



Product Overview

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

- [Layer 2 Software Features, page 1-1](#)
- [Layer 3 Software Features, page 1-13](#)
- [Management Features, page 1-23](#)
- [Security Features, page 1-34](#)

Layer 2 Software Features

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

- [802.1Q Tunneling, VLAN Mapping, and Layer 2 Protocol Tunneling, page 1-2](#)
- [Cisco IOS Auto Smartport Macros, page 1-2](#)
- [Cisco Discovery Protocol, page 1-2](#)
- [Cisco Group Management Protocol \(CGMP\) server, page 1-3](#)
- [EtherChannel Bundles, page 1-3](#)
- [Ethernet CFM, page 1-3](#)
- [Ethernet OAM Protocol, page 1-3](#)
- [Flex Links and MAC Address-Table Move Update, page 1-3](#)
- [Flexible NetFlow \(Supervisor Engine 7-E, 7L-E, and 8-E only\), page 1-4](#)
- [Internet Group Management Protocol \(IGMP\) Snooping, page 1-4](#)
- [IPv6 Multicast BSR and BSR Scoped Zone Support, page 1-5](#)
- [IPv6 Multicast Listen Discovery \(MLD\) and Multicast Listen Discovery Snooping, page 1-6](#)
- [Jumbo Frames, page 1-6](#)
- [Link Aggregation Control Protocol, page 1-7](#)
- [Link Layer Discovery Protocol, page 1-7](#)
- [Link State Tracking, page 1-8](#)
- [Location Service, page 1-8](#)
- [Multiple Spanning Tree, page 1-8](#)

- [Per-VLAN Rapid Spanning Tree](#), page 1-8
- [Quality of Service](#), page 1-9
- [Resilient Ethernet Protocol](#), page 1-10
- [SmartPort Macros](#), page 1-10
- [Spanning Tree Protocol](#), page 1-10
- [Stateful Switchover](#), page 1-10
- [SVI Autostate](#), page 1-11
- [Unidirectional Link Detection](#), page 1-11
- [VLANs](#), page 1-11
- [Virtual Switching Systems \(Catalyst 4500-X and Supervisor Engine 7-E, 7L-E, and 8-E\)](#), page 1-12
- [Virtual Switch System Client](#), page 1-12
- [Y.1731 \(AIS and RDI\)](#), page 1-13

802.1Q Tunneling, VLAN Mapping, 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.

To map customer VLANs to service-provider VLANs, you can configure VLAN mapping (or VLAN ID translation) on trunk ports connected to a customer network. Packets entering the port are mapped to a service provider VLAN (S-VLAN) based on the port number and the original customer VLAN-ID (C-VLAN) of the packet.

For information on configuring 802.1Q tunneling and VLAN Mapping, see [Chapter 30, “Configuring 802.1Q Tunneling, VLAN Mapping, and Layer 2 Protocol Tunneling.”](#)

Cisco IOS Auto Smartport Macros

Cisco IOS Auto SmartPort macros dynamically configure ports based on the device type detected on the port. When the switch detects a new device on a port it applies the appropriate Cisco IOS Auto Smartports macro. When a link-down event occurs on the port, the switch removes the macro. For example, when you connect a Cisco IP phone to a port, Cisco IOS Auto SmartPorts automatically applies the IP phone macro. The IP phone macro enables quality of service (QoS), security features, and a dedicated voice VLAN to ensure proper treatment of delay-sensitive voice traffic.

For information on configuring SmartPort macros, see [Chapter 22, “Configuring Cisco IOS Auto Smartport Macros.”](#)

Cisco Discovery Protocol

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 31, “Configuring CDP.”](#)

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.

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 27, “Configuring EtherChannel and Link State Tracking.”](#)

Ethernet CFM

Ethernet CFM is an end-to-end per-service-instance (per-VLAN) Ethernet layer OAM protocol that includes proactive connectivity monitoring, fault verification, and fault isolation. End-to-end can be provider-edge-to-provider-edge (PE-to-PE) device or customer-edge-to-customer-edge (CE-to-CE) device. Ethernet CFM, as specified by IEEE 802.1ag, is the standard for Layer 2 ping, Layer 2 traceroute, and end-to-end connectivity check of the Ethernet network.

For information about CFM, see [Chapter 75, “Configuring Ethernet OAM and CFM.”](#)

Ethernet OAM Protocol

Ethernet Operations, Administration, and Maintenance (OAM) is a protocol for installing, monitoring, and troubleshooting Ethernet networks to increase management capability within the context of the overall Ethernet infrastructure. You can implement Ethernet OAM on any full-duplex, point-to-point, or emulated point-to-point Ethernet link for a network or part of a network (specified interfaces).

For information about OAM, see [Chapter 75, “Configuring Ethernet OAM and CFM.”](#)

Flex Links and MAC Address-Table Move Update

Flex Links are a pair of Layer 2 interfaces (switch ports or port channels) where one interface is configured to act as a backup to the other. The feature provides an alternative solution to the Spanning Tree Protocol (STP). Flex Links are typically configured in service provider or enterprise networks where customers do not want to run STP on the switch.

MAC Address-Table Move Update allows a switch to provide rapid bidirectional convergence when a primary (forwarding) link goes down and the standby link begins forwarding traffic.

For information about Flex Links and MAC Address-Table Move Update, see [Chapter 24, “Configuring Flex Links and MAC Address-Table Move Update.”](#)

Flexible NetFlow (Supervisor Engine 7-E, 7L-E, and 8-E only)

Flow is defined as unique set of key fields attributes, which might include fields of packet, packet routing attributes, and input and output interface information. A NetFlow feature defines a flow as a sequence of packets that have the same values for the feature key fields. Flexible NetFlow (FNF) allows a flow record that specifies various flow attributes to be collected and optionally exported. NetFlow collection supports IP, IPv6 and Layer 2 traffic.

For information on configuring Flexible NetFlow, see [Chapter 74, “Configuring Flexible NetFlow.”](#)

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.

The IGMP Snooping Querier is a Layer 2 feature required to support IGMP snooping in a VLAN where PIM and IGMP are not configured because the multicast traffic does not require routing.

With SSO support, Stateful IGMP Snooping 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.

Beginning with Release IOS XE 3.5.0E and IOS 15.2(1)E, the Catalyst 4500 series switch supports an application of local IGMP snooping, Multicast VLAN Registration (MVR). MVR is designed for applications using wide-scale deployment of multicast traffic across an Ethernet ring-based service provider network (for example, the broadcast of multiple television channels over a service-provider network). MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on the network-wide multicast VLAN. It allows the single multicast VLAN to be shared in the network while subscribers remain in separate VLANs.

For information on configuring IGMP snooping and MVR, see [Chapter 28, “Configuring IGMP Snooping and Filtering, and MVR.”](#)

IPv6 Multicast BSR and BSR Scoped Zone Support

The bootstrap router (BSR) protocol for PIM-SM provides a dynamic, adaptive mechanism to distribute group-to-RP mapping information rapidly throughout a domain. With the IPv6 BSR feature, if an RP becomes unreachable, it will be detected and the mapping tables will be modified so that the unreachable RP is no longer used, and the new tables will be rapidly distributed throughout the domain.

BSR provides scoped zone support by distributing group-to-RP mappings in networks using administratively scoped multicast. The user can configure candidate BSRs and a set of candidate RPs for each administratively scoped region in the user's domain.

For information on BSR and BSR Scoped Zone Support, see this URL with the following caveats related to support on a Catalyst 4500 series switch:

http://www.cisco.com/en/US/docs/ios-xml/ios/ipmulti_pim/configuration/xe-3s/ip6-mcast-bsr.html

- In the section “IPv6 BSR: Scoped Zone Support,” a paragraph starts as follows:

Unless the C-RP is configured with a scope, it discovers the existence of the administratively scoped zone and its group range through reception of a BSM from the scope zone's elected BSR containing the scope zone's group range.

A C-RP can no longer be configured with a scope. So, the sentence should read:

A C-RP discovers the existence of the administratively scoped zone and its group range through reception of a BSM from the scope zone's elected BSR containing the scope zone's group range.

- In the section "Configuring a BSR and Verifying BSR Information" in Step 3 under Summary Steps and Detailed Steps, the command for configuring a C-BSR is listed as:

```
ipv6 pim [vrf vrf-name] bsr candidate bsr ipv6-address [hash-mask-length] [priority priority-value]
```

Because the original syntax mistakenly excludes **scope scope-value** and the “new” option (**accept-rp-candidate access-list-name**) is supported with this release.

```
ipv6 pim [vrf vrf-name] bsr candidate bsr ipv6-address [hash-mask-length] [priority priority-value] [scope scope-value] [accept-rp-candidate access-list-name]
```

- In the section "Sending PIM RP Advertisements to the BSR" in Step 3 under Summary and Detailed Steps, the keyword **scope scope-value** should be removed. The **scope** keyword no longer exists for C-RPs.
- In the section "Configuring BSR for Use Within Scoped Zones,” several changes apply.

The following paragraph:

If scope is specified on the candidate RP, then this device will advertise itself as C-RP only to the BSR for the specified scope. If the group list is specified along with the scope, then only prefixes in the access list with the same scope as that configured will be advertised.

Should read:

The candidate RP will advertise the different ranges it serves to the respective elected BSRs. If a group list is specified, for each of the prefixes in the group list, it will verify that there is an elected scoped BSR for the scope of the prefix. If none exists, the prefix will be announced to the elected non-scoped BSR, provided one is present.

Note: If a prefix is not scope specific (for example, FF00::/8), it will only be announced to a non-scoped BSR. If the candidate RP is not configured with a group list, it will behave as if a group list with only the prefix FF00::/8 is configured.

Under the Summary Steps, steps 3 and 4 should read as follows:

```
ipv6 pim [vrf vrf-name] bsr candidate bsr ipv6-address [hash-mask-length] [priority priority-value] [scope scope-value] [accept-rp-candidate access-list-name]
```

```
ipv6 pim [vrf vrf-name] bsr candidate rp ipv6-address [group-list access-list-name] [priority priority-value] [interval seconds] [bidir]
```

Under the Details Steps, Step 3 should read:

```
ipv6 pim [vrf vrf-name] bsr candidate bsr ipv6-address [hash-mask-length] [priority priority-value] [scope scope-value] [accept-rp-candidate access-list-name]
```

Example:

```
Device(config)# ipv6 pim bsr candidate bsr 2001:DB8:1:1:4 scope 6
```

Under the Details Steps, Step 4 should read:

```
ipv6 pim [vrf vrf-name] bsr candidate rp ipv6-address [group-list access-list-name] [priority priority-value] [interval seconds] [bidir]
```

Example:

```
Device(config)# ipv6 pim bsr candidate rp 2001:DB8:1:1:1 group-list list
```

- In the section "Configuring BSR Devices to Announce Scope-to-RP Mappings," the keyword **scope scope-value** should be removed from Step 3, both under Summary and Detail Steps.
- In the section "Additional References section," it would be helpful to reference RFC 5059.

IPv6 Multicast Listen Discovery (MLD) and Multicast Listen Discovery Snooping

MLD is a protocol used by IPv6 multicast devices to discover the presence of multicast listeners (nodes that want to receive IPv6 multicast packets) on its directly attached links and to discover which multicast packets are of interest to neighboring nodes. MLD snooping is supported in two different versions: MLD v1 and MLD v2. Network switches use MLD snooping to limit the flood of multicast traffic, causing IPv6 multicast data to be selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN. This lessens the load on devices in the network, minimizing unnecessary bandwidth on links, enabling efficient distribution of IPv6 multicast data.

For information on configuring multicast services, see [Chapter 38, "Configuring IP Multicast."](#)

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 frames feature can be configured on a per-port basis on Layer 2 and Layer 3 interfaces. The feature is supported only on the following hardware:

- WS-X4306-GB: all ports
- WS-X4232-GB-RJ: ports 1-2
- WS-X4418-GB: ports 1-2
- WS-X4412-2GB-TX: ports 13-14

- WS-4648-RJ45V-E
- WS-X4648+RJ45V+E
- WS-C4500X-16
- WS-C4500X-32
- WS-X4706-10GE linecards
- supervisor engine uplink ports

For information on Jumbo Frames, see [Chapter 9, “Configuring Interfaces.”](#)

Link Aggregation Control Protocol

LACP supports the automatic creation of EtherChannels by exchanging LACP packets between LAN ports. LACP packets are exchanged only between ports in passive and active modes. The protocol "learns" the capabilities of LAN port groups dynamically and informs the other LAN ports. After LACP identifies correctly matched Ethernet links, it facilitates grouping the links into an EtherChannel. Then the EtherChannel is added to the spanning tree as a single bridge port.

Cisco IOS XE IP Application Services Features in Cisco IOS XE 3.1.0SG

This section lists the IP Application Services software features that are supported in Cisco IOS XE 3.1.0SG. Links to the feature documentation are included.

Feature guides may contain information about more than one feature. To find information about a specific feature within a feature guide, see the Feature Information table at the end of the guide.

Feature guides document features that are supported on many different software releases and platforms. Your Cisco software release or platform may not support all the features documented in a feature guide. See the Feature Information table at the end of the feature guide for information about which features in that guide are supported in your software release. Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

IEEE 802.3ad Link Aggregation (LACP)

http://www.cisco.com/en/US/docs/ios/cether/configuration/guide/ce_inkbndl.html

Link Aggregation Control Protocol (LACP) (802.3ad) for Gigabit Interfaces

http://www.cisco.com/en/US/docs/ios/ios_xe/cether/configuration/guide/ce_inkbndl_xe.html

Link Layer Discovery Protocol

To support non-Cisco devices and to allow for interoperability between other devices, the switch supports the IEEE 802.1AB LLDP. Link Layer Discovery Protocol (LLDP) is a neighbor discovery protocol that is used for network devices to advertise information about themselves to other devices on the network. This protocol runs over the data-link layer, which allows two systems running different network layer protocols to learn about each other.

LLDP supports a set of attributes that it uses to discover neighbor devices. These attributes contain type, length, and value descriptions and are referred to as *TLVs*. LLDP supported devices can use TLVs to receive and send information to their neighbors. Details such as configuration information, device capabilities, and device identity can be advertised using this protocol.

For information on configuring LLDP, see [Chapter 32, “Configuring LLDP, LLDP-MED, and Location Service.”](#)

Link State Tracking

Link-state tracking, also known as trunk failover, is a feature that binds the link state of multiple interfaces. For example, link-state tracking provides redundancy in the network when used with server NIC adapter teaming. When server network adapters are configured in a primary or secondary relationship known as teaming, if the link is lost on the primary interface, connectivity is transparently changed to the secondary interface.

For information on configuring Link State Tracking, see [Chapter 27, “Configuring EtherChannel and Link State Tracking.”](#)

Location Service

The location service feature allows the switch to provide location and attachment tracking information for its connected devices to a Cisco Mobility Services Engine (MSE). The tracked device can be a wireless endpoint, a wired endpoint, or a wired switch or controller. The switch informs device link up and link down events through encrypted Network Mobility Services Protocol (NMSP) location and attachment notifications to the MSE.

For information on configuring LLDP, see [Chapter 32, “Configuring LLDP, LLDP-MED, and Location Service.”](#)

Multiple Spanning Tree

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 23, “Configuring STP and MST.”](#)

Per-VLAN Rapid Spanning Tree

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 23, “Configuring STP and MST.”](#)

Quality of Service

**Note**

QoS functionality on Catalyst 4900M, Catalyst 4948E, Catalyst 4948E-F, Supervisor Engine 6-E and Supervisor Engine 6L-E are equivalent.

The quality of service (QoS) feature prevents 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, including per-port per-VLAN policing
- Sharing and shaping

The 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.

Cisco Modular QoS Command-Line-Interface

Cisco Modular QoS CLI (MQC) is the framework that implements Cisco IOS software QoS. MQC allows the user to define a traffic class, create a traffic policy (containing the QoS feature to be applied to the traffic class), and attach the traffic policy to an interface. MQC is a cross-Cisco baseline that provides a consistent syntax and behavior of QoS features across multiple product families. Cisco IOS Software Release 12.2(40)SG complies to MQC for configuration of QoS features on the Supervisor Engine 6-E. MQC enables rapid deployment of new features and technology innovations and facilitates the management of network performance with respect to bandwidth, delay, jitter, and packet loss, enhancing the performance of mission-critical business applications. The rich and advanced QoS features are enabled using Cisco MQC.

Two-Rate Three-Color Policing

The Two-Rate Three-Color Policing feature (also termed *Hierarchical QoS*) limits the input or output transmission rate of a class of traffic based on user-defined criteria and marks or colors packets by setting the applicable differentiated services code point (DSCP) values. This feature is often configured on the interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. Using this feature, traffic that conforms to user-defined criteria can be sent through the interfaces, while traffic that exceeds or violates these criteria is sent out with a decreased priority setting or even dropped.

For information on QoS and Auto QoS, see [Chapter 44, “Configuring Quality of Service.”](#)

Resilient Ethernet Protocol

Resilient Ethernet Protocol (REP) is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP) to control network loops, handle link failures, and improve convergence time. REP controls a group of ports connected in a segment, ensures that the segment does not create any bridging loops, and responds to link failures within the segment. REP provides a basis for constructing more complex networks and supports VLAN load balancing.

For information on REP, see [Chapter 25, “Configuring Resilient Ethernet Protocol.”](#)

SmartPort Macros

SmartPort macros provide a convenient way to save and share common configurations. You can use SmartPort macros to enable features and settings based on the location of a switch in the network and for mass configuration deployments across the network.

For information on configuring SmartPort macros, see [Chapter 21, “Configuring SmartPort Macros.”](#)

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 23, “Configuring STP and MST.”](#)

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 26, “Configuring Optional STP Features.”](#)

Stateful Switchover

Stateful switchover (SSO) enables you to propagate 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.

For information about SSO, see [Chapter 13, “Configuring Cisco NSF with SSO Supervisor Engine Redundancy.”](#)

SVI Autostate

When an SVI has multiple ports on a VLAN, normally the SVI will go down when all the ports in the VLAN go down. You can design your network so that some ports are not counted in the calculation of SVI “going up or down.” SVI Autostate provides a knob to mark a port so that it is not counted in the SVI “going up and down” calculation and applies to all VLANs that are enabled on that port.

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.

With standard UDLD, the time to detect a unidirectional link can vary from a few seconds to several minutes depending on how the timers are configured. Link status messages are exchanged every couple of seconds. With Fast UDLD, you can detect unidirectional links in under one second (this also depends on how the timers are configured). Link status messages are exchanged every couple of hundred milliseconds.

For information about UDLD and Fast UDLD, see [Chapter 33, “Configuring UDLD.”](#)

VLANs

A VLAN configures switches and routers according to logical, rather than physical, topologies. Using VLANs, you 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, VTP, and Dynamic VLAN Membership, see [Chapter 17, “Configuring VLANs, VTP, and VMPS.”](#)

The following VLAN-related features also are 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.

- **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 47, “Configuring Private VLANs.”](#)
- **Private VLAN Trunk Ports**—Private VLAN trunk ports allow a secondary port on a private VLAN to carry multiple secondary VLANs.
- **Private VLAN Promiscuous Trunk Ports**—Private VLAN promiscuous trunk extends the promiscuous port to a 802.1Q trunk port, carrying multiple primary VLANs (hence multiple subnets). Private VLAN promiscuous trunk is typically used to offer different services or content on different primary VLANs to isolated subscribers. Secondary VLANs can not be carried over the private VLAN promiscuous trunk.
- **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.

Virtual Switching Systems (Catalyst 4500-X and Supervisor Engine 7-E, 7L-E, and 8-E)

Network operators increase network reliability by configuring switches and by provisioning links to the redundant pairs. Redundant network elements and redundant links can add complexity to network design and operation. Virtual switching simplifies the network by reducing the number of network elements and hiding the complexity of managing redundant switches and links.

A VSS combines a pair of Catalyst 4500 or 4500-X series switches into a single network element. The VSS manages the redundant links, which externally act as a single port channel. Starting with Cisco Release IOS XE 3.4.0SG, the Catalyst 4500 or 4500-X series switches support VSS.



Note

Smart Install Director is not supported with VSS.

For information on VSS, see [Chapter 5, “Configuring Virtual Switching Systems.”](#)

Virtual Switch System Client

Catalyst 4500 series switches support enhanced PAgP. If a Catalyst 4500 series switch is connected to a Catalyst 6500 series Virtual Switch System (VSS) with a PAgP EtherChannel, the Catalyst 4500 series switch will automatically serve as a VSS client, using enhanced PAgP on this EtherChannel for dual-active detection. This VSS client feature has no impact on the performance of Catalyst 4500 series switches and does not require any user configuration.

For more details, see [Chapter 27, “Configuring EtherChannel and Link State Tracking.”](#)

Y.1731 (AIS and RDI)

Y.1731 ETH-AIS (Ethernet Alarm Indication Signal function) and ETH-RDI (Ethernet Remote Defect Indication function) provides fault and performance management for service providers in large networks.

ETH-AIS suppresses alarms following detection of defect conditions at the server (sub) layer. Due to independent restoration capabilities provided within the Spanning Tree Protocol (STP) environments, ETH-AIS is not expected to be applied in the STP environments. In this case, AIS is configurable, and the administrator describes how to enable and disable AIS in STP environment or not.

ETH-RDI can be used by a MEP to communicate to its peer MEPs that a defect condition has been encountered. ETH-RDI is used only when ETH-CC transmission is enabled.

For information about Y.1731, see [Chapter 76, “Configuring Y.1731 \(AIS and RDI\).”](#)

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 wire-speed 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 by using application-specific integrated circuit (ASIC) hardware instead of microprocessor-based engines.

The following sections describe the key Layer 3 switching software features on the switch:

- [Bidirectional Forwarding Detection, page 1-14](#)
- [Cisco Express Forwarding, page 1-14](#)
- [Device Sensor, page 1-14](#)
- [EIGRP Stub Routing, page 1-14](#)
- [Enhanced Object Tracking, page 1-15](#)
- [GLBP, page 1-15](#)
- [HSRP, page 1-16](#)
- [In Service Software Upgrade, page 1-20](#)
- [IP Routing Protocols, page 1-17](#)
- [IPv6, page 1-20](#)
- [Multicast Services, page 1-20](#)
- [NSF with SSO, page 1-21](#)
- [OSPF for Routed Access, page 1-22](#)
- [Policy-Based Routing, page 1-22](#)
- [Unicast Reverse Path Forwarding, page 1-22](#)
- [Unicast Reverse Path Forwarding, page 1-22](#)
- [Unidirectional Link Routing, page 1-23](#)
- [VRF-lite, page 1-23](#)
- [Virtual Router Redundancy Protocol, page 1-23](#)

Bidirectional Forwarding Detection

**Note**

Starting with Cisco IOS Release IOS 15.1(1)SG, Bidirectional Forwarding Detection (BFD) support was introduced on Catalyst 4900M, Catalyst 4948E, and Catalyst 4948E-F Ethernet switches. With Cisco IOS XE 3.5.0E and IOS 15.2(1)E, supported was extended to Supervisor Engine 7-E, and Supervisor Engine 7L-E. With Cisco IOS XE 3.6.0E and IOS 15.2(2)E, supported was extended to Supervisor Engine 8-E.

Bidirectional Forwarding Detection (BFD) protocol. BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. It includes a description of how to configure multihop BFD sessions. BFD provides a consistent failure detection method for network administrators in addition to fast forwarding path failure detection.

For information on configuring BFD, see [Chapter 40, “Configuring Bidirection Forwarding Detection.”](#)

Cisco Express Forwarding

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 36, “Configuring Cisco Express Forwarding.”](#)

Device Sensor

Device Sensor uses protocols such as Cisco Discovery Protocol (CDP), Link Layer Discovery Protocol (LLDP), and DHCP to obtain endpoint information from network devices and make this information available to its clients. Device Sensor has internal clients, such as the embedded Device Classifier (local analyzer), Auto Smartports (ASP), MediaNet Service Interface (MSI)-Proxy, and EnergyWise. Device Sensor also has an external client, Identity Services Engine (ISE), which uses RADIUS accounting to receive and analyze endpoint data. When integrated with ISE, Device Sensor provides central policy management and device-profiling capabilities.

For more information on Device Sensor, see [Chapter 49, “Configuring 802.1X Port-Based Authentication.”](#)

EIGRP Stub Routing

The EIGRP stub routing feature, available in all images, reduces resource utilization by moving routed traffic closer to the end user.

The IP base image contains only EIGRP stub routing. The IP services image contains complete EIGRP routing.

In a network using EIGRP stub routing, the only route for IP traffic to follow to the user is through a switch that is configured with EIGRP stub routing. The switch sends the routed traffic to interfaces that are configured as user interfaces or are connected to other devices.

For information on configuring EIGRP Stub Routing, see [Chapter 35, “Configuring Layer 3 Interfaces.”](#)

Enhanced Object Tracking

Before the introduction of the Enhanced Object Tracking feature, the Hot Standby Router Protocol (HSRP) had a simple tracking mechanism that allowed you to track the interface line-protocol state only. If the line-protocol state of the interface went down, the HSRP priority of the router was reduced, allowing another HSRP router with a higher priority to become active.

The Enhanced Object Tracking (EOT) feature separates the tracking mechanism from HSRP and creates a separate standalone tracking process that can be used by other Cisco IOS processes as well as HSRP. This feature allows tracking of other objects in addition to the interface line-protocol state.

A client process, such as HSRP, Virtual Router Redundancy Protocol (VRRP), or Gateway Load Balancing Protocol (GLBP), can now register its interest in tracking objects and then be notified when the tracked object changes state.

For details on EOT, refer to this URL:

For platform specific information on Enhanced Object Tracking, see [Chapter 69, “Configuring Enhanced Object Tracking.”](#)

For more detailed information on Enhanced Object Tracking, see the URL:

<http://www.cisco.com/en/US/docs/ios-xml/ios/ipapp/configuration/12-4t/iap-eot.html>

GLBP

The Gateway Load Balancing Protocol (GLBP) feature provides automatic router backup for IP hosts configured with a single default gateway on a LAN. Multiple first hop routers on the LAN combine to offer a single virtual first hop IP router while sharing the IP packet forwarding load. GLBP devices share packet-forwarding responsibilities, optimizing resource usage, thereby reducing costs. Other routers on the LAN may act as redundant GLBP routers that will become active if any of the existing forwarding routers fail. This improves the resiliency of the network and reduces administrative burden.

For details on GLBP, refer to this URL:

http://www.cisco.com/en/US/docs/ios/12_2t/12_2t15/feature/guide/ft_glbp.html

Cisco IOS XE IP Application Services Features in Cisco IOS XE 3.1.0SG

This section lists the IP Application Services software features that are supported in Cisco IOS XE 3.1.0SG. Links to the feature documentation are included.

Feature guides may contain information about more than one feature. To find information about a specific feature within a feature guide, see the Feature Information table at the end of the guide.

Feature guides document features that are supported on many different software releases and platforms. Your Cisco software release or platform may not support all the features documented in a feature guide. See the Feature Information table at the end of the feature guide for information about which features in that guide are supported in your software release. Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Gateway Load Balancing Protocol (GLBP), GLBP MD5 Authentication

http://www.cisco.com/en/US/docs/ios/12_2sx/12_2sxh/feature/guide/sxglbpm.html

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/tech/tk648/tk362/tk321/tsd_technology_support_sub-protocol_home.html

Cisco IOS XE IP Application Services: HSRP Features in Cisco IOS XE 3.1.0SG

This section lists the IP Application Services:HSRP software features that are supported in Cisco IOS XE 3.1.0SG. Links to the feature documentation are included.

Feature guides may contain information about more than one feature. To find information about a specific feature within a feature guide, see the Feature Information table at the end of the guide.

Feature guides document features that are supported on many different software releases and platforms. Your Cisco software release or platform may not support all the features documented in a feature guide. See the Feature Information table at the end of the feature guide for information about which features in that guide are supported in your software release. Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

HSRP—Hot Standby Router Protocol

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

HSRP MD5 Authentication

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

HSRP Support for ICMP Redirects

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

IP Precedence Accounting

http://www.cisco.com/en/US/docs/ios/12_2/ipaddr/command/reference/1rfip2.html

ISSU—HSRP

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

SSO—HSRP

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

SSO Aware HSRP

SSO Aware HSRP offers continuous data packet forwarding during a supervisor engine switchover without a path change to the standby HSRP router. During supervisor engine switchover, NSF with SSO continues forwarding data packets along known routes using the HSRP virtual IP address. When both supervisor engines fail on the active HSRP router, the standby HSRP router takes over as the active HSRP router. It further extends reliability and availability offered by the NSF with SSO to Layer 3. SSO aware HSRP is available for Supervisor Engine IV, V, and V-10GE on Catalyst 4507R and 4510R chassis with supervisor redundancy.

NHRP

Beginning in Cisco IOS XE Release 3.7.1E, NHRP is supported on Catalyst 4500 series switches.

The Next Hop Resolution Protocol (NHRP) is an ARP-like protocol that dynamically maps a Non-Broadcast Multi-Access (NBMA) network. With NHRP, systems attached to an NBMA network can dynamically learn the NBMA (physical) address of the other systems that are part of that network, allowing these systems to directly communicate.

NHRP is a client and server protocol where the hub is the Next Hop Server (NHS) and the spokes are the Next Hop Clients (NHCs). Catalyst 4500 series switches act as NHRP clients that communicate with the NHRP hub for registration, and to request the resolution of other spoke addresses.

**Note**

Catalyst 4500 series switches cannot be used as NHRP hubs.

For information on configuring NHRP, see

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_nhrp/configuration/xe-3s/asr1000/nhrp-xe-3s-asr1000-book/config-nhrp.html

IP Routing Protocols

The following routing protocols are supported on the Catalyst 4500 series switch:

- [BGP, page 1-17](#)
- [EIGRP, page 1-18](#)
- [IS-IS, page 1-18](#)
- [OSPF, page 1-19](#)
- [RIP, page 1-19](#)

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.

BGP Route-Map Continue

The BGP Route-Map Continue feature introduces the continue clause to the BGP route-map configuration. The continue clause provides more programmable policy configuration and route filtering. It introduces the capability to execute additional entries in a route map after an entry is executed with successful match and set clauses. Continue clauses allow configuring and organizing more modular policy definitions to reduce the number of policy configurations that are repeated within the same route map.

For details on BGP, refer to this URL:

http://www.cisco.com/en/US/docs/ios/12_4t/ip_route/configuration/guide/t_brbbas.html

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).



Note

Customers can configure Enhanced Interior Gateway Routing Protocol (EIGRP) to route IPv6 prefixes. EIGRP configuration and protocol behavior for both IPv4 and IPv6 prefixes are similar, providing operational familiarity and continuity. EIGRP support for IPv6 will enable customers to use their existing EIGRP knowledge and processes, allowing them to deploy an IPv6 network at a low cost.

For details on EIGRP, refer to this URL:

http://www.cisco.com/en/US/products/ps6630/products_ios_protocol_option_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.

For details on IS-IS, refer to this URL:

http://www.cisco.com/en/US/products/ps6632/products_ios_protocol_option_home.html

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 uses 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.



Note

For Catalyst 4500 series switch, usage of OSPF aggressive timers may lead to session transition and packet loss during CPU-intensive operation. It is recommended to use either default OSPF timers or BFD with hardware offload functionality.

For details on OSPF, refer to this URL:

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

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.

For details on RIP, refer to this URL:

http://www.cisco.com/en/US/tech/tk365/tk554/tsd_technology_support_sub-protocol_home.html

In Service Software Upgrade

SSO requires the same version of Cisco IOS on both the active and standby supervisor engines. Because of version mismatch during an upgrade or downgrade of the Cisco IOS software, a Catalyst 4500 series switch is forced into operating in RPR mode. In this mode, after the switchover you can observe link-flaps and a disruption in service. This issue is solved by the In-Service Software Upgrade (ISSU) feature that enables you to operate in SSO/NSF mode while performing software upgrade or downgrade.

ISSU allows an upgrade or downgrade of the Catalyst IOS or IOS XE images at different release levels on the both the active and standby supervisor engines by utilizing the Version Transformation Framework between the stateful components running on each supervisor engine.

For details on Cisco IOS ISSU, refer to [Chapter 7, “Configuring the Cisco IOS In-Service Software Upgrade Process.”](#)

For details on Cisco IOS XE ISSU, refer to [Chapter 8, “Configuring the Cisco IOS XE In Service Software Upgrade Process.”](#)

IPv6

IPv6 provides services such as end-to-end security, quality of service (QoS), and globally unique addresses. The IPv6 address space reduces the need for private addresses and Network Address Translation (NAT) processing by border routers at network edges.

For more information about IPv6 services supported on a Catalyst 4500 series switch, see [Chapter 63, “Support for IPv6.”](#)

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:

- ANCP Client —ANCP Multicast enables you to control multicast traffic on a Catalyst 4500 switch using either ANCP (rather than IGMP) or direct static configuration on the CLI.
- 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.

The IGMP Snooping Querier is a Layer 2 feature required to support IGMP snooping in a VLAN where PIM and IGMP are not configured because the multicast traffic does not require routing.

For information on configuring IGMP snooping, see [Chapter 28, “Configuring IGMP Snooping and Filtering, and MVR.”](#)

- IPv6 Multicast Listen Discovery (MLD) and Multicast Listen Discovery snooping—MLD is a protocol used by IPv6 multicast devices to discover the presence of multicast listeners (nodes that want to receive IPv6 multicast packets) on its directly attached links and to discover which multicast packets are of interest to neighboring nodes. MLD snooping is supported in two different versions: MLD v1 and MLD v2. Network switches use MLD snooping to limit the flood of multicast traffic, causing IPv6 multicast data to be selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN. This lessens the load on devices in the network, minimizing unnecessary bandwidth on links, enabling efficient distribution of IPv6 multicast data.

For information on configuring multicast services, see [Chapter 29, “Configuring IPv6 Multicast Listener Discovery Snooping.”](#)

- 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 PIM-SSM mapping, see the URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.1SY/config_guide/sup720/ude_udlr.html

- IP Multicast Load Splitting (Equal Cost Multipath (ECMP) Using S, G and Next Hop)—IP Multicast Load Splitting introduces more flexible support for ECMP multicast load splitting by adding support for load splitting based on source and group address and on source, group, and next-hop address. This feature allows multicast traffic from devices that send many streams to groups or that broadcast many channels, such as IPTV servers or MPEG video servers, to be more effectively load shared across equal-cost paths.

For information on configuring multicast services, see [Chapter 38, “Configuring IP Multicast.”](#)

NSF with SSO

Non-Stop Forwarding with Stateful Switchover (NSF/SSO) offers continuous data packet forwarding in a Layer 3 routing environment during supervisor engine switchover. During supervisor engine switchover, NSF/SSO continues forwarding data packets along known routes while the routing protocol information is recovered and validated, avoiding unnecessary route flaps and network instability. With NSF/SSO, IP phone calls do not drop. NSF/SSO is supported for OSPF, BGP, EIGRP, IS-IS, and Cisco Express Forwarding (CEF). NSF/SSO is typically deployed in the most critical parts of an enterprise or service provider network, such as Layer 3 aggregation/core or a resilient Layer 3 wiring closet design. It is an essential component of single chassis deployment for critical applications. NSF/SSO is available for all shipping supervisor engines on Catalyst 4507R and 4510R chassis with supervisor redundancy.



Note

With the IP Base image, NSF is supported with EIGRP-stub routing and OSPF.



Note

With the Enterprise Services image, NSF is supported on all routing protocols except for RIP.



Note

The LAN Base image does not support NSF.

For information on NSF with SSO, see [Chapter 13, “Configuring Cisco NSF with SSO Supervisor Engine Redundancy.”](#)

OSPF for Routed Access

OSPF for Routed Access is designed specifically to enable customers to extend Layer 3 routing capabilities to the access or wiring closet.

**Note**

OSPF for Routed Access supports only one OSPFv2 and one OSPFv3 instance with a maximum number of 1000 dynamically learned routes.

With the typical topology (hub and spoke) in a campus environment, where the wiring closets (spokes) are connected to the distribution switch (hub) forwarding all nonlocal traffic to the distribution layer, the wiring closet switch does not need to hold a complete routing table. Ideally, the distribution switch sends a default route to the wiring closet switch to reach inter-area and external routes (OSPF stub or totally stub area configuration).

Refer to the following link for more details:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/routed-ex.html>

With Cisco IOS Release 12.2(53)SG, the IP Base image supports OSPF for routed access. The Enterprise Services image is required if you need multiple OSPFv2 and OSPFv3 instances without route restrictions. Enterprise Services also is required to enable the VRF-lite feature.

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.

Starting with Release IOS XE 3.4.0SG and IOS 15.1(2)SG, the PBR Recursive Next Hop feature enhances route maps to enable configuration of a recursive next-hop IP address. The recursive next-hop IP address can be a subnet that is not directly connected. The routing table will be looked up to find the directly connected next-hop to which to send the packet so that it is routed via the recursive next-hop that has been configured.

For more information on policy-based routing, see [Chapter 42, “Configuring Policy-Based Routing.”](#)

Unicast Reverse Path Forwarding

The Unicast Reverse Path Forwarding (Unicast RPF) feature helps to mitigate problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address.

For information on URPF, see [Chapter 37, “Configuring Unicast Reverse Path Forwarding.”](#)

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 URL

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/15.1SY/config_guide/sup720/ude_udlr.html

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.

Prior to Release IOS XE 3.5.0E and IOS 15.2(1)E, only IPv4 was available. With Release IOS XE 3.5.0E and IOS 15.2(1)E, VRF-lite support has been extended to IPv6.

For information on VRF-lite, see [Chapter 43, “Configuring VRF-lite.”](#)

Virtual Router Redundancy Protocol

Virtual Router Redundancy Protocol (VRRP) is a standard based first-hop redundancy protocol. With VRRP, a group of routers function as one virtual router by sharing one virtual IP address and one virtual MAC address. The master router performs packet forwarding, while the backup routers stay idle. VRRP is typically used in the multi-vendor first-hop gateway redundancy deployment.

For details on VRRP, refer to this URL:

http://www.cisco.com/en/US/docs/ios/ipapp/configuration/guide/ipapp_vrrp_ps6441_TSD_Products_Configuration_Guide_Chapter.html

Management Features

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

- [Cisco Call Home, page 1-24](#)
- [Cisco Energy Wise, page 1-25](#)
- [Cisco IOS IP Service Level Agreements, page 1-25](#)
- [Cisco Media Services Proxy, page 1-25](#)
- [Cisco Medianet AutoQoS, page 1-26](#)
- [Cisco Medianet Flow Metadata, page 1-26](#)

- [Cisco IOS Mediatrace and Performance Monitor, page 1-27](#)
- [Cisco Network Assistant, page 1-28](#)
- [Dynamic Host Control Protocol, page 1-28](#)
- [Easy Virtual Network, page 1-29](#)
- [Embedded CiscoView, page 1-29](#)
- [Embedded Event Manager, page 1-29](#)
- [Ethernet Management Port, page 1-30](#)
- [File System Management \(Supervisor Engine 7-E, 7L-E, and 8-E\), page 1-30](#)
- [FAT File Management System on Supervisor Engine 6-E, Supervisor Engine 6L-E, Catalyst 4948E, Catalyst 4948E-F, and Catalyst 4900M, page 1-30](#)
- [Forced 10/100 Autonegotiation, page 1-30](#)
- [Intelligent Power Management, page 1-31](#)
- [MAC Address Notification, page 1-31](#)
- [MAC Notify MIB, page 1-31](#)
- [NetFlow-lite, page 1-31](#)
- [Power over Ethernet, page 1-31](#)
- [Secure Shell, page 1-32](#)
- [Simple Network Management Protocol, page 1-32](#)
- [Smart Install, page 1-32](#)
- [SPAN and RSPAN, page 1-32](#)
- [Universal Power over Ethernet, page 1-33](#)
- [Web Content Coordination Protocol, page 1-33](#)
- [Wireshark, page 1-33](#)
- [XML-PI, page 1-34](#)

Cisco Call Home

Call Home provides e-mail-based and web-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications. Common uses of this feature may include direct paging of a network support engineer, e-mail notification to a Network Operations Center, XML delivery to a support website, and utilization of Cisco Smart Call Home services for direct case generation with the Cisco Systems Technical Assistance Center (TAC).

The Call Home feature can deliver alert messages containing information on configuration, diagnostics, environmental conditions, inventory, and syslog events.

For more information on Call Home, see [Chapter 77, “Configuring Call Home.”](#)

Cisco Energy Wise

Cisco EnergyWise is an energy-management technology added onto Cisco switching solutions to help you measure, report, and reduce energy consumption across your entire infrastructure. With EnergyWise's management interface, network management applications can communicate with endpoints and each other, using the network as the unifying fabric.

For details refer to the URLs:

http://www.cisco.com/en/US/docs/switches/lan/energywise/phase2/ios/configuration/guide/ew_v2.html

http://www.cisco.com/en/US/docs/switches/lan/energywise/phase2_5/ios/release/notes/ol23554.html#wp604941

Cisco IOS IP Service Level Agreements

Cisco IOS IP Service Level Agreements (SLAs) allows Cisco customers to analyze IP service levels for IP applications and services by using active traffic monitoring—the generation of traffic in a continuous, reliable, and predictable manner—for measuring network performance. With Cisco IOS IP SLA, service provider customers can measure and provide service level agreements, and enterprise customers can verify service levels, verify outsourced service level agreements, and understand network performance. Cisco IOS IP SLAs can perform network assessments, verify quality of service (QoS), ease the deployment of new services, and assist with network troubleshooting.

For platform-specific information on Cisco IOS IP SLA, see [Chapter 78, “Configuring Cisco IOS IP SLA Operations.”](#)

For more detail on Cisco IOS IP SLAs, see the *Cisco IOS IP SLAs Configuration Guide, Release 12.4T*: http://www.cisco.com/en/US/docs/ios/ipsla/configuration/guide/12_4t/sla_12_4t_book.html

A Catalyst 4500 series switch also supports a Built-in Traffic Simulator using Cisco IOS IP SLAs video operations to generate synthetic traffic for a variety of video applications, such as Telepresence, IPTV and IP video surveillance camera. You can use the simulator tool:

- for network assessment before deploying applications that have stringent network performance requirements.
- along with the Cisco IOS Mediatrace for post-deployment troubleshooting for any network related performance issues.

The traffic simulator includes a sophisticated scheduler that allows the user to run several tests simultaneously or periodically and over extended time periods. (Supported only on switches running the Enterprise Services feature set.)

For information on configuring this feature, see the *Configuring Cisco IOS IP SLAs Video Operations* document at:

http://www.cisco.com/en/US/docs/ios-xml/ios/ipsla/configuration/12-2se/sla_video.html

Cisco Media Services Proxy

The Media Services Proxy (MSP) feature identifies various media end points in the network automatically and renders appropriate media services. It acts as a layer that connects appropriate devices with their respective network services automatically.

MSP follows a network-centric model, where the access switches and routers learn information about devices and flow by using mechanisms such as Cisco Discovery Protocol (formerly known as CDP) and DHCP, or by snooping on key protocol packets such as the Session Initiation Protocol (SIP) and H.323. Modifications to the endpoints are not required to achieve the information learning. After the information is gleaned, MSP provides appropriate services to the network devices.

Following are the benefits of MSP:

- Automatic identification of devices and flow in the network.
- Application of appropriate services to the endpoints.
- Configuration control for the administrator, thereby reducing the manual configuration and management of services. For example, configuring the Resource Reservation Protocol (RSVP) in the network for video applications requiring guaranteed bandwidth.



Note The system cannot scale to greater than 512 SIP flows with MSP and Flow Metadata enabled.

For information on configuring this feature, refer to the following documents:

<http://www.cisco.com/en/US/docs/ios-xml/ios/msp/configuration/15-1sg/med-ser-prxy.html>

<http://www.cisco.com/en/US/docs/ios-xml/ios/msp/configuration/xe-3sg/med-ser-prxy-xe.html>

Cisco Medianet AutoQoS

Cisco Medianet AutoQoS provides a default configuration to ease the process of enabling QoS on switches. This process can be difficult given the functional/behavioral differences in QoS across different platforms. This functionality extends AutoQoS functionality for the Catalyst 4500 to support video traffic as well as other kinds of traffic.

The goal of AutoQoS is to simplify the work customers have to undertake while configuring their networks to support QoS. This is done by automating QoS configurations to handle various classes of traffic. AutoQoS for Medianet provides commands, which act as macros that call existing CLI commands to implement desired configurations. You are required to specify the type of device (PC, another switch, ip camera, etc.) connected to a given interface. AutoQoS for Medianet applies a default QoS configuration to that interface, which you can later fine-tune as needed.

For details, refer to [Chapter 44, “Configuring Quality of Service.”](#)

Cisco Medianet Flow Metadata

Flow Metadata is the data that qualifies other data. Flow Metadata aids in supporting an intelligent network by making the network aware about the type, nature, and characteristics of the media stream that flows in the network. Flow Metadata also allows for the network to apply policies on the media streams. Across the Medianet system, Flow Metadata is produced, transported, stored, retrieved, and acted on consistently by a wide variety of Medianet services.

The Flow Metadata infrastructure provides a framework that allows data from one component be available to another component on the same network element as well as across network elements.

Flow Metadata is supported on releases prior to Cisco IOS Release 15.1(1)SG. Flow metadata is the data that describes a flow in the network. This Flow Metadata describes the five tuple flow along with the attributes. Network elements can take action based on the Flow Metadata generated by the endpoints.

The Flow Metadata infrastructure consists of two major components: producers and consumers.

- Flow Metadata producer is any source of Flow Metadata. The producer propagates all the attributes of a given flow. Producers can be anywhere in the network: endpoint, proxy agents, or intermediate nodes. Currently, Flow Metadata generated by the endpoints is supported. Producers use a specific transport protocol, such as RSVP for signalling the Flow Metadata attributes and store the information in a database, referred to as the control plane database, which can then be used by the consumers.
- Flow Metadata consumer is any network element that uses the flow tuple and Flow Metadata provided by the producers. The flow tuple and Flow Metadata can also be propagated along the media path to consumers in different network elements via a transport infrastructure.

For configuration details, refer to the following URLs:

<http://www.cisco.com/en/US/docs/ios-xml/ios/mdata/configuration/xe-3sg/metadata-framework.html>

<http://www.cisco.com/en/US/docs/ios-xml/ios/mdata/configuration/15-1sg/metadata-framework.html>

For details on the Flow Metadata commands, refer to the following URL:

<http://www.cisco.com/en/US/docs/ios-xml/ios/qos/command/qos-cr-book.html>

Cisco IOS Mediatrace and Performance Monitor

Cisco IOS Mediatrace helps to isolate and troubleshoot network degradation problems by enabling a network administrator to discover an IP flow's path, dynamically enable monitoring capabilities on the nodes along the path, and collect information on a hop-by-hop basis. This information includes, among other things, flow statistics; utilization information for incoming and outgoing interfaces, CPUs, and memory; as well as any changes to IP routes or the Cisco IOS Mediatrace monitoring state.

For details, see the following URLs:

http://www.cisco.com/en/US/docs/ios-xml/ios/media_monitoring/configuration/15-1sg/mm-pasv-mon.html

http://www.cisco.com/en/US/docs/ios-xml/ios/media_monitoring/configuration/xe-3sg/mm-pasv-mon.html

http://www.cisco.com/en/US/docs/ios/media_monitoring/command/reference/mm_book.html

http://www.cisco.com/en/US/docs/ios-xml/ios/media_monitoring/configuration/15-1sg/mm-mediatriace.html

http://www.cisco.com/en/US/docs/ios-xml/ios/media_monitoring/configuration/xe-3sg/mm-mediatriace.html

Configuration guidelines for Cisco IOS Mediatrace and Performance Monitor include the following:

- Video monitoring is supported only on physical ports.

Limitations for Cisco IOS Mediatrace and Performance Monitor on a Catalyst 4500 series switch include the following:

- Both features can only be configured to monitor ingress traffic.
- Packets cannot be monitored by both CEure and the rxSPAN session with encapsulation. The first-applied configuration takes precedence.
- Not all packets received by an interface can be monitored. After a packet is received by an ingress interface, it might be either unable to make a forwarding decision or dropped at various stages because of configured security features (like IP Source Guard). The switch attempts to monitor packets close to the switch, but only those that are not dropped before the input classification stage can be monitored.

- CPU utilization is impacted when you monitor a high traffic rate. After the internally-determined threshold is crossed, monitored packets are dropped although the original packet is forwarded in hardware intact. Starting with Release IOS XE 3.3.0SG and IOS 15.1(1)SG, monitored packets might be dropped if any of the following apply:
 - The packet rate exceeds 512 PPS per flow.
 - The aggregated bandwidth of the monitor traffic exceeds 10Mbps.
 - Resources are insufficient to enqueue a new monitored packet.

When monitored packets are dropped, the *monitor event* will be set to TRUE if the flow record contains *collect monitor event*. If one minute passes with no new drops, the *monitor event* is set to FALSE but is not reflected in the output of the **show performance monitor status** until the new monitor interval starts.

monitor event is a global flag. This means that any packet drops that would trigger "monitor event" be set to TRUE for all monitored flows at that monitor interval. If a metric depends on the collection of continuous packets, the accuracy of that metric might be impacted when a *monitor event* is TRUE.

Cisco Network Assistant

Cisco Network Assistant manages standalone devices, clusters of devices, or federations of devices from anywhere in your intranet. Using its graphical user interface, you can perform multiple configuration tasks without having to remember command-line interface commands. Embedded 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.

For more information on Cisco Network Assistant, see [Chapter 16, “Configuring the Catalyst 4500 Series Switch with Cisco Network Assistant.”](#)

Dynamic Host Control Protocol

The Catalyst 4500 series switch uses DHCP in the following ways:

- 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.
- Dynamic Host Control Protocol autoconfiguration—With this feature your switch (the DHCP client) is automatically configured at startup with IP address information and a configuration file.

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

Easy Virtual Network

Easy Virtual Network (EVN) is an IP-based virtualization technology that provides end-to-end virtualization of the network. You can use a single IP infrastructure to provide separate virtual networks whose traffic paths remain isolated from each other. Configure Easy Virtual Network to configure two or more virtual IP networks.

For details on EVN, refer to the following URLs:

<http://www.cisco.com/en/US/docs/ios-xml/ios/evn/configuration/xe-3s/evn-xe-3s-book.html>

The following restrictions/feature interactions apply:

- Multicast

When multicast traffic traverses VRFs, where source and receivers are in different VRFs, multicast counters will not increment on receivers' VRFs.
- NetFlow

When configured on an EVN trunk interface, NetFlow captures traffic information for all VRFs but does not preserve the VRF information.
- SPAN
 - When an EVN trunk interface is configured as a SPAN source, traffic belonging to all VRFs carried by the EVN trunk is spanned. By default, the VNET tag is not preserved. To preserve it, configure SPAN destination with the **encapsulation dot1q** option.
 - To span traffic belonging to specific VRFs on an EVN trunk, configure **filter vlan** on the SPAN session with the corresponding VNET tags as `vlan_ids` and configure VLANs specified in **filter vlan**.
 - To span traffic in specific VRFs on all interfaces, configure **vlan** as the SPAN source with VNET tags as `vlan_ids` and configure VLANs specified as sources.
 - If `cpu` is configured as a SPAN source, then transmit packets that will be spanned are tagged by default. If the **encapsulation dot1q** option is set on the SPAN session, then the `cpu` transmitted packets, which are spanned, are double tagged.

Refer to [Chapter 66, “Configuring SPAN and RSPAN,”](#) for information on configuring SPAN sessions.

Embedded CiscoView

A web-based tool to configure the Catalyst 4500 series switch. Embedded 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.

For more information on Embedded CiscoView, see [Chapter 4, “Administering the Switch.”](#)

Embedded Event Manager

Embedded Event Manager (EEM) is a distributed and customized approach to event detection and recovery offered directly in a Cisco IOS device. EEM offers the ability to monitor events and take informational, corrective, or any desired EEM action when the monitored events occur or when a threshold is reached. An EEM policy is an entity that defines an event and the actions to be taken when that event occurs.

For information on EEM, see the URL:

http://www.cisco.com/en/US/products/ps6815/products_ios_protocol_group_home.html

Ethernet Management Port

The Ethernet management port, also referred to as the *Fal* or *fastethernet1* port, is a Layer 3 host port to which you can connect a PC. You can use the Ethernet management port instead of the switch console port for network management. When managing a switch stack, connect the PC to the Ethernet management port on a Catalyst 4500 series switch.

For more information on Ethernet management port, see the “Using the Ethernet Management Port” section in [Chapter 9, “Configuring Interfaces.”](#)

File System Management (Supervisor Engine 7-E, 7L-E, and 8-E)

The format command for IOS XE 3.1.0SG changed slightly compared to the classic IOS format because the later does not support ext2 format.

For USB flash under IOS XE 3.1.0SG, there are 3 optional formats, i.e. FAT16, FAT32 and EXT2:

```
Switch# format usb0: ?
      FAT16  FAT16 filesystem type
      FAT32  FAT32 filesystem type
      ext2   ext2  filesystem type
```

For SD card under IOS XE 3.1.0SG, the default format is FAT16:

```
Switch# format slaveusb0: ?
      FAT16  FAT16 filesystem type
      FAT32  FAT32 filesystem type
      ext2   ext2  filesystem type
```

FAT File Management System on Supervisor Engine 6-E, Supervisor Engine 6L-E, Catalyst 4948E, Catalyst 4948E-F, and Catalyst 4900M

The FAT file system is widely used to manage files on devices disks and flash. The support of the FAT file system allows you to easily remove, add, and/or transfer images to and from the flash.

Forced 10/100 Autonegotiation

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.

Intelligent Power Management

Working with powered devices (PDs) from Cisco, this feature uses power negotiation to refine the power consumption of an 802.3af-compliant PD beyond the granularity of power consumption provided by the 802.3af class. Power negotiation also enables the backward compatibility of newer PDs with older modules that do not support either 802.3af or high-power levels as required by IEEE standard.

For more information on Intelligent Power Management, see the “Intelligent Power Management” section in [Chapter 15, “Configuring Power over Ethernet.”](#)

MAC Address Notification

MAC address notification monitors the MAC addresses that are learned by, aged out, or removed from the Catalyst 4500 series switch. Notifications are sent out or retrieved by using the CISCO-MAC-NOTIFICATION MIB. It is typically used by a central network management application to collect such MAC address notification events for host moves. User-configurable MAC table utilization thresholds can be defined to notify any potential DoS or man-in-the-middle attack.

For information on MAC Address Notification, see [Chapter 4, “Administering the Switch.”](#)

MAC Notify MIB

The MAC Notify MIB feature monitors network performance, utilization, and security conditions enabling a network administrator to track the MAC addresses that are learned or removed on the switch forwarding the Ethernet frames.

NetFlow-lite

**Note**

NetFlow-lite is only supported on the Catalyst 4948E and Catalyst 4948E-F Ethernet switches.

The NetFlow-lite feature is based on ingress packet sampling at a monitoring point that can be an interface on the switch. By exporting NetFlow sampled packets, it provides visibility into traffic that is switched through the device. The rate at which input packets are sampled is configurable and a wide range of sampling rates are supported. Each sampled packet is exported as a separate NetFlow data record in the data path. NetFlow V9 and V10(IPFIX) export formats are supported.

For more information on NetFlow statistics, see [Chapter 73, “Configuring NetFlow-lite.”](#)

Power over Ethernet

Power over Ethernet (PoE) allows the LAN switching infrastructure to provide power to an endpoint (“powered device”) over a copper Ethernet cable. This capability, once referred to as “inline power,” was originally developed by Cisco in 2000 to support emerging IP telephony deployments.

IP telephones need power for operation, and Power over Ethernet supports scalable, manageable power delivery and simplifies IP telephony deployments. As wireless networking emerged, Power over Ethernet began powering wireless devices in locations where local power access did not exist.

For more information on Power over Ethernet, see [Chapter 15, “Configuring Power over Ethernet.”](#)

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.

Simple Network Management Protocol

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, such as cat4000-i5k91s-mz)
- SNMP trap message enhancements—Additional information with certain SNMP trap messages, including spanning-tree topology change notifications and configuration change notifications

For more information on SNMP, see [Chapter 72, “Configuring SNMP.”](#)

Smart Install

Beginning with Cisco IOS XE 3.4.0SG and 15.1(2)SG, Catalyst 4500 series switches supported Smart Install, which is a plug-and-play configuration and image-management feature that provides zero-touch deployment for new switches. You can ship a switch to a location, place it in the network and power it on with no configuration required on the device.

For details on Smart Install, see the URL:

http://www.cisco.com/en/US/docs/switches/lan/smart_install/configuration/guide/smart_install.html

SPAN and RSPAN

Switched Port Analyzer (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 66, “Configuring SPAN and RSPAN.”](#)

Remote SPAN (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 66, “Configuring SPAN and RSPAN.”](#)

Universal Power over Ethernet

The IEEE 802.3 Power over Ethernet (PoE) standard sets the maximum power that can be sourced by data terminal equipment (DTE) at 30W. This power is sourced over two pairs out of the four twisted pairs of conductors in a Class D, or better, cabling as specified in ISO/IEC 11801:1995.

Cisco® Universal Power over Ethernet (UPOE) is a Cisco proprietary technology that extends the IEEE 802.3 PoE standard to provide the capability to source up to 60W of power over standard Ethernet cabling infrastructure (Class D or better).

For more information on UPOE, see the “Configuring Universal PoE” section in [Chapter 15, “Configuring Power over Ethernet.”](#)

Web Content Coordination Protocol



Note

WCCP version 1 is *not* supported.

Web Content Communication Protocol (WCCP) Version 2 Layer 2 redirection enables Catalyst 4500 series switches to transparently redirect content requests to the directly connected content engines by using a Layer 2 and MAC address rewrite. The WCCPv2 Layer 2 redirection is accelerated in the switching hardware, and is more efficient than Layer 3 redirection using Generic Routing Encapsulation (GRE). The content engines in a cache cluster transparently store frequently accessed content, and then fulfill successive requests for the same content, eliminating repetitive transmissions of identical content from the original content servers. It supports the transparent redirection of HTTP and non-HTTP traffic with ports or dynamic services, such as Web caching, HTTPS caching, File Transfer Protocol (FTP) caching, proxy caching, media caching, and streaming services. WCCPv2 Layer 2 redirection is typically deployed for transparent caching at network edge, such as regional or branch sites. WCCPv2 Layer 2 redirection cannot be enabled on the same input interface with PBR or VRF-lite. ACL-based classification for Layer 2 redirection is not supported.

For information on WCCP, see [Chapter 81, “Configuring WCCP Version 2 Services.”](#)

Wireshark



Note

Wireshark is supported only on Supervisor Engine 7-E, Supervisor Engine 7L-E, Supervisor Engine 8-E, and Catalyst 4500X.

Starting with Cisco IOS Release XE 3.3.0SG and the IP Base and Enterprise Services feature sets, the Catalyst 4500 series switch supports Wireshark. This is a packet analyzer program, formerly known as Ethereal that supports multiple protocols and presents information in a graphical and text-based user interface. Wireshark is applied or enabled on an individual interface; global packet capture is not supported.

For information on Wireshark, see [Chapter 68, “Configuring Wireshark.”](#)

XML-PI

eXtensible Markup Language Programmatic Interface (XML-PI) Release 1.0 leverages the Network Configuration Protocol (NETCONF). It provides new data models that collect running configurations and **show** command output down to the keyword level without requiring the technologies or external XML-to-command line interface (CLI) gateways. XML-PI allows you to develop XML-based network management applications to control any number of network devices simultaneously.

Refer to the following link for more details:

http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_xmlpi_v1.html

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 security features:

- [802.1X Identity-Based Network Security, page 1-35](#)
- [Dynamic ARP Inspection, page 1-37](#)
- [Cisco TrustSec Security Architecture, page 1-36](#)
- [Cisco TrustSec Security Groups, SGTs and SGACLs, page 1-37](#)
- [Dynamic ARP Inspection, page 1-37](#)
- [Dynamic Host Configuration Protocol Snooping, page 1-37](#)
- [Flood Blocking, page 1-38](#)
- [Hardware-Based Control Plane Policing, page 1-38](#)
- [IP Source Guard, page 1-38](#)
- [IP Source Guard for Static Hosts, page 1-39](#)
- [IPsec VPN, page 1-40](#)
- [IPv6 First Hop Security, page 1-39](#)
- [Local Authentication, RADIUS, and TACACS+ Authentication, page 1-40](#)
- [Network Admission Control, page 1-41](#)
- [Network Security with ACLs, page 1-41](#)
- [Port Security, page 1-42](#)
- [PPPoE Intermediate Agent, page 1-42](#)
- [Session Aware Networking, page 1-42](#)
- [Storm Control, page 1-43](#)

- [uRPF Strict Mode, page 1-43](#)
- [Utilities, page 1-43](#)
- [Web-based Authentication, page 1-44](#)

802.1X Identity-Based Network Security

This security feature consists of the following:

- 802.1X Authentication for Guest VLANs—Allows you to use VLAN assignment to limit network access for certain users.
- 802.1X Authentication Failed Open Assignment—Allows you to configure a switch to handle the case when a device fails to authenticate itself correctly through 802.1X (for example, not providing the correct password).
- 802.1X Authentication with ACL Assignment—Downloads per-host policies such as ACLs and redirect URLs to the switch from the RADIUS server during 802.1X or MAB authentication of the host.
- 802.1X Authentication with Per-User ACL and Filter-ID ACL—Allows ACL policy enforcement using a third-party AAA server.
- 802.1X Convergence—Provides consistency between the switching business units in 802.1X configuration and implementation.
- 802.1X Protocol—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 RADIUS accounting—Allows you to track the use of network devices.
- 802.1X Supplicant and Authenticator Switches with Network Edge Access Topology (NEAT)—Extends identity to areas outside the wiring closet (such as conference rooms). NEAT is designed for deployment scenarios where a switch acting as 802.1X authenticator to end-hosts (PC or Cisco IP-phones) is placed in an unsecured location (outside wiring closet); the authenticator switch cannot always be trusted.
- 802.1X with Authentication Failed VLAN Assignment—Allows you to provide access for authentication failed users on a per-port basis. Authentication failed users are end hosts that are 802.1X-capable but do not have valid credentials in an authentication server or end hosts that do not give any username and password combination in the authentication pop-up window on the user side.
- 802.1X with Inaccessible Authentication Bypass—Applies when the AAA servers are unreachable or nonresponsive. In this situation, 802.1X user authentication typically fails with the port closed, and the user is denied access. Inaccessible Authentication Bypass provides a configurable alternative on the Catalyst 4500 series switch to grant a critical port network access in a locally specified VLAN.
- 802.1X with Port Security—Allows port security on an 802.1X port in either single- or multiple-host mode. When you enable port security and 802.1X on a port, 802.1X authenticates the port, and port security manages the number of MAC addresses allowed on that port, including that of the client.
- 802.1X with MAC Authentication Bypass—Provides network access to agentless devices without 802.1X supplicant capabilities, such as printers. Upon detecting a new MAC address on a switch port, the Catalyst 4500 series switch will proxy an 802.1X authentication request based on the device's MAC address.
- 802.1X with RADIUS-Provided Session Timeouts—Allows you to specify whether a switch uses a locally configured or a RADIUS-provided reauthentication timeout.

- 802.1X with Unidirectional Controlled Port—Allows the Wake-on-LAN (WoL) magic packets to reach a workstation attached to an unauthorized 802.1X switch port. Unidirectional Controlled Port is typically used to send operating systems or software updates from a central server to workstations at night.
- 802.1X with Violation Mode—This feature allows you to configure 802.1X security violation behavior as either shutdown, restrict, or replace mode, based on the response to the violation.
- 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 with VLAN user distribution—An alternative to dynamically assigning a VLAN ID or a VLAN name, this feature assign a VLAN Group name. It enables you to distribute users belonging to the same group (and characterized by a common VLAN Group name) across multiple VLANs. Ordinarily, you do this to avoid creating an overly large broadcast domain.
- 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.
- Multi-Domain Authentication—This feature allows both a data device and a voice device, such as an IP phone (Cisco or non-Cisco), to authenticate on the same switch port, which is divided into a data domain and a voice domain.
- RADIUS Change of Authorization—This feature employs Change of Authorization (CoA) extensions defined in RFC 5176 in a push model to allow for the dynamic reconfiguring of sessions from external authentication, authorization, and accounting (AAA) or policy servers.

For more information on 802.1X identity-based network security, see [Chapter 49, “Configuring 802.1X Port-Based Authentication.”](#)

Cisco TrustSec MACsec Encryption

MACsec (Media Access Control Security) is the IEEE 802.1AE standard for authenticating and encrypting packets between two MACsec-capable devices. The Catalyst 4500 series switch supports 802.1AE encryption with MACsec Key Agreement (MKA) on downlink ports for encryption between the switch and host devices. The switch also supports MACsec link layer switch-to-switch security by using Cisco TrustSec Network Device Admission Control (NDAC) and the Security Association Protocol (SAP) key exchange. Link layer security can include both packet authentication between switches and MACsec encryption between switches (encryption is optional).

For more information on TrustSec MACsec encryption, see [Chapter 48, “Configuring MACsec Encryption.”](#)

Cisco TrustSec Security Architecture

The Cisco TrustSec security architecture builds secure networks by establishing domains of trusted network devices. Each device in the domain is authenticated by its peers. Communication on the links between devices in the domain is secured with a combination of encryption, message integrity check, and data-path replay protection mechanisms. Cisco TrustSec uses the device and user credentials acquired during authentication for classifying the packets by security groups (SGs) as they enter the network. This packet classification is maintained by tagging packets on ingress to the Cisco TrustSec network so that they can be properly identified for the purpose of applying security and other policy criteria along the data path. The tag, called the security group tag (SGT), allows the network to enforce the access control policy by enabling the endpoint device to act upon the SGT to filter traffic.

For more information, refer to the following URL:

<http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/trustsec.html>

Cisco TrustSec Security Groups, SGTs and SGACLs



Note

This support is provided only on Supervisor Engine 7-E, Supervisor Engine 7L-E, Supervisor Engine 8-E, and Catalyst 4500X.

A security group is a grouping of users, endpoint devices, and resources that share access control policies. Security groups are defined by the administrator in the Cisco ISE or Cisco Secure ACS. As new users and devices are added to the Cisco TrustSec domain, the authentication server assigns these new entities to appropriate security groups. Once a device is authenticated, Cisco TrustSec tags any packet that originates from that device with a security group tag (SGT) that contains the security group number of the device. The packet carries this SGT throughout the network.

Using security group access control lists (SGACLs), you can control the operations that users can perform based on the security group assignments of users and destination resources.

For more information, refer to the following URL:

<http://www.cisco.com/en/US/docs/switches/lan/trustsec/configuration/guide/trustsec.html>

For Cisco TrustSec SGFT and SGACL guidelines and restrictions that apply on the Catalyst 4500 series switch, refer to "Appendix B. Notes for the Catalyst 4500 Series Switches" in the *Cisco TrustSec Switch Configuration Guide*.

Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) 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.

For more information on dynamic ARP inspection, see [Chapter 58, "Configuring Dynamic ARP Inspection."](#)

Dynamic Host Configuration Protocol 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.

With SSO support, DHCP Snooping 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.

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

For information on configuring DHCP snooping, see [Chapter 60, “Configuring DHCP Snooping, IP Source Guard, and IPSG for Static Hosts.”](#)

Flood Blocking

Flood blocking 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.

For information on flood blocking, see [Chapter 64, “Port Unicast and Multicast Flood Blocking.”](#)

Hardware-Based Control Plane Policing

Control Plane Policing provides a unified solution to limit the rate of CPU bound control plane traffic in hardware. It enables users to install system wide control plane ACLs to protect the CPU by limiting rates or filtering out malicious DoS attacks. Control plane policing ensures the network stability, availability and packet forwarding, and prevents network outages such as loss of protocol updates despite an attack or heavy load on the switch. Hardware-based control plane policing is available for all Catalyst 4500 supervisor engines. It supports various Layer 2 and Layer 3 control protocols, such as CDP, EAPOL, STP, DTP, VTP, ICMP, CGMP, IGMP, DHCP, RIPv2, OSPF, PIM, TELNET, SNMP, HTTP, and packets destined to 224.0.0.* multicast link local addresses. Predefined system policies or user-configurable policies can be applied to those control protocols.

Through Layer 2 Control Packet QoS, you can police control packets arriving on a physical port or VLAN; it enables you to apply QoS on Layer 2 control packets

For information on control plane policing and Layer 2 control packet QoS, see [Chapter 57, “Configuring Control Plane Policing and Layer 2 Control Packet QoS.”](#)

IP Source Guard

Similar to DHCP snooping, this feature is enabled on an untrusted Layer 2 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 60, “Configuring DHCP Snooping, IP Source Guard, and IPSG for Static Hosts.”](#)

IP Source Guard for Static Hosts

This feature allows you to secure the IP address learned from static hosts by using ARP packets and then bind that IP address to a given MAC address using the device tracking database, allowing entries to survive through link down events.

IP Source Guard (IPSG) for static hosts allows multiple bindings per-port per-MAC address for both DHCP and static hosts, in both device tracking database and DHCP snooping binding database. The feature allows you to take action when a limit is exceeded.

For information on configuring IPSG for static hosts, see [Chapter 60, “Configuring DHCP Snooping, IP Source Guard, and IPSG for Static Hosts.”](#)

IPv6 First Hop Security

**Note**

IPv6 First Hop Security is supported only on Catalyst 4948E, Catalyst 4948E-F, Catalyst 4500-X, Supervisor Engine 6-E, 6L-E, 7-E, 7L-E, and 8-E.

IPv6 FHS is a suite of features designed to secure link operations in an IPv6 enabled network as well as address certain scalability issues seen in large L2 domains. IPv6 FHS provides effective counter measures for the following types of attacks or misconfiguration errors that could result in DoS or information theft:

- Router impersonation (MiM attacks)
- Address theft
- Address spoofing
- Remote address resolution cache exhaustion (DoS attacks)

These attacks can come from malicious or mis-configured users and could result in severe disruption to users of the Layer 2 domain and to the network in general.

The following features are supported:

- DAD Proxy
- Data Glean
- Destination Guard
- IPv6 Snooping (DHCP Data Gleaning, per-limit Address Limit)
- IPv6 Address Glean
- IPv6 Device Tracking
- Lightweight DHCPv6 Relay Agent (LDRA)
- NDP Inspection
- Per ND Cache Limit
- Per Port Address Limit
- Source and Prefix Guard

**Note**

IPv6 LDRA is the only FHS feature supported on EtherChannels.

**Note**

Configuring IPv6 FHS on secondary VLANs is not allowed; they inherit the policy from the primary VLAN configuration. Whatever policy is applied on the primary VLANs is programmed automatically on the associated secondary VLANs. The applied policy, however, always overrides the VLAN level configuration.

The following caveats are specific for Data Glean, Prefix Guard, and Source Guard enabled on a Catalyst 4500 series switch:

- First Hop Security (FHS) cannot be configured on the same port or VLAN as dot1X, because the latter asserts control over the MAC table and FHS requires similar control to allow only valid NDP or DHCPv6 hosts.
- If unicast Rpf (unicast reverse path forwarding; uRPF) is configured on box and FHS is enabled, Forward Lookup CAM is populated with routes from FHS and uRPF. Packets that normally fail the uRPF check are admitted provided it passes the Source Guard or Prefix Guard check.
- If Data Glean policy and Source Guard (or Prefix Guard) are applied such that VLAN policies and port policies differ, neither VLAN nor port policy are effective.
- All ICMP and DHCP version 6 control packets are permitted even when Source Guard or Prefix Guard is enabled.

For a brief overview of FHS, see the URL:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6553/aag_c45-707354.pdf

For detailed information on how to implement FHS, see the URL:

<http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/12-4t/ip6-first-hop-security.html>

IPsec VPN

When a growing organization expands to multiple locations, one of the challenges it faces is how to interconnect remote sites to the corporate network. As network security risks increase and regulatory compliance becomes essential, it is important to address these critical needs.

You can dramatically increase the reach of your network without significantly expanding your infrastructure by using Cisco IOS IPsec VPNs. IPsec is a standards-based encryption technology that enables you to securely connect branch offices and remote users and provides significant cost savings compared to traditional WAN access such as Frame Relay or ATM. IPsec VPNs provide high levels of security through encryption and authentication, protecting data from unauthorized access.

For additional information, refer to the following URL:

http://www.cisco.com/en/US/products/ps6635/products_ios_protocol_group_home.html

Local Authentication, RADIUS, and TACACS+ Authentication

Local Authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+) 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_ps6350_TSD_Products_Configuration_Guide_Chapter.html

Network Admission Control

Network Admission Control consists of two features:

- NAC Layer 2 IP validation

NAC Layer 2 IP is an integral part of Cisco Network Admission Control. It offers the first line of defense for infected hosts (PCs and other devices attached to a LAN port) attempting to connect to the corporate network. NAC Layer 2 IP on the Catalyst 4500 series switch performs posture validation at the Layer 2 edge of the network for non-802.1x-enabled host devices. Host device posture validation includes antivirus state and OS patch levels. Depending on the corporate access policy and host device posture, a host may be unconditionally admitted, admitted with restricted access, or quarantined to prevent the spread of viruses across the network.

For more information on Layer 2 IP validation, see the URL:

http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.1/configuration/guide/nac_conf.html

- NAC Layer 2 802.1X authentication

The Catalyst 4500 series switch extends NAC support to 802.1x-enabled devices. Like NAC Layer 2 IP, the NAC Layer 2 802.1x feature determines the level of network access based on endpoint information.

For more information on 802.1X identity-based network security, see [Chapter 49, “Configuring 802.1X Port-Based Authentication.”](#)

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 Catalyst 4500 series switch supports three types of ACLs:

- IP ACLs, which filter IP traffic, including TCP, the User Datagram Protocol (UDP), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol (ICMP)
- IPv6 ACLs
- MAC ACLs which match based on Ethernet addresses and Ether Type

The switch supports the following applications of ACLs to filter traffic:

- 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.
- Router ACLs, which are applied to Layer 3 interfaces to control the access of routed traffic between VLANs.
- VLAN ACLs or VLAN maps to control the access of all packets (bridged and routed).

For information on ACLs, MACLs, VLAN maps, MAC address filtering, and Port ACLs, see [Chapter 62, “Configuring Network Security with ACLs.”](#)

Port Security

Port security restricts traffic on a port based upon the MAC address of the workstation that accesses the port. Trunk port security extends this feature to trunks, including private VLAN isolated trunks, on a per-VLAN basis.

Sticky port security extends port security by saving the dynamically learned MAC addresses in the running configuration to survive port link down and switch reset. It enables a network administrator to restrict the MAC addresses allowed or the maximum number of MAC addresses on each port.

Voice VLAN sticky port security further extends the sticky port security to the voice-over-IP deployment. Voice VLAN sticky port security locks a port and blocks access from a station with a MAC address different from the IP phone and the workstation behind the IP phone.

For information on port security, see [Chapter 55, “Configuring Port Security.”](#)

PPPoE Intermediate Agent

PPPoE Intermediate Agent (PPPoE IA) is placed between a subscriber and BRAS to help the service provider BRAS distinguish between end hosts connected over Ethernet to an access switch. On the access switch, PPPoE IA enables Subscriber Line Identification by appropriately tagging Ethernet frames of different users. (The tag contains specific information such as which subscriber is connected to the switch and VLAN.) PPPoE IA acts as mini-security firewall between host and BRAS by intercepting all PPPoE Active Discovery (PAD) messages on a per-port per-VLAN basis. It provides specific security feature such as verifying the intercepted PAD message from untrusted port, performing per-port PAD message rate limiting, inserting and removing VSA tags into and from PAD messages, respectively.

For information on PPPoE IA, see [Chapter 51, “Configuring the PPPoE Intermediate Agent.”](#)

Session Aware Networking

Session Aware Networking provides an identity-based approach to access management and subscriber management. It offers a consistent way to configure features across technologies, a command interface that allows easy deployment and customization of features, and a robust policy control engine with the ability to apply policies defined locally or received from an external server to enforce policy in the network.

Session Aware Networking allows a single session identifier to be used for web authentication sessions in addition to all 802.1X and MAB authenticated sessions for a client. This session ID is used for all reporting purposes such as show commands, MIBs, and RADIUS messages and allows users to distinguish messages for one session from messages for other sessions. This common session ID is used consistently across all authentication methods and features applied to a session.

**Note**

IPv6 is not supported for web authentication, dot.1X, or MAB.

For additional information, refer to the following URL:

<http://www.cisco.com/en/US/docs/ios-xml/ios/san/configuration/xe-3se/3850/san-overview.html>

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.

Starting with Cisco IOS Release 12.2(40)SG, the Catalyst 4500 series switch allows suppression of broadcast and multicast traffic on a per-port basis.

For information on configuring broadcast suppression, see [Chapter 65, “Configuring Storm Control.”](#)

uRPF Strict Mode

The uRPF feature mitigates problems caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address. uRPF deflects denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by forwarding only packets that have source addresses that are valid and consistent with the IP routing table. This helps to protect the network of the customer, the ISP, and the rest of the Internet. When using uRPF in strict mode, the packet must be received on the interface that the router uses to forward the return packet. uRPF strict mode is supported for both IPv4 and IPv6 prefixes.

For information on configuring broadcast suppression, see [Chapter 37, “Configuring Unicast Reverse Path Forwarding.”](#)

Utilities

Supported utilities include the following:

Layer 2 Traceroute

Layer 2 traceroute 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 10, “Checking Port Status and Connectivity.”](#)

Time Domain Reflectometry

Time Domain Reflectometry (TDR) is a technology used for diagnosing the state and reliability of cables. TDR can detect open, shorted, or terminated cable states. The calculation of the distance to the failure point is also supported.

For information about TDR, see [Chapter 10, “Checking Port Status and Connectivity.”](#)

Debugging Features

The switch has several commands to help you debug your initial setup. These commands are included in the following command groups:

- **platform**
- **debug platform**

For more information, refer to the command reference guide.

Web-based Authentication

The web-based authentication feature, known as Web Authentication Proxy, enables you to authenticate end users on host systems that do not run the IEEE 802.1X supplicant. When you initiate an HTTP session, this feature intercepts ingress HTTP packets from the host and sends an HTML login page to your. You key in the credentials, which the web-based authentication feature sends to the AAA server for authentication. If authentication succeeds, web-based authentication sends a Login-Successful HTML page to the host and applies the access policies returned by the AAA server.

For information on configuring web-based authentication, see [Chapter 52, “Configuring Web-Based Authentication.”](#)

