Overview

This chapter provides an overview of the Cisco NX-OS devices that support Layer 2 features.

This chapter includes the following sections:

Layer 2 Ethernet Switching Overview

The device supports simultaneous, parallel connections between Layer 2 Ethernet segments. Switched connections between Ethernet segments last only for the duration of the packet. New connections can be made between different segments for the next packet.

The device solves congestion problems caused by high-bandwidth devices and a large number of users by assigning each device (for example, a server) to its own domain. Because each LAN port connects to a separate Ethernet collision domain, servers in a switched environment achieve full access to the bandwidth.

Because collisions cause significant congestion in Ethernet networks, an effective solution is full-duplex communication. Typically, 10/100-Mbps Ethernet operates in half-duplex mode, which means that stations can either receive or transmit. In full-duplex mode, which is configurable on these interfaces, two stations can transmit and receive at the same time. When packets can flow in both directions simultaneously, the effective Ethernet bandwidth doubles. 1/10-Gigabit Ethernet operates in full duplex only.

VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment.

Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Each VLAN is considered as a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a bridge or a router.

All ports are assigned to the default VLAN (VLAN1) when the device first comes up. A VLAN interface, or switched virtual interface (SVI), is a Layer 3 interface that is created to provide communication between VLANs.

The devices support 4094 VLANs in accordance with the IEEE 802.1Q standard. These VLANs are organized into several ranges, and you use each range slightly differently. Some of these VLANs are reserved for internal use by the device and are not available for configuration.

Note Inter-Switch Link (ISL) trunking is not supported on the Cisco NX-OS.

Private VLANs

Private VLANs provide traffic separation and security at the Layer 2 level.

A private VLAN is one or more pairs of a primary VLAN and a secondary VLAN, all with the same primary VLAN. The two types of secondary VLANs are isolated and community VLANs. Hosts on isolated VLANs communicate only with hosts in the primary VLAN. Hosts in a community VLAN can communicate only among themselves and with hosts in the primary VLAN but not with hosts in isolated VLANs or in other community VLANs.

Regardless of the combination of isolated and community secondary VLANs, all interfaces within the primary VLAN comprise one Layer 2 domain, and therefore, require only one IP subnet.

Spanning Tree

This section discusses the implementation of the Spanning Tree Protocol (STP) on the software. Spanning tree is used to refer to IEEE 802.1w and IEEE 802.1s. When the IEEE 802.1D Spanning Tree Protocol is referred to in the publication, 802.1D is stated specifically.

STP Overview

STP provides a loop-free network at the Layer 2 level. Layer 2 LAN ports send and receive STP frames, which are called Bridge Protocol Data Units (BPDUs), at regular intervals. Network devices do not forward these frames but use the frames to construct a loop-free path.

802.1D is the original standard for STP, and many improvements have enhanced the basic loop-free STP. You can create a separate loop-free path for each VLAN, which is named Per VLAN Spanning Tree (PVST+). Additionally, the entire standard was reworked to make the loop-free convergence process faster to keep up with the faster equipment. This STP standard with faster convergence is the 802.1w standard, which is known as Rapid Spanning Tree (RSTP). Now, these faster convergence times are available as you create STP for each VLAN, which is known as Per VLAN Rapid Spanning Tree (Rapid PVST+).

Finally, the 802.1s standard, Multiple Spanning Tree (MST), allows you to map multiple VLANs into a single spanning tree instance. Each instance runs an independent spanning tree topology.

Although the software can interoperate with legacy 802.1D systems, the system runs Rapid PVST+ and MST. You can use either Rapid PVST+ or MST in a given VDC; you cannot mix both in one VDC. Rapid PVST+ is the default STP protocol for Cisco NX-OS devices.

Note Cisco NX-OS uses the extended system ID and MAC address reduction; you cannot disable these features.

In addition, Cisco has created some proprietary features to enhance the spanning tree activities.

Rapid PVST+

Rapid PVST+ is the default spanning tree mode for the software and is enabled by default on the default VLAN and all newly created VLANs.

A single instance, or topology, of RSTP runs on each configured VLAN, and each Rapid PVST+ instance on a VLAN has a single root device. You can enable and disable STP on a per-VLAN basis when you are running Rapid PVST+.

MST

The software also supports MST. The multiple independent spanning tree topologies enabled by MST provide multiple forwarding paths for data traffic, enable load balancing, and reduce the number of STP instances required to support a large number of VLANs.

MST incorporates RSTP, so it also allows rapid convergence. MST improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).

Note Changing the spanning tree mode disrupts the traffic because all spanning tree instances are stopped for the previous mode and started for the new mode.

You can force specified interfaces to send prestandard, rather than standard, MST messages using the command-line interface.

STP Extensions

The software supports the following Cisco proprietary features:

- Spanning tree port types—The default spanning tree port type is normal. You can configure interfaces connected to Layer 2 hosts as edge ports and interfaces connected to Layer 2 switches or bridges as network ports.
- Bridge Assurance—Once you configure a port as a network port, Bridge Assurance sends BPDUs on all ports and moves a port into the blocking state if it no longer receives BPDUs. This enhancement is available only when you are running Rapid PVST+ or MST.
- BPDU Guard—BPDU Guard shuts down the port if that port receives a BPDU.
- BPDU Filter—BPDU Filter suppresses sending and receiving BPDUs on the port.
- Loop Guard—Loop guard helps prevent bridging loops that could occur because of a unidirectional link failure on a point-to-point link.

- Root Guard—The root guard feature prevents a port from becoming root port or blocked port. If a port configured for root guard receives a superior BPDU, the port immediately goes to the root-inconsistent (blocked) state.

Virtualization

Cisco NX-OS devices introduce support for multiple virtual device contexts (VDCs) on a single switching device. Each VDC is treated as a standalone device with specific resources, such as physical interfaces, allocated to each VDC by the network admin role. An administrator is assigned to each VDC and that administrator has a limited view of the system within that specific VDC. Faults are also isolated to within the specific VDC.

This VDC concept applies to all features on Cisco NX-OS, including all Layer 2 switching features.

Figure 1: VDCs with Layer 2 Services

All processes work independently in each VDC. You can reuse the process identification numbers in different VDCs. This figure shows how to reuse the VLAN 100 identifier in each separate VDC.

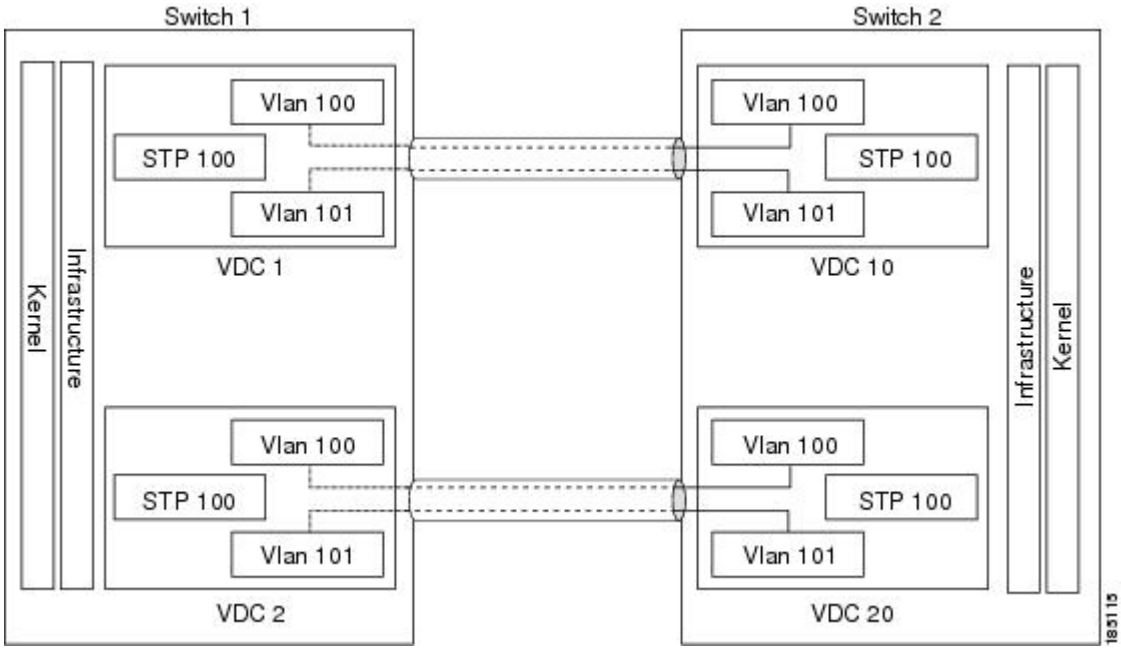

Each VDC acts as a standalone device with Layer 2 services available. VDCs allow you to share a physical device among several logical functions. You can provision and assign entirely separate Layer 2 resources to individual VDCs.

You can configure several VDCs, and each VDC is a group of physical device resources. You assign resources and user roles for each VDC. VDCs allows flexible resources as well as enhanced fault isolation.

VDCs allow the separation of processes and management environments, providing well-defined fault and administrative boundaries between logical devices. Each VDC can be considered as a separate device with its own configuration, resources, users, and management interface.

VDCs define different administrator levels, or roles, that can access and administer each VDC. Commands outside the scope of a given user role are either hidden from that user's view or can return an error if the

command is entered. This feature limits the number of users who can access the entire physical device and introduce traffic-disrupting misconfigurations.

Note See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for complete information on virtual device contexts (VDCs) and assigning resources.

See the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide* for information on restartability and seamless transitions.

Related Topics

The following documents are related to the Layer 2 switching features:

- *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*
- *Cisco DCNM Layer 2 Switching Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*
- *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*
- *Cisco NX-OS Licensing Guide*
- *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide*
- *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide*

Related Topics