IP SAN: Low on Fibre

Smriti Bhagat
CS Dept, Rutgers University

Abstract

The era of Fibre Channel interoperability is slowly dawning. In a world where Internet Protocol (IP) dominates local and wide area networks, and data storage requirements grow unabated, it seems inevitable that these two forces converge. Internet Protocol Storage Area Networks (IP SANs) [4] unite storage and IP networking enabling IP and Ethernet infrastructure to be used for expanding access to SAN storage and extending SAN connectivity across any distance.

This paper provides an overview of storage networking and discusses its need. It goes deeper into the storage area network technology and analyzes its two types: Fibre Channel SAN [1] (FC SAN) and IP SAN. The paper then compares the two, making a case for IP SAN and highlighting that IP SANs will take the centre stage, possibly replacing Fibre Channel SANs completely.

1 Introduction

The exploding volume of information being processed and stored today increases the demand for new techniques to access massive amounts of data quickly and reliably. The information that was centralized in a host-centric environment is now dispersed across networks. The amount of data stored on removable storage doubles every year. Studies [4] show that from 5.6 exabytes[1] of data stored in the year 2000, the amount is expected to be 99.5 exabytes by 2005. Also that 67 percent of all storage will be networked. With such a tremendous increase in storage and networked storage requirements, the Storage Area Network (SAN) paradigm has emerged as the most desirable storage solution.

This paper focuses on comparing the two storage area network technologies: Fiber Channel SAN and Internet Protocol SAN. Section 1 describes the fundamental features of SANs compared with other storage subsystems. Section 2 discusses FC SANs in detail followed by Section 3 which discusses IP SANs. The discussion highlights the promising IP SAN trends and is followed by the conclusions

1.1 Storage Subsystems

Traditional storage subsystems like Server Attached Storage (SAS) [1] and Network Attached Storage (NAS) [1] use disks directly attached to servers over a Small Computer Systems Interface (SCSI) bus. However, these failed to meet the challenges posed by the rapidly expanding storage needs.

The server attached storage model is based on each server having its own SCSI connection to the storage. This SCSI has its own limitations namely, the maximum length of the cable is 25m and the maximum data transfer rate achievable on this bus is 80 Mbps. Another serious drawback of SAS is the unavailability of data during a server failure.

[1] One Exabyte = 2^{60} bytes

Network attached storage is typically a data server on a network that provides file storage accessed via the network. The overall performance of a NAS decreases with new storage devices added to the network. Also, servers backing up data on the storage devices and client-server interactions both using the same LAN, leads to bandwidth congestion and hence throughput degradation.

Storage area networks on the other hand, link storage devices (disks, disk arrays etc) to create a pool of storage that users can access. This relieves the network of the high intensity I/O traffic. SAN technology has opened up new opportunities in simplified connectivity, scalability, and cost and capacity manageability.

1.2 Storage Area Networks: An Overview

Keeping pace with the growing storage demands and providing solutions to the limitations of traditional subsystems, storage area networks have emerged as the leader of storage solutions. The fundamental premise of a SAN is the ability of host servers to share storage resources in the network, enabling high degree of utilization.

A Storage Area Network is essentially a high speed dedicated, scalable and secure information infrastructure, which enables universal storage connectivity [1]. Further, storage devices on a SAN may be heterogeneous in nature, thus allowing *any-to-any* connectivity. SANs support true data sharing [6] by allowing access to common data from multiple computer platforms and servers. *(See Figure 1)*

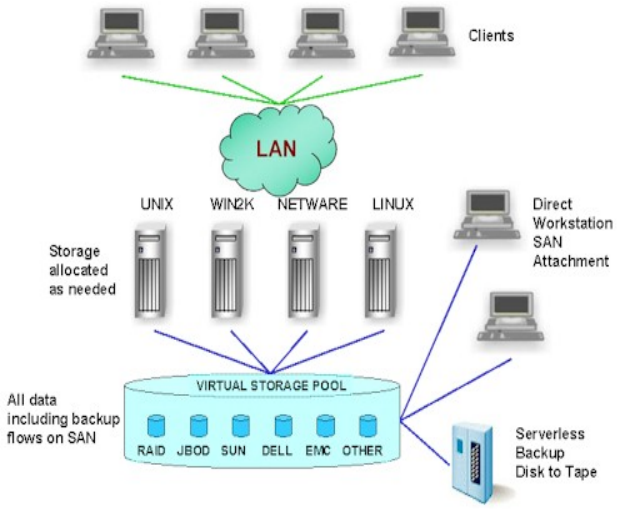

Fig 1: A typical SAN. It consists of clients on a network, a set of servers running different operating systems and a pool of storage devices including disk arrays. Source: IBM Redbooks [2]

SANs allow the separation of storage traffic from general network traffic. This is beneficial from a security, performance and management standpoint. Increased bandwidth can be achieved by providing additional I/O paths. Multi-pathing and mirroring (using RAID disk arrays) enable higher availability. To provide unlimited storage, SANs can be seamlessly scaled by adding more disk arrays or other storage devices.

Storage consolidation [6] allows a storage device to be attached to multiple servers concurrently. Once this device is attached to multiple servers, it can be partitioned or shared between servers facilitating storage sharing.

1.3 SAN technologies

SAS has traditionally used SCSI technology to link storage devices. To overcome the transmission speed and distance limitations of SCSI, Fibre Channel [9] protocol was

adopted as the de facto standard for implementing SANs.

Fibre Channel technology allows SANs to transmit data faster and over larger distances than SCSI, using fibre optic cable or copper cable as the medium. It exhibits higher scalability and performance.

IP Storage Networking [4] refers to computer systems and storage elements that are connected via the existing IP networks. IP SANs have redefined storage area networks and are proving to be highly reliable, scalable and cost-effective.

Among IP storage protocols, the most popular one is Internet SCSI (iSCSI), which enables transmission of SCSI commands over TCP/IP. In iSCSI, the SCSI command set is mapped to IP for transmission of block-level data. Like Fibre Channel, iSCSI preserves the SCSI command set. Native iSCSI removes Fibre Channel out of the picture, and runs over a standard Ethernet connection.

2 Fibre Channel SAN

2.1 Motivation

Fibre Channel [9] is the name given to a set of standards that define mechanisms for connecting network nodes using serial communication over copper or optical cable. Fibre Channel was developed to overcome the shortcomings of the SCSI infrastructure, and to provide high-speed connections between servers and storage devices. FC SANs were designed to use this fast serial transport system that would move almost any kind of data reliably over long distances.

2.2 FC SAN: the up side

FC SAN was born out of a need for a fast, dedicated and robust network for storage. It has taken the spotlight, promising greater *performance* and connectivity distances than SCSI. FC technology and product deployment has evolved from 1Gbps to 2Gbps and 10Gbps has been successfully tested.

Fibre Channel speeds can be attributed to its ability to communicate with other network systems using *channel technology*. This means that using Fibre Channel, a point-to-point connection can be established between two devices creating a *fast and low overhead* connection. Once a connection is established a series of data sequences are transmitted and as they propagate through the fabric without the overhead of checking for routing. Again, since the connection is established at the lower layers, there is less overhead to transport the data to its destination.

Fibre Channel provides a reliable, low latency, high throughput transport mechanism. Fibre Channel Protocol (FCP) provides high end-to-end performance as it uses credit-based *congestion control* and implements *zero-copy* send and receive semantics. The congestion control mechanism ensures that frames are not dropped at switches during congestion. Zero copy ensures a low CPU utilization.

It took a while for the storage industry to settle with the Fibre Channel standards, with the result that now virtually any I/O load can run through a Fibre Channel adapter with little impact on CPU utilization [3]. High performance is one of the key features of FC SAN.

While Fibre Channel SANs promise performance, they come with a price. The following section discusses the pitfalls of FC SANs.

2.3 FC SAN: limitations

Fibre Channel SAN topology is the high-speed storage option at a top dollar price. Its *cost* is its fundamental limitation. Fibre Channel SANs have always been prohibitively expensive, both from an infrastructure and a support perspective. This has limited the deployment of FC SANs only to large enterprises.

To adopt FC SANs, a new network infrastructure has to be installed. Fibre Channel requires a *separate fiber-optic network* for the SAN to run. Plumbing and installing new Fibre Channel pipes may turn out to be too cumbersome.

FC SANs also require administrators to learn a new set of network management skills. Skilled labor and the cost of training are also considerations before deploying FC SANs.

Fibre Channel was designed to operate in secure environments and hence its *security infrastructure* is not as well defined.

Fibre Channel was not designed to be a wide area network protocol and is not scalable with respect to *distance*. It is limited to just 50 miles, which is an inherent limitation of optic fiber. Hence SANs are geographically restricted. This is a key drawback of FC SANs. If the data backups and the data are in the same geographical location which faces a catastrophe, *disaster recovery* may not be possible. Especially after 9/11, when organizations are spending a fortune for data storage and backup, it would defeat the purpose if we are unable to keep copies of data at distant locations.

Finally, the *interoperability* record of Fibre Channel devices from different vendors has been a cause of concern.

All these limitations of FC SANs urge us to carefully examine the other options, primarily Internet Protocol SAN.

3 Internet Protocol SAN

3.1 Motivation

The notion of placing storage traffic on IP networks has been well studied. IP was originally designed to support the US government in maintaining communications between R&D facilities, military bases, and other locations. Keeping in mind the packet loss nature of existing networks, very efficient retransmission mechanisms were built into TCP. Thus error recovery and retransmission mechanisms of TCP are more robust compared to other protocols.

IP SANs were developed to leverage the existing IP network management protocols and tools.

3.2 IP SAN: In limelight

Long distance networking is standardizing on the Internet Protocol (IP) due to its inherent routing capabilities and its industry acceptance. Advances in Ethernet technology, along with the desire to simplify training, hardware and support, have made IP SANs possible.

The real selling point is being able to reuse existing, easily understood TCP/IP infrastructure to build SANs. With the

advances in QoS (Quality of Service) and *security*, the opportunity to share storage with the *existing infrastructure* represents a significant cost-savings in hardware, training and implementation.

TCP/IP makes use of several advanced cryptographic algorithms to implement robust authorization and authentication schemes.

TCP/IP is most widely deployed *reliable* protocol that is supported by all operating systems.

The theme of IP SANs is to get everything on a *single type of network platform*. This does not necessarily mean that storage traffic will be flowing on the same network as the general data traffic. From an implementation standpoint, an enterprise SAN exists as a separate entity from the general data network.

There are many reasons to choose IP SANs over its competing technologies. The key factor which makes IP SANs an attractive proposition is the cost. The *cost-savings* come about in several ways: Common technologies between your general data and SAN networks reduce training and personnel-acquisition costs, and the large installed base of Ethernet drives the total cost of ownership (TCO) down. This obviates the need for expensive Fibre Channel infrastructure. IP SANs can typically be deployed for around 20 percent of the cost of a fibre channel solution.

Riding the IP wave of technology development and enhancements like the introduction of *10 Gigabit Ethernet*, IP SANs have achieved a different horizon and a level of practicality and acceptance.

The advanced *scalability* of IP SANs is proven by the Internet and the World Wide Web. IP networks deliver unrestricted topologies that can scale to very large populations.

Since there are no limitations on *distance* with an IP infrastructure, it is easier to implement disaster-tolerant solutions. IP is being increasingly seen as a storage interconnect for long-haul data-replication applications because of its greater distance capabilities compared to Fibre Channel.

While Fibre Channel host-bus adapters (HBA) must be connected directly to the SAN switch [8]. An IP HBA can connect to a storage router anywhere on the Gigabit Ethernet SAN. This provides more flexibility when it comes to building a complex SAN.

Inherent multi-path and failover technology in iSCSI makes the solution *highly available* with no single point of failure.

IP SANs leverage the large pool of IP network management professionals averting the need for training requirements.

Since IP SANs are built upon existing technology, improvements to these technologies will automatically carry over into improvements in the functionality of IP SANs.

While FC SANs are suited for only data centers, IP SANs are widely applicable for primary storage, server consolidation, disaster recovery, and disk-to-disk backup applications.

Given IP's growing popularity, supporters of IP storage believe the technology will sweep away everything that came before. Bill Miller, chief technology officer at

StorageNetworks Inc., says SANs are too complex for most businesses and predicts that IP storage will eventually dominate the market, although it may take a couple of years. (As highlighted by *Figure 2*)

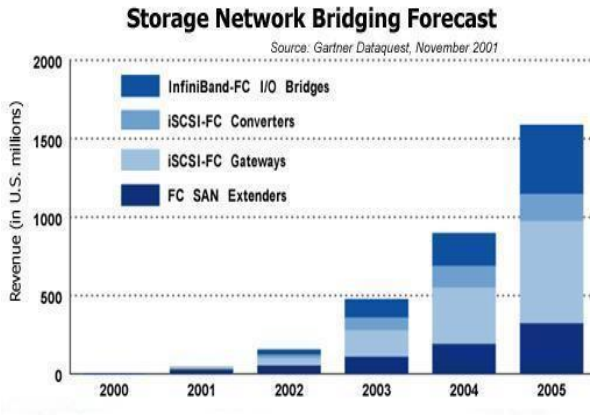

Figure 2: Shows the sales of bridging devices for FC SAN and IP SAN (iSCSI).

According to studies by International Data Corporation, IP SANs are expected to garner more than 25 percent of the global storage market by 2007.

4 Other emerging trends

4.1 Fibre Channel over TCP/IP

Fibre Channel over TCP/IP (FCIP) involves encapsulating Fibre Channel frames within TCP/IP specifically for joining two fabrics or linking Fibre Channel SANs over wide areas. Instead of a native IP connection between individual storage devices, FCIP provides IP connections between Fibre Channel SAN islands. In FCIP, if congestion takes place, all the communication between the SAN islands is affected. The performance of the FCIP link is directly dependent on the bandwidth and packet-loss characteristics of IP network used for connecting the SAN islands.

4.2 Internet Fibre Channel Protocol

Internet Fibre Channel Protocol (iFCP) is a gateway-to-gateway protocol for providing Fibre Channel fabric services to Fibre Channel end devices over a TCP/IP network. It uses TCP for congestion control, error detection and recovery. iFCP can be used in place of fibre channel fabrics in order to provide a more integrated IP storage solution. In iFCP, congestion may only affect a single communicating pair of devices due to availability of multiple routes. Thus the performance of the IP network does not affect the performance of iFCP to as great an extent as FCIP.

5 Conclusion

Storage networks are a natural extension of IP networks. Therefore, IP-based storage networks are inevitable and will come at the expense of other technologies, such as Fibre Channel.

An IP SAN can provide all the benefits of a higher-cost FC SAN: infinite scalability of storage for a server or group of servers; capacity consolidation across storage devices on the subnet; aggregation of storage allowing higher utilization; offloading of storage traffic from the main LAN; local or remote disaster recovery; disk-to-disk backup, snapshots, and failover. FC SANs offer better performance but this performance comes at a very high cost. IP SANs offer slightly less performance at a comparatively low cost. Considering tradeoffs between cost and performance, IP SAN are gaining more popularity.

However, IP SANs can achieve the benefits without the cost, infrastructure hassle, and distance limitations of a Fibre Channel SAN.

Fibre Channel SAN will eventually be replaced by IP SANs it was only a matter of time.

6 References

[1] Jon Tate, Angelo Bernasconi, Peter Mescher, Fred Scholten, Introduction to Storage Area Networks, SAN. IBM Redbooks, 2003

[2] Ravi Khattar, Mark S Murphy, Giulio John Tarella, Kjell E. Nystrom. Introduction to Storage Area Networks, SAN. IBM Redbooks, 1999

[3] K.Voruganti, and P. Sarkar, "An Analysis of Three Gigabit Networking Protocols for Storage Area Networks'. 20th IEEE International Performance, Computing, and Communications Conference", April 2001.

[4] P. Sarkar and K. Voruganti. IP Storage: The Challenges Ahead. In 19th IEEE Symposium on Mass Storage Systems, 2002

[5] University of California at Berkeley and Horison Information Strategies (http://www.horison.com)

[6] McLendon K, Babbitt C. Saying goodbye to optical storage technology. Top Health Info Manage. 2002 Aug;23(1):16-25.

[7] Barry Phillips. Have Storage Area Networks Come of Age? Computer, IEEE Computer Society. July 1998 (Vol. 31, No. 7) Pg10-12

[8] Dan McConnell. IP Storage: A Technology Overview. Dell White Paper, Oct 2001.

[9] Fibre Channel and IP SAN Integration. Proceedings of MSST2004 NASA/IEEE Conference on Mass Storage Systems and Technologies.