

Cisco Nexus 7000 Series Switches: Ethernet Storage Directors

What You Will Learn

Cisco Nexus 7000 Series Switches and Ethernet Storage

Ethernet's simplicity, resilience, and cost effectiveness allow it to complement traditional Fibre Channel networks, thereby making IP storage solutions increasingly popular in enterprise deployments. IP-based storage solutions provide low-cost flexible options when deploying SANs.

New standards are giving data center architects options for designing new data centers and optimizing current installations. These new technologies allow disparate storage protocols to co-exist on the same infrastructure, thereby dramatically increasing the utilization of installed equipment and reducing the need to duplicate infrastructures. Cisco Nexus® 7000 Series Switches, critical components of the data center network, provides a solid foundation for building a converged network. With support for multiple storage protocols, fewer points of management, reduced network size, and a streamlined deployment model, the Cisco Nexus 7000 Series can lead to lower operating expenses.

Challenges of Creating a Successful Data Center Network

There are many challenges to deploying a successful data center. Architects must take into account how to handle not just the existing data that traverses the network, but also how to handle growth over time. Additionally, they are concerned with protecting the data from corruption and unauthorized access, now and in the future.

These concerns lead to several questions:

- How does the administrative team manage all these systems, as well as the information that traverses them?
- How can the data be safely isolated to address concerns about security and resiliency?
- How does a successful data center manage not only the technology (both current and new), but also its security and growth over time?

This document discusses how Cisco® data center innovations such as the Cisco Nexus 7000 Series Switches and Cisco NX-OS Software address these questions and create flexible yet robust multiprotocol storage networks that provide:

- Highly aggregated bandwidth
- Clean, deliberate segregation of storage traffic
- Data security and authentication
- High availability and resiliency of data
- Multiprotocol I/O consolidation for comprehensive platform manageability

As business needs change and additional demands are placed on data centers, today's IT operators and administrators are looking for ways to get extra value from their investments. This document examines each of these features and helps you identify how the Cisco Nexus 7000 Series can provide that value both now and in the future.

Bandwidth Aggregation

Higher available bandwidth and low data latency enhances the performance of IP storage solutions such as Small Computer System Interface over IP (iSCSI) and network-attached storage (NAS). However, the Cisco Nexus 7000


Series offers additional Layer 2 solutions such as Cisco FabricPath and virtual PortChannel (vPC) technology that allow data center designers to build highly scalable IP storage networks.

Traditionally, Layer 2 network designs have required loop avoidance, and the technology that achieves this is Spanning Tree Protocol. With Spanning Tree Protocol, redundant links in a fully meshed network are pruned until only one path exists to each host on the network. Network designers purposely build redundant links to increase bandwidth and resiliency. Spanning Tree Protocol, by creating a single path, limits available bandwidth and prevents operators from deploying scalable and complex network architectures.


Additional features available on the Cisco Nexus 7000 Series enable all links to operate in an active mode, while still providing loop avoidance. These features allow the operator to use all the available aggregate bandwidth and load-share traffic among all the available links. The two features that enable this bandwidth utilization are vPCs and Cisco FabricPath.

Virtual PortChannel is a Cisco Layer 2 technology that addresses the need to use all the bandwidth provided by redundant links between different switching devices. In essence, a vPC bundles a number of active links and represents them as one single link (PortChannel) to the Spanning Tree Protocol. In a traditional PortChannel, the PortChannel can be between only two devices. However, in a vPC, the vPC can span more than two devices as the source or destination of the links. This capability allows a loop-free network while also achieving link-based redundancy since the PortChannel stays active as long as at least one link member is available and active. The Cisco vPC implementation does not put any limitations on the location of the ports belonging to the vPC, allowing them to be spread across multiple line cards. In addition, a vPC can be linked to multiple devices, further spreading the traffic load and adding switch-level (hardware) redundancy (Figure 1).

Figure 1. Greater Port Availability: vPC Compared to Spanning Tree Protocol




Cisco FabricPath is Cisco's innovative Layer 2 multipathing (L2MP) protocol available in Cisco NX-OS. Here a link-state protocol is used to build a routing table of all relevant Layer 2 endpoints at the access layer. Cisco FabricPath uses all the available links when building its Layer 2 routing table, allowing the network designers to create an equal-cost multipath (ECMP) data plane in which traffic can be load-balanced across up to 16 paths. This approach drastically increases the cross-sectional bandwidth available at the access layer and avoids blocking while allowing path resiliency to be built through the Layer 2 domain (Figure 2).

Figure 2. Greater Availability of Links: Cisco FabricPath Compared to Spanning Tree Protocol

Use of vPC or Cisco FabricPath on the Cisco Nexus 7000 Series at the access layer is extremely beneficial for any IP storage deployment because it increases the aggregate network bandwidth available to the servers while providing a flexible yet highly available access network. This type of deployment provides better convergence because the network adjusts itself to changes in link availability. Eliminating Spanning Tree Protocol from the network reduces the number of hops at this layer and creates a more deterministic path to the servers. These innovative technologies available on the Cisco Nexus 7000 Series benefit all IP storage traffic by increasing the efficiency of the data path at the access layer.

Storage Traffic Segregation

Storage network architects have traditionally isolated storage traffic from other LAN traffic to reduce collateral damage and help guarantee throughput for the storage traffic. Historically, this was achieved by dedicating a VLAN to the IP storage traffic. However, on Cisco Nexus 7000 Series Switches, the virtual device context (VDC) feature allows a switch to be virtualized into multiple logical switches, each with its own separate set of processes and runtime configuration. Each VDC operates as an independent switch. Figure 3 shows the VDC model in an IP storage network.

Figure 3. Dedicated VDC for IP Storage Connectivity

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. Each VDC has a separate configuration, routing processes, and process memory. In addition, security and management policies can be established on a per-VDC basis. This switching isolation allows iSCSI and NAS traffic, the two most common forms of IP storage traffic, to be segregated from other LAN traffic and monitored or managed accordingly.

As shown in Figure 3, each VDC is isolated from the others. This isolation allows the storage traffic to be unaffected by faults or changes in other parts of the system (or other VDCs). In addition, a separate management plane for configuration and debugging enhances the operator's capability to control certain resources in the network, unaffected by actions taken on other parts of the network despite sharing the same physical resource. This higher level of redundancy and management capability available on the Cisco Nexus 7000 Series is beneficial to any IP storage network.

Data Security and Authentication

In addition to the management isolation that VDCs provide, the Cisco Nexus 7000 Series platform offers other management and security features that make it particularly favorable for IP storage deployments.

Role-Based Access Control

Cisco NX-OS enables administrators to restrict and control access to a certain group of users. Users, defined by role, are permitted or denied access to network resources, thereby enhancing the security of the network. This feature, called role-based access control (RBAC), allows administrators to be in charge of a specific VLAN without having any access to or visibility into other VLANs. Storage network operators can define independent policies that create much stricter access rules for IP storage infrastructure and resources and provide separate access policies for their Ethernet resources.

Authentication, Authorization, and Accounting

Authentication, authorization, and accounting (AAA) network security services provide the basic secure authentication framework for network operators to control access to the switches in the data center. These services enable administrators to identify and authenticate users through secure login mechanisms. After authentication, users are authorized for remote access control, which determines the network resources to which authenticated users have access. Administrators then can collect statistics on such accesses for billing and auditing services and track what the authorized and authenticated users did while logged into the switch.

Cisco NX-OS on the Cisco Nexus 7000 Series provides AAA services for switches in the data center network. In addition, each VDC has its own set of RBAC and AAA policies independent of other VDCs on the switch. AAA in combination with RBAC allows administrators to bolster network security by limiting access to sensitive storage-related network resources, further securing storage traffic that is traversing the network.

Security

Securing access to storage network devices and storage traffic is of paramount concern to data center network operators. Security is even more important in an environment that has consolidated infrastructure carrying multiprotocol traffic. Cisco NX-OS on the Cisco Nexus 7000 Series provides advanced security features that alleviate some of these concerns for IP storage deployments:

- The Cisco TrustSec® solution provides flexible RBAC service that spans multiple devices and operating systems to bring holistic 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 Cisco TrustSec operations on the Cisco Nexus 7000 Series, refer to [Cisco TrustSec Solution Overview](#). The Cisco TrustSec solution uses security group tags to create role-aware networks in which the role information is available at each point in the network. The link between switches can be authenticated to give individuals access to certain security groups and in turn gain access to specific resources within the

network. The Cisco TrustSec solution enables an end-to-end security policy to be applied to the data center network.

- Dynamic Host Configuration Protocol (DHCP) snooping is a security feature that filters out untrusted DHCP requests, which 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 helps prevent IP spoofing attacks, which occur when a host tries to 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) are an ordered set of rules that enable the operator to filter traffic based on certain attributes. Actions then may be performed on this filtered traffic. The flexibility of ACLs are favored by administrators because ACLs allows them to perform complicated security functions on the switch and choose the traffic processed by the switch.
- Bridge Protocol Data Unit (BPDU) guard is a feature that prevents unnecessary Spanning Tree Protocol recalculations. When a switch joins a bridge domain, it causes the Spanning Tree Protocol to run and identify a new root. This action could indicate 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.

Bandwidth and Congestion Management

Congestion management features enable network administrators to determine the order in which packets are sent from an interface. The order is determined based on the priority given to the traffic type. Bandwidth management features, in contrast, allow a percentage of bandwidth to be guaranteed for certain traffic types. After deployment, these quality-of-service (QoS) features together provide the means by which the performance and integrity of Ethernet storage traffic can be guaranteed, especially during periods of high traffic load.

In a consolidated environment in which a mix of traffic is traversing the same switch infrastructure, important traffic such as storage traffic can be assigned higher priority and guaranteed bandwidth. Enhanced transmission selection (ETS), defined by IEEE 802.1Qaz, is used to assign traffic to a particular priority group, each representing a different class of service (CoS). Traffic belonging to a particular priority group has its bandwidth guaranteed as it is aggregated onto the link.

Layer 2 networks can be prone to traffic storms because they use broadcast and multicast frames during the discovery phases of other devices in the network. For example, in broadcast and multicast storms, frames may bounce forever between interconnected switches. These storms waste Layer 2 bandwidth and may block other traffic in the network, including storage traffic. Traffic storm controls can be enabled on Cisco Nexus Series Switches to suppress broadcast and multicast traffic storms if they reach a certain configured threshold.

The various bandwidth and congestion management tools available on the Cisco Nexus 7000 Series increase the stability of the Layer 2 network and thereby reduce any collateral damage that may affect the IP storage traffic.

High Availability and Resiliency

High availability and network resiliency are important requirements for successful deployment of an IP storage network. Cisco Nexus 7000 Series Switches provide redundancy and high availability in several ways, some of them relevant to IP storage.

Supervisor-Level High Availability

The Cisco Nexus 7000 Series provides a director-class level of high availability. The platform includes dual redundant power supplies, hot-swappable transceivers, and dual supervisors in active-standby mode. When a switch boots up, the active and standby roles are defined for the two available supervisors. The standby switch is in a hot-standby state, monitoring the active supervisor for failures. If a failure occurs, the standby switch will assume the

active role, and the active supervisor is reset. This switchover occurs using nonstop forwarding (NSF), helping ensure the correct forwarding of traffic in transit.

Process restart is another feature of Cisco NX-OS. With this feature, a process failure causes the failing process to restart without triggering a system wide failure.

Cisco NX-OS also offers In-Service Software Upgrade (ISSU), which enables administrators to perform nondisruptive software upgrades. This feature allows upgrades to supervisor and switching modules with little to no negative effect on the data forwarding plane.

Process-Level Availability Through VDCs

Deployment of VDCs on a Cisco Nexus 7000 Series Switch allows the operator to provide process-level high availability in conjunction with any other device-level redundancy offered on the switch. VDCs segment the process space so that failure in one VDC does not affect other VDCs. These logical contexts can be configured to operate in different failure modes according to a predetermined policy, providing additional flexibility for the storage network operator. The actions performed when a failure occurs are defined by the following modes:

- **Bringdown:** The VDC stays in the failure mode. The switch needs to be reloaded for recovery.
- **Reset:** If the switch is occupied by two supervisors in active-standby mode, a supervisor switchover to the standby module is performed. If only a single supervisor is present, that module will be reloaded, thereby affecting all VDCs.
- **Restart:** The failed VDC is deleted and re-created.

In the case of a failure on a switch with dual supervisors, the storage administrators can use the Reset mode and perform a NSF switchover to the standby supervisor. This process will not cause traffic loss, and it will help ensure the integrity of the IP storage traffic traversing the Ethernet network.

Path-Level High availability

Another level of high availability can be achieved at the port level. Deploying vPC technology, administrators can physically connect one switch to multiple other switches yet represent them as one logical connection. This logical connection stays active as long as at least one operational port in the group is still active. Failures on connected switches and line cards that may bring down individual ports are mitigated because the vPC stays active, forwarding traffic. More important, upper-layer protocols see the vPC simply as an active interface and are not affected by failures on the individual ports underlying that logical interface. This feature helps reduce state-change activity at the upper layers and provides a more stable operating environment because neighboring switch failures are isolated.

Cisco FabricPath, with its basis in routing protocols, provides a full mesh of connectivity among multiple switches. Cisco FabricPath allows redundancy to be built in and makes multiple paths available to allow traffic to be forwarded on a different path despite failure.

These software and hardware resiliency features reduce the possibility of network downtime and enhance the profile of the Cisco Nexus 7000 Series as true Ethernet director switches.

Multiprotocol I/O Consolidation

Enabling multiple storage protocols to traverse the same physical infrastructure is important in achieving higher utilization of established networks. The Cisco Nexus 7000 Series empowers data center designers to deploy multiple storage protocols on the same Ethernet-based infrastructure.

In addition to file-based protocols such as Network File System (NFS) and Common Internet File System (CIFS), the Cisco Nexus 7000 Series supports block-based storage protocols such as Fibre Channel over Ethernet (FCoE). FCoE is a protocol that natively encapsulates a Fibre Channel frame into an Ethernet frame for transport across an

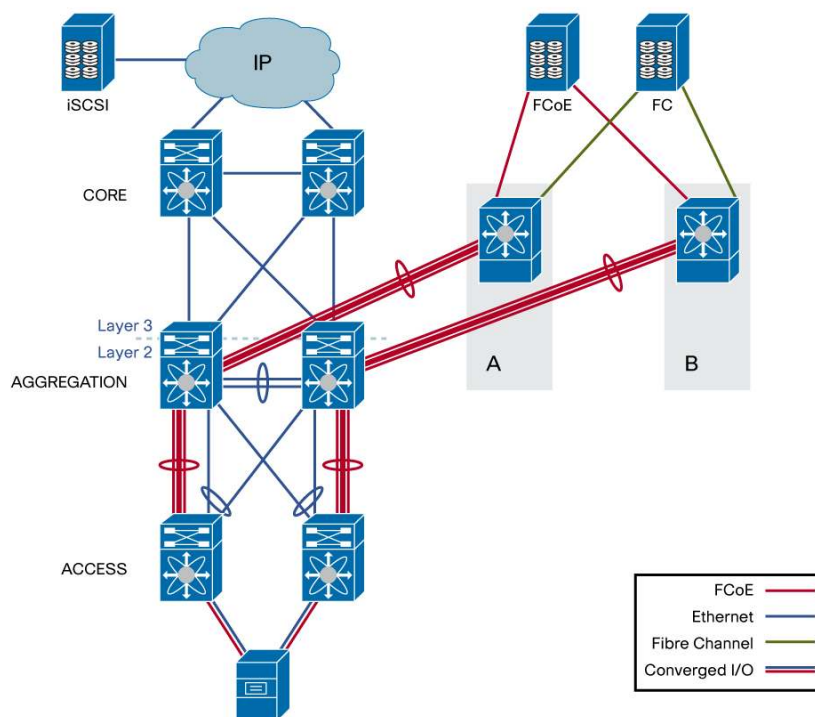
Ethernet network. Deployment of FCoE in an Ethernet network enables consolidation of the LAN and SAN traffic on the same physical Ethernet infrastructure. Storage device manufacturers are increasingly developing FCoE-enabled devices, helping the technology gain a place in end-to-end data center deployments.

Storage networks traditionally are designed with an application's I/O requirements taken into consideration. With the advent of server virtualization and virtualized applications, many purpose-built networks are being deployed and managed in the same data center. A single infrastructure that can meet all the I/O requirements is an attractive prospect, potentially reducing both capital expenditures (CapEx) and operating expenses (OpEx). A consolidated infrastructure would allow administrators to:

- Simplify the management of ports and I/O across servers
- Build redundancy once for all storage traffic
- Commission or decommission LAN and SAN ports without server reconfiguration
- Create much tighter security
- Easily create and manage authentication boundaries
- Use the same management tools and personnel across the network

Cisco Nexus 7000 Series Switches with FCoE give data center architects the flexibility to create one physical network running a multitude of storage protocols. In particular, it allows administrators to deploy a flexible iSCSI storage solution alongside their existing Fibre Channel-based SAN and meet the resiliency and high-availability requirements of each. Figure 4 shows a typical consolidated network.

Figure 4. Consolidated Network with Multiple Storage Protocol Types



Conclusion

Cisco Nexus 7000 Series Switches are full-featured, manageable, secure, highly available, and resilient Ethernet switches that give data centre designers and administrators the flexibility to deploy any storage device, anywhere in the network. In particular, IP storage devices benefit from the increased cross-sectional bandwidth, segregation of storage traffic, and high availability provided by these switches. These features enable data center designers to deploy IP storage arrays for mission-critical applications. In addition, by enabling the consolidation of I/O traffic, this platform allows administrators to reduce their overall CapEx and OpEx while taking advantage of both mature and future Ethernet features (40- and 100-Gbps Ethernet). This flexibility allows the Cisco Nexus 7000 Series to dynamically meet the changing business and performance requirements of the data center.

In particular, deployment of the Cisco Nexus 7000 Series as Ethernet storage director switches meets and exceeds the requirements for a successful data center.

For More Information



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.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned 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. (1005R)