



Product Overview

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

- [Layer 2 Software Features, page 1-2](#)
- [Layer 3 Software Features, page 1-6](#)
- [QoS Features, page 1-11](#)
- [Management and Security Features, page 1-11](#)
- [Configuring Embedded CiscoView Support, page 1-13](#)

**Note**

For complete syntax and usage information for the switch commands used in this chapter, look at the *Cisco Catalyst 4500 Series Switch Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html>

If the command is not found in the Catalyst 4500 Command Reference, it will be found in the larger Cisco IOS library. Refer to the *Catalyst 4500 Series Switch Cisco IOS Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/ps6350/index.html>

Layer 2 Software Features

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

- [802.1Q and Layer 2 Protocol Tunneling, page 1-2](#)
- [Storm Control, page 1-3](#)
- [CDP, page 1-3](#)
- [DHCP Snooping, page 1-3](#)
- [EtherChannel Bundles, page 1-3](#)
- [IP Source Guard, page 1-4](#)
- [Jumbo Frames, page 1-4](#)
- [Layer 2 Traceroute, page 1-4](#)
- [MST, page 1-4](#)
- [PVRST+, page 1-4](#)
- [Spanning Tree Protocol, page 1-5](#)
- [UDLD, page 1-5](#)
- [Unidirectional Ethernet, page 1-5](#)
- [VLANs, page 1-5](#)

802.1Q and Layer 2 Protocol Tunneling

802.1Q tunneling is a Q-in-Q technique that expands the VLAN space by retagging the tagged packets that enter the service provider infrastructure. 802.1Q tunneling allows service providers to assign a VLAN to each customer without losing the original customer VLAN IDs inside the tunnel. All data traffic that enters the tunnel is encapsulated with the tunnel VLAN ID. Layer 2 Protocol Tunneling is a similar technique for all Layer 2 control traffic. 802.1Q tunneling and Layer 2 Protocol Tunneling are supported on Supervisor Engine V only.

For information on configuring 802.1Q tunneling, see [Chapter 17, “Configuring 802.1Q and Layer 2 Protocol Tunneling.”](#)

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. Multicast and 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 36, “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 18, “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 information on configuring DHCP snooping, see [Chapter 31, “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 15, “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 31, “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 instances 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 within 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 19, “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 20, “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 25, “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 34, “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. With the VMPS Client feature, you can convert a dynamic access port to a VMPS client. VMPS clients can use VQP queries to communicate with the VMPS server to obtain a VLAN assignment for the port based on the MAC address of the host attached to that port. 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 an intranet, and it 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 4500 series switch:

- [CEF, page 1-7](#)
- [HSRP, page 1-7](#)
- [IP Routing Protocols, page 1-7](#)
- [Multicast Services, page 1-9](#)
- [Network Security with ACLs, page 1-10](#)
- [Policy-Based Routing, page 1-10](#)
- [Unidirectional Link Routing, page 1-10](#)
- [VRF-lite, page 1-10](#)

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 22, “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 4500 series 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.

For details on OSPF, refer to this URL:

http://www.cisco.com/en/US/tech/tk365/tk480/tsd_technology_support_sub-protocol_home.html

IS-IS

The Intermediate System-to-Intermediate System Protocol (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 Protocol requires each router to maintain a full topology map of the network (that is, which intermediate systems and end systems are connected to which other intermediate systems and end systems). Periodically, the router runs an algorithm over its map to calculate the shortest path to all possible destinations.

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

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

Only Ethernet frames are supported. The IS-IS Protocol 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).

For details on EIGRP, refer to this URL

http://www.cisco.com/en/US/products/ps6630/products_ios_protocol_option_home.html

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.

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 forwards packets based on their content. Multicast traffic is forwarded only to ports with attached hosts that request multicast traffic.

Support for IGMPv3 provides constrained flooding of multicast traffic in the presence of IGMPv3 hosts or routers. IGMPv3 snooping listens to IGMPv3 query and membership report messages to maintain host-to-multicast group associations. It enables a switch to propagate multicast data only to ports that need it. IGMPv3 snooping is fully interoperable with IGMPv1 and IGMPv2.

Explicit Host Tracking (EHT) is an extension to IGMPv3 snooping. EHT enables immediate leave operations on a per-port basis. EHT can be used to track per host membership information or to gather statistics about all IGMPv3 group members.

For information on configuring IGMP snooping, see [Chapter 16, “Configuring IGMP Snooping and Filtering.”](#)

- 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 4500 series 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, MACLs, VLAN maps, MAC address filtering, and Port ACLs, see [Chapter 33, “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, and so on. This feature allows network managers more flexibility in how they configure and design their networks.

For more information on policy-based routing, see [Chapter 24, “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*.

VRF-lite

VPN routing and forwarding (VRF-lite) is an extension of IP routing that provides multiple routing instances. Along with BGP, it enables the creation of a Layer 3 VPN service by keeping separate IP routing and forwarding tables for each VPN customer. VRF-lite uses input interfaces to distinguish routes for different VPNs. It forms virtual packet-forwarding tables by associating one or more Layer 3 interfaces with each VRF, allowing the creation of multiple Layer 3 VPNs on a single switch. Interfaces in a VRF could be either physical, such as an Ethernet port, or logical, such as a VLAN switch virtual interface (SVI). However, interfaces cannot belong to more than one VRF at any time.

For information on VRF-lite, see [Chapter 26, “Configuring VRF-lite.”](#)

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 through automatic configuration.

For information on QoS and Auto QoS, see [Chapter 27, “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:

- Intelligent Power Management—This feature enables a 802.3af compliant module port that is attached to a powered device (PD) to detect the PD’s 802.3af class and provision power accordingly. The switch and the PD exchange a CDP packet that identifies the precise maximum power required by the PD.
- Stateful switchover (SSO)—This feature propagates configuration and state information from the active to the redundant supervisor engine so that sub-second interruptions in Layer 2 traffic occur when the active supervisor engine switches over to the redundant supervisor engine.

- Stateful IGMP Snooping

This feature propagates the IGMP data learned by the active supervisor engine to the redundant supervisor engine so that when a switchover occurs, the newly active supervisor engine is aware of the multicast group membership, which alleviates a disruption to multicast traffic during a switchover.

- Stateful DHCP Snooping

This feature propagates the DHCP-snooped data from the active supervisor engine to the redundant supervisor engine so that when a switchover occurs, the newly active supervisor engine is aware of the DHCP data that was already snooped, and the security benefits continue uninterrupted.

- 802.1X protocol—This feature provides a means for a host that is 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.
 - 802.1X RADIUS accounting—This feature allows you to track the usage of network devices.
 - 802.1X with Voice VLAN—This feature allows you to use 802.1X security on a port while enabling it to be used by both Cisco IP phones and devices with 802.1X supplicant support.
 - 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.
 - Visual port status information—The switch LEDs provide visual management of port- and switch-level status.
 - Auto 10/100/1000 negotiation—This feature allows you to configure a port to limit the speed at which it will autonegotiate to a speed lower than the physically maximum speed. This method of reducing the throughput incurs much less overhead than using an ACL.
 - 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 switch may not initiate SSH connections: SSH will be limited to providing a remote login session to the switch and will only function as a server.
 - 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:
 - Configure ACLs on SPAN sessions.
 - 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.
 - Mirror packets sent to or from the CPU out of a SPAN destination port for troubleshooting purposes.
- For information on SPAN, see [Chapter 38, “Configuring SPAN and RSPAN.”](#)
- Remote SPAN (RSPAN)—RSPAN is an extension of SPAN, where source ports and destination ports are distributed across multiple switches, allowing remote monitoring of multiple switches across the network. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session on all participating switches.

For information on RSPAN, see [Chapter 38, “Configuring SPAN and RSPAN.”](#)

- 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 v3—Security framework with three levels: noAuthNoPriv, authNoPriv, and authPriv (available only on a crypto image, like cat4000-i5k91s-mz)
 - SNMP trap message enhancements—Additional information with certain SNMP trap messages, including spanning-tree topology change notifications and configuration change notifications

- Dynamic Host Control Protocol server—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.

With DHCP-based autoconfiguration, your switch (the DHCP client) is automatically configured at startup with IP address information and a configuration file.

- 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, refer to the *Catalyst 4500 Series Switch Cisco IOS Command Reference*.

Configuring Embedded CiscoView Support

The Catalyst 4500 series switch supports CiscoView web-based administration through the Catalyst Web Interface (CWI) tool. CiscoView is a device management application that can be embedded on the switch flash and provides dynamic status, monitoring, and configuration information for your switch.

CiscoView displays a physical view of your switch chassis, with color-coded modules and ports, and monitoring capabilities that display the switch status, performance, and other statistics. Configuration capabilities allow comprehensive changes to devices, given that the required security privileges have been granted. The configuration and monitoring capabilities for the Catalyst 4500 series of switches mirror those available in CiscoView in all server-based CiscoWorks solutions, including CiscoWorks LAN Management Solution (LMS) and CiscoWorks Routed WAN Management Solution (RWAN).

These sections describe the Embedded CiscoView support available with Release 12.1(20)EW and later releases:

- [Understanding Embedded CiscoView, page 1-14](#)
- [Installing and Configuring Embedded CiscoView, page 1-14](#)
- [Displaying Embedded CiscoView Information, page 1-16](#)

Understanding Embedded CiscoView

The Embedded CiscoView network management system is a web-based interface that uses HTTP and SNMP to provide a graphical representation of the switch and to provide a GUI-based management and configuration interface. You can download the Java Archive (JAR) files for Embedded CiscoView at this URL:

<http://www.cisco.com/cgi-bin/tablebuild.pl/cview-cat4000>

Installing and Configuring Embedded CiscoView

To install and configure Embedded CiscoView, perform this task:

	Command	Purpose
Step 1	Router# dir <i>device_name</i>	Displays the contents of the device. If you are installing Embedded CiscoView for the first time, or if the CiscoView directory is empty, skip to Step 5 .
Step 2	Switch# delete <i>device_name:cv/*</i>	Removes existing files from the CiscoView directory.
Step 3	Switch# squeeze <i>device_name:</i>	Recovers the space in the file system.
Step 4	Switch# acopy tftp bootflash	Copies the tar file to bootflash.
Step 5	Switch# archive tar /xtract tftp:// ip address of tftp server/ciscoview.tar device_name:cv	Extracts the CiscoView files from the tar file on the TFTP server to the CiscoView directory.
Step 6	Switch# dir <i>device_name:</i>	Displays the contents of the device. In a redundant configuration, repeat Step 1 through Step 6 for the file system on the redundant supervisor engine.
Step 7	Switch# configure terminal	Enters global configuration mode.
Step 8	Switch(config)# ip http server	Enables the HTTP web server.
Step 9	Switch(config)# snmp-server community string ro	Configures the SNMP password for read-only operation.
Step 10	Switch(config)# snmp-server community string rw	Configures the SNMP password for read/write operation.



Note

The default password for accessing the switch web page is the enable-level password of the switch.

The following example shows how to install and configure Embedded CiscoView on your switch:

```
Switch# dir
Directory of bootflash:

 1  -rw-      8620304  Dec 23 2002 23:27:49 +00:00  wickwire.EW1
 2  -rw-      9572396  Dec 30 2002 01:05:01 +00:00  cat4000-i9k2s-mz.121-19.EW
 3  -rw-      9604192   Jan 3 2003 07:46:49 +00:00  cat4000-i5k2s-mz.121-19.EW
 4  -rw-      1985024  Jan 21 2003 03:31:20 +00:00  Cat4000IOS.v4-0.tar
 5  -rw-      1910127  Jan 23 2003 04:23:39 +00:00  cv/Cat4000IOS-4.0.sgz
 6  -rw-           7258  Jan 23 2003 04:23:46 +00:00  cv/Cat4000IOS-4.0_ace.html
 7  -rw-           405  Jan 23 2003 04:23:46 +00:00  cv/Cat4000IOS-4.0_error.html
 8  -rw-           2738  Jan 23 2003 04:23:46 +00:00  cv/Cat4000IOS-4.0_install.html
```

```

 9 -rw-      20450 Jan 23 2003 04:23:46 +00:00 cv/Cat4000IOS-4.0_jks.jar
10 -rw-      20743 Jan 23 2003 04:23:46 +00:00 cv/Cat4000IOS-4.0_nos.jar
11 -rw-     12383 Jan 23 2003 04:23:46 +00:00 cv/applet.html
12 -rw-       529 Jan 23 2003 04:23:46 +00:00 cv/cisco.x509
13 -rw-     2523 Jan 23 2003 04:23:46 +00:00 cv/identitydb.obj
14 -rw-    9630880 Feb 27 2003 01:25:16 +00:00 kurt70.devtest-enh
15 -rw-       1173 Mar 19 2003 05:50:26 +00:00 post-2003.03.19.05.50.07-passed.txt
16 -rw-    10511956 Mar 26 2003 04:24:12 +00:00 kurt_alpha_bas_crypto_103

61341696 bytes total (9436548 bytes free)
Switch#
Switch# del cv/*
Delete filename [cv/*]?
Delete bootflash:cv/Cat4000IOS-4.0.sgz? [confirm]y
Delete bootflash:cv/Cat4000IOS-4.0_ace.html? [confirm]y
Delete bootflash:cv/Cat4000IOS-4.0_error.html? [confirm]y
Delete bootflash:cv/Cat4000IOS-4.0_install.html? [confirm]y
Delete bootflash:cv/Cat4000IOS-4.0_jks.jar? [confirm]y
Delete bootflash:cv/Cat4000IOS-4.0_nos.jar? [confirm]y
Delete bootflash:cv/applet.html? [confirm]y
Delete bootflash:cv/cisco.x509? [confirm]y
Delete bootflash:cv/identitydb.obj? [confirm]y
Switch#

Switch# squeeze bootflash:
All deleted files will be removed. Continue? [confirm]y
Squeeze operation may take a while. Continue? [confirm]y
Squeeze of bootflash complete
Switch#

Switch# copy tftp bootflash
Address or name of remote host []? 10.5.5.5
Source filename []? Cat4000IOS.v5-1.tar
Destination filename [Cat4000IOS.v5-1.tar]?
Accessing tftp://10.5.5.5/Cat4000IOS.v5-1.tar...
Loading Cat4000IOS.v5-1.tar from 10.5.5.5 (via FastEthernet1):
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
[OK - 2031616 bytes]

2031616 bytes copied in 11.388 secs (178400 bytes/sec)
Switch#
Switch# dir
Directory of bootflash:/

 2 -rw-      8620304 Dec 23 2002 23:27:49 +00:00 wickwire.EW1
 2 -rw-      9572396 Dec 30 2002 01:05:01 +00:00 cat4000-i9k2s-mz.121-19.EW
 3 -rw-      9604192 Jan  3 2003 07:46:49 +00:00 cat4000-i5k2s-mz.121-19.EW
 4 -rw-      1985024 Jan 21 2003 03:31:20 +00:00 Cat4000IOS.v4-0.tar
 5 -rw-      9630880 Feb 27 2003 01:25:16 +00:00 kurt70.devtest-enh
 6 -rw-         1173 Mar 19 2003 05:50:26 +00:00 post-2003.03.19.05.50.07-passed.txt
 7 -rw-    10511956 Mar 26 2003 04:24:12 +00:00 kurt_alpha_bas_crypto_103
 8 -rw-      2031616 Mar 26 2003 05:33:12 +00:00 Cat4000IOS.v5-1.tar

61341696 bytes total (9383128 bytes free)
Switch#
Switch# archive tar /xtract Cat4000IOS.v5-1.tar /cv
extracting Cat4000IOS-5.1.sgz (1956591 bytes)
extracting Cat4000IOS-5.1_ace.html (7263 bytes)
extracting Cat4000IOS-5.1_error.html (410 bytes)
extracting Cat4000IOS-5.1_install.html (2743 bytes)
extracting Cat4000IOS-5.1_jks.jar (20450 bytes)
extracting Cat4000IOS-5.1_nos.jar (20782 bytes)
extracting applet.html (12388 bytes)
extracting cisco.x509 (529 bytes)
extracting identitydb.obj (2523 bytes)

```

```

Switch#
Switch# dir
Directory of bootflash:/

 1  -rw-      8620304  Dec 23 2002 23:27:49 +00:00  wickwire.EW1
 2  -rw-      9572396  Dec 30 2002 01:05:01 +00:00  cat4000-i9k2s-mz.121-19.EW
 3  -rw-      9604192   Jan 3 2003 07:46:49 +00:00  cat4000-i5k2s-mz.121-19.EW
 4  -rw-      1985024   Jan 21 2003 03:31:20 +00:00  Cat4000IOS.v4-0.tar
 5  -rw-      9630880   Feb 27 2003 01:25:16 +00:00  kurt70.devtest-enh
 6  -rw-         1173   Mar 19 2003 05:50:26 +00:00  post-2003.03.19.05.50.07-passed.txt
 7  -rw-     10511956   Mar 26 2003 04:24:12 +00:00  kurt_alpha_bas_crypto_103
 8  -rw-     2031616   Mar 26 2003 05:33:12 +00:00  Cat4000IOS.v5-1.tar
 9  -rw-     1956591   Mar 26 2003 05:36:11 +00:00  cv/Cat4000IOS-5.1.sgz
10  -rw-         7263   Mar 26 2003 05:36:19 +00:00  cv/Cat4000IOS-5.1_ace.html
11  -rw-         410   Mar 26 2003 05:36:19 +00:00  cv/Cat4000IOS-5.1_error.html
12  -rw-         2743   Mar 26 2003 05:36:19 +00:00  cv/Cat4000IOS-5.1_install.html
13  -rw-        20450   Mar 26 2003 05:36:19 +00:00  cv/Cat4000IOS-5.1_jks.jar
14  -rw-        20782   Mar 26 2003 05:36:19 +00:00  cv/Cat4000IOS-5.1_nos.jar
15  -rw-        12388   Mar 26 2003 05:36:19 +00:00  cv/applet.html
16  -rw-         529   Mar 26 2003 05:36:19 +00:00  cv/cisco.x509
17  -rw-         2523   Mar 26 2003 05:36:19 +00:00  cv/identitydb.obj

61341696 bytes total (7358284 bytes free)
Switch#
Switch# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# ip http server
Switch(config)# snmp-server community public ro
Switch(config)# snmp-server community public rw
Switch(config)# exit
Switch# wr
Building configuration...
Compressed configuration from 2735 bytes to 1169 bytes[OK]
Switch# show ciscoview ?
  package  ADP Package Details
  version  ADP version
  |        Output modifiers
  <

```

Displaying Embedded CiscoView Information

To display the Embedded CiscoView information, enter the following commands:

Command	Purpose
Switch# show ciscoview package	Displays information about the Embedded CiscoView files.
Switch# show ciscoview version	Displays the Embedded CiscoView version.

The following example shows how to display the Embedded CiscoView file and version information:

```
Switch# show ciscoview package
File source:
CVFILE                               SIZE(in bytes)
-----
Cat4000IOS-5.1.sgz                   1956591
Cat4000IOS-5.1_ace.html               7263
Cat4000IOS-5.1_error.html             410
Cat4000IOS-5.1_install.html           2743
Cat4000IOS-5.1_jks.jar                20450
Cat4000IOS-5.1_nos.jar                20782
applet.html                           12388
cisco.x509                             529
identitydb.obj                        2523

Switch# show ciscoview version
Engine Version: 5.3.4 ADP Device: Cat4000IOS ADP Version: 5.1 ADK: 49
Switch#
```

