International Journal of Ethics in Engineering & Management Education

Website: www.ijeee.in (ISSN: 2348-4748, Volume 1, Issue 4, April 2014)

Advance Network Storage Systems (SAN & NAS)

Nagesh Salimath
Software Test Engineer,
Minds Solvit Pvt.Ltd,"Keonics IT park",
A1 Block,Opp.High Court,Ring Road,Gulbarga-585 103
nash.salimath@gmail.com

Mohammed Shoukat Ali
Pricipal Lead Engineer
Minds Solvit Pvt.Ltd,"Keonics IT park",
A1 Block,Opp.High Court,Ring Road,Gulbarga-585 103
sisali7@gmail.com

Abstract: **From this paper we are going to discuss about the SAN and NAS over the network. The storage networking is a dedicated, high-speed network established to directly connect storage peripherals. It allows the direct communication between the storage devices and the client machines on the network. It provides a efficient solution for historical view of server attached storage. The emergence of storage networking brings a new trend in today's networking technology. The two-major components of storage networking are SAN and NAS. SAN is a storage area network. This is a high performance network deployed between the servers and storage devices. Typically SAN uses the fibre channel technology. It provides the various advance facilities over the LAN and WAN. NAS is a network-attached storage. It is a storage element, which is directly connected to network to provide file access to the servers. NAS is generally used on TCP/IP networks. NAS devices are mainly used for security purposes.**

Keywords: SAN, NAS, TCP/IP.

1. INTRODUCTION

Since the first computer was developed, there has been an increasing demand to make them faster, cheaper, and more applicable to everyday lives. The early mainframe computers have evolved from large, centralized systems to more nimble, enterprise-class servers brought on by cheaper and more efficient computing technologies. In turn, networking technologies have had an effect on the evolution of computing platforms. The maturation of two technologies and the increasing hunger for computer processing power and associated data has driven the need for faster, more accessible data storage techniques and the advent of storage networking. To achieve the best use of the storage there are two major components of storage is available those are Storage Area Network (SAN) and Network Attached Storage (NAS).

2. STORAGE AREA NETWORK (SAN)

Fig 1: SAN Architecture.

A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. The term SAN is usually (but not necessarily) identified with block I/O services rather than file access services.

a. Component of SAN

When the term SAN is used in connection with Fibre Channel technology, use of a qualified phrase such as "Fibre Channel SAN" is encouraged. According to this definition an Ethernet-based network whose primary purpose is to provide access to storage elements would be considered a SAN. SANs are sometimes also used for system interconnection in cluster.
A SAN (storage area network) connects a group of servers (or hosts) to their shared storage devices (such as disks, disk arrays and tape drives) through an interconnection fabric consisting of hubs, switches and links.

b. Emergence of SAN

The emergence of storage area networks (SANs) has created the need for new storage management tools and capabilities. While SANs provide many benefits such as lower cost of ownership and increased configuration flexibility, SANs are more complex than traditional storage environments. This inherent complexity associated with storage area networks creates new storage management challenges.

c. Technology used by SAN

The prominent technology for implementing storage area networks is Fibre Channel. Fibre Channel technology offers a variety of topologies and capabilities for interconnecting storage devices, subsystems, and server systems. These varying topologies and capabilities allow storage area networks to be designed and implemented that range from simple to complex configurations. Due to the potential complexity and diverse configurations of the Fibre Channel SAN environment, new management services, policies, and capabilities need to be identified and addressed.

A SAN environment typically consists of four major classes of components. These four classes are:

- · End-user platforms such as desktops and/or thin clients;
- · Server systems;
- · Storage devices and storage subsystems;
- · Interconnect entities.

3. NETWORK ATTACHED STORAGE (NAS)

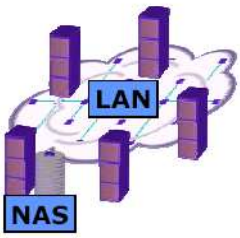

Fig:2 NAS Architecture.

A term used to refer to storage element that connect to a network and provide file access services to computer systems. Abbreviated NAS. A NAS Storage Element consists of an engine, which implements the file services, and one or more devices, on which data is stored. NAS elements may be attached to any type of network. When attached to SANs, NAS elements may be considered to be members of the SAS class of storage elements.

A class of systems that provide file services to host computers. A host system that uses network-attached storage uses a *file system device driver* to access data using file access protocols such as NFS or CIFS. NAS systems interpret these commands and perform the internal file and device I/O operations necessary to execute them.

Network-attached storage (NAS) is a concept of shared storage on a network. It communicates using Network File System (NFS) for UNIX® environments, Common Internet File System (CIFS) for Microsoft Windows environments, FTP, http, and other networking protocols. NAS brings platform independence and increased performance to a network, as if it were an attached appliance. A NAS device is typically a dedicated, high-performance, high-speed communicating, single-purpose machine or component. NAS devices are optimized to stand-alone and serve specific storage needs with their own operating systems and integrated hardware and software. Think of them as types of plug-and-play appliances, except with the purpose of serving your storage requirements. The systems are simplified to address specific needs as quickly as possible-in real time. NAS devices are well suited to serve networks that have a mix of clients, servers, and operations and may handle such tasks as Web cache and proxy, firewall, audio-video streaming, tape backup, and data storage with file serving. These highly optimized servers enable file and data sharing among different types of clients. It also defines NAS benefits with respect to storage area networks (SANs).

a. Managing network attached storage

The fundamental goal of our network-attached storage research is to enable scalable storage systems while minimizing the file manager bottleneck. One solution is to use homogeneous clusters of trusted clients that issue unchecked commands to shared storage. However, few environments can tolerate such weak integrity and security guarantees. Even if only for accident prevention, file protections and data/metadata boundaries should be checked by small number of administrator-controlled file manager machines. To provide this more appropriate degree of integrity and security, we identify two basic architectures for direct network-attached storage.[2]

The first, NetSCSI, makes minimal changes to the hardware and software of SCSI disks, while allowing NetSCSI disks to send data directly to clients, similar to the support for third-party transfers already supported by SCSI. Drives' efficient data-transfer engines ensure that each drive's sustained bandwidth is available to clients. Further, by eliminating file management from the data path, manager workload per active client decreases. Cryptographic hashes and encryption, verified by the NetSCSI disks, can provide for integrity and privacy. The principal limitation of NetSCSI is that the file manager is still involved in each storage access; it translates namespaces and sets up the third-party transfer on each request.[2]

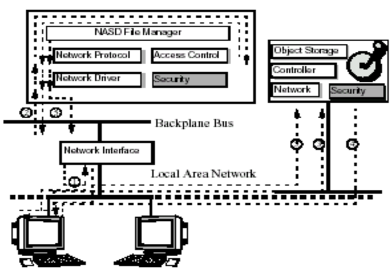

Figure 3. Network attached secure disk

The second architecture, Network-Attached Secure Disks (NASD), relaxes the constraint of minimal change from the existing SCSI interface. The NASD architecture provides a command interface that reduces the number of client-storage interactions that must be relayed through the file manager,

International Journal of Ethics in Engineering & Management Education
Website: www.ijeee.in (ISSN: 2348-4748, Volume 1, Issue 4, April 2014)

thus avoiding a file manager bottleneck without integrating file system policy into the disk. In NASD, data-intensive operations (e.g., reads and writes) go straight to the disk, while less-common policy making operations (e.g., namespace and access control manipulations) go to the file manager.

b. Network support for NAS

The success of the NASD architecture depends critically on its networking environment. Clearly, support for high-bandwidth, large data transfers is essential. Unfortunately, traditional client-server communication paths do not support efficient network transport. For example, measurements of our NASD prototype drive (running DCE/RPC over UDP/IP) show that non-cached read or write requests can easily be serviced by modest hardware. However, requests that hit in the drive cache incur order-of magnitude increases in service time due to the NASD drive and client both spending up to 97% of their time in the network stack [Gibson98]. This problem with traditional protocol stacks forces network-attached storage to explore alternative techniques for delivering scalable bandwidth to client applications. Several other network issues are also important to consider in a NASD environment.

4. STORAGE AREA NETWORK VS NETWORK ATTACHED STORAGE

Some people confuse NAS with storage area networks (SANs); after all NAS is SAN spelled backwards. The technologies also share a number of common attributes. Both provide optimal consolidation, centralized data storage, and efficient file access. Both allow you to share storage among a number of hosts, support multiple different operating systems at the same time, and separate storage from the application server. In addition, both can provide high data availability and can ensure integrity with redundant components and redundant array of independent disks (RAID).

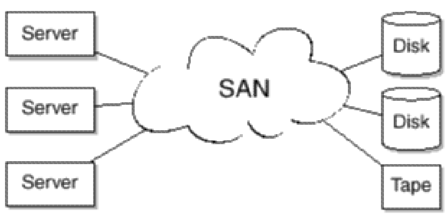

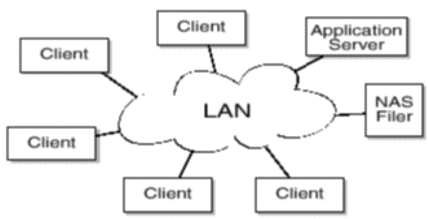

Others may view NAS as competitive to SAN, when both can, in fact, work quite well in tandem. Their differences? NAS and SAN represent two different storage technologies and they attach to your network in very different places. NAS is a defined product that sits between your application server and your file system. SAN is a defined architecture that sits between your file system and your underlying physical storage. A SAN is its own network, connecting all storage and all servers. For these reasons, each lends itself to supporting the storage needs of different areas of your business.

a. NAS Think network user

NAS is network-centric. Typically used for client storage consolidation on a LAN, NAS is a preferred storage capacity solution for enabling clients to access files quickly and directly. This eliminates the bottlenecks users often encounter when accessing files from a general-purpose server. NAS provides security and performs all file and storage services through standard network protocols, using TCP/IP for data transfer, Ethernet and Gigabit Ethernet for media access, and CIFS, http, and NFS for remote file service. In addition, NAS can serve both UNIX and Microsoft Windows users seamlessly, sharing the same data between the different architectures. For client users, NAS is the technology of choice for providing storage with unen-cumbered access to files.

Although NAS trades some performance for manageability and simplicity, it is by no means a lazy technology. Gigabit Ethernet allows NAS to scale to high performance and low latency, making it possible to support a myriad of clients through a single interface. Many NAS devices support multiple interfaces and can support multiple networks at the same time. As networks evolve, gain speed, and achieve latency (connection speed between nodes) that approaches locally attached latency, NAS will become a real option for applications that demand high performance.

b. SAN – Think back-end / computer room storage

A SAN is data-centric - a network dedicated to storage of data. Unlike NAS, a SAN is separate from the traditional LAN or messaging network. Therefore, a SAN is able to avoid standard network traffic, which often inhibits performance. Fibre channel-based SANs further enhance performance and decrease latency by combining the advantages of I/O channels with a distinct, dedicated network.
SANs employ gateways, switches, and routers to facilitate data movement between heterogeneous server and storage environments. This allows you to bring both network connectivity and the potential for semi-remote storage (up to 10 km distances are feasible) to your storage management efforts. SAN architecture is optimal for transferring storage blocks. Inside the computer room, a SAN is often the

International Journal of Ethics in Engineering & Management Education

Website: www.ijeee.in (ISSN: 2348-4748, Volume 1, Issue 4, April 2014)

preferred choice for addressing issues of bandwidth and data accessibility as well as for handling consolidations.

Due to their fundamentally different technologies and purposes, you need not choose between NAS and SAN. Either or both can be used to address your storage needs. In fact, in the future, the lines between the two may blur a bit according to Evaluator Group, Inc. analysts. For example, down the road you may choose to back up your NAS devices with your SAN, or attach your NAS devices directly to your SAN to allow immediate, nonbottlenecked access to storage.

5. THE DIFFERENCE TABLE OF SAN & NAS

	SAN	NAS
Protocol	• Fibre Channel • Fibre Channel-to-SCSI	• TCP/IP
Applications	• Mission-critical transaction-based database application processing • Centralized data backup • Disaster recovery operations • Storage consolidation	• File sharing in NFS and CIFS • Small-block data transfer over long distances • Limited read-only database access
Advantages	• High availability • Data transfer reliability • Reduced traffic on the primary network • Configuration flexibility • High performance • High scalability • Centralized management • Multiple vendor offerings	• Few distance limitations • Simplified addition of file sharing capacity • Easy deployment and maintenance

Application of SAN:

i. Data mining on PC cluster
ii. Disk to disk copy function is very efficient using PC connected SAN clusters than the PC connected LAN clusters.
iii. Distributed file system.
iv. Distributed file system using SAN is having large number of advantages over a historically storage attached systems. In this the data transfer takes place without file bottlenecks.

6. APPLICATION OF NAS:

Integrity and security management

7. CONCLUSION

The emergence storage networking provides a very high performance networks. The major components of the storage networking are SAN and NAS. These components performs various functions as follows

SAN:

- Data transfer reliability
- Configuration flexibility
- High performance
- High scalability
- Centralized management of storage

NAS:

- Simplified addition of file sharing capacity
- Easy maintenance

REFERENCES

[1]. Garth A. Gibson and Rodney Van Meter “Network attached storage architecture ”November 2000

[2]. David f. nagle, Gregory R. ganger, Jeff butler, Garth Goodson and Chris sabol “Network supports for network-attached storage” 18-20 august 1999

[3]. Masato oguchi and Masaru kitsuregawa “Data mining on PC cluster connected with storage area network” May 2000

[4]. Brad Kline “Distributed file systems for storage area network”