

Cisco Virtual SAN Advantages and Use Cases

The Challenge

Application data is accumulating at an increasingly fast pace, and to control costs, IT departments are looking at the benefits of economies of scale. By consolidating the data of multiple user communities onto a single SAN infrastructure, IT departments can achieve significant cost reduction.

A shared storage infrastructure must provide benefits and flexibility that are as good as those previously available to each individual user community. At the same time, the consolidation must be transparent to the user.

Security, for instance, is an implicit requirement of a shared storage infrastructure. The need for data privacy is of strategic importance. Different departments within an enterprise, as well as different clients of a common service provider, require assurances that data on shared storage infrastructure cannot be accessed by other users. Even within the same organization, compliance rules often determine who is authorized to access particular data, and this access control requires appropriate network segregation.

Advanced storage hardware that is capable of advanced storage management cannot provide a complete solution alone. Corresponding advanced capabilities in the SAN are required to achieve a complete solution.

The Solution: Virtual SANs

Cisco foresaw these requirements and developed Virtual SANs (VSANs) in 2002 to address the challenges of the data center. VSANs are integral to the Cisco® MDS 9000 Family of switches and have been engineered into the product line as it was designed. VSANs are logical SANs built on a common physical fabric. Each VSAN provides its own fabric services and is functionally isolated from other VSANs on the same switch. Sharing of common resources across VSANs can be accomplished using Inter-VSAN Routing (IVR). IVR is performed internally on all line cards and requires no additional resources or specialized hardware on the switch.

By using VSANs, customers can consolidate separate physical SAN fabrics into one large fabric for ease of management and reduced Total Cost of Ownership (TCO) while providing secure and reliable isolation of different storage islands. Additionally, the use of IVR to selectively join VSANs allows the sharing of common infrastructure and resources such as tape libraries, further reducing storage and network expenditures.

Cisco's VSAN implementation complies with ANSI T.11 standards and forms the basis for the ANSI Virtual Fabric standard.

Cisco VSAN Technology

Cisco VSAN technology provides secure hardware-based network segmentation, similar to the VLAN technology that is widely deployed in LANs.

The Cisco MDS 9000 Family maintains a VSAN membership attribute for each interface in the fabric. By changing this attribute, any interface in an entire Cisco MDS 9000 Family fabric can be assigned to any VSAN in seconds with a simple configuration command.

Servers and storage devices belonging to the same user community can be deployed in the same VSAN, even though they may be physically connected to any port anywhere in the SAN. Adding or removing a physical server from a VSAN is a simple configuration step, eliminating the need for rewiring. Note that because any fabric port can belong to any VSAN, you do not need to plan in advance which physical server or storage device will be associated

with which user community or segment; after cabling is complete; the assignment is performed with a configuration command. This process is simpler and less disruptive than the re-cabling required with physical SAN islands.

As a Fibre Channel frame from an attached device enters a Cisco MDS 9000 Family switch, it is given a hardware VSAN tag. This tag is maintained as the frame traverses the Cisco SAN fabric. As they traverse the fabric, frames belonging to different VSANs can share Inter-Switch Links (ISLs). The frame maintains its VSAN identifier until it reaches the destination interface, where the tag is stripped off.

From a fabric services perspective, Cisco MDS 9000 NX-OS Software runs an independent set of fabric services (zone server, name server, domain manager, Fabric Shortest Path First [FSPF] routing, etc.) for each VSAN. For IBM Fiber Connection (FICON) VSANs, for the IBM System z environment, each has its own instance of the IBM Control Unit Port (CUP) in-band management server. Configuration and operation of fabric services is confined to an individual VSAN, isolating other VSANs from configuration changes and errors. For instance, a separate instance of the zoning fabric service runs in each VSAN's infrastructure, so a configuration error or a protocol violation that could adversely affect operation has no effect outside that particular VSAN.

Furthermore, fabric events such as registered state change notifications (RSCNs) are contained within the VSAN, in which they originate, isolating incidents and increasing overall fabric uptime. For these and other reasons, Cisco MDS 9000 Family customers have deployed VSANs at a very high rate.

Management

Cisco MDS 9000 NX-OS Software management offers several levels of role-based access control (RBAC). RBAC allows an administrator to be in charge of a specific VSAN without having any visibility into other VSANs.

The administrator with higher privileges can create roles and user accounts and then assign each user account to a role. The role configuration is very granular. For instance, an administrator can have read-only access to a subset of the configuration information of a specific VSAN. This feature allows infrastructure to be shared across a large number of user groups, each of them assisted by a team of administrators having specific skill sets and job assignments.

Security

Whether it occurs within a single enterprise or at a multi-tenant service provider, consolidation of different business functions on the same physical infrastructure must provide the same level of security and isolation as was provided on the physically isolated networks. Although the foundation of fabric security is embedded in the VSAN-capable hardware, the holistic approach to security provided by the Cisco MDS 9000 Family provides peace of mind, even in a highly consolidated environment.

The main elements are:

- Hardware-enforced, standards-based, VSAN segmentation to isolate data traffic
- Secure, individual instances of the fabric services and protocols in each VSAN
- Management performed through secure protocols
- Option to connect to the enterprise authentication, authorization, and accounting (AAA) infrastructure
- Configuration management, distribution, and consistency analysis
- Fabric access security (fabric binding and port security), support for Fibre Channel Security Protocol (FC-SP; an ANSI T.11 standard developed with substantial contribution by Cisco), and enhanced support for zoning (for example, logical unit number [LUN] zoning and read-only zone)
- Integrated protection for data in flight (SAN extension over IP [Fibre Channel over IP, or FCIP] and Small Computer System Interface over IP [iSCSI] can use IP Security [IPSec]) and for data at rest (Cisco Storage Media Encryption [SME])

The Cisco MDS 9000 Family of Fibre Channel switches and directors provide all the tools for implementing the required level of security in an enterprise environment.

VSAN Use Cases


Since the introduction of the Cisco MDS 9000 Family, VSANs have provided value to customers and helped them achieve numerous business goals. This section describes the most common solutions based on Cisco VSAN technology.

Backup SAN Consolidation: Deploying Backup VSANs Compared to Deploying Dedicated SAN Switches

It is common practice for backup traffic to be sent to an independent and physically isolated backup or tape SAN. This solution provides the desired level of isolation and performance for the tape traffic, but it is complex and expensive.

By deploying the VSAN technology, the SAN administrator can create a dedicated VSAN to carry only tape traffic. This design alleviates the cost of building a physically isolated SAN for backup while achieving the same level of isolation (Figure 1).

Figure 1. Consolidating Backup SAN with VSANs



Replication SAN Consolidation: Deploying Replication VSANs Compared to Deploying Dedicated SAN Switches

Data replication between storage arrays is crucial to business continuance and disaster recovery capabilities.

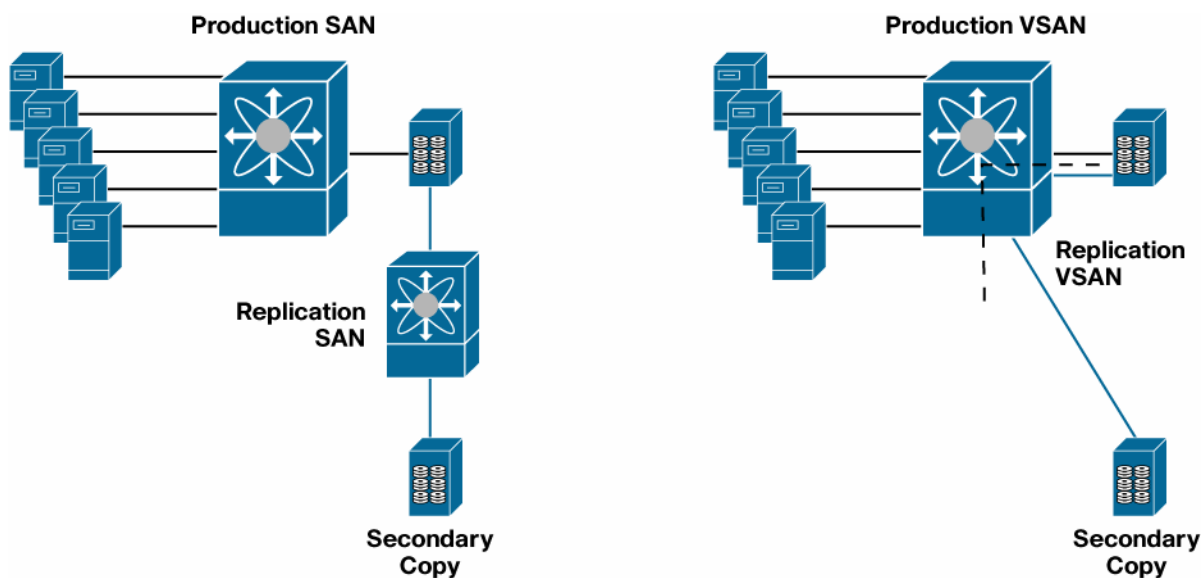
Typical solutions involve a multi-stage approach in which the first data replication occurs synchronously between two disk arrays located in the same data center or in a nearby location, and then a second data replication stage occurs asynchronously, to a remote location to protect against regional disasters.

Data replication is performed by dedicated Fibre Channel ports on the disk array, which are configured to act in pairs as a data source (initiator) or a data destination (target). In the simplest case, the interconnection between arrays is a direct fiber connection, but in more complex deployments, the replication devices are connected to the fabric, and source and destination ports are zoned together like any other initiator and target pair.

Even for replication within the same data center, it has been common to dedicate an isolated SAN to the replication ports, to achieve better configuration control, isolation, and performance for this valuable traffic.

Using Cisco VSAN technology, a VSAN can be dedicated to the replication traffic, providing a level of isolation equivalent to a dedicated SAN, but with the benefits of a consolidated SAN infrastructure (Figure 2).

Figure 2. Consolidating Replication SANs with VSANs



The replication VSAN provides the same level of isolation of faults and fabric services, such as zoning, as a small, isolated physical SAN. RBAC reserves administration privileges to the administrators who focus on the replication infrastructure, limiting the opportunities for human errors.

MAN and WAN Network Consolidation: Consolidating Multiple Types of SANs—Backup, Replication, and FICON—Using VSANs

Another important attribute of VSANs is their ability to be carried along a shared ISL, over various transports, while still maintaining traffic isolation.

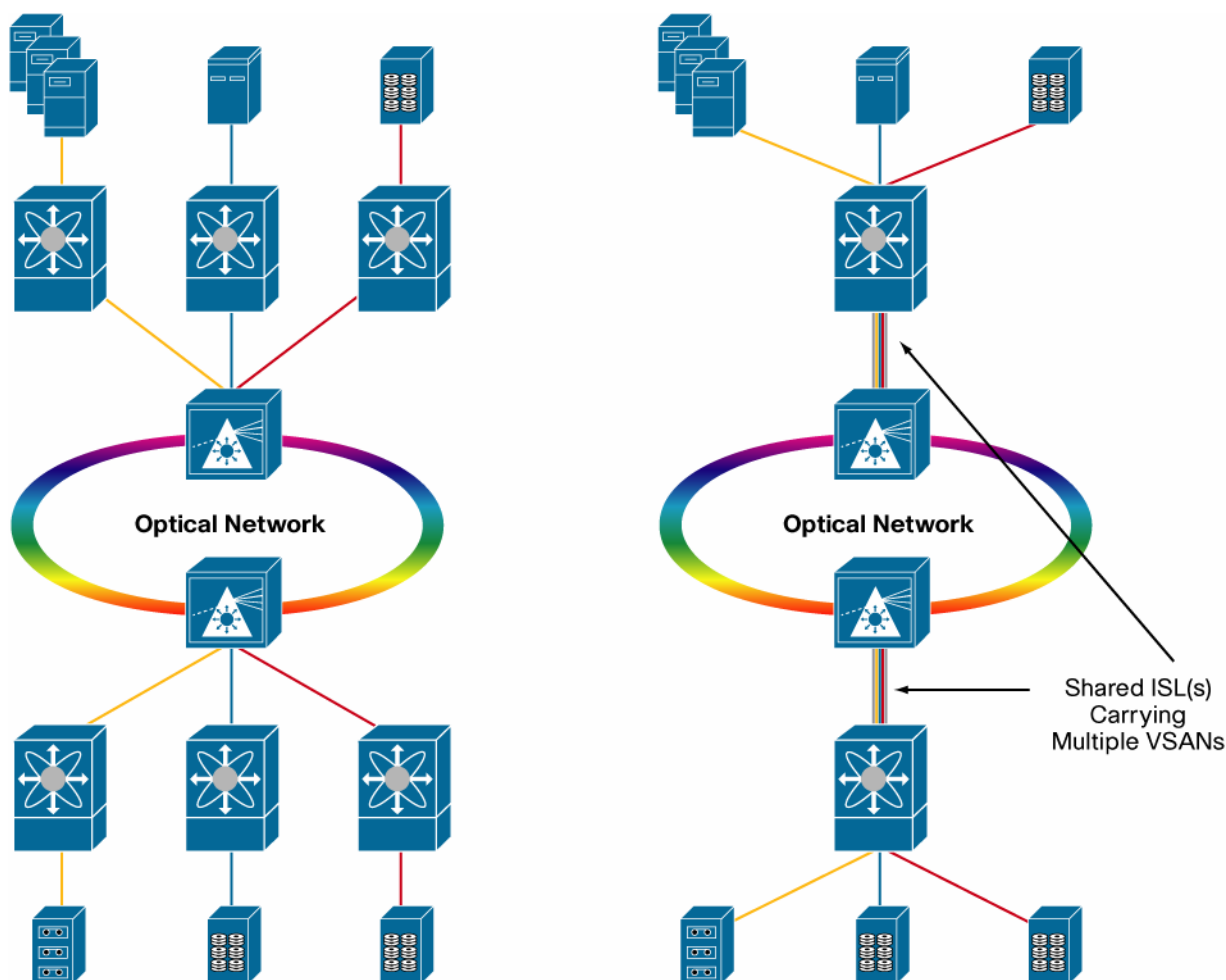
As customers continue to build distributed application environments across multiple data centers, whether for geographically distributed applications or disaster recovery reasons, traffic segregation is absolutely required. Data centers were traditionally connected with multiple Metropolitan Area Network (MAN) and WAN links for the sole purpose of maintaining isolation between applications.

Frames maintain their VSAN tags and properties, even when they are transported across a variety of media. Such media may include optical transports such as Dense Wave Division Multiplexing (DWDM) and Coarse Wave Division Multiplexing (CWDM) along with FCIP, a method using any transport that can carry IP (Packet over SONET [PoS], ATM, Frame Relay, etc.).

VSANs can provide isolation of data from multiple applications, which can be transported across a common MAN or WAN to the disaster recovery site.

In Figure 3, while the traffic originated from the mainframe system, backup servers, and data replication infrastructure is sharing the same physical transport, it belongs to different VSANs and thus is fully isolated. The independence of Fibre Channel services, such as zoning, is maintained across the data centers, and the desired service level can be selected with quality of service (QoS). The amount of SAN and optical transport equipment is also reduced, saving capital, power, cooling, and management costs.

Figure 3. Consolidation of SAN Extension SANs Using VSANs and Shared ISLs



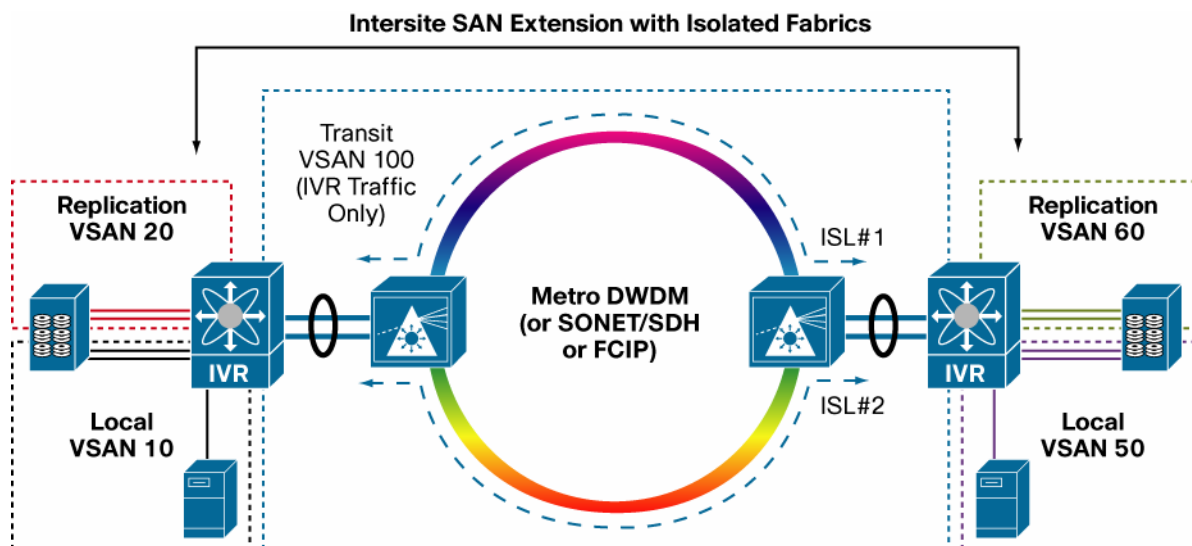
Resilient SAN Extension Solution for Disaster Recovery: Connecting Data Centers without Merging Fabrics (IVR)

To perform data replication across two data centers, some source storage ports in the data replication VSAN in the primary data center must communicate with the corresponding target storage ports in the data replication VSAN in the secondary data center.

If the local and the remote fabrics are merged, because of the tight interaction of Fibre Channel services such as zoning, addressing, and routing, a configuration operation in the local fabric may have an unexpected effect on the remote fabric. When connecting two data centers, communication between remote devices should not occur through merging the fabrics at each site. Keeping the configuration and administration of the two sites independent is the safest approach and provides the best protection against errors.

The Cisco IVR feature can establish tightly controlled communication between devices in different VSANs without the need to merge the individual VSANs, and hence management and fault domains (Figure 4).

Figure 4. Using IVR and a Transit VSAN to Manage Data Center Connectivity



IVR is established by creating an IVR zone. IVR zones can span multiple devices in multiple VSANs; any device pair across VSANs that requires connectivity must belong to the same IVR zone. The local and remote VSANs can be independently managed by different administrators. Configuring a transit VSAN—a VSAN that contains only the ISL and not any server or storage—is not strictly required but can help maintain a clear separation between the local and remote data centers.

Switch-to-Switch Interoperability: Connecting to Other Vendor's Switches without Merging Fabrics or Losing Capabilities

A large SAN is not always a monolithic fabric of homogeneous switches; it may be upgraded in steps or composed by merging smaller existing fabrics, leading to the coexistence of different technologies and code levels.

Fibre Channel standards are open to interpretation, and vendors may choose not to implement optional features or extend features specified in the standards. Some older equipment implements a subset of the standards and cannot interoperate with devices that are standards compliant.

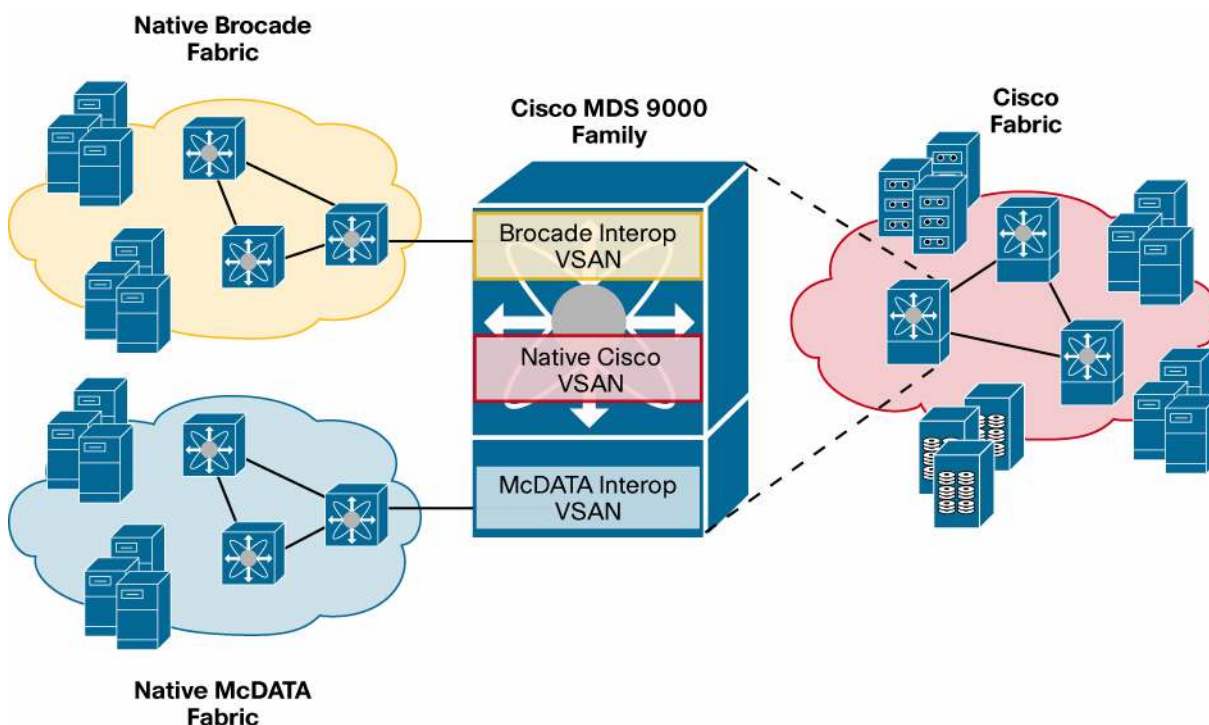
Because these profiles are often proprietary and prevent interoperability, vendors have had to implement interoperability modes to accommodate heterogeneous environments.

When deploying a fabric that requires interoperability between groups of switches of different vendors, one option is to select the interoperability mode that represents the least common denominator between all vendors, usually the basic Fibre Channel standard or even less, losing all the vendor-specific features.

Another option is to force all switches to operate in a vendor-native mode, preserving the features provided by that specific vendor, but not allowing connectivity between different vendor's switches.

A Cisco MDS 9000 Family fabric segmented into VSANs can operate in as many different interoperability modes as VSANs. One VSAN can operate in the most advanced mode, while another VSAN can be connected to an older vendor-specific SAN island. This capability allows the Cisco MDS 9000 Family switch to consolidate multiple physical fabrics, working in different modes, into a single physical fabric, to create multiple-vendor fabrics that may include different and otherwise incompatible firmware levels (Figure 5).

Figure 5. VSANs Selectively Connect Fabrics from Multiple Vendors



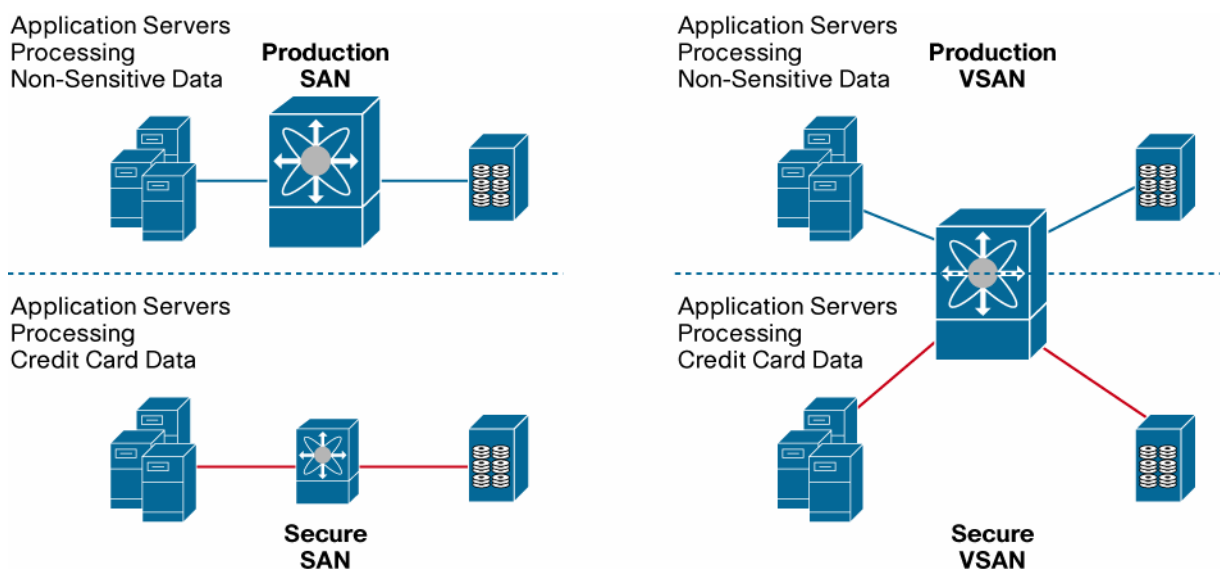
In some cases, the heterogeneous deployment is limited in time to a migration scenario, when a new generation of devices is added to the fabric and will eventually replace the older devices. In other cases, as in a merger or acquisition, the interoperation between different vendors is a long-term requirement, and interconnecting two or more fabrics with a significant loss of features is a significant drawback.

Consolidation of Fabrics with Different Security Scopes: Providing Security and Isolation on a Single Switch

A number of security industry standards and legal requirements are applicable to the SAN infrastructure. In general, in a large IT infrastructure, a portion of the infrastructure is used for sensitive data, and a larger portion carries data that does not require any additional precautions. The portion with sensitive data is subject to specific management procedures, and isolation is required. Historically, this isolation has been physical, resulting in increased capital expenditures (CapEx), operating expenses (OpEx), and complexity.

In Ethernet networking, it is accepted best practice to achieve the required isolation by implementing a separate VLAN. Using the same approach with the Cisco MDS 9000 Family, SAN isolation can be achieved by deploying separate VSANs.

For example, to achieve Payment Card Industry (PCI) Data Security Standard (DSS) compliance, a VSAN can be deployed as a compensating control to isolate the SAN segment that carries credit card data in clear text. In Figure 6, administration of the secure VSAN is assigned to a different administrator, who must be authenticated and authorized through the organization's AAA server. Because the secure VSAN has its own port security, device authentication, independent zoning database, and other security features, it can be physically consolidated with the primary SAN infrastructure without sacrificing management security or traffic isolation.

Figure 6. Achieving Security More Economically with a VSAN

Blade Server Deployments: Using VSANs to Provide Different Levels of SAN Services per Blade Server

Blade servers come in a specific form factor in a compact chassis designed for assembling a large number of servers in a given rack space. To provide access to the storage fabric, the blade server chassis hosts one or two blade server switches. These are fabric switches specially designed for the blade server environment. The blade switch has interfaces internal to the chassis that are hardwired to each individual blade server, and external interfaces to connect to the core fabric.

The Cisco MDS 9000 Family blade switches can operate in two modes:

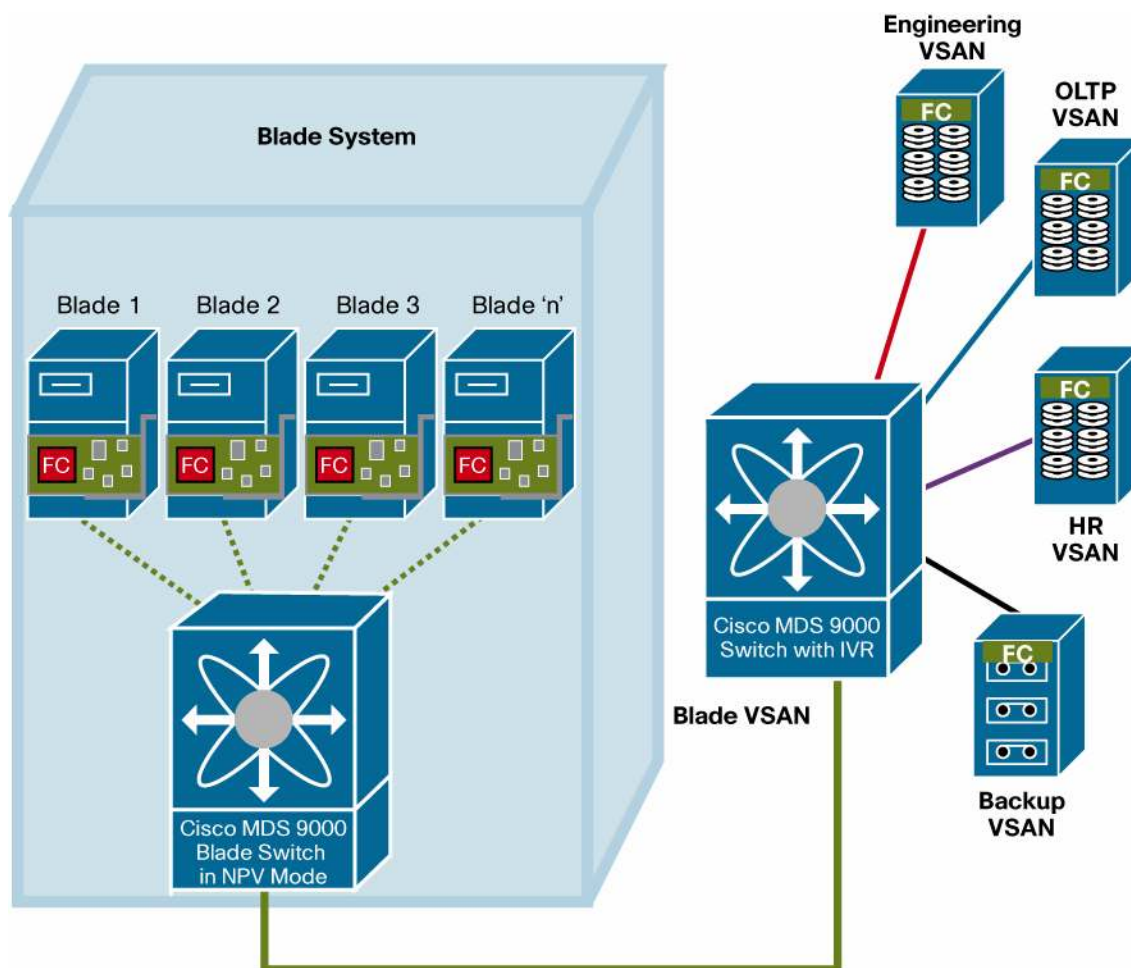
- **Switch mode:** The blade switch connects to the core fabric as a fully functional Fibre Channel switch. The interfaces to the core work in E-port mode (ISL mode), and the device must be managed like any other switch in the fabric.
- **Cisco Network Port Virtualizer (NPV) mode:** The blade switch behaves like a regular fabric port for each server, but it looks like a set of host bus adapters (HBAs) to the core network. Each server blade appears to the core network as a virtual HBA, with the Network Port Identifier Virtualization (NPIV) Fibre Channel standard used to share the fabric-facing ports of the blade switch.

Working in switch mode, the blade switch is no different than any other Fibre Channel switch; only the form factor is different. Each blade can be assigned to a different VSAN simply by setting the properties of the interface to which it is connected. This option is not commonly deployed, because any data center that deploys blade servers on a large scale can end up managing a large number of small switches.

The NPV mode is very useful in reducing the management overhead for each blade switch. The SAN fabric size is dramatically reduced since the switch does not take part in the fabric services, but acts like an aggregator of the HBAs.

Each blade can be individually routed to a specific target VSAN on the core switch by using IVR. The core switch dynamically routes traffic from each blade to a specific target group VSAN as needed. This approach enables the sharing of blades between different user groups in the same blade enclosure (Figure 7).

Figure 7. Using VSANs with Blade Switches in NPV Mode




Virtual Machine Deployments: Extending VSANs to Virtual Machines and Assigning Different Levels of SAN Services on a per-Virtual Machine Basis

In an enterprise-class virtual machine deployment, different virtual machines on a given physical server can provide a variety of services (QoS, security, RBAC, etc.), and they may need to access storage on different VSANs. If all the virtual machines share the identity of the physical HBA, virtual machine access can be provided only by configuring IVR to give the physical machine access to all the VSANs, a solution that also negates some of the isolation benefits of VSANs.

If NPIV is implemented, a physical HBA can support multiple identities to be mapped to virtual HBAs. The physical HBAs and all the virtual HBAs still belong to the VSAN assigned to the physical ingress port of the switch, but assigning a virtual HBA to a virtual machine enables selective IVR routing of virtual machines to the desired VSAN.

This function, known as nested NPIV, is available in the Cisco MDS 9000 Family blade or fabric switch in NPV mode. Each virtual HBA is recognized by the core switch as a unique entity and can be routed to the appropriate target using IVR (Figure 8).

Figure 8. Nested NPIV and IVR for Virtual Machines

By deploying NPIV, the VSAN-based fault isolation and segmentation can be propagated back to the individual virtual machine level. Additionally, there is no difference in the capabilities available to a virtual machine running on a blade, a virtual machine running on a blade server, or a physical application server connected directly to a switch.

Conclusion

Cisco VSANs implemented with the Cisco MDS 9000 Family of Fibre Channel switches and directors are fully compliant with the ANSI T.11 standard requirements for hardware isolation of fabric segments and provide each segment with independent fabric services (such as zoning and routing), QoS, and role-based management.

The high degree of security and flexibility offered by VSANs enables a wide variety of solutions, and VSANs constitute a fundamental building block for the end-to-end virtualized data center of the future.

Applications include consolidation of SAN islands traditionally dedicated to specialized services such as backup and data replication, sharing of the MAN and WAN interconnection infrastructure between data centers, deployment of multiple-vendor SANs requiring multiple Fibre Channel protocol modes, and compliance with data security standards requirements.

VSANs are well suited for the flexibility and granularity offered by server virtualization technologies, providing the framework to fully virtualize both the processing power and the storage networking of the third-generation data center.

For More Information

For additional information on Cisco MDS storage solutions, please visit:

- <http://www.cisco.com/go/storage>
- <http://www.cisco.com/go/whymds>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flip Video, Flip Video (Design), Flipshare (Design), Flip Ultra, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Store, and Flip Gift Card are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CQVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0907R)