



Product Overview

This chapter provides an overview of Catalyst 4500 series switches and includes the following major sections:

- [Layer 2 Software Features, page 1-1](#)
- [Layer 3 Software Features, page 1-5](#)
- [QoS Features, page 1-9](#)
- [Management and Security Features, page 1-9](#)



Note

For more information about the chassis, modules, and software features supported by the Catalyst 4500 series switch, refer to the *Release Notes for the Catalyst 4500 Series Switch, 12.1(19)EW* at http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/release/note/OL_2170.html

Layer 2 Software Features

The following subsections describe the key Layer 2 switching software features on the Catalyst 4500 series switch:

- [Storm Control, page 1-2](#)
- [CDP, page 1-2](#)
- [DHCP Snooping, page 1-2](#)
- [EtherChannel Bundles, page 1-2](#)
- [IP Source Guard, page 1-3](#)
- [Jumbo Frames, page 1-3](#)
- [Layer 2 Traceroute, page 1-3](#)
- [MST, page 1-3](#)
- [PVRST+, page 1-3](#)
- [Spanning Tree Protocol, page 1-4](#)
- [UDLD, page 1-4](#)
- [Unidirectional Ethernet, page 1-4](#)
- [VLANs, page 1-4](#)

Storm Control

Broadcast suppression is used to prevent LANs from being disrupted by a broadcast storm on one or more switch ports. A LAN broadcast storm occurs when broadcast packets flood the LAN, creating excessive traffic and degrading network performance. Errors in the protocol-stack implementation or in the network configuration can cause a broadcast storm. Broadcast suppression measures how much broadcast traffic is passing through a port and compares the broadcast traffic with some configurable threshold value within a specific time interval. If the amount of broadcast traffic reaches the threshold during this interval, broadcast frames are dropped, and optionally the port is shut down.

For information on configuring broadcast suppression, see [Chapter 28, “Configuring Port-Based Traffic Control.”](#)

CDP

The Cisco Discovery Protocol (CDP) is a device-discovery protocol that is both media- and protocol-independent. CDP is available on all Cisco products, including routers, switches, bridges, and access servers. Using CDP, a device can advertise its existence to other devices and receive information about other devices on the same LAN. CDP enables Cisco switches and routers to exchange information, such as their MAC addresses, IP addresses, and outgoing interfaces. CDP runs over the data-link layer only, allowing two systems that support different network-layer protocols to learn about each other. Each device configured for CDP sends periodic messages to a multicast address. Each device advertises at least one address at which it can receive Simple Network Management Protocol (SNMP) messages.

For information on configuring CDP, see [Chapter 16, “Understanding and Configuring CDP.”](#)

DHCP Snooping

Dynamic Host Configuration Protocol (DHCP) snooping is a security feature that is a component of a DHCP server. DHCP snooping provides security by intercepting untrusted DHCP messages and by building and maintaining a DHCP snooping binding table. An untrusted message is a message that is received from outside the network or firewall that can cause traffic attacks within your network.

DHCP snooping acts like a firewall between untrusted hosts and DHCP servers. It also provides a way to differentiate between untrusted interfaces connected to the end-user and trusted interfaces connected to the DHCP server or another switch.

For DHCP server configuration information, refer to the chapter, “Configuring DHCP,” in the *Cisco IOS IP and IP Routing Configuration Guide* at the following URL:

http://www.cisco.com/en/US/docs/ios/ipaddr/configuration/guide/iad_dhcp_rdmf_ps6350_TSD_Products_Configuration_Guide_Chapter.html

For information on configuring DHCP snooping, see [Chapter 19, “Configuring DHCP Snooping and IP Source Guard.”](#)

EtherChannel Bundles

EtherChannel port bundles allow you to create high-bandwidth connections between two switches by grouping multiple ports into a single logical transmission path.

For information on configuring EtherChannel, see [Chapter 17, “Understanding and Configuring EtherChannel.”](#)

IP Source Guard

Similar to DHCP snooping, this feature is enabled on an untrusted 12 port that is configured for DHCP snooping. Initially all IP traffic on the port is blocked except for the DHCP packets, which are captured by the DHCP snooping process. When a client receives a valid IP address from the DHCP server, a PVACL is installed on the port, which restricts the client IP traffic only to clients with assigned IP addresses; so any IP traffic with source IP addresses other than those assigned by the DHCP server will be filtered out. This filtering prevents a malicious host from attacking a network by hijacking neighbor host's IP address.

For information on configuring IP Source Guard, see [Chapter 19, “Configuring DHCP Snooping and IP Source Guard.”](#)

Jumbo Frames

The jumbo frames feature allows the switch to forward packets as large as 9216 bytes (larger than the IEEE Ethernet MTU), rather than declare those frames “oversize” and discard them. This feature is typically used for large data transfers. The jumbo feature can be configured on a per-port basis on Layer 2 and Layer 3 interfaces and is supported only on non-blocking GB front ports.

For information on Jumbo Frames, see [Chapter 4, “Configuring Interfaces.”](#)

Layer 2 Traceroute

The Layer 2 traceroute feature allows the switch to identify the physical path that a packet takes from a source device to a destination device. Layer 2 traceroute supports only unicast source and destination MAC addresses.

For information about Layer 2 Traceroute, see [Chapter 5, “Checking Port Status and Connectivity.”](#)

MST

IEEE 802.1s Multiple Spanning Tree (MST) allows for multiple spanning tree instantiations within a single 802.1q or Inter-Switch Link (ISL) VLAN trunk. MST extends the IEEE 802.1w Rapid Spanning Tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment.

MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This new architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).

For information on configuring MST, see [Chapter 14, “Understanding and Configuring Multiple Spanning Trees.”](#)

PVRST+

Per-VLAN Rapid Spanning Tree (PVRST+) is the implementation of 802.1w on a per-VLAN basis. It is the same as PVST+ with respect to STP mode and runs RSTP protocol based on 802.1w.

For information on configuring PVRST+, see [Chapter 12, “Understanding and Configuring STP.”](#)

Spanning Tree Protocol

The Spanning Tree Protocol (STP) allows you to create fault-tolerant internetworks that ensure an active, loop-free data path between all nodes in the network. STP uses an algorithm to calculate the best loop-free path throughout a switched network.

For information on configuring STP, see [Chapter 12, “Understanding and Configuring STP.”](#)

The Catalyst 4500 series switch supports the following STP enhancements:

- Spanning tree PortFast—PortFast allows a port with a directly attached host to transition to the forwarding state directly, bypassing the listening and learning states.
- Spanning tree UplinkFast—UplinkFast provides fast convergence after a spanning-tree topology change and achieves load balancing between redundant links using uplink groups. Uplink groups provide an alternate path in case the currently forwarding link fails. UplinkFast is designed to decrease spanning-tree convergence time for switches that experience a direct link failure.
- Spanning tree BackboneFast—BackboneFast reduces the time needed for the spanning tree to converge after a topology change caused by an indirect link failure. BackboneFast decreases spanning-tree convergence time for any switch that experiences an indirect link failure.
- Spanning tree root guard—Root guard forces a port to become a designated port so that no switch on the other end of the link can become a root switch.

For information on the STP enhancements, see [Chapter 13, “Configuring STP Features.”](#)

UDLD

The UniDirectional Link Detection (UDLD) protocol allows devices connected through fiber-optic or copper Ethernet cables to monitor the physical configuration of the cables and detect a unidirectional link.

For information about UDLD, see [Chapter 18, “Configuring UDLD.”](#)

Unidirectional Ethernet

Unidirectional Ethernet uses only one strand of fiber for either transmitting or receiving one-way traffic for the Gigaport, instead of two strands of fiber for a full-duplex Gigaport Ethernet.

For information about Unidirectional Ethernet, see [Chapter 7, “Configuring Unidirectional Ethernet.”](#)

VLANs

A VLAN configures switches and routers according to logical, rather than physical, topologies. Using VLANs, a network administrator can combine any collection of LAN segments within an internetwork into an autonomous user group, such that the segments appear as a single LAN in the network. VLANs logically segment the network into different broadcast domains so that packets are switched only between ports within the VLAN. Typically, a VLAN corresponds to a particular subnet, although not necessarily.

For more information about VLANs, see [Chapter 8, “Understanding and Configuring VLANs.”](#)

The following VLAN-related features are also supported.

- **VLAN Trunking Protocol (VTP)**—VTP maintains VLAN naming consistency and connectivity between all devices in the VTP management domain. You can have redundancy in a domain by using multiple VTP servers, through which you can maintain and modify the global VLAN information. Only a few VTP servers are required in a large network.

For more information about VTP, see [Chapter 11, “Understanding and Configuring VTP.”](#)

- **Private VLANs**—Private VLANs are sets of ports that have the features of normal VLANs and also provide some Layer 2 isolation from other ports on the switch.

For information about private VLANs, see [Chapter 10, “Configuring Private VLANs.”](#)

- **Private VLAN Trunk Ports**—Private VLAN trunk ports allow a secondary port on a private VLAN to carry multiple secondary VLANs.
- **Dynamic VLAN Membership**—Dynamic VLAN Membership allows you to assign switch ports to VLANs dynamically, based on the source Media Access Control (MAC) address of the device connected to the port. When you move a host from a port on one switch in the network to a port on another switch in the network, that switch dynamically assigns the new port to the proper VLAN for that host.

For more information about Dynamic VLAN Membership, see [Chapter 9, “Configuring Dynamic VLAN Membership.”](#)

Layer 3 Software Features

A Layer 3 switch is a high-performance switch that has been optimized for a campus LAN or intranet and that provides both wirespeed Ethernet routing and switching services. Layer 3 switching improves network performance with two software functions—route processing and intelligent network services.

Compared to conventional software-based switches, Layer 3 switches process more packets faster; they do so by using application-specific integrated circuit (ASIC) hardware instead of microprocessor-based engines.

The following subsections describe the key Layer 3 switching software features on the Catalyst 4000 family switch:

- [CEF, page 1-6](#)
- [HSRP, page 1-6](#)
- [IP Routing Protocols, page 1-6](#)
- [Multicast Services, page 1-8](#)
- [Network Security with ACLs, page 1-8](#)
- [Policy-Based Routing, page 1-9](#)
- [Unidirectional Link Routing, page 1-9](#)

CEF

Cisco Express Forwarding (CEF) is an advanced Layer 3 IP-switching technology. CEF optimizes network performance and scalability in networks with large and dynamic traffic patterns, such as the Internet, and on networks that use intensive web-based applications, or interactive sessions. Although you can use CEF in any part of a network, it is designed for high-performance, highly resilient Layer 3 IP-backbone switching.

For information on configuring CEF, see [Chapter 21, “Configuring Cisco Express Forwarding.”](#)

HSRP

The Hot Standby Router Protocol (HSRP) provides high network availability by routing IP traffic from hosts on Ethernet networks without relying on the availability of any single Layer 3 switch. This feature is particularly useful for hosts that do not support a router discovery protocol and do not have the functionality to switch to a new router when their selected router reloads or loses power.

For information on configuring HSRP, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/ipapp/configuration/guide/ipapp_hsrp_ps6350_TSD_Products_Configuration_Guide_Chapter.html

IP Routing Protocols

The following routing protocols are supported on the Catalyst 4000 family switch:

- [RIP](#)
- [OSPF](#)
- [IS-IS](#)
- [IGRP](#)
- [EIGRP](#)
- [BGP](#)

RIP

The Routing Information Protocol (RIP) is a distance-vector, intradomain routing protocol. RIP works well in small, homogeneous networks. In large, complex internetworks, it has many limitations, such as a maximum hop count of 15, lack of support for variable-length subnet masks (VLSMs), inefficient use of bandwidth, and slow convergence. (RIP II does support VLSMs.)

OSPF

The Open Shortest Path First (OSPF) protocol is a standards-based IP routing protocol designed to overcome the limitations of RIP. Because OSPF is a link-state routing protocol, it sends link-state advertisements (LSAs) to all other routers within the same hierarchical area. Information on the attached interfaces and their metrics is used in OSPF LSAs. As routers accumulate link-state information, they use the shortest path first (SPF) algorithm to calculate the shortest path to each node. Additional OSPF features include equal-cost multipath routing and routing based on the upper-layer type of service (ToS) requests.

OSPF employs the concept of an area, which is a group of contiguous OSPF networks and hosts. OSPF areas are logical subdivisions of OSPF autonomous systems in which the internal topology is hidden from routers outside the area. Areas allow an additional level of hierarchy different from that provided by IP network classes, and they can be used to aggregate routing information and mask the details of a network. These features make OSPF particularly scalable for large networks.

IS-IS

The IS-IS protocol uses a link-state routing algorithm. It closely follows the Open Shortest Path First (OSPF) routing protocol used within the TCP/IP environment. The operation of ISO IS-IS requires each router to maintain a full topology map of the network (that is, which ISs and ESs are connected to which other ISs and ESs). Periodically, the router runs an algorithm over its map to calculate the shortest path to all possible destinations.

IS-IS is a two-level hierarchy. Intermediate Systems (or routers) are classified as Level 1 and Level 2. Level 1 ISs deal with a single routing area. Traffic is relayed only within their area. Any other internetwork traffic is sent to nearest Level 2 ISs, which also acts as a Level 1 ISs. Level 2 ISs move traffic between different routing areas within the same domain.

An IS-IS with multiarea support allows multiple Level 1 areas within in a single IS, thus allowing an IS to be in multiple areas. A single Level 2 area is used as backbone for interarea traffic.

Only Ethernet frames are supported. The IS-IS does not support IPX.

IGRP

The Interior Gateway Routing Protocol (IGRP) is a robust distance-vector Interior Gateway Protocol (IGP) developed by Cisco to provide for routing within an autonomous system (AS). Distance vector routing protocols request that a switch send all or a portion of its routing table data in a routing update message at regular intervals to each of its neighboring routers. As routing information proliferates through the network, routers can calculate distances to all nodes within the internetwork. IGRP uses a combination of metrics: internetwork delay, bandwidth, reliability, and load are all factored into the routing decision.

EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a version of IGRP that combines the advantages of link-state protocols with distance-vector protocols. EIGRP incorporates the Diffusing Update Algorithm (DUAL). EIGRP includes fast convergence, variable-length subnet masks, partially bounded updates, and multiple network-layer support. When a network topology change occurs, EIGRP checks its topology table for a suitable new route to the destination. If such a route exists in the table, EIGRP updates the routing table instantly. You can use the fast convergence and partial updates that EIGRP provides to route Internetwork Packet Exchange (IPX) packets.

EIGRP saves bandwidth by sending routing updates only when routing information changes. The updates contain information only about the link that changed, not the entire routing table. EIGRP also takes into consideration the available bandwidth when determining the rate at which it transmits updates.

**Note**

Layer 3 switching does not support the Next Hop Resolution Protocol (NHRP).

BGP

The Border Gateway Protocol (BGP) is an exterior gateway protocol that allows you to set up an interdomain routing system to automatically guarantee the loop-free exchange of routing information between autonomous systems. In BGP, each route consists of a network number, a list of autonomous systems that information has passed through (called the autonomous system path), and a list of other path attributes.

The Catalyst 4500 series switch supports BGP version 4, including classless interdomain routing (CIDR). CIDR lets you reduce the size of your routing tables by creating aggregate routes, resulting in supernets. CIDR eliminates the concept of network classes within BGP and supports the advertising of IP prefixes. CIDR routes can be carried by OSPF, EIGRP, and RIP.

For BGP configuration information, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/iproute/configuration/guide/irp_bgp_overview_ps6350_TSD_Products_Configuration_Guide_Chapter.html

Multicast Services

Multicast services save bandwidth by forcing the network to replicate packets only when necessary and by allowing hosts to join and leave groups dynamically. The following multicast services are supported:

- Cisco Group Management Protocol (CGMP) server—CGMP server manages multicast traffic. Multicast traffic is forwarded only to ports with attached hosts that request the multicast traffic.
- Internet Group Management Protocol (IGMP) snooping—IGMP snooping manages multicast traffic. The switch software examines IP multicast packets and makes forwarding decisions based on their content. Multicast traffic is forwarded only to ports with attached hosts that request multicast traffic. Support for IGMPv3 allows a user to establish a reporting interest in a specific channel rather than just a multicast group.
- Protocol Independent Multicast (PIM)—PIM is protocol-independent because it can leverage whichever unicast routing protocol is used to populate the unicast routing table, including EIGRP, OSPF, BGP, or static route. PIM also uses a unicast routing table to perform the Reverse Path Forwarding (RPF) check function instead of building a completely independent multicast routing table.

For information on configuring multicast services, see [Chapter 23, “Understanding and Configuring IP Multicast.”](#)

Network Security with ACLs

An access control list (ACL) filters network traffic by controlling whether routed packets are forwarded or blocked at the router interfaces. The Catalyst 4000 family switch examines each packet to determine whether to forward or drop the packet, based on the criteria you specified within the access lists.

MAC access control lists (MACLs) and VLAN access control lists (VACLs) are supported. VACLs are also known as VLAN maps in Cisco IOS.

The following security features are supported:

- MAC address filtering, which enables you to block unicast traffic for a MAC address on a VLAN interface.
- Port ACLs, which enable you to apply ACLs to Layer 2 interfaces on a switch for inbound traffic.

For information on ACLs, MACs, VLAN maps, MAC address filtering, and Port ACLs, see [Chapter 24](#), “Configuring Network Security with ACLs.”

Policy-Based Routing

Traditional IP forwarding decisions are based purely on the destination IP address of the packet being forwarded. Policy Based Routing (PBR) enables forwarding based upon other information associated with a packet, such as the source interface, IP source address, Layer 4 ports, etc. This feature allows network managers more flexibility in how they configure and design their networks.

For more information on policy-based routing, see [Chapter 22](#), “Configuring Policy-Based Routing.”

Unidirectional Link Routing

Unidirectional link routing (UDLR) provides a way to forward multicast packets over a physical unidirectional interface (such as a satellite link of high bandwidth) to stub networks that have a back channel.

For information on configuring unidirectional link routing, refer to the chapter “Configuring Unidirectional Link Routing” in the *Cisco IP and IP Routing Configuration Guide*.

QoS Features

The quality of service (QoS) features prevent congestion by selecting network traffic and prioritizing it according to its relative importance. Implementing QoS in your network makes network performance more predictable and bandwidth use more effective.

The Catalyst 4500 series switch supports the following QoS features:

- Classification and marking
- Ingress and egress policing
- Sharing and shaping

Catalyst 4500 series switch supports trusted boundary, which uses the Cisco Discovery Protocol (CDP) to detect the presence of a Cisco IP phone (such as the Cisco IP Phone 7910, 7935, 7940, and 7960) on a switch port. If the telephone is not detected, the trusted boundary feature disables the trusted setting on the switch port and prevents misuse of a high-priority queue.

The Catalyst 4500 series switch also supports QoS Automation (Auto QoS), which simplifies the deployment of existing QoS features via automatic configuration.

For information on QoS and Auto QoS, see [Chapter 29](#), “Configuring QoS.”

Management and Security Features

The Catalyst 4500 series switch offers network management and control through the CLI or through alternative access methods, such as SNMP. The switch software supports these network management and security features:

- 802.1x protocol—This feature provides a means for a host connected to a switch port to be authenticated before it is given access to the switch services.

- 802.1x with VLAN assignment—This feature allows you to enable non-802.1x capable hosts to access networks that use 802.1x authentication.
- 802.1x authentication for guest VLANs—This feature allows you to use VLAN assignment to limit network access for certain users.
- Dynamic ARP inspection—This feature intercepts all ARP requests, replies on untrusted ports, and verifies each intercepted packet for valid IP to MAC bindings. Dynamic ARP Inspection helps to prevent attacks on a network by not relaying invalid ARP replies out to other ports in the same VLAN. Denied ARP packets are logged by the switch for auditing.
- Password-protected access (read-only and read-write)—This feature protects management interfaces against unauthorized configuration changes.
- Flood Blocking— This feature enables users to disable the flooding of unicast and multicast packets on a per-port basis. Occasionally, unknown unicast or multicast traffic from an unprotected port is flooded to a protected port because a MAC address has timed out or has not been learned by the switch.
- Port Security—This feature restricts traffic on a port based upon the MAC address of the workstation that accesses the port.
- Local Authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+) authentication—These authentication methods control access to the switch. For additional information, refer to the following URL: http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_cfg_authentifcn_ps_6350_TSD_Products_Configuration_Guide_Chapter.html
- Visual port status information—The switch LEDs provide visual management of port- and switch-level status.
- Secure Shell—Secure Shell (SSH) is a program that enables you to log into another computer over a network, to execute commands remotely, and to move files from one machine to another. The implementation will be limited to providing a remote login session to the switch, and will only function as a server; that is, the switch may not initiate SSH connections.
- NetFlow statistics—This feature is a global traffic monitoring feature that allows flow-level monitoring of all IPv4-routed traffic through the switch.
- Switched Port Analyzer (SPAN)—SPAN allows you to monitor traffic on any port for analysis by a network analyzer or Remote Monitoring (RMON) probe. You also can do the following:
 - Allow incoming traffic on SPAN destination ports to be switched normally.
 - Explicitly configure the encapsulation type of packets that are spanned out of a destination port.
 - Restrict ingress sniffing depending on whether the packet is unicast, multicast, or broadcast, and depending on whether the packet is valid.
 - For troubleshooting purposes, mirror packets sent to or from the CPU out of a SPAN destination port.

For information on SPAN, see [Chapter 32, “Configuring SPAN.”](#)

- Simple Network Management Protocol—SNMP facilitates the exchange of management information between network devices. The Catalyst 4500 series switch supports these SNMP types and enhancements:
 - SNMP—A full Internet standard
 - SNMP v2—Community-based administrative framework for version 2 of SNMP
 - SNMP trap message enhancements—Additional information with certain SNMP trap messages, including spanning-tree topology change notifications and configuration change notifications

For information on SNMP, refer to the *Cisco IOS Configuration Fundamentals Configuration Guide* at the following URL:

http://www.cisco.com/en/US/docs/ios/fundamentals/configuration/guide/12_4/cf_12_4_book.html

- Dynamic Host Control Protocol server—DHCP server enables you to automatically assign reusable IP addresses to DHCP clients. The Cisco IOS DHCP server feature is a full DHCP server implementation that assigns and manages IP addresses from specified address pools within the router to DHCP clients. If the Cisco IOS DHCP server cannot satisfy a DHCP request from its own database, it can forward the request to one or more secondary DHCP servers defined by the network administrator.

For DHCP server configuration information, refer to the chapter, “Configuring DHCP,” in the *Cisco IOS IP and IP Routing Configuration Guide* at the following URL:

http://www.cisco.com/en/US/docs/ios/ipaddr/configuration/guide/iad_dhcp_rdm_p6350_TSD_Products_Configuration_Guide_Chapter.html

- Debugging features—The Catalyst 4500 series switch has several commands to help you debug your initial setup. These commands include the following groups:
 - **platform**
 - **debug platform**

For more information on these commands, refer to the *Catalyst 4500 Series Switch Cisco IOS Command Reference*.

