Dominique A. Heger, Fortuitous Technology, (dom@fortutuous.com), Austin, TX, 2006

SAN and NAS Solutions - Introduction, Topology, and Terminology

Introduction

At a first glance, SAN and NAS solutions seem almost identical, and as a matter of fact, many times either solution may be applicable in any given environment. In general, NAS and SAN solutions utilize RAID systems that are connected to an interconnect (network). However, there are *significant differences* between SAN & NAS solutions, which do have a profound impact on how the actual data is being accessed and utilized (see Figure 1).

Figure 1: SAN & NAS Topology

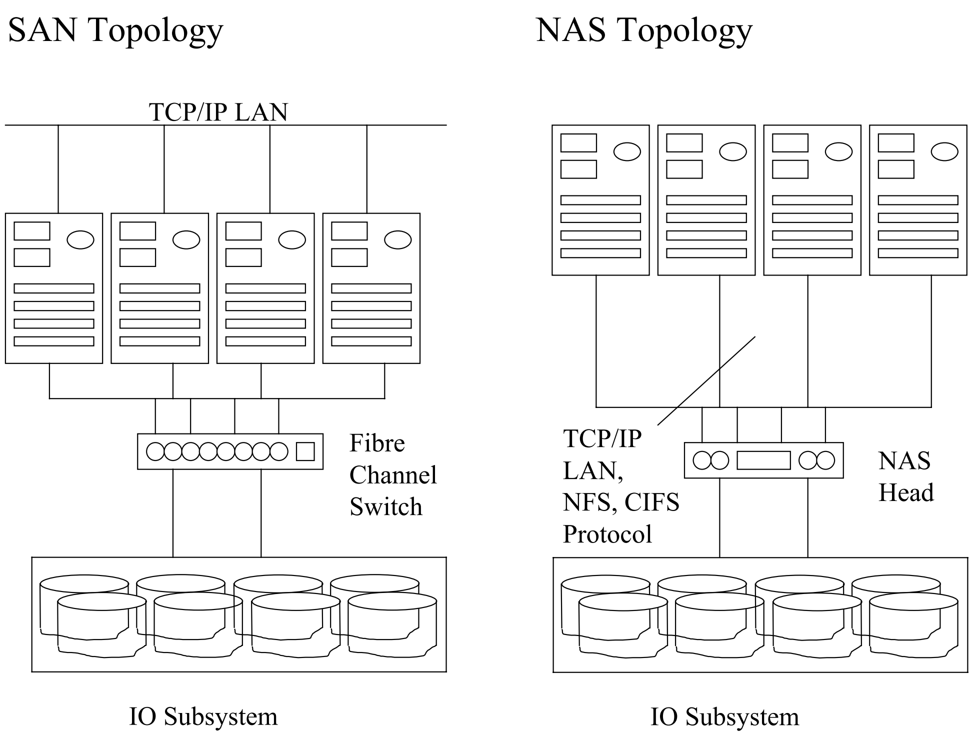

Most SAN verses NAS comparisons focus on the actual wires, but the argument can be made that the protocols are the most important factor that distinguish the two solutions. To illustrate, one thesis is that SCSI is faster than Ethernet, and hence reflects a better solution from a performance perspective. This statement is normally based on that fact that the TCP/IP overhead significantly impacts the efficiency of any data transfer scenario.

The Wires being used:

- NAS solutions utilize TCP/IP based networks, such as Ethernet, FDDI, or ATM
- SAN solutions use Fibre Channel connections

The Protocols being used:

- NAS solutions use TCP/IP and NFS/CIFS/HTTP based networks
- SAN solutions utilizes Fibre Channel encapsulated SCSI setups

Table 1: Additional SAN NAS Differences

SAN and NAS Distinctions		
NAS		***SAN***
Almost any machine that connects to a LAN (or is interconnected to a LAN via a WAN) may utilize NFS, CIFS or HTTP protocol to connect to a NAS		Server class devices that are equipped with SCSI Fibre Channel adapters connect to a SAN. A Fibre Channel based solution has a distance limit of approximately 6 miles
A NAS identifies the data by file name and byte offset, transfers file data or metadata, and handles security, user authentication, file locking		A SAN addresses the data by logical block numbers, and transfers the data in (raw) disk blocks.
A NAS allows greater sharing of information, especially among different operating systems		File Sharing is operating system dependent, and may not exist for all operating systems that are being used
File system is managed by the NAS head unit		The SAN servers manage the file system
Backups and mirrors are generated on files, not blocks (this may save bandwidth and time)		Backups and mirrors require a block by block copy operation. A mirrored system has to be either identical, or greater in capacity (compared to the source)

NAS Terminology

- The NAS head represents the part of the NAS solution required for the clients to connect to the IO subsystem. Behind the NAS head, hundreds or thousands of GB of available IO storage may exist, but the clients have to access the IO space via the NAS head. A NAS head is also called a NAS Gateway (a system), which serves as the actual control function of a NAS

- NFS (Network File System) is one of the communications protocols usually supported by NAS heads (for the communication with the network clients); particular in UNIX or Linux based solutions. It has to be pointed out though that NFS clients are available for just about any operating systems these days

- The CIFS (Common Internet File System) protocol is primarily responsible for file sharing and communication with Windows (and Linux-based Samba) servers, and represents another commonly supported protocol for most NAS heads. Most Windows clients utilize CIFS to communicate with the NAS head. Both, NFS and CIFS utilize TCP/IP as the underlying communication facility.

SAN Terminology

- A SAN solution can be loosely described as a network of storage disks. In large environments, a SAN connects multiple server systems to a centralized pool of disk storage. Compared to managing hundreds of servers (each with its own disk subsystem), SAN's simplify system administration tasks. By treating all the company's storage as a single resource, disk maintenance and backups are easier to schedule and control.

- SAN solutions provide high-speed disk access capabilities. The SAN network allows data transfers among server systems and IO subsystems at the same (high peripheral channel) speeds, as if the IO subsystems were directly attached to the serve systems. The Fibre Channel technology is the driving force behind SAN's, and is typically used to encapsulate SCSI commands. SSA and ESCON channels are also supported in SAN environments.

- The question of centralized or distributed solution may arise in a SAN design study. A centralized SAN connects multiple server systems to a collection of disks, whereas a distributed SAN typically uses one or more

Fibre Channel or SCSI switches to connect the nodes from several buildings or campuses. For longer distances, SAN traffic may be transferred across ATM or SONET fabrics

- SAN over IP. Another valuable SAN option is an IP storage based solution, which enables data transfers via IP over fast Gigabit Ethernet (locally) or via the Internet (remotely).

- A Channel Attached verses Network Attached discussion may surface in the SAN/NAS design phase as well. The NAS solution reflects basically a file server solution that attaches to the LAN (like any other node on the network). Rather than supporting a full-blown operating system, a NAS utilizes a slim micro-kernel that is specialized to handle only file reads and writes (CIFS, NFS, and NCP). However, a NAS may be subject to the variable performance behavior and overhead scenarios of a network that may contain and serve thousands of users.

SAN & Fibre Channel

Fibre Channel (FC) solutions reflect point-to-point, serial bi-directional interfaces. *4G Fibre Channel solutions will replace 1G and 2G storage networks in the near future.* The current technology used in Fibre Channel SANs reflects the 2G technology, which allows for a maximum throughputs of up to 2 Gbps. However, as demand for bandwidth-intensive applications (such as CAD/CAM, real-time computing, data warehousing or video streaming) grows, this speed will be insufficient, and hence has to be increased. That is where 4G Fibre Channel solutions emerge, a technology, which doubles the maximum throughput to 4 Gbps. The new specs for 4G were approved by FCIA (Fibre Channel Industry Association) in 2003, and it is widely supported by most vendors today. Originally, the specs sought to deal with internal connectivity, connecting disk drives to the server. Later it was decided to extend it for interconnecting the switching Fabric in SAN solutions. This basically includes the Fibre Channel switches, which intelligently manage the interconnectivity among various devices and nodes in a SAN. The 4G technology maintains backward compatibility with both older specs (1G and 2G). It also supports the loop architecture that is common to both older solutions. Customers can incrementally upgrade their systems to 4G. The technology reduces the number of connections among IO storage systems and server nodes, while significantly improving the throughput (this is workload dependent of course).

4G products are starting to become more and more of a mainstream solution. Cisco focuses on their MSD 9000 family, whereas PMC-Sierra produces some of the 4G switches. Emulex has its HBA's (Host Bus Adapters) and embedded storage switches benchmarked and tested for 4G. Broadcom has launched the BCM 8421 repeater for 4G switches and storage arrays. Customers with high performance computing needs may find the technology useful, as faster backups and data recovery can be performed. Scientific environments that need to access large amounts of data to solve complex problems have the opportunity to take advantage of a 4G Fibre Channel solution as well. High quality graphics, such as animated movies, can be produced in lesser time. The 4G technology also offers reliable transmission of digital audio/video applications. It has to be pointed out though, that at the moment, a significant number of customers are not taking full advantage of the existing 1 Gbps and 2 Gbps Fibre-Channel solutions. Fibre Channel is rather expensive to set up (especially for smaller environments) and hence, iSCSI may be a valuable alternative. It has to be reemphasized though that iSCSI is not as fast as a Fibre Channel based solution.

Summary

In a nutshell, NAS devices are storage appliances, large, single-purpose servers that plug into a network. These appliances perform one task, and they do it well. The capacity of large NAS appliances is typically in the TB range, while SAN solutions reflect multi-server multi-storage networks that may grow to hundreds of TB's. A SAN acts as a secondary network to a LAN, and every connected server that requires access to the SAN has to have a Fibre Channel connection to the SAN. Another key distinction between NAS and SAN solutions is the concept of heterogeneous file sharing. NAS appliances define heterogeneity at the file or data-block level, while SAN solutions normally define heterogeneity at the volume level. Out of the box, a NAS appliance allows UNIX and Windows clients to share the same file. In a SAN, a storage system, not that actual data, is shared. The storage is shared at the cabinet level by partitioning a physical storage device, assigning logical volumes to a given server. In most cases, a Windows server is prevented from accessing a UNIX volume and vice versa. Both methods allow performing data backups and restores, each providing their own benefits. In a NAS environment, a client makes a file system call over the network, as opposed to execute device-oriented commands from a server to a storage device in a SAN setup.

References

Joachim, "Storage Pipeline: A Look at the Storage Professional," Network Computing, 16 September 2004.

Gray, "HPC at BP: Storage Strategy," Storage on the Lunatic Fringe, SC'03 Panel, 19 November 2003.

Menon, Pease, Rees, Duyanovich, Hillsberg, "IBM Storage Tank—A heterogeneous scalable SAN file system," IBM Systems Journal, July 2003.

Rodeh, Teperman, "zFS - A Scalable Distributed File System Using Object Disks," IEEE Mass Storage Systems & Technologies, April 2003.

Callaghan, " NFS Illustrated". Addison-Wesley Publishing Company, 1999, ISBN 0-20-132570-5."

Shepler, Callaghan, Robinson, Thurlow, Beame, Eisler, Noveck, "Network File System (NFS) version 4 Protocol", IETF RFC 3530, April 2003.

Huang. "Understanding Gigabit Ethernet Performance on Sun Fire™ Systems", Sun BluePrints™ OnLine, February 2003.

Benner. "Fibre channel: Gigabit Communications and I/O for Computer Networks". McGraw-Hill, 1996

Ruwart. "Performance Characterization of Large and Long Fibre Channel Arbitrated Loops". Mass Storage Systems, 1999. IEEE Symposium, 1999

Heath, Yakutis. "High-Speed Storage Area Networks Using a Fibre Channel Arbitrated Loop Interconnect". IEEE Network, 2000