to perform network tasks by choosing options from menus.

You can complete some tasks, such as creating users, only in a menu utility.

The utilities are designed so that some options are available only to SUPERVISOR or users designated as SUPERVISOR-equivalent or as operators.

Menu utilities make it easier to remember how to complete tasks by displaying options.

The menu utilities are listed and described in Table 21.

Table 21 Menu utilities

Utility name	Tasks
COLORPAL	Use to modify the color schemes for NetWare menu utilities.
FILER	Use to view and modify volume, directory, and file information; grant and revoke trustee assignments; assign and modify directory and file attributes.
MAKEUSER	Use to create large groups of users.
SESSION	Use to perform a variety of network tasks. SESSION combines the functions of several command line utilities into a single menu utility and is intended primarily for network users.
SYSCON	Use to create users and groups, specify workgroup and user account managers, and designate trustee assignments.
USERDEF	Use to create templates to define and then create large groups of users. Works in conjunction with MAKEUSER.

Router Utilities

Router utilities are executed at the router console to monitor and regulate the router's resources. You can run VAPs on the router.

The following router commands are documented in *Installation*:

CONFIG
CONSOLE
DOS
DOWN
MONITOR
OFF

21 V

Value-Added Process

See “VAP” on page 291.

Value-added server

A separate, specialized, dedicated computer (such as a print server or a database server) that fulfills a specific function for network users.

See also “VAP” on page 291.

VAP

A process that ties enhanced operating system features to a NetWare v2.x operating system without interfering with the network's normal operation.

VAPs run on top of the operating system in much the same way a word-processing or spreadsheet application runs on top of DOS.

NetWare Loadable Modules provides this type of enhancement for the NetWare v3.12 operating system.

Virtual Loadable Module

(VLM) A modular executable program that runs at each DOS workstation and enables communication with the NetWare server.

A VLM file has a .VLM filename extension. For example, the IPX VLM file is IPXNCP.VLM.

The NetWare DOS Requester is composed of several VLMs. These VLMs replace, and provide backward compatibility with, NetWare shells used in previous NetWare versions.

There are two types of VLMs: child VLMs and multiplexor VLMs.

Child VLMs

Child VLMs handle a particular implementation of a logical grouping of functionality. For example, each NetWare server type has its own child VLM:

- NDS.VLM, for NetWare Directory Services™-based (NetWare v4.0) servers.
- BIND.VLM, for bindery-based servers (including NetWare v3.12).
- PNW.VLM, for NetWare desktop-based servers.

Various implementations of transport protocols also have their individual child VLMs. For example, IPXNCP.VLM handles IPX services, and TCPNCP.VLM handles TCP functions.

Multiplexor VLMs

A multiplexor is a VLM that routes calls to the proper child VLM. Requester multiplexors can be considered parent VLMs, ensuring that requests to child VLMs reach the appropriate VLM module.

See also “NetWare DOS Requester” on page 156.

VLM

See “Virtual Loadable Module” on page 292.

Volume

A physical amount of hard disk storage space, fixed in size.

A NetWare volume is the highest level in the NetWare directory structure (on the same level as a DOS root directory). A NetWare file server supports up to 64 volumes.

NetWare volumes are subdivided in two ways:

- *Logically*, volumes are divided into directories by network supervisors and users (having the appropriate rights).
- *Physically*, volumes are divided into volume segments; different segments of a volume can be stored on one or more hard disks.

  A single hard disk can contain up to eight volume segments belonging to one or more volumes, and each volume can consist of up to 32 volume segments.

  Placing segments of the same volume on multiple hard disks allows different parts of the same volume to be read from or written to simultaneously, thus speeding up disk input/output.

  However, when segments of a volume are spread over multiple disks, the volumes should be protected against disk failure by mirroring; otherwise, if a single disk fails, one or more entire volumes are shut down.

  A volume's size can be increased by adding a hard disk to the file server, setting up a NetWare partition on the disk, or adding the new NetWare partition to the existing volume as one or more new volume segments.

  Volume sizes can be increased, in many cases, while the file server is running and the volume still mounted.

The first network volume is named SYS:, and contains the SYSTEM, PUBLIC, LOGIN, MAIL, and DELETED.SAV directories. Additional volumes can be defined with INSTALL, and are assigned volume names between 2 and 15 characters long.

Several NetWare utilities (including MAP and VOLINFO) list a file server's volume names in one form or another.

In addition, the DOS DIR command lists the volume name for the specified network drive (for example, "Volume in drive F is SYS"). This corresponds to the DOS volume label shown by the DOS DIR command for local disks (floppy disks or workstation hard disks). (A local disk can be given a volume label during formatting or with the DOS LABEL command.)

When a volume is used as part of a directory path, either in NetWare documentation or on the screen, the volume name is followed by a colon (:), as in SYS:PUBLIC.

Limiting Volume and Directory Size

A volume is defined during installation; the default is the entire hard disk.

With NetWare you can limit how much disk space a user can use in a volume and in individual directories. This is useful when you need to assign the Supervisory right in a directory (rather than in a volume) for a workgroup manager.

You can use the DSPACE utility to set disk space restrictions for that directory in kilobytes. You may also want to limit username directories, particularly if your workstations have hard disks.

See also "Directory structure" on page 61; "Volume definition table" on page 294.

Related utility: "INSTALL" (*System Administration*).

Volume definition table

A table that keeps track of volume segment information such as the volume name, the volume size, and where the volume segments are located on the various network hard disks.

Each NetWare volume contains a volume definition table in its NetWare partition.

22 W

Wait state

A period of time when the processor does nothing; it simply waits.

A wait state is used to synchronize circuitry or devices operating at different speeds. For example, wait states used in memory access slow down the CPU so all components seem to be running at the same speed.

Wait time

In a NetWare UPS system, the number of seconds the UPS waits before signaling to the file server that the normal power supply is off.

The file server then alerts users at any operational workstations that they should log out.

See also “Power conditioning” on page 180; “Uninterruptible power supply” on page 271.

WAN

(Wide Area Network) A computer network not confined to a single location.

A WAN communication system is distinguished from a local area network because of its longer-distance communication capabilities.

Workgroup Manager

An assistant network supervisor with rights to create and delete bindery objects (such as users or groups) and to manage user accounts.

A Workgroup Manager has supervisory privileges over a part of the bindery. When several groups share a file server, you can use Workgroup Managers over groups that want autonomous control over their own users and data.

Workgroup Managers supplement, but do not replace, the network SUPERVISOR. SUPERVISOR retains absolute control over the network.

Workgroup Managers do not automatically acquire rights to the directory structure and file system. These rights must be granted by SUPERVISOR.

We recommend that you assign Workgroup Managers the Supervisory right in a specific volume or directory so that they can grant directory and file rights to users in their workgroups.

If you assign file rights to a directory (rather than to a volume), we also recommend that you make a disk space assignment to the directory.

Workgroup Manager Tasks

Workgroup Managers can

- Create users and manage their accounts.
- Delete users they have created.
- Create groups and add users they manage.
- With file rights, make or modify
- Trustee directory assignments;
- Trustee file assignments;
- Volume restrictions;
- Disk space restrictions.

Workgroup Managers cannot

- Create a user and make that user security equivalent to SUPERVISOR;
- Create a Workgroup Manager;
- Manage users or groups they have not created unless they are also designated as the User Account Manager;
- Assign any rights they have not been assigned by SUPERVISOR;
- Modify the login restrictions of their own accounts (unless they are also assigned management of their own accounts);
- Create or delete print queues.

Designating a Workgroup Manager

A Workgroup Manager can be a user or a group.

The user or group must first be created on the network. Then in SYSCON you select "Supervisor Options"; then select "Workgroup Managers" and insert the username or groupname.

If you designate a group as Workgroup Manager, each member of the group has Workgroup Manager privileges. By using a group you can change the assignment when you add or delete group members.

Workgroup Manager is not a bindery object. A set property called "Managers" is added to the SUPERVISOR object.

Any object (user or group) that is security equivalent to an object listed in the "Managers" property set is considered to be a Workgroup Manager.

When a Workgroup Manager creates users, the bindery uses the Workgroup Manager's ID number to indicate that the Workgroup Manager becomes the User Account Manager for these users.

Security Considerations

For *maximum security*, we suggest the following.

- Specify passwords in SYSCON for all Workgroup Manager accounts.
- Require Workgroup Managers to create a second username to use when they work as regular network users. They should log in under the Workgroup Manager username only when performing management tasks.
- Use caution in making anyone security equivalent to a Workgroup Manager.

See also “SUPERVISOR” on page 255; “User” on page 274; “User Account Manager” on page 282.

Related utility: "SYSCON" (*Utilities Reference*).

Workstation

A personal computer that is connected to a NetWare network and used to perform tasks through application programs or utilities.

Workstation boot files

See “Boot files” on page 33 (Workstation boot files).

Workstation shell

See “NetWare DOS Requester” on page 156.

Write Audit attribute

See “Attributes” on page 21; “Security” on page 221 (Attribute Security).

Write right

See “Rights” on page 207; “Security” on page 221 (Rights Security).

23 X

XON/XOFF

See “Serial communication” on page 246.