



Product Overview

This chapter provides an overview of the Catalyst 4006 switch with Supervisor Engine III. This chapter includes these two major sections:

- [Understanding Layer 3 Switching, page 1-1](#)
- [Supported Features, page 1-1](#)

Understanding Layer 3 Switching

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.

Supported Features

The following sections describes the key supported features:

- [Supported Hardware, page 1-2](#)
- [Layer 2 Software Features, page 1-2](#)
- [Layer 3 Software Features, page 1-4](#)
- [QoS Features, page 1-7](#)
- [Management and Security Features, page 1-7](#)



Note

For more information about the chassis, modules, and software features supported by the Catalyst 4006 switch with Supervisor Engine III, refer to the *Release Notes for Catalyst 4006 switch with Supervisor Engine III* at <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/relnotes/>

Supported Hardware

The Catalyst 4006 switch supports the Supervisor Engine III and the following modules:

- WS-X4124-FX-MT—24-port 100BASE-FX Fast Ethernet switching module
- WS-X4148-FX-MT—48-port 100BASE-FX Fast Ethernet switching module
- WS-X4148-RJ—48-port 10/100 Fast Ethernet RJ-45
- WS-X4148-RJ21—48-port 10/100-Mbps Fast Ethernet switching module
- WS-X4148-RJ45-V—48-port inline power 10/100BASE-TX switching module
- WS-X4232-GB-RJ—32-port 10/100 Fast Ethernet RJ-45, plus 2-port 1000BASE-X (GBIC) Gigabit Ethernet
- WS-X4232-RJ-XX—32-port 10/100 Fast Ethernet RJ-45
- WS-X4306-GB—6-port 1000BASE-X (GBIC) Gigabit Ethernet
- WS-X4418-GB—18-port 1000BASE-X (GBIC) Gigabit Ethernet switching module
- WS-X4412-2GB-T—12-port 1000BASE-T Gigabit Ethernet switching module
- WS-X4424-GB-RJ45—24-port 10/100/1000BASE-T Gigabit Ethernet switching module
- WS-X4448-GB-LX—48-port 1000BASE-LX Gigabit Ethernet Fiber Optic interface module
- WS-X4448-GB-RJ45—48-port 10/100/1000BASE-T Gigabit Ethernet switching module
- WS-X4095-PEM—Catalyst 4000 DC Power Entry Module
- WS-P4603-2PSU—Catalyst 4000 Auxiliary Power Shelf (3-slot) including two WS-X4608 power supplies
- WS-X4608—Catalyst 4603 Power Supply Unit for WS-P4603

Layer 2 Software Features

The following sections describe the key Layer 2 switching software features on the Catalyst 4006 switch with Supervisor Engine III:

- [CDP, page 1-2](#)
- [EtherChannel Bundles, page 1-3](#)
- [Spanning Tree Protocol, page 1-3](#)
- [VLANs, page 1-3](#)
- [UniDirectional Link Detection, page 1-4](#)

CDP

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.

EtherChannel Bundles

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

For more information on configuring EtherChannel, see [Chapter 11, “Understanding and Configuring EtherChannel.”](#)

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.

The Catalyst 4006 switch with Supervisor Engine III supports the following spanning-tree enhancements:

- Spanning tree PortFast—PortFast allows a port with a directly attached host to transition to the forwarding state immediately, 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.

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, which 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 7, “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 9, “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 more information about private VLANs, see [Chapter 8, “Understanding and Configuring Private VLANs.”](#)

UniDirectional Link Detection

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 more information about UDLD, see [Chapter 19, “Understanding and Configuring UDLD.”](#)

Layer 3 Software Features

The following sections describe the key Layer 3 switching software features on the Catalyst 4006 switch with Supervisor Engine III:

- [IP Routing Protocols, page 1-4](#)
- [Multicast Services, page 1-6](#)
- [CEF, page 1-6](#)
- [Network Security \(ACLs\), page 1-7](#)

IP Routing Protocols

The following sections describe the routing protocols supported on the Catalyst 4006 switch with Supervisor Engine III:

- [RIP, page 1-4](#)
- [OSPF, page 1-4](#)
- [IGRP, page 1-5](#)
- [EIGRP, page 1-5](#)
- [BGP, page 1-5](#)
- [HSRP, page 1-6](#)

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 grouping 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.

IGRP

The Interior Gateway Routing Protocol (IGRP) is a distance vector interior-gateway routing protocol developed by Cisco. Distance vector routing protocols request that a switch send all or a portion of its routing table in a routing update message at regular intervals to each of its neighboring routers. As routing information proliferates through the network, routers can calculate distance to all the 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, partial 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 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 (EGP) 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 4006 switch with Supervisor Engine III 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 “Configuring BGP” in the *Cisco IOS IP and IP Routing Configuration Guide* at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt2/1cdbgp.htm

- For a complete description of the BGP commands, refer to “BGP Commands” in the *Cisco IOS IP and IP Routing Command Reference* at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_r/iprprt2/1rdbgp.htm

- For multiprotocol BGP configuration information, refer to “Configuring Multiprotocol BGP Extensions for IP Multicast” in the *Cisco IOS IP and IP Routing Configuration Guide* at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt2/1cdmbgp.htm

- For multiprotocol BGP command descriptions, refer to “Multiprotocol BGP Extensions for IP Multicast Commands” in the *Cisco IOS IP and IP Routing Command Reference* at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_r/iprprt2/1rdmbgp.htm

HSRP

Hot Standby Routing Protocol (HSRP) provides high network availability by routing IP traffic from hosts on Ethernet networks without relying on the availability of any single 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 more information on configuring HSRP, refer to the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ip_c/ipcprt1/1cdip.htm

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.
- Protocol Independent Multicast (PIM)—PIM is a protocol-independent multicast service because it can leverage whichever unicast routing protocol is used to populate the unicast routing table, including EIGRP, OSPF, BGP, or static route, to support IP multicast. 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 15, “Understanding and Configuring IP Multicast.”](#)

CEF

The Catalyst 4006 switch with Supervisor Engine III features Cisco Express Forwarding (CEF). 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 characterized by 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.

Network Security (ACLs)

Access control lists (ACLs) filter network traffic by controlling whether routed packets are forwarded or blocked at the router's interfaces. The Catalyst 4006 switch with Supervisor Engine III examines each packet to determine whether to forward or drop the packet, based on the criteria you specified within the access lists.

For information on configuring access lists, see [Chapter 16, “Configuring Network Security.”](#)

QoS Features

The quality of service (QoS) features select network traffic, prioritize it according to its relative importance, and use this information to avoid congestion. Implementing QoS in your network makes network performance more predictable and bandwidth use more effective.

The Catalyst 4006 switch with Supervisor Engine III supports the following QoS features:

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

For information on configuring QoS, see [Chapter 20, “Understanding and Configuring QoS.”](#)

Management and Security Features

The Catalyst 4006 switch with Supervisor Engine III 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:

- Password-protected access (read-only and read-write)—This feature protects management interfaces against unauthorized configuration changes.
- Local, RADIUS, and TACACS+ authentication—These authentication methods control access to the switch. For additional information, refer to “Authentication, Authorization, and Accounting (AAA),” in *Cisco IOS Security Configuration Guide*, Release 12.1, at the following URL: http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgr/secur_c/scprt1/index.htm
- Visual port status information—The switch LEDs provide visual management of port- and switch-level status.
- Switched Port Analyzer (SPAN)—SPAN allows you to monitor traffic on any port for analysis by a network analyzer or remote monitoring (RMON) probe. For information on SPAN, see [Chapter 21, “Configuring SPAN.”](#)
- Simple Network Management Protocol (SNMP)—This protocol facilitates the exchange of management information between network devices. The Catalyst 4006 switch with Supervisor Engine III 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 *IOS Configuration Fundamentals Configuration Guide* and *Configuration Fundamentals Command Reference* at the following URL:
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/index.htm>

- Dynamic Host Control Protocol (DHCP) server— DHCP 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 more information on configuring the DHCP Server, refer to the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t1/120t1/easyip2.htm>

- Debugging features—The Catalyst 4006 switch with Supervisor Engine III has several commands to help you debug your initial setup. These commands include the following groups:
 - **show platform**
 - **platform**
 - **debug platform**

For more information on these commands, refer to the *Command Reference for the Catalyst 4006 Switch with Supervisor Engine III*.