The Transition to Networked Storage

Jim Metzler
Ashton, Metzler & Associates

Table of Contents

1.0 Executive Summary

Driven by dramatic improvements in both performance and cost, networked storage is rapidly gaining momentum in the enterprise IT market. For example, companies that deploy a Storage Area Network (SAN) can expect to reduce the cost of storage by greater than fifty percent when compared with the cost of traditional storage solutions However, the technologies that enable the efficient and effective deployment of storage are rapidly evolving. Enterprise IT organizations that deploy a SAN need to guarantee that they will continue to leverage the benefits of this deployment, even as the technologies change. In order to provide this guarantee, these organizations must develop a storage architecture. This architecture needs to properly position storage, computing, and the storage network. The architecture needs to also articulate best practices in SAN design.

This is the first paper in a three part series that is intended to address the key challenges facing the enterprise IT organization as they strive to deploy SANs that will be able to both protect their initial investment and still leverage the ongoing improvements in storage related technologies. This paper will provide background on both the movement to deploy SANs, as well as some of the key challenges relating to ongoing SAN deployment. The second paper in the series will provide insight into best practices in SAN design. The third paper in the series will discuss some of the technological innovation that is impacting the deployment of SANs.

This paper will specifically address the following questions:

1. What role does Fibre Channel play in a SAN?
2. What factors are driving enterprises to deploy a SAN?
3. How does a SAN enable enhanced applications availability?
4. What is the business case for deploying a SAN?
5. What are the key challenges associated with ongoing SAN deployment?

2.0 The Emergence of the Storage Area Network

The demand for storage is increasing, and is expected to grow for the foreseeable future. While the estimates of the rate of growth in the demand for storage vary, the statistics suggest that storage is growing at least forty to fifty percent per year. At this rate of growth, the demand for storage doubles roughly every eighteen to twenty-four months. In order to successfully manage the cost of storage during this time of increasing demand, and at the same time improve application availability, IT professionals need to transition from traditional methods of providing storage to networked storage.

The traditional way to provide storage is called Direct Attached Storage (DAS) , and refers to a couple of slightly different approaches. For example, DAS refers to having storage that is internal to a computer, as is the case with most PCs. DAS also refers to

having hosts and storage devices attached over point-to-point connections, as shown in Figure 1.

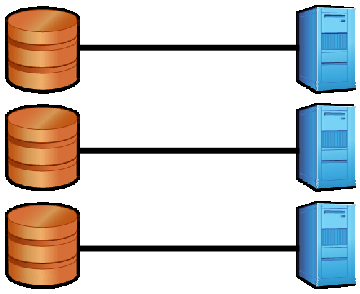

Storage Devices *Hosts*

A DAS Environment
Figure 1

In the early DAS environments, the Small Computer Systems Interface (SCSI) was the standard technique used to connect computers to storage. However, in the 1990s the speed and distance limitation of SCSI lead to the development of Fibre Channel. Currently, the most common use of Fibre Channel is to transport SCSI traffic, and in so doing, overcome SCSI's previous distance and performance limitations.

Until recently, the maximum speed that Fibre Channel could support has been 1 Gbps. However, many Fibre Channel networks are currently running at 2 Gbps and work is underway that will lead to Fibre Channel running at 10 Gbps. The primary organization that develops standards for Fibre Channel is T11 (www.t11.org). Among other functions, the 190 member Fibre Channel Industry Association (www.fibrechannel.org) is responsible for a program that tests the conformance of Fibre Channel devices to the relevant standards.

In the late 1990s, a number of factors drove the need for a networked approach to providing storage. These factors include that in a DAS environment:

- Data is local to each system and is not accessible by other systems
- Storage resources are statically assigned to a given system, which tends to waste storage resources
- Data backup is both time consuming and difficult
- Management is highly complex and resource intensive

These factors lead to the development of Storage Area Networking. As shown in Figure 2, a Storage Area Network (SAN) consists of hosts and storage devices, connected via one or more Fibre Channel switches.

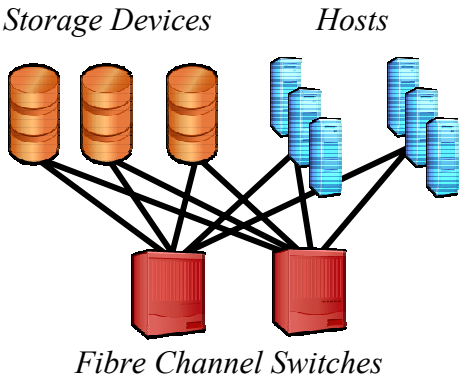

A SAN Environment
Figure 2

One of the very clear advantages of a SAN is the ability to flexibly scale the amount of storage that is available to a company's key applications. In particular, in a DAS environment, the only way to increase storage is to increase the storage that is attached to a given host. However, that storage is only available to applications that run on that particular host. If the IT organization wanted to ensure that there was increased storage for all of the applications, it would have to increase the storage that is attached to each host. This is in sharp contrast to a SAN environment. In a SAN, storage devices can be shared among multiple applications.

It should be noted that there are many organizations, in addition to the FCIA, that are devoted to advancing storage networking. One such organization, the Storage Networking Industry Association (www.snia.com), has the mission to ensure that storage networks become efficient, complete, and trusted solutions across the IT community.

3.0 The Link Between Business and Storage Area Networks

On an ever-increasing basis, business success is closely tied to the success of the IT organization. However, when business and functional managers measure the success of the IT organization, they seldom if ever think about the individual components of the IT infrastructure; i.e., storage devices, servers, Wide Area Network links. The primary metric that business and functional managers think about when they are evaluating the success of the IT organization is the performance of the applications that they see and use on a regular basis.

Because of this increasing focus on application performance, some leading edge IT organizations are beginning to change how they communicate with the company's business and functional managers. In particular, many IT organizations used to regularly report on the various components of the IT infrastructure using a wide variety of metrics; i.e., availability, delay, packet loss. While having all of this detail on the health of the IT infrastructure is important to the IT organization, it actually made it more difficult for the company's business and functional managers to understand how their key applications

were performing. Because of this, some leading edge IT organizations are now reporting on a single metric, which is the availability of each of the key applications.

What is interesting, however, is the manner in which these IT organizations are defining application availability. For example, assume that a given inquiry/response application is supposed to respond to a query in two seconds or less. If, for whatever reason, that application does not respond to a query in that time frame, then the application is deemed to be non-available. Defining application availability this way allows the IT organization to focus on the requirement that the company's business and functional managers have for application performance.

The first generation of SANs delivered on the dual goals of enhanced application availability and lower cost primarily by facilitating:

- Server clustering
- Storage consolidation
- More efficient data backups

In the last several months there has been increased discussion in the industry relative to the importance of Business Continuity, Disaster Recovery, and data backup. However, since these terms are used somewhat interchangeably, there appears to be some confusion as to what each term means. For the sake of this document, Business Continuity will refer to keeping some number of key business processes functioning in spite of a catastrophe. Disaster Recover refers to that component of Business Continuity that focuses solely on computer systems and storage. Data backup is the component of Disaster Recovery that focuses just on storage.

While there is new emphasis in the industry on Business Continuity and Disaster Recovery, the issue of data backup has always been important to IT Organizations. The importance of data backup is a result of the fact that many of the key applications that are critical to a company's success involve data that is highly dynamic. For example, consider a company whose sales order entry system is tied into both their inventory and their distribution systems. On a minute-by-minute basis, sales orders are being entered into the system, the amount of inventory is changing, and shipping orders are being created. In this type of situation there is increased pressure on IT organizations to regularly backup the data that these applications either utilize or create.

However, whether it is part of a broad Disaster Recovery plan or not, backing up a company's data can create somewhat of a "Catch 22" situation for an IT organization. If they do not backup the data regularly, they run the risk that a lot of data that is central to running the business would be lost if there were some sort of outage. However, if they do backup the data regularly, it can swamp the production LAN and negatively impact the performance of the very applications the IT organization is working hard to protect.

The deployment of a SAN resolves this dilemma. SANs accomplish this by taking the backup traffic off of the production LAN and placing it on a high speed Fiber Channel

network devoted just to moving data. There are two sets of related technologies that are enhancing the value that SANs bring to data backup. The first set of related technologies is a collection of hardware and software that moves backup traffic directly onto the SAN, and hence conserves CPU cycles. The second set of related technologies is the burgeoning collection of Metropolitan Area Network (MAN) technologies. By bridging Fibre Channel through MANs, a SAN can be used to facilitate backup either within a campus environment, or over distances of tens of kilometers.

4.0 The Results of SAN Deployment

Many industry analysts have developed a business case analysis of the viability of migrating from a DAS environment to a SAN environment. The results of these analyses have been so consistent that there is little added value in creating yet another SAN business case analysis. Instead, this section of the document will present a high level synthesis of the results of the existing analyses.

These existing business case analyses point to the following factors as driving the large cost savings that are associated with SAN deployment.

- The amount of storage handled/manager is increased by 75%
- Disk utilization is increased by 70%
- The cost of backup hardware is reduced by 50% or more due to the centralization of this hardware
- The data center floor space requirements are reduced because there are fewer storage devices
- There is a reduced need to deploy relatively expensive general purpose servers to function as storage devices

Last year, McKinsey & Company and Merrill Lynch's Technology Research Group wrote a report entitled "The Storage Report – Customer Perspectives & Industry Evolution". In this document, that report will be referred to as The McKinsey Report.

> *According to The McKinsey Report, over a three-year time frame, the Total Cost of Ownership (TCO) of storage in a DAS environment is $0.84 per megabyte. In a SAN environment, the three-year TCO is only $0.38 per megabyte. This is a reduction of 55%.*

The largest component of the cost reduction of moving from a DAS environment to a SAN environment is a result of saving on staffing costs. However, even if staffing savings are ignored, the data in The McKinsey Report indicates that the TCO of a SAN environment is still 25% lower than the TCO of a DAS environment.

The common explanation for the extra staffing costs that are associated with a DAS environment vs. a SAN environment is that in a DAS environment the staffing requirements increase with both the number of storage subsystems, as well as the total amount of storage. In a SAN environment, the staffing requirements increase with an

increase in the number of storage subsystems, but are independent of the total amount of storage.

There appears to be widespread acceptance that a SAN environment requires significantly less staff resources than does a DAS environment. That point was made quite clearly in a report published by the Enterprise Storage Group entitled "Enterprise Storage Report Technical Review Volume 8". That report stated that companies that had implemented a SAN believed that their staff could manage 3.7 times as much storage in a SAN environment than they could in a DAS environment.

Another reason that a SAN environment is more cost effective than a DAS environment is that disk utilization is typically higher in a SAN environment than it is in a DAS environment. According to The McKinsey Report the typical SAN environment has a disk utilization of 85%, while the typical DAS environment has a disk utilization of only 50%. The low disk utilization associated with a DAS environment drives up cost. In a report published in June of 2002, Gartner stated that the average data center's unused capacity costs over $925,000 per year.

In order understand why a DAS environment tends to have a low disk utilization, consider the DAS environment depicted in Figure 1. Assume that each of the three hosts is connected to a 100 Gigabyte storage device. Further assume that one of the three disk drives is running at virtually 100% utilization while the other two are running at 25% utilization. In this scenario, the average utilization of the three disks is only 50%, and yet a disk upgrade is required!

5.0 The Next Wave of SANs

As was previously mentioned, the first generation of SANs delivered on the dual goals of enhanced application availability and lower cost primarily by facilitating:

- Server clustering
- Storage consolidation
- More efficient data backups

As the use of SANs becomes increasingly commonplace, the next wave of SAN deployment will be notably broader in scope. This wave will support entire applications suites as well as entire enterprises.

There are a number of industry wide initiatives aimed at resolving some of the issues, such as storage virtualization and interoperability, which are associated with the next wave of SAN deployment. However, far less attention is being paid to the fact that in the next wave of SAN deployment, the network itself will be more of an independent entity that is was in previous SAN deployments.

The emergence of the network as an independent entity introduces the following SAN challenges:

- Management
- Performance
- Scalability

The management challenge stems in large part from the fact that it is very common to see a number of SANs within a given data center. Typically each of these SANs was deployed as a result of separate, non-coordinated initiatives to provide storage for different applications. Since a number of SAN deployments involve proprietary technologies from a variety of storage or server suppliers, these individual SANs are often referred to as "SAN Islands".

Having multiple SAN islands comprised of different technologies greatly increases the complexity of management. As a result, one of the critical success factors relative to the successful deployment of a SAN environment is whether or not there are effective tools to proactively manage that environment. At this point in time, many SAN vendors have gotten to where they provide a remote monitoring functionality, including the ability to report on an event. However, what is needed is the ability to go past basic monitoring to using the network feedback to proactively manage SAN performance.

In order to understand the performance challenge associated with the next wave of SAN deployment, it is necessary to understand some of the Best Practices in SAN design. As shown in Figure 3, SAN's are designed around the principal of many hosts trying to gain access to relatively few storage devices.

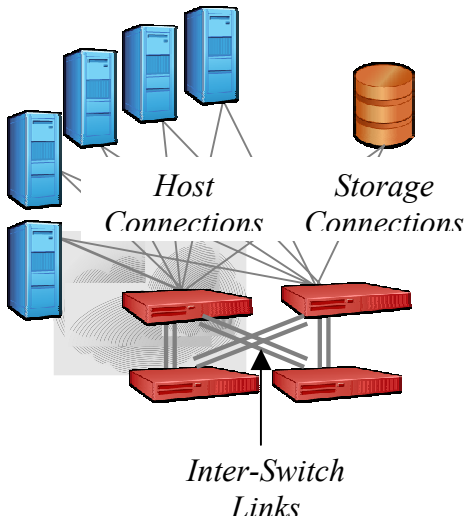

5:1 Ratio Servers-to-Storage
Figure 3

Typical SAN deployments result in five times as many hosts as storage devices that are all connected to a fabric built using Interswitch Links (ISLs) that are over subscribed by a seven to one (7:1) ratio. These practices clearly result in the storage devices being a bottleneck to overall system performance. The documented impact of this bottleneck is that performance can degrade significantly, and that this degradation can appear to be random. In order to provide predictable performance, the network must enable intelligent control over contended-for resources.

When analyzing the scalability of a SAN, three factors have to be considered. Those factors are:

I. The ability of the switch to easily support higher speed interfaces.

As was mentioned previously in this document, development work is currently underway that will enable Fibre Channel to run at 10 Gbps. IT organizations looking to deploy a SAN environment need to ensure that whatever switch they deploy can easily support Fibre Channel running at 10 Gbps.

II. The ability to support an expanding number of hosts and storage devices.

Using Figure 2 as a reference, as the number of hosts and/or the number of storage devices increases, it is important that the Fibre Channel switch can support the additional hosts and storage devices.

III. The ability to support a hierarchical SAN architecture

In the majority of cases, efficient network design has typically meant deploying a network with a hierarchical structure. One of the benefits of deploying a hierarchical SAN architecture is that it will enable an IT organization to efficiently interconnect multiple SAN Islands. The requirement to support a hierarchical SAN architecture translates into needing multiple classes of SAN switches based both on size and functionality.

To understand the need for multiple classes of SAN switches based on size, assume that only one size of SAN switch existed, and that it had 16 ports. Figure 4 depicts how such a switch would be deployed in an Edge - Core design.

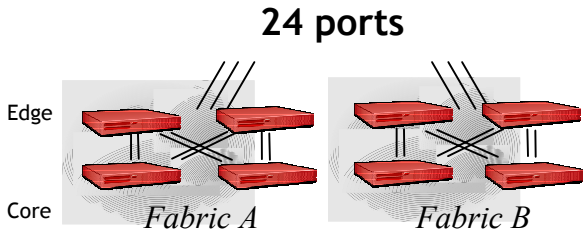

Standard Edge – Core Design
Figure 4

In the design depicted in Figure 4, each of the Edge switches has two ISLs to each of two Core switches. This means that each Edge switch dedicates four ports in order to connect to the two Core Switches. This leaves twelve ports per Edge switch, or twenty-four ports in total.

Table 2 depicts how the number of available ports grows with the number of switches. It should be noted that growing the number of switches from one to ten involves some significant management challenges. However, expanding beyond ten switches would either add an order of magnitude more complexity or require a forklift upgrade. It should also be noted that the network shown in Figure 4 conforms to the best practice of having redundant fabrics. As such, the number of switches that is required to provide the given number of available ports is actually twice the number that is shown in Table 2.

Number of Switches	Number of Available Ports
1	16
4	24
6	48
8	72
10	96

The Limit to Scaling the Edge-Core Design
Table 2

6.0 Summary

Inside the vast majority of companies, storage requirements are doubling every eighteen to twenty four months. Driven by the need to manage the cost of providing this additional storage, and at the same time improve the availability of applications, enterprise IT organizations are rethinking their approach to providing storage. In particular, enterprise IT organizations are limiting their investments in DAS, and are accelerating their investments in SANs.

Driven by factors such as increasing disk utilization, companies that deploy a SAN can expect to reduce the cost of storage by greater than fifty percent when compared with the cost of a DAS solution. However, storage technologies are rapidly evolving. In order to continue to leverage the benefits of SAN deployment, organizations must develop a storage architecture. This architecture needs to properly position storage, computing, and the storage network. The architecture needs to also articulate best practices in SAN design.