Cisco NX-OS Software: Business-Critical Cross-Platform Data Center OS

What You Will Learn

Modern data centers power businesses through a new generation of applications and services. Virtualization, cloud computing, high-performance computing, data warehousing, and disaster recovery strategies, among others prevalent in the current environment, are prompting a whole new set of requirements for network infrastructure. To meet the needs of the modern data center. a network device - or more particularly, the operating system that powers that device - must be:

- Resilient: To provide critical business-class availability
- Modular: To be capable of extension to evolve with business needs and provide extended lifecycles
- Portable: For consistency across platforms
- Secure: To protect and preserve data and operations
- Flexible: To integrate and enable new technologies
- Scalable: To accommodate and grow with the business and its requirements
- Easy to use: To reduce the amount of learning required, simplify deployment, and ease manageability

Cisco® NX-OS Software is designed to meet all these criteria across all Cisco platforms that run it. This document describes the features of Cisco NX-OS operating system to help you understand how it can meet the needs of your organization.

Building on a Proven Foundation

Cisco NX-OS is a highly-evolved modular operating system that builds on more than 15 years of innovation and experience in high-performance switching and routing. Cisco NX-OS finds its roots in the Cisco SAN-OS operating system used worldwide in business-critical loss-intolerant SAN networks. As a direct result of having been deployed and evolving from nearly a decade in the extremely critical storage area networking space, NX-OS can deliver the performance, reliability, and lifecycle expected in the data center.

Cisco NX-OS is built on a Linux kernel. By using a version 2.6 based Linux kernel as its foundation, Cisco NX-OS gains the following benefits:

- Established and widely field-proven core kernel code
- Efficient multithreaded preemptive multitasking capabilities
- Native multiprocessor and multicore support

The choice of Linux over other operating systems and of the Linux 2.6 kernel over other versions of Linux was strategic. The following are some of the reasons for this choice:

- By inheritance, this implies that NX-OS shares the largest installed base of any UNIX-like operating system in existence. The community development model and widespread use of Linux provide the benefits of rigorous code review, rapid community-based defect resolution, and exceptional real-world field testing.
- The kernel is a near-real-time OS kernel, which is actually preferred over a true-real-time OS for applications such as networking, which may have many parallel critical activities running on a platform.

- This particular kernel version currently provides the best balance of advanced features and maturity and stability. It is currently the most widely deployed version of the Linux kernel.
- As of version 2.6 the kernel introduced an advanced and scalable kernel architecture leveraging multiple run-queues for handling multi-core and multiple-CPU system configurations.

These characteristics provide the solid foundation of resilience and robustness necessary for any network device OS powering the mission-critical environment of today's enterprise-class data centers.

The multithreaded preemptive multitasking capability provides protected fair access to kernel and CPU resources. This approach helps ensure that critical system processes and services are never starved for processor time. This feature, in turn, helps preserve system and network stability by helping ensure that routing protocols, spanning tree, and internal service processes get access to the CPU cores as needed.

Scalability for Future Growth

With Cisco NX-OS, scalability is integral and effectively built-in. As environments grow, the software's native support for multiprocessor and multicore hardware platforms helps simplify scalability through the effective use of current and future hardware.

With its pre-emptive multitasking multi-threaded kernel, Cisco NX-OS also provides advanced multicore and multiple-CPU processing. NX-OS incorporates a highly scalable CPU queue and process management architecture which employs multiple processor thread run-queues. This in turn enables more efficient use of modern multi-core CPUs. Combined with its memory mapping techniques and path to 64-bit, NX-OS provides simplified scalability both upwards and downwards to accommodate both control plane growth and multiple platforms.

Modular Code Base

Several categories of modular system code are built on top of the Linux kernel (Figure 1). These can be generally described as:

- Platform-dependent hardware-related modules
- System-infrastructure modules
- Feature modules

Figure 1. Cisco NX-OS Employs a Highly Granular Modular Architecture

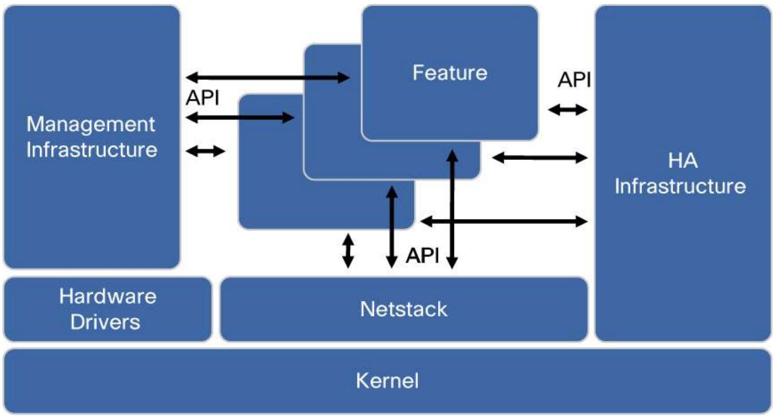

The platform-dependent hardware-related modules consist of subsystems such as hardware and chipset drivers specific to a particular hardware platform on which Cisco NX-OS runs. This portion of the OS is the part that must be ported across hardware platforms and allow the other subsystems within the OS to communicate with and tie into the specific hardware features of a platform. The platform-dependent modules typically provide standardized APIs and messaging capabilities to upper-layer subsystems. The modules essentially constitute a hardware abstraction layer to enable consistent development at higher layers in the OS, improving overall OS portability. The defined nature of the platform-dependent modules enables the overall reduction of the code base that specifically requires porting to deliver Cisco NX-OS on other hardware platforms. The result is greater consistency in implementation, reduced complexity in defect resolution, and faster implementation of cross-platform features across the various Cisco NX-OS platforms.

The system infrastructure modules provide essential base system services that enable system process management, fault detection, fault recovery, and interservice communication. The system management component of the system infrastructure provides service management for other features of the OS. It is also the component responsible for fault detection for the feature services, and it is fully capable of performing fault recovery of a feature service as needed. Working together with other infrastructure modules services, it can provide stateful fault recovery of a feature, enabling recovery of a fault within a specific feature in less than a second, while preserving the runtime state of that feature. This capability enables transparent and nondisruptive fault recovery within the system, increasing overall network stability and service uptime.

The individual feature modules consist of the actual underlying services responsible for delivering a particular feature or protocol capability. Open Shortest Path First (OSPF), Enhanced (Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS) Protocol, Border Gateway Protocol (BGP), Spanning Tree Protocol, Fibre Channel over Ethernet (FCoE), the routing information base (RIB), Overlay Transport Virtualization (OTV), and NetFlow export are all examples of system-level features embodied in modular components.

Each feature is implemented as an independent, memory-protected process spawned as needed based on the overall system configuration. This approach differs from that of traditional network operating systems in that only the specific features that are configured are automatically loaded and started. This highly granular approach to modularity enables benefits such as:

- Compartmentalization of fault domains within the OS and its services, resulting in significantly improved overall system resiliency and stability
- Simplified portability for cross-platform consistency through reusable components, or building blocks. and little use of platform-specific code
- More efficient defect prioritization and repair through the isolation of specific functions to particular modules
- Improved long-term platform extensibility through the capability to easily integrate new feature modules into the OS infrastructure through established and consistent OS interfaces
- More efficient resource utilization because only features specifically enabled through configuration are loaded into memory, present command-line interface (CLI) elements, and consume CPU cycles
- Improved security because features that are not configured or enabled do not run, thus reducing the exposure of the OS to attacks

Intelligent Fault Detection and Recovery

In addition to the resiliency gained from architectural improvements, Cisco NX-OS provides internal hierarchical and multilayered system fault detection and recovery mechanisms. No software system is completely immune to problems, so an effective strategy for detecting and recovering from faults quickly and with as little effect as possible is essential. Cisco NX-OS is designed from the start to provide this capability.

Individual service and feature processes are monitored and managed by the Cisco NX-OS System Manager, an intelligent monitoring service with integrated high-availability logic. The system manager can detect and correct a failure or lockup of any feature service within the system. The system manager is in turn monitored and managed for health by the Cisco NX-OS kernel. A specialized portion of the kernel is designed to detect failures and lockups of the Cisco NX-OS System Manager. The kernel itself is monitored through hardware. A hardware process constantly monitors the kernel health and activity. Any fault, failure, or lockup at the kernel level is detected by hardware and will trigger a supervisor switchover. Figure 2 shows the fault detection and recovery process.

Figure 2. Cisco NX-OS Provides Multilevel Hierarchical Fault Detection and Recovery

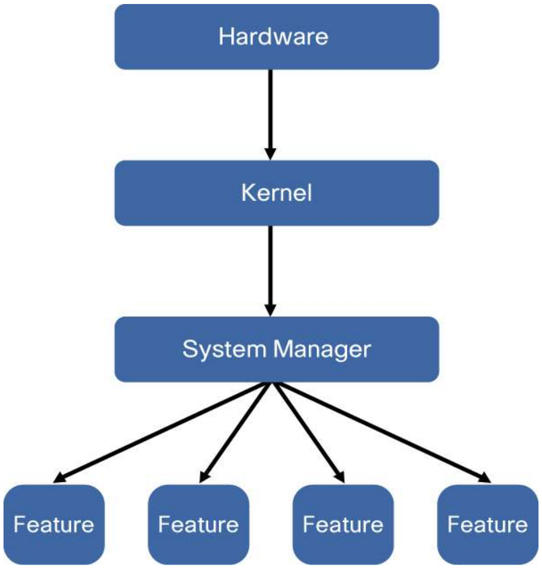

The combination of these multilevel detection and health monitoring systems provides creates a robust and resilient operating environment that can reduce the overall effect of internal faults and, more importantly, preserve the stability of the overall network by internalizing these types of events.

Continuous Operation and High Availability

Cisco NX-OS is designed from the start to provide consistent, predictable, and reliable high availability. The design goal for the data center is continuous operation: no service disruption. Cisco NX-OS provides a high-availability architecture that moves toward this goal with fully nondisruptive stateful supervisor switchover (SSO) for control-plane redundancy in modular platforms, and nondisruptive In-Service Software Upgrade (ISSU) for all Cisco Nexus® platforms.

When running on platforms that offer redundant control-plane hardware, Cisco NX-OS is designed to provide efficient event-based state synchronization between active and standby control-plane entities. This approach allows the system to rapidly perform a fully stateful control-plane switchover with little system disruption and no service disruption.

For platforms without redundant control planes, Cisco NX-OS can implement ISSU by retaining the software state throughout the upgrade process and retaining packet-forwarding intelligence through its hardware subsystems, preventing service disruption.

Cisco NX-OS is also designed to take full advantage of the distributed environment on platforms with distributed hardware forwarding, so that data-plane forwarding is not affected during redundant control-plane switchover. This architecture effectively delivers true nondisruptive control-plane failover that has been verified to date by several independent third-party sources.

At its core, Cisco NX-OS is designed to take advantage of distributed platforms to reduce any effects on the data plane during control-plane operations, including software upgrades. Again, on platforms designed to be highly distributed, it uses this same high-availability infrastructure and distributed architecture to deliver fully non disruptive ISSU. The control planes of a distributed system are upgraded without affecting data-plane forwarding, and after the control planes are successfully upgraded, the control-plane portions of any forwarding hardware that can be nondisruptively upgraded are serviced. This approach effectively transforms planned maintenance windows so that they longer automatically imply a service outage. This increased level of continuous operation successfully accommodates business-critical environments, in which little downtime or degradation of service is essential.

Enhanced Usability and Familiar Operation

Cisco IOS® Software is already the recognized leader in internetworking device operating systems. For decades, Cisco IOS Software has been the foundation for routing and switching configuration in all environments. The Cisco IOS CLI has essentially become the standard for configuration in the networking industry.

To reduce the amount of time needed to learn Cisco NX-OS and to accelerate adoption, Cisco NX-OS maintains the familiarity of the Cisco IOS CLI. Users comfortable with the Cisco IOS CLI will find themselves equally comfortable with Cisco NX-OS. In addition, Cisco NX-OS has integrated numerous user interface enhancements on top of the familiar Cisco IOS CLI to make configuration and maintenance more efficient. These are just some of the simple but effective UI enhancements found in Cisco NX-OS:

- Nonhierarchical CLI: Almost any command can be run from any mode. You can run show commands from the interface and global configuration modes. Global commands can be run from the interface configuration mode. The system is intelligent enough to determine the nature of a command and process it regardless of whether the current mode the configuration or execution mode.
- Configuration mode contextual CLI history: Separate CLI command histories are maintained for each configuration mode. Reentry of commands in a given configuration mode is simplified; you can use the up and down arrow to cycle through the command history stored for that particular configuration mode.
- Advanced multilevel and multicommand output piping: The capability to stream CLI output through advanced filters and parsing commands enables complex formatting and manipulation of information for easier parsing and processing.
- More verbose and descriptive status output: The show command output tends to be more informative and less obscure or opaque in Cisco NX-OS, allowing more effective troubleshooting and status monitoring.

Cisco IOS Software users will quickly find themselves familiar with the Cisco NX-OS CLI and its enhancements. Typically, most networking professionals will also quickly find themselves seeking those functional enhancements in other operating systems.

Virtual Device Contexts

Cisco NX-OS also provides the capability to virtualize the platform on which it is running. Using Cisco NX-OS virtual device contexts (VDCs), a single physical device can be virtualized into many logical devices, each operating independently, effectively approximating separate physical devices (Figure 3).

Figure 3. VDCs Enable One-to-Many Virtualization of Logical Devices from a Single Platform

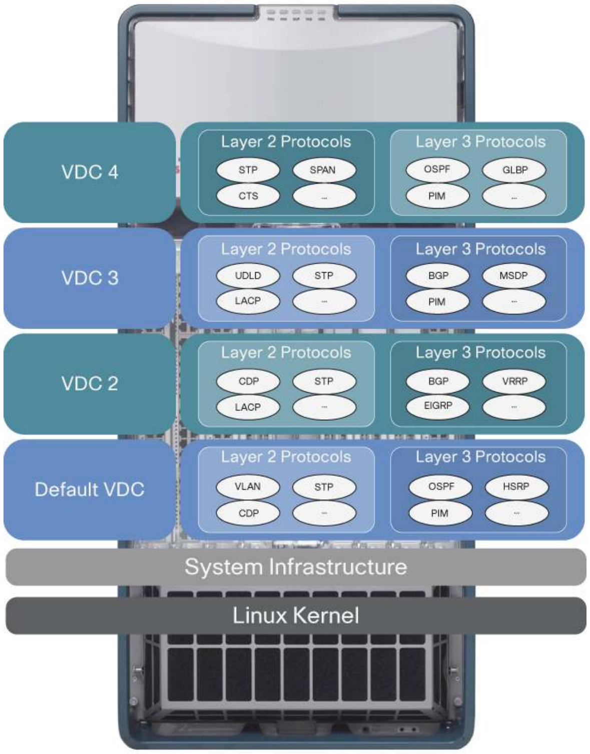

Cisco NX-OS VDCs differ from other implementations of virtualization found in most networking devices by applying in-depth virtualization across multiple planes of operation:

- Virtualization at the data plane: Physical interfaces are associated with a specific VDC instance. Data-plane traffic transiting the physical device can be switched only from one interface in a VDC to another within the same VDC. This virtualization is internalized in the switching system and is not subject to external influence, providing very strong data-plane separation between VDCs. The only means of integrating traffic between VDCs is through a physical cross-connection of ports between two or more VDCs.
- Virtualization at the control plane: All control-plane functions are virtualized within the operating system at the process level. This approach effectively creates separate failure domains between VDCs, reducing the fate-sharing between them. Network or system instability within the domain of a single VDC does not affect other VDCs or the network domain in which they are operating.
- Virtualization at the management plane: Virtualization of the management plane is where VDCs truly stand out compared to other network device virtualization solutions. VDCs virtualize the configuration of each logical device, and they also virtualize all supporting management environment services and capabilities. Each VDC maintains separate configurations and operational relationships with typical common support and security services such as:
    - Separate independent syslog servers configurable per VDC
    - Separate independent authorization, authentication, and accounting (AAA) servers configurable per VDC
    - Independent addressable management IP addresses per VDC
    - Separate independent NetFlow export targets per VDC
    - Independent per-VDC local authentication user lists with per-VDC role-based access control (RBAC)

The end result is extremely effective separation of data traffic and operational management domains suitable for cost-effective infrastructure consolidation in security-sensitive environments.

Security

Cisco NX-OS provides the tools required to enable the advanced security features needed to protect the network infrastructure as well as the actual platform on which it is running.

Security is designed using a two-pronged approach: security at the data plane and security at the control plane. This approach effectively secures transient traffic passing through a Cisco NX-OS device as well as the traffic destined for the device itself.

Cisco NX-OS currently enables the deployment of both common and more advanced data-plane and infrastructure-level security features, including:

- IEEE 802.1ae and Cisco TrustSec™ platform
- IP source guard (IPSG)
- Dynamic Host Control Protocol (DHCP) snooping
- Unicast reverse-path forwarding (uRPF)
- Hardware-based IP packet validity checking
- Port, router, and VLAN access control lists (ACLs)
- IEEE 802.1x
- Bridge Protocol Data Unit (BPDU) guard

These controls enable effective protection against most man-in-the-middle, common resource, and spoofing attacks.

At the control plane, Cisco NX-OS provides a robust security toolset to prevent attacks directed at the Cisco NX-OS device itself or active sessions on the device. Tools include capabilities such as:

- Control-plane policing (CoPP)
- RBAC with RADIUS and TACACS+ integration
- Strong password checking
- Secure Shell (SSH) Protocol and Secure FTP (SFTP)

More important, the overall Cisco NX-OS system architecture provides the interfaces and structures necessary to easily implement future security features consistently and cleanly.

Unified I/O and Unified Fabric

Cisco NX-OS delivers unified I/O and unified fabric architecture capabilities, introducing a new model of data center design. Cisco Unified Fabric is a critical building block for traditional and virtualized data centers, unifying storage networking, data networking, and services to achieve transparent multiprotocol convergence, multidimensional scale, and distributed intelligence and enabling customers to derive greater value from their network platform investments. Complementing the Cisco Unified Computing and Unified Network Services, Cisco Unified Fabric is a foundational element of the Cisco Data Center Business Advantage architectural framework.

Cisco Unified Fabric provides the flexibility to run Fibre Channel, IP-based storage such as network-attached storage (NAS) and Small Computer System Interface over IP (iSCSI), or FCoE, or a combination of these technologies, on a converged network. Providing the best of both LAN and SAN capabilities, Cisco Unified Fabric enables storage network users to take advantage of the economies of scale, robust vendor community, and aggressive roadmap of Ethernet while providing high-performance, lossless characteristics of a Fibre Channel storage network. Cisco Unified Fabric deployment can easily be implemented through a phased approach; because FCoE is fully interoperable with Fibre Channel, existing networks can gradually evolve to unified fabrics.

Cisco Nexus 5000 and 4000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders enable a single-hop FCoE architecture at the access layer. This capability combined with Cisco Fabric Extender Link (FEX-Link) technology provides a nearly immediate, low-cost high-density entry point into unified I/O with relatively few overall network design changes.

Cisco Nexus 7000 Series Switches and Cisco MDS 9000 Family products are also integral parts of this architecture, with the introduction of the Cisco Nexus F-Series Modules, enriching the platforms that supporting Cisco Unified Fabric and the design possibilities, including multihop FCoE. The FCoE capabilities on the Cisco Nexus 7000 Series support a number of flexible designs, enabling unified network fabric deployment benefits from the access layer all the way through the aggregation layer and core of the data center network.

Cisco Unified Fabric delivers reliable, agile, and cost-effective network services to servers, storage, and applications while improving the user experience across the distributed enterprise. It provides many benefits for users, reducing capital expenditures (CapEx) through infrastructure reduction and operating expenses (OpEx) through network simplification, saving power, cooling, and space costs while protecting the organization’s existing investment in tools, training, and infrastructure.

Cisco FabricPath

The ability to build larger, more scalable, Layer 2 domains while preserving stability, resiliency, and robustness is becoming a crucial criterion for modern data center network design. Another key demand on the network infrastructure, stemming from the drive towards large-scale virtualization, is flexibility. Business applications require flexibility of provisioning any workload anywhere in the data center; private “cloud” type environments require the flexibility of transparent resource allocation.

Cisco NX-OS is architecturally designed and ready to support the evolution in Ethernet topology management and forwarding logic needed to deliver these capabilities. This evolution encompasses an industry wide shift away from traditional Spanning-Tree protocol to Link-State routing at the link layer for topology and forwarding management. By building on the key benefits offered by technologies based on decades of research and development in IP routing, Cisco FabricPath seeks to deliver the infrastructure required to form the foundation of the data center “cloud” that powers all enterprise applications.

Multiple components contribute to this capability:

- An advanced approach to topology management that leverages link-state routing at the Ethernet link layer to provide advanced multi-pathing and network resilience, providing some of the benefits obtained in IP routing.
- The introduction of hierarchical abstraction in the Layer 2 infrastructure to reduce the amount of state that needs to be stored across devices in order to improve scalability.
- More intelligent source and destination MAC address learning and forwarding - allowing greater scalability while utilizing resources more efficiently.
- Effective traffic load balancing and engineering capabilities to enable more efficient utilization of already deployed bandwidth in the infrastructure.

Figure 4. Cisco FabricPath Architecture Provides Freedom from Typical Layer 2 Spanning-Tree Design Constraints, Enabling Network, Application, and Business Flexibility Across the Data Center

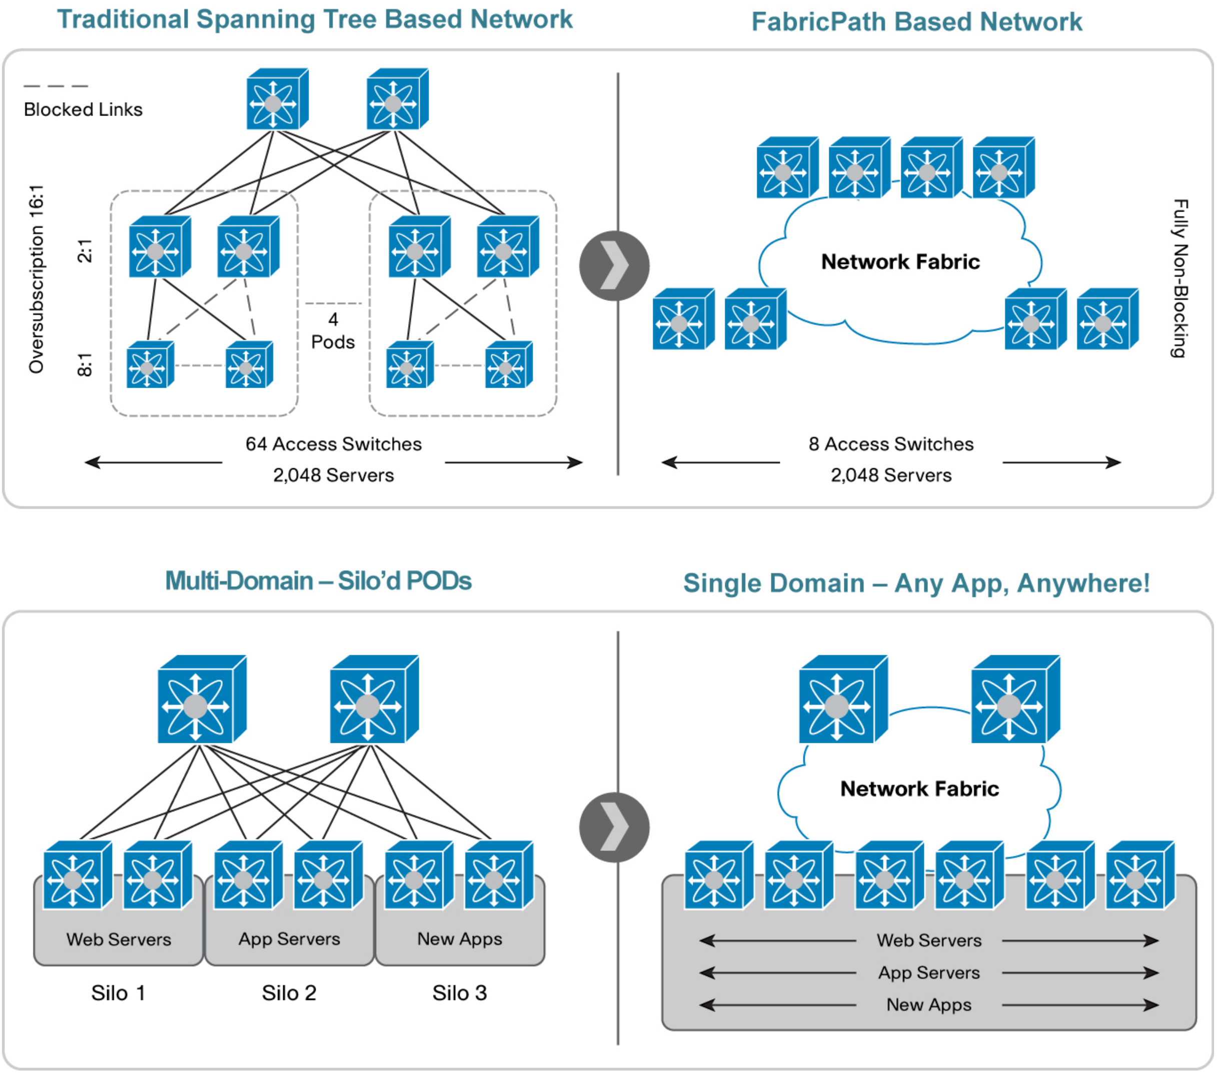

Cisco vPC and vPC+ Technology

Cisco vPC technology enables the deployment of a link aggregation from a generic downstream network device to two individual and independent Cisco NX-OS devices (vPC peers). This diverse, multichassis link aggregation path provides both link redundancy and active-active link throughput scaling with high-performance failover characteristics. vPC is delivered in the form of an industry standard IEEE 802.3ad Link Aggregation Control Protocol (LACP) PortChannel interface. No special handling or intelligence is required by the generic downstream device other than support for IEEE 802.1ag LACP. The logical link bundling across two nodes is handled by Cisco NX-OS at the pair of upstream Cisco NX-OS devices that provide vPC capabilities.

A significant benefit to vPC is reduced reliance on the Spanning Tree Protocol to provide topology redundancy and loop management. Since the virtual PortChannels are presented as a single logical link, the actual spanning-tree topology is logically loop free, thereby reducing the number of links that are blocked by spanning tree. All link failures are rerouted to redundant active paths using PortChannel hashing logic instead of spanning tree, which results in much faster failover times. The reduction in logical looping also reduces the complexity of the overall spanning-tree domain.

vPC, when coupled with Configuration Sync, a feature that allows intelligent synchronization of configurations between the vPC peers, drastically reduces management complexity.

Cisco FEX-Link Technology

Cisco FEX-Link technology enables data center architects to gain new design flexibility while simplifying cabling infrastructure and management complexity. Cisco FEX-Link uses the Cisco Nexus 2000 Series Fabric Extenders to extend the capacities and benefits offered by upstream Cisco Nexus switches.

Fabric extenders are essentially extensions of the parent Cisco Nexus switch fabric, with the fabric extenders and the parent Cisco Nexus switch together forming a distributed modular system. This architecture enables flexible physical topologies, combining the flexibility and benefits of both top-of-rack (ToR) and end-of-row (EoR) deployments.

Cisco FEX-Link provides a technology platform for highly scalable unified server access across a range of 100 Megabit Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, unified fabric, copper, and fiber connectivity rack and blade server environments. The platform is well suited to support today's Gigabit Ethernet and 10 Gigabit Ethernet environments and allows transparent migration to 10 Gigabit Ethernet, virtual machine-aware unified fabric technologies (Figure 4).

Figure 5. Cisco FEX-Link Architecture Provides Highly Scalable Unified Server Access Connectivity

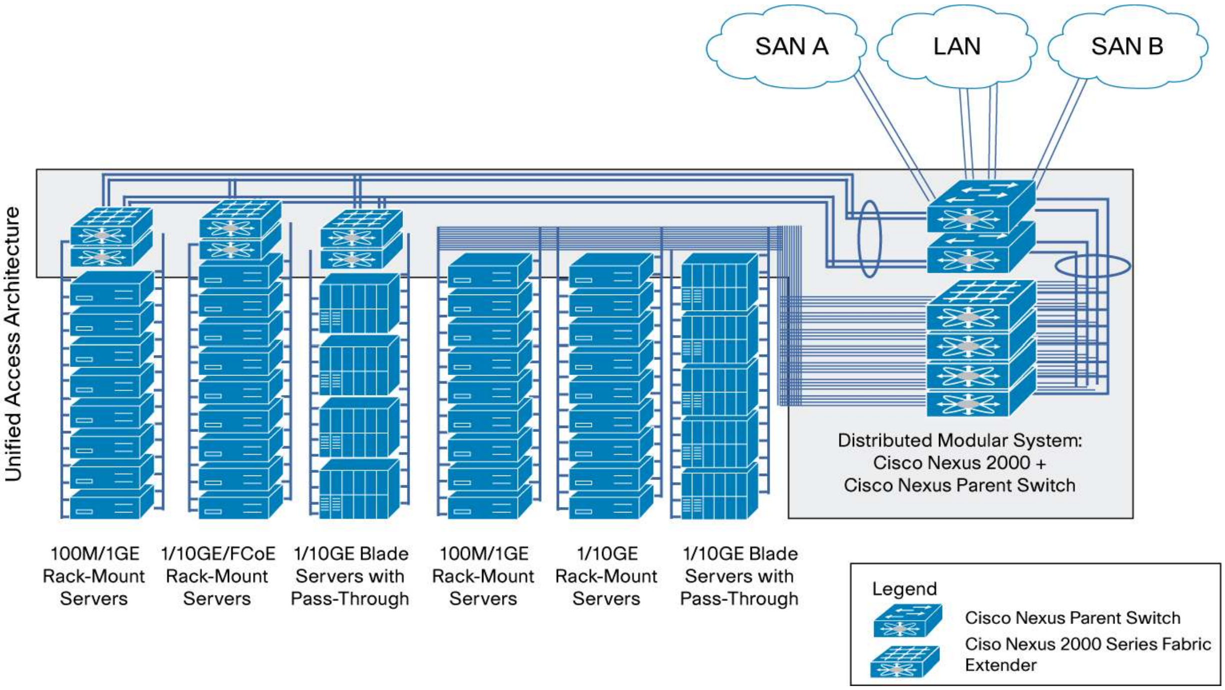
SAN A
LAN
SAN B
Unified Access Architecture
Distributed Modular System:
Cisco Nexus 2000 +
Cisco Nexus Parent Switch
100M/1GE Rack-Mount Servers
1/10GE/FCoE Rack-Mount Servers
1/10GE Blade Servers with Pass-Through
100M/1GE Rack-Mount Servers
1/10GE Rack-Mount Servers
1/10GE Blade Servers with Pass-Through
Legend
Cisco Nexus Parent Switch
Ciso Nexus 2000 Series Fabric Extender

Cisco FEX-Link architecture provides the following benefits:

- Architecture flexibility: Common, scalable, and adaptive architecture across data center racks and points of delivery (PoDs)[1] supports a variety of server options, connectivity options, physical topologies, and evolving needs.

[1] A PoD is a module or group of network, computing, storage, and application components that work together to deliver a network service. The PoD is a repeatable pattern, and its components increase the modularity, scalability, and manageability of data centers.

- Simplified operations: Simplified operations are provided through a single point of management and policy enforcement using Cisco Nexus switches. This feature eases the commissioning and decommissioning of server racks with zero-touch installation and automatic configuration of fabric extenders.
- Breakthrough business benefits: Scalable 10 Gigabit Ethernet provides 10 times the bandwidth for approximately twice the price of Gigabit Ethernet. Consolidation, cabling reduction, rack-space reduction, reduced power and cooling, investment protection through feature inheritance from the parent switch, and the capability to add functions without the need for a major equipment upgrade of server-attached infrastructure all contribute to reduced OpEx and CapEx.

Virtualization Ready

A network infrastructure that is more integrated and aware in environments that use server virtualization is a critical element of the modern data center. Cisco NX-OS incorporates several technologies and an effective roadmap for support of virtualization-aware networking. This approach allows Cisco NX-OS to enable tighter integration of the network into virtualized server environments to simplify the management, orchestration, and provisioning of data center resources.

The Cisco VN-Link technology for virtualization awareness in the network delivers several features to provide flexible methods of coupling the network configuration with virtualized endpoints. The portfolio of Cisco VN-Link products provides a variety of options that satisfy a range of customer needs. Cisco VN-Link provides advanced hypervisor switching as well as high-performance hardware switching; it is flexible, extensible, and service enabled (Figure 5).

Figure 6. Cisco VN-Link Architecture Provides Virtualization-Aware Networking and Policy Control

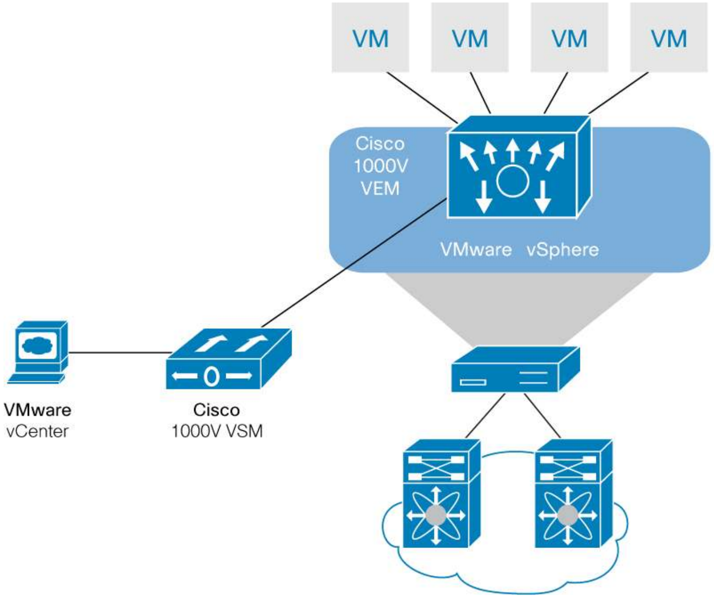

The technologies that constitute the building blocks of Cisco VN-Link include:

- Port profiles: These intelligent configuration templates allow dynamic provisioning of relevant configuration parameters to affected interfaces or ports.
- Virtual Ethernet interfaces: These soft interfaces can be associated with a virtualized endpoint, allowing interface parameters associated with that endpoint to fluidly move with the endpoint as needed.
- Virtual Ethernet module (VEM): This lightweight software component represents the data plane and runs inside the hypervisor. It enables advanced networking and security features and performs switching between directly attached virtual machines.

- Virtual supervisor module (VSM): This Cisco NX-OS software-based physical or virtual appliance provides command and control along with management and monitoring of virtual Ethernet interfaces through a traditional network CLI. The VSM also provides integration with the hypervisor management tools.
- VN-tag: Virtualized endpoints need to be identified with a tag when an external hardware switch is used for virtual machine traffic forwarding instead of the software switch within the hypervisor. Such an identifier is used inline to granularly identify specific traffic destined for or originated by a given virtualized endpoint. This approach allows specific streams of traffic within an aggregate to be identified and associated with a specific virtualized endpoint for more granular and efficient application of network-based services and policies.

While Cisco actively provides contributions and feedback to standards and industry bodies to improve virtualization awareness in the network, Cisco NX-OS remains a proactive leader in this space by providing capabilities today across the Cisco Nexus Family of products to help provide customers with solutions that satisfy immediate near-term requirements.

Overlay Transport Virtualization

OTV enables high-performance with simplified and scalable multipoint extension of Layer 2 domains across underlying segments of Layer 3 routed networks. This feature allows the network to retain the fault isolation and scalability characteristics of Layer 3 routing in the core of the network, but still enables Layer 2 adjacency for applications that require it (Figure 7).

Figure 7. Cisco Overlay Transport Virtualization Provides Simplified Yet Scalable Layer 2 Extension for Applications That Require Layer 2 Adjacency

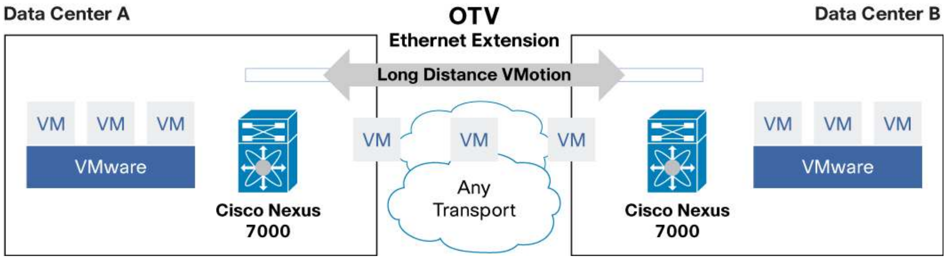

OTV intelligently connects Layer 2 domains without unnecessarily extending and joining spanning tree domains. This feature preserves individual fault domains and improves overall scalability. Additionally, OTV can provide intelligent conversation-based learning of link layer addresses (MAC addresses) without flooding participating domains, reducing the amount of traffic sent to and from the core and the amount of traffic that must be processed by OTV participating nodes.

Comprehensive IPv6 Support

From the start, Cisco NX-OS was designed to comprehensively support IPv6. Its broad portfolio of IPv6 features and extensible architecture allows it to integrate flexibly into world-class enterprise and service provider IPv6 environments.

Support for common IPv6 routing protocols, features, and functions and the capability to easily add more make Cisco NX-OS an excellent OS for IPv6 critical deployments.

For detailed IPv6 feature support information, please refer to the data sheet for the specific device supported by the Cisco NX-OS platform.

Advanced System Management

The operation, administration, and monitoring (OAM) of network elements are critical to long-term data center infrastructure sustainability. Cisco NX-OS includes a number of traditional and advanced features to ease OAM in the data center.

- Simple Network Management Protocol (SNMP): Traditional SNMP Versions 1, 2c, and 3 are supported for read operations, allowing integration into existing systems for effective monitoring. Cisco NX-OS on the Cisco Nexus Family is also certified in the latest versions of the EMC Ionix management platform.
- NETCONF and XML: Cisco NX-OS provides integration with IETF NETCONF-compliant systems through XML transactions over a secure SSH interface.
- Cisco Generic Online Diagnostics (GOLD): Cisco NX-OS supports Cisco GOLD online diagnostics for active component and subsystem testing.
- Cisco Embedded Event Manager: Cisco NX-OS provides a scripted interface that enables the configuration of automated event-triggered actions to be run by the system autonomously.
- Single-image download: Because of its simplified licensing and image structure, every image of Cisco NX-OS contains all the features of Cisco NX-OS available at that release. Individual features are loaded, enabled, and made available to the system platform based on Cisco NX-OS electronic licensing. Therefore, only a single image is available for download for a given version on a given platform. No decoder or version and feature chart is required to determine which image is appropriate for a given environment.
- Scheduler: Cisco NX-OS includes a generic system scheduler that can be configured to run CLI-based system commands at a given time, on a one-time or recurring basis.
- CallHome: The CallHome feature in Cisco NX-OS allows an administrator to configure one or more contact email addresses that are notified when the CallHome function is triggered. The notification process is triggered during certain events that are configurable, such as Cisco GOLD test results and scripts that run based on Cisco EEM events. CallHome enables rapid escalation of events and proactive prevention of a pending failure.
- Configuration checkpoint and rollback: Cisco NX-OS incorporates an advanced configuration and rollback facility to preserve and protect the configuration state. Configuration snapshots, or checkpoints, can be created manually at the CLI or initiated automatically by the system at major configuration events (such as the disabling of a feature). Checkpoints can be stored locally on the device in the local checkpoint database or in a file in integrated or removable storage. Using the rollback capability, the current running configuration can be restored to a particular state stored in a checkpoint.

Conclusion

The architecture and features discussed here are only some of the characteristics that make Cisco NX-OS the most advanced data center device OS available. The reliability, resiliency, availability, and extensibility of Cisco NX-OS provide a solid foundation on which to build business-critical data center environments. Using that foundation to advance technologies in the data center to meet the requirements of current and future generations of applications and services effectively positions Cisco NX-OS as the internetworking device operating system for the next decade.

For More Information

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

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

C11-622511-00 12/10