

Using VSANs and Zoning in the Cisco MDS 9000 Family of Multilayer Fibre Channel Switches

Purpose

The purpose of this white paper is to provide an overview of VSAN and zoning features within the Cisco MDS 9000 Family of Multilayer storage switches and their practical use within storage networking environments. VSANs and zones are very complementary technologies and represent two powerful tools for the SAN designer. This paper explores the benefits and limitations of zoning along with the enhancements offered by VSANs in terms of cost, scalability, availability, and security within a storage network.

Introduction

There are several factors to be considered today when designing and deploying storage area networks (SANs). Attributes such as high availability, scalability, and security of the network all must be carefully considered when designing and implementing the network itself. The Cisco MDS 9000 Family of multilayer directors and fabric switches provide many enabling software and hardware features to achieve these critical design factors thereby offering the SAN designer a more cost-effective and reliable SAN deployment.

Fabric zoning is a fundamental feature in today's Fibre Channel switching products. Zoning provides a means of restricting visibility and connectivity between devices connected to a common Fibre Channel

SAN. While zoning provides a basic security function within a fabric, it does not provide any enhancements in terms of the scalability and availability of the fabric.

To enhance fabric scalability and availability, and further augment the security services offered by fabric zoning, Cisco has developed a new technology called Virtual SANs or VSANs within the Cisco MDS 9000 Family of Multilayer directors and fabric switches. VSANs combined with hardware-enforced zoning provide the SAN designer with new tools to highly optimize SAN deployments in terms of scalability, availability, security and management. VSANs provide the ability to create completely isolated fabric topologies, each with its own set of fabric services, on top of a scalable common physical infrastructure. As each VSAN possesses its own zoning service, zoning is then configured within each VSAN independently and has no affect on any other VSAN and zoning service.

The purpose of this paper is to provide an understanding of the differences between VSANs and zoning and how the two technologies are complementary to one another. Each technology is designed to address different problems within the Fibre Channel network space. Understanding the benefits of both technologies is key to effectively deploying them with maximum benefit.



Fabric Zoning Service

The zoning service within a Fibre Channel fabric was designed to provide security between devices sharing the same fabric. The primary goal was to prevent certain devices from accessing other devices within the fabric. With many different types of servers and storage devices on the network the need for security is critical. For example, if a host were to gain access to a disk being used by another host, potentially with a different operating system, the data on this disk could become corrupted. To avoid any compromise of critical data within the SAN, zoning allows the user to overlay a security map dictating which devices, namely hosts, can see which targets thereby reducing the risk of data loss.

Zoning does however have its limitations. Zoning was designed to do nothing more than prevent devices from communicating with other unauthorized devices. It is a distributed service that is common throughout the fabric. Any installed changes to a zoning configuration are therefore disruptive to the entire connected fabric. Zoning also was not designed to address availability or scalability of a Fibre Channel infrastructure. Therefore while zoning provides a necessary service within a fabric, the use of VSANs along with zoning provides an optimal solution.

Benefits of Zoning

- **Enhanced Device Security**—By deploying zoning within a Fibre Channel fabric, device access is limited to devices within the zone. This allows the user to segregate devices based on access to a particular storage device (target). This is generally an absolute requirement when dealing with multi-OS environments accessing the same physical fabric. Most operating systems cannot read or understand the block layout or file system structure of different operating systems. Should, for example, a Windows-based host see a disk being utilized by an AIX host, the file structure would look foreign and may be viewed by the Windows system as being a corrupt Windows volume that is available and must be repaired. Once the Windows server writes its 'signature' to that disk it will now look foreign to the original AIX server and data corruption will occur. This is exactly the type of scenario zoning was designed to prevent and the reason zoning is so critical. An even further zoning capability is provided by the Cisco MDS 9000 Family of products in that all zoning is implemented in hardware. Each time a zone configuration is installed, whether it is port-based or WWN-based, the configuration is actually downloaded into hardware and a frame-by-frame hardware-based filter is installed to ensure no frames shall pass that are disallowed by the zoning configuration.
- **RSCN Suppression**—By implementing zoning within a Fibre Channel network, Registered State Change Notifications (RSCNs) are isolated to the devices within the zone in which the state change happened. This causes fewer disruptions to devices within the fabric, especially to devices that have no reliance or are not in the proximity of the actual state change. In some instances however, such as a storage array interface that may reside in several zones simultaneously, zoning cannot provide complete protection.
- **Multiple Zones per Port**—When utilizing zoning to isolate device connectivity, a single shared device may have to exist in multiple zones at the same time. This typically occurs when multiple servers require access to the same disk subsystem interface. Each server usually resides in a separate zone with the disk subsystem interface and hence the disk subsystem interface resides in a zone with each server. This allows access to the common device, the storage subsystem interface, while preventing the servers themselves from communicating with each other.



Limitations of Zoning

- **Scalability**—While Fibre Channel networks today are relatively limited in size compared to a typical IP network, the increasing scale of storage networks in the future will likely reveal design limitations. Fibre Channel not only dictates a limit on the number of domains (typically one switch per domain) within the fabric, but also the number of ports within the domain. Zoning does nothing to provide a higher level of scalability for the network. Even though devices may be in separate zones, such zones still reside in one common fabric bound by the addressing limitations and routing scalability within the fabric. In addition, as fabrics grow larger, so too will the number of deployed zones and the control plane bandwidth required to manage the higher number of zones.
- **Fabric Availability**—While zoning, when enforced by hardware, can provide significant security benefits within a network it does nothing to provide enhanced availability within the fabric itself. Should the name server stop responding or FSPF begin to route erratically this will have a very disruptive effect on the entire storage network as common fabric services are shared across all zones within a fabric. In addition, due to the distributed nature of zoning, any changes made to an active zone set have fabric-wide impacts and potential disruptions as the zone set changes are installed.
- **Traffic Management**—Zoning is very effective in managing connectivity between end devices within a storage network. However, zoning does not offer any capability to control the path selection and flow through a storage network between zoned devices. As long as two devices are allowed to communicate as dictated by the zoning configuration, their flows can traverse any path within the SAN as determined by routing protocols within the fabric. The requirement to engineer traffic flows becomes increasingly important as storage networking environments grow to ensure fabric bandwidth is utilized in an optimal manner.
- **Manageability**—As zoning is a common distributed service throughout the fabric, it is managed as a common service. Therefore as organizations look to collapse multiple applications on fewer larger SAN fabrics, zoning will still exist as a common distributed service. Zoning is not a service that was designed to have its management partitioned amongst different management groups. Therefore, if an organization were to collapse an HR application infrastructure along with an Engineering application infrastructure onto the same SAN fabric, extensive and critical coordination is required between the application owners to ensure the one common zone set shared by the two application groups doesn't become corrupted or unintentionally altered.
- **Zone Management Security**—The zoning service is implemented as a distributed service within a Fibre Channel storage network. A distributed database is synchronized and maintained within all member switches of the storage network. As a distributed service, the zoning database can be updated by any member switch of the fabric. Therefore, although zoning is a security service, its very implementation is relatively insecure. If a member switch within the fabric were to be compromised, fabric-wide zoning configurations could be adversely affected. While new proposals exist within the ANSI T11 community to secure the zoning service, today's zoning implementations have their weaknesses.
- **Accounting**—Zoning is a relatively dynamic service that can be modified frequently for a variety of application-oriented reasons. As such, it is exceedingly difficult to account for bandwidth usage amongst application groups that may be defined by their zoning definition. Therefore if a common fabric is built to contain multiple applications, each divided by a zoning definition, it is generally impossible with today's hardware to isolate and account for usage on a per-zone basis.



Virtual SANs (VSAN)

Virtual SANs offer the ability to scale SANs beyond current limitations in a resilient, secure, cost-effective, and manageable fashion. Using VSANs, SAN designers are able to build larger consolidated fabrics and still maintain the required security and isolation between applications beyond what is currently offered through zoning.

Today, SAN designers build separate fabrics, otherwise known as SAN islands, for a variety of reasons. A SAN island refers to a completely physically isolated switch or group of switches used to connect hosts to storage devices. Reasons for building SAN islands may include the desire to isolate different applications into their own fabric or to raise availability by minimizing the impact of fabric-wide disruptive events. In addition, physically separate SAN islands also offer a higher degree of security as each physical infrastructure contains its own separate set of fabric services and management access. While these are valid reasons for building separate fabrics, this practice can quickly become costly and wasteful in terms of fabric ports and resources. The prospect of additional separate fabrics means more hardware, more ports, more cost, more devices to manage, and typically underutilized hardware. Another drawback to building out separate SAN islands is the inflexible nature of completely isolated islands in terms of resource re-allocation. If one fabric has many unused ports and another fabric is short of ports, one cannot simply reassign the unused ports where they are required.

To help achieve the same isolated environments while eliminating the added expense of building physically separate fabrics, Cisco has introduced the Virtual SAN (VSAN) within the Cisco MDS 9000 Family of Multilayer directors and fabric switches. A VSAN provides the ability to create separate virtual fabrics on top of the same redundant physical infrastructure.

Using VSANs, SAN designers can raise the efficiency of a SAN fabric and alleviate the need to build multiple physically isolated fabrics to meet organizational or application needs. Instead, fewer less-costly redundant fabrics can be built, each housing multiple applications, and still provide island-like isolation. Spare ports within the fabric can be quickly and non-disruptively assigned to existing VSANs thereby providing a clean method of growing application-specific SAN islands virtually.

Another major benefit of the VSAN feature is in terms of its contribution to high availability. VSANs provide not only hardware-based isolation, but also a full replicated set of Fibre Channel services for each VSAN. Therefore, when a VSAN is created, a completely separate set of fabric services, configuration management capability, and policies are created within the new VSAN. Examples of fabric services that are created include *name server*, *zone server*, *domain controller*, *alias server*, and *login server*. This replica of services provides the ability to build the isolated environments needed to address HA concerns on top of the same physical infrastructure. For example, an installment of an active zone set within VSAN1 does not affect the fabric in any way within VSAN2.

VSANs offer a great deal of flexibility to the SAN designer in addition to an infrastructure cost saving. For example, the MDS 9000 Family of products support 1000 VSANs per physical fabric. Each VSAN can be selectively added or pruned from a trunk link so as to control the propagation of VSANs through the fabric. In addition, special traffic counters are provided to track statistics on a per-VSAN basis.


VSANs also provide a method to interconnect isolated fabrics in remote data centers over a common long-haul infrastructure. Since the frame tagging is done in hardware and included in every EISL frame, it can be carried across transports including DWDM, CWDM, and FCIP. Therefore, traffic from several VSANs can be multiplexed across a single pair of fibers and transported a greater distance and yet still remain completely isolated. VSANs bring scalability to a new level by leveraging a common redundant physical infrastructure to build flexible isolated fabrics to achieve HA goals.



Each separate virtual fabric is isolated from one another using a hardware-based frame tagging mechanism on VSAN member ports and EISL links. The Enhanced ISL (EISL) link type has been created and includes added tagging information for each frame within the fabric. The EISL link is supported on links interconnecting any MDS 9000 Family switch products. Membership to a VSAN is based on physical port and no physical port may belong to more than one VSAN. Therefore, whatever node is connected to a physical port becomes a member of that port's VSAN.

Benefits of VSANs

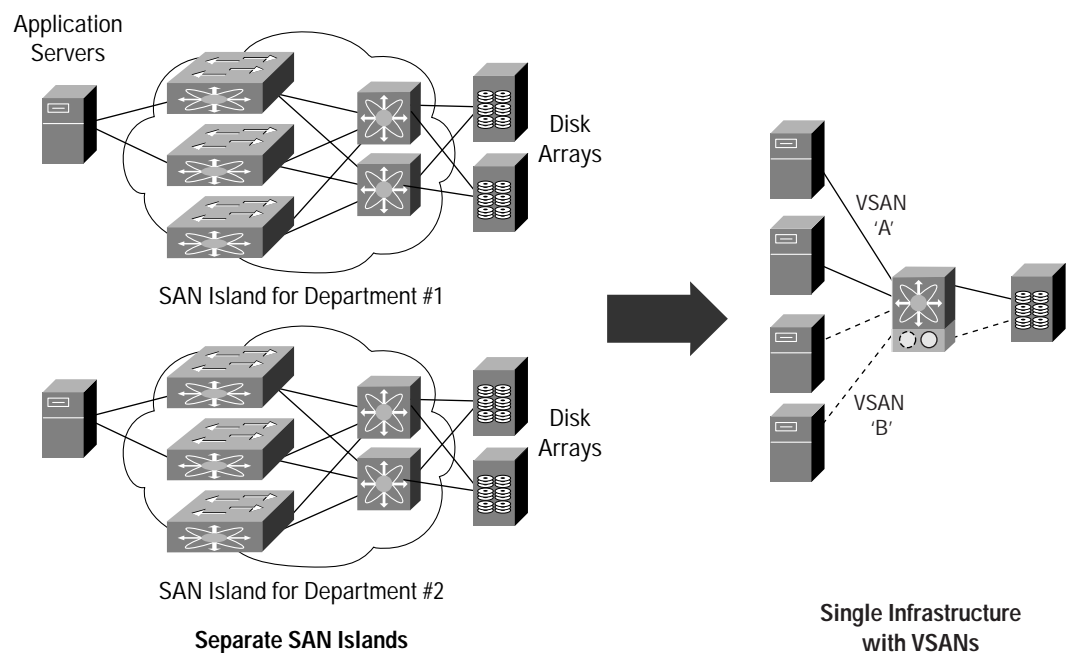
- **Virtual SAN Islands**—Many reasons still exist today in storage networking that mandate SAN designers build separate fabrics for different applications. Many such reasons have been discussed in this paper including departmental security concerns and system OS conflicts. VSANs offer the SAN designer a way to consolidate what would amount to multiple costly physical SAN islands onto a more cost-effective common redundant SAN fabric. Using VSANs the same security and isolation as achieved by building physically separate islands can be replicated virtually on the same physical infrastructure.
- **Transparent to End Devices**—VSANs do not require any special awareness, configuration, or software on the SAN end devices such as hosts/HBAs or disk subsystems. Traffic is tagged as it enters a switch and the tagging is removed once a frame leaves a switch for an Nx_Port.
- **ISL Trunking**—Even though each VSAN represents a separate fabric and traffic cannot cross VSANs, the MDS 9000 Family supports 'trunking' of VSANs over a Trunking E_Port (TE_Port). Using TE_Ports provides several advantages. While traffic cannot span VSANs, multiple VSANs can share the same ISLs. Multiple VSANs can share the bandwidth of ISLs for increased ISL utilization. This could significantly reduce the number of ISLs needed in a given deployment. Trunking of VSANs across TE_Ports also allow for a basic form of traffic shaping. Because VSANs can be individually assigned to a Trunk, VSANs with lower priority traffic could be assigned to ISLs that may have a higher path metric thus leaving shorter paths for higher priority traffic.



- **Fabric Availability**—Each VSAN includes separate instances of all fabric services. This provides for a much more stable fabric as not only are fabric service failures isolated per VSAN but fabric level events such as *Build Fabric* or *Reconfigure Fabric* are also isolated per VSAN. Should a switch need to be added to an existing network, only the VSAN(s) required on the new switch will experience a fabric rebuild or reconfigure and not the remaining VSANs on the entire network. The VSAN capability limits any possible disruption to devices that need to be in an isolated environment without the need for physical isolation. The resultant increase in availability offered by VSANs also allows the SAN designer to build larger and more cost efficient SANs instead of smaller SAN islands.



- **Fabric Scalability**—Fibre Channel has several limitations when it comes to the scalability of the network. However, VSANs provide a way to scale a fabric. When deploying VSANs within a physical infrastructure, the Fibre Channel addressing scheme must only be unique per VSAN. Within a standard fabric only 239 domains (switches) are allowed. This effectively limits the scalability of the fabric. By deploying VSANs the Fibre Channel addressing scheme is implemented on a per VSAN basis. Now, up to 239 domains can exist per VSAN, thus extending the scalability within the physical infrastructure.
- **Collapsed Physical Infrastructure**—A common implementation of Storage Area Networks is to deploy 'SAN Islands'. Each application, operating system, or business unit has its own SAN fabric. When deploying with this design, hardware is typically underutilized and wastes costly hardware and management resources. In contrast, deploying VSANs provides the ability to collapse the many individual SANs into a larger single infrastructure. This reduces hardware costs, and increases the manageability of the entire network while maintaining the stability and traffic isolation of the SAN island model.



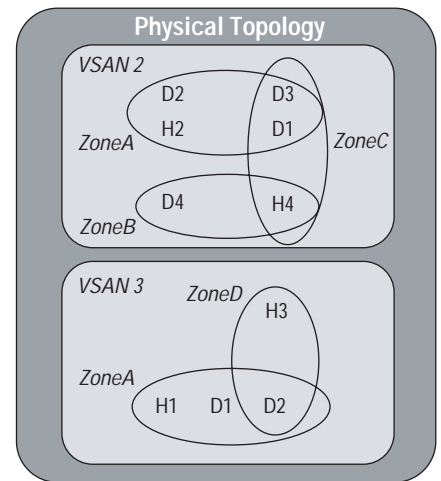
- **Traffic Management and Service Differentiation**—The implementation of VSANs gives the SAN designer more control over the flow of traffic and its prioritization through the network. Using the VSAN capability, different VSANs can be prioritized and given access to specific paths within the fabric on a per-application basis. Using VSANs, traffic flows can be engineered to provide an efficient usage of network bandwidth. One level of traffic engineering allows the SAN designer to selectively enable or disable a particular VSAN from traversing any given common VSAN trunk (EISL) thereby creating a restricted topology for the particular VSAN. A second level of traffic engineering is derived from independent routing configurations per VSAN. As discussed, the implementation of VSANs dictates that each configured VSAN support a separate set of fabric services. One such service is the FSPF routing protocol which can be independently configured per VSAN. Therefore, within each VSAN topology, FSPF can be configured to provide a unique routing configuration and resultant traffic flow. Using the traffic engineering capabilities offered by VSANs allows a greater control over traffic within the fabric and a higher utilization of the deployed fabric resources.



- **VSAN Management Security**—Unlike zoning, the VSAN service is not a distributed service within the fabric. VSAN configuration is local to each switch and configuring VSANs on one switch does not affect the configuration of any other switch within the network. Although one could use the Cisco Fabric Manager to configure VSANs across many switches within the network, each switch is individually configured by the tool. Therefore, each switch only enforces the VSAN configuration locally configured on the switch itself. Using the roles-based configuration security within the Cisco MDS 9000 Family of products, VSAN configuration can further be limited to selected users on selected switches.

Using VSANs and Zoning

VSANs and zoning within the MDS 9000 Family of products are two powerful tools to aid the SAN designer in building robust, secure, and manageable networking environments while optimizing the use and cost of switching hardware. In general, VSANs are used to divide a redundant physical SAN infrastructure into separate virtual SAN islands each with its own set of Fibre Channel fabric services. By each VSAN supporting an independent set of Fibre Channel services, a VSAN-enabled infrastructure can house numerous applications without the concern for fabric resource or event conflicts between these virtual environments. Once the physical fabric has been divided, zoning is then used to implement a security layout within each VSAN that is tuned to the needs of each application within each VSAN. The following chart summarizes the primary differences between VSANs and zones. The figure to the right also shows the relationship between VSANs in zones in a common physical fabric. VSANs are first created as isolated fabrics within a common physical topology. Once VSANs have been created, individual unique zone sets can then be applied as necessary within each VSAN.



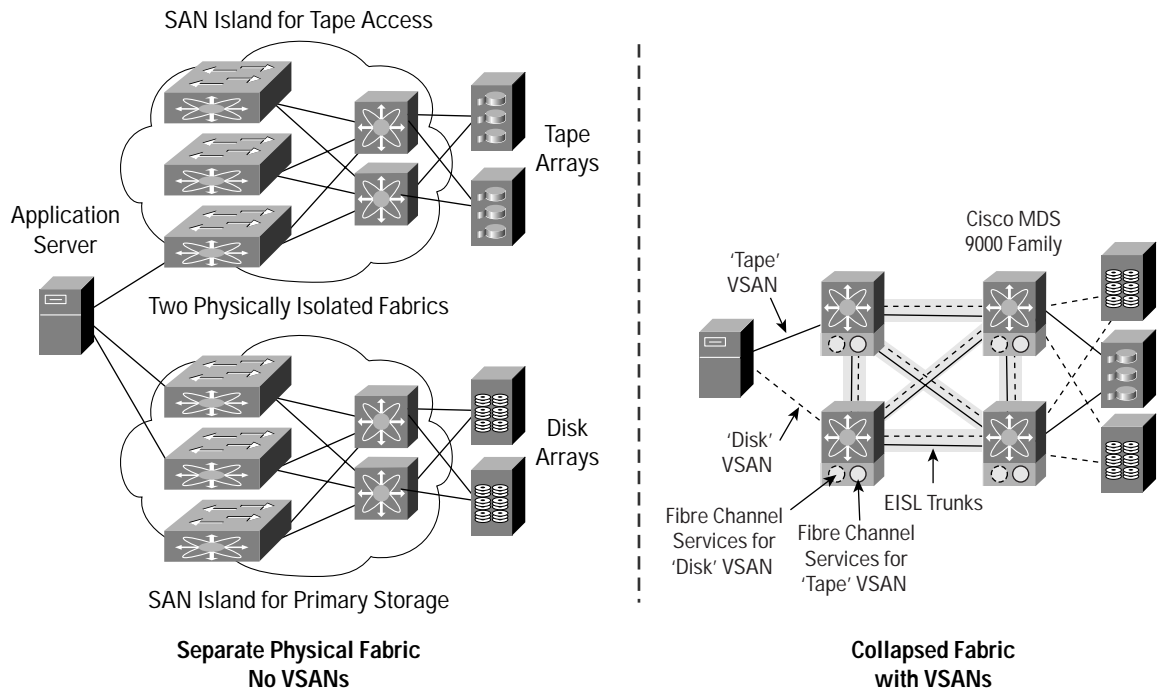
	VSANs	Zoning
Maximum Per Switch/Fabric	1024 per switch	1000+ zones per fabric (VSAN)
Membership Criteria	Physical Port	Physical Port, WWN
Isolation Enforcement Method	Hardware	Hardware
Fibre Channel Service Model	New set of services per VSAN	Same set of services for entire fabric
Traffic Isolation Method	Hardware-based tagging	Implicit using hardware ACLs
Traffic Accounting	Yes per VSAN	No
Separate Manageability	Yes per VSAN (future)	No
Traffic Engineering	Yes per VSAN	No

VSANs can be used wherever the call for SAN islands exists. The following section outlines some common design examples using VSANs. Within each of the following designs are customized and independent zoning configurations per VSAN based on the individual application requirements.



The Backup VSAN

It is common practice for an independent and physically isolated backup SAN to be created strictly to carry backup traffic. Using VSANs, an additional VSAN in a common physical infrastructure could be created to carry tape traffic only. Using this design the SAN designer can alleviate the cost of building a physically isolated SAN for backup and yet still achieve the same level of isolation.



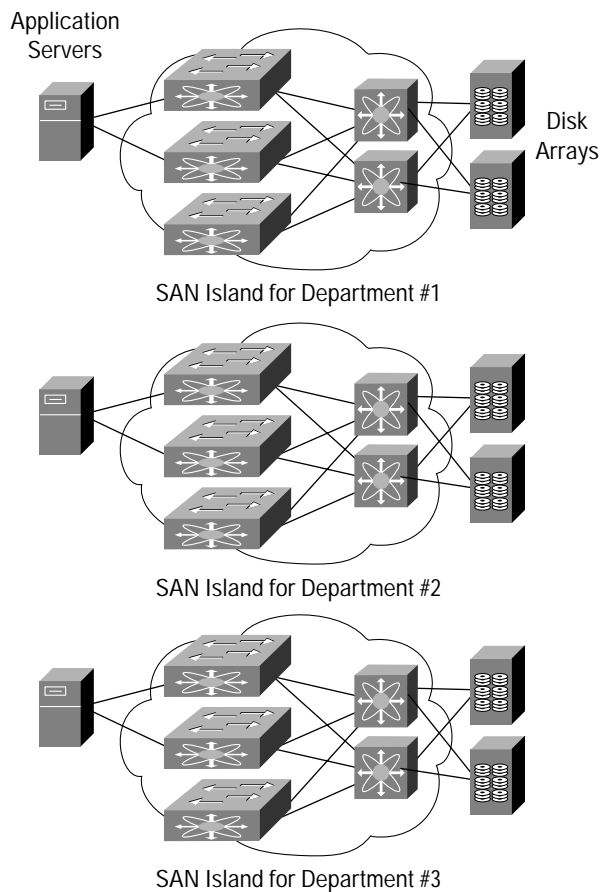
Departmental VSANs

Many reasons still exist today that warrant building separate SAN islands. Many such reasons have been mentioned in this document. The need to build physically separate fabrics has a hefty price tag associated with it due to the excessive hardware required. Due to the fact that each SAN island is treated as its own physically separate fabric, it must be managed independently and also capacity-planned independently. Each SAN island is typically over provisioned to satisfy additional growth requirements of an application and, until such growth requirements are realized, these extra ports will remain unused. Just as storage consolidation has enabled better utilization of storage, so to does consolidation of SANs enable better utilization of SAN network hardware.



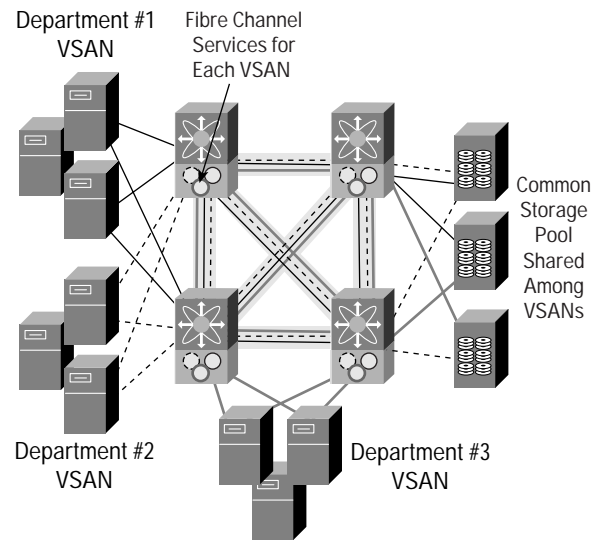
As an example, if a departmental application requires 42 SAN connectivity ports, typically 64 ports will be provisioned in an isolated SAN island to accommodate growth. However, should it be determined at a later date that this application will not require all these extra ports, it is physically impossible to move these ports and re-deploy them for another application. Therefore, the effective cost per port in this example is 152% higher when the unused ports are factored into the design (64 ports deployed for 42 connection points). The cost becomes even higher when you factor in the cost of any ISLs required to build the 64 port SAN.

Using VSANs in this example, the initial deployment can be exactly 42 ports out of the larger redundant fabric. A VSAN would be created and 42 ports would be assigned to the VSAN thereby optimizing the effective cost per SAN port. In the future, should more ports be required for growth, they can simply be assigned non-disruptively from the pool of unused ports in the physical fabric. Unused ports can easily be re-deployed for other applications in other VSANs through a simple software configuration.



Application/Department-Based SAN Island

- Separate physical fabrics
- Over-provisioning ports on each island
- High number of switches to manage
- No flexibility to move unused ports



Collapsed Fabric with VSANs

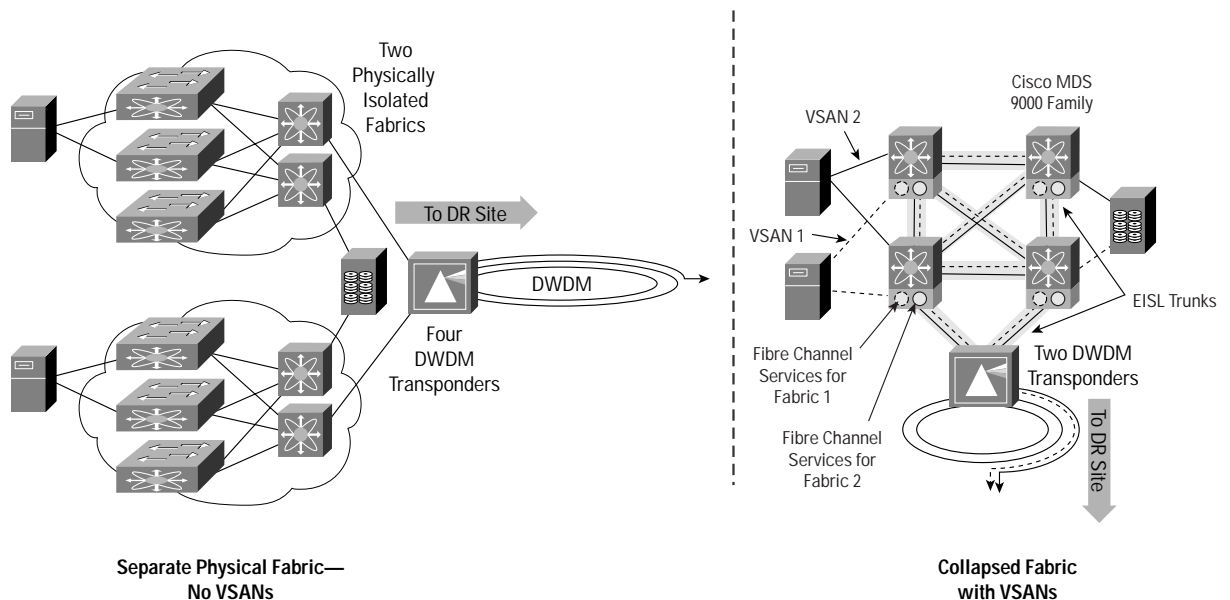
- Common redundant physical infrastructure
- Less over-provisioning required—lower \$\$
- Fewer switches to manage
- Move unused ports non-disruptively



Extended VSANs

Another important attribute of VSANs is their ability to be carried between data centers over various transports while still maintaining isolation. As customers continue to build distributed application environment across multiple data centers, whether for geographic applications or disaster recovery reasons, the need to maintain VSAN segregation is absolutely required.

Since VSANs involve the explicit tagging of each frame with a membership identifier, these tagged frames can be transported across a variety of mediums. Such mediums may include optical transports such as DWDM (Dense Wave Division Multiplexing) and CWDM (Coarse Wave Division Multiplexing) along with Fibre Channel over IP (FCIP) methods using any transport that can carry IP (Packet-over-SONET, ATM, Frame Relay, fixed wireless, etc.). As an example, using VSANs can provide virtual isolation of data from multiple applications which can be transported across a metro or wide area to a common disaster recovery site. At the disaster recovery site, separate implementations of DR configurations can be used in isolation of one another using the VSAN-tagged data.



Leveraging the VSAN technology can significantly reduce the costs associated with transporting storage traffic from multiple SANs across a metro or wide-area network. Today, multiple isolated SONET circuits or DWDM wavelengths are required to transport storage traffic associated with isolated SAN islands. However, using the isolation capability offered by the VSAN technology, traffic from multiple SAN islands can be transported through fewer circuits and/or wavelengths while still maintaining strict isolation. Therefore isolation capability offered by the VSAN technology can directly reduce the cost of additional transport hardware including SONET client-side interfaces and DWDM transponders.

Conclusion

VSANs provide a solution to many of the challenges experienced by SAN designers today. As a superset of the zoning functionality, VSANs provide a complementary and necessary function in addition to zoning.

Many challenges exist today for the SAN designer to build resilient storage networking infrastructures while keeping a close eye on costs. However, cost savings cannot be at the expense of availability, security, or manageability of a SAN infrastructure. With current Fibre Channel switching products, SAN designers have been forced to build multiple physically separate SAN infrastructures to help ensure availability and security between application environments. This resultant design practice has complicated the task of capacity planning and resource optimization due to the physical separation of switch, disk and tape resources. Resources, especially costly switch ports, cannot be easily migrated from one application to another. Therefore a resultant over provisioning of ports has arisen thereby significantly increasing the effective cost per port of the storage networking infrastructure. While zoning features help contribute to an overall secure SAN infrastructure, it does not do anything to address the above challenges.

The introduction of Virtual SANs within the Cisco MDS 9000 Family of multilayer directors and fabric switches now gives SAN designers the flexibility to build consolidated and cost-optimized storage networks while adhering to the security and availability requirements of mission critical applications. The VSAN capability provides a method to optimize Fibre Channel infrastructure usage and maintain the ability to reassign and migrate fabric and storage resources from virtual fabric to virtual fabric non-disruptively. In addition, detailed VSAN-based statistical accounting is provided to aid in the capacity planning and possible charge-back exercises associated with fabric resources. The VSAN technology is just one of the tools in the Cisco MDS 9000 Family of switches that help SAN designers build the multilayer storage network.

For More information

Contact the authors:

Dan Hersey (dahersey@cisco.com)

Tom Nosella (tnosella@cisco.com)



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www.europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992-2002, Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0208R) EW/LW3749 12/02