Network Storage for Business Continuity and Disaster Recovery and Home Media

White Paper
July, 2008

Abstract

Network storage is a complex IT discipline that includes a multitude of concepts and technologies, like back-up, mirroring, replication, disaster recovery, RAID, NAS, and SAN. Companies of all sizes and shapes need to consider network storage issues when upgrading systems or designing new storage solutions. This white paper explores today's network storage challenges, defines common network storage terms, compares technologies, and provides a brief introduction to D-Link's storage solutions.

D-Link®
Building Networks for People

Business Continuity and IT Challenges

Business continuity is critical. Organizations depend on information to stay competitive, manage operations, produce goods and services, and report on progress. To that end, file accessibility, disaster recovery, back-up capabilities and the IT systems that support these functions need to be in place and tuned for efficiency. Fortunately, today's network storage hardware and software enable ubiquitous access to enterprise data anytime, anywhere.

In this day and age, information and intellectual property (IP) constitutes the majority of a company's assets. By extension, a company's storage infrastructure and strategy becomes a very valuable asset. To maintain communication, production, reporting, compliance and archiving practices, companies need to have continuous access to information stored on their computing resources. They need to deliver data and documents to users quickly, recover lost and corrupted files, search for internal information, and maintain vigilance against fire, flood, theft and accidental data loss.

Not all companies have the IT resources and expertise to make the sound decisions about storage, however. Assembling the appropriate network storage system at a reasonable cost is one of the big IT challenges of the day.

File Serving, Back-Up, Replication and Disaster Recovery

Most of the world's digital information resides on network storage devices. Information has to be somewhere to get somewhere, right? Data really doesn't come out of an 'Internet cloud' as some diagrams would have you believe. Low cost disk drives are the medium of choice for disk to disk backup, with tape as a common archive backup medium.

This wasn't always the typical architecture, however. Not long ago, much of the world's data resided on decentralized servers and desktops with individual disk drives within each computer. Much of small business and individual computing still operates this way. Laptops still have disk drives or large flash memory. Technically, you could consider a computer in a small office a 'server' that links up with the rest of the network via tools like GoToMyPC and Microsoft Remote Desktop.

In any event, the data center model – where centrally managed disk farms provide storage for widely flung client computers – is where everything is headed. The model provides optimal data loss risk reduction and disaster recovery benefits. Eventually every size company will utilize the data center model – whether via a remote service or on the premises.

To make it work, the model requires efficiently managed network storage. Files need to be efficiently 'served' to end-users. Data has to be backed up. The term backup refers to the process of copying files in order to restore them in the event of data loss. Backup serves two primary purposes – one, to restore after a disaster (disaster recovery), and two, to restore small numbers of files after they have been accidentally or intentionally deleted or corrupted. Disaster recovery is usually associated with natural or human-induced disasters. Most companies employ tested disaster recovery plans to ensure data integrity and security. In larger enterprises these plans are part a larger process known as business continuity planning (BCP).

Replication refers to the ability to keep distributed data synchronized by routinely copying entire databases to other servers or storage systems in the network. Hot backup, sometimes called failover replication, maintains copies of data at a different site for backup. Remote replication refers to a similar process where data is duplicated onto another storage device at another location. All these practices are best managed with leading-edge network storage solutions.

SAN or NAS?

The terms SAN and NAS – for Storage Area Network and Network Attached Storage – are often confused with each other. The acronym similarity doesn't help, and the fact that each can do many of the things the other can do obscures matters. In most situations, each technology works equally well.

Each storage system connects disk arrays (RAID) to a network. The way they connect is different, however. (See diagram below). The main difference is in the wiring and related protocols that connect the systems. Traditionally, SAN uses Fibre Channel wiring with the SCSI protocol, and NAS uses the TCP/IP protocol over common networking cables.

There has been an ongoing debate as to speed discrepancies between the two systems. The common argument said that SAN was faster than NAS because TCP/IP overhead reduces the efficiency of data transfer. There is a cost discrepancy as well. SAN solutions are generally more expensive to implement than NAS.

Both technologies are advancing rapidly, however, with new versions bumping up speeds exponentially. And, as SAN and NAS technologies evolve, the boundaries between the two are blurring. SCSI over IP and Open Storage Networking (OSN) technologies combine the best of NAS and SAN into one coherent data management solution, for example. Internet SCSI (iSCSI) encapsulates SCSI packets in TCP then routes them using IP. This allows block-level storage data to be transported over widely

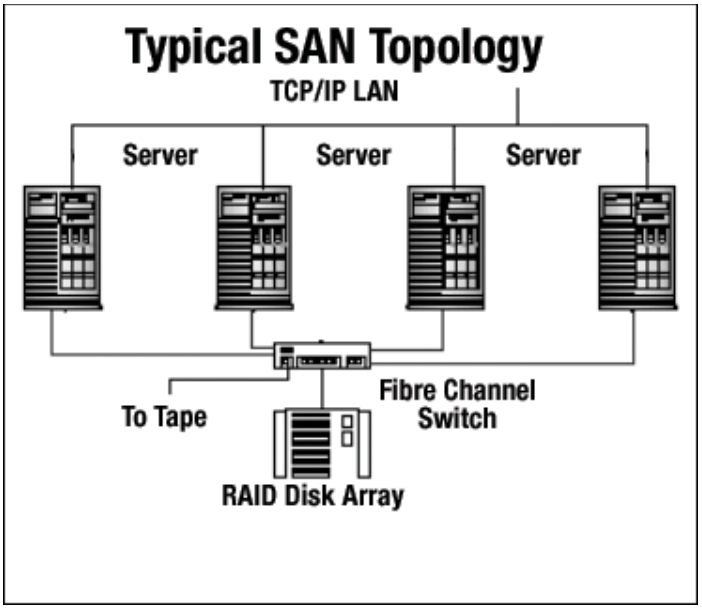

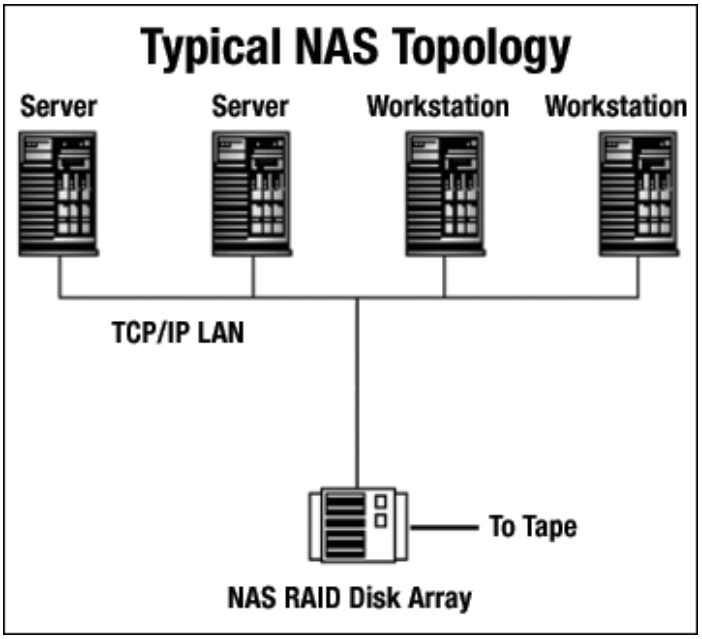

Traditionally, a typical SAN topology uses fibre channel switch wiring while a typical NAS topology uses TCP/IP protocol over common networking cables.

used IP networks. As a result, enterprises and service providers can manage SAN storage networks from a central location by using the ubiquitous IP infrastructure – which is similar to a NAS architecture.

iSCSI vs. Fibre Channel SAN Solutions

An iSCSI IP storage network is generally viewed as less complex to implement and maintain than Fibre Channel. There is also a significant cost differential between iSCSI and Fibre Channel, though iSCSI currently does not scale to the same performance level as a Fibre Channel SAN. iSCSI shows storage as a virtual block-level device, which allows applications to operate under the illusion that the storage space being used is exclusive. Data-intensive applications need this structure to perform optimally.

Fibre Channel, on the other hand, is a high-performance transmission technology that uses the same block storage format the attached disk drive uses. Data transfer is same to same. Fibre Channel Total Cost of Ownership (TCO) is fairly high in comparison to iSCSI. Many IT staffs have limited Fibre Channel expertise. Small businesses would need to hire specialists to manage a Fibre Channel SAN. As a result, Fibre Channel installation and maintenance costs (in terms of time, difficulty and real dollars) are frequently prohibitive for SMBs. Fibre Channel cabling distances are limited, as well. The theoretical range is 10km, however multi-mode fiber links in Fibre Channel SANs are closer to 250 to 500 meters.

Choosing iSCSI SAN or NAS

Enterprises with fairly robust computing environments will want to lean toward an iSCSI SAN solution. Data intense applications that are distributed among remote users or many users on a network would benefit from an iSCSI SAN, for example. However, many small businesses, like law offices, doctors, coffee shops, small retail businesses and SMBs, can get plenty of value from a NAS while saving quite a bit of money. D-Link's NAS devices for example are often 1/10th the cost of our SAN solutions.

NAS or USB External Hard Drive?

For most businesses and home networks, a NAS is preferable to off-the-shelf external USB hard drives. One of the biggest advantages of NAS is data accessibility. Users of D-Link NAS solutions (the D-Link DNS-323, for example) also cite much better performance measurements in comparison to standard NAS devices. One specific user on forums.techguy.org cited read/write speeds of 15Mbps with the D-Link DNS-323. Other NAS units tested achieved less than 3Mbps write speeds and 5Mbps read speeds.

NAS RAID configurations provide enhanced mirroring, system-wide performance and fault-tolerance. These systems also include built-in Internet and FTP access, tool-less drive bays, and easy user permission configurations. This is ideal for office environments where employee-specific data is sensitive, or for home networks where children need to be prevented from accessing confidential data and sensitive material. D-Link NAS solutions also make storage available to any computer (PC, MAC, or Linux-based) on the network – without installing any software on the computers.

D-Link Solutions

D-Link offers a wide variety of SAN, NAS and mixed-capability solutions that leverage the technologies and protocols mentioned above. We help organizations regain control of data and reduce exposure to data loss disaster. Our integrated solutions protect your intellectual property, enable instantaneous restores, and ensure business continuity.

The products listed below offer flexible SAN and NAS network storage solutions. For more detailed product descriptions, please visit our network storage solution site.

DSN-3200 8x1GbE iSCSI SAN Array, 15 Bays, 3U

Features:

- Integrated iSCSI System-on-a-Chip (SoC) solution that can handle over 80,000 I/Os per second
- Eight 1GbE ports
- 15TB capacity with 1TB hard drives (Supports higher capacity drives as they are introduced)
- Embedded disk controller supporting RAID levels 0, 1, 1+0, and 5

DSN-3400 1x10GbE iSCSI SAN Array, 15 Bays, 3U

Features:

- Integrated iSCSI System-on-a-Chip (SoC) solution that can handle over 80,000 I/Os per second
- Single 10GbE port
- 15TB capacity with 1TB hard drives (Supports higher capacity drives as they are introduced)
- Embedded disk controller supporting RAID levels 0, 1, 1+0, and 5

DSN-2100 4x1GbE iSCSI SAN Array, 8 Bays, 2U

Features:

- Integrated iSCSI System-on-a-Chip (SoC) solution that can handle over 80,000 I/Os per second
- Four 1GbE ports
- 8TB capacity with 1TB hard drives (Supports higher capacity drives as they are introduced)
- Embedded disk controller supporting RAID levels 0, 1, 1+0, and 5

D-Link DNS-323 2-Bay (NAS) Network Attached Storage Enclosure

Features:

- Insert either one or two internal 3.5” SATA hard drives
- Protect your important files with mirrored hard drives using RAID 1 Technology
- Access stored files over the Internet
- Supports Standard, JBOD, RAID 0, RAID1

DNS-343 4-Bay (NAS) Network Attached Storage Enclosure

Features:

- Easily Insert up to 4 SATA hard drives without using tools
- Protect important files using RAID 1 or RAID 5 Technology
- Access stored files over the Internet
- OLED display screen
- Supports standard, JBOD, RAID 0, 1, 5

For more information about D-Link Network Storage solutions, please visit our Network Storage Site or call **1-800-326-1688**.

DEFINITIONS

RAID
(Redundant Array of Independent Disks) A disk subsystem used to increase disk access and system-wide performance while providing fault tolerance. A RAID array typically includes two or more ordinary hard disks and a RAID disk controller.

SAN
(Storage Area Network) A network of storage disks that connect multiple servers to a centralized storage pool via Fibre Channel cabling typically on a dedicated network. As opposed to managing hundreds of servers, each with their own disks, SANs simplify resource administration. As a result, disk maintenance and routine backups are easier to schedule and control.

NAS
(Network Attached Storage) A file server system that attaches to the LAN like any other client or server in the network (typically via Ethernet). Rather than utilize a standard operating system, NAS uses a lightweight micro-kernel that is designed for file I/O. NAS shares the bandwidth on the production network. In simple comparative terms, a NAS is a file server, while a SAN is extended disk storage. See the section titled “SAN or NAS?” for more in-depth SAN-NAS comparisons.

iSCSI
A protocol that serializes SCSI commands and encapsulates them into TCP/IP package. iSCSI supports Gigabit Ethernet interface at the physical layer, which allows systems supporting iSCSI interfaces to connect directly to standard Gigabit Ethernet switches and/or IP routers.