Guide

Director-Class FCoE: Converge the Network with Cisco Nexus 7000 Series Switches and Cisco MDS 9500 Series Multilayer Directors

What You Will Learn

This document is intended for storage administrators and sales engineers who want to understand Cisco's director-class Fibre Channel over Ethernet (FCoE) offerings. It explores basic requirements, deployment scenarios, and best practices for a successful deployment of multi-hop FCoE.

FCoE Overview

Data centers are typically designed for multiple networks that have different requirements for latency and resiliency and so deploy a separate, loss-intolerant network dedicated to storage (SAN) and another network dedicated to traditional loss-tolerant Ethernet (LAN). Managing multiple networks presents administrators with a challenging and inflexible environment, with redundant network adapters, cabling, and switches.

FCoE is a standards-based protocol that natively maps Fibre Channel to Ethernet for transport in a lossless Ethernet LAN. FCoE allows I/O consolidation in the data center, which addresses many of the challenges faced in today's data center. For a more detailed introduction to FCoE, please consult Cisco FCoE Fundamentals.

FCoE Terminology

This document uses the following FCoE terminology:

- **Data Center Bridging (DCB):** A collection of standards that extend classical Ethernet protocols for use in the data center
- **DCB Exchange (DCBX) Protocol:** Capabilities through which peers can exchange DCB parameters to help ensure lossless Ethernet behavior
- **Enhanced transmission selection (ETS):** A feature defined in the IEEE 802.1Qaz standard; enables a percentage of available bandwidth on a link to be divided among specified priority groups
- **Fibre Channel forwarder (FCF):** The function in an FCoE-capable device that allows FCoE frames to be switched across multiple hops
- **FCoE Initialization Protocol (FIP):** An initialization protocol that enables establishment of point-to-point virtual Fibre Channel links over a multi-access network such as Ethernet
- **Fibre Channel:** The dominant protocol for use in the SAN; its deterministic and lossless reliability make it well suited for storage traffic
- **SAN:** A network designated for storage traffic
- **Virtual SAN (VSAN):** A grouping of logical SANs residing on the same physical network

Cisco FCoE Requirements

Software Requirements

F1 FCoE capabilities on Cisco Nexus® 7000 Series Switches and Cisco® MDS 9500 Series Multilayer Directors are supported in Cisco NX-OS Software Release 5.2 and later. FCoE capabilities on F2 will be supported in Edinburgh-2 (NX-OS 6.1) release.

Hardware Requirements

Hardware requirements vary based on the platform being used for FCoE connectivity. The platform-specific requirements to build the multi-hop topologies discussed in this document are outlined in Table 1. Note that a multi-hop FCoE topology also requires converged network adaptors (CNAs) for server connectivity; currently, CNAs are provided by EMC, QLogic, and the Cisco Unified Computing System™.

Table 1. Cisco FCoE Hardware Requirements

Platform	Part Number	Description
Cisco MDS 9500 Series	DS-X9708-K9	Cisco MDS 9000 10-Gbps, 8-Port FCoE Module
Cisco MDS 9500 Series	DS-X9530-SF2AK9	Cisco MDS 9500 Series Supervisor-2A Module
Cisco Nexus 7000 Series	N7K-F132XP-15	32-port 1 and 10 Gigabit Ethernet module (Small Form-Factor Pluggable [SFP] and Enhanced SFP [SFP+])
Cisco Nexus 7000 Series	N7K-F248XP-25	48-Port 1 and 10 Gigabit Ethernet F2-Series Module (Small Form-Factor Pluggable [SFP] and Enhanced SFP [SFP+])

Licensing Requirements

To enable FCoE capability on a Cisco Nexus 7000 Series line card, an FCoE license is required. One license is required for each line card (Table 2).

Table 2. Cisco FCoE Licensing Requirements

Platform	Part Number	Description
Cisco Nexus 7000 Series	N7K-FCOEF132XP	FCoE License for Cisco Nexus 7000 Series 32-port 10 Gigabit Ethernet SFP+ (F1)
Cisco Nexus 7000 Series	N7K-FCOEF248XP	FCoE License for Cisco Nexus 7000 Series 48-port 10 Gigabit Ethernet SFP+ (F2)

Storage High Availability

High availability is a crucial requirement in any data center design. It is especially important in a converged network, an environment in which loss-intolerant storage traffic shares the network infrastructure with loss-tolerant Ethernet traffic. High availability is implemented at many different levels of the network: chassis level, link level, and process level. Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches have hardware and software features that help network designers meet the high-availability requirements of their data centers and secure the uptime and integrity of the storage traffic in a consolidated environment. For more information about these features, see Fibre Channel over Ethernet Fault Tolerance: Achieve High Availability in a Converged Network.

Traditionally, SAN designers have used redundant network links and network equipment to create high availability in their networks. These designs provision and manage parallel, disparate SANs. These parallel, independent SANs are often referred to as SAN A and SAN B.

Data center designers need to consider high-availability options when designing a consolidated network.

Storage Traffic Security

Security is another important factor to consider when designing a data center network. This is not just about securing the data traffic. It includes securing access to network resources and establishing roles and managing these roles. Security for storage traffic has many elements. The management plane needs to be secured by preventing unauthorized access to storage-specific resources; this can be achieved using zoning and VSAN separation. In a converged network, in which storage and LAN traffic share the same infrastructure, the storage traffic also needs to be protected against anomalies in the Ethernet network; traffic patterns caused by Ethernet broadcast storms should be quickly mitigated and segregated so that they do not cause loss or delay for intolerant storage traffic. In addition, the storage control plane needs to be protected from unauthorized modifications that could introduce instability into the network.

As increasing numbers of devices are consolidated in the data center, the management and administration of these devices is also being consolidated. Consolidation means fewer physical devices in the network performing multiple functions and needing multiple groups of administrators to manage them. In such scenarios more administrators become responsible for fewer devices in the network. Therefore, defining role-based policy for access to these network resources is an important security function of the operating system.

The Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches provide the hardware and software features to address these and other security concerns of data center administrators. For more information about these features, see Fibre Channel over Ethernet Fault Tolerance: Achieve High Availability in a Converged Network.

Oversubscription

Oversubscription is the practice of having multiple devices share a common resource. In networks, bandwidth is commonly oversubscribed by using a set of end devices that have a greater aggregate bandwidth than the links and end devices through which they are connecting. This setup is often described in terms of an oversubscription ratio, which is the ratio between the bandwidth available to one set of ports, devices, and links and another set of ports, devices, and links. In a SAN, the most common ratio is host bandwidth to storage bandwidth. The amount of oversubscription on a link depends on the I/O capabilities of the hosts and targets and also the applications using them. Designs with more highly oversubscribed links may be satisfactory for applications with light I/O requirements, and conversely links with lower oversubscription ratios may be required for applications that have greater I/O requirements.

Oversubscribing the network allows higher port-count deployments without compromising performance, because the performance requirements of the application may be far less than that of the negotiated port speed of the switch or host. Network designers should keep their applications' requirements in mind when designing end-to-end FCoE networks. Figure 1 shows a typical network design with an end-to-end host-to-storage oversubscription ratio of 10:1.

Figure 1. Oversubscription in the SAN

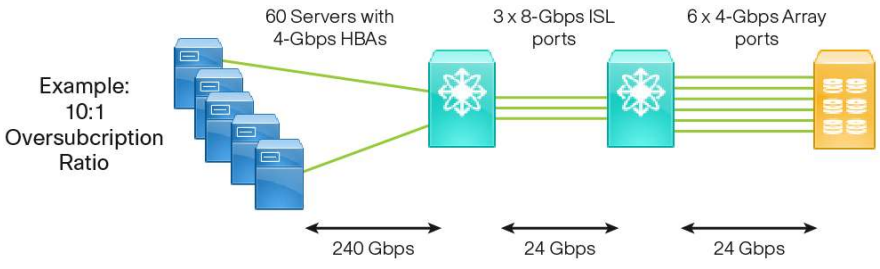

Network Topologies and Deployment Scenarios

In a typical data center, a server is connected to two different and independent networks: a LAN to provide IP connectivity, and a Fibre Channel SAN to provide storage connectivity. Figure 2 shows such a data center network, with a host connected to the LAN and the Fibre Channel SANs in a redundant fashion. This network design has scalability limitations because of the number of adapters used at the host and the amount of dedicated data switching equipment used in the network.

Figure 2. Traditional Data Center Network

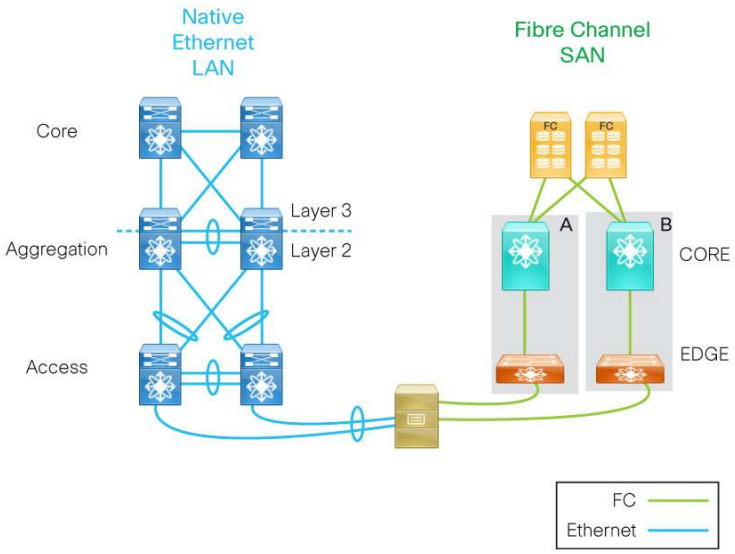

FCoE-capable line cards available on the Cisco Nexus 7000 Series and Cisco MDS 9500 Series director switches allow the creation of different topologies to meet any requirement. These FCoE-capable devices enable reduction in the number of protocol-specific host adapters and discrete switching platforms deployed in the data center. Some of the most common end-to-end FCoE deployments and the value they offer network administrators and operators are discussed in the following sections.

Converging the Access Layer

The first logical step in reducing the number of host adapters and discrete switching platforms is to converge the access-layer LAN and SAN switches into a single set of host adapters and single switching platform. Access-layer convergence significantly reduces capital expenditures (CapEx) and operating expenses (OpEx), and extending convergence beyond the access layer increases those benefits. Figure 3 shows a converged access layer. By deploying FCoE-capable line cards in both the Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches, the need for dedicated Fibre Channel hardware is reduced, with such hardware required only in the Fibre Channel SAN's core. This approach largely benefits IT administrators as they manage to reap the benefits of convergence while using the same management tools that they use on the LAN to monitor their mission-critical applications in the SAN.

Figure 3. Converged Access Network

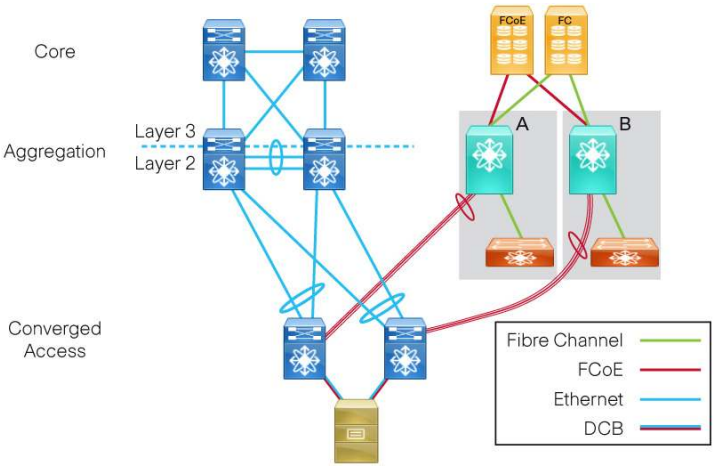

The capability to bridge converged access designs with Fibre Channel SANs preserves existing and continued investments made in the Fibre Channel SAN. This design also allows the IT operators to start their migration to convergence with director classes of service (CoSs) offered by these switches (such as high availability, high port density, and intelligent SAN services). A phased approach to convergence enables data center operators to enhance their deployments as more advanced FCoE storage solutions become available and at the same time move toward a converged network that benefits from the economics of Ethernet.

Figure 4 shows a sample core-edge design using the Cisco Nexus 7000 18-Slot Switch and Cisco MDS 9513 Multilayer Director to achieve an overall storage oversubscription ratio of 7:1.

Figure 4. Sample Multi-hop FCoE Core-Edge Design

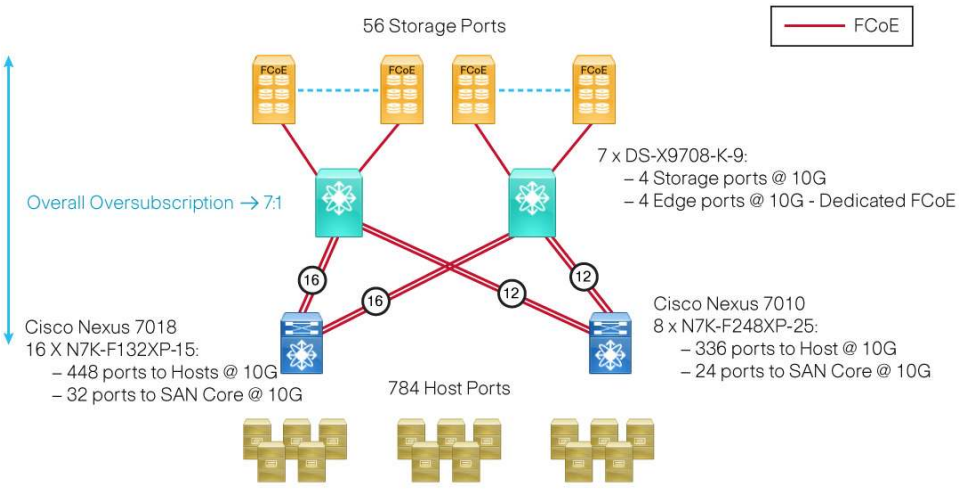

Notes:
1. Classical Ethernet LAN connections not shown.
2. This sample is only SAN A (or B) portion of the network.
3. Shared Links to the hosts allocate 50% of bandwidth to FCoE traffic.
4. From a LAN network perspective, Cisco Nexus 7010 and 7018 switches are access switches.

Converging the Aggregation Layer

Cisco's multi-hop FCoE solution allows the use of a Cisco Nexus 7000 Series Switch as a converged Ethernet aggregation and SAN core switch. This switch can then be connected to a Cisco MDS 9500 Series switch in a three-tier edge-core-edge SAN design. Figure 5 shows a typical converged network design. Here FCoE traffic traverses a dedicated link from the access layer to the aggregation switches and then on to the edge SAN network. Use of dedicated links to connect the access layer to the aggregation layer in this scenario provides a separation for the storage traffic and isolates each fabric, which is a requirement for many SAN operators.

Figure 5. Converged Network Design

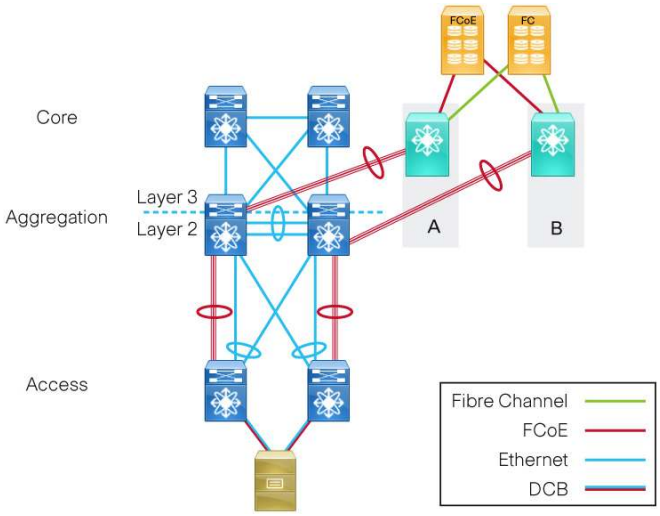

Additionally, use if the Cisco Nexus 7000 Series at the access layer allows designers and administrators to deploy Cisco virtual PortChannel (vPC) or FabricPath technology on LAN ports in the access layer, thereby enabling higher port utilization between the aggregation and access layers and increasing access bandwidth to the servers for IP storage (Small Computer System Interface over IP [iSCSI] or network-attached storage [NAS]) and LAN traffic while keeping the FCoE storage traffic segregated on separate links. This approach allows all Ethernet features to be applied to the IP storage and LAN traffic while preserving the traditional three-tier edge-core-edge SAN network design.

Edge-core-edge topologies are well-suited for environments in which the network is predicted to grow so that the number of storage ports exceeds the number of ports available at the core. Figure 6 shows a sample design.

Figure 6. Sample Multi-hop FCoE Edge-Core-Edge Design

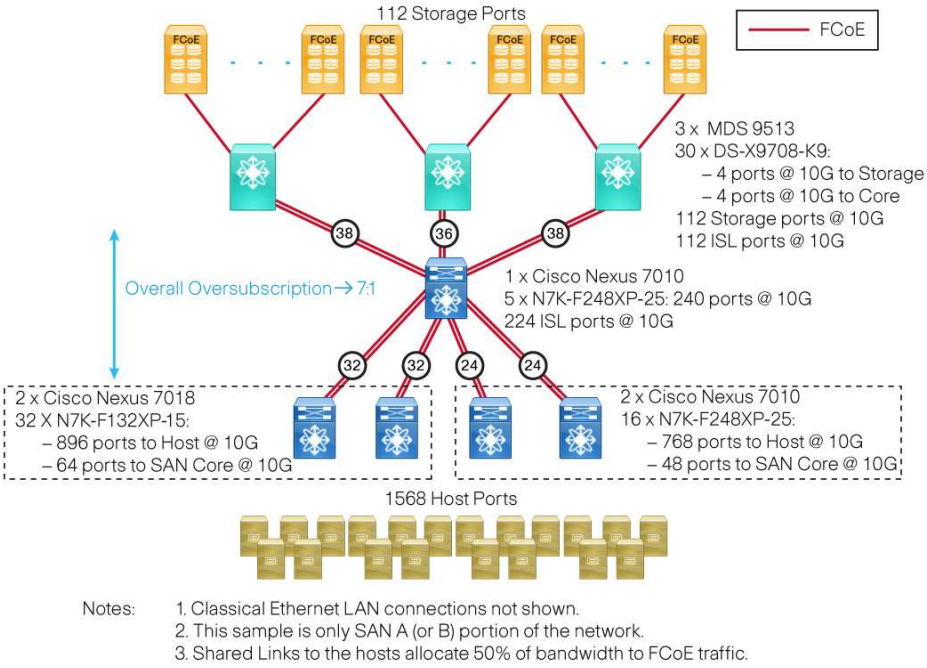

Providing an End-to-End Ethernet Solution

Another option for deploying a SAN is to create an end-to-end FCoE-based network using the Cisco Nexus 7000 Series as the core building block. As in the previous examples, a segregated Ethernet SAN is an attractive option particularly in environments in which separation of LAN and SAN traffic is desired or strict organizational boundaries exist. This approach, shown in Figure 7, allows SAN administrators to take advantage of Ethernet's scale and economic model and to grow as higher-capacity Ethernet speeds become available.

Figure 7. Converged Access with Dedicated Ethernet SAN Core

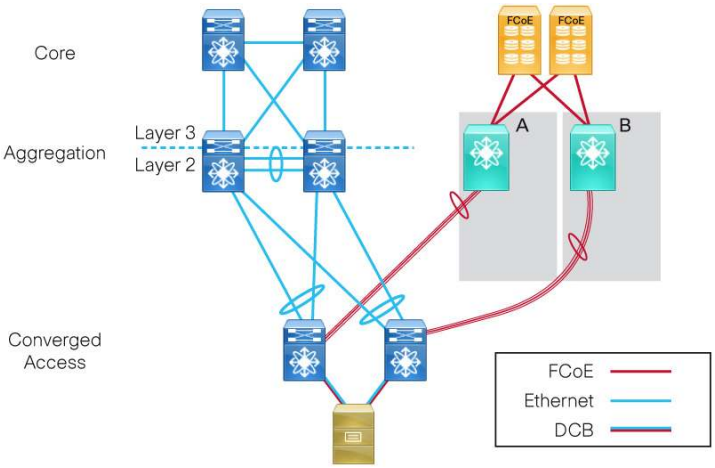

The port density offered by the Cisco Nexus 7000 Series Switches also enables the LAN and SAN designers to directly attach FCoE storage arrays to Cisco Nexus 7000 Series aggregation switches in their networks (Figure 8). SAN separation beyond the access layer still is achieved using dedicated FCoE links. The capability of the network designer to directly lay the FCoE network over the existing LAN increases the number of devices that can be enabled to access FCoE storage while maintaining the SAN isolation that may be required, through the use of virtual device contexts (VDCs) and dedicated Inter-Switch Links (ISLs). Directly attaching storage to aggregation switches creates a core-edge SAN topology and allows storage operators to take advantage of director-class services available on these core switches such as high availability, security, and port scalability.

Figure 8. Converged Network Design with Storage Attached

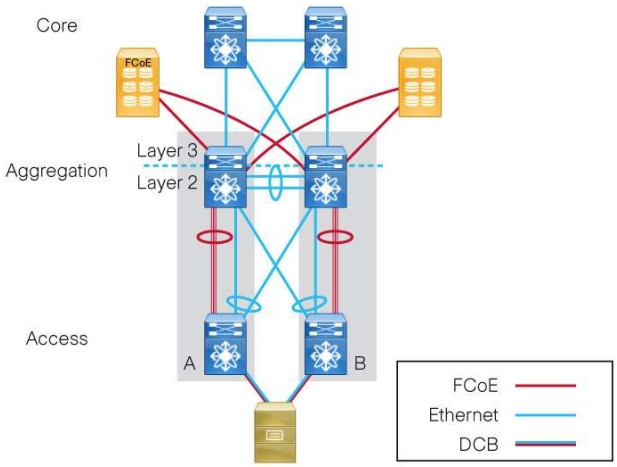

Figure 9 shows a sample core-edge network for this converged network. This figure shows how scalability similar to that of a Cisco MDS 9500 Series-based edge-core-edge design (as shown in Figure 6) can be achieved when using the Cisco Nexus 7000 18- Slot Switch as a scalable Ethernet storage director.

Figure 9. Sample Multi-Hop FCoE Edge-Core Design with Storage Attached

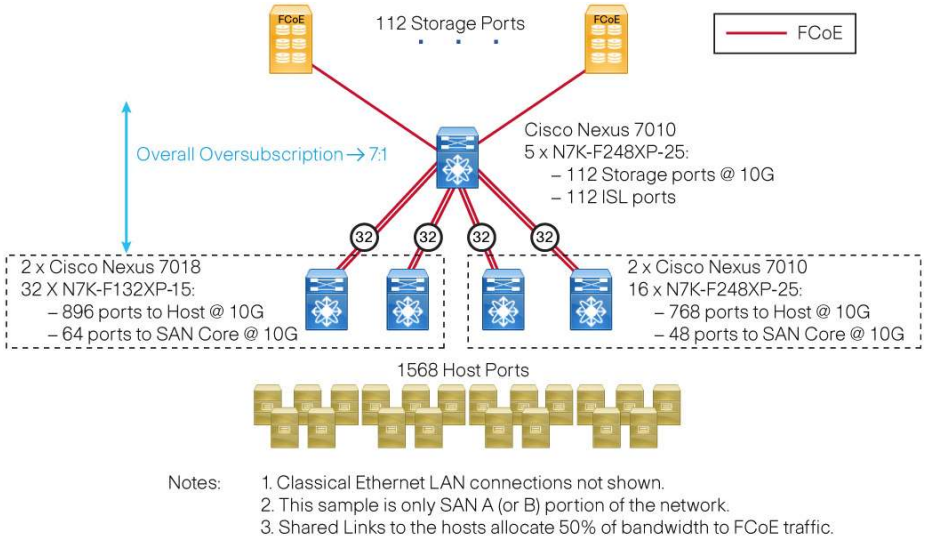

Converged Network Considerations

When designing an end-to-end FCoE-enabled network, designers need to consider network and end-device requirements: in particular, the number of VSANs to be deployed, zoning, and the scale of the fabric logins required. The Cisco Nexus 7000 Series and Cisco MDS 9500 Series products share the same software that provides these capabilities.

VSANs

VSAN technology enables a physical storage fabric to be divided into distinct logical entities in which traffic in each VSAN is segregated. This segregation improves network security and management because VSANs can be provisioned without the need for a physical change to the network. There are multiple use cases for VSANs: for example, allocation of a specific VSAN for a particular type of operating system, or separation of traffic based on business or department function.

In converged environments, on Cisco Nexus 7000 Series Switches, Ethernet VLANs used to carry FCoE traffic can be directly mapped to Fibre Channel VSANs on Cisco MDS 9500 Series switches. This mapping allows classical Ethernet and FCoE traffic to share the same ports and yet be segregated logically. Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches can support up to 4000 VSANs in a physical fabric.

Zoning

Zoning complements VSANs in securing the network. Zoning provides the means to restrict access and visibility between devices using the same FCoE VLAN or Fibre Channel VSAN but in different functional areas. Zoning defines which targets can be accessed by which hosts if they are all logged into the same FCoE VLAN or Fibre Channel VSAN. Zoning also limits state-change notifications to devices in the same zone, thereby reducing the amount of control traffic and unnecessary disruption to mission-critical devices in the network. In all cases, Cisco switches enforce these zones using software-based methods as well as enforcing them in hardware.

In a converged environment in which the transit ports are Ethernet, zoning is configured and is applied only to the FCoE traffic. With each FCoE VLAN having its own zone sets and zones, all the perceived benefits of zoning in a traditional Fibre Channel SAN are also available in a converged FCoE environment. Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches support up to 8000 zones and 20,000 zone members in a physical fabric.

Fabric Logins

Every end node in the Fibre Channel network (servers and storage ports mainly) performs a login procedure every time it joins the network. The main purpose of these logins is the establishment of service parameters and assignment of port identifiers.

A major consideration that places a limitation on design and deployment is the number of fabric logins that need to be supported. The number of fabric logins that a device in the SAN can support is orthogonal to the number of VSANs, port scale, and zoning. Important considerations here are the supported number of logins per port, line card, and switch. These limitations affect the design and future growth of the SAN whether it is based on FCoE or Fibre Channel. Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches support up to 10,000 fabric logins in a physical fabric regardless of the number of VSANs deployed in the network.

Table 3 summarizes the limitations in Cisco NX-OS Software Release 5.2. (These limitations may vary by release; see the relevant data sheet for specific limitations for a release.)

Table 3. Fabric Scalability Limits: Cisco NX-OS 5.2

Limits per Physical Fabric	Cisco Validated Limits	Cisco Maximum Limits
Number of VSANs	80 per fabric	4000 per fabric
Number of fabric logins	10,000	10,000
Number of fabric logins per switch	2000	2000
Number of fabric logins per line card	400	400
Number of fabric logins per port	256	256
Zone members	16,000 per fabric (includes all VSANs)	20,000 per fabric (includes all VSANs)
Zones	8000 per switch (includes all VSANs)	8000 per switch (includes all VSANs)
Zone sets	500 per switch (includes all VSANs)	1000 per switch (includes all VSANs)
Number of FCoE hops	7	7

Other Design Considerations

- In the preceding scenarios, the Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches deployed will be acting as FCFs. The FCF in the multiprotocol switch or router is the feature that enables FCoE forwarding: the basic function of encapsulating and de-encapsulating FCoE frames.
- A virtual E-port (VE-port) is a port that emulates an E-port over an Ethernet link. In the design scenarios depicted, a VE-port on an FCF can be connected to another FCF's VE-port to form an ISL between the two FCFs.
- VDCs, available on the Cisco Nexus 7000 Series Switches, allow virtualization at the switch level, where logical entities are created to provide process separation and fault tolerance. The dedicated FCoE links can be configured so that ingress traffic is processed in a separate and distinct VDC, the storage VDC. For more information about storage VDCs on the Cisco Nexus 7000 Series, refer to Virtual Device Contexts on Cisco NX-OS.
- Up to 16 dedicated FCoE links can be placed in a PortChannel to increase the aggregated bandwidth available between switches. These links can be in a PortChannel regardless of their location in the chassis. This approach is very attractive from a high-availability standpoint and allows administrators to further segregate their storage traffic if they want to do so.

Conclusion

As data centers continue to grow, data center architects are increasingly challenged to design networks that not only meet current needs but have the potential to expand to meet future requirements. Converging disparate data center networks into one ubiquitous infrastructure both reduces CapEx and OpEx and allows the data center to grow with business requirements. Cisco Nexus 7000 Series and Cisco MDS 9500 Series switches with FCoE capabilities enable this convergence to an Ethernet infrastructure to take place. These switches extend FCoE from the host all the way to the existing Fibre Channel-based SAN, allowing data center operators to take advantage of already installed storage devices and thereby reducing the total cost of ownership TCO). The port density offered by these switches also positions them well for any future growth in the network. Architects who want to design such data centers can use the guidelines in this document to design a flexible and scalable data center network.

For More Information

http://www.cisco.com/go/unfiedfabric

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.

Printed in USA C07-648629-01 07/12