Fibre Channel over Ethernet Management: Secure Storage Traffic in a Converged Network

What You Will Learn

As increasing numbers of data center designers and operators consider consolidating their I/O operations, legitimate concerns about the manageability and security of their storage traffic arise. The networking equipment that enables this consolidation must provide the tools necessary to segregate, manage, and secure the storage traffic.

To guard against intruders and malicious attacks, an additional layer of security is essential. Security features integrated into switches must be able first to authenticate the users and then allow them access to data center network resources based on predefined policies.

This document discusses two main types of network security:

- Control-plane protection: Access protection for the switches
- Data-plane protection: Security for traffic passing through the switches

Both types of protection are important in securing converged storage traffic.

Fibre Channel over Ethernet (FCoE) is a standards-based protocol that natively maps Fibre Channel to Ethernet for transport in a lossless Ethernet LAN. FCoE allows the consolidation of LAN and Fibre Channel SAN traffic over a single switching infrastructure in the data center.

This document describes the features of the Cisco Nexus® 7000 Series Switches and Cisco® MDS 9500 Series Multilayer Directors that help operators and data center designers successfully secure storage traffic in a converged network.

Authentication

The Cisco Nexus 7000 Series and Cisco MDS 9500 Series provide two types of authentication:

- Authentication, authorization, and accounting (AAA): AAA network security services provide the basic secure authentication framework that allows network operators to control access to the switches in the data center. These services enable administrators to identify users through secure logins, provide authorization for remote access control, and collect statistics on these accesses for billing and auditing services. Cisco NX-OS Software, the Operating System that runs the Cisco Nexus and Cisco MDS 9000 Family switches, provides AAA services for switches used to build the converged network. These services function in combination with role-based access control to give administrators access control regardless of the composition of the network infrastructure in the deployment.
- Role-based access control (RBAC): Cisco NX-OS enables administrators to restrict and control access to a certain group of users. Users, defined by roles, are permitted or denied access to resources in a network, thereby enhancing the security of the network. This feature allows administrators to be in charge of a specific virtual SAN (VSAN) without having any access or visibility to other VSANs. In a converged network, storage network operators could define independent sets of policies, creating much stricter access policies for storage-related infrastructure and resources, and separate, distinct access policies for Ethernet-related resources.

Security

Securing access to storage network devices and storage traffic is a principal concern for data center network operators. Security is especially critical in an environment in which consolidated infrastructure carries multiprotocol traffic. Cisco NX-OS provides advanced security features that address some of these concerns for the storage administrators:

- Cisco TrustSec® service: This flexible, policy-based service provides access control that spans multiple devices and operating systems to bring a holistic approach to security that is not bounded by physical device attributes. It allows multiple security policies to be merged, thereby providing tighter security for the system as a whole. For a more detailed description of how the Cisco TrustSec service operates on Cisco Nexus 7000 Series Switches, refer to Cisco TrustSec Solution Overview. The Cisco TrustSec service uses security group tags to create role-aware networks, with the role information available at each point in the network. Figure 1 shows how the link between switches can be authenticated, with individuals having access to certain security groups and in turn getting access to specific resources in the network. The Cisco TrustSec service enables an end-to-end security policy to be applied to the data center network.
- In addition to TrustSec®, Fibre Channel Security Protocol (FCSP) provides switch to switch and switch to host authentication. Available on the Cisco NX-OS software, FCSP enables the authentication of Cisco Nexus 7000 switches and Cisco MDS 9500 directors in the same network and prevents switches with malicious intentions to be part of the network.

Figure 1. Authenticating Users on a Link Between Nexus 7000 Series Switches

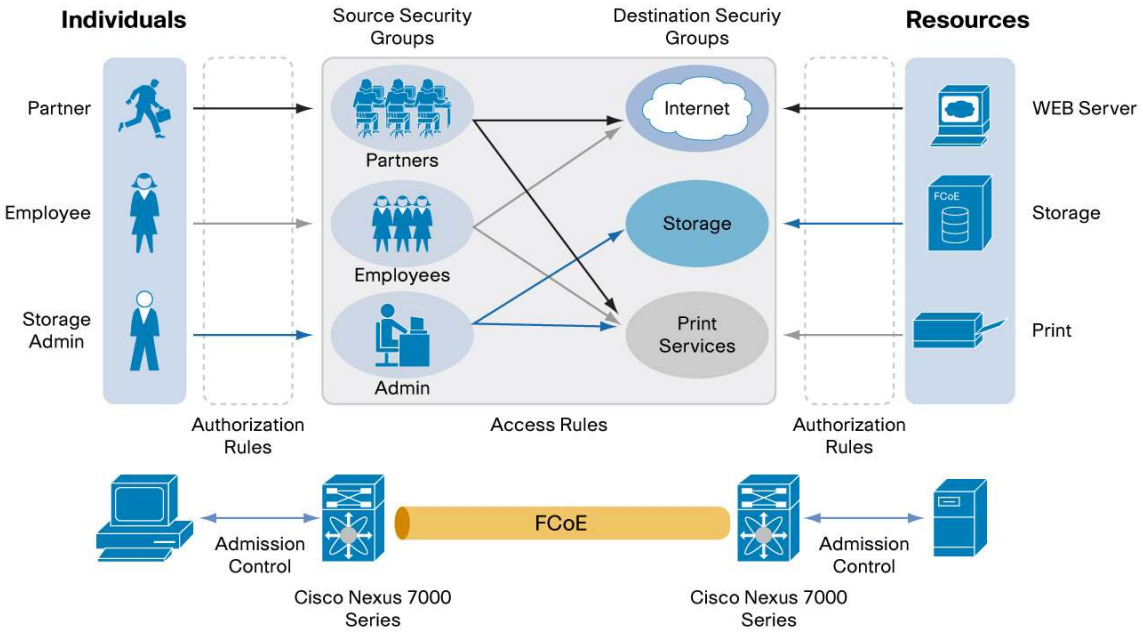

- Dynamic Host Configuration Protocol (DHCP) snooping: This security feature filters out untrusted DHCP requests that could cause the switch to peer with malicious clients on the network. DHCP snooping can be enabled on the switch as a whole or on individual VLANs.
- IP source guard: This feature helps prevent IP spoofing attacks, in which a host tries to spoof and use the IP address of another host. Malicious clients could attempt to attack the switch using a legitimate client's IP address.

- Access control lists (ACLs): ACLs are an ordered set of rules that enable the operator to filter traffic based on certain attributes. Actions then can be performed on this filtered traffic. ACLs are favored by administrators because their flexibility allows administrators to perform complicated security functions on the switch and choose the traffic processed by the switch.
- Bridge Port Data Unit (BPDU) guard: This feature prevents the occurrence of an unnecessary Spanning Tree Protocol recalculation. When a switch joins a bridge domain, it can cause the Spanning Tree Protocol to run and determine a new root. This event could be a denial-of-service (DoS) attack on the network, in which a low-priority switch is added and subsequently removed causing a permanent Spanning Tree Protocol recalculation to occur. BPDU guard configured on an ingress port protects the network against such attacks.

All these features of Cisco NX-OS increase the stability and security of switches that have consolidated I/O traffic traversing them. In a converged network, in which storage traffic is also sharing the same infrastructure, these features help ensure that the switch contributes to the security of storage traffic traversing the shared infrastructure. The capability to prevent malicious attacks and contain threats makes Cisco Nexus 7000 Series Switches running Cisco NX-OS attractive choices for consolidated networks in which the security of storage traffic is of vital importance.

Protection Against Broadcast Storms

Storage-related traffic needs to be protected from anomalies in native Ethernet traffic behavior if FCoE is to become a viable option for SANs. Even though applications may be able to function with lost storage traffic, this loss creates inefficiencies in the network (for instance, retransmissions) and creates application interruptions (through timeouts). SAN designers and operators must help ensure allocation of adequate input buffering when deploying an FCoE solution, because shortfalls in such resources can affect switch performance and use in the SAN. In particular, in the event of Ethernet broadcast storms, if the buffer design is inadequate, buffers may become exhausted and no longer capable of accepting storage traffic.

Each platform addresses this situation differently. On the Cisco MDS 9500 Series directors, Ethernet traffic ingress to the switch is dropped at the port because the switches cannot handle classic Ethernet packets (only FCoE traffic is accepted and processed). On the Cisco Nexus 7000 Series Switches, however, the mix of traffic (both classic Ethernet and FCoE) is efficiently processed through an intelligent buffering scheme that is designed to meet the service requirements of each traffic type (for example, achieving the lossless behavior needed for FCoE traffic). This buffering capability is facilitated by priority-based flow control (PFC) and buffering and storage protection.

Priority-Based Flow Control

In a Fibre Channel network, a credit-based flow-control (buffer-to-buffer [B2B]) mechanism is used to control the flow and provide lossless behavior on the Fibre Channel link. The receiver gives credit to the transmitter, specifying the amount of buffering available for reception of Fibre Channel frames. The transmitter is therefore bound by the number of credits available to it when transmitting frames. This interaction is summarized in Figure 2. Credit-based flow control helps ensure that the receiver always has enough buffer space to process the incoming frames and avoid traffic loss in a SAN.

Figure 2. Fibre Channel Flow Control: B2B Credit Mechanism

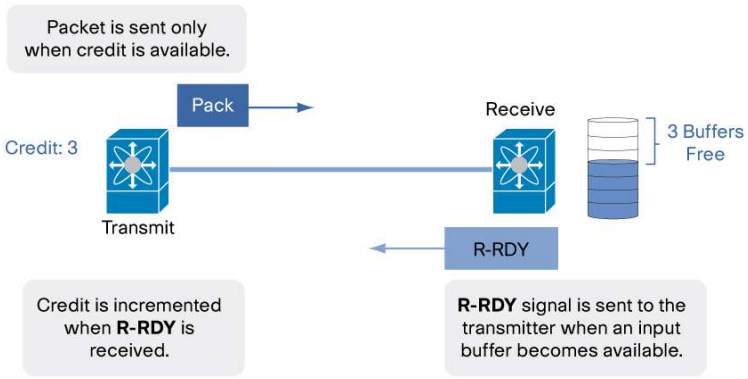

Functionally, PFC is very similar to the Fibre Channel B2B credit system in which the sender keeps track of buffer availability at the receiver's end and will transmit packets only if it knows that buffers are available. A traditionally lossy Ethernet network can be made lossless through the use of a Pause mechanism. The Pause mechanism is an Ethernet frame sent by the receiver to indicate exhaustion of resources. Upon receipt of a Pause frame, the sender will withhold transmission until either a timer has expired (a timeout defined by the receiver) or the receiver allows transmission to proceed. Use of a Pause mechanism enables flow control on the physical link.

PFC, standardized as part of IEEE 802.1Qbb, enables the Ethernet Pause capability on a priority basis. This mechanism provides detailed congestion control on the link, allowing administrators to provide no-drop (lossless) behavior for FCoE traffic that shares the link with Ethernet traffic (for which drops can be gracefully handled by upper-level applications). PFC allows lossless behavior to be assigned to certain priority groups, according to the requirements of the mixed traffic types on a link-by-link basis. This feature essentially enables I/O consolidation on the same infrastructure carrying a mix of traffic with different requirements. A total of eight such priority groups, representing virtual lanes of traffic, can be available in an Ethernet network, as shown in Figure 3.

Figure 3. Priority-Based Flow Control: Eight Virtual Lanes of Traffic

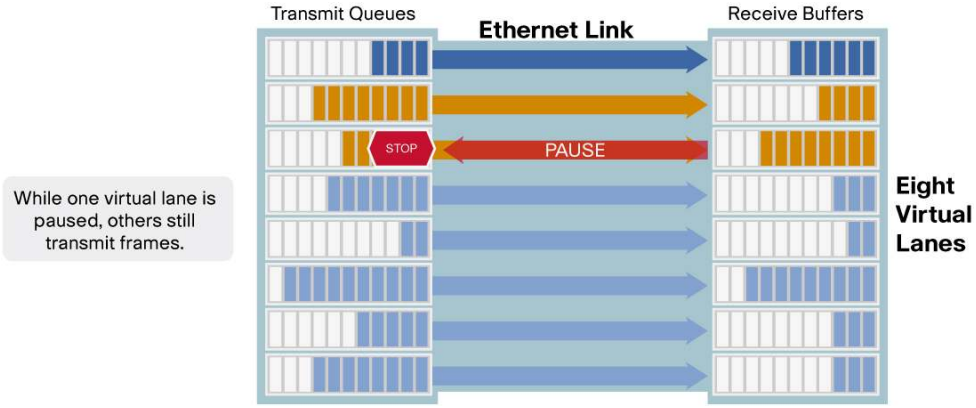

Buffering and Storage Protection

The hardware that supports this lossless behavior on the Cisco Nexus 7000 Series Switches and Cisco MDS 9500 directors has dedicated buffers for each lane of traffic ingress to the port. Each virtual lane is mapped to an output buffer. Traffic is classified based on the priority group to which it belongs and queued for the appropriate output buffer for further processing. This classification into distinct physical buffers allows the traffic behavior to be determined by priority group, as shown in Figure 4.

Figure 4. Output Queues and Their Mappings to Virtual Lanes

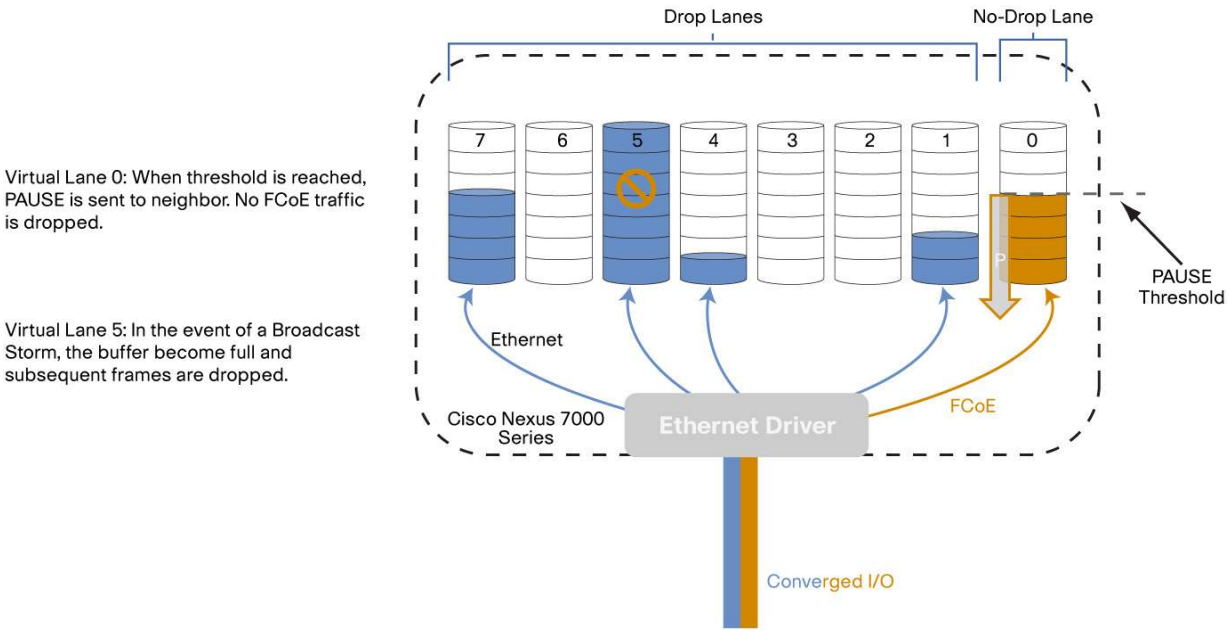

Assignment of an output buffer to a particular traffic type allows a more precise application of the Pause mechanism on that particular traffic stream. In this case, FCoE-related traffic is inserted into a separate queue and achieves lossless behavior through application of the Pause mechanism on that queue. Buffer exhaustion caused by other traffic types (for example, classic Ethernet) will cause frame drops without affecting the FCoE frames waiting to be processed.

Layer 2 networks can be subject to traffic storms because they use broadcast and multicast frames during the discovery phases of other devices in the network. Broadcast and multicast storms are examples of this type of storm, creating frames that may bounce forever between interconnected switches. In a converged network consisting of Cisco Nexus and Cisco MDS 9000 Family switches, accumulation of packets from these storms will cause packet drops at the switch without having any effect on the FCoE traffic that shares the same infrastructure.

Management: Virtual Device Contexts

The virtual device context (VDC) feature of the Cisco Nexus 7000 Series Switches allows the switch to be virtualized into multiple logical devices, each with its own separate processes. Each local entity operates as an independent switch.

Each VDC can contain its own unique and independent set of VLANs and Virtual Route Forwarding (VRF) instances. Multiple interfaces can belong to a single VDC, which allows the forwarding data plane to be independent and virtualized. In addition, each VDC has a separate configuration, routing processes, and process memory. Security and management policies (for example, RBAC and AAA, discussed earlier in this document) can be established on a per-VDC basis.

Traditionally, storage operators manage and administer their Fibre Channel switches separately from the Ethernet part of the data center network. This autonomy creates security and a segregated management plane. On the Cisco Nexus 7000 Series Switches, which are predominantly Ethernet switches, a storage VDC is created to handle only FCoE traffic. This special application of VDCs allows storage traffic to be segregated from the rest of the traffic and monitored accordingly. This model gives enormous power to the storage operators; when operators function in the context of a VDC, all debugging, management, and configuration is applied to the VDC. In a converged environment, storage administrators can be responsible for the storage VDCs, and network operators can be responsible for LAN VDCs, all residing on the same infrastructure.

On a Cisco Nexus 7000 Series Switch, each VDC is isolated from all others. This isolation allows the storage traffic to be unaffected by faults in other parts of the system (or other VDCs). In addition, use of a separate management plane for configuration and debugging enhances the storage network operator's capability to control resources in the network unaffected by provisions applied on the LAN network, despite sharing the same physical resource. Debugging and troubleshooting is also independent on each VDC. From a control-plane perspective, traffic processed in one VDC cannot be routed or switched to another VDC unless a physical, external connection exists between the VDCs. This higher level of segregation, redundancy, and management available on the Cisco Nexus 7000 Series is beneficial to any converged network.

In addition to providing process-level isolation and therefore resiliency, VDCs can be configured so that if a failure occurs, nonstop forwarding (NSF) switchover is performed. In this case, a new control plane is instantiated, and the data plane forwarding is not affected. This mechanism therefore provides switch-level redundancy to help ensure the integrity of the storage traffic.

Overall, VDCs create an effective separation of data traffic and management domains suitable for infrastructure consolidation.

VSANs and Zoning

VSAN is a technology by which a physical storage fabric can 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.

In a converged network, the FCoE-enabled ports can be configured as part of a VSAN, therefore preserving the same capabilities as achieved by Fibre Channel ports.

Zoning complements VSANs in securing the network. Zoning provides the means to restrict access and visibility between devices using the same VSAN but in different functional areas. Zoning defines which storage devices, such as disk arrays, can be accessed by which hosts, who are all logged into the same VSAN. Zoning also limits state-change notifications to devices in the same zone, thereby reducing the amount of noise and unnecessary disruption on mission-critical devices in the network.

In a converged network environment in which the transit ports are Ethernet, zoning is configured in the storage VDC and will be applied only to FCoE traffic. All the benefits that zoning brings in a SAN also apply to a converged FCoE network.

Conclusion

Ethernet-based LANs and Fibre Channel-based SANs may have different I/O requirements because they each are aimed at different sets of applications; however, they also have many similar requirements, such as security for data in transit through the network. A successful deployment of a converged network must address these security requirements.

In a converged network, a separate and distinct management plane gives power to the storage administrators and enables them to provision storage-related resources and control access to critical pieces of infrastructure. As discussed in this document, deploying certain control-plane features increases the resiliency of Ethernet devices and protects them from malicious attacks that could have catastrophic consequences for the SAN.

Cisco Nexus 7000 Series Switches and Cisco MDS 9500 Series directors offer comprehensive sets of features, and their deployment in a converged network helps ensure secure and reliable transmission of storage traffic

For More Information

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

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 C11-647658-00 03/11