

# **Getting the Most from Your Storage**

## **How to Deploy a SAN**

**Whitepaper**

**Erik Ottem**






## 1.0 Introduction

Enterprise storage continues to grow at a phenomenal rate, with some estimates seeing a doubling in storage every 12-24 months. Although the cost of hard disk drives continues to drop, the rapid increase in the amount of data being stored creates increased management costs. The costs of providing data in terms of backup and recovery, high availability, high performance access and storage management creates a complex web of software and hardware to support clients in the enterprise.

Today's SCSI bus architecture is hard pressed to keep pace with this growth in capacity and functional requirements. In many enterprises there are multitudes of underutilized servers resulting in wasted hardware, software and human resources. To provide the advanced data services and configuration flexibility required in today's high growth environment, a new way of connecting servers and storage is required. The new model is the Storage Area Network (SAN) and it brings networking technology to the data storage environment. The SAN provides better availability, reliability and performance, with more configuration flexibility for improved access at a lower cost than today's traditional server-based captive storage approach.

## 2.0 Traditional Captive Storage Architecture

In large enterprises today there may be hundreds or even thousands of servers distributed over many sites. There may well be terabytes of data that contain vital information for that enterprise. Traditionally, this data resides on storage arrays that are controlled by the server that hosts the application using that data (see figure below). Things get more interesting if server A needs data from storage array B. To get to data today, the client request must go through server B to access the data. When the application running on server A needs data on storage array B, it accesses a local area network (LAN) which connects servers and clients.



**Figure 1 - Traditional storage configuration**

The server will authenticate the client's request for data (check the authorization of the client to read or write data) and, if approved, provide access to the data. The server is often doing other things as well, processing large number-crunching applications, managing a printer or modem network, and many other tasks. Servers are therefore often rather complicated devices requiring significant management effort and expense. The traditional captive storage architecture creates several issues for the use and management of data in the enterprise.

The availability of data depends upon the reliability of the mechanism to access that data, which involves hardware and software. In the traditional captive storage model above, if server B is broken, there is no way to access the data on storage array B since the server controls all access to the data. In many cases the value of the data exceeds that of the application or the server itself, since they can be replaced. Availability of data is one problem with traditional captive storage.

Performance is another point to evaluate with traditional captive storage. Since the server that controls the storage is doing other things, the workload of the server can affect the ability to respond to data requests from other servers. If server A needs data on storage array B, it must go through server B. If server B is busy, server A must wait. The move to enterprise resource planning will only exacerbate this situation, where more of an enterprise's operations need to coordinate efforts for maximum effectiveness. Therefore, poor performance in terms of access to data on another server can affect overall enterprise efficiencies.

Configuration flexibility in traditional captive storage is limited by the interface, Small Computer Standard Interface (SCSI). The SCSI bus has been around for years but is severely limited by distance since it can span only 25 meters at best, and 12 meters or less distance in most typical circumstances. This limits the ability to link together devices in the same building, not to mention the building next door. Distance limitations require extra cost and management to circumvent using complex SCSI bridge and expander products. Additionally, SCSI can only link 15 devices to a host, so large configurations require several buses, which in turn require several adapter cards that plug into several server expansion slots, which are often limited. Greater configuration flexibility could save management effort and cost as well as the capital costs of servers, adapters and software.

### 3.0 SAN Architecture

The SAN approach is a shared storage alternative to the traditional captive storage approach. Network concepts can now be extended from the LAN to storage by using a SAN. It allows the storage and the server to grow independently of each other. The SAN architecture uses the Fibre Channel (FC) interface to provide flexibility that allows the SAN to do storage networking. In the traditional captive storage environment SCSI is the legacy attachment, which requires a dedicated server to manage the connection between the storage and the rest of the world. In the captive storage model SCSI has one way to get the data, while a SAN's FC provides a different relationship between the server and the data, a peer-to-peer topology instead of linear as in captive storage use of SCSI. This means that multiple servers can access multiple peripherals creating a more flexible approach than with SCSI. This approach is less susceptible to server and network constraints than that found in traditional captive storage. Storage now has a network connection, a LAN type topology brought to the SAN as shown in Figure 2. Just as the LAN allows the client easy access to many servers, the SAN provides easy access for many servers to many storage devices.

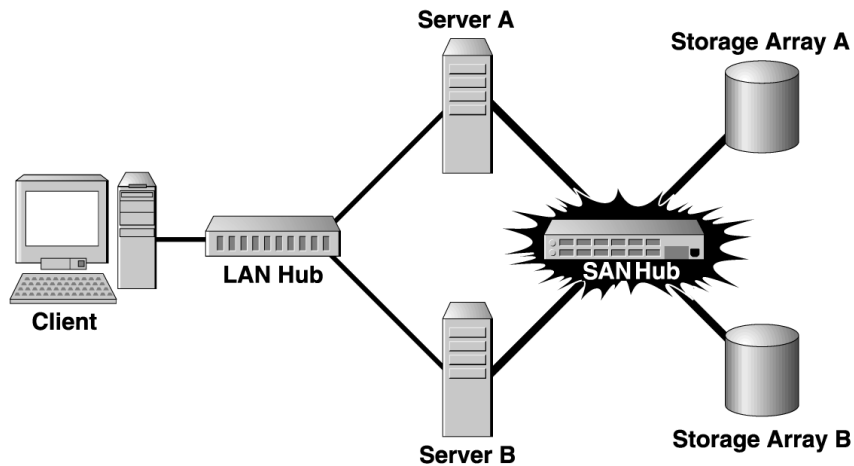


Figure 2 - SAN storage configuration

## 4.0 Getting Started on a SAN


As with any new technology, a phased implementation is the best for several reasons. Primary reasons include risk management, maximized return on investment, improved learning of a new technology, and ease of integration. By allowing the IS staff to incorporate a new technology in a staged and thoughtful way enables not only the SAN architecture but allows its optimization for each unique environment. As the IS staff gains experience, more components can be added with increased storage management to enjoy more advantages of the architecture. SAN deployment is no riskier than many other IS projects, but like all significant IS projects it should be carefully managed for best results. Indeed, using SAN for data replication may reduce risk for other critical IS projects like year 2000 application testing and migration projects.

Implementation should start with a simple configuration using a single FC hub and a few devices in a local environment replacing the SCSI bus architecture. Once this is mastered, and extended to greater numbers and types of devices, the next logical progression would be to link two FC hubs together with a storage switch for additional benefits of distance and logical isolation. The storage switch could then be tied to a backbone switch to gain a more enterprise-wide view of storage for improved storage management and cost effectiveness.

### 4.1 Start with Hubs

The best way to integrate any new technology is to start small, master the technology, and grow. Start the SAN topology in a department with either a managed or unmanaged hub connecting the existing server to the new SAN, which connects either a RAID or simple disk array. This will provide a low-cost entry with great flexibility that will be used as a building block for subsequent configurations. A reasonable first step is migrating SCSI bus devices to the new Fibre Channel way of doing things. Attaching existing SCSI arrays to a protocol converter for Fibre Channel can provide investment protection, improved management and easy implementation. Even this initial stage of SAN will provide faster access to data, improved scalability, greater distance capability and improved data availability. It can also save money. Hub implementation will greatly ease the ability to enable a LAN-free/SAN-based backup and restore capability with the appropriate application software. This results in a backup procedure that minimizes its impact on the LAN, by putting the backup traffic on the the SAN.

The hub stage should then be repeated in different departments or facilities as learning improves and the procedures have been tailored to fit the particular installation environment. The end result is a series of SAN hubs installed throughout the enterprise providing improved access, reliability and performance. It also provides the structure to reduce the number of underutilized servers saving hardware, software and management costs.



**Figure 3 -** Adding legacy storage to SAN


This phase is also known as the SCSI bus migration phase, and will provide immediate benefits while positioning for subsequent implementations. The money and time spent implementing this phase will be leveraged since the next phase can incorporate the changes made here. Additionally, if a Gadzoox® Gibraltar™ managed hub is used,

the ability to manage the SAN begins here and continues to develop in subsequent stages. With a managed hub, it is possible to construct a SAN with high availability, performance and configuration management.

#### 4.2 Link Hubs to a Storage Switch

After the initial hubs and the storage system management software have been installed, departments or other logical groups should be linked together. This phase is called the Storage Infrastructure phase and is the beginning of clustering. This phase provides the first glimpse into the broader benefits of the SAN, while still protecting the investment in the hub phase, and providing a smooth learning curve to ensure the enterprise gets the highest return on investment from each phase of the SAN implementation. Linking hubs with a storage switch provides complementary SAN services such as allowing the departmental hubs to be connected with a backup link for higher availability. In the event that one of the primary links between departments has problems, the alternate link can be utilized to provide improved availability. Additionally, a storage switch allows a long distance connection, up to 10km between two nodes, which provides greater disaster tolerance.

The performance (concurrency) and department segmentation that are provided with a storage switch will allow the departments to share valuable storage resources and develop common management of the storage resource. One of the largest benefits of a SAN is the simple sharing of storage assets such as disk arrays and tape libraries. One significant problem with traditional captive storage architecture is assuring consistent management of the data in distributed departments. The storage switch allows hubs to link together while maintaining logical separation, but with a high-performance connection. The storage switch provides a consistent protocol for easy integration when migrating from the hub phase to the switch phase. The switch also allows more disaster tolerant applications since the departmental storage can now be backed up remotely, improving the availability of important data in case of an emergency. The logical separation of paths in a switch prevents network problems in one department from spilling over to another department.



**Figure 4 -** Storage switch SAN configuration

In the above example the client has access to a valuable tape library for backup via his SAN. The Capellix™ is the Gadzoos storage switch connecting the tape library to the SAN hub and the servers attached to them. The SAN allows a tape library to be used by more clients in different departments. This provides better protection for enterprise data with less server overhead and cost, setting the stage for common management of storage assets. This can be done at distances up to 10km to allow groups at some distance to access these resources. It also allows data replication for better disaster tolerance by mirroring data from one site to another via the storage switch without the distance limitations of SCSI or the overhead of traditional data networks. As the amount of critical data grows and the need for continuous access increases, data replication on the campus with the above configuration provides a simple way to protect your operations. This becomes even more feasible as the cost of storage continues to decline. (Data replication, mirroring and other storage management tasks require the implementation of the appropriate software to ensure data integrity is preserved.) Storage sharing or pooling becomes possible here for more efficient use of any storage asset allowing lower overall storage costs. If the Gadzoos Capellix chassis switch is used, its ability to scale can grow with the enterprise, also providing extensibility to leverage your switch investment, in addition to the easy integration of any existing hub architecture.

### 4.3 Backbone Switch Phase

Once the hubs and storage switches are installed together with the storage management software, it is time for a larger switch. This environment will allow the enterprise to extend the benefits of the SAN, and it will have been implemented in an incremental way with high reliability, legacy equipment utilization and maximum return on investment of time and capital. It is not necessary to risk the integrity of your data to incorporate a new technology; but it is necessary to manage its introduction. This final stage is referred to as the universal data center since data is now available throughout the enterprise with a variety of applications, locations and storage devices. During the phased-in introduction, experience has already been gained in the hub and storage switch topologies. Experience gained with storage management software has allowed consistent storage management in a diverse environment. The incremental implementation being recommended allows the realization of improved availability, reliability, performance, and configuration flexibility at an attractive cost. The SAN allows the enterprise to maximize the utilization of its LAN, servers and storage by differentiating the network to provide for its different functions. It allows easier data sharing among different systems, and backup that will not impact LAN performance of the clients gaining access to their data. In this final stage we find consistent data management, access security and shared access managed by the SAN in a cost-effective manner that can scale to the largest data centers.

The switch environment shown below allows many departments to be linked together with high-capacity dedicated bandwidth for the traffic associated with a backbone storage network. For example, large clustered environments are enabled with the Capellix switch and are linked through a series of backbone switches.

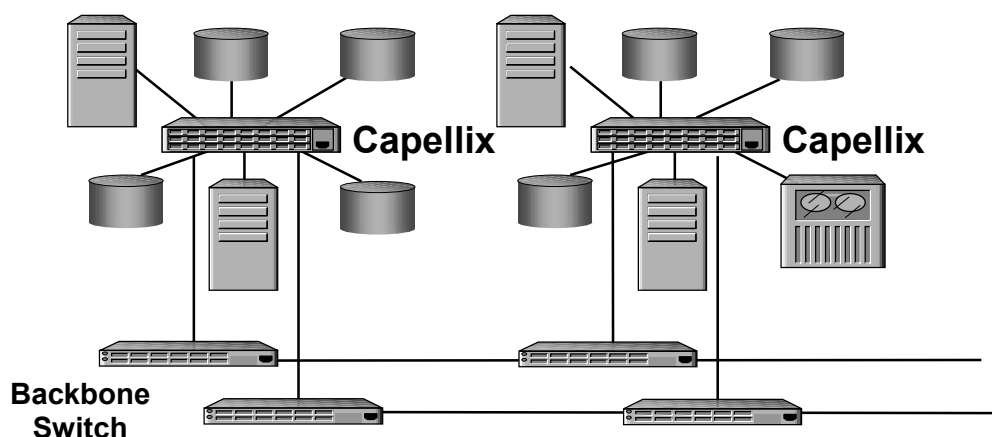


Figure 5 - Capellix enabling cluster configurations

## 5.0 Summary

SAN implementation can be expected to progress through three phases for the highest range of benefits. Hubs like the Gadzoom Gibraltar managed hub for the SCSI bus migration phase, storage switches like the Capellix chassis switch for building a storage infrastructure, and backbone switches for the universal data center make the SAN story real. Chassis-based switches such as Capellix with interchangeable plug-in modules allow simple storage switch configurations now, and a smooth, lower risk transition to future plug-ins for backbone connectivity and services as data centers are consolidated. By building the SAN in well-defined and thoughtfully considered plan, each element is more likely to build a low-risk, high-return implementation. SAN architecture provides a superior solution to the concerns of traditional captive storage; reliability, availability, performance, configuration flexibility and cost. In each area of concern SAN provides a better way of doing things than traditional captive storage. The SAN architecture is based on Fibre Channel development that started in 1988 and a vast amount of experience has been gained since that time. It is time for your enterprise to enjoy the benefits of SAN technology. This low-risk incremental approach preserves the investment in the equipment from the previous phase, and builds on the experience and management techniques used in each successive step. The phased approach outlined in this paper will allow a greater opportunity for a successful adoption of the SAN technology so your enterprise can reap the rewards of better reliability, performance, and configuration flexibility at a more reasonable cost.