Beginner's Guide to Storage Area Networks

Overview

With the announcement of the new Dell Storage SCv2000 Series arrays, customers now have a highly capable, enterprise class, entry-level storage solution that brings many of the benefits of Dell's more advanced SC4000 and SC8000 class storage area network (SAN) based products. For many customers that have traditionally used storage solutions tied directly to physical servers — often referred to as "direct-attached storage" or DAS — the SCv2000 presents a great opportunity to consider new storage management and deployment paradigms that introduce beneficial new functionality and features that promote data scalability, security, and protection.

To help new potential SAN buyers understand the differences between a DAS and SAN solution, this paper explores the differences between the DAS and SAN solutions, the benefits of moving from a DAS to a SAN architecture, and provide a high-level discussion of building a first SAN solution.

Target audience

This paper is designed to help small and medium business systems administrators who are considering becoming, or are first-time buyers of a SAN-based solution who have traditionally used DAS solutions.

What is a storage area network?

For those administrators that have traditionally used DAS solutions, they may not be familiar with what a SAN really is and how it compares to their previous experience with DAS solutions. The following sections discuss how SAN is similar or different than DAS solutions, then we look at SAN from various aspects including the SAN protocols most commonly used.

DAS vs. SAN

The two predominant storage solution technologies are the DAS and SAN solution. How these two technologies are similar and different can be used to determine which is better suited for the application being deployed.

Many customers like the simplicity of DAS solutions for a variety of reasons. DAS implementations typically rely on basic, OS native or vendor provided drivers resulting in DAS being readily available with little additional configuration outside of the OS provided tools. Whether individual disks attached to an internal storage connector or larger solutions consisting of a group of disks tied together and presenting a single pool of storage capacity — often referred to as "just a bunch of disks" (JBOD) or "redundant arrays of inexpensive disks" (RAID) — attached to a PCIe bus adapter, the simplicity of DAS solutions can be a very valuable selling point to customers.

A SAN is typically a specialized network that provides one or more high-performance, scalable paths between one or more servers and one or more storage resources. These storage resources can be remote from the servers that are using the resources, but typically are located in the same data center. This differs from a DAS solution, in which a single server is directly attached to a single (or a small number) storage resource(s) that is only visible to that server, and each server has its own dedicated disk resources. Both SAN and DAS provide data in the same way to the server, but simply differ in the transmission medium used. Both solution types do not provide higher-level services such as file abstraction, which is provided by the server's operating system.

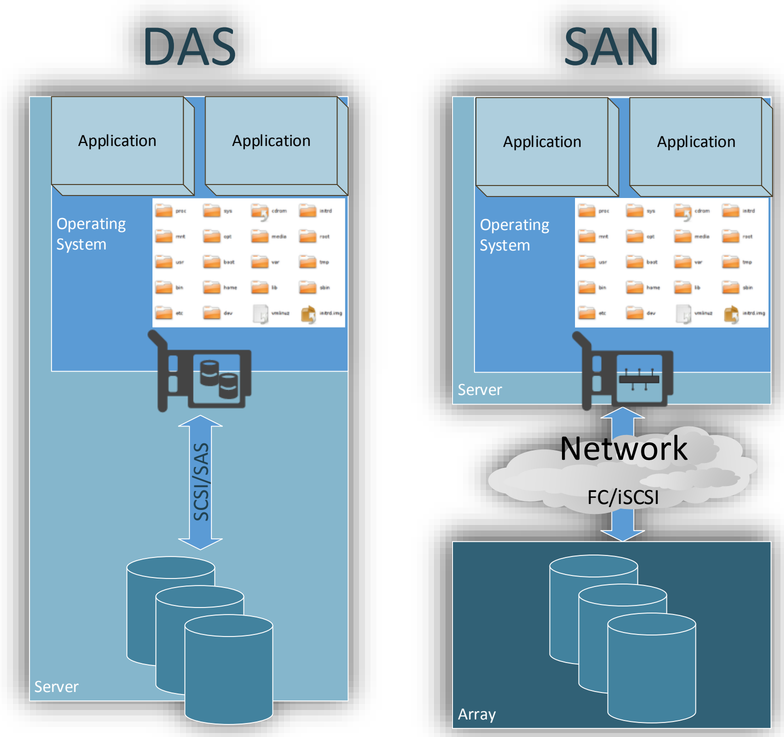

Figure 1 - DAS versus SAN

Figure 1 illustrates the high-level differences between SAN and DAS solutions. This figure shows how both solutions use special devices called storage controllers that interface between the server and the storage solution, and both use a communications protocol to move data between the application and the physical disks in the storage solution. In the DAS solution, the protocol is typically the small computer system interface (SCSI), and the SAN uses either Fibre Channel (FC) or Internet SCSI (iSCSI) protocols.

Advantages of SAN storage solutions over DAS

SAN has several advantages over a typical DAS solution. These advantages can be categorized into the following categories:

- Scalability – the ability to grow the solution as needed
- Management – centralized, secure control of the solution
- Efficiency – maximizing the use of storage resources for the least cost
- Availability – maximizing uptime of the storage solution, even in the event of a failure of a component or data center

Scalability

Scalability can take several forms: capacity growth, performance scalability and server growth. The ability to start small and grow the storage solution really allows smaller customers to take advantage of all of the other benefits of deploying a SAN at a lower entry price, then adding compute, capacity, networking, or higher performing media — in the form of higher speed spinning media or by adding solid-state drives (SSD) — as needed to support the application services that are using the SAN.

In terms of capacity scalability, a SAN solution like the Dell Storage SC Series family, can scale from only a few terabytes (TB) to petabytes (PB) of storage within the same solution. The entry-level SCv2000 array offers capacities ranging from a few terabytes up to hundreds of terabytes.

Scalability is also important in terms of overall performance. While storage solutions can increase performance by adding disks — especially higher performing disks like 15K spinning disks or SSD — the SAN solution can scale in terms of throughput. You can increase throughput by adding additional network connections between the array and the storage networking devices or by adding additional storage network ports to one or more hosts. You can expand throughput by moving from a lower bandwidth networking infrastructure to a higher bandwidth network infrastructure. Examples are updating arrays that were initially deployed using 1Gb Ethernet with 10Gb Ethernet interfaces, or upgrading from 8Gb Fibre Channel to 16Gb Fibre Channel.

The SAN also enables giving additional servers access to the storage resources simply by attaching new servers to the storage network and creating and assigning storage resources to those servers using the provided storage management tools without any effect on the other servers' storage access.

Management

Management of the disk resources deployed using the DAS storage architecture requires storage be managed (allocated or de-allocated, backups performed, and so on) on a server-by-server basis due to the fact that each DAS implementation, by definition, is physically connected to each individual server. If the administrator wants to determine the free space on DAS storage on multiple servers, the administrator must log into each server then run the correct report or script on each server individually.

Contrast this to a SAN where the centralized nature of the storage deployment allows for centralized, and in many cases, automated management. You can assign storage to new volumes then assign those volumes to any number of servers in one place. This centralized management paradigm means that any conceivable storage management function can be accomplished using a limited set of tools, and can be accomplished by using these tools from locations different than the physical storage solution. Changing RAID types, adjusting performance parameters, backing up data, and so on can all be done with a few clicks of a mouse or by running a simple script.

Efficiency

Efficiency can take many forms. This paper discussed management efficiency and budget efficiency in the previous section, but what about efficient usage of the storage already in service? In a DAS solution, efficient usage of deployed capacity is not possible because each server has its own, directly attached storage. That storage is only available for use by the specific server to which the storage is attached. In this usage model, if there are 10 servers, and the free space on each server averages around 100GB, then there is 1TB of storage space that is not being used. What happens if an application on one of those servers then runs out of free space? Additional storage must be purchased, attached to that server, and made available to the application. The deployed capacity attached to the other servers is not available for use.

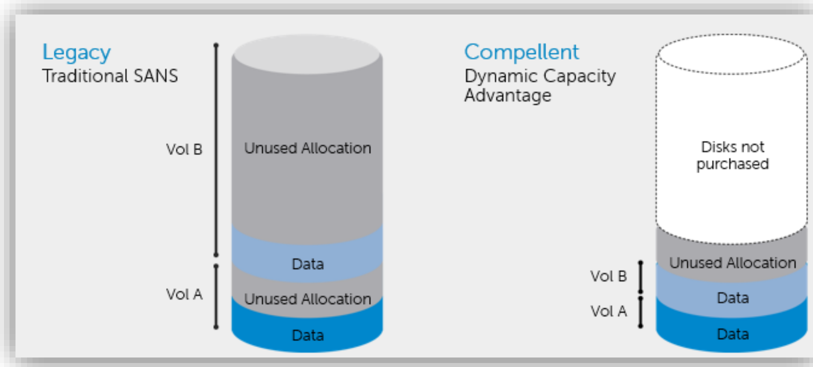

Figure 2 - Thin Provisioning

In a SAN solution, all of the storage is centrally located, centrally managed, and can be allocated — in most cases — on the fly by the storage administrator. In this environment, that 1TB of free storage space can be held in a reserve "free pool" that can be allocated as needed.

In fact, most modern SAN storage solutions, like the Dell Storage SC Series arrays, can actually configure a special type of volume (a disk drive visible to the server) called a "thin volume". This type of virtual volume can be configured to report to its assigned server that its capacity is much larger than is actually allotted to the volume by the administrator. Then when the volume begins to run out of physically assigned free space, the array operating system can automatically allocate more physical capacity to that volume. Figure 3 illustrates a comparison to traditional volume allocation versus a thin provisioned volume allocation. Ultimately, thin provisioning allows administrators to better optimize the usage of their total physical storage capacity over many volumes and many servers.

Availability

SANs provide a much more fault tolerant storage solution than the typical DAS solution. This is primarily because networks, in general, provide mechanisms to allow two devices (such as a host server and its SAN attached storage resources) to communicate via multiple data transmission paths. These paths can, depending on the implementation, be configured to be failover paths, or configured as parallel, simultaneous data paths. This latter implementation, not only provides fault tolerance — like the former — but also provides enhanced performance by increasing the available bandwidth available to the host.

Also, due to the fact that SAN storage resources are centralized allows for better control over data protection by supporting more advanced forms of RAID (disk failure protection) as well as supporting the ability to back data with the SAN solution to a remote location — typically hosting another SAN solution — through data replication technologies. Together with parallel host-to-storage, path management truly can provide a fully redundant storage solution with no single points of failure.

Storage area networking protocols

As previously shown, the two most visible values of SAN is the ability to share the storage resources with many servers at one time providing a new level of usage efficiency, along with the ability to centralize the storage — sometimes at a distance from one or more of the servers. But how does the SAN actually do this? By using well understood networking standards, a SAN can do both of these feats as well as providing the infrastructure to support the other advantages discussed.

Regardless of whether the storage solution is a DAS or SAN, the primary storage transport "language" being used is SCSI. Almost all enterprise-class storage devices in use today utilize SCSI as its low-level control and transport. In a DAS solution, this SCSI protocol is the only protocol used to communicate between the disks and the server they are attached to. The actual implementation may be parallel (SCSI) or serial (Serial Attached SCSI or SAS).

A modern array architecture like the SC Series and the SCv2000 actually use multiple protocols and "networks". The SCv2000 typically uses two types of networks: the "front-end" network (network used to connect hosts to storage resources), and a "back-end" network (typically used for the array to talk to its included disk drivers).

For completeness, it is important to understand that the SCv2000 uses front-end networks that leverage either Fibre Channel, iSCSI, or SAS. While this paper is focused on the more scalable FC and iSCSI front-end networks, the array also has a SAS front-end option to support DAS solutions. Also, the SCv2000 uses SAS on its back-end network to enable communications from the storage controllers and the disk drives in the SCv2000 base unit and any expansion disk enclosures.

A SAN, regardless of which SAN protocol is being used, also uses SCSI as the underlying, low-level communications protocol. The reason for using SCSI is primarily compatibility. Also, by using SCSI, the servers have no idea that the storage resources are not attached directly to the server, that the resources are shared, or that the resources may be some distance away from the server itself.

In a SAN, these SCSI data packets are embedded — or encapsulated — in a second protocol that can be transported, routed, and controlled by various networking devices such as switches and routers. This second protocol is typically either iSCSI or FC. For this reason, most people don't really even refer to the SCSI protocol contained in the FC or iSCSI protocol, but simply refer to FC and iSCSI as *the* protocol.

iSCSI

ISCSI is probably one of the most versatile and cost effective of the SAN protocols. This is because iSCSI uses industry standard Ethernet switches and adapters leveraging the same series of standards used for traditional client-server/Internet networks. For this reason, iSCSI is a great protocol to use in small to medium businesses implementing their first SAN solutions. A business can deploy a SAN for a lower initial cost leveraging almost any on-hand networking devices. For a more technical understanding of iSCSI, refer to the RFC7143 specification on the ietf.org website.

Fibre Channel

When IT professionals discuss the topic of SAN, FC is generally assumed as the protocol used to build an enterprise class SAN. While in the past, this might have been true due to early implementations of iSCSI relying on networking technologies that were typically much slower than FC networking technologies. As stated earlier, the raw speed of the networking technologies has become all but irrelevant. At the time of the release of this white paper, FC networks relying on 16Gb technology are becoming common with 32Gb technologies expected in another year or so. Even so, the most common implementations today are still based around 8Gb technology. For a more thorough understanding of Fibre Channel, refer to http://fibrechannel.org/.

Which protocol should be used?

As a storage vendor, Dell offers SAN storage solutions that use iSCSI and/or FC protocols for each array model. The performance between the two is pretty much the same in terms of throughput, latency, and so on. Because each array model offers an identical set of functionality regardless of protocol used, there is no special feature that is dependent on one protocol or the other. The decision point may simply come down to what networking solutions are already deployed. Almost everyone has an Ethernet network in place, and experience managing that network may be a tipping point. If the customer already has an existing FC based storage network then, unless there is a reason to move away from that technology, it might be best to stick with FC. If the customer is new to SAN and only has experience with DAS, it might be easier to leverage existing skills and purchasing power to deploy an Ethernet based iSCSI solution.

Dell believes you should have the freedom to choose the network fabric that’s right for you — not only within one SAN, but across your enterprise. This includes the flexibility to change technologies as business needs change or to implement new technologies as they become available. In other words, choose the technology that’s right based on business needs — not on what a specific storage vendor recommends — keeping in mind how those needs may change over time.

Components of a storage area network

Every SAN consists of one or more hosts, one or more storage arrays/devices, each of which is connected together using some form of networking technology as Figure 4 illustrates. The SAN network may consist of a single or dual fabric implementation, which is discussed later in this paper. The array component is straightforward, but the following sections covers a little more detail around the host/server and network components typically used in a SAN.

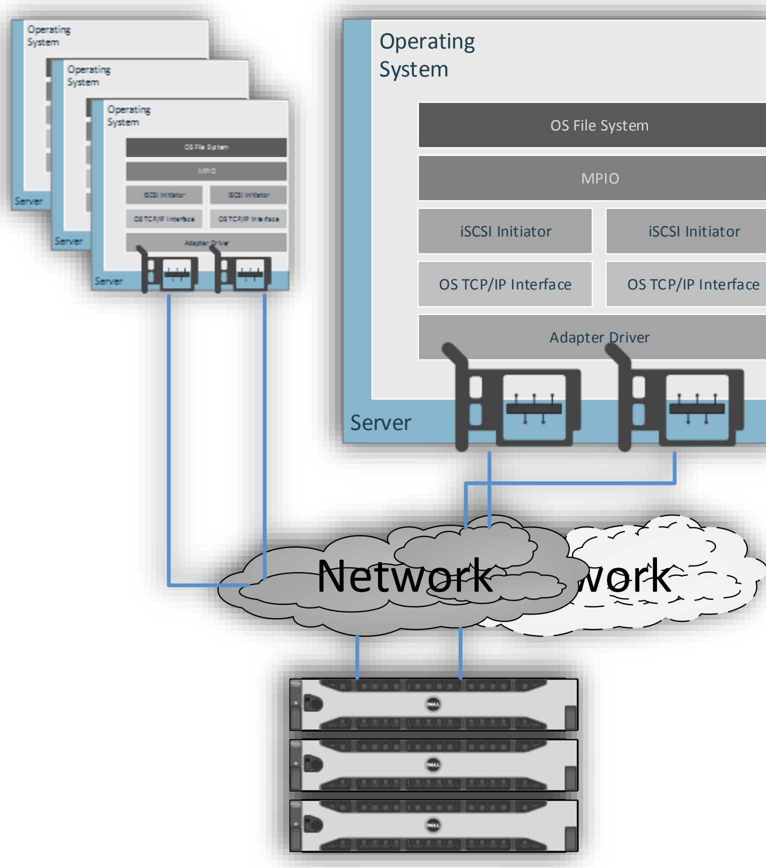

Figure 3 - Typical Storage Area Network

Host

Fibre Channel host networking devices are called storage host bus adapters (HBA) and almost exclusively provide hardware-based, self-contained Fibre Channel protocol implementations. This means that an HBA does all of the FC specific work on behalf of the server requiring no FC protocol awareness by the host operating system. The host operating systems sees these HBAs as "storage devices" and not as networking devices. For example, in Windows, an FC HBA is listed in Device Manager under the Storage Controllers sub-tree.

iSCSI uses industry standard Ethernet and therefore can use almost any standard network interface card (NIC) that supports TCP/IP over Ethernet. NICs can come in a variety of implementations that range from providing no iSCSI awareness to full HBA functionality — just like the FC HBAs just discussed. In the former case, a "basic NIC" will have no awareness of iSCSI and must rely on the operating system to provide all of the iSCSI protocol management. The OS will see these devices as "network adapter". This is called "software-based iSCSI" and is the most common deployments in the small to medium customer environments. In the latter case, just like the FC HBA, the iSCSI HBA is a self-contained (or semi-self-contained) iSCSI protocol device and the server operating system requires no (or little) iSCSI protocol awareness. This implementation is called "hardware-based iSCSI" or "iSCSI hardware offload". As with FC HBAs, iSCSI HBAs will also be seen by most operating systems as "storage controllers". Those devices that are semi-self-contained are typically seen as both network adapter and as storage controllers simultaneously.

There are advantages to each implementation. The software-based solution is lower cost, provides excellent performance, and can be optimized by the operating system to meet its needs. The hardware-based solution, while more expensive, can be implemented across a broader range of operating systems and does not use critical CPU resources to manage the iSCSI data flow. Many first-time iSCSI SAN implementers use software-based iSCSI since it is easier to find networking devices and most modern operating systems — including Microsoft® Windows®, VMware®, and almost all Linux® distributions provide the required iSCSI protocol software.

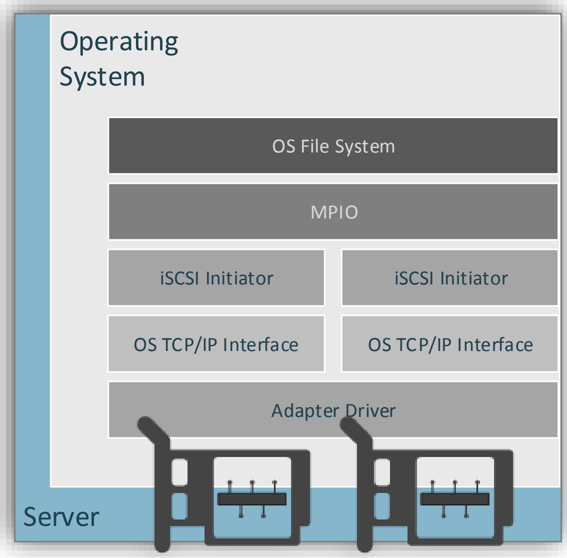

Figure 4 – iSCSI SAN Host Components

As shown in Figure 5, a typical host consists of several sub-components that enable the operating system, and the applications requiring storage, to access remote SAN storage resources. This example illustrates the software-based iSCSI as the protocol, but FC SAN hosts will have equivalent functionality at some level either in software or hardware. Each host should be configured with two or more network interface adapters that support the desired storage protocol. This ensures that access to the storage resources in the SAN will survive hardware failures that might prevent that access by providing multiple paths for the server to access the storage resources (we will discuss multiple paths in a moment).

MPIO

The MPIO block in Figure 5 stands for "multi-path input/output". In terms of fault tolerance, high availability and performance maximization, the MPIO component is extremely important. It is this component that is responsible for identifying and creating different paths between the host and the storage device (a "disk" or "volume") on the remote array. These paths are built on top of the multiple hardware adapters in the server. In the illustration, the host has two physical paths between the operating system and the array — each path using a different adapter. By creating a multipath environment, the MPIO component can optimize the movement of data between the OS file system and the volume(s) on the array. For those who are familiar with traditional Ethernet networking, this MPIO architecture is different than the "NIC teaming" concept used for the client-server data network (commonly referred to as the "LAN" or Local Area Network). Basically, LANs leverage NIC teaming configurations, and SANs use MPIO configurations.

Network

In general, the network is designed in one of two ways: single-fabric or dual-fabric. Either implementation can support a redundant, fault-tolerant multipath solution. Traditionally, dual-path solutions have been the norm for most storage array implementations, but most can also operate within a single-fabric configuration. Single-fabric tends to be more prevalent in Ethernet-based iSCSI networks, while dual-fabric implementation tends to be prevalent in FC or FC/iSCSI multi-protocol networks.

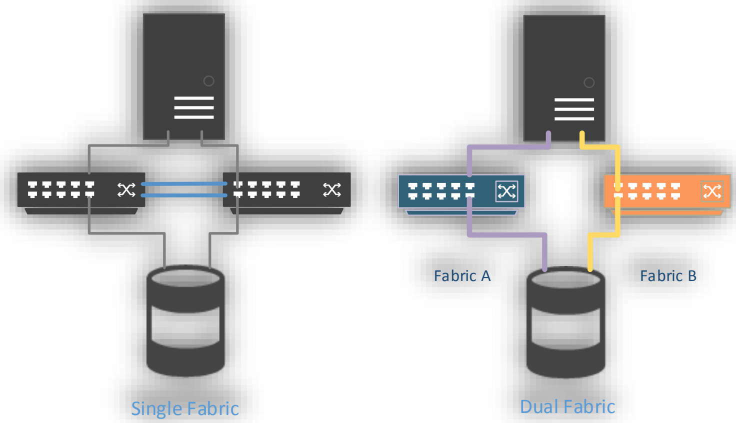

Figure 5 - Single vs. Dual Network Fabrics

iSCSI

iSCSI uses Ethernet, a general purpose networking technology that is the dominant networking technology currently available. As a general purpose network, these networks carry almost every type of data traffic imaginable, from web traffic, video, audio, data, and storage. Ethernet can be used to transmit almost any kind of storage traffic including FC traffic. Today, Ethernet is used primarily for iSCSI traffic, but also can, with specialized components, transport FC traffic using a new protocol called Fibre Channel over Ethernet (FCoE), though this is not typically a consideration for a first SAN solution due to additional costs and complexity when compared to iSCSI.

Fibre Channel

Fibre Channel networks are specialized, high-speed networks specifically designed to interconnect hosts/servers and storage arrays. There are several different types of FC network "topologies" such as point-to-point and arbitrated loop that were initially used in the early years of FC deployments, but the most common topology today is the switched fabric, and is the topology supported for SC Series storage arrays. Figure 6 illustrates a typical FC SAN architecture consisting of two fabrics and each fabric contains two switches.

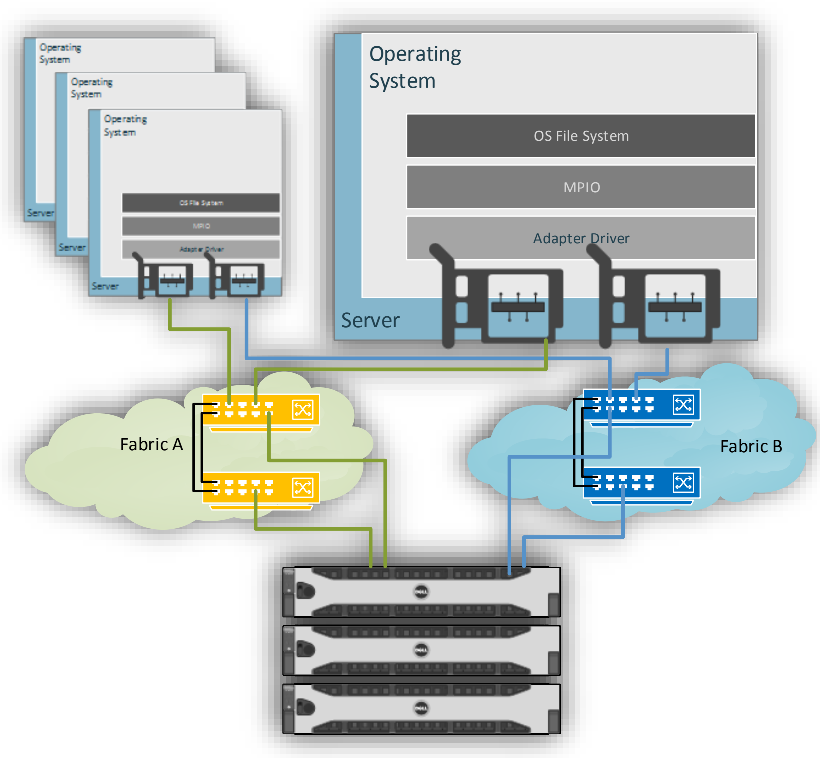

Figure 6 - Fibre Channel SAN

Array

Dell's Storage Center family of storage products — the SC Series — provide a broad set of features and are designed from the ground up for enterprise workloads common in today's virtualized data center. Providing a range of performance and capacity options, the SC Series is flexible, modular, and scalable. The SCv2000 is a new, entry-level offering in the SC Series that specifically targets the first-time SAN purchaser by offering several key features from its bigger cousins, the SC4000 Series and SC8000 Series.

The following two figures provide an overview of the SCv2000 controller options. Each controller provides two front-end ports for connection to either an iSCSI or FC SAN network and two back-end serially attached SCSI ports to support additional disks by allowing the addition of up to two additional disk enclosures. The SCv2000 comes with ports providing native support for connecting to out-of-band management and iSCSI-based replication to another SC array.

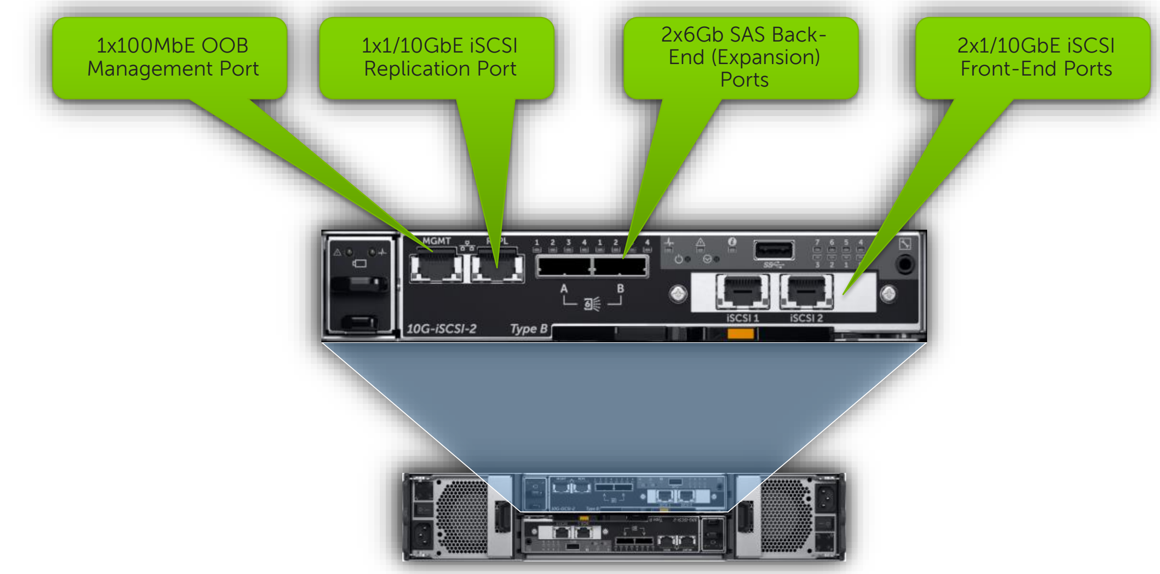

Figure 7 - SCv2000 iSCSI Controller

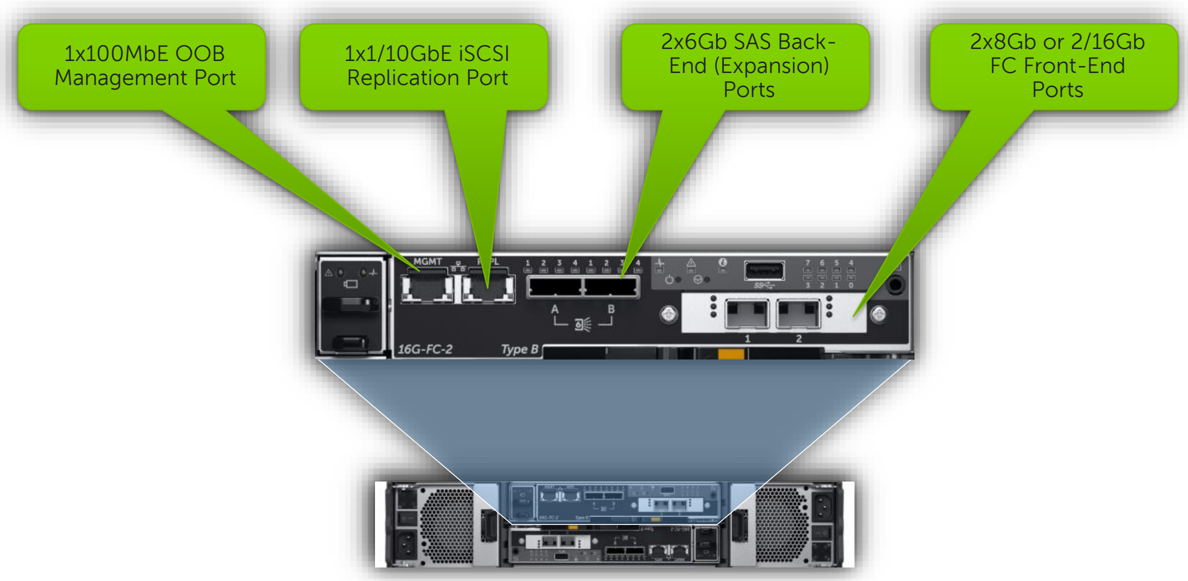

Figure 8 - SCv2000 Fibre Channel Controller

Building your first storage area network with the SCv2000

The following sections provide a set of recommendations for building a SAN. Assumptions made in this section are that the reader does not currently have a deployed SAN, but may have experience with DAS storage solutions. The following sections cover both FC and iSCSI based SAN solutions.

Fault tolerance

An important aspect of any SAN deployment to consider is that of fault tolerance. To ensure constant connection to important data stored in the storage solution — even if a NIC/HBA or a SAN switch fails — at least two different paths must exist between any host and the desired target volume. Most storage vendor's array offerings typically use two physical network fabrics as described earlier in this paper, but many iSCSI based arrays can also operate using a single network fabric that might typically be found in an existing data center. The bottom line is that every host needs to have two network adapters, the network must have two switches, and an array must have two storage controllers installed to ensure fault tolerance. The remainder of this discussion focuses on building a completely fault-tolerant SAN.

Choose your protocol

To some extent, choosing a protocol is a decision based on several factors as described earlier in "Which Protocol Should be Used?" and the decision of either protocol will enable the creation of a SC Series SAN with the same basic set of features and services.

Creating the network

Once the protocol decision has been made, it is time to build the network. As discussed earlier, there are two basic strategies in building a SAN: single-fabric and dual-fabric. Because it is the most common deployment method and works for both FC and iSCSI SCv2000 SANs, this discussion is be focused on building a dual-fabric SAN. Use a single-fabric design if the network is Ethernet, and the SAN will be leveraging existing LAN switching infrastructure (which by design are single-fabric).

The main tasks here include:

- Selecting and buying appropriate switches
- Configuring each switch according to best practices recommendations

Selecting and buying SAN switches

The good thing to remember here is that all switch vendors for the most part follow a very clear set of industry standards when they design and build switches. Each protocol, FC and Ethernet/iSCSI, are defined and controlled by standards maintained by various standards.

Standards association	Area of responsibility
Institute of Electrical and Electronics Engineers (IEEE)	Ethernet (IEEE 802.1, 802.2, and 802.3)
Internet Engineering Task Force (IETF)	TCP/IP (RFC 793, 794, 1349, 2460, 2474, 6864) iSCSI (RFC 3720, 3721, 5048)
International Committee for Information Technology Standards (INCITS)	SCSI (T10) Fibre Channel (T11)

Table 1 - SAN Industry Standards

The fact that switch and network adapter vendors build products that meet these standards requirements is very valuable. It allows customer to build SANs from almost any set of components that claim to follow these FC and Ethernet/iSCSI standards.

When building a SAN, the easiest way to ensure success is to look at switches that the array vendor has already tested. While a vendor probably cannot test every switch offering in the market, they can focus on testing switches they think will complement the features and performance requirements of the array. Dell has a long history of testing various networking devices (switches and network adapters) with their storage array offerings. Dell takes this information and publishes it in the Dell Storage Compatibility Matrix (DSCM). While this is not an exhaustive list of switches and adapters that are supported with the SC Series of array products, it is a great place to start. Even if the desired switch is not listed in the DSCM, the fact that these devices use the same standards and technologies means that there is very little risk that the selected device will not work.

Configure the switches

Building a dual-fabric SAN is a straight-forward process. For iSCSI based SANs, Dell has built a library of Switch Configuration Guides for various Dell Networking switches. These can be found on Dell's TechCenter community. Regardless of the switch vendor, these guides are helpful in guiding the reader in building an iSCSI based SAN for the SCv2000.

For Fibre Channel and iSCSI SANs, it is also recommended that switches be configured using the guidelines in the SCv2000 Deployment Guide. All recommendations made in this document adhere to these standard recommended configuration guidelines.

Attaching the array

The SCv2000 comes standard with two controllers and each controller supplies two primary (front-end) network ports (either FC or iSCSI) to connect to the SAN switch infrastructure. These controllers and ports must be connected to two different switches with a port from each controller connected to a switch in each SAN fabric. In normal operating mode, all ports are active and process IO from hosts.

If a storage controller fails, then host connections using that controller will migrate to the other controller. If a single port fails on a controller, then the hosts using that port will migrate to another port on the same controller.

Table 2 lists the connections that connect the array to the two switch fabrics.

Controller/Port	Fabric A	Fabric B
Controller 0/Port 0	X	
Controller 0/Port 1		X
Controller 1/Port 0		X
Controller 1/Port 1	X	

Table 2 - Array connections to SAN fabrics

Figure 9 illustrates a typical, redundant connection scenario.

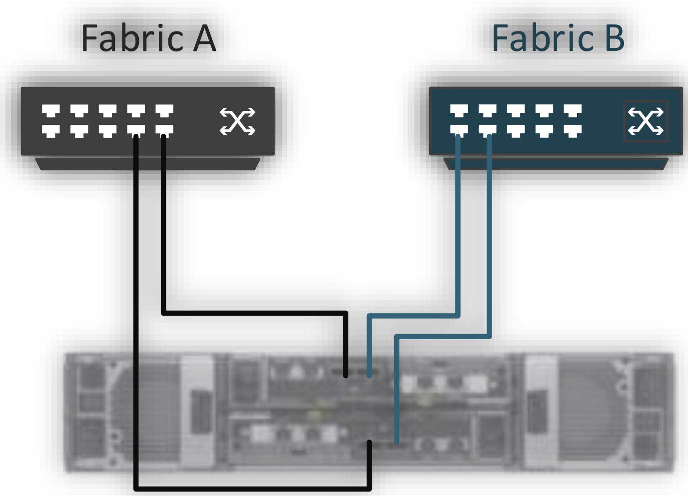

Figure 9 - Attaching array to switches

Configuring your host

Once the array(s) are connected to the SAN fabrics, it is time to connect each host server to the same two fabrics. Regardless of the protocol used, the concepts are the same — install two adapters, either NICs (iSCSI) or FC HBAs (FC), into each server, then connect at least one port from each adapter to one SAN fabric. To ensure a fully redundant, highly available SAN design, the administrator should install two separate adapter cards rather than simply installing a single dual-port adapter and using both ports from that single adapter. While not required, by using ports from two different adapters, the solution will continue to be fully redundant.

The basic steps for configuring the hosts are:

1) Install at least two SAN network adapters in each host
2) Install drivers for adapter
3) Configure ports
    a. If using iSCSI/Ethernet switches, at least one port should be configured with an IP address associated with the SAN Fabric A subnet, and at least one port should be configured with an IP address associated with SAN B subnet
    b. If using FC switches, add the World Wide Port Name (WWN) for at least one port from an adapter into the SAN A zone and the WWN for at least one port from a second adapter into SAN B zone

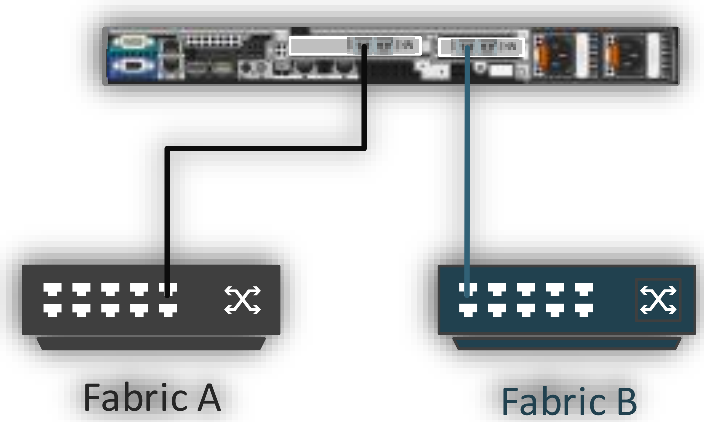

Figure 10 - Attaching hosts to switches

Table 3 lists the most commonly used Ethernet and FC adapters for SCv2000 SANs.

Manufacturer	Adapter Family	Protocol	Speed	Notes
Emulex	OneConnect OC*x*14100 Family	iSCSI	1/10Gb	SW & HW iSCSI capable
	Lightpulse LP*x*12000 Family	FC	8Gb	
	Lightpulse LP*x*16000 Family	FC	16Gb	
QLogic	8200 Family	iSCSI	1/10Gb	SW & HW iSCSI capable
	BCM57800 Family	iSCSI	1/10Gb	SW & HW iSCSI capable
	2500 Family	FC	8Gb	
	2600 Family	FC	16Gb	
Intel	X520/X540 (Intel 82599) Family	Ethernet	1/10Gb	SW iSCSI only
	i350 Family	Ethernet	1Gb	SW iSCSI only

Table 3 - Validated host adapters for SCv2000

Creating and assigning volumes

Once the SAN has been configured, the real value of the SAN becomes apparent. During the volume creation process, the administrator has the opportunity to add valuable controls and features to each volume such as:

- Data protection by creating snapshots
- Disaster recovery via site-to-site replication between SCv2000 arrays
- Efficient storage utilization by using thin provisioning
- Improved data security by assigning volumes to specific servers
- Configuring data optimization to move data to different performing RAID sets

These are just a few of the things that can be done during volume creation and assignment.

Creating and assigning a volume to a server takes just a few steps:

1) Start the "Create Volume" process in System Manager
2) Select whether volume will be put on RAID (redundant) storage or non-redundant
3) Specify Size of volume
4) Decide if Replay (snapshot protection) will be used
5) Map the volume to one or more servers

Once these steps are completed, the server will then be able to "discover" all volumes assigned to it in System Manager.

Conclusion

SANs are really not much different from storage directly attached to each server, and should definitely be considered as the next step in scaling the enterprise. They provide some key advantages in efficiency, scalability and management over direct attached storage that may be valuable as the business grows. While the idea of a SAN may seem intimidating due to the addition of a specialized network separating the servers and storage resources, in reality, using standards based networking components takes much of the risk out of the solution. By adding a few extra steps in initially deploying a SAN, the administrator gains a lot of advantages and maintaining a SAN imposes no additional workload.

The release of the new Dell Storage SCv2000 Series offers a great opportunity for those needing to scale their legacy server and storage infrastructure model from "local" storage and create a shared, high-performance storage network. Its ease of use, affordable price, and advanced features allows administrators to move their enterprise to the next level with confidence.