White Paper

# Architecting SANs for Successful Deployments

**Brent Knight, System Engineer**
QLogic Corporation

# Architecting SANs for Successful Deployments

Brent Knight, System Engineer
QLogic Corporation

## Introduction

Investing in architecting a Storage Area Network (SAN) today will not only solve your current issues of storage consolidation, improvements in backup and recovery or high performance data access, it will facilitate day-to-day support and maintenance in the data center. A SAN investment should save support dollars, not require additional support resources. Investing in a discovery process of both your current needs and future expansion requirements will save lots of headaches and could save costly upgrades in the future. This paper examines the strategies in building a SAN architecture with considerations on fabric allocation, software needs and troubleshooting.

## Architecting the SAN

The biggest impact to SAN performance will be realized through a solid SAN architecture and design. The infrastructure should be chosen carefully to ensure that the customer’s requirements are addressed in the most efficient and resilient way possible.

## User Port Counts

First, determine the number of user-attached ports required. Then add 20% to the calculation for near-term expansion. Experience tells us that we always underestimate the number of users that need to be attached to the SAN. Even if these ports aren’t used at the outset of the installation, these extra ports will prove to be invaluable when it comes time to troubleshoot connectivity problems should they occur. The extra ports can also be utilized as ISLs (Interswitch Links) when expanding a fabric with multiple chassis. This allows production traffic to remain stable while the SAN can be scaled for future needs.

## Take Stock of the Initiators

Secondly, take an inventory of the types of devices that will be attached to the SAN. Understand their characteristics and the physical location of this equipment. Initiator's first – explore and document your servers. Determine the type of server, the I/O bus type, performance expectations; it's application usage and the degree of redundancy required. The answers to these questions will begin to supply us with the port requirements for the SAN. Next, we must determine the level of server redundancy required and add that number of ports to our port count total. Some servers may or may not require a redundant connection. It may well depend on the number of users accessing that server and their demand for getting to the data storage pool. Redundant switches do not require ISLs between the chassis; redundant fail-over pathing can take place on the server or on the storage unit itself.

## Target Devices

Third to be examined is the target devices or storage devices. What type of storage devices is being attached, what servers will require access to it and how is it going to be controlled and by which component of the SAN? Storage should be viewed as the most dynamic portion of the network. One of the key advantages of a SAN is the ability to add "storage on demand." This storage on demand capability insures that applications and servers that need the storage have immediate access to it. Depending on the type of storage, additional ports may need to be factored into the original design to incorporate the need for additional storage without redesigning the infrastructure. And again, the physical location of the storage is an important factor in the design. Ports for expansion have to be in the right physical location and if the design is a multi-chassis design, we must plan for the ISL requirements also.

## Centralized Backup

In addition to the storage devices themselves, other target devices like Tape Libraries play an important part in your SAN infrastructure. Backup is often cited as the primary reason for architecting and installing a SAN. A proper SAN infrastructure can increase the effectiveness of the backup application and shorten the window needed to complete backups. Additionally, through the high-speed bandwidth of Fibre Channel, it may be appropriate to backup during production times without little if any degradation in SAN performance. This allows for greater flexibility in getting valuable data secured regularly. Tape libraries and devices can be attached to the SAN infrastructure and have visibility to all devices attached to the SAN allowing for all devices to be incorporated into the backup strategy.

## Scaling the SAN

The larger the SAN, the more planning is needed; however the better the planning, the greater the return. Return with regards to SAN supportability and easy, scalable SAN expansion in the future. If the SAN fabric requires multiple switch chassis to be interconnected to achieve a larger user port, the architecting of this is very critical to a successful implementation and deployment.

There are a number of methods that can be used to expand the user port count in a SAN. They range from deploying a large port count switch or director (64 to 128 port solutions) to interconnecting the 8 and 16 port chassis together to form larger port count fabrics. The large port count switches or directors are an effective way of building a large SAN. The InRange FC/9000 that's powered by QLogic technology provides a high level of redundancy into the architecture. There is no single point of failure in the director-class switch product. Each component is redundant and hot swappable providing non-stop performance. Another great advantage is ease-of-management. All 64/128 ports are managed from a single management console. The obvious consideration in selecting a director is the cost. However, if customers require 99.999% uptime, this product meets the criteria and is the best choice. The other option is to expand the fabric beyond a single chassis by interconnecting chassis using ISLs. This method is a great method for customers who have devices at multiple locations. The geographically remote sites can be connected together to allow all the devices to be incorporated into a central backup scenario. The distance between switches can be up to 10 km and additional distance can be supported (up to 50 km) on the 8 and 16 port switches through long-wave SFPs/GBICs. This method gives the customer a much cheaper entry point to creating a scaleable SAN. Factored into this type of architecture is the ability to scale the SAN without disruption of normal data traffic.

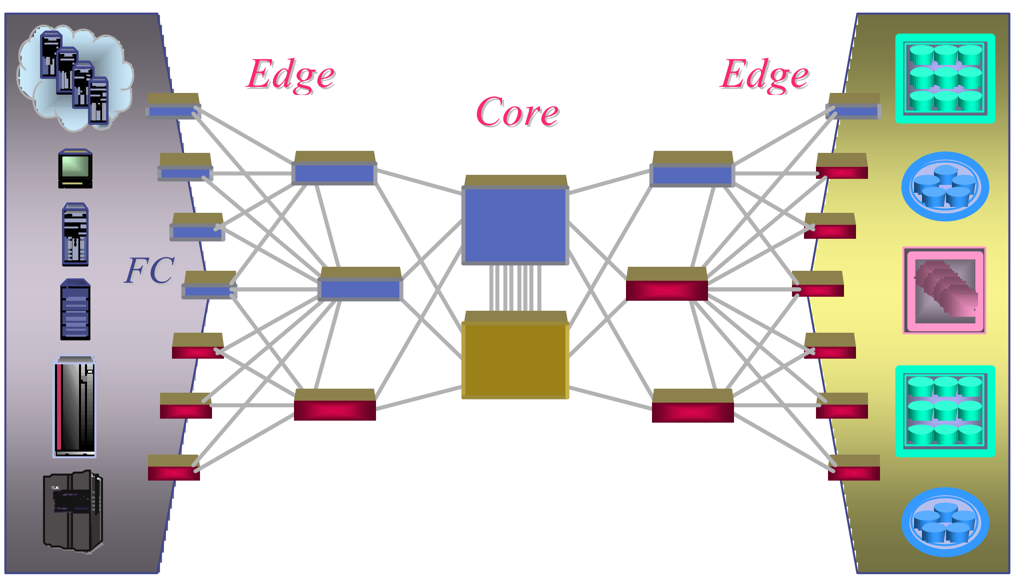

Keep in mind when designing SANs that servers never (rarely) talk to each other, so if you can physically attach servers and storage devices that normally work together on the same chassis, you can get the best throughput possible. Architecting a SAN with this in mind will keep you from making the biggest mistake found in multi-chassis SANs today. If servers that access storage are all connected to the same chassis and all storage that they access is connected on another chassis, all traffic must go through the ISLs. This creates a potential debugging nightmare for the IT department. If this is physically the way it has to be, design-in extra ports for ISLs. That way, when the question is brought up, "How come at 9:15am my file access is slow", the IS department has the ability to increase the bandwidth on the fly to quickly resolve the issue. Also, if the configuration is running fine, additional users can be added to the ports as the requests come in without disrupting the SAN environment.

## Fabric Allocation

The next step in building an effective and supportable SAN is looking at the methods of allocating and managing the resources that we've added to the SAN. There are a couple of different ways to do this and each has some advantages. What follows is some guidelines to help select the method that's right for your SAN.

The current methods for controlling disk access are LUN masking and switch zoning. The LUN masking can be controlled from either the host or server (initiator) side or from the storage (target) device itself. LUN masking is a means of slicing up the physical storage into logical partitions that can be presented and accessed by the servers they are intended for. Most of the high-end storage devices such as Hitachi Data Systems (HDS) and EMC support this method through their RAID controller. The switch can be zoned to further insure the separation of storage devices and servers further securing access to unauthorized devices. By combining switch zoning and LUN masking at the device level, this prevents human error from happening and causing storage to be presented to the wrong server inadvertently. This becomes critical when Microsoft Windows NT and UNIX servers are accessing the same storage device.

## Application Considerations

The applications that are implemented to solve the end users' problems must be understood and examined. Their requirements may dictate some infrastructure requirements that must be architect into the SAN design. If the applications require 100% availability, then all devices that support this application must be designed with redundant connections. Servers will require multiple HBA cards to ensure multiple fail-over paths are available and switches that are deployed must support the high availability requirements that the applications demand. When looking at the total solution provided, some application servers may be attached to the SAN but not carry the high availability requirements of other devices. This can lead to a design that includes high availability director-class switches at the core that attach the high-availability application devices to the storage and also allows for edge switches to interconnect less sensitive departmental configurations. This architectural design facilitates central backup and resource sharing.

## Summary

The architecting of the SAN is one of the most important steps and if done correctly can have long-lasting positive effects on the implementation. The first step to the proper SAN implementation is understanding exactly what problem the end user is trying to solve by implementing a SAN. With that information in hand, we can begin to build the proper infrastructure to support the desired solution set. Also, by planning for future expansion up-front, additions can be integrated without disruption to the installed environment. A thorough understanding of the traffic loads and bandwidth expectations can facilitate a design to support these expectations and avoid a hasty and disruptive patch down the road. Pre-planning is often overlooked and can lead to implementations that don't fit customer expectations and requirements. Time and resources must be allocated to plan and design the infrastructure that will support the ultimate SAN solution.