Quantum. | ATL

Network Data Management Protocol Primer

White Paper

A discussion of the technology that standardizes network backup in the enterprise

May17, 2000
Tel: 949-856-7800
Tel: 800-248-5101
Fax: 949-856-7799

Introduction

The common goal of enterprise storage management systems is to provide centrally managed enterprise-wide data protection for users in a heterogeneous environment. For the enterprise, the ability to backup and restore data is as obligatory as a dial tone. Unfortunately, the heterogeneous nature of the network storage environment requires backup solutions to comply with a myriad of operating system platforms and media devices. This daunting task is one of the driving forces behind the Network Data Management Protocol (NDMP), a scalable solution that defines a common architecture for backing up data located on heterogeneous file servers in the enterprise.

This white paper examines NDMP and its features and benefits. At the end of this white paper is a glossary of terms.

Overview

NDMP is an open-standard protocol for conducting network-based backup of storage devices in the enterprise. NDMP offers the first opportunity to achieve true enterprise-wide heterogeneous storage management solutions.

Currently, storage management vendors adapt (or 'port') their architecture to the operating systems and hardware they support. This approach results in bloated implementations strewn with layers of operating system and hardware dependencies. Organizations that use these backup solutions are forced to work around these dependencies to accommodate departmental systems across the enterprise.

NDMP is designed to address these issues by creating a "universal agent" that can be used by all centralized backup applications and file servers. This plug-and-play approach lets administrator's backup data throughout the enterprise using any combination of NDMP-compliant network-attached servers, backup devices and management applications. In this way, NDMP shifts the current paradigm so that backup is driven at the enterprise level while platforms can be driven at the departmental level.

By deviating from the current backup model, NDMP:

- Addresses the need for centralized control of enterprise-wide network data management across disparate operating systems and hardware platforms.
- Minimizes network traffic, reduces costs and eliminates concessions to departmental systems in the enterprise.
- Allows backup and network-attached storage (NAS) vendors to concentrate on their core competencies.
- Provides users with an unparalleled level of choice and interoperability.

Answering the Challenge of Increasing Data Storage Requirements

In today's information-rich society, spiraling amounts of information are streaming into organizations from the Internet, e-commerce sites, extranets, intranets, Unix® and Windows NT®-based networks, and other data-rich conduits. This avalanche of information is causing enterprise storage to grow at phenomenal rates.

The following trends illustrate the increasing amounts of data in the workplace:

- Industry experts predict that organizations will double their storage every 12 to 18 months.
- Some e-commerce (.com) organizations are expanding their data storage capacities by as much as 100% each month .
- According to International Data Corporation (September 1998): "Multiuser disk storage grew significantly from approximately 10,000 trillion bytes in 1994 to approximately 116,000 trillion bytes in 1998, and will reach approximately 1,400,000 trillion bytes in 2002. This expansion is driven by the growth in data-intensive environments and applications such as the Internet, e-commerce, data warehousing, data mining, and enterprise resource planning."
- A recent report by Strategic Research estimated the cost of managing storage to be over 10 times the initial cost of the storage device itself.
- According to Dataquest, the worldwide RAID storage market for the UNIX and Windows NT operating platforms was approximately $15.0 billion in 1998 and is expected to grow at an annual compound rate of 22% to approximately $39.8 billion by 2003.

With these daunting statistics, today's enterprise requires a comprehensive, centralized, and simplified solution for backing up data residing on disparate operating systems and NAS devices (or dedicated file server "appliances"). To address these issues efficiently and effectively, backup solutions must:

- Support backup of all operating systems and media device platforms used for storing data in the enterprise.
- Be scalable.
- Meet or exceed an organization's performance expectations and requirements.

NDMP meets the strategic need to centrally manage and control distributed data, while minimizing network traffic. As an embedded protocol, NDMP separates the data path and the control path so network data can be backed up locally and yet managed from a central location.

Answering the Challenges of Network Backup

Backup is an issue that affects the following solution providers:

- Server vendors
- Backup software vendors
- Backup device vendors
- Tape library (jukebox) vendors

Backup also effects users who demand best-of-breed solutions for their particular requirements. User demands can best be filled by defining, promoting and adopting an open-standard network-backup protocol within the backup community. Agreeing to support a well defined, standard network protocol allows solution providers to create products that interoperate with all other products that adhere to the same standard. As a result, vendors can eliminate interoperability efforts and concentrate on their core competencies. For this reason, an increasing number of solution providers are turning to NDMP to provide centrally managed, enterprise-wide data protection solutions for users in heterogeneous environments.

The following table summarizes the key challenges faced by vendors and users, and shows how they are met by standardizing on NDMP.

Category	The Challenge	NDMP's Answer
File Server Vendor	To meet customer demands, file server vendors must support a wide range of backup media devices and backup software applications. To achieve this, file server vendors work with other vendors and develop proprietary APIs to ensure interoperability with backup hardware and software. These ancillary development and support efforts impact development and support cycles that could be used to develop new features and improve performance.	Standardizing on NDMP allows vendors to focus on performance enhancements, without concern for backup or interoperability. It also eliminates the need to divert valuable resources for developing APIs.
Backup Software Vendor	To ensure complete compatibility and interoperability with various platforms and versions, vendors must expend large amounts of time, manpower, and expertise. These requirements lengthen product development timelines, and redirect resources from feature enhancements and support to compatibility assurance.	Standardizing on NDMP minimizes the need for cross-platform porting and allows vendors to focus on improving features and functions. It also shortens product development timelines, allowing vendors to improve their products' time-to-market schedules.
NAS Vendor	Because NAS devices are dedicated appliances designed to optimize file service, they do not usually support backup software solutions. Also, performing NFS (or CIFS)	NDMP provides a direct disk-to-tape backup path right out of the box, eliminating installation procedures, network traffic, and the overhead

Category	The Challenge	NDMP's Answer
	mounts of their file systems and backing up over the network is inefficient and is a time and resource consuming activity.	associated with NFS Mount. This also eliminates much of the traditional server overhead associated with backup operations.
Tape Library Vendor	The heterogeneous nature of the enterprise requires vendors to ensure their tape devices work with a myriad of operating systems and media devices.	NDMP standardizes the interface by leveraging the SCSI-2 interface or allowing for network-attached tape devices.
Users	The enterprise is extremely heterogeneous. Operating systems are routinely upgraded and there is a broad range of backup media devices and technologies. Backup software solutions must be fully compliant with all platforms on which the data is stored and with the media devices being used.	Standardizing on NDMP provides users with plug-and-play solutions for their heterogeneous environments. Users can choose the servers, backup devices, and backup software solutions best suited for their network environment, with interoperability assured.

Comparing Conventional and NDMP Architectures

To appreciate the NDMP standard, it is important to understand the complexities involved with enterprise backup.

The current procedure for backing up data on enterprise servers involves copying data stored in file systems to a backup device, such as tape in a tape library. The backup software controls the data being backed up and manages the database or catalog of the data being backed up. Individual backup software vendors implement their own protocols to manage this data transfer. Therefore, although this architecture remains constant, the system calls differ by software package and, sometimes, by server platform. Also, the data format on tape varies from vendor to vendor.

NDMP is an open network protocol that standardizes the functional interfaces used for data transfers conducted during the backup-and-restore process. With NDMP, vendors follow common interfaces for common architecture data transfers. In this way, file system data is copied from the file system to the backup device using a common interface, regardless of the platform or device. Similarly, control or file meta data (such as index data) uses common commands to be passed to and from the backup software, regardless of the software package being utilized.

Because NDMP integrates with a standard "dump" backup program on the NDMP server, the data format on tape is consistent no matter which backup vendor is acting as the NDMP client. In traditional backup systems, by contrast, each backup application writes its own proprietary format to tape.

Figures 1 and 2 compare the common backup architecture using NFS or CIFS with NDMP-based backup. In figure 1, file systems on the Network Appliance™ (NetApp®) filers are backed up to the backup server by first being mounted (via NFS or CIFS), creating a lot of overhead and heavy traffic on the LAN. This traditional method does not use NDMP.

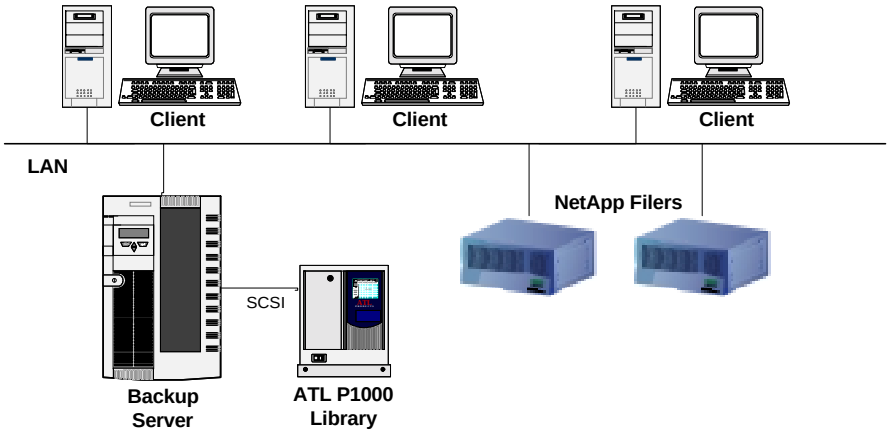

Figure 1: Traditional Backup Over the LAN

In Figure 2, only NDMP commands flow on the LAN, providing significantly faster throughput to tape. However, it is expensive and error prone when staff has to manually change tapes on a daily basis.

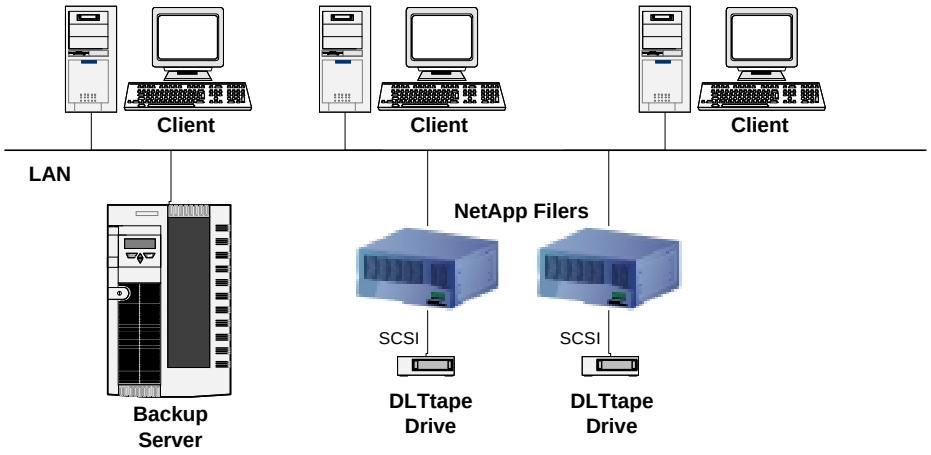

Figure 2: NDMP-Based Backup Using Individual Tape Drives Connected to Filers

More about NDMP Architecture

NDMP is a network protocol that defines communication between a server and backup software. Communication is achieved using a series of defined interfaces. These interfaces are XDR-encoded messages that pass over a bi-directional TCP/IP connection.

The NDMP architecture follows a client/server model, where the backup software is a client to the NDMP server. For every connection between the client on the backup

software host and the NDMP host, there is a virtual state machine on the NDMP host that is controlled using NDMP. This virtual state machine is referred to as the NDMP server.

NDMP History and Roadmap

Intelliguard™ and Network Appliance jointly developed NDMP several years ago, writing the initial NDMP specification and implementing the first version of the protocol. From there, support from other solution providers was sought to make NDMP an open-standard backup protocol. Initially, 17 vendors backed the initiative. The complete NDMP specification was submitted to the Internet Engineering Task Force (IETF) in October, 1996 through the Internet Draft and Request for Comment (RFC) processes.

Version 2 of the protocol has been approved by the IETF and version 3 has been posted to gather input. Now becoming a standard sponsored by the Storage Networking Industry Association (SNIA), NDMP will continue to grow and attract solution providers. A working group of SNIA has now started developing NDMP version 4. As a result, solution providers will be implementing an increasing number of NDMP-compliant solutions in their future product releases. The full NDMP specification and additional information is published at www.ndmp.org.

Quantum|ATL's NDMP-Compliant Product Offerings

Quantum|ATL is proud to embrace NDMP and believes its endorsement of NDMP will increase the protocol's acceptance among industry leaders as an open approach to building enterprise storage networking and backup solutions. During calendar year 2000, two NDMP-based solutions will be brought to market by Quantum|ATL.

Figure 3 illustrates Quantum|ATL's initial NDMP-compliant product offering. In this figure, data residing on two Network Appliance storage arrays is backed up to an Quantum|ATL P1000 tape library. The P1000 supports 4 tape drives and 5 Logical Unit Numbers (LUNs); one LUN corresponds to each drive, with the one additional LUN reserved for the robot.

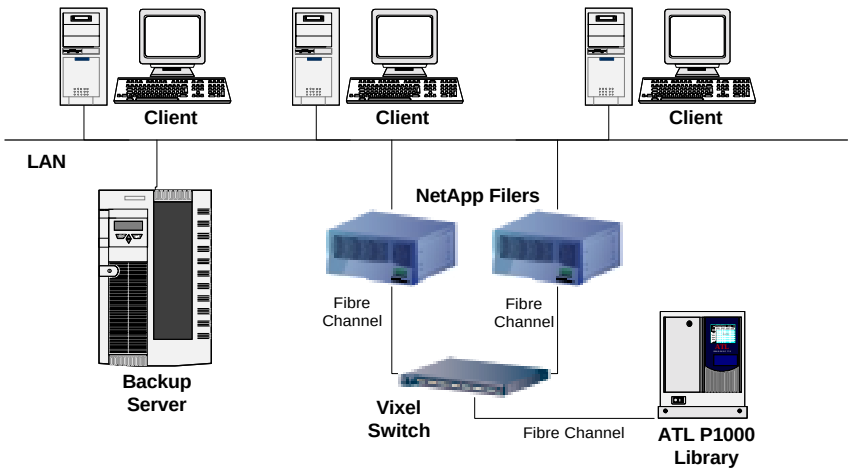

Figure 3: Quantum|ATL's NDMP-Compliant Backup Solution

The backup data path is a 100 MB/sec Fibre Channel Storage Area Network (SAN) that traverses a Vixel™ switch. The P1000 library connects to port 8 on the Vixel switch. The entire backup process is externalized from the Ethernet® network.

The Application Manager on the messaging network is responsible for managing the backup operation. It maintains the status and media inventory, making sure cartridges are in the correct slots or drives and are available when needed. The Application Manager sends its control commands across the network to the Network Appliance filers that route the commands appropriately. The filers also receive information from the storage array and forward it to the Application Manager.

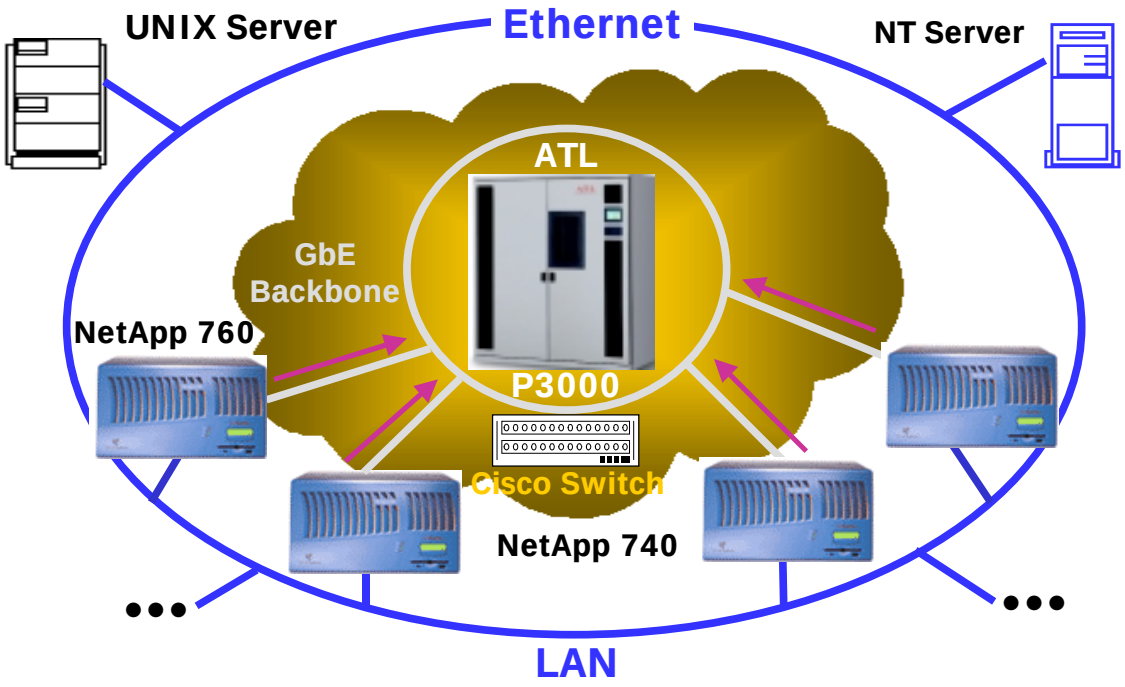

Figure 4: Quantum|ATL's Revolutionary NDMP Operated Tape Library

Figure 4 illustrates Quantum|ATL's second NDMP initiative. This new technology was announced at Networld+Interop in Las Vegas on May 10, 2000 with participation from Network Appliance, Cisco Systems™, Foundry Networks™, VERITAS® and Legato®.

The Quantum|ATL P3000 enterprise library (or any other library incorporating Quantum|ATL's Prism Storage Architecture™) is directly addressable as a NDMP storage unit. Network Appliance filers can send data via their native "dump" command to the library at high speed over Ethernet (100 baseT or Gigabit Ethernet) upon being commanded to do so via NDMP. Use of NDMP in this way provides the following benefits:

- Based on the maturity and interoperability of Ethernet.

- Leverages existing infrastructure, tools and skills, thereby ensuring investment protection.
- Permits heterogeneous data and device sharing.

Conclusion

Every software and hardware backup product interacts with other management tools and applications in various ways. A standard protocol can facilitate the interoperability of these products, allowing them to operate more efficiently and with greater simplicity. That is the objective of NDMP.

NDMP is designed to ease the burden of backup and restore operations and provide interoperability in the heterogeneous enterprise. To achieve this, NDMP sets forth a series of well-defined interfaces that address the transfer of file system and control data for backup and restore activities. Defining such a standard accelerates the deployment of interoperable backup solutions and fosters a cooperative environment in which all solution providers work toward the same goal. It also frees solution providers from dedicating resources to ensure interoperability, allowing them to redirect efforts to their areas of expertise.

NDMP ensures that compliant products can be used by all customers, regardless of operating system and media device platforms. NDMP's intelligent, open-standards approach also encourages users to take full advantage of the latest architecture, while protecting their investment in their existing data management solution.

NDMP compliance provides true plug-and-play interoperability to users that are accustomed to plug-and-play solutions. As a result, users can now select the best enterprise-wide backup software solutions and hardware to meet the demands of their particular environment, without concern for interoperability.

Examples of products that can benefit from the interoperability afforded by NDMP include NDMP-compliant servers and tape libraries that can be delivered “backup-ready” to users. Once attached to the network, the NDMP-compliant backup software provides backup protection to the servers as part of an enterprise-wide solution.

For more information about NDMP, visit www.ndmp.org. For additional information about Quantum|ATL’s NDMP-compliant enterprise storage solutions, please contact Quantum|ATL via phone at (800) 284-5101 or 949-856-7800, via e-mail at atlinfo@atlp.com, or by visiting our web site at www.QuantumATL.com.

NDMP FAQs

This section provides answers to frequently asked questions about NDMP.

What is NDMP?

An abbreviation for Network Data Management Protocol, NDMP is an open-standard protocol for enterprise-wide backup of heterogeneous network-attached storage. NDMP was jointly developed by Network Appliance and Intelliguard software (now Legato) and submitted to the IETF for approval.

What NDMP resources are currently available?

You can click the URL http://www.ndmp.org/download/ to sign up for the NDMP technical discussion group and the NDMP announcement list, request that your company join the NDMP initiative and download the following documentation and software:

- NDMP specification document
- NDMP workflow analysis document
- NDMP rpcgen input (.x) file
- NDMP software development kit (SDK)

You can also contact the following FTP site (ftp.ndmp.org) to download the documentation and SDK mentioned above.

What tape format does NDMP use?

The NDMP data server produces an NDMP stream that the NDMP tape server writes directly to tape. Controlling this stream of data and control characters controls the format on the tape.

How does NDMP control a tape library?

NDMP has a SCSI interface that allows an NDMP client to pass SCSI command data blocks to a tape library attached to the NDMP server via SCSI or Fibre Channel.

How does NDMP control a tape?

NDMP has a tape interface that allows an NDMP client to perform tasks such as positioning the tape and reading and writing tape labels.

Glossary

The following glossary defines the technical terms in this white paper. A bold italicized word is itself defined in this glossary.

Backbone	A network transport system (such as Ethernet or Fibre Channel) that interconnects intermediate systems (bridges, switches, and/or servers).
Bandwidth	The amount of data that can be transmitted in a fixed amount of time. For digital devices, the bandwidth is usually expressed in bits per second (bps) or bytes per second. For analog devices, the bandwidth is expressed in cycles per second, or Hertz (Hz).
Bit	Short for binary digit, the smallest unit of information on a machine. A single bit can hold a value of 0 or 1. Combining consecutive bits produces more meaningful information. For example, a byte comprises 8 consecutive bits.
Byte	Abbreviation for "binary term," a unit of storage capable of holding a single character. On nearly all computers, a byte equals 8 bits. Large amounts of memory are indicated in terms of kilobytes, megabytes, and gigabytes. For example, a disk that can hold 1.44 megabytes can store approximately 1.4 million characters (about 3,000 pages of information).
Client	The client part of a client-server architecture. Typically, a client is an application that runs on a personal computer or workstation and relies on a server to perform some operations. For example, an e-mail client is an application that enables you to send and receive electronic mail.
Client-server Architecture	A network architecture in which each computer or process on the network is either a client or a server. Clients rely on servers for resources, such as files and devices.
Cluster	A group of independent computers managed as a single system for higher availability, easier manageability, and greater scalability.
CIFS	Common Internet File System. Used by Windows NT for the same purpose as NFS is used by Unix systems.
Ethernet	A Local Area Network (LAN) protocol developed by Xerox Corporation, DEC, and Intel. Ethernet uses a bus or star topology and supported data transfer rates of 10. Newer versions of Ethernet include 100Base-T (or Fast Ethernet) which supports data transfer rates of 100 Mbps and Gigabit Ethernet, which supports data transfer rates of 1000 Mbps.
Fabric	A network topology composed of switches. Each switch port provides full bandwidth. Since concurrent conversions between ports may occur on each switch in the fabric, the total capacity is measured in multiple hundreds of megabytes per second.

Fast Ethernet A networking standard that supports data transfer rates up to 100 megabits per second. Fast Ethernet is based on the older Ethernet standard and is officially referred to as 100BASE-TX.

Fibre Channel	A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers. Fibre Channel was designed for new mass storage devices and other peripheral devices that require very high bandwidth. Using optical fiber to connect devices, Fiber Channel supports full-duplex data transfer rates of 100. Fiber Channel is compatible with, and is expected to eventually replace, SCSI for high-performance storage systems.
Fibre Channel – Arbitrated Loop topology	A Fibre Channel topology where the Fibre Channel is connected in a loop with devices all connecting to the loop. This topology has a similar structure to a Token-Ring network. Up to 126 nodes can be connected to the loop.
Host Bus Adapter (HBA)	A device that is similar to SCSI host bus adapters and Network Interface Cards (NICs). HBAs are available for copper and fiber optical media. A typical Fibre Channel PCI HBA is half-length, and utilizes a highly integrated Fibre Channel ASIC for processing the Fibre Channel protocol and managing the I/O with the host. Adapters are also available for EISA, GIO, HIO, MCA, PCI, PMC, SBus, and Compact PCI.
Hub	Devices that connect nodes in a loop. Logically, the hub is similar to a Token Ring hub with "ring in" and "ring out." Hubs are generally used in arbitrated loop topologies, and can be used in complementary roles to switches in fabric topologies. Each port on a hub contains a port bypass circuit (PBC) to automatically open and close the loop. Hubs support hot insertion and removal from the loop. If an attached node is not operational, a hub will detect this and bypass the node. A typical hub has seven to 10 ports and can be stacked to the maximum loop size of 126 ports.
Local Area Network (LAN)	A computer network that spans a relatively small area. Most LANs are confined to a single building or group of buildings. However, one LAN can be connected to other LANs over any distance via telephone lines and radio waves. A system of LANs connected in this way is called a Wide Area Network (WAN). Most LANs connect workstations, personal computers, and devices such as printers. This allows a single device, such as a printer, to be shared by many users. Users can also use a LAN to communicate with each other by sending e-mail. There are many types of LANs, Ethernet being the most common for personal computers.
Mbps	An abbreviation for megabits per second. When used to describe data storage, a megabit refers to 1,024 kilobits or one million bits. Networks are often measured in megabits per second.
MBps	An abbreviation for megabytes per second. Data transfer to and from storage devices is often measured in megabytes per second.
NFS	Network File System. A method of sharing files across a computer network. Pioneered by Sun Microsystems, it is now a defacto standard in the Unix environment. NFS is built on Transmission Control Protocol/Internet Protocol (TCP/IP) and Ethernet.
100Base-TX	A networking standard that supports data transfer rates up to 100 Mbps. 100BASE-TX is based on the older Ethernet standard. Because it is 10 times faster than Ethernet, it is often referred to as Fast Ethernet. Officially, the 100BASE-TX standard is IEE802.3u.

Point-to-point Topology	A topology that provides a simple, direct connection between just two nodes. This topology is the closest approximation to the traditional SCSI topology.
Protocol	A common set of rules and signals that computers on the network use to communicate. A couple popular protocols for Local Area Networks (LANs) are Ethernet and Token-Ring.
Router	An interface device that allows legacy storage systems (layer 2 and /or 3 devices) to use Fibre Channel as a reliable, gigabit backbone.
SCSI	Abbreviation for Small Computer System Interface, pronounced "scuzzy." SCSI is a parallel interface standard for attaching peripheral devices to computers. SCSI interfaces provide for faster data transmission rates (up to 80 megabytes per second) than serial and parallel ports. In addition, you can attach many devices to a single SCSI port.
Segment	A section of a network that is bounded by bridges, routers, hubs, or switches. Dividing an Ethernet network into multiple segments increases bandwidth on the Local Area Network (LAN). If segmented correctly, most network traffic will remain within a single segment, enjoying the full 10 Mbps bandwidth. Hubs and switches are used to connect each segment to the rest of the LAN.
Server	A computer or device on a network that manages network resources. Servers are often dedicated, meaning that they perform no other tasks besides their server tasks. For example, a file server is a computer used to store files, while a network server is a computer that manages network traffic.
Static Switch	A device that provides point-to-point connections. A static switch (or link switch) offers a low-cost alternative for applications that do not require the fast, dynamic switching capability inherent in the Fibre Channel protocol.
Storage Area Network (SAN)	A high-speed subnetwork of shared storage devices.
Switch	A device used to create Fibre Channel fabrics. Switches provide more resiliency features than hubs and enable the fabric's aggregate bandwidth to scale as nodes are added. With hub-connected Storage Area Networks, aggregate bandwidth remains constant as nodes are added, but bandwidth per node decreases as more nodes share a fixed amount of bandwidth. With switched fabrics, bandwidth per node remains constant as nodes are added, but the aggregate bandwidth of the fabric increases proportionally to the number of nodes.
Tape	A magnetically coated strip of plastic on which data can be encoded. Storing data on tapes is considerably less expensive than storing data on disks. Tapes also have large storage capacities, ranging from a few hundred kilobytes to several gigabytes. Accessing data on tapes, however, is much slower than accessing data on disks. Tapes are sequential-access media: to get to a particular location on the tape, the tape must go through all the preceding points. In contrast, disks are random-access media because a disk drive can access any point at random without passing through intervening points. Tapes come in a variety of sizes and formats.

10Base-T	One of several Ethernet adaptations for Local Area Networks (LANs). The 10Base-T standard uses a twisted-pair cable with maximum lengths of 100 meters.
Topology	The logical and/or physical arrangement of stations on a network.
Wide Area Network (WAN)	A computer network that spans a relatively large geographical area. Typically, a WAN consists of two or more Local Area Networks (LANs) Computers connected to a WAN are often linked through public networks, such as the telephone system. They can also be connected through leased lines or satellites. The largest WAN in existence is the Internet.