



Recent Trends on DAS and SAN Protocols

M. Rao Batchanaboyina¹, Namoori Prakash Babu², Devarapalli Anandam³, Ramaiah Thalluri⁴

¹Assoc. Prof, Department of CSE, ^{2,4}Asst. Prof, Department of ECE, ³Asst. Prof, Department of CSE
^{1,2,3,4}PACE Institute of Technology & Sciences, Ongole, JNTUK, Andhra Pradesh, India

Abstract— This paper provides an overview of key technologies that have evolved around data storage and storage networking. And it focuses on analyzing the system architectures of the different building blocks of storage networks. With the help of SAN and DAS, storage can be efficiently utilized. Both technologies alleviate the need to eliminate direct attached storage to aid more flexible storage access. SAN and DAS use open industry-standard network protocols to provide storage facility. DAS uses the IDE, SCSI, SATA, etc protocols and SAN uses the Fibre Channel, iSCSI, FoE, etc protocols to provide the storage facility. And this paper provides the basics and survey on new technologies of these protocols.

Keywords— SAN Technology, DAS Technology, iSCSI, FoE, DAS Architecture

I. INTRODUCTION

In recent years, enterprise data storage has seen explosive growth in demand from users. This growth is driven by increasingly more sophisticated applications that generate more rich and numerous quantities of content data, and an increasingly larger number of users/consumers of this rich content data. The rapid advancement of networking technology both in the LAN and the WAN has enable new applications that generate large demands on data storage. Humans and applications use data in various format including texts, video, audio and images. It is very much vital to the network storage system that should be high capacity, high availability and scalable. Storage-Area Networks (SANs) constitute an important element in modern IT infrastructures due to their client management of the underlying storage capabilities in environments shared by several servers, applications, users or even organisations. These new applications drive the data storage demand in the following areas:

- Availability
- Scalability
- Capacity
- Cost
- Performance

II. STORAGE MODELS

The three categories of data storage technologies including Direct Attached Storage (DAS), NAS (Network Attached Storage), and SAN (Storage Area Networks). Out of this comparison will appear the benefit of sharing storage resources over the network, and how different schemes can be used to accomplish the task of sharing storage resources.

1. **DAS** is a block device from a disk which is physically [directly] attached to the host machine.
 - You must place a file system upon it before it can be used.
 - Technologies to do this include IDE, SCSI, SATA, etc.
2. **SAN** is a block device which is delivered over the network.
 - Like DAS you must still place a file system upon it before it can be used.
 - Technologies to do this include Fibre Channel, iSCSI, FoE, etc.
3. **NAS** is a file system delivered over the network.
 - It is ready to mount and use.
 - Technologies to do this include NFS, CIFS, AFS, etc.

III. DAS TECHNOLOGY

DAS devices can utilize traditional spinning hard disk drives or solid state media. Hard disk drives (HDDs) are less expensive than solid state drives (SSDs) on a per gigabyte basis. However, they aren't as fast as SSDs. The majority of DAS devices sold today use hard disk drives, although SSDs and hybrid devices are becoming more popular, particularly for applications where performance is paramount, such as Big Data analytics. Conversely, SSDs offer the fastest performance, but they are also the most expensive. However, in recent years, they have become more affordable. In addition, because SSDs don't have a spinning disk, they last longer and offer better reliability than HDDs. Many vendors now offer SSDs as an option on mid- to high-end servers for part or all of the storage capacity.

The simplest storage is a Hard Disk connected to your computer/server. Hard disk drives need to communicate somehow, and they need to use a well defined physical interface and protocol in order for your computer to understand them, the most common interface and protocol employed today for PC is the S-ATA or Serial ATA or Serial Advanced Technology Attachment. So a hard disk which is physically connected through the same interface to your PC is identified as a **Direct Attached Storage** or **DAS** for short. To make it clear, any **block** device which is directly connected to your

system makes a DAS, even an USB stick is a DAS (but the interface is USB in this case). And unless you access your disk by block numbers you need a file system on top of it in order to put it at good use.

A typical DAS configuration consists of a computer that is directly connected to one or several hard disk drives (HDDs) or disk arrays. Standard buses are used between the HDDs and the computers, such as SCSI, ATA, Serial-ATA (SATA), or Fibre Channel (FC). Some of the bus cabling definitions allow for multiple HDDs to be daisy chained together on each host bus adapter (HBA), host channel adapter, or integrated interface controller on the host computer.



Figure 1: DAS diagram

Two examples of DAS systems are given in Figure 2. In the first example, 3 SCSI HDDs are attached to the host computer via a daisy chain of SCSI cabling. The second example uses Fibre Channel cabling to connect the host computer and a RAID/JBOD storage system together.

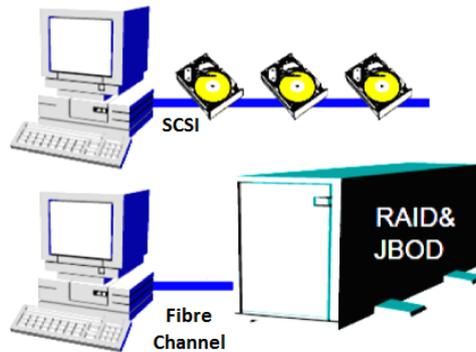


Figure 2: Direct Attached Storage

DAS is a widely deployed technology in enterprise networks. It is easy to understand, acquire and install, and is low cost. For home PC and small enterprise network applications, DAS is still the dominant choice, as the low-end requirements for growth in capacity, performance and reliability can be easily addressed by the advancements in HDD and bus technologies. It is well suited to the purpose of attaching data storage resources to a computer or a server when capacity, administration, backup, high-availability, high performance are not key requirements

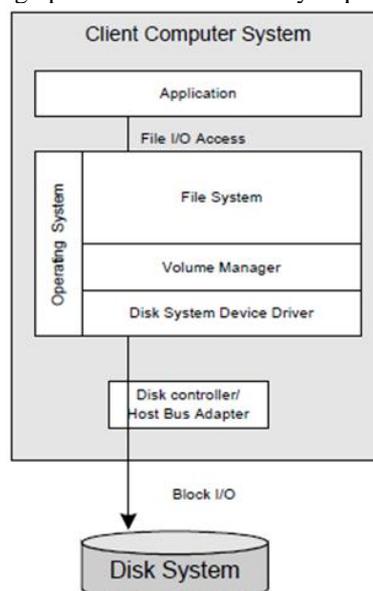


Figure: 3 DAS Software Architecture.

The past few years have seen 2x sequential increase in HDD capacity per year, while maintaining the low cost point of HDDs targeting the personal computer market. The advent of Ultra-ATA, SATA, SATA-2, and Serial Attached SCSI (SAS) and FC bus standards alleviates the performance bottleneck on the bus interface. The quality of the HDDs has also much improved over the years. These technology advancements have helped DAS systems address the requirements of low-end data storage users.

The software layers of a DAS system are illustrated in Figure 3. The directly attached storage disk system is managed by the client operating system. Software applications access data via file I/O system calls into the Operating System.

The file I/O system calls are handled by the *File System*, which manages the directory data structure and mapping from files to disk blocks in an abstract logical disk space.

The *Volume Manager* manages the block resources that are located in one or more physical disks in the *Disk System* and maps the accesses to the logical disk block space to the physical volume/cylinder/sector address.

The *Disk System Device Driver* ties the *Operating System* to the *Disk controller* or *Host Bus Adapter* hardware that is responsible for the transfer of commands and data between the client computer and the *disk system*.

The file level I/O initiated by the client application is mapped into block level I/O transfers that occurred over the interface between the client computer and the disk system. The efficiency of the storage resource is low, as the storage capacity is bound to a given computer/server. The distributed nature of the storage resource not only means more content replication, but also means the free resources on one computer cannot be used by another computer/server whose disk space is running low.

A typical DAS system is made of a data storage device (for example enclosures holding a number of hard disk drives) connected directly to a computer through a host bus adapter (HBA). Between those two points there is no network device (like hub, switch, or router), and this is the main characteristic of DAS.

The main protocols used for DAS connections are ATA, SATA, eSATA, SCSI, SAS, and Fibre Channel.

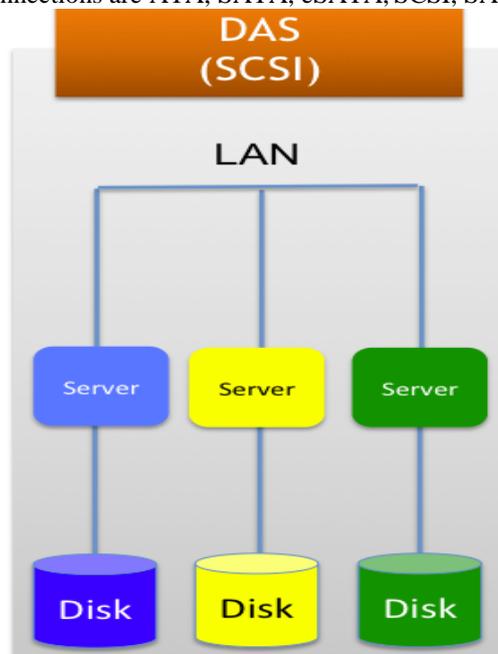


Figure : 4 DAS diagram with SCSI protocol

IV. SAN TECHNOLOGY

Storage Area Networks Storage area networks (SANs) represent the latest stage in the evolution of mass data storage for enterprises and other large institutions. In host-centric environments, storage, as well as applications, was centralized and centrally managed. With the advent of client/server environments, information that was previously centralized became distributed across the network. The management problems created by this decentralization are addressed in two principal ways: network attached storage (NAS), where storage devices are directly attached to the LAN, and SANs. Composed of servers, storage devices (tapes, disk arrays), and network devices (multiplexers, hubs, routers, switches, and so on), a SAN constitutes an entirely separate network from the LAN (see Figure 3-3). As a separate network, the SAN can relieve bottlenecks in the LAN by providing the resources for applications such as data mirroring, transaction processing and backup and restoration.

A. SAN Protocols:

- **SCSI** – The SCSI is a low level protocol is used encapsulated into secondary or transport protocols to send commands to controllers or disks
- **SAS** - Serial Attached SCSI This allows to create entire SANs as SAS supports routing and addressing
- **FC** - Fibre Channel
- **iSCSI** - Internet SCSI Or simply SCSI over Ethernet (in IP networks)

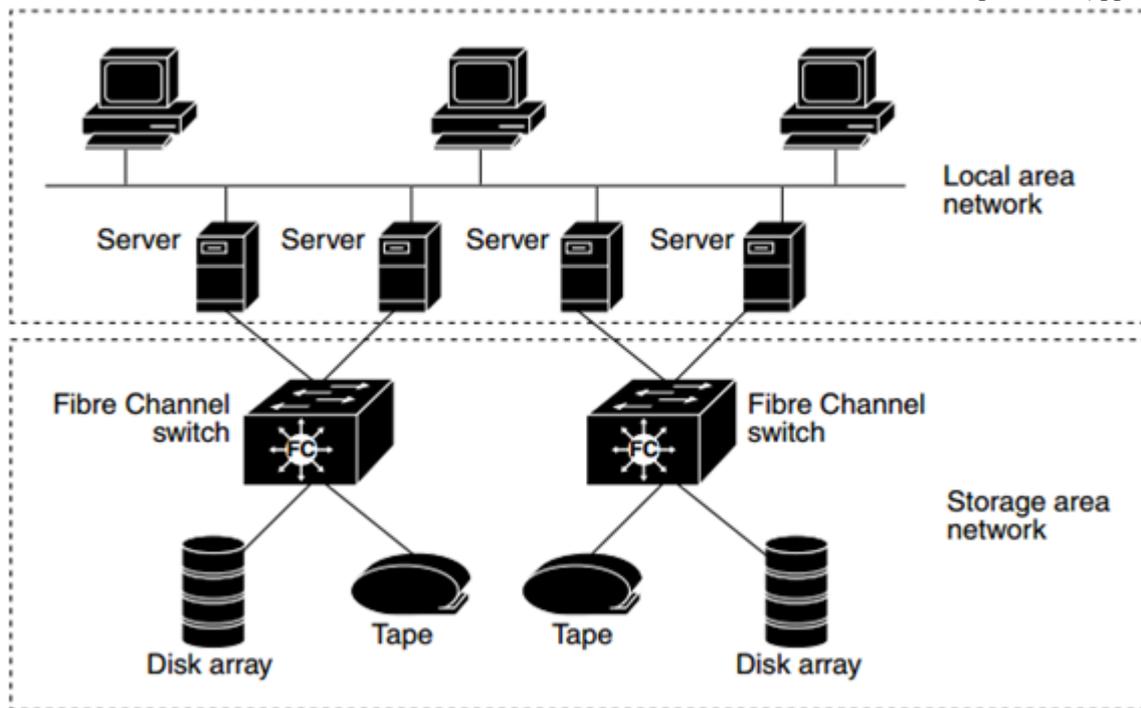


Figure 5: SAN Architecture

A SAN can be composed of very different hardware, but can usually be broken down into various components:

Physical Disks: they, of course, archive the data. Enterprise-level disks are used, which means they usually have lower per-disk capacity, but much higher performance and reliability; also, they are a lot more expensive than consumer-class disks. The disks can use a wide range of connections and protocols (SATA, SAS, FC, etc.) and different storage media (Solid-State Disks are becoming increasingly common), depending on the specific SAN implementation.

Storage Controllers/Processors: these manage disk I/O, RAID and caching. Again, enterprise-level controllers are used, so they have much better performance and reliability than consumer-class hardware. They can, and usually are, configured in pair for redundancy.

Fibre Channel: a technology which uses fibre-optics for high-speed connections to shared storage. It includes host bus adapters, fibre-optic cables and FC switches, and can achieve transfer speeds ranging from 1 Gbit to 20 Gbit. Also, multipath I/O can be used to group several physical links together, allowing for higher bandwidth and fault tolerance.

iSCSI: an implementation of the SCSI protocol over IP transport. It runs over standard Ethernet hardware, which means it can achieve transfer speeds from 100 Mbit (generally not used for SANs) to 100 Gbit. Multipath I/O can also be used (although the underlying networking layer introduces some additional complexities).

Fibre Channel over Ethernet (FCoE): a technology in-between full FC and iSCSI, which uses Ethernet as the physical layer but FC as the transport protocol, thus avoiding the need for an IP layer in the middle.

Infini Band: a very high-performance connectivity technology, less used and quite expensive, but which can achieve some *impressive* bandwidth.

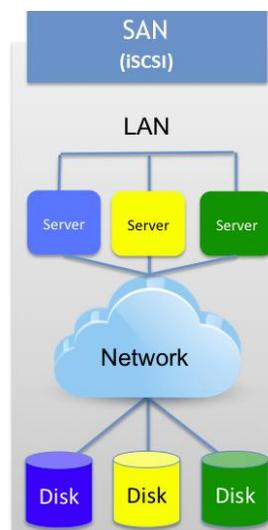


Figure 6: SAN with iSCSI protocol

A number of types of interfaces have been used to connect servers to devices in a SAN. The most prevalent is IBM's Enterprise System Connection (ESCON), a 17-MBps half-duplex protocol over fibre. Fibre Channel, on which IBM's FICON is based, is also frequently employed in SANs and has a much higher capacity than ESCON. Both technologies, however, have significant distance limitations. For example, the standard maximum distance without repeaters is around 3 km (1.9 mi) for half duplex ESCON and around 10 km (6.2 mi) for full duplex 100-MBps Fibre Channel. There is performance degradation as distances increase beyond these numbers. This distance limitation can be overcome by transporting data between one or more enterprise locations and one or more SANs over the optical layer using DWDM.

V. CONCLUSION

Finally this paper, we have summarized the important technologies for both DAS and SAN. We also provided brief explanation for DAS and SAN protocols and recent research and development of these protocols. This paper provides details and storage network protocols information to the storage network beginners.

REFERENCES

- [1] <http://serverfault.com/questions/425335/what-is-a-storage-area-network-and-which-benefits-does-it-have-over-different-s?lq=1>
- [2] <http://serverfault.com/questions/81723/what-is-the-difference-between-san-nas-and-das>
- [3] John Gantz and David Reinsel , — The Digital Universe in 2020: Big
- [4] Data, Bigger Digital Shadows, and Biggest Growth in the Far East, White Paper, EMC, December 2012.
- [5] De-Zhi Han, —SNINS:A Storage Network Integrating NAS and SAN, IEEE, Vol 1, page 488 – 493, ISBN- 0-7803-9091-1, 18-21 Aug. 2005
- [6] Guo liqing, Huo limin, Cui liwei, —Construction of Virtual Private Network for remote disaster-tolerant channel based on SG186 Network, IEEE, Electrical and Control Engineering (ICECE), 2011 International Conference, pp 2077 – 2080, Sept. 2011
- [7] Roger, —Future Hard Disk Drive System, Journal of Magnetism and Magnetic Materials, pp. 551-556, 2009
- [8] Xu Xuedong , Su Jiyuan , —Research and Implementation of iSCSIbased SAN Static Data Encryption System, Computer Science and Network Technology, 2nd International Conference, IEEE, pp 257– 260, Dec. 2012
- [9] Brocade Communications Systems, White Paper, —Comparing Storage Area Networks And Network Attached Storage, May 2013
- [10] Bo Li, Panyong Zhang, Zhigang Huo, Dan Mang. —Early Experience with Write-Write Design of NFS over RDMA, Networking, Architecture, and Storage NAS. IEEE International Conference, pp 303 – 308, July 2009
- [11] Claudio DeSanti, Joy Jiang, —FCoE in Perspective, International Conference on Advanced Infocomm Technology pp 29-31, July 2008.
- [12] Renato Recio et al., —A Case for Convergence Enhanced Ethernet: Requirements and Applications, ICC Proceedings, 2008
- [13] Yueping Cai, Yao Yan, Zhenghao Zhang, and Yuanyuan Yang, — Survey on Converged Data Center Networks with DCB and FCoE: Standard and protocols, Network, IEEE , Vol. 27, pp 27-32, July- August 2013
- [14] Jiwu Shu, Bigang Li, and Weimin Zheng , —Design and Implementation of an SAN System Based on the Fiber Channel Protocol , IEEE, Vol. 54, pp. 439 – 448, April 2005
- [15] Yong Chen, Lionel M.Ni and Mingyao Yang , — A Storage Cluster Architecture using Network Attached Storage Devices, Parallel and Distributed Systems, Ninth International Conference, IEEE, pp 301 – 306, Dec. 2002