Cisco NX-OS Software

Product Overview

Cisco® NX-OS Software is a data center-class operating system built with modularity, resiliency, and serviceability at its foundation. Based on the industry-proven Cisco MDS 9000 SAN-OS Software, Cisco NX-OS helps ensure continuous availability and sets the standard for mission-critical data center environments. The self-healing and highly modular design of Cisco NX-OS makes zero-impact operations a reality and enables exceptional operational flexibility.

Focused on the requirements of the data center, Cisco NX-OS provides a robust and comprehensive feature set that fulfills the switching and storage networking needs of present and future data centers. With an XML interface and a command-line interface (CLI) like that of Cisco IOS® Software, Cisco NX-OS provides state-of-the-art implementations of relevant networking standards as well as a variety of true data center-class Cisco innovations.

Cisco NX-OS offers reliability, innovation, and operational consistency across data center platforms. Cisco NX-OS runs on the Cisco Nexus® Family of hardware-based network switches, which include Cisco Nexus 7000, 5000, 4000, and 1000V Series Switches and Cisco Nexus 2000 Series Fabric Extenders; Cisco MDS 9000 Family storage switches; and Cisco UCS 6100 Series Fabric Interconnects.

Features and Benefits

Built as the foundation of the Cisco Data Center Business Advantage (DCBA) solution, Cisco NX-OS helps ensure continuous availability and sets the standard for mission-critical environments. The main attributes, which constitute the Cisco NX-OS foundation, are summarized here.

Resiliency

Cisco NX-OS is designed from the start to deliver continuous operation with failure detection, fault isolation, self-healing features, and small maintenance windows.

- **Modular software design:** Cisco NX-OS is designed to support distributed multithreaded processing on symmetric multiprocessors (SMPs), multicore CPUs, and distributed line-card processors. Cisco NX-OS modular processes are instantiated on demand, each in a separate protected memory space. Thus, processes are started and system resources allocated only when a feature is enabled. The modular processes are governed by a real-time preemptive scheduler that helps ensure the timely processing of critical functions.
- **Continuous system operation:** Cisco NX-OS provides continuous system operation, permitting maintenance, upgrades, and software certification without service interruption. The combination of process modularity, Cisco In-Service Software Upgrade (ISSU) capability, and stateful graceful restart mitigates the effects of software upgrades and other network operations.
- **Cisco ISSU:** Cisco ISSU provides the capability to perform transparent software upgrades on platforms with redundant supervisors, reducing downtime and allowing customers to integrate the newest features and functions with little or no negative effect on network operation.
- **Quick development of enhancements and problem fixes:** The modularity of Cisco NX-OS allows new features, enhancements, and problem fixes to be quickly integrated into the software. These updated images can then be installed without disruption using Cisco ISSU.

- **Process survivability:** Critical processes are run in protected memory space and independently of each other and the kernel, providing granular service isolation and fault containment and enabling modular patching and upgrading and rapid restartability. Individual processes can be restarted independently without loss of state information and without affecting data forwarding, so that after an upgrade or failure, processes restart in milliseconds without negatively affecting adjacent devices or services. Processes with large amounts of state such as IP routing protocols are restarted using standards-based nonstop forwarding (NSF) graceful restart mechanisms; other processes use a local persistent storage service (PSS) to maintain their state.
- **Reliable interprocess communication:** Cisco NX-OS facilitates reliable communication between processes to help ensure that all messages are delivered and properly acted on during failures and adverse conditions. This communication helps ensure process synchronization and state consistency across processes that may be instantiated on processors distributed over multiple supervisors and I/O modules.
- **Stateful supervisor failover:** Redundant supervisors are kept synchronized at all times to enable rapid stateful supervisor failover. Sophisticated checks are in place to help ensure that the state is consistent and reliable throughout the entire distributed architecture after failover occurs.
- **Network-based availability:** Network convergence is optimized by providing tools and functions to make both failover and fallback transparent and fast. For example, Cisco NX-OS provides Spanning Tree Protocol enhancements such as Bridge Protocol Data Unit (BPDU) guard, loop guard, root guard, BPDU filters, and bridge assurance to help ensure the health of the Spanning Tree Protocol control plane; Unidirectional Link Detection (UDLD) Protocol; NSF graceful restart of routing protocols; IEEE 802.3ad link aggregation with adjustable timers; virtual Port Channel (vPC); Cisco FabricPath; and Bidirectional Forwarding Detection (BFD).

Extensibility

Cisco NX-OS is highly scalable and can easily integrate with and adapt to ongoing innovation, technologies, and evolving standards.

- **Software compatibility:** Cisco NX-OS interoperates with Cisco products running any variant of the Cisco IOS Software operating system. Cisco NX-OS also interoperates with any networking OS that conforms to the networking standards listed as supported in this data sheet.
- **Common software throughout the data center:** Cisco NX-OS simplifies the data center operating environment and provides a unified OS designed to run all areas of the data center network, including storage, virtualization, and Layer 3 network protocols
- **Ethernet switching:** Cisco NX-OS is built to support high-density, high-performance Ethernet systems and provides a complete data center-class Ethernet switching feature set. Table 1 summarizes the Layer 2 feature set.

Table 1. Layer 2 Feature Set

Layer 2 Features
Layer 2 switch ports and VLAN trunks
IEEE 802.1Q VLAN encapsulation
Support for up to 4000 VLANs
Support for up to 32 virtual SANs (VSANs) per switch
Rapid Per-VLAN Spanning Tree Plus (PVRST+) (IEEE 802.1w compatible)
Multiple Spanning Tree Protocol (MSTP) (IEEE 802.1s): 64 instances
Spanning Tree PortFast, Root Guard, and Bridge Assurance
Cisco EtherChannel technology (up to 16 ports per EtherChannel)

Layer 2 Features
Cisco vPC technology
vPC configuration synchronization
LACP: IEEE 802.3ad
Advanced PortChannel hashing based on Layer 2, 3, and 4 information
Jumbo frames on all ports (up to 9216 bytes)
Pause frames (IEEE 802.3x)
Storm control (unicast, multicast, and broadcast)
Link Layer Discovery Protocol (LLDP; IEEE 802.1AB),
UDLD in aggressive and standard modes
VLAN Trunking Protocol (VTP) Versions 1 and 2 in client, server, pruning, and transparent modes
Private VLANs
Private VLAN over trunks (isolated and promiscuous)
Private VLANs over vPC and EtherChannels
Cisco FabricPath

- **Virtual PortChannel:** The vPC feature allows one end of a PortChannel to be split across a pair of Cisco Nexus switches. vPC provides Layer 2 multipathing through the elimination of Spanning Tree Protocol and enables fully utilized bisectional bandwidth and simplified Layer 2 logical topologies without the need to change the existing management and deployment models.
- **Cisco Overlay Transport Virtualization (OTV):** OTV is a “MAC address in IP” technique for supporting Layer 2 VPNs over any transport, whether it is Layer 2 based or Layer 3 based. By using the principles of MAC address routing, OTV provides an overlay that enables Layer 2 connectivity between separate Layer 2 domains while preserving the fault-isolation benefits of an IP-based interconnection. The core principles on which OTV operates are the use of a control protocol to advertise MAC address reachability information (instead of using data-plane learning) and packet switching of IP encapsulated Layer 2 traffic (instead of using circuit switching). Some of the main benefits achieved with OTV include:
  - **Zero impact on existing network design:** OTV is a transport-agnostic Layer 2 interconnect technology. The configuration is transparent to the sites under consideration.
  - **Failure isolation:** Failure boundaries and site independence are preserved. OTV does not rely on traffic flooding to propagate reachability information for MAC addresses; instead, a control protocol is used to distribute such information, sites remain independent of each other, and failures do not propagate beyond the OTV edge device.
  - **Optimized operations:** OTV enables single-touch site additions and removals. This feature has major operational benefit given that the configuration is succinct and uses a single protocol with no add-ons.
  - **Optimal bandwidth utilization, resiliency, and scalability:** OTV allows multipathing (cross-sectional bandwidth and end-to-end Layer 2 multipathing), transparent multihoming with built-in loop prevention, and multipoint connectivity in an easy-to-manage point-to-cloud model. It does not require the creation of closed tunnels, and the only state maintained is that of a MAC address routing table. The state is distributed and can be programmed in the hardware conditionally to allow the overlay to handle larger numbers of MAC addresses.
  - **Transparent migration path:** Since OTV is agnostic to the core and transparent to the sites, it can be incrementally deployed over any existing topology without the need to alter the network design.
- **Cisco FabricPath:** Cisco FabricPath is a set of multipath Ethernet technologies that combine the reliability and scalability benefits of Layer 3 routing with the flexibility of Layer 2 networks, enabling IT to build massively scalable data centers. Cisco FabricPath offers a topology-based Layer 2 routing mechanism that

provides an equal-cost multipath (ECMP) forwarding model. Cisco FabricPath implements an enhancement that solves the MAC address table scalability problem characteristic of switched Layer 2 networks. Furthermore, Cisco FabricPath supports vPC+, a technology similar to vPC that allows redundant interconnection of the existing Ethernet infrastructure to Cisco FabricPath without using Spanning Tree Protocol. Benefits introduced by the Cisco FabricPath technology include:

  - **Operational simplicity:** Cisco FabricPath embeds an autodiscovery mechanism that does not require any additional platform configuration. By offering Layer 2 connectivity, the “VLAN anywhere” characteristic simplifies provisioning and offers workload flexibility across the network.
  - **High resiliency and performance:** Since Cisco FabricPath is a Layer 2 routed protocol, it offers stability, scalability, and optimized resiliency along with network failure containment.
  - **Massively scalable fabric:** By building a forwarding model on 16-way ECMP routing, Cisco FabricPath helps prevent bandwidth bottlenecks and allows organizations to add capacity dynamically, without network disruption.

- **Locator/ID Separation Protocol (LISP):** LISP is an evolutionary routing architecture designed for Internet scale and global reach across organizations. The scalability of the routing system and the exhaustion of the IPv4 address space have motivated several proposals based on a common concept: the separation of the locator and identifier in the numbering of Internet devices, often called the locator ID (Loc/ID) split. LISP defines this protocol. The basic idea behind the Loc/ID split is that the current Internet routing and addressing architecture combines two functions: routing locators (RLOCs), which describe how a device is attached to the network, and endpoint identifiers (EIDs), which define “who” the device is, in a single numbering space: the IP address. The advantages include improved scalability of the routing system through greater aggregation of RLOCs. Cisco LISP Virtual Machine Mobility (VM-Mobility) is designed to enable global IP endpoint mobility across private networks as well as the Internet to provide a flexible connectivity continuum and enable global cloud computing across organizational boundaries.
- **IP routing:** Cisco NX-OS supports a wide range of IPv4 and IPv6 services and routing protocols, providing state-of-the-art implementations of the following routing protocols:
  - Open Shortest Path First (OSPF) Protocol Versions 2 (IPv4) and 3 (IPv6)
  - Intermediate System-to-Intermediate System (IS-IS) Protocol for IPv4
  - Border Gateway Protocol (BGP) for IPv4 and IPv6
  - Enhanced Interior Gateway Routing Protocol (EIGRP) for IPv4 and IPv6
  - Routing Information Protocol Version 2 (RIPv2)

    The implementations of these protocols are fully compliant with the latest standards, providing modern enhancements and parameters such as 4-byte autonomous system numbers (ASNs), while shedding unutilized older functions in favor of a lean implementation that increases feature velocity and enhances system stability. NSF graceful restart (NSF-GR) is supported by all unicast protocols. All protocols support all interface types, including Ethernet interfaces, switched virtual interfaces (SVIs) and subinterfaces, PortChannels, tunnel interfaces, and loopback interfaces. The abundant variety of routing protocols and functions is complemented by a broad collection of IP services, including the following:

    Virtual Route Forwarding (VRF) (All routing protocols and IP services are VRF aware. Note that VRF support in this context does not imply support for BGP or Multiprotocol Label Switching [MPLS] IP VPNs, as described in RFCs 2547 and 4364.)
  - Dynamic Host Configuration Protocol (DHCP) Helper
  - Unicast Reverse Path Forwarding (uRPF) for IPv4 and IPv6
  - Hot-Standby Routing Protocol (HSRP) for IPv4 and IPv6

- Virtual Router Redundancy Protocol (VRRP) for IPv4
- Gateway Load Balancing Protocol (GLBP) for IPv4
- Enhanced object tracking
- Policy-based routing (PBR) for IPv4 and IPv6
- Generic routing encapsulation (GRE) tunneling
- Unicast graceful restart for all protocols in IPv4
- Unicast graceful restart for OSPFv3 in IPv6

- **IP Multicast:** Cisco NX-OS provides an industry-leading IP Multicast feature set. The Cisco NX-OS implementation lays the foundation for the future development of a comprehensive portfolio of multicast-enabled network functions. Similar to the unicast routing protocols, Cisco NX-OS includes state-of-the-art implementations of the following multicast protocols and functions:
  - Protocol Independent Multicast Version 2 (PIMv2)
  - Source-Specific Multicast (SSM) for IPv4 and IPv6
  - PIM Sparse Mode (Any-Source Multicast [ASM] for IPv4 and IPv6)
  - Bidirectional PIM (Bidir PIM) for IPv4 and IPv6
  - Anycast Rendezvous Point (Anycast-RP)
  - Multicast NSF for IPv4 and IPv6
  - RP-Discovery using bootstrap router (BSR): Auto-RP and static
  - Internet Group Management Protocol (IGMP) Versions 1, 2, and 3 router role
  - IGMPv2 host mode
  - IGMP snooping
  - Multicast Listener Discovery (MLD) Protocol Version 2 (for IPv6)
  - Multicast Source Discovery Protocol (MSDP) (for IPv4 only)
  - IGMP cache on non-disaster recovery for fast convergence
  - Policies for multicast configuration (ip pim rp-addr and ip igmp join-group/static-group)
  - IGMP group-specific queries to router ports only
  - Debug filters for IGMP snooping
- **Data Center Bridging (DCB)** enables Ethernet fabrics to support lossless transmission to increase network scalability, support I/O consolidation, ease management of multiple traffic flows, and optimize performance. Although SAN consolidation requires only the lossless fabric provided by the Ethernet Pause mechanism, the Cisco NX-OS provides additional features that create an even more easily managed, high-performance, unified network fabric. DCBX is a protocol that simplifies network deployment and reduces configuration errors by providing autonegotiation of DCB features between the NIC and the switch and between switches.
- **Cisco TrustSec® security:** As part of the Cisco TrustSec security suite, Cisco NX-OS provides outstanding data confidentiality and integrity, supporting standard IEEE 802.1AE link-layer cryptography with 128-bit Advanced Encryption Standard (AES) cryptography. Link-layer cryptography helps ensure end-to-end data privacy while allowing the insertion of security service devices along the encrypted path. Security group access control lists (SGACLs), a new model in network access control, are based on security group tags instead of IP addresses, enabling implementation of policies that are more concise and easier to manage due to their topology independence.

  In addition to Cisco TrustSec security, Cisco NX-OS delivers the following security features:
  - Data path intrusion detection system (IDS) for protocol conformance checks

- Control-plane policing (CoPP)
- Message-digest algorithm 5 (MD5) routing protocol authentication
- Cisco integrated security features, including Dynamic Address Resolution Protocol (ARP) Inspection (DAI), DHCP snooping, and IP source guard
- Authentication, authorization, and accounting (AAA) and TACACS+
- Secure Shell (SSH) Protocol Version 2
- Simple Network Management Protocol Version 3 (SNMPv3 support)
- Port security
- IEEE 802.1x authentication and RADIUS support
- Layer 2 Cisco Network Admission Control (NAC) LAN port IP
- Policies based on MAC addresses and IPv4 and IPv6 addresses supported by named ACLs (port-based ACLs [PACLs], VLAN-based ACLs [VACLs], and router-based ACLs [RACLs])

- **Switch and host authentication:** Fibre Channel Security Protocol (FC-SP) capabilities in Cisco NX-OS provide switch-to-switch and host-to-switch authentication for enterprisewide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) is used to perform authentication locally or remotely through RADIUS or TACACS+. If authentication fails, a switch or host cannot join the fabric.
- **Port security and fabric binding:** Port security locks down the mapping of an entity to a switch port. The entities can be hosts, targets, or switches that are identified by their World Wide Names (WWNs). This locking mechanism helps ensure that unauthorized devices connecting to the switch port do not disrupt the SAN fabric. Fabric binding extends port security to allow ISLs only between specified switches.

Efficiency

Cisco NX-OS provides operation tools that reduce complexity and offer consistent features and operations without compromising capabilities.

- **Cisco Fabric Extender Link (FEX-Link) support:** Cisco FEX-Link architecture provides a highly scalable, unified server-access platform 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 traditional Gigabit Ethernet technology while allowing transparent migration to 10 Gigabit Ethernet, virtual machine-aware unified fabric technologies. The combination of Cisco FEX-Link and the Cisco Nexus parent switch enables data centers to scale the number of Gigabit Ethernet access ports, reducing cable runs and the number of management points in the network.
- **Protocol offload:** To reduce the load on the control plane of the device in a Cisco FEX-Link design, Cisco NX-OS provides the capability to offload link-level protocol processing to the fabric extender CPU. The following protocols are supported:
  - Link Layer Discovery Protocol (LLDP) and Data Center Bridging Exchange (DCBX)
  - Cisco Discovery Protocol
  - Link Aggregation Control Protocol (LACP)
- **IEEE 802.1Qbb priority-based flow control (PFC)** offers point-to-point flow control of Ethernet traffic based on IEEE 802.1p class of service (CoS). With a flow=control mechanism in place, congestion does not result in drops, transforming Ethernet into a reliable medium. A networking device implementing PFC makes an implicit agreement with the other end of the wire: any accepted packet will be delivered to the next hop and never be locally dropped. To keep this promise, the device must signal the peer when no more packets can reliably be accepted, and that, essentially, is the flow-control function performed by PFC. The benefits are

significant for any protocol that assumes reliability at the media level, such as Fibre Channel over Ethernet (FCoE).

- **Troubleshooting and diagnostics:** Cisco NX-OS is built with unique serviceability functions to enable network operators to take early action based on network trends and events, enhancing network planning and improving network operations center (NOC) and vendor response times. Smart Call Home and Cisco Generic Online Diagnostics (GOLD) are some of the features that enhance the serviceability of Cisco NX-OS.
- **Cisco Switched Port Analyzer (SPAN):** The Cisco SPAN feature allows an administrator to analyze all traffic between ports (called the SPAN source ports) by nonintrusively directing the SPAN session traffic to a SPAN destination port that has an external analyzer attached to it.
- **Cisco Fabric Analyzer:** The embedded Cisco Fabric Analyzer can save Fibre Channel control traffic inside the switch for text-based analysis or can send IP-encapsulated Fibre Channel control traffic to a remote PC for decoding and display using the open-source Ethereal network-analyzer application. Fibre Channel control traffic therefore can be captured and analyzed without an expensive Fibre Channel analyzer.
- **Etheranalyzer:** Cisco NX-OS includes a built-in packet analyzer to monitor and troubleshoot control-plane and data-plane traffic. The packet analyzer is based on the popular Wireshark open source network protocol analyzer.
- **Smart Call Home:** The Smart Call Home feature continuously monitors hardware and software components to provide email-based notification of critical system events. A versatile range of message formats is available for optimal compatibility with pager services, standard email, and XML-based automated parsing applications. It offers alert grouping capabilities and customizable destination profiles. This feature can be used, for example, to directly page a network support engineer, send an email message to a NOC, and employ Cisco AutoNotify services to directly generate a case with the Cisco Technical Assistance Center (TAC). This feature is a step toward autonomous system operation, enabling networking devices to inform IT when a problem occurs and helping ensure that the problem is acted on quickly, reducing time to resolution and increasing system uptime.
- **Cisco GOLD:** Cisco GOLD is a suite of diagnostic facilities to verify that hardware and internal data paths are operating as designed. Boot-time diagnostics, continuous monitoring, standby fabric loopback tests, and on-demand and scheduled tests are part of the Cisco GOLD feature set. This industry-leading diagnostics subsystem allows rapid fault isolation and continuous system monitoring critical in today's continuously operating environments.
- **Cisco IOS Embedded Event Manager (EEM):** Cisco IOS EEM is a powerful device and system management technology integrated into Cisco NX-OS. Cisco IOS EEM helps customers make use of the network intelligence intrinsic to the Cisco software and customize behavior based on network events as they occur.
- **Cisco NetFlow:** The Cisco NX-OS implementation of NetFlow supports Version 5 and 9 exports as well as the Flexible NetFlow configuration model and hardware-based Sampled NetFlow for enhanced scalability. In addition to Layer 3 NetFlow, Layer 2 NetFlow is supported.
- **Traffic redirection:** Cisco NX-OS supports Web Cache Control Protocol (WCCP) Version 2 in Layer 2 forwarding mode. WCCP allows the use of cache engines to localize web traffic patterns in the network, enabling content requests to be fulfilled locally. Traffic localization reduces transmission costs and download time. WCCP allows to transparently redirect content requests. The main benefit of transparent redirection is that users need not configure their browsers to use a web proxy. Instead, they can use the target URL to request content and have their requests automatically redirected to a cache engine. WCCP enables a series of cache engines, called a cache engine cluster, to provide content to a router or multiple routers. Clustering cache engines greatly improves the scalability, redundancy, and availability of the caching solution. Clustering of up to 32 cache engines per service group is supported.

- **Programmatic XML interface:** Based on the NETCONF industry standard, the Cisco NX-OS XML interface provides a consistent API for devices, enabling rapid development and creation of tools to enhance the network.
- **Simple Network Management Protocol (SNMP):** Cisco NX-OS complies with SNMPv1, v2c, and v3. A rich collection of MIBs is supported.
- **Configuration verification and rollback:** With Cisco NX-OS, the system operator can verify the consistency of a configuration and the availability of necessary hardware resources prior to committing the configuration. A device can thus be preconfigured and the verified configuration applied at a later time. Configurations also include checkpoints, to allow operators to roll back to a known good configuration as needed.
- **Configuration synchronization:** Configuration synchronization (config-sync) mode allows users to create switch profiles to synchronize local and peer switches. Config-sync allows administrators to make configuration changes on one switch and have the system automatically synchronize the switch's peers. This feature eliminates user errors and reduces the administrative overhead of having to configure both members of a vPC simultaneously.
- **Port profiles:** Port profiles enable customers to define a policy once and then apply it many times across virtual and physical ports, significantly increasing both efficiency and flexibility in today's virtual data centers.
- **Role-based access control (RBAC):** With RBAC, Cisco NX-OS enables administrators to limit access to switch operations by assigning roles to users. Administrators can customize access and restrict it to the users who require it. Cisco NX-OS also provides a mechanism for distributing the configuration of RBAC roles across devices running Cisco NX-OS, for simplified deployment.
- **Cisco Fabric Services:** Cisco NX-OS incorporates many management features that facilitate effective management of growing storage environments with existing resources. Cisco fabric services simplify SAN provisioning by automatically distributing configuration information to all switches in a storage network. Distributed device alias services provide fabricwide alias names for host bus adapters (HBAs), storage devices, and switch ports, eliminating the need to reenter names when devices are moved.
- **Cisco NPV technology:** Cisco NX-OS supports industry-standard N-port ID virtualization (NPIV), which allows multiple N-port fabric logins concurrently on a single physical Fibre Channel link. HBAs that support NPIV can help improve SAN security by enabling configuration of zoning and port security independently for each virtual machine (OS partition) on a host. In addition to being useful for server connections, NPIV is beneficial for connectivity between core and edge SAN switches.

  Cisco NPV is a complementary feature that reduces the number of Fibre Channel domain IDs in core-edge SANs. It is used by edge switches in the NPV mode to log in to multiple end devices that share a link to the core switch.
- **Autolearn feature for network security configuration:** The autolearn feature allows the Nexus switches to automatically learn about devices and switches that connect to it. The administrator can use this feature to configure and activate network security features such as port security without having to manually configure the security for each port.
- **Connectivity management processor (CMP) support:** Cisco NX-OS supports the use of a CMP for lights-out, remote management of the platform. The CMP aids operations by providing an out-of-band access channel to the Cisco NX-OS console. IPv6 support for the CMP interface is also available, including ping6 and traceroute6.
- **Fibre Channel ping and Fibre Channel traceroute:** Cisco NX-OS brings to storage networks features such as Fibre Channel ping and Fibre Channel traceroute, which are essential for IP network troubleshooting. With Fibre Channel ping, administrators can check the connectivity of an N-port and determine its round-trip

latency, and with Fibre Channel traceroute, administrators can check the reachability of a switch by tracing the path followed by frames and determining hop-by-hop latency.

Virtualization

Cisco NX-OS enhances virtual machine portability and converges multiple services, platforms, and networks to simplify and reduce infrastructure sprawl and total cost of ownership (TCO).

- **Virtual device contexts (VDCs):** Cisco NX-OS offers the capability to segment OS and hardware resources into virtual contexts that emulate virtual devices. Each VDC has its own software processes, dedicated hardware resources (physical interfaces, VLANs, routing table size, VRF instances, etc.), and independent management environment. VDCs are instrumental in the consolidation of separate networks onto a common infrastructure, maintaining the administrative boundary separation and fault isolation characteristics of physically separate networks while providing many of the operating cost benefits of a single infrastructure. Each VDC can be restarted without affecting the control, data, and management planes of other VDCs in the system.
- **Multiprotocol Label Switching (MPLS):** Cisco NX-OS supports a comprehensive set of MPLS features, including label switching, Layer 3 VPNs, MPLS Traffic Engineering with Fast Reroute (FRR), Multicast VPNs for IPv4, and IP v6 provider edge (6PE) and IPv6 VPN provider edge (6VPE). These features interoperate with Cisco IOS Software. These features provide the foundation for network consolidation and centralization of services and policy control for a securely segmented network fabric, enabling reduced capital expenditures (CapEx) and operating expenses (OpEx) for IT managers.
- **VSANs:** VSAN technology partitions a single physical SAN into multiple VSANs. Each VSAN is a logically and functionally separate SAN with its own set of Fibre Channel fabric services. This partitioning of fabric services greatly reduces network instability by containing fabric reconfigurations and error conditions within an individual VSAN. The strict traffic segregation provided by VSANs helps ensure that the control and data traffic of a given VSAN are confined within the VSAN's own domain, increasing SAN security, scalability, and resilience. VSANs help reduce costs by facilitating the consolidation of isolated SAN islands into a common infrastructure without compromising availability, security, or scalability.

  Users can create SAN administrator roles that are limited in scope to certain VSANs. For example, a SAN administrator role can be set up to allow configuration of all platform-specific capabilities, and other roles can be set up to allow configuration and management within specific VSANs only. This approach improves the manageability of large SANs and reduces disruptions resulting from human errors by isolating the effect of a SAN administrator's action to a specific VSAN whose membership can be isolated based on switch ports or WWNs of attached devices. VSANs are supported across Fibre Channel over IP (FCIP) links between SANs, extending VSANs to include devices at a remote location.

  Cisco NX-OS also implements trunking for VSANs. Trunking allows Inter-Switch Links (ISLs) to carry traffic for multiple VSANs on the same physical link. F-port trunking allows multiple VSANs on a single uplink in Cisco N-Port Virtualization (NPV) mode.
- **Cisco Adapter Fabric Extender:** The Cisco Adapter Fabric Extender extends the current benefits of the Cisco FEX-Link architecture to the server network interface cards (NICs), providing architecture flexibility and high scalability with 4000 logical interfaces with a single point of management and policy enforcement. The adapter fabric extender is logically an extension of the parent switch inside the server.

  Interfaces of an adapter fabric extender are local logical ports on the parent switch. An adapter fabric extender uses an innovative server connectivity (I/O connectivity) technology that enables on-demand creation of virtual NICs (vNICs) on a single NIC. With the adapter fabric extender, a single physical adapter is presented as multiple logical adapters to the server OS and the network as if they were multiple physical adapters. A dual-port 10 Gigabit Ethernet adapter fabric extender can support hundreds of Peripheral

Component Interconnect Express (PICe) standards-compliant virtual interfaces that can be configured as needed by the server administrator.

Product Specifications

Supported Standards

Tables 2 and 3 provide standards compliance information for Cisco NX-OS.

Table 2. IEEE Compliance

Standard	Description
IEEE 802.1D	MAC bridges
IEEE 802.1s	Multiple Spanning Tree Protocol
IEEE 802.1w	Rapid Spanning Tree Protocol
IEEE 802.1ab	LLDP
IEEE 802.1AE	MAC security (link-layer cryptography)
IEEE 802.3ad	Link aggregation with LACP
IEEE 802.3ab	1000BASE-T (10/100/1000 Ethernet over copper)
IEEE 802.3z	Gigabit Ethernet
IEEE 802.3ae	10 Gigabit Ethernet
IEEE 802.1Q	VLAN tagging
IEEE 802.1p	CoS tagging for Ethernet frames
IEEE 802.1x	Port-based network access control
Fibre Channel Standards	
T11 FC-BB-5	Fibre Channel over Ethernet (FCoE)

Table 3. RFC Compliance

Standard	Description
BGP	
RFC 1997	BGP Communities Attribute
RFC 2385	Protection of BGP Sessions with the TCP MD5 Signature Option
RFC 2439	BGP Route Flap Damping
RFC 2519	Framework for Inter-Domain Route Aggregation
RFC 2545	Use of BGPv4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
RFC 2858	Multiprotocol Extensions for BGPv4
RFC 3065	Autonomous System Confederations for BGP
RFC 3392	Capabilities Advertisement with BGPv4
RFC 4271	BGPv4
RFC 4273	BGPv4 MIB: Definitions of Managed Objects for BGPv4
RFC 4456	BGP Route Reflection
RFC 4486	Subcodes for BGP Cease Notification Message
RFC 4724	Graceful Restart Mechanism for BGP
RFC 4893	BGP Support for Four-Octet as Number Space
ietf-draft	Bestpath Transition Avoidance (draft-ietf-idr-avoid-transition-05.txt)
ietf-draft	Peer Table Objects (draft-ietf-idr-bgp4-mib-15.txt)
ietf-draft	Dynamic Capability (draft-ietf-idr-dynamic-cap-03.txt)
OSPF	
RFC 2370	OSPF Opaque LSA Option

Standard	Description
RFC 2328	OSPF Version 2
RFC 2740	OSPF for IPv6 (OSPFv3)
RFC 3101	OSPF Not-So-Stubby-Area (NSSA) Option
RFC 3137	OSPF Stub Router Advertisement
RFC 3509	Alternative Implementations of OSPF Area Border Routers
RFC 3623	Graceful OSPF Restart
RFC 4750	OSPF Version 2 MIB
RIP	
RFC 1724	RIPv2 MIB Extension
RFC 2082	RIPv2 MD5 Authentication
RFC 2453	RIP Version 2
IS-IS	
RFC 1142 (OSI 10589)	OSI 10589 Intermediate System-to-Intermediate System (IS-IS) Intradomain Routing Exchange Protocol
RFC 1195	Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
RFC 2763	Dynamic Hostname Exchange Mechanism for IS-IS
RFC 2966	Domainwide Prefix Distribution with Two-Level IS-IS
RFC 2973	IS-IS Mesh Groups
RFC 3277	IS-IS Transient Black-Hole Avoidance
RFC 3373	Three-Way Handshake for IS-IS Point-to-Point Adjacencies
RFC 3567	IS-IS Cryptographic Authentication
RFC 3847	Restart Signaling for IS-IS
ietf-draft	Internet Draft Point-to-Point Operation over LAN in Link-State Routing Protocols (draft-ietf-isis-igp-p2p-over-lan-06.txt)
IP Services	
RFC 768	User Datagram Protocol (UDP)
RFC 783	Trivial File Transfer Protocol (TFTP)
RFC 791	IP
RFC 792	Internet Control Message Protocol (ICMP)
RFC 793	TCP
RFC 826	ARP
RFC 854	Telnet
RFC 959	FTP
RFC 1027	Proxy ARP
RFC 1305	Network Time Protocol (NTP) Version 3
RFC 1519	Classless Interdomain Routing (CIDR)
RFC 1542	BootP Relay
RFC 1591	Domain Name System (DNS) Client
RFC 1812	IPv4 Routers
RFC 2131	DHCP Helper
RFC 2338	VRRP
RFC 2784	GRE
IP Multicast	
RFC 2236	IGMPv2
RFC 2710	Multicast Listener Discovery (MLD) for IPv6
RFC 3376	IGMPv3
RFC 3446	Anycast Rendezvous Point Mechanism Using PIM and MSDP

Standard	Description
RFC 3569	Overview of SSM
RFC 3618	MSDP
RFC 3810	MLDPv2 for IPv6
RFC 4601	Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)
RFC 4607	Source-Specific Multicast for IP
RFC 4610	Anycast-RP using PIM
RFC 5132	IP Multicast MIB
ietf-draft	Traceroute Facility for IP Multicast (draft-ietf-idmr-traceroute-ipm-07.txt)
ietf-draft	Bidirectional Protocol Independent Multicast (BIDIR-PIM, draft-ietf-pim-bidir-09.txt)
Ietf-draft	Bidirectional Forwarding Detection
OTV	
Ietf-draft	Overlay Transport Virtualization (draft-hasmit-otv-00)
MPLS	
RFC 3031	MPLS Architecture
RFC 3032	MPLS Label Stack-Encoding
RFC 3036	LDP Specification
RFC 3478	Graceful Restart Mechanism for Label Distribution Protocol
RFC 3812	Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)
RFC 3813	Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)
RFC 4382	MPLS/BGP Layer 3 Virtual Private Network (VPN) Management Information Base
RFC 3815	Definitions of Managed Objects for Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP)
IETF DRAFT	draft-ietf-mpls-fastreroute-mib: Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base for Fast Reroute
RFC 5036	LDP Specification (obsoletes RFC3036): Partial Support
RFC 5443	LDP IGP Synchronization
IETF DRAFT	LDP Capabilities (draft-ietf-mpls-ldp-capabilties-04.txt draft)
IETF DRAFT	LDP Typed Wildcard FEC (draft-ietf-mpls-ldp-typed-wildcard-03.txt)
RFC 2685	Virtual Private Networks Identifier
RFC 2858	Multiprotocol Extensions for BGP-4
RFC 3107	Carrying Label Information in BGP-4
RFC 3630	Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 4364	BGP or MPLS IP VPNs (No InterAS support)
RFC 4365	Applicability Statement for BGP or MPLS IP VPNs
RFC 4382	MPLS or BGP Layer 3 VPN MIB
RFC 4576	Using LSA Options Bit to Prevent Looping in BGP or MPLS IP VPNs (DN Bit)
RFC 4577	OSPF as the PE or CE Protocol in BGP or MPLS IP VPNs
RFC 4659	BGP-MPLS IP VPN Extension for IPv6 VPN (No InterAS support)
RFC 4760	Multi-protocol Extensions for BGP-4
RFC 4781	Graceful Restart Mechanism for BGP with MPLS
RFC 5305	IS-IS Extensions for Traffic Engineering
RFC 5307	IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)
IETF DRAFT	BGP Custom Decision Process
RFC 2205	RSVP v1 Functional Specification
RFC 2209	RSVP v1 Message Processing Rules
RFC 2702	TE over MPLS
RFC 2747	RSVP Cryptographic Authentication

Standard	Description
RFC 2961	RSVP Refresh Overhead Reduction Extensions
RFC 3209	RSVP-TE
RFC 3270	MPLS Support of Differentiated Services
RFC 3784	ISIS-TE
RFC 4090	Fast Re-Route for RSVP-TE Extensions
RFC 4569	BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
RFC 4798	Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
LISP	
IETF DRAFT	LISP Canonical Address Format (LCAF)
IETF DRAFT	Locator/ID Separation Protocol (LISP)
IETF DRAFT	Interworking LISP with IPv4 and IPv6
IETF DRAFT	LISP Map-Versioning
IETF DRAFT	LISP Map Server
IETF DRAFT	LISP for Multicast Environments

Cisco Services

Cisco offers a wide range of services to help accelerate your success in deploying and optimizing Cisco NX-OS and Nexus Switches in your data center. Cisco's innovative services are delivered through a unique combination of people, processes, tools, and partners and are focused on helping you increase operation efficiency and improve your data center network. Cisco Advanced Services uses an architecture-led approach to help you align your data center infrastructure with your business goals and achieve long-term value. Cisco SMARTnet® Service helps you resolve mission-critical problems with direct access at any time to Cisco network experts and award-winning resources. With this service, you can take advantage of the Smart Call Home service capability, which offers proactive diagnostics and real-time alerts on your Cisco Nexus Switches. Spanning the entire network lifecycle, Cisco Services helps protect your investment, optimize network operations, support migration, and strengthen your IT expertise. For more information about Cisco Data Center Services, visit http://www.cisco.com/go/dcservices.

For More Information

Cisco NX-OS Software: http://www.cisco.com/go/nxos

Cisco NX-OS Licensing: http://www.cisco.com/en/US/products/ps9402/products_licensing_information_listing.html

Cisco Nexus 7000 Series Switches: http://www.cisco.com/go/nexus7000

Cisco Nexus 5000 Series Switches: http://www.cisco.com/go/nexus5000

Cisco Nexus 2000 Series Fabric Extenders: http://www.cisco.com/go/nexus2000

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 C78-652063-00 03/11