



Release Notes for the Catalyst 4900M Series Switch, Cisco IOS Release 12.2(53)SG

Current Release

12.2(53)SG3—August 30, 2010

Previous Release

12.2(53)SG2, 12.2(53)SG, 12.2(52)SG, 12.2(50)SG6, 12.2(50)SG5, 12.2(50)SG4, 12.2(50)SG3, 12.2(50)SG2, 12.2(50)SG1, 12.2(50)SG, 12.2(46)SG, 12.2(40)XO

These release notes describe the features, modifications, and caveats for Cisco IOS software on the Catalyst 4900M switch.

Cisco Systems announces the Cisco Catalyst 4900M Series, a premium extension to the widely deployed Catalyst 4948 Series top of rack Ethernet switches for data center server racks. Optimized for ultimate deployment flexibility, the Catalyst 4900M Series can be deployed for 10/100/1000 server access with 1:1 uplink to downlink oversubscription, mix of 10/100/1000 and 10 GbE servers or all 10GbE servers in the same rack. The Catalyst 4900M is a 320Gbps, 250Mpps, 2RU fixed configuration switch with 8 fixed wire speed X2 ports on the base unit and 2 optional half card slots for deployment flexibility and investment protection. Low latency, scalable buffer memory and high availability with 1+1 hot swappable AC or DC power supplies and field replaceable fans optimize the Catalyst 4900M for any size of data center.

Support for Cisco IOS Software Release 12.2(53)SG, the default image, follows the standard Cisco Systems® support policy, available at http://www.cisco.com/en/US/products/products_end-of-life_policy.html

For more information about the Cisco Catalyst 4900M Series, visit: <http://www.cisco.com/go/cat4900/docs>.



Note

Although their *Release Notes* are unique, the 4 platforms (Catalyst 4500, Catalyst 4900, Catalyst ME 4900, and Catalyst 4900M) use the same *Software Configuration Guide*, *Command Reference Guide*, and *System Message Guide*. Refer to this location:

<http://www.cisco.com/go/cat4500/docs>



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2008-2010 Cisco Systems, Inc. All rights reserved.

Contents

This publication consists of these sections:

- [Cisco IOS Software Packaging for the Cisco Catalyst 4900M Switch, page 2](#)
- [System Requirements, page 4](#)
- [New and Changed Information, page 15](#)
- [Minimum and Recommended ROMMON Release, page 21](#)
- [Limitations and Restrictions, page 22](#)
- [Caveats, page 24](#)
- [Troubleshooting, page 143](#)
- [Related Documentation, page 144](#)
- [Notices, page 146](#)
- [Obtaining Documentation and Submitting a Service Request, page 148](#)

Cisco IOS Software Packaging for the Cisco Catalyst 4900M Switch

Catalyst 4900M software features based on Cisco IOS Software 12.2(53)SG will support the IP Base image and the entservices image.

The IP Base image does not support enhanced routing features such as Nonstop Forwarding/Stateful Switchover (NSF/SSO), BGP, Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Internetwork Packet Exchange (IPX), AppleTalk, Virtual Routing Forwarding (VRF-lite), GLBP, and policy-based routing (PBR). The IP Base image supports Static routes, RIPv1/v2 for IP BASE, and EIGRP-Stub for limited routing on Cisco Catalyst 4900 Series Switches.

The Enterprise Services image supports Cisco Catalyst 4900M Series software features based on Cisco IOS Software 12.2(53)SG, including enhanced routing. BGP capability is included in the Enterprises Services package.

**Note**

The recommended Cisco IOS image on the Catalyst 4900M is 12.2(50)SG3.

Orderable Product Numbers:

- S49MES-12253SG - Cisco IOS Software for Cisco Catalyst 4900M Switches (Enterprise Services image with BGP support)
- S49MESK9-12253SG - Cisco IOS Software for Cisco Catalyst 4900M Switches (Enterprise Services image with 3DES and BGP support)
- S49MIPB-12253SG - Cisco IOS Software for Cisco Catalyst 4900M Switches (IP Base image)
- S49MIPBK9-12253SG - Cisco IOS Software for Cisco Catalyst 4900M Switches (IP Base image with 3DES)
- S45EIPB-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (IP Base Image)

- S45IPBK9-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (IP Base Image with 3DES) (cat4500-ipbasek9-mz)
- S45EES-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (Enterprise Services image) (cat4500-ipbasek9-mz)
- S45EESK9-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (Enterprise Services image) (cat4500-ipbasek9-mz)
- S45EIPB-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (IP Base Image)
- S45IPBK9-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (IP Base Image with 3DES) (cat4500-ipbasek9-mz)
- S45EES-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (Enterprise Services image) (cat4500-ipbasek9-mz)
- S45EESK9-12250SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (Enterprise Services image) (cat4500-ipbasek9-mz)
- S45EIPB-12246SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (IP Base Image)
- S45IPBK9-12246SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (IP Base Image with 3DES) (cat4500-ipbasek9-mz)
- S45EES-12246SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (Enterprise Services image) (cat4500-ipbasek9-mz)
- S45EESK9-12246SG—Cisco IOS software for the Catalyst 4500 Series Supervisor Engine 6-E (Enterprise Services image) (cat4500-ipbasek9-mz)
- S49IPB-12252SG—Cisco IOS software for the Catalyst 4900 Series (IP Base image) (cat4500-ipbase-mz)
- S49IPBK9-12252SG—Cisco IOS software for the Catalyst 4900 Series (IP Base image with Triple Data Encryption Standard (3DES)) (cat4500-ipbasek9-mz)
- S49ES-12252SG—Cisco IOS software for the Catalyst 4900 Series (Enterprise Services image with BGP support) (cat4500-entservices-mz)
- S49ESK9-12252SG—Cisco IOS software for the Catalyst 4900 Series (Enterprise Services image with 3DES and BGP) (cat4500-entservicesk9-mz)

Cisco 4900M Series Ethernet Switch Cisco IOS Release Strategy

Customers with Catalyst 4900M switches who need the latest hardware support and software features should migrate to Cisco IOS Release 12.2(53)SG.


For more information on the Cisco 4900M Switch, visit the following URL:
www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/index.htm

Cisco IOS Software Migration

Figure 1 displays the two active, 12.2(31)SGA and 12.2(50)SG, and newly introduced 12.2(53)SG extended maintenance trains.

Support for the Catalyst 4900M platform was introduced in 12.2(40)XO. Moving forward, the Cisco Catalyst 4900M platform has two maintenance trains. The Cisco IOS Release 12.2(53)SG is the latest maintenance train and includes the most recent features including support for OSPF for routed Access

Figure 1 Software Release Strategy for the Catalyst 4900M Series Switch



Summary of Migration Plan

- Customers requiring the latest Cisco Catalyst 4900M Switch hardware and software features should migrate to Cisco IOS Software Release 12.2(53)SG.

Support

Support for Cisco IOS Software Release 12.2(53)SG follows the standard Cisco Systems® support policy, available at http://www.cisco.com/en/US/products/products_end-of-life_policy.html

System Requirements



Note

The recommended Cisco IOS image on the Catalyst 4900M is 12.2(50)SG3.

This section describes the system requirements:

- [Supported Hardware, page 4](#)
- [Supported Features, page 6](#)
- [Unsupported Features, page 14](#)

Supported Hardware

The following tables lists the hardware supported on the Catalyst 4900M series switch.

Table 1 Supported Hardware

Product Number (append with “=” for spares)	Product Description
Small Form-Factor Pluggable Modules (supported only in WS-X4908-10GE(=) half-card)	
GLC-SC-MM	Gigabit Ethernet SFP, LC connector, and SX transceiver small form-factor pluggable module
GLC-LH-SM	Gigabit Ethernet SFP, LC connector, and LX/LH transceiver small form-factor pluggable module
GLC-ZX-SM	1000BASE-ZX small form-factor pluggable module
GLC-T	1000BASE-T small form-factor pluggable module
CWDM-SFP-xxxx	CWDM small form-factor pluggable module (See Table 2 on page 6 for a list of supported wavelengths.)
10 Gigabit Ethernet X2 Pluggable Modules	
X2-10GB-LR	10GBASE-LR X2 transceiver module for SMF, 1310-nm wavelength, SC duplex connector
X2-10GB-ER	10GBASE-ER X2 transceiver module for SMF, 1550-nm wavelength, SC duplex connector
X2-10GB-CX4	10GBASE-CX4 X2 transceiver module for CX4 cable, copper, Infiniband 4X connector
X2-10GB-LX4	10GBASE-LX4 X2 transceiver module for MMF, 1310-nm wavelength, SC duplex connector
X2-10GB-LRM	10GBASE-LRM X2 transceiver module for MMF, 1310-nm wavelength, SC duplex connector
X2-10GB-SR	10GBASE-SR X2 transceiver module for MMF, 850-nm wavelength, SC duplex connector
X2-10GB-ZR	10GBASE-ZR X2 transceiver module for SMF, 1550 nm wavelength up to 80 km. DOM is not supported.
X2-10GB-DWDM	10GBASE-ZR X2 transceiver module for SMF, 32 nontunable ITU 100-GHz wavelengths up to 80 km are supported. DOM is supported. Dual SC/PC connectors are supported.
CVR-X2-SFP10G	Hot-swappable input/output (I/O) converter module that fits into a 10-Gigabit Ethernet X2 slot on a switch or line card module. Hosts one 10-Gigabit Ethernet SFP+ transceiver module.
SFP+ Modules	
SFP-10G-SR	Cisco 10GBASE-SR SFP+ Module for MMF

Table 2 briefly describes the supported wavelengths in the Catalyst 4900M series switches.

Table 2 CWDM SFP Supported Wavelengths

Product Number (append with “=” for spares)	Product Description
CWDM-SFP -1470	Longwave 1470 nm laser single-mode
CWDM- SFP -1490	Longwave 1490 nm laser single-mode
CWDM-SFP -1510	Longwave 1510 nm laser single-mode
CWDM-SFP -1530	Longwave 1530 nm laser single-mode
CWDM-SFP -1550	Longwave 1550 nm laser single-mode
CWDM-SFP -1570	Longwave 1570 nm laser single-mode
CWDM-SFP -1590	Longwave 1590 nm laser single-mode
CWDM-SFP -1610	Longwave 1610 nm laser single-mode

The following tables lists the hardware supported on the Catalyst 4900M series switch.

Table 3 Supported Hardware

Product Number (append with “=” for spares)	Product Description
WS-C4900M	Catalyst 4900M 8-port base system
WS-X4920-GB-RJ45 (=)	Catalyst 4900M 20-port 10/100/1000 RJ-45 half card
WS-X4904-10GE (=)	Catalyst 4900M 4 port 10GbE half card with X2 interfaces
WS-X4908-10GE (=)	Catalyst 4900M 8 port 10GbE half card with X2 interfaces
PWR-C49M-1000AC(=)	Catalyst 4900M AC Power Supply
PWR-C49M-1000AC/2	Catalyst 4900M AC Power Supply Redundant
PWR-C49M-1000DC(=)	Catalyst 4900M DC Power Supply
PWR-C49M-1000DC/2	Catalyst 4900M DC Power Supply Redundant
WS-X4992=	Catalyst 4900M Spare Fan Tray
CVR-X2-SFP=	TwinGig module

Supported Features



Note

The default image for the Catalyst 4900M series switch is Cisco IOS Release 12.2(50)SG5.

Table 4 lists the Cisco IOS software features for the Catalyst 4900M series switch.

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch

Layer 2 Switching Features

Storm control

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch (continued)

Storm Control: Per-Port Multicast Suppression
Multicast storm control
IP Source Guard
IP Source Guard for Static Hosts
PVRST+
Layer 2 transparent bridging ¹
Layer 2 MAC ² learning, aging, and switching by software
Unicast MAC address filtering
VMPS ³ Client
Layer 2 hardware forwarding up to 102 Mpps
Layer 2 Control Policing (Not supported on Supervisor Engine 6-E)
Layer 2 switch ports and VLAN trunks
Spanning-Tree Protocol (IEEE 802.1D) per VLAN
802.1s and 802.1w
Layer 2 traceroute
Unidirectional Ethernet port
Per-VLAN spanning tree (PVST) and PVST+
Spanning-tree root guard
Spanning-tree Loop guard and PortFast BPDU Filtering
Support for 9216 byte frames
Port security
Port security on Voice VLAN
Port security MAC Aging
Trunk Port Security
Unicast MAC Filtering
802.1X with Port Security
Private VLANs
Private VLAN DHCP snooping
Private VLAN trunks
IEEE 802.1Q-based VLAN encapsulation
Multiple VLAN access port
VLAN Trunking Protocol (VTP) and VTP domains
VTP v3
Support for 4096 VLANs per switch
Unidirectional link detection (UDLD) and aggressive UDLD
SNMP V3 support for Bridge-MIB with VLAN indexing
Ethernet CFM

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch (continued)

Ethernet OAM Protocol
Supported Protocols
DTP ⁴
RIPv1 ⁵ and RIPv2, Static Routing
EIGRP ⁶
EIGRP Stub Routing
OSPF ⁷
BGP4 ⁸
BGP route-map Continue
BGP Neighbor Policy
MBGP ⁹
MSDP ¹⁰
ICMP ¹¹ Router Discovery Protocol
Static routes
Classless interdomain routing (CIDR)
DVMRP ¹²
NTP ¹³
STP - Portfast BPDU Guard
STP- BPDU Filtering
STP - Root Guard
SCP ¹⁴
EtherChannel Features
Cisco EtherChannel technology - 10/100/1000 Mbps, 10 Gbps
Load balancing for routed traffic, based on source and destination IP addresses
Load sharing for bridged traffic based on MAC addresses
IEEE 802.1Q on all EtherChannels
Bundling of up to eight Ethernet ports
Trunk Port Security over EtherChannel
Additional Protocols and Features
Secure Copy Protocol (SCP)
Routed Jumbo Frame support
SPAN CPU port mirroring
SPAN packet-type filtering
SPAN destination in-packets option
SPAN ACL filtering
Enhanced VLAN statistics
Secondary addressing
Bootstrap protocol (BOOTP)

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch (continued)

Authentication, authorization, and accounting using TACACS+ and RADIUS protocol
Cisco Discovery Protocol (CDP)
CDP 2nd Port Status TLV
FlexLink and MAC Address-Table Move Update
Network Mobility Services Protocol
Sticky port security
Voice VLAN Sticky Port Security
Cisco Group Management Protocol (CGMP) server support
HSRP ¹⁵ over Ethernet, EtherChannels - 10/100/1000Mbps, 10 Gbps
GLBP
VRRP
IGMP ¹⁶ snooping version 1, version 2, and version 3 (Full Support)
IGMP filtering
IGMP Querier
Multicast VRF-lite
VRF-aware IP services
Configurable IGMP Leave Timer
Multicast Source Discovery Protocol (MSDP)
Smartports I custom macros
Smartports II default macros
Smartports III global macros
Port Aggregation Protocol (PagP)
802.3ad LACP
SSH version 1 and version 2 ¹⁷
show interface capabilities command
IfIndex persistence
Enhanced SNMP MIB support
SNMP ¹⁸ version 1, version 2, and version 3
SNMP version 3 (with encryption)
DHCP server and relay-agent
DHCP Snooping Statistics and SYSLOG
DHCP client autoconfiguration
DHCP Option 82 data Insertion
DHCP Option 82 Pass Through
DHCP Relay Agent for IPv6
DHCP Option 82 - Configurable Remote ID and Circuit ID
Port flood blocking

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch (continued)

Router standard and extended ACLs ¹⁹ on all ports with no performance penalty
Downloadable ACL
VLAN ACL
PAACL ²⁰
VACL
RACL
Unicast RPF
Local Proxy ARP
Dynamic ARP Inspection on PVLANS
Dynamic ARP Inspection
Per-VLAN CTI
ARP QoS
MQC
Ingress/Egress Policing
Ingress Rate Limiting
Egress Bandwidth Limiting/port shaping
Per VLAN Policy & Per Port Policer
802.1p Priority
Strict Priority Scheduling
Ingress/Egress Strict Priority Queuing (Expedite)
Shaped Round Robin (SRR)
Egress Shaped Queues
Ingress/egress Shared Queues
DSCP Mapping
DSCP Filtering
AutoQoS - VoIP
PBR ²¹
Auto QoS 1.5
Trust Boundary Configuration
Dynamic Buffer Limiting (DBL)
Per-VLAN Control Traffic Intercept
Table Map Based Classification
Interface Index Persistence
UDI - Unique Device Identifier
Per-port QoS ²² rate-limiting and shaping
Per-port Per-VLAN QoS
Energy Wise

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch (continued)

Two-Rate Three-Color Policing
Dynamic Multi-Protocol Ternary Content Addressable Memory
SmartPort macros
802.1s standards compliance
Flexible Authentication Sequencing
Multi-Authentication
Open Authentication
Web Authentication
Local Web Authentication (EPM syslog and Common session ID)
PPPoE Intermediate Agent
Identity ACL Policy Enforcement ²³
IPv6 routing - unicast routing "RIPng"
IPv6 Neighbor Discovery Throttlingly
IPv6 MLDv1 & v2 Snooping
IPv6 Host support (- IPv6 support: Addressing; IPv6: Option processing, Fragmentation, ICMPv6, TCP/UDP over IPv6; Applications: Ping/Traceroute/VTY/SSH/TFTP, SNMP for IPv6 objects)
IPv6 ACLs
IPv6 Management Services (CDP over IPv6, SSHv2 over IPv6)
IPv6: MLDv1/v2
IPv6:CEFv6
IPv6:MLD Snooping
Non-stop Forwarding Awareness
Non-stop Forwarding Awareness for EIGRP-stub in IP base for all supervisor engines
BGP MIB
OSPF Fast Convergence ²⁴
AutoRP
Service-Aware Resource Allocation
TwinGig Converter Module
FAT File System
EEM ²⁵
VSS client with PagP+
Ethernet Management Port

Table 4 Cisco IOS Software Feature Set for the Catalyst 4900M series Switch (continued)

Enhanced Object Tracking subfeatures:

- HSRP with EOT
 - VRRP with EOT
 - GLBP with EOT
 - IP SLA with EOT
 - Reliable Backup Static Routing with EOT
-

ANCP Client

Bidirectional PIM

OSPF and EIGRP Fast Convergence

Inactivity Timer

boot config command

Crashdump enhancement

Unicast MAC filtering

Energy Wise

DHCPv6 Ethernet Remote ID option

DHCPv6 Relay - Persistent Interface ID option DHCPv6 Relay Agent notification for Prefix Delegation

PIM SSM Mapping

VRF lite NSF support with routing protocols OSPF/EIGRP/BG

Layer 2 Tunneling Protocol

Online Diagnostics

PIM Accept Register - Rogue Multicast Server Protection²⁶

Configuration Rollback

IP Multicast Load Splitting (Equal Cost Multipath (ECMP) using S, G and Next-hop)

OSPF for Routed Access

Archiving crashfiles

Cisco Network Assistant (CNA)

1. Hardware-based transparent bridging within a VLAN
2. MAC = Media Access Control
3. VMPS = VLAN Management Policy Server
4. DTP = Dynamic Trunking Protocol
5. RIP = Routing Information Protocol
6. EIGRP = Enhanced Interior Gateway Routing Protocol
7. OSPF = Open Shortest Path First
8. BGP4 = Border Gateway Protocol 4
9. MBGP = Multicast Border Gateway Protocol
10. MSDP = Multicast Source Discovery Protocol
11. ICMP = Internet Control Message Protocol
12. DVMRP = Distance Vector Multicast Routing Protocol
13. NTP = Network Time Protocol
14. SCP = Secure Copy Protocol

15. HSRP = Hot Standby Router Protocol
16. IGMP = Internet Group Management Protocol
17. SSH = Secure Shell Protocol
18. SNMP = Simple Network Management Protocol
19. ACLs = Access Control Lists
20. PACL = Port Access Control List
21. Policy-based Routing
22. QoS = Quality of Service
23. filter-ID and per-user ACL
24. The Catalyst 4500 series switch supports Fast Hellos, ISPF, and LSA Throttling.
25. EEM = Embedded Event Manager
26. The route-map keyword is not supported.

Following Features are Supported only on the Catalyst 4900M

With Cisco IOS Release 12.2(52)SG, the following features are available only with Supervisor Engine 6-E:

- IPv6
 - IPv6 Addressing Architecture
 - CDP IPv6 Address Family
 - CEFv6
 - DNS resolver for AAAA over an IPv4 transport
 - DNS resolver for AAAA over an IPv6 transport
 - Extended ACL
 - Hop-by-Hop option header
 - ICMP Rate Limiting
 - ICMPv6
 - ICMPv6 Redirect
 - IPv6 MIB
 - IPv6 over IEEE 802.1Q
 - IPv6 over IPv4 GRE tunnel
 - ISATAP
 - Loopback
 - MFIB for IPv6
 - MLD Snooping (will show up as a new chapter in the Config Guide)
 - MLDv1/v2
 - MTU Path Discovery for IPv6
 - OSPFv3
 - RIPng
 - EIGRPv6
 - BGPv4

- FAT filesystem
- PIM (SM, DM, SDM)
- QoS
 - Two Rate three Color Policing
 - Table map support for marking
 - Class based queuing actions (shaping/bandwidth/queue-limit/dbl/strict priority)
- Voltage Margining CLI
- QoS for IPv6
- ARP QoS

Unsupported Features

These features are not supported in Cisco IOS Release 12.2(53)SG for the Catalyst 4900M switch:

- The following ACL types:
 - Standard Xerox Network System (XNS) access list
 - Extended XNS access list
 - DECnet access list
 - Protocol type-code access list
- ADSL and Dial access for IPv6
- AppleTalk EIGRP (use native AppleTalk routing instead)
- Auto RP
- AutoQoS - VoIP
- Bridge groups
- CEF Account
- CER for E-911 Support
- CFM CoS
- Cisco IOS software IPX ACLs:
 - <1200-1299> IPX summary address access list
- Cisco IOS software-based transparent bridging (also called “fallback bridging”)
- Cisco-Port-QoS-MIB
- Connectionless (CLNS) routing; including IS-IS routing for CLNS. IS-IS is supported for IP routing only.
- DLSw (data-link switching)
- Global QoS (enable QoS)
- HTTP Software Upgrade
- IGRP (use EIGRP instead)
- IP SLA
- IS-IS

- IS-IS MIB
- **isis network point-to-point** command
- ISSU
- Kerberos support for access control
- Lock and key
- MAC Address Notification
- MAC notification MIB support
- NAC L2 IP - Inaccessible authentication bypass
- NAT-PT for IPv6
- NSF with SSO
- Packet Based Storm Control
- PIM Stub in IP Base
- Real Time DiagNosis (GOLD-Lite)
- Reflexive ACLs
- Routing IPv6 over an MPLS network
- RPR
- Time Domain Reflectometry
- Two-way community VLANs in private VLANs
- UniDirectional Link Routing (UDLR)
- WCCP v1 and v2

New and Changed Information

These sections describe the new and changed information for the Catalyst 4900M series switch running Cisco IOS software:

- [New Hardware Features in Release 12.2\(53\)SG, page 16](#)
- [New Software Features in Release 12.2\(53\)SG, page 16](#)
- [New Hardware Features in Release 12.2\(52\)SG, page 16](#)
- [New Software Features in Release 12.2\(52\)SG, page 17](#)
- [New Hardware Features in Release 12.2\(50\)SG5, page 18](#)
- [New Software Features in Release 12.2\(50\)SG5, page 18](#)
- [New Software Features in Release 12.2\(50\)SG4, page 18](#)
- [New Software Features in Release 12.2\(50\)SG4, page 18](#)
- [New Hardware Features in Release 12.2\(50\)SG3, page 18](#)
- [New Software Features in Release 12.2\(50\)SG3, page 18](#)
- [New Hardware Features in Release 12.2\(50\)SG2, page 19](#)
- [New Software Features in Release 12.2\(50\)SG2, page 19](#)
- [New Hardware Features in Release 12.2\(50\)SG1, page 19](#)

- [New Software Features in Release 12.2\(50\)SG1, page 19](#)
- [New Hardware Features in Release 12.2\(50\)SG, page 19](#)
- [New Software Features in Release 12.2\(50\)SG, page 19](#)
- [New Hardware Features in Release 12.2\(46\)SG, page 20](#)
- [New Software Features in Release 12.2\(46\)SG, page 20](#)

New Hardware Features in Release 12.2(53)SG

Release 12.2(53)SG provides no new hardware for the Catalyst 4900M switch.

New Software Features in Release 12.2(53)SG

Release 12.2(53)SG provides the following Cisco IOS software features for the Catalyst 4900M switch:

- IP Multicast Load Splitting (Equal Cost Multipath (ECMP) using S, G and Next-hop)
- Cisco Network Assistant (CNA)
- 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 200 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 need not hold a complete routing table. A best practice design, where 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) should be used when OSPF for Routed Access is used in the wiring closet.

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. Additionally, Enterprise Services is required to enable the VRF-lite feature.

New Hardware Features in Release 12.2(52)SG

Release 12.2(52)SG provides no new hardware for the Catalyst 4900M series switch.

New Software Features in Release 12.2(52)SG

Release 12.2(52)SG provides the following new Cisco IOS software features for the Catalyst 4900M series switch:

- DHCPv6 Enhancements
 - DHCPv6 Ethernet Remote ID option
 - DHCPv6 Relay - Persistent Interface ID option DHCPv6 Relay Agent notification for Prefix Delegation
- EnergyWise
- HSRPv2 for IPv6
- Identity ACL Policy Enforcement Enhancement
 - Filter-ID
 - Per-user ACL
- Local WebAuth Enhancement
- Network Mobility Services Protocol
- Online Diagnostics
- PIM Accept Register - Rogue Multicast Server Protection (**route-map** option is not supported)
- QinQ Tunneling and Layer 2 Protocol Tunneling (“Configuring 802.1Q and Layer 2 Protocol Tunneling” chapter)
- Smart Call Home
- SSM Mapping
- Supported MIBs
 - Cisco Enhanced Image MIB
 - Cisco HSRP extension MIB
 - CISCO-CALLHOME-MIB.my
 - EnergyWise MIB
 - POE MIB
 - POE ext MIB
 - Entity-Diag-MIB
 - Bridge MIB
- VRF lite NSF support with routing protocols OSPF/EIGRP/BGP

New Hardware Features in Release 12.2(50)SG5

Release 12.2(50)SG5 provides no new hardware for the Catalyst 4900M series switch.

New Software Features in Release 12.2(50)SG5

Release 12.2(50)SG5 provides no new software for the Catalyst 4900M series switch.

New Hardware Features in Release 12.2(50)SG4

Release 12.2(50)SG4 provides no new hardware for the Catalyst 4900M series switch.

New Software Features in Release 12.2(50)SG4

Release 12.2(50)SG4 provides no new software for the Catalyst 4900M series switch.

New Hardware Features in Release 12.2(50)SG3

Release 12.2(50)SG3 provides the following hardware for the Catalyst 4500 series switch:

- CVR-X2-SFP10G
Hot-swappable input/output (I/O) converter module that fits into a 10-Gigabit Ethernet X2 slot on a switch or line card module. Hosts one 10-Gigabit Ethernet SFP+ transceiver module.
- SFP-10G-SR
Cisco 10GBASE-SR SFP+ Module for MMF

New Software Features in Release 12.2(50)SG3

Release 12.2(50)SG3 provides no new features for the Catalyst 4500 series switch.

New Hardware Features in Release 12.2(50)SG2

Release 12.2(50)SG2 provides no new hardware for the Catalyst 4900M series switch.

New Software Features in Release 12.2(50)SG2

Release 12.2(50)SG2 provides no new software for the Catalyst 4900M series switch.

New Hardware Features in Release 12.2(50)SG1

Release 12.2(50)SG1 provides no new hardware for the Catalyst 4900M series switch.

New Software Features in Release 12.2(50)SG1

Release 12.2(50)SG1 provides the following new Cisco IOS software features for the Catalyst 4900M series switch:

- EEM version 2

New Hardware Features in Release 12.2(50)SG

Release 12.2(50)SG provides the following new hardware for the Catalyst 4900M series switch:

- X2-10GB-ZR optical module
- X2-10GB-DWDM optical module

New Software Features in Release 12.2(50)SG

Release 12.2(50)SG provides the following Cisco IOS software features for the Catalyst 4900M series switch:



Note The following chapter references are for the *Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide*.

- Multicast VRF-lite (“Configuring VRF-Lite” chapter)
- IGMP Querier (“Configuring IGMP Snooping” chapter)
- Bidirectional PIM (“Configuring IP Multicast” chapter)
- Private VLAN trunks (“Configuring Private VLANs” chapter)
- DHCP Relay Agent for IPv6 (refer to Cisco IOS Release 12.2 mainline documentation)
- OSPF and EIGRP fast convergence and protection (Refer to the Cisco IOS Release 12.4 documentation)
- CDP 2nd Port Status TLV (no configuration required on the switch)
- Flexible Authentication Sequencing (“Configuring 802.1X” chapter)

- Multi-Authentication (“Configuring 802.1X” chapter)
- Open Authentication (“Configuring 802.1X” chapter)
- Web Authentication (“Configuring Web Authentication” chapter)
- Inactivity Timer (“Configuring 802.1X” chapter)
- Downloadable ACLs (“Configuring Network Security with ACLs” chapter)
- ANCP Client (“Configuring ANCP Client” chapter)
- PPPoE Intermediate Agent (“PPPoE Circuit-Id Tag Processing” chapter)
- VTP version 3 (“Configuring VLANs, VTP, and VMPS” chapter)
- VRF-aware IP services (“Configuring VRF-Lite” chapter)
- Control Plane Policing (“Configuring CPP” chapter)
- **boot config** command (Refer to the Cisco IOS Release 12.4 documentation)
- Archiving Crashinfo Files (“Configuring Command-Line Interfaces” chapter)
- Unicast MAC filtering (“Configuring Network Security with ACLs” chapter)
- Configuration Rollback

New Hardware Features in Release 12.2(46)SG

Release 12.2(46)SG provides no new hardware for the Catalyst 4900M series switch.

New Software Features in Release 12.2(46)SG



Note

All features supported in Release 12.2(44)SG on Supervisor Engine 6-E (except for SSO) apply to this chassis.

Release 12.2(46)SG provides the following Cisco IOS software features for the Catalyst 4500 series switch:



Note

The following chapter references are for the *Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide*.

- 802.1X Catchup (Refer to the “Configuring 802.1X” chapter)
 - 802.1X Guest VLAN
 - 802.1X Critical Authentication
 - Wake on LAN
 - Radius Accounting
 - Radius Supplied Timeout
- ARP QoS (Refer to the “Configuring QoS” chapter)
- Per-VLAN CTI (Refer to the “Configuring QoS” chapter)
- Flash support for Layer 3 features

- FlexLink and FlexLink+ with MAC Address-Table Move Update (Refer to the “Configuring FlexLink” chapter)
- Ethernet Management Port (Refer to the “Configuring Interfaces” chapter)
- LLDP-MED: location TLV and MIB (Refer to the “Configuring LLDP and LLDP-MED” chapter)
- Enhanced Object Tracking (EOT) ((Refer to the Cisco IOS Release 12.2 documentation)
 - HSRP with EOT
 - VRRP with EOT
 - GLBP with EOT
 - IP SLA with EOT
 - Reliable Backup Static Routing with EOT
- RSPAN (Refer to the “Configuring SPAN and RSPAN” chapter)
- CFM 802.1ag (Refer to the “Configuring Ethernet CFM and OAM” chapter)
- E-OAM 802.3ah (Refer to the “Configuring Ethernet CFM and OAM” chapter)
- Ethernet Management Port (Refer to the “Configuring Interfaces” chapter)
- Embedded management (Refer to the Cisco IOS Release 12.4 documentation)
- MAC notify MIB (Refer to the Cisco IOS Release 12.4 documentation)
- BGP (Refer to the Cisco IOS Release 12.4 documentation)
- 802.1X Dynamic VLAN Assignment (Refer to the “Configuring 802.1X” chapter)
- 802.1X MAC Authentication Bypass (Refer to the “Configuring 802.1X” chapter)
- 802.1X with VVID/PVID (Refer to the “Configuring 802.1X” chapter)
- Eight configurable queues per port (Refer to the “Configuring QoS” chapter)
- VSS client with PagP+

After configuring VSS dual-active on a Catalyst 6500 switches, the Catalyst 4500 series switch can detect VSS dual-active with PagP+ support.
- IP SLA (Refer to the Cisco IOS Release 12.2 documentation)
- 802.1ab LLDP and 802.1ab LLDP-MED (Refer to the “Configuring LLDP and LLDP-MED” chapter)
- X2 Link Debounce Timer (Refer to the “Configuring Interfaces” chapter)
- Resilient Ethernet Protocol (REP) (Refer to the “Configuring REP” chapter)

Minimum and Recommended ROMMON Release

Table 5 Minimum and Recommended ROMMON Release for Catalyst 4900M

Minimum ROMMON Release	Recommended ROMMON Release
12.2(40r)XO	12.2(44r)SG5

Limitations and Restrictions

- The WS-X4920-GB-RJ45 card performs at wire speed until it operates at 99.6% utilization. Beyond this rate, the card will lose some packets.
- Compact Flash is not supported on a Cisco Catalyst 4900M switch running Cisco IOS Release 12.2(40)XO. Attempting to use Compact Flash may corrupt your data.
- IP classful routing is not supported; do not use the **no ip classless** command; it will have no effect, as only classless routing is supported. The command **ip classless** is not supported as classless routing is enabled by default.
- A Layer 2 LACP channel cannot be configured with the spanning tree PortFast feature.
- Netbooting using a boot loader image is not supported. See the [“Troubleshooting” section on page 143](#) for details on alternatives.
- An unsupported default CLI for mobile IP is displayed in the HSRP configuration. Although this CLI will not harm your system, you might want to remove it to avoid confusion.

Workaround: Display the configuration with the **show standby** command, then remove the CLI. Here is sample output of the **show standby GigabitEthernet1/1** command:

```
switch(config)# interface g1/1
switch(config)# no standby 0 name (0 is hsrp group number)
```

- For HSRP “preempt delay” to function consistently, you must use the **standby delay minimum** command. Be sure to set the delay to more than 1 hello interval, thereby ensuring that a hello is received before HSRP leaves the initiate state.
Use the **standby delay reload** option if the router is rebooting after reloading the image.
- You can run only .1q-in-.1q packet pass-through with Catalyst 4900M switch.
- For PVST and Catalyst 4900M switch VLANs, Cisco IOS Release 12.2(40)XO and higher support a maximum of 3000 spanning tree port instances. If you want to use more than this number of instances, you should use MST rather than PVST.
- Because the Catalyst 4900M switch supports the FAT filesystem, the following restrictions apply:
 - The **verify** and **squeeze** commands are not supported.
 - The **rename** command is supported in FAT file system.
For the Catalyst 4900M switch, the **rename** command has been added for bootflash and slot0. For all other supervisor engines, the **rename** command is supported for nvram devices only.
 - the **fsck** command is supported for slot0 device. It is not supported in the file systems on supervisor engines other than 6-E.
 - In the FAT file system, the IOS **format bootflash:** command erases user files only. It does not erase system configuration.
 - The FAT file system supports a maximum of 63 characters for file/directory name. The maximum for path length is 127 characters.
 - The FAT file system does not support the following characters in file/directory names: { } # % ^ and space characters.
 - The FAT file system honors the Microsoft Windows file attribute of "read-only" and "read-write", but it does not support the Windows file "hidden" attribute.
 - Supervisor Engine 6-E uses the FAT file system for compact flash (slot0). If a compact flash is not formatted in FAT file system (such as compact flash on a supervisor engine other than 6-E), the switch does not recognize it.

- The Fast Ethernet port (10/100) on the supervisor module is active in ROMMON mode only.
- If an original packet is dropped due to transmit queue shaping and/or sharing configurations, a SPAN packet copy can still be transmitted on the SPAN port.
- All software releases support a maximum of 16,000 IGMP snooping group entries.
- Use the **no ip unreachable** command on all interfaces with ACLs configured for performance reasons.
- The threshold for the Dynamic Arp Inspection err-disable function is set to 15 ARP packets per second per interface. You should adjust this threshold depending on the network configuration. The CPU should not receive DHCP packets at a sustained rate greater than 1000 pps.
- If you first configure an IP address or IPv6 address on a Layer 3 port, then change the Layer 3 port to a Layer 2 port with the **switchport** command, and finally change it back to a Layer 3 port, the original IP/IPv6 address will be lost.
- If a Catalyst 4900M switch requests information from the Cisco Secure Access Control Server (ACS) and the message exchange times out because the server does not respond, a message similar to this appears:

```
00:02:57: %RADIUS-4-RADIUS_DEAD: RADIUS server 172.20.246.206:1645,1646 is not responding.
```

If this message appears, check that there is network connectivity between the switch and the ACS. You should also check that the switch has been properly configured as an AAA client on the ACS.

- For IP Port Security (IPSG) for static hosts, the following apply:
 - As IPSG learns the static hosts on each interface, the switch CPU may hit 100 per cent if there are a large number of hosts to learn. The CPU usage will drop once the hosts are learned.
 - IPSG violations for static hosts are printed as they occur. If multiple violations occur simultaneously on different interfaces, the CLI displays the last violation. For example, if IPSG is configured for 10 ports and violations exist on ports 3,6 and 9, the violation messages are printed only for port 9.
 - Inactive host bindings will appear in the device tracking table when either a VLAN is associated with another port or a port is removed from a VLAN. So, as hosts are moved across subnets, the hosts are displayed in the device tracking table as INACTIVE.
 - Autostate SVI does not work on EtherChannel.
- When ipv6 is enabled on an interface via any CLI, it is possible to see the following message:

```
% Hardware MTU table exhausted
```

In such a scenario, the ipv6 MTU value programmed in hardware will be different from the ipv6 interface MTU value. This will happen if there is no room in the hw MTU table to store additional values.

You must free up some space in the table by unconfiguring some unused MTU values and subsequently disable/re-enable ipv6 on the interface or reapply the MTU configuration.

- To stop IPSG with Static Hosts on an interface, use the following commands in interface configuration submode:

```
Switch(config-if)# no ip verify source
Switch(config-if)# no ip device tracking max"
```

To enable IPSG with Static Hosts on a port, issue the following commands:

```
Switch(config)# ip device tracking ****enable IP device tracking globally
Switch(config)# ip device tracking max <n> ***set an IP device tracking maximum on int
```

```
Switch(config-if)# ip verify source tracking [port-security] ****activate IPSG on port
```

**Caution**

If you only configure the **ip verify source tracking [port-security]** interface configuration command on a port without enabling IP device tracking globally or setting an IP device tracking maximum on that interface, IPSG with Static Hosts will reject all the IP traffic from that interface.

**Note**

The issue above also applies to IPSG with Static Hosts on a PVLAN Host port.

- IPv6 ACL is not supported on a Catalyst 4900M switchport. IPv6 packets cannot be filtered on switchports using any of the known methods (PACL, VACL, or MACLs).
- Class-map match statements using **match ip prec | dscp** match only IPv4 packets whereas matches performed with **match prec | dscp** match both IPv4 and IPv6 packets.
- IPv6 QoS hardware switching is disabled if the policy-map contains IPv6 ACL and match cos in the same class-map with the ipv6 access-list has any mask range between /81 and /127. It results in forwarding packets to software which efficiently disable the QoS.
- Management port does not support *non-VRF* aware features.
- When you enter the **permit any any ?** command you will observe the **octal** option, which is unsupported in Cisco IOS Release 12.2(52)SG.
CSCsy31324
- A Span destination of fa1 is not supported.
- The "keepalive" CLI is not supported in interface mode on the switch, although it will appear in the running configuration. This behaviour has no impact on functionality.

Caveats

Caveats describe unexpected behavior in Cisco IOS releases. Caveats listed as open in a prior release are carried forward to the next release as either open or resolved.

**Note**

For the latest information on PSIRTS, refer to the Security Advisories on CCO at the following URL:
<http://tools.cisco.com/security/center/publicationListing.x>

Open Caveats in Cisco IOS Release 12.2(53)SG3

This section lists the open caveats in Cisco IOS Release 12.2(53)SG3:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.
On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCs139767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.

- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- a. Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- b. **Shut** any one REP port in the segment to cause a failure in that segment.
- c. **No-shut** that port to restore normal REP topology with one ALT port.
- d. Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- When two WS-C4900M chassis are attached to an optical ring and an optical switchover is performed to choose a different path, you might see CRC Align Errors and Sequence Errors after performing an end to end ping. The ping success rate ranges from 90% to 100%.

The errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of the Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

The issue is seen with release 12.2(44)XO and later releases.

Workaround: Enter **shut**, then **no shut**.

You may need to do this multiple times until the issue is resolved.

CSCsx80612

- When multiple streams of CRC errors are encountered on WS-C4900M configured with OAM Configuration of monitoring the frame errored seconds, OAM does not always report the value of errored frame seconds correctly.

To observe this issue, the following CLIs are configured with window size as the period for monitoring the errors and a low threshold equal to the number of CRC errored seconds seen/expected.

```
ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low
```

Workaround: Configure a lower value of low threshold such that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If RSPAN is configured on a WS-C4900M running Cisco IOS 12.2(46)SG, CPU utilization will be high.

Workaround: Disable RSPAN.

CSCsu81046

- When two Catalyst 4900M switches are attached to an optical ring and you perform an optical switchover to choose a different path, you might observe CRC Align Errors and Sequence after performing an end to end ping. The ping success rate ranges from 90% to 100%. The interface errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of a Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

Workaround: Enter the commands **shut** then **no shut**.

Occasionally, you need to re-enter the commands.

CSCsx80612

- When .1X with MDA is set in host mode and guest VLAN is enabled, when you pump traffic from a traffic generator at a high rate, a Security violation is wrongly flagged.

Workaround: None.

CSCsy38640

- When you enter the **show adjacency x.x.x.x internal** command for an adjacency, the packet counters are increment correctly but the byte counters remain 0.

Workaround: None.

CSCsu35604

- On a redundant switch running Cisco IOS Release 12.2(52)SG, after a ports is authorized through 802.1X, the **show dot1x interface statistics** command may display empty values on the standby supervisor engine.

The statistics are displayed properly on the active supervisor.

Workaround: None.

CSCsx64308

- When multiple streams of CRC errors are encountered on a WS-C4900M chassis configured with OAM monitoring of frame errored seconds, OAM does not report the value of errored frame seconds correctly if you configure the following CLIs:

```

ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low

```

Workaround: Configure a lower value for the low threshold so that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- When you request an on demand Call Home message send without specifying a profile name & the specified module returns an unknown diagnostic result, the following error message displays:

```

Switch# call-home send alert-group diagnostic module 2
Sending diagnostic info call-home message ...
Please wait. This may take some time ...
Switch#
*Jan  3 01:54:24.471: %CALL_HOME-3-ONDEMAND_MESSAGE_FAILED: call-home on-demand
message failed to send (ERR 18, The alert group is not subscribed)

```

Workaround: Specify a profile name when you enter the diagnostic command.

You might want to avoid requesting on demand send for invalid modules. First, enter the **show module** command to check for valid or present modules.

CSCsz05888

- If you simultaneously apply a service-policy to a port in the output direction and a service-policy to a vlan-range under that port in the output direction, the class-map hit counters in the output of the **show policy-map interface** command are wrong.

Workaround: None.

The queue transmit counters as well as the policing statistics (if any) are correct.

CSCsz20149

- When you configure EnergyWise power control on PoE ports with a time-based execution schedule, time entry executes without adjusting for daylight savings time.

Workaround: Manually re-enter all entries with new time settings.

CSCsy27389

- On wireless control system (WCS), some device information is incorrectly displayed for PCs sitting behind an lldp-med capable phone. Specifically, WCS displays the phone's serial number, model number, and software version in the PC's device information. All other information about the PC is correctly displayed on WCS.

This only happens when the switch is running network mobility service protocol (nmsp). It does not happen if the phone is CDP enabled.

Workaround: Use VLAN ID or name to differentiate the IP phone and the PC sitting behind the phone on the WCS. Specifically, the IP phone is detected on the voice VLAN, and the displayed information of serial number, model number, and software version is correct. However, a PC sitting behind the phone is detected on a data VLAN, and the displayed device information is wrong and should be ignored.

CSCsz34522

- When port-security is configured on normal trunks carrying primary and secondary private VLANs, its configuration can be erased from the running-config under the following circumstances:

Entering **shut/no shut** on the port after deleting a secondary VLAN. (CSCsz73895)

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** after deleting the VLAN.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, you can reconfigure port-security on the port only after reloading the switch.

CSCsz73895

- If a host is authenticated in the data VLAN, the STP state of the VLAN is blocked.

Assuming that you configured authentication open on the port and a host is authenticated on that port, if you unconfigure open auth (no authentication open), the STP state becomes blocked on an authenticated port.

The connected host is authenticated so it should be able to send traffic and the STP state should be Forwarding.

Workaround: Enter **shut**, then **no shut** on the port.

CSCta04665

- When CX1 or SFP+ are plugged into a OneX converter (CVR-X2-SFP10G) in a WS-X4908-10GE, the later requires 1 minute to boot the link.

Workaround: None.

CSCtc46340

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

- If *time* is not specified in the **link debounce** command, the default value depends on the supervisor engine. The default is 10 mS for C4900M, Supervisor Engine 6-E, and Supervisor Engine 6L-E. The default is 100 mS for all other supervisor engines.

Workaround: None.

Despite the different default value, you can configure any value in the time range.

CSCte51948

- If you observe a periodic increase in call or packet drops and a constant decrease in free memory available in your switch, you could use the **show memory debug leak** command. However, this command is CPU intensive; it might tear down your call or data session if used on live network.

The **show memory debug leak lowmem** command can work in extremely low memory conditions but might crash the switch due to its very high CPU intensity. It also takes between 20 and 90 minutes to complete.

Workaround: If call or packet drops persist, contact TAC rather than entering these commands on your own. CSCsi48986

- A switch may crash while loading BGP routes if the **ip cef accounting non-recursive** command is already configured.

Workaround: Disable the **ip cef accounting non-recursive** command.
(CSCtn68186)

Resolved Caveats in Cisco IOS Release 12.2(53)SG3

This section lists the resolved caveats in Release 12.2(53)SG3:

- The 10Gig uplink on a standby supervisor WS-X45-SUP6-E stops transmitting or receiving traffic after the old standby engine becomes active through an OIR (if the OIR is done quickly, within 5 seconds) of the active supervisor engine.

Workaround: Reload the active and standby supervisor engine.

While performing OIR of the supervisor engines, the engines must be removed completely before re-insertion. CSCsy70428

- When you configure **switchport block multicast** on a switch running Cisco IOS Release 12.2(53)SG1 or 12.2(50)SG6, Layer 2 multicast is not blocked.

Prior to Cisco IOS Release 12.2(53)SG1, 12.2(50)SG6, the **switchport block multicast** command would block IP Multicast, Layer 2 multicast, and broadcast traffic (CSCta61825).

Workaround: None CSCtb30327

Open Caveats in Cisco IOS Release 12.2(53)SG2

This section lists the open caveats in Cisco IOS Release 12.2(53)SG2:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map subtype. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.
Workaround: Enter the **show policy-map interface** command. (CSCsi71036)
- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.
Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)
- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.
Workaround: None. (CSCsl39767)
- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.
Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.
- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.
Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)
- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.
Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.
(CSCsq84796)
- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.
Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)
- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)
Workaround: None. (CSCsq99468)
- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:


```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:


```
config# interface interface-number
config-if# switchport
```

 (CSCsq47116)
- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.
You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.
Workaround: Unconfigure, then reconfigure the IFM on the port.
- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootstrap, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the `snmp mib target list vrf` command. This CLI is rejected even if the `vrf` is present in the DUT.

Workaround: None. (CSCsr95941)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- When two WS-C4900M chassis are attached to an optical ring and an optical switchover is performed to choose a different path, you might see CRC Align Errors and Sequence Errors after performing an end to end ping. The ping success rate ranges from 90% to 100%.

The errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of the Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

The issue is seen with release 12.2(44)XO and later releases.

Workaround: Enter **shut**, then **no shut**.

You may need to do this multiple times until the issue is resolved.

CSCsx80612

- When multiple streams of CRC errors are encountered on WS-C4900M configured with OAM Configuration of monitoring the frame errored seconds, OAM does not always report the value of errored frame seconds correctly.

To observe this issue, the following CLIs are configured with window size as the period for monitoring the errors and a low threshold equal to the number of CRC errored seconds seen/expected.

```
ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low
```

Workaround: Configure a lower value of low threshold such that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If RSPAN is configured on a WS-C4900M running Cisco IOS 12.2(46)SG, CPU utilization will be high.

Workaround: Disable RSPAN.

CSCsu81046

- When two Catalyst 4900M switches are attached to an optical ring and you perform an optical switchover to choose a different path, you might observe CRC Align Errors and Sequence after performing an end to end ping. The ping success rate ranges from 90% to 100%. The interface errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of a Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

Workaround: Enter the commands **shut** then **no shut**.

Occasionally, you need to re-enter the commands.

CSCsx80612

- When .1X with MDA is set in host mode and guest VLAN is enabled, when you pump traffic from a traffic generator at a high rate, a Security violation is wrongly flagged.

Workaround: None.

CSCsy38640

- When you enter the **show adjacency x.x.x.x internal** command for an adjacency, the packet counters are increment correctly but the byte counters remain 0.

Workaround: None.

CSCsu35604

- On a redundant switch running Cisco IOS Release 12.2(52)SG, after a ports is authorized through 802.1X, the **show dot1x interface statistics** command may display empty values on the standby supervisor engine.

The statistics are displayed properly on the active supervisor.

Workaround: None.

CSCsx64308

- When multiple streams of CRC errors are encountered on a WS-C4900M chassis configured with OAM monitoring of frame errored seconds, OAM does not report the value of errored frame seconds correctly if you configure the following CLIs:

```

ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low

```

Workaround: Configure a lower value for the low threshold so that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- The 10Gig uplink on a standby supervisor WS-X45-SUP6-E stops transmitting or receiving traffic after the old standby engine becomes active through an OIR (if the OIR is done quickly, within 5 seconds) of the active supervisor engine.

Workaround: Reload the active and standby supervisor engine.

While performing OIR of the supervisor engines, the engines must be removed completely before re-insertion.

CSCsy70428

- When you request an on demand Call Home message send without specifying a profile name & the specified module returns an unknown diagnostic result, the following error message displays:

```

Switch# call-home send alert-group diagnostic module 2
Sending diagnostic info call-home message ...
Please wait. This may take some time ...
Switch#
*Jan 3 01:54:24.471: %CALL_HOME-3-ONDEMAND_MESSAGE_FAILED: call-home on-demand
message failed to send (ERR 18, The alert group is not subscribed)

```

Workaround: Specify a profile name when you enter the diagnostic command.

You might want to avoid requesting on demand send for invalid modules. First, enter the **show module** command to check for valid or present modules.

CSCsz05888

- If you simultaneously apply a service-policy to a port in the output direction and a service-policy to a vlan-range under that port in the output direction, the class-map hit counters in the output of the **show policy-map interface** command are wrong.

Workaround: None.

The queue transmit counters as well as the policing statistics (if any) are correct.

CSCsz20149

- When you configure EnergyWise power control on PoE ports with a time-based execution schedule, time entry executes without adjusting for daylight savings time.

Workaround: Manually re-enter all entries with new time settings.

CSCsy27389

- On wireless control system (WCS), some device information is incorrectly displayed for PCs sitting behind an lldp-med capable phone. Specifically, WCS displays the phone's serial number, model number, and software version in the PC's device information. All other information about the PC is correctly displayed on WCS.

This only happens when the switch is running network mobility service protocol (nmsp). It does not happen if the phone is CDP enabled.

Workaround: Use VLAN ID or name to differentiate the IP phone and the PC sitting behind the phone on the WCS. Specifically, the IP phone is detected on the voice VLAN, and the displayed information of serial number, model number, and software version is correct. However, a PC sitting behind the phone is detected on a data VLAN, and the displayed device information is wrong and should be ignored.

CSCsz34522

- When port-security is configured on normal trunks carrying primary and secondary private VLANs, its configuration can be erased from the running-config under the following circumstances:

Entering **shut/no shut** on the port after deleting a secondary VLAN. (CSCsz73895)

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** after deleting the VLAN.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, you can reconfigure port-security on the port only after reloading the switch.

CSCsz73895

- If a host is authenticated in the data VLAN, the STP state of the VLAN is blocked.

Assuming that you configured authentication open on the port and a host is authenticated on that port, if you unconfigure open auth (no authentication open), the STP state becomes blocked on an authenticated port.

The connected host is authenticated so it should be able to send traffic and the STP state should be Forwarding.

Workaround: Enter **shut**, then **no shut** on the port.

CSCta04665

- When you configure **switchport block multicast** on a switch running Cisco IOS Release 12.2(53)SG1 or 12.2(50)SG6, Layer 2 multicast is not blocked.

Prior to Cisco IOS Release 12.2(53)SG1, 12.2(50)SG6, the **switchport block multicast** command would block IP Multicast, Layer 2 multicast, and broadcast traffic (CSCta61825).

Workaround: None

CSCtb30327

- When CX1 or SFP+ are plugged into a OneX converter (CVR-X2-SFP10G) in a WS-X4908-10GE, the later requires 1 minute to boot the link.

Workaround: None.

CSCtc46340

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

- If *time* is not specified in the **link debounce** command, the default value depends on the supervisor engine. The default is 10 mS for C4900M, Supervisor Engine 6-E, and Supervisor Engine 6L-E. The default is 100 mS for all other supervisor engines.

Workaround: None.

Despite the different default value, you can configure any value in the time range.

CSCte51948

- A switch may crash while loading BGP routes if the **ip cef accounting non-recursive** command is already configured.

Workaround: Disable the **ip cef accounting non-recursive** command.

(CSCtn68186)

Resolved Caveats in Cisco IOS Release 12.2(53)SG2

This section lists the resolved caveats in Release 12.2(53)SG2:

- A 802.1X port enabled for multi-authentication might not begin learning the MAC address of a successfully authenticated phone.

Workaround: Configure the port in multi-domain mode (rather than multi-auth mode) with the **authentication host-mode multi-domain** command

CSCtb28114

- When using subsecond timers for protocols like HSRP or OSPF, writing to bootflash causes high CPU, and potentially, protocol flapping.

Workaround: Avoid lengthy bootflash operations, like copying large files in IOS.

CSCsw84727

- The 4500-E and 4900M switches running IOS Release 12.2(53)SG1 or 12.2(50)SG6 may crash when the only Qos service-policy in a given VLAN is at the VLAN level.

The problem occurs when the following three conditions are met:

- A software-generated or software-switched packet exits an interface (P), which is a member of a VLAN (V).
- The packet is not a high priority; PAK_PRIORITY is not set.
- Of the three possible targets, port P, VLAN V, and port-VLAN PV in the output direction, a qos policy-map is attached only to the VLAN V in the output direction.

Workaround:

- Provided the VLAN-only policy-map has only marking actions., replace the VLAN-only policy-map with a port-VLAN policy-map on all the ports in the VLAN.
- Provided the VLAN-only policy-map has a policing action, retain the VLAN output policymap and attach a queuing action-only output policymap to all the ports in that VLAN.

The port level policy-map should appear as follows.

```
policy-map p1
  class class-default
    bandwidth percent 100
```


CSCte12571

- A PBR policy is not honored on a Supervisor Engine 6 running Cisco IOS Release 12.2(53)SG or 12.2(52)SG. Packets are forwarded through the normal routing table instead of through policy based routing.

This is a side effect of a heavily shared path.

Workaround: None.

CSCtc90702

- Upon upgrading to Cisco IOS Releases 12.2(52)SG, 12.2(52)XO, 12.2(53)SG, or 12.2(53)SG1, if the flash device name differs from the default name *flash:*, you might observe the following message continuously on your console:

```
%Error copying flash:/eem_pnt_2 (Invalid path)
```

Workaround: Rename the flash device to the default name *flash:*.

CSCte05909

- MAC learning does not work with Guest VLAN, Wake-on-LAN, and port security. When these features are enabled simultaneously in an interface, MAC learning does not work; none of the packets are forwarded.

Workaround: None.

You will need to disable Wake-on-LAN on the interface.

CSCtc58982

- When you delete and recreate an interface, the tacking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- After a .1X port is enabled for Guest VLAN, if you shut down the port connected to the RADIUS server so that the server goes dead and EAPOL packets are sent on that port, it is authorized in the access VLAN although the server is unreachable.

Workaround: Enter **shut**, then **no shut** on the port.

CSCsz63355

- When a switch enabled for explicit host tracking runs IGMPv3, ports that stopped sending IGMPv3 reports are displayed in the IGMPv3 table until a timeout. This behavior didn't exist in Cisco IOS Release 12.2(50)SG.

Workaround: Disable explicit host tracking in the affected VLANs.

CSCsz28612

- On a Catalyst 4900M, on each reload or power off/on, the system clock may lose (decrease) up to 59 seconds.

All software releases up to and including Cisco IOS Releases 12.2(31)SGA9, 12.2(50)SG6 and 12.2(53)SG1 are affected.

Workaround: After rebooting the switch, adjust the system clock with the **clock set** command.

CSCtc65375

- A switch running Cisco IOS Release 12.2(53)SG displays the message

%C4K_EBM-4-HOSTFLAPPING: happening between master loopback port and the source port during layer3 (IPv4 and IPv6) packets loop using ethernet oam (EOAM)

This message is does not impact performance.

Workaround: None.

CSCtc26043

- EnergyWise is enabled and you use the **energywise level level recurrence importance importance at minute hour day_of_month month day_of_week** interface configuration command to configure a recurring event on a switch. After the time changes from daylight savings time to standard time, the switch might

- Restart when it tries to power a PoE device
- Power on or off the PoE device at an incorrect time
- Fail

This occurs when the time change for the next year occurs after the time change for the current year.

Before the time change occurs, use one of these workarounds:

- Remove the recurring events from the EnergyWise configuration, do not use recurring events for a week, and reconfigure them a week after the time change occurs.
- Use the **energywise level level recurrence importance importance time-range time-range-name** interface configuration command to reschedule the events.
- Use the **power inline auto** interface configuration command to power on the PoE port.

CSCtc91312

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- If the router has a (*,G) entry for the group, then a fastdrop entry is not created to block the non-RPF packets from hitting the CPU.

Workaround: Create an ACL to block non-RPF packets from entering non-RPF ports.

CSCta93522

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

Open Caveats in Cisco IOS Release 12.2(53)SG1

This section lists the open caveats in Cisco IOS Release 12.2(53)SG1:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCs139767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QOS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tacking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- When two WS-C4900M chassis are attached to an optical ring and an optical switchover is performed to choose a different path, you might see CRC Align Errors and Sequence Errors after performing an end to end ping. The ping success rate ranges from 90% to 100%.

The errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of the Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

The issue is seen with release 12.2(44)XO and later releases.

Workaround: Enter **shut**, then **no shut**.

You may need to do this multiple times until the issue is resolved.

CSCsx80612

- When multiple streams of CRC errors are encountered on WS-C4900M configured with OAM Configuration of monitoring the frame errored seconds, OAM does not always report the value of errored frame seconds correctly.

To observe this issue, the following CLIs are configured with window size as the period for monitoring the errors and a low threshold equal to the number of CRC errored seconds seen/expected.

```
ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low
```

Workaround: Configure a lower value of low threshold such that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If RSPAN is configured on a WS-C4900M running Cisco IOS 12.2(46)SG, CPU utilization will be high.

Workaround: Disable RSPAN.

CSCsu81046

- When two Catalyst 4900M switches are attached to an optical ring and you perform an optical switchover to choose a different path, you might observe CRC Align Errors and Sequence after performing an end to end ping. The ping success rate ranges from 90% to 100%. The interface errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of a Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

Workaround: Enter the commands **shut** then **no shut**.

Occasionally, you need to re-enter the commands.

CSCsx80612

- When .1X with MDA is set in host mode and guest VLAN is enabled, when you pump traffic from a traffic generator at a high rate, a Security violation is wrongly flagged.

Workaround: None.

CSCsy38640

- When you enter the **show adjacency x.x.x.x internal** command for an adjacency, the packet counters are increment correctly but the byte counters remain 0.

Workaround: None.

CSCsu35604

- On a redundant switch running Cisco IOS Release 12.2(52)SG, after a ports is authorized through 802.1X, the **show dot1x interface statistics** command may display empty values on the standby supervisor engine.

The statistics are displayed properly on the active supervisor.

Workaround: None.

CSCsx64308

- When the ports connecting a RADIUS server and a client are placed in different VLANs, and you enter the **ip radius source-interface** command and perform two SSO switchovers, the authenticated session is lost.

Workaround: Re-authenticate the client.

CSCsx94066

- When multiple streams of CRC errors are encountered on a WS-C4900M chassis configured with OAM monitoring of frame errored seconds, OAM does not report the value of errored frame seconds correctly if you configure the following CLIs:

```

ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low

```

Workaround: Configure a lower value for the low threshold so that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- The 10Gig uplink on a standby supervisor WS-X45-SUP6-E stops transmitting or receiving traffic after the old standby engine becomes active through an OIR (if the OIR is done quickly, within 5 seconds) of the active supervisor engine.

Workaround: Reload the active and standby supervisor engine.

While performing OIR of the supervisor engines, the engines must be removed completely before re-insertion.

CSCsy70428

- When you request an on demand Call Home message send without specifying a profile name & the specified module returns an unknown diagnostic result, the following error message displays:

```

Switch# call-home send alert-group diagnostic module 2
Sending diagnostic info call-home message ...
Please wait. This may take some time ...
Switch#
*Jan 3 01:54:24.471: %CALL_HOME-3-ONDEMAND_MESSAGE_FAILED: call-home on-demand
message failed to send (ERR 18, The alert group is not subscribed)

```

Workaround: Specify a profile name when you enter the diagnostic command.

You might want to avoid requesting on demand send for invalid modules. First, enter the **show module** command to check for valid or present modules.

CSCsz05888

- When an access-list is attached to an interface under extreme hardware resource exhaustion, the ACL may not be automatically loaded into the hardware even if hardware resources later become available.

No TCAM entries are available for the new access-list.

Workaround: Manually remove and reapply the ACL after freeing hardware TCAM resources by removing or shortening other classification policies on the switch.

CSCsy85006

- If you simultaneously apply a service-policy to a port in the output direction and a service-policy to a vlan-range under that port in the output direction, the class-map hit counters in the output of the **show policy-map interface** command are wrong.

Workaround: None.

The queue transmit counters as well as the policing statistics (if any) are correct.

CSCsz20149

- After a .1X port is enabled for Guest VLAN, if you shut down the port connected to the RADIUS server so that the server goes dead and EAPOL packets are sent on that port, it is authorized in the access VLAN although the server is unreachable.

Workaround: Enter **shut**, then **no shut** on the port.

CSCsz63355

- When a switch enabled for explicit host tracking runs IGMPv3, ports that stopped sending IGMPv3 reports are displayed in the IGMPv3 table until a timeout. This behavior didn't exist in Cisco IOS Release 12.2(50)SG.

Workaround: Disable explicit host tracking in the affected VLANs.

CSCsz28612

- When you configure EnergyWise power control on PoE ports with a time-based execution schedule, time entry executes without adjusting for daylight savings time.

Workaround: Manually re-enter all entries with new time settings.

CSCsy27389

- On wireless control system (WCS), some device information is incorrectly displayed for PCs sitting behind an lldp-med capable phone. Specifically, WCS displays the phone's serial number, model number, and software version in the PC's device information. All other information about the PC is correctly displayed on WCS.

This only happens when the switch is running network mobility service protocol (nmsp). It does not happen if the phone is CDP enabled.

Workaround: Use VLAN ID or name to differentiate the IP phone and the PC sitting behind the phone on the WCS. Specifically, the IP phone is detected on the voice VLAN, and the displayed information of serial number, model number, and software version is correct. However, a PC sitting behind the phone is detected on a data VLAN, and the displayed device information is wrong and should be ignored.

CSCsz34522

- When port-security is configured on normal trunks carrying primary and secondary private VLANs, its configuration can be erased from the running-config under the following circumstances:

Entering **shut/no shut** on the port after deleting a secondary VLAN. (CSCsz73895)

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** after deleting the VLAN.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, you can reconfigure port-security on the port only after reloading the switch.

CSCsz73895

- If a host is authenticated in the data VLAN, the STP state of the VLAN is blocked.

Assuming that you configured authentication open on the port and a host is authenticated on that port, if you unconfigure open auth (no authentication open), the STP state becomes blocked on an authenticated port.

The connected host is authenticated so it should be able to send traffic and the STP state should be Forwarding.

Workaround: Enter **shut**, then **no shut** on the port.

CSCta04665

- If the router has a (*,G) entry for the group, then a fastdrop entry is not created to block the non-RPF packets from hitting the CPU.

Workaround: Create an ACL to block non-RPF packets from entering non-RPF ports.

CSCta93522

- A 802.1X port enabled for multi-authentication might not begin learning the MAC address of a successfully authenticated phone.

Workaround: Configure the port in multi-domain mode (rather than multi-auth mode) with the **authentication host-mode multi-domain** command

CSCtb28114

- On a Catalyst 4900M, on each reload or power off/on, the system clock may lose (decrease) up to 59 seconds.

All software releases up to and including Cisco IOS Releases 12.2(31)SGA9, 12.2(50)SG6 and 12,2(53)SG1 are affected.

Workaround: After rebooting the switch, adjust the system clock with the **clock set** command.

CSCtc65375

- When you configure **switchport block multicast** on a switch running Cisco IOS Release 12.2(53)SG1 or 12.2(50)SG6, Layer 2 multicast is not blocked.

Prior to Cisco IOS Release 12.2(53)SG1, 12.2(50)SG6, the **switchport block multicast** command would block IP Multicast, Layer 2 multicast, and broadcast traffic (CSCta61825).

Workaround: None

CSCtb30327

- When CX1 or SFP+ are plugged into a OneX converter (CVR-X2-SFP10G) in a WS-X4908-10GE, the later requires 1 minute to boot the link.

Workaround: None.

CSCtc46340

- A switch running Cisco IOS Release 12.2(53)SG displays the message

%C4K_EBM-4-HOSTFLAPPING: happening between master loopback port and the source port during layer3 (IPv4 and IPv6) packets loop using ethernet oam (EOAM)

This message is does not impact performance.

Workaround: None.

CSCtc26043

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

- When using subsecond timers for protocols like HSRP or OSPF, writing to bootflash causes high CPU, and potentially, protocol flapping.

Workaround: Avoid lengthy bootflash operations, like copying large files in IOS.

CSCsw84727

- EnergyWise is enabled and you use the **energywise level level recurrence importance importance at minute hour day_of_month month day_of_week** interface configuration command to configure a recurring event on a switch. After the time changes from daylight savings time to standard time, the switch might
 - Restart when it tries to power a PoE device
 - Power on or off the PoE device at an incorrect time
 - Fail

This occurs when the time change for the next year occurs after the time change for the current year.

Before the time change occurs, use one of these workarounds:

- Remove the recurring events from the EnergyWise configuration, do not use recurring events for a week, and reconfigure them a week after the time change occurs.

- Use the **energywise level *level* recurrence importance *importance* time-range *time-range-name*** interface configuration command to reschedule the events.
- Use the **power inline auto** interface configuration command to power on the PoE port.

CSCtc91312

- Upon upgrading to Cisco IOS Releases 12.2(52)SG, 12.2(52)XO, 12.2(53)SG, or 12.2(53)SG1, if the flash device name differs from the default name *flash:*, you might observe the following message continuously on your console:

```
%Error copying flash:/eem_pnt_2 (Invalid path)
```

Workaround: Rename the flash device to the default name *flash:*.

CSCte05909

- If *time* is not specified in the **link debounce** command, the default value depends on the supervisor engine. The default is 10 mS for C4900M, Supervisor Engine 6-E, and Supervisor Engine 6L-E. The default is 100 mS for all other supervisor engines.

Workaround: None.

Despite the different default value, you can configure any value in the time range.

CSCte51948

- A switch may crash while loading BGP routes if the **ip cef accounting non-recursive** command is already configured.

Workaround: Disable the **ip cef accounting non-recursive** command.

(CSCtn68186)

Resolved Caveats in Cisco IOS Release 12.2(53)SG1

This section lists the resolved caveats in Release 12.2(53)SG1:

- When you configure **switchport block multicast** on a port to block unknown multicast traffic, broadcast traffic is also blocked. Therefore, the port will receive neither unknown multicast or broadcast traffic.

All broadcast traffic (such as ARP request and DHCP discovery) are not received by the port. So, protocols that use such broadcasts stop working.

Workaround:None

CSCta61825

- On a WS-C4900M chassis running Cisco IOS Release 12.2(50)SG1, EIGRP adjacency breaks provided you do either of the following:
 - Enable ip pim sparse-mode on a VLAN interface in a vrf without enabling multicast routing on the vrf.
 - Enable multicast routing on the vrf and setting the STP threshold to infinity.

Workaround: Use static neighbors.

CSCsz61756

- When a service-policy is attached to a port-channel and that service-policy is configured to match CPU generated packets, the classification statistics do not increment for the CPU generated packets.

Workaround: Configure an access-list to permit the CPU generated packets and apply the ACL to the class-map.

CSCsy43967

- When you edit a policy-map to add a policer configuration, entering either the **do show policy-map interface** or **do show policy-map control-plane** command causes a system reload.

Workaround: Enter either the **show policy-map interface** and **show policy-map control-plane** commands in Exec mode and not in policy-map config mode.

CSCsy43261

- If a policy map is applied on an interface and the interface is inactive (i.e. the port is running in 10GE mode instead of twin gig mode), your WS-C4900M might crash with Vector 0xD00 when you enter the **show policy-map interface** command.

Workaround: Ensure that the port is active before apply the policy-map or entering the **show policy-map** command.

The command to activate a previously inactive interface is the following:

hw-module module [module number] port-group [group number] select [gigabitethernet]

CSCtb90328

- When you configure EnergyWise power control on the PoE ports of a WS-C4900M with a time-based execution schedule, time entry executes without adjusting for daylight savings time.

Workaround: Manually re-enter all entries with new time settings.

CSCsy27389

- If many ARP entries (47k) exist and you clear the ARP table, the system reloads and the switch crashes with the message:

```
ROM by abort at PC 0x0
```

Workaround: None.

Downgrade to Cisco IOS Release 12.2(50)SG3 if needed.

CSCta49512

- If you configure RSPAN on a WS-C4900M chassis running Cisco IOS Release 12.2(46)SG, CPU utilization will be high when monitored traffic is sent to the estination port.

Workaround: Disable RSPAN.

CSCsu81046

- When you configure a large number of ACLs on a Catalyst 4900M switch and enable statistics, the switch might exhibit high CPU utilization.

Certain applications such as IP Source Guard and QoS enable ACL statistics by default. Configuring such features trigger the high CPU.

High CPU usage is observed through the **show proc cpu** command. The output of the **show platform health** command reveals that the process using a high percentage of CPU is "K5AcICamStatsMan hw".

This issue can occur in any release after Cisco IOS Release 12.2(40)SG.

This issue is resolved in Cisco IOS Release 12.2(53)SG1 and 12.2(50)SG6.

Workaround: Reduce the size of the ACL, IPSG, and QoS configurations. If statistics are enabled explicitly for ACLs, disable them with the CLI.

If the high CPU is due to ACLs and IPSG, upgrade to the new software.

If the high CPU is due to the QoS configuration, upgrade the IOS image and enter the **no qos statistics classification** command.

CSCta54369

- If you enable VTP pruning after a switch is moved to VTP version 3, VLAN pruning does not happen on the trunks.

Workaround: Change the VTP version from 3 to version 2 or 1 and then revert to version 3.

CSCsy66803

- On a switch running Cisco IOS Release 12.2(50)SG or 12.2(52)SG, when an 802.1X port configured with PVLAN community VLAN receives a new PVLAN assignment from the AAA server, resetting the configuration on this interface may cause the switch to reload.

Workaround: None.

CSCsz38442

- When the vlan-port state changes on flexlink ports, the following two messages appear on the console:

```
A syslog warning message "%SM-4-BADEVENT: Event 'forward' is invalid for the current
state 'present': pm_vp .."
```

```
A traceback error message
```

This issue happens only on flexlink ports under the following two scenarios:

- You configure flexlink vlan load balancing before changing the port mode of a backup interface to trunk mode.
- Flexlink recovers from per vlan-port error disable states.

Workaround: None

The syslog and Traceback do not impact functionality. Flexlink states end up with correct states and there is no impact on traffic forwarding.

CSCta05317

- Per vlan-port error disable features (dhcp-rate-limit and arp-inspection) do not work on flexlink (without VLAN load balancing). When a violation occurs on the Active link, the corresponding vlan-port will not be error disabled.

The existing per-port error disable (that is, when a violation happens, the entire port will be error disabled) still works on flexlink.

Workaround: Use flexlink with VLAN load balancing.

If you do not want to use vlan load balancing, then enter the

switchport backup interface prefer vlan command on the Active interface, where vlan z is set to an unused vlan on the system

CSCta76320

- High CPU utilization might be observed on a switch for a prolonged period of time when a large number of packets on a VLAN/SVI are processed by software.

Workaround: None. Functionality is unaffected.

CSCsy32312

Open Caveats in Cisco IOS Release 12.2(53)SG

This section lists the open caveats in Cisco IOS Release 12.2(53)SG:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When a Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

(CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan** *vlan* command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map name**, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCs139767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QOS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QOS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tacking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- When two WS-C4900M chassis are attached to an optical ring and an optical switchover is performed to choose a different path, you might see CRC Align Errors and Sequence Errors after performing an end to end ping. The ping success rate ranges from 90% to 100%.

The errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of the Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

The issue is seen with release 12.2(44)XO and later releases.

Workaround: Enter **shut**, then **no shut**.

You may need to do this multiple times until the issue is resolved.

CSCsx80612

- When multiple streams of CRC errors are encountered on WS-C4900M configured with OAM Configuration of monitoring the frame errored seconds, OAM does not always report the value of errored frame seconds correctly.

To observe this issue, the following CLIs are configured with window size as the period for monitoring the errors and a low threshold equal to the number of CRC errored seconds seen/expected.

```
ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low
```

Workaround: Configure a lower value of low threshold such that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If RSPAN is configured on a WS-C4900M running Cisco IOS 12.2(46)SG, CPU utilization will be high.

Workaround: Disable RSPAN.

CSCsu81046

- When two Catalyst 4900M switches are attached to an optical ring and you perform an optical switchover to choose a different path, you might observe CRC Align Errors and Sequence after performing an end to end ping. The ping success rate ranges from 90% to 100%. The interface errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of a Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

Workaround: Enter the commands **shut** then **no shut**.

Occasionally, you need to re-enter the commands.

CSCsx80612

- When .1X with MDA is set in host mode and guest VLAN is enabled, when you pump traffic from a traffic generator at a high rate, a Security violation is wrongly flagged.

Workaround: None.

CSCsy38640

- When you enter the **show adjacency x.x.x.x internal** command for an adjacency, the packet counters are increment correctly but the byte counters remain 0.

Workaround: None.

CSCsu35604

- On a redundant switch running Cisco IOS Release 12.2(52)SG, after a ports is authorized through 802.1X, the **show dot1x interface statistics** command may display empty values on the standby supervisor engine.

The statistics are displayed properly on the active supervisor.

Workaround: None.

CSCsx64308

- When the ports connecting a RADIUS server and a client are placed in different VLANs, and you enter the **ip radius source-interface** command and perform two SSO switchovers, the authenticated session is lost.

Workaround: Re-authenticate the client.

CSCsx94066

- When multiple streams of CRC errors are encountered on a WS-C4900M chassis configured with OAM monitoring of frame errored seconds, OAM does not report the value of errored frame seconds correctly if you configure the following CLIs:

```

ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low

```

Workaround: Configure a lower value for the low threshold so that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If you enable VTP pruning after a switch is moved to VTP version 3, VLAN pruning does not happen on the trunks.

Workaround: Change the VTP version from 3 to version 2 or 1 and then revert to version 3.

CSCsy66803

- The 10Gig uplink on a standby supervisor WS-X45-SUP6-E stops transmitting or receiving traffic after the old standby engine becomes active through an OIR (if the OIR is done quickly, within 5 seconds) of the active supervisor engine.

Workaround: Reload the active and standby supervisor engine.

While performing OIR of the supervisor engines, the engines must be removed completely before re-insertion.

CSCsy70428

- When you request an on demand Call Home message send without specifying a profile name & the specified module returns an unknown diagnostic result, the following error message displays:

```

Switch# call-home send alert-group diagnostic module 2
Sending diagnostic info call-home message ...
Please wait. This may take some time ...
Switch#
*Jan 3 01:54:24.471: %CALL_HOME-3-ONDEMAND_MESSAGE_FAILED: call-home on-demand
message failed to send (ERR 18, The alert group is not subscribed)

```

Workaround: Specify a profile name when you enter the diagnostic command.

You might want to avoid requesting on demand send for invalid modules. First, enter the **show module** command to check for valid or present modules.

CSCsz05888

- When an access-list is attached to an interface under extreme hardware resource exhaustion, the ACL may not be automatically loaded into the hardware even if hardware resources later become available.

No TCAM entries are available for the new access-list.

Workaround: Manually remove and reapply the ACL after freeing hardware TCAM resources by removing or shortening other classification policies on the switch.

CSCsy85006

- If you simultaneously apply a service-policy to a port in the output direction and a service-policy to a vlan-range under that port in the output direction, the class-map hit counters in the output of the **show policy-map interface** command are wrong.

Workaround: None.

The queue transmit counters as well as the policing statistics (if any) are correct.

CSCsz20149

- On a switch running Cisco IOS Release 12.2(50)SG or 12.2(52)SG, when an 802.1X port configured with PVLAN community VLAN receives a new PVLAN assignment from the AAA server, resetting the configuration on this interface may cause the switch to reload.

Workaround: None.

CSCsz38442

- After a .1X port is enabled for Guest VLAN, if you shut down the port connected to the RADIUS server so that the server goes dead and EAPOL packets are sent on that port, it is authorized in the access VLAN although the server is unreachable.

Workaround: Enter **shut**, then **no shut** on the port.

CSCsz63355

- When a switch enabled for explicit host tracking runs IGMPv3, ports that stopped sending IGMPv3 reports are displayed in the IGMPv3 table until a timeout. This behavior didn't exist in Cisco IOS Release 12.2(50)SG.

Workaround: Disable explicit host tracking in the affected VLANs.

CSCsz28612

- When you configure EnergyWise power control on PoE ports with a time-based execution schedule, time entry executes without adjusting for daylight savings time.

Workaround: Manually re-enter all entries with new time settings.

CSCsy27389

- On wireless control system (WCS), some device information is incorrectly displayed for PCs sitting behind an lldp-med capable phone. Specifically, WCS displays the phone's serial number, model number, and software version in the PC's device information. All other information about the PC is correctly displayed on WCS.

This only happens when the switch is running network mobility service protocol (nmsp). It does not happen if the phone is CDP enabled.

Workaround: Use VLAN ID or name to differentiate the IP phone and the PC sitting behind the phone on the WCS. Specifically, the IP phone is detected on the voice VLAN, and the displayed information of serial number, model number, and software version is correct. However, a PC sitting behind the phone is detected on a data VLAN, and the displayed device information is wrong and should be ignored.

CSCsz34522

- When port-security is configured on normal trunks carrying primary and secondary private VLANs, its configuration can be erased from the running-config under the following circumstances:

Entering **shut/no shut** on the port after deleting a secondary VLAN. (CSCsz73895)

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** after deleting the VLAN.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, you can reconfigure port-security on the port only after reloading the switch.

CSCsz73895

- High CPU utilization might be observed on a switch for a prolonged period of time when a large number of packets on a VLAN/SVI are processed by software.

Workaround: None. Functionality is unaffected.

CSCsy32312

- If a host is authenticated in the data VLAN, the STP state of the VLAN is blocked.

Assuming that you configured authentication open on the port and a host is authenticated on that port, if you unconfigure open auth (no authentication open), the STP state becomes blocked on an authenticated port.

The connected host is authenticated so it should be able to send traffic and the STP state should be Forwarding.

Workaround: Enter **shut**, then **no shut** on the port.

CSCta04665

- When the vlan-port state changes on flexlink ports, the following two messages appear on the console:

```
A syslog warning message "%SM-4-BADEVENT: Event 'forward' is invalid for the current
state 'present': pm_vp .."
```

```
A traceback error message
```

This issue happens only on flexlink ports under the following two scenarios:

- You configure flexlink vlan load balancing before changing the port mode of a backup interface to trunk mode.
- Flexlink recovers from per vlan-port error disable states.

Workaround: None

The syslog and Traceback do not impact functionality. Flexlink states end up with correct states and there is no impact on traffic forwarding.

CSCta05317

- Per vlan-port error disable features (dhcp-rate-limit and arp-inspection) do not work on flexlink (without VLAN load balancing). When a violation occurs on the Active link, the corresponding vlan-port will not be error disabled.

The existing per-port error disable (that is, when a violation happens, the entire port will be error disabled) still works on flexlink.

Workaround: Use flexlink with VLAN load balancing.

If you do not want to use vlan load balancing, then enter the

switchport backup interface perfer vlan command on the Active interface, where vlan z is set to an unused vlan on the system

CSCta76320

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
```



```
permit icmp 2020::/96 any nd-ns sequence 10
deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

- When using subsecond timers for protocols like HSRP or OSPF, writing to bootflash causes high CPU, and potentially, protocol flapping.

Workaround: Avoid lengthy bootflash operations, like copying large files in IOS.

CSCsw84727

- EnergyWise is enabled and you use the **energywise level level recurrence importance importance at minute hour day_of_month month day_of_week** interface configuration command to configure a recurring event on a switch. After the time changes from daylight savings time to standard time, the switch might
 - Restart when it tries to power a PoE device
 - Power on or off the PoE device at an incorrect time
 - Fail

This occurs when the time change for the next year occurs after the time change for the current year.

Before the time change occurs, use one of these workarounds:

- Remove the recurring events from the EnergyWise configuration, do not use recurring events for a week, and reconfigure them a week after the time change occurs.
- Use the **energywise level level recurrence importance importance time-range time-range-name** interface configuration command to reschedule the events.
- Use the **power inline auto** interface configuration command to power on the PoE port.

CSCtc91312

- Upon upgrading to Cisco IOS Releases 12.2(52)SG, 12.2(52)XO, 12.2(53)SG, or 12.2(53)SG1, if the flash device name differs from the default name *flash:*, you might observe the following message continuously on your console:

```
%Error copying flash:/eem_pnt_2 (Invalid path)
```

Workaround: Rename the flash device to the default name *flash:*.

CSCtc05909

- If *time* is not specified in the **link debounce** command, the default value depends on the supervisor engine. The default is 10 mS for C4900M, Supervisor Engine 6-E, and Supervisor Engine 6L-E. The default is 100 mS for all other supervisor engines.

Workaround: None.

Despite the different default value, you can configure any value in the time range.

CSCte51948

- A switch may crash while loading BGP routes if the **ip cef accounting non-recursive** command is already configured.

Workaround: Disable the **ip cef accounting non-recursive** command.

(CSCtn68186)

Resolved Caveats in Cisco IOS Release 12.2(53)SG

This section lists the resolved caveats in Release 12.2(53)SG:

- On a Catalyst 4900M switch running Cisco IOS Release 12.2(46)SG, if you configure RSPAN, the CPU utilization will be high. This problem can occur when capturing traffic.

Workarounds: Disable RSPAN.

CSCsu81046

- When two Catalyst 4900M switches are attached to an optical ring and an optical switchover is performed on the ring to choose a different path, CRC Align Errors and Sequence Errors might be observed when you issue an end to end ping after the switchover. The ping success rate is between 90 and 100 per cent. The interface errors can occur with data traffic as well.

This issue is seen with the Ten-Gigabit ports of a Catalyst 4900M base board but not with the Ten-Gigabit ports of a WS-X4908-10GE line card.

Workaround: Enter **shut**, then **no shut**.

Sometimes you need to do this multiple times before the issue is resolved.

CSCsx80612

- Entering **shut/no shut** on the port after configuring **port-security vp err disable** and a violation occurs.

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** to recover the port.
- Configure **clear errdisable interface name vlan [range]** instead of entering **shut/no shut**.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, reconfigure port-security on the port after reloading the switch.

(CSCsy80415)

- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command.

(CSCsu03507)

Open Caveats in Cisco IOS Release 12.2(52)SG

This section lists the open caveats in Cisco IOS Release 12.2(52)SG:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCsl39767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QOS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QOS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tacking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- When two WS-C4900M chassis are attached to an optical ring and an optical switchover is performed to choose a different path, you might see CRC Align Errors and Sequence Errors after performing an end to end ping. The ping success rate ranges from 90% to 100%.

The errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of the Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

The issue is seen with release 12.2(44)XO and later releases.

Workaround: Enter **shut**, then **no shut**.

You may need to do this multiple times until the issue is resolved.

CSCsx80612

- When multiple streams of CRC errors are encountered on WS-C4900M configured with OAM Configuration of monitoring the frame errored seconds, OAM does not always report the value of errored frame seconds correctly.

To observe this issue, the following CLIs are configured with window size as the period for monitoring the errors and a low threshold equal to the number of CRC errored seconds seen/expected.

```
ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low
```

Workaround: Configure a lower value of low threshold such that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If RSPAN is configured on a WS-C4900M running Cisco IOS 12.2(46)SG, CPU utilization will be high.

Workaround: Disable RSPAN.

CSCsu81046

- When two Catalyst 4900M switches are attached to an optical ring and you perform an optical switchover to choose a different path, you might observe CRC Align Errors and Sequence after performing an end to end ping. The ping success rate ranges from 90% to 100%. The interface errors can also occur with data traffic.

This issue is seen with the TenGigabit ports of a Catalyst 4900M base board. It is not seen with the TenGigabit ports of a WS-X4908-10GE line card.

Workaround: Enter the commands **shut** then **no shut**.

Occasionally, you need to re-enter the commands.

CSCsx80612

- When .1X with MDA is set in host mode and guest VLAN is enabled, when you pump traffic from a traffic generator at a high rate, a Security violation is wrongly flagged.

Workaround: None.

CSCsy38640

- When you enter the **show adjacency x.x.x.x internal** command for an adjacency, the packet counters are increment correctly but the byte counters remain 0.

Workaround: None.

CSCsu35604

- On a redundant switch running Cisco IOS Release 12.2(52)SG, after a ports is authorized through 802.1X, the **show dot1x interface statistics** command may display empty values on the standby supervisor engine.

The statistics are displayed properly on the active supervisor.

Workaround: None.

CSCsx64308

- When the ports connecting a RADIUS server and a client are placed in different VLANs, and you enter the **ip radius source-interface** command and perform two SSO switchovers, the authenticated session is lost.

Workaround: Re-authenticate the client.

CSCsx94066

- When multiple streams of CRC errors are encountered on a WS-C4900M chassis configured with OAM monitoring of frame errored seconds, OAM does not report the value of errored frame seconds correctly if you configure the following CLIs:

```

ethernet oam link-monitor frame-seconds window
ethernet oam link-monitor frame-seconds threshold low

```

Workaround: Configure a lower value for the low threshold so that the frame errors are seen divided into the expected number of frame errored seconds.

CSCsy37181

- If you enable VTP pruning after a switch is moved to VTP version 3, VLAN pruning does not happen on the trunks.

Workaround: Change the VTP version from 3 to version 2 or 1 and then revert to version 3.

CSCsy66803

- The 10Gig uplink on a standby supervisor WS-X45-SUP6-E stops transmitting or receiving traffic after the old standby engine becomes active through an OIR (if the OIR is done quickly, within 5 seconds) of the active supervisor engine.

Workaround: Reload the active and standby supervisor engine.

While performing OIR of the supervisor engines, the engines must be removed completely before re-insertion.

CSCsy70428

- When you request an on demand Call Home message send without specifying a profile name & the specified module returns an unknown diagnostic result, the following error message displays:

```

Switch# call-home send alert-group diagnostic module 2
Sending diagnostic info call-home message ...
Please wait. This may take some time ...
Switch#
*Jan 3 01:54:24.471: %CALL_HOME-3-ONDEMAND_MESSAGE_FAILED: call-home on-demand
message failed to send (ERR 18, The alert group is not subscribed)

```

Workaround: Specify a profile name when you enter the diagnostic command.

You might want to avoid requesting on demand send for invalid modules. First, enter the **show module** command to check for valid or present modules.

CSCsz05888

- When an access-list is attached to an interface under extreme hardware resource exhaustion, the ACL may not be automatically loaded into the hardware even if hardware resources later become available.

No TCAM entries are available for the new access-list.

Workaround: Manually remove and reapply the ACL after freeing hardware TCAM resources by removing or shortening other classification policies on the switch.

CSCsy85006

- If you simultaneously apply a service-policy to a port in the output direction and a service-policy to a vlan-range under that port in the output direction, the class-map hit counters in the output of the **show policy-map interface** command are wrong.

Workaround: None.

The queue transmit counters as well as the policing statistics (if any) are correct.

CSCsz20149

- On a switch running Cisco IOS Release 12.2(50)SG or 12.2(52)SG, when an 802.1X port configured with PVLAN community VLAN receives a new PVLAN assignment from the AAA server, resetting the configuration on this interface may cause the switch to reload.

Workaround: None.

CSCsz38442

- After a .1X port is enabled for Guest VLAN, if you shut down the port connected to the RADIUS server so that the server goes dead and EAPOL packets are sent on that port, it is authorized in the access VLAN although the server is unreachable.

Workaround: Enter **shut**, then **no shut** on the port.

CSCsz63355

- When a switch enabled for explicit host tracking runs IGMPv3, ports that stopped sending IGMPv3 reports are displayed in the IGMPv3 table until a timeout. This behavior didn't exist in Cisco IOS Release 12.2(50)SG.

Workaround: Disable explicit host tracking in the affected VLANs.

CSCsz28612

- When you configure EnergyWise power control on PoE ports with a time-based execution schedule, time entry executes without adjusting for daylight savings time.

Workaround: Manually re-enter all entries with new time settings.

CSCsy27389

- On wireless control system (WCS), some device information is incorrectly displayed for PCs sitting behind an lldp-med capable phone. Specifically, WCS displays the phone's serial number, model number, and software version in the PC's device information. All other information about the PC is correctly displayed on WCS.

This only happens when the switch is running network mobility service protocol (nmosp). It does not happen if the phone is CDP enabled.

Workaround: Use VLAN ID or name to differentiate the IP phone and the PC sitting behind the phone on the WCS. Specifically, the IP phone is detected on the voice VLAN, and the displayed information of serial number, model number, and software version is correct. However, a PC sitting behind the phone is detected on a data VLAN, and the displayed device information is wrong and should be ignored.

CSCsz34522

- When port-security is configured on normal trunks carrying primary and secondary private VLANs, its configuration can be erased from the running-config under the following circumstances:

Entering **shut/no shut** on the port after deleting a secondary VLAN. (CSCsz73895)

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** after deleting the VLAN.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, you can reconfigure port-security on the port only after reloading the switch.

CSCsz73895

Entering **shut/no shut** on the port after configuring **port-security vp err disable** and a violation occurs. (CSCsz80415)

Workarounds:

- Configure error recovery for port-security violation instead of entering **shut/no shut** to recover the port.
- Configure **clear errdisable interface name vlan [range]** instead of entering **shut/no shut**.
- Configure port-security aging time to age out the MAC addresses before entering **shut/no shut**. Then, reconfigure port-security on the port after reloading the switch.
- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

- When using subsecond timers for protocols like HSRP or OSPF, writing to bootflash causes high CPU, and potentially, protocol flapping.

Workaround: Avoid lengthy bootflash operations, like copying large files in IOS.

CSCsw84727

- EnergyWise is enabled and you use the **energywise level level recurrence importance importance at minute hour day_of_month month day_of_week** interface configuration command to configure a recurring event on a switch. After the time changes from daylight savings time to standard time, the switch might
 - Restart when it tries to power a PoE device
 - Power on or off the PoE device at an incorrect time
 - Fail

This occurs when the time change for the next year occurs after the time change for the current year.

Before the time change occurs, use one of these workarounds:

- Remove the recurring events from the EnergyWise configuration, do not use recurring events for a week, and reconfigure them a week after the time change occurs.
- Use the **energywise level level recurrence importance importance time-range time-range-name** interface configuration command to reschedule the events.
- Use the **power inline auto** interface configuration command to power on the PoE port.

CSCtc91312

- Upon upgrading to Cisco IOS Releases 12.2(52)SG, 12.2(52)XO, 12.2(53)SG, or 12.2(53)SG1, if the flash device name differs from the default name *flash:*, you might observe the following message continuously on your console:

```
%Error copying flash:/eem_pnt_2 (Invalid path)
```

Workaround: Rename the flash device to the default name *flash:*.

CSCte05909

- If *time* is not specified in the **link debounce** command, the default value depends on the supervisor engine. The default is 10 mS for C4900M, Supervisor Engine 6-E, and Supervisor Engine 6L-E. The default is 100 mS for all other supervisor engines.

Workaround: None.

Despite the different default value, you can configure any value in the time range.

CSCte51948

Resolved Caveats in Cisco IOS Release 12.2(52)SG

This section lists the resolved caveats in Release 12.2(52)SG:

- Under normal operation, you will observe the following messages in the logs:

```
001298: .Oct  8 01:38:50.968: %C4K_SWITCHINGENINEMAN-4-TCAMINTERRUPT: flCam0
aPErr interrupt. errAddr: 0x2947 dPErr: 1 mPErr: 0 valid: 1
001299: .Oct  8 01:51:20.100: %C4K_SWITCHINGENINEMAN-4-TCAMINTERRUPT: flCam0
aPErr interrupt. errAddr: 0x2B59 dPErr: 1 mPErr: 0 valid: 1
```

Workaround: None

CSCsv17545

- Under control plane policing, control plane classes (the classes that are auto created by the **macro global apply system-cpp** command and use predefined ACLs to match traffic) increment both their packet and byte count. So, both counters are non-zero.

In contrast, data plane classes (the classes that are configured manually by user written ACLs), the byte counter increments as expected, but the packet count remains 0.

Workaround: None.

CSCsw16557

- On a Catalyst 4500, if an isolated private VLAN trunk interface flaps, the ingress and egress per-port per-vlan service policies are no longer applied on the port.

This impacts Cisco IOS Releases 12.2(31)SGA08, 12.2(37)SG, 12.2(40)SG, 12.2(44)SG, 12.2(46)SG, 12.2(50)SG, and 12.2(50)SG1.

Workarounds:

For a Classic Series Supervisor Engine, disable and configure QoS on the port.

For example, to configure Gig 2/1 as an isolated private VLAN trunk port, do the following:

```
Switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitEthernet 2/1
Switch(config-if)# no qos
Switch(config-if)# qos
Switch(config-if)# end
Switch#
```

You can configure the following EEM script to automate this workaround. QoS will be disabled and re-enabled whenever a port flaps.

```
logging event link-status global

event manager applet linkup-reqos
  event syslog pattern "changed state to up"
  action 1 cli command "enable"
  action 2 cli command "conf t"
  action 3 cli command "interface gigabitEthernet 2/1"
  action 4 cli command "no qos"
  action 5 cli command "qos"
```

CSCsw19087

- When you run an SNMP (getmany) query on cbQosPoliceStatsTable and cbQosREDClassStatsTable with a single SSH window (session), CPU utilization achieves 99 per cent. If you query cbQosPoliceStatsTable and cbQosREDClassStatsTable from 18 SSH sessions, a CPU-HOG error message displays.

Workaround: None, other than stopping the query.

CSCsw89720

- On a supervisor engine running Cisco IOS Release 12.2(50)SG or later releases with one or more ports configured for single-host mode, MAB, and authentication control-direction in, hosts are not authenticated through MAB when a port is configured for single-host mode and you enter the **unidirectional control in** command (Wake-on-LAN).

Workaround: Disable the **authentication control-direction in** command.

If you require **authentication control-direction in**, configure the port for multi-authentication or Multi-Domain Authentication (MDA).

CSCsx98360

- On a redundant switch running Cisco IOS Releases 12.2(50)SG or 12.2(50)SG1 where 802.1X VVID and port security are configured on a port, CDP MAC from the non 802.1X capable Cisco IP phone might not be added to the port security table on the standby supervisor engine.

Workaround: None.

This problem is fixed in Cisco IOS Releases 12.2(50)SG2 and 12.2(52)SG.

CSCsw29489

- On a switch running Cisco IOS Release 12.2(50)SG or 12.2(50)SG1 where 802.1X VVID and port security are configured on a port, inserting a non 802.1X capable Cisco IP phone with LLDP capability and a PC behind it may trigger a security violation.

Workaround: Turn off LLDP (on the switch) and the phone (from Call Manager).

This problem is fixed in 12.2(50)SG2 and 12.2(52)SG.

CSCsy21167

- Parity errors in the CPU's cache cause IOS to crash with a crashdump file like the following:

```
Switch# show platform crashdump

VECTOR 0

*** CRASH DUMP ***
02/09/2009 10:10:30
Last crash: 02/09/2009 10:10:30

Build: 12.2(20090206:234053) IPBASE
buildversion addr: 13115584

MCSR: 40000000 <--- non-zero value!
.
```

The key pieces of data are "VECTOR 0" and a MCSR value of 40000000, 20000000, or 10000000.

Workaround: Enter the **show platform cpu cache** command to launch an IOS algorithm that detects and recovers from parity errors in the CPU's cache. You will obtain a running count of the number of CPU cache parity errors that have been successfully detected and corrected on a running system:

```
Switch# show platform cpu cache
L1 Instruction Cache: ENABLED
L1 Data Cache: ENABLED
L2 Cache: ENABLED
Machine Check Interrupts: 5
L1 Instruction Cache Parity Errors: 3
L1 Instruction Cache Parity Errors (CPU30): 1
L1 Data Cache Parity Errors: 1
```

CSCsx15372

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name (device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An Unhandled Rommon Exception occurs while booting a WS-X4013+10GE for Cisco IOS Releases 12.2(31)SGA8, 12.2(31)SGA9, 12.2(46)SG, 12.2(46)SG1, 12.2(50)SG, 12.2(50)SG1.

Workaround: Upgrade to ROMMON version 1.2(31r)SGA4.

CSCsw91043

- On a switch running Cisco IOS Release 12.2(50)SG, supplicants authorized on PVLAN in multi-auth host mode are not moved to an Unauthorized state when the PVLAN is removed.

This problem occurs only when a port is configured with PVLAN and 802.1X multi-auth.

Workaround: Shut down then reopen the interface. (CSCsr58573)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)

- The switch may reload after destroying the `expExpressionTable` row via SNMP when you enable the **debug management expression evaluator** command.

Workaround: Remove the above debug command. (CSCsu67323)

- A router may crash when a privilege-level 15 user logs on with the callback or callback-dialstring attribute.

This problem is seen on all Catalyst 4500 or 4900 chassis running Cisco IOS Release 12.2.(50)SG. The problem occurs when the following conditions are present:

- The router is configured with AAA authentication and authorization.
- The AAA server runs CiscoSecure ACS 2.4.
- The callback or callback-dialstring attribute is configured on the AAA server for the user.

Workarounds: Do not configure the callback or callback-dialstring attribute for the user. If you use the callback-dialstring attribute in the TACACS+ profile, ensure that the NULL value is not configured. (CSCei62358)

- When you attempt an ISSU upgrade or downgrade between Cisco IOS Release 12.2(50)SG and 12.2(44)SG or 12.2(46)SG, the switch displays a traceback.

Workaround: None. (CSCsw32519)

- Entering the `channel-group x` mode or `channel-protocol` followed by `lacp` or `pagp` command on an `fa1` management interface causes the active supervisor engine to reload.

Port-channel functionality is not supported on the management interface.

This is a configuration error.

Workaround: None. (CSCsv91302)

- On classic series supervisors and Supervisor Engine 6-E running Cisco IOS Release 12.2(50)SG and later releases, egress traffic is not allowed on ports configured for Wake-on-LAN (through the **authentication control-direction in** command) and Multi-domain Authentication (MDA) (through the **authentication host-mode multi-domain** command) before the port is authorized.

Workaround: None. CSCsy29140

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- An Unhandled Rommon Exception occurs while booting a WS-X4013+10GE for Cisco IOS Releases 12.2(31)SGA8, 12.2(31)SGA9, 12.2(46)SG, 12.2(46)SG1, 12.2(50)SG, 12.2(50)SG1.

Workaround: Upgrade to ROMMON version 1.2(31r)SGA4.

CSCsw91043

- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command.
(CSCsu03507)

This issue may occur on switches with Supervisor 6(L)-E and 4900M running Cisco IOS Releases 12.2(40)SG, 12.2(44)SG, 12.2(46)SG, 12.2(50)SG-SG5. This issue does not affect switches with Supervisor V-10GE.

Resolved in 12.2(52)SG and beyond and 12.2(50)SG6 and beyond.

- Attempting to use the nested policy-map feature on Supervisor Engine 6-E or 6L-E can cause the switch to reboot.

This issue may occur on switches running Cisco IOS Releases 12.2(40)SG, 12.2(44)SG, 12.2(46)SG, 12.2(50)SG-SG5. This issue does not affect switches with Supervisor V-10GE.

This issue is resolved in 12.2(52)SG (and later) and 12.2(50)SG6 (and later) releases.

Workaround: Do not use the nested policy-map feature in Cisco IOS Release 12.2(40)SG and 12.2(44)SG. (CSCsy80664)

- On a switch running Cisco IOS 12.2(52)SG, when a port configured with 802.1X enters per vp errdisable mode because of a violation triggered by port security, DAI, DHCP snooping, or BPDU guard, the port's 802.1X sessions are not cleared despite the linkdown.

Workaround: None.

Do not configure 802.1X with other per vp errdisable features.

CSCsx74871

Open Caveats in Cisco IOS Release 12.2(50)SG6

This section lists the open caveats in Cisco IOS Release 12.2(50)SG6:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps `system-cpp-dhcp-cs`, `system-cpp-dhcp-sc`, and `system-cpp-dhcp-ss` may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

(CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.
Workaround: None. (CSCs139767)
- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.
Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.
- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.
Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)
- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.
Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family. (CSCsq84796)
- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.
Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)
- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.
Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)
- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)
Workaround: None. (CSCsq99468)
- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:


```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:


```
config# interface interface-number
config-if# switchport
```

 (CSCsq47116)
- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.
Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)
- IPv6 EIGRP routes are not learned through the port channel.
Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)
- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name (device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootstrap, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.
Workaround: None. (CSCsu42775)
- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.
Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)
- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.
Workarounds: Do one of the following:
 - Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
 - Enter **shut** then **no shut** on a 802.1X port.
 (CSCsv05205)
- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.
Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:
 - Enter any commands for that port.
 - Insert an SFP+ in that port.
 - Reinsert the removed SFP+ in any other port.
 (CSCsv90044)
- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.
Workaround: None. (CSCsr95941)
- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.
Workaround: None. (CSCsr95941)
- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.
Workaround: Do one of the following:
 - Do not attach routers to PVLAN isolated ports.
 - Disable igmp snooping (either globally or on the VLAN).
 - Do not use a router connected to PVLAN isolated port as a multicast source.
 (CSCsu39009)
- When you delete and recreate an interface, the tacking process is unable to track its state track.
Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)
- The switch may reload after destroying the expExpressionTable row via SNMP when you enable the **debug management expression evaluator** command.
Workaround: Remove the above debug command. (CSCsu67323)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- On a Catalyst 4900 series switch running Cisco IOS Release 12.2(46)SG and later versions, if you enter the **clear port-security dynamic interface fastethernet1** command, the switch reloads.

Do not enter this command if port security is not configured on the interface.

Do not enter this command on fa1.

Workaround: None. CSCtb16586

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RAACL to malfunction:

- ACL are applied on the output direction of the interface.

- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- On PVLAN trunk ports, learned MAC addresses age out unconditionally, resulting in flooding not only at the initial phase of frame delivery, but periodically at every MAC age interval. This behavior makes use of the **switchport block unicast** command risky, because it prevents communication.

Workaround: None. However, you cannot enter the **switchport block unicast** command on PVLAN trunk ports.

CSCtd49056

- When port security is configured or have a static MAC address on an isolated trunk port, the adjacencies for the port are resolved on the primary VLAN rather than on the secondary VLAN.

Workaround: None.

CSCtc79119

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

- If *time* is not specified in the **link debounce** command, the default value depends on the supervisor engine. The default is 10 mS for C4900M, Supervisor Engine 6-E, and Supervisor Engine 6L-E. The default is 100 mS for all other supervisor engines.

Workaround: None.

Despite the different default value, you can configure any value in the time range.

CSCte51948

Resolved Caveats in Cisco IOS Release 12.2(50)SG6

This section lists the resolved caveats in Release 12.2(50)SG6:

- When you run Supervisor Engine 6 with a large number of Layer 3 routes in the system, high CPU utilization may occur when minimal persistent ARP activity exists.

The **show processes cpu** command indicates that Cat4k Mgmt LoPri consumes a significant amount of CPU. The **show platform health** command indicates that K5L3FlcMan FwdEntry, K5L3Unciast IFE Review, and K5L3UnicastRpf IFE Review processes are running above their target utilization.

Note that large amounts of incomplete ARP entries may result from a scanning device or virus.

Workarounds:

- Reduce the number of Layer 3 routes.
- Prevent the ARP activity that triggers the high CPU utilization.

CSCta77487

- When you configure a large number of ACLs on a Supervisor 6-E/6L-E and enable statistics, the switch might exhibit high CPU utilization.

Certain applications such as IP Source Guard and QoS enable ACL statistics by default. Configuring such features trigger the high CPU.

High CPU usage is observed through the **show proc cpu** command. The output of the **show platform health** command reveals that the process using a high percentage of CPU is "K5AclCamStatsMan hw".

This issue can occur in any release after Cisco IOS Release 12.2(40)SG.

This issue is resolved in Cisco IOS Release 12.2(53)SG1 and 12.2(50)SG6.

Workaround: Reduce the size of the ACL, IPSG, and QoS configurations. If statistics are enabled explicitly for ACLs, disable them with the CLI.

If the high CPU is due to ACLs and IPSG, upgrade to the new software.

If the high CPU is due to the QoS configuration, upgrade the IOS image and enter the **no qos statistics classification** command.

CSCta54369

- If many ARP entries (47k) exist and you clear the ARP table, the system reloads and the switch crashes with the message:

```
ROM by abort at PC 0x0
```

Workaround: None.

Downgrade to Cisco IOS Release 12.2(50)SG3 if needed.

CSCta49512

- When using subsecond timers for protocols like HSRP or OSPF, writing to bootflash causes high CPU, and potentially, protocol flapping.

Workaround: Avoid lengthy bootflash operations, like copying large files in IOS.

CSCsw84727

- ARP entries learned on PVLAN SVIs are not aged out even if the **no ip sticky arp** command is configured globally.

ARP entries learned on normal SVIs are unaffected.

Workaround: Clear these ARP entries with the **clear ip arp** command.

CSCtb37718

- When port security and ARP inspection are configured together, the first ARP packet from a host, which is connected to the switch, could bypass the ARP inspection and be bridged out mistakenly.

Workaround: Disable port security.

CSCtb40187

- When you exit policy-map configuration mode without making changes to a policy-map on a switch configured with a service-policy for QoS, configuring an output service policy on an EtherChannel interface causes a link flap.

Workarounds: Configure identical policy-maps with different names so that each EtherChannel has its own policy. This action restricts the effect of this link flap to a limited number of EtherChannels.

CSCsz82795

- When a service-policy is attached to a port-channel and that service-policy is configured to match CPU generated packets, the classification statistics do not increment for the CPU generated packets.

Workaround: Configure an access-list to permit the CPU generated packets and apply the ACL to the class-map.

CSCsy43967

Open Caveats in Cisco IOS Release 12.2(50)SG5

This section lists the open caveats in Cisco IOS Release 12.2(50)SG5:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When a Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submenu. Then, apply the new class-map with the updated changes.

(CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCs139767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QOS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name (device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootstrap, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tacking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)

- The switch may reload after destroying the expExpressionTable row via SNMP when you enable the **debug management expression evaluator** command.

Workaround: Remove the above debug command. (CSCsu67323)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- a. Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- b. **Shut** any one REP port in the segment to cause a failure in that segment.
- c. **No-shut** that port to restore normal REP topology with one ALT port.

- d. Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- On a Catalyst 4900 series switch running Cisco IOS Release 12.2(46)SG and later versions, if you enter the **clear port-security dynamic interface fastethernet1** command, the switch reloads.

Do not enter this command if port security is not configured on the interface.

Do not enter this command on fa1.

Workaround: None. CSCtb16586

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RAACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RAACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

Resolved Caveats in Cisco IOS Release 12.2(50)SG5

This section lists the resolved caveats in Release 12.2(50)SG5:

- Under extremely rare conditions, a switch may silently stop forwarding traffic.
This caveat occurs when a register value is corrupted and you subsequently enable a Layer 3 feature.
Workaround: None (CSCsz48273)

Open Caveats in Cisco IOS Release 12.2(50)SG4

This section lists the open caveats in Cisco IOS Release 12.2(50)SG4:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGallInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

(CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan** *vlan* command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map name**, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCsl39767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name (device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.
Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)
- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.
Workarounds: Do one of the following:
 - Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
 - Enter **shut** then **no shut** on a 802.1X port.
 (CSCsv05205)
- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.
Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:
 - Enter any commands for that port.
 - Insert an SFP+ in that port.
 - Reinsert the removed SFP+ in any other port.
 (CSCsv90044)
- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.
Workaround: None. (CSCsr95941)
- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.
Workaround: None. (CSCsr95941)
- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.
Workaround: Do one of the following:
 - Do not attach routers to PVLAN isolated ports.
 - Disable igmp snooping (either globally or on the VLAN).
 - Do not use a router connected to PVLAN isolated port as a multicast source.
 (CSCsu39009)
- When you delete and recreate an interface, the tacking process is unable to track its state track.
Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)
- The switch may reload after destroying the expExpressionTable row via SNMP when you enable the **debug management expression evaluator** command.
Workaround: Remove the above debug command. (CSCsu67323)
- IP Router Option may not work with IGMP version 2.
Workaround: None. (CSCsv42869)
- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- a. Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- b. **Shut** any one REP port in the segment to cause a failure in that segment.
- c. **No-shut** that port to restore normal REP topology with one ALT port.
- d. Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- On a Catalyst 4900 series switch running Cisco IOS Release 12.2(46)SG and later versions, if you enter the **clear port-security dynamic interface fastethernet1** command, the switch reloads.

Do not enter this command if port security is not configured on the interface.

Do not enter this command on fa1.

Workaround: None. CSCtb16586

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
 permit icmp any any nd-ns sequence 10
```

```
deny ipv6 any any sequence 20

IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

Resolved Caveats in Cisco IOS Release 12.2(50)SG4

This section lists the resolved caveats in Release 12.2(50)SG4:

- A router may crash when a privilege-level 15 user logs on with the callback or callback-dialstring attribute.

'This problem is seen on all Catalyst 4500 or 4900 chassis running Cisco IOS Release 12.2.(50)SG. The problem occurs when the following conditions are present:

- The router is configured with AAA authentication and authorization.
- The AAA server runs CiscoSecure ACS 2.4.
- The callback or callback-dialstring attribute is configured on the AAA server for the user.

Workarounds: Do not configure the callback or callback-dialstring attribute for the user. If you use the callback-dialstring attribute in the TACACS+ profile, ensure that the NULL value is not configured. (CSCei62358)

- On a switch running Cisco IOS Release 12.2(50)SG, supplicants authorized on PVLAN in multi-auth host mode are not moved to an Unauthorized state when the PVLAN is removed.

This problem occurs only when a port is configured with PVLAN and 802.1X multi-auth.

Workaround: Shut down then reopen the interface. (CSCsr58573)

- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command. (CSCsu03507)

- Ordinarily, you observe the following messages frequently in the logs:

```
001298: .Oct  8 01:38:50.968: %C4K_SWITCHINGENINEMAN-4-TCAMINTERRUPT: flCam0
aPErr interrupt. errAddr: 0x2947 dPErr: 1 mPErr: 0 valid: 1
001299: .Oct  8 01:51:20.100: %C4K_SWITCHINGENINEMAN-4-TCAMINTERRUPT: flCam0
aPErr interrupt. errAddr: 0x2B59 dPErr: 1 mPErr: 0 valid: 1
```

They imply no impact to performance.

Workaround: None. (CSCsv17545)

- Entering the channel-group x mode or channel-protocol followed by lacp or pagp command on an fa1 management interface causes the active supervisor engine to reload.

Port-channel functionality is not supported on the management interface.

This is a configuration error.

Workaround: None. (CSCsv91302)

- When you attempt an ISSU upgrade or downgrade between Cisco IOS Release 12.2(50)SG and 12.2(44)SG or 12.2(46)SG, the switch displays a traceback.

Workaround: None. (CSCsw32519)

- On classic series supervisors and Supervisor Engine 6-E running Cisco IOS Release 12.2(50)SG and later releases, egress traffic is not allowed on ports configured for Wake-on-LAN (through the **authentication control-direction in** command) and Multi-domain Authentication (MDA) (through the **authentication host-mode multi-domain** command) before the port is authorized.

Workaround: None. CSCsy29140

- On a Catalyst 4900M switch, when you use a WS-X4908-10GE card with CVR-X2-SFP twin gig converters, the giga ports do not link up to the peer device that sends a remote fault. The **show int status | inc gi x/y** command indicates notconnect.

Similar behavior is observed with Supervisor Engine 6-E uplinks and the WS-X4706-10GE line card.

This behavior is seen in Cisco IOS Releases 12.2(50)SG thru 12.2(50)SG3 when the peer device sends a remote fault.

Workaround: Disable auto negotiation at both ends.

(CSCta02425)

Open Caveats in Cisco IOS Release 12.2(50)SG3

This section lists the open caveats in Cisco IOS Release 12.2(50)SG3:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.
On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps `system-cpp-dhcp-cs`, `system-cpp-dhcp-sc`, and `system-cpp-dhcp-ss` may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submenu. Then, apply the new class-map with the updated changes.

(CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan vlan** command, unconditional marking actions that are configured on the VLAN are not shown.

- Workaround:** None. However, if you enter the **show policy-map name**, the unconditional marking actions are displayed. (CSCsi94144)
- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCsl39767)
 - IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.
 - IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)
 - IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family. (CSCsq84796)
 - Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)
 - In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)
 - In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)
 - Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:


```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)
 - The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)
 - IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name (device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+, SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command. (CSCsu03507)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configured with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- On a switch running Cisco IOS Release 12.2(50)SG, supplicants authorized on PVLAN in multi-auth host mode are not moved to an Unauthorized state when the PVLAN is removed.

This problem occurs only when a port is configured with PVLAN and 802.1X multi-auth.

Workaround: Shut down then reopen the interface. (CSCsr58573)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbors.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.

- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tacking process is unable to track its state track.
Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)
- The switch may reload after destroying the expExpressionTable row via SNMP when you enable the **debug management expression evaluator** command.

Workaround: Remove the above debug command. (CSCsu67323)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- A router may crash when a privilege-level 15 user logs on with the callback or callback-dialstring attribute.

'This problem is seen on all Catalyst 4500 or 4900 chassis running Cisco IOS Release 12.2.(50)SG. The problem occurs when the following conditions are present:

- The router is configured with AAA authentication and authorization.
- The AAA server runs CiscoSecure ACS 2.4.
- The callback or callback-dialstring attribute is configured on the AAA server for the user.

Workarounds: Do not configure the callback or callback-dialstring attribute for the user. If you use the callback-dialstring attribute in the TACACS+ profile, ensure that the NULL value is not configured. (CSCEi62358)

- When you attempt an ISSU upgrade or downgrade between Cisco IOS Release 12.2(50)SG and 12.2(44)SG or 12.2(46)SG, the switch displays a traceback.

Workaround: None. (CSCsw32519)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- Entering the channel-group x mode or channel-protocol followed by lacp or pagp command on an fa1 management interface causes the active supervisor engine to reload.

Port-channel functionality is not supported on the management interface.

This is a configuration error.

Workaround: None. (CSCsv91302)

- Ordinarily, you observe the following messages frequently in the logs:

```
001298: .Oct  8 01:38:50.968: %C4K_SWITCHINGENINEMAN-4-TCAMINTERRUPT: f1Cam0
aPErr interrupt. errAddr: 0x2947 dPErr: 1 mPErr: 0 valid: 1
001299: .Oct  8 01:51:20.100: %C4K_SWITCHINGENINEMAN-4-TCAMINTERRUPT: f1Cam0
aPErr interrupt. errAddr: 0x2B59 dPErr: 1 mPErr: 0 valid: 1
```

They imply no impact to performance.

Workaround: None. (CSCsv17545)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- On classic series supervisors and Supervisor Engine 6-E running Cisco IOS Release 12.2(50)SG and later releases, egress traffic is not allowed on ports configured for Wake-on-LAN (through the **authentication control-direction in** command) and Multi-domain Authentication (MDA) (through the **authentication host-mode multi-domain** command) before the port is authorized.

Workaround: None. CSCsy29140

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- On a Catalyst 4900M switch, when you use a WS-X4908-10GE card with CVR-X2-SFP twin gig converters, the giga ports do not link up to the peer device that sends a remote fault. The **show int status | inc gi x/y** command indicates notconnect.

Similar behavior is observed with Supervisor Engine 6-E uplinks and the WS-X4706-10GE line card.

This behavior is seen in Cisco IOS Releases 12.2(50)SG thru 12.2(50)SG3 when the peer device sends a remote fault.

Workaround: Disable auto negotiation at both ends.

(CSCta02425)

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RAACL to malfunction:

- ACL are applied on the output direction of the interface.

- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

Resolved Caveats in Cisco IOS Release 12.2(50)SG3

This section lists the resolved caveats in Release 12.2(50)SG3:

- A Catalyst 4900M switch might crash if you insert/remove a TwinGig converter or boot it with installed TwinGig converters.

TwinGig converters are only supported on E-series supervisors and line cards. This bug does not affect systems without installed converters.

Workaround: None.

Once the switch has booted successfully and has detected all installed TwinGig converters, it is unlikely to crash unless you insert a converter. CSCsz49331

Open Caveats in Cisco IOS Release 12.2(50)SG2

This section lists the open caveats in Cisco IOS Release 12.2(50)SG2:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.
On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCs139767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name(device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootstrap, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFp+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+ , SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command. (CSCsu03507)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configurd with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- On a switch running Cisco IOS Release 12.2(50)SG, supplicants authorized on PVLAN in multi-auth host mode are not moved to an Uauthorized state when the PVLAN is removed.

This problem occurs only when a port is configured with PVLAN and 802.1X multi-auth.

Workaround: Shut down then reopen the interface. (CSCsr58573)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed , do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- The switch does not accept the `snmp mib target list vrf` command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbours.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tracking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)

- The switch may reload after destroying the `expExpressionTable` row via SNMP when you enable the **debug management expression evaluator** command.

Workaround: Remove the above debug command. (CSCsu67323)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- A router may crash when a privilege-level 15 user logs on with the `callback` or `callback-dialstring` attribute.

This problem is seen on all Catalyst 4500 or 4900 chassis running CiscoIOS Release 12.2.(50)SG. The problem occurs when the following conditions are present:

- The router is configured with AAA authentication and authorization.
- The AAA server runs CiscoSecure ACS 2.4.
- The `callback` or `callback-dialstring` attribute is configured on the AAA server for the user.

Workarounds: Do not configure the `callback` or `callback-dialstring` attribute for the user. If you use the `callback-dialstring` attribute in the TACACS+ profile, ensure that the NULL value is not configured. (CSCei62358)

- When you attempt an ISSU upgrade or downgrade between Cisco IOS Release 12.2(50)SG and 12.2(44)SG or 12.2(46)SG, the switch displays a traceback.

Workaround: None. (CSCsw32519)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- Entering the channel-group x mode or channel-protocol followed by lacp or pagp command on an fa1 management interface causes the active supervisor engine to reload.

Port-channel functionality is not supported on the management interface.

This is a configuration error.

Workaround: None. (CSCsv91302)

- Ordinarily, you observe the following messages frequently in the logs:

```
001298: .Oct 8 01:38:50.968: %C4K_SWITCHINGENGINEMAN-4-TCAMINTERRUPT: flCam0
aPErr interrupt. errAddr: 0x2947 dPErr: 1 mPErr: 0 valid: 1
001299: .Oct 8 01:51:20.100: %C4K_SWITCHINGENGINEMAN-4-TCAMINTERRUPT: flCam0
aPErr interrupt. errAddr: 0x2B59 dPErr: 1 mPErr: 0 valid: 1
```

They imply no impact to performance.

Workaround: None. (CSCsv17545)

- The host's MAC address is not synchronized to the standby supervisor engine after you unconfigure 802.1X on the port and reconnect the host to a IP phone (with CDP port status TLV support) that is connected to the switch.

If the switch were to run a supervisor switchover while in this state, the host's MAC address would not be present in the new active supervisor engine's MAC address table, causing possible connectivity interruption on the host.

Workaround: Enter the **shutdown** command, followed by the **no shutdown** command on the interface. This triggers relearning and synchronizing of the host's MAC to the standby supervisor engine. CSCsw91661

- On classic series supervisors and Supervisor Engine 6-E running Cisco IOS Release 12.2(50)SG and later releases, egress traffic is not allowed on ports configured for Wake-on-LAN (through the **authentication control-direction in** command) and Multi-domain Authentication (MDA) (through the **authentication host-mode multi-domain** command) before the port is authorized.

Workaround: None. CSCsy29140

- Class-map hit counters do not increment on the egress policy-map when it is attached to the primary VLAN on a PVLAN trunk ports. However, the traffic is properly classified and the actions configured in the policy are applied properly.

Workaround: None. CSCsy72343

- On a Catalyst 4900M switch, when you use a WS-X4908-10GE card with CVR-X2-SFP twin gig converters, the giga ports do not link up to the peer device that sends a remote fault. The **show int status | inc gi x/y** command indicates notconnect.

Similar behavior is observed with Supervisor Engine 6-E uplinks and the WS-X4706-10GE line card.

This behavior is seen in Cisco IOS Releases 12.2(50)SG thru 12.2(50)SG3 when the peer device sends a remote fault.

Workaround: Disable auto negotiation at both ends.

(CSCta02425)

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

Resolved Caveats in Cisco IOS Release 12.2(50)SG2

This section lists the resolved caveats in Release 12.2(50)SG2:

- Packets for traffic destined to SNAP host might be dropped if the ARP table indicates that the MAC entry is SNAP.

Workarounds:

1. Configure a static ARPA entry for host.
2. Upgrade to a future IOS release containing the fix.

CSCsu90780

- On a Catalyst 4500 switch running 12.2(50)SG or 12.2(50)SG1, when 802.1X VVID and port security are configured together on a switch port, inserting a non 802.1x capable Cisco IP phone with a PC behind it may trigger a security violation.

Workaround: None. CSCsv63638

- If you configure multiple REP segments, pre-emption in one segment brings down all REP segments.

Workaround: None. CSCsv91297

- On a Catalyst 4500 series switch, if an isolated private VLAN trunk interface flaps, the ingress per-port per-vlan policer is no longer applied on the port.

Affected Cisco IOS releases include 12.2(31)SGA08, 12.2(37)SG, 12.2(40)SG, 12.2(46)SG, and 12.2(50)SG.

Workaround: Disable and configure QoS, as follows:

```
Switch# conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# no qos
Switch(config)# qos
Switch(config)# end
Switch#
```

CSCsw19087

- On a Catalyst 4500 redundant switch running Cisco IOS Release 12.2(50)SG or 12.2(50)SG1, when 802.1X VVID and port security are configured together on a switch port, the CDP MAC from the non 802.1X capable Cisco IP phone may not be added to the port security table on the standby supervisor engine.

Workaround: None. CSCsw29489

- A crash occurs when you enter the **show idprom interface FastEthernet 1** command.

Workaround: None. CSCsw77413

- Hosts are not authenticated through MAB when you configure a port for single-host mode (with the **authentication host-mode single-host** command) and Wake-on-LAN (with the **authentication control-direction in** command).

Workarounds: Disable Wake-on-LAN with the **no authentication control-direction in** command.

CSCsx98360

- On a Catalyst 4500 series switch running Cisco IOS Release 12.2(50)SG or 12.2(50)SG1, when you configure both 802.1X VVID and port security together on a switch port, then insert a non-802.1X capable Cisco IP phone with LLDP capability and a PC behind it, you might trigger a security violation. The violation is triggered when the PC behind the phone gets authorized on the port before the IP phone sends LLDP packet.

Workaround: Turn off LLDP on the switch and Cisco IP phone from Call Manager.

CSCsy21167

Open Caveats in Cisco IOS Release 12.2(50)SG1

This section lists the open caveats in Cisco IOS Release 12.2(50)SG1:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCsl39767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QOS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QOS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
```

```
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name(device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+ , SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)
- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command. (CSCsu03507)
- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configurd with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface. (CSCso50921)
- On a switch running Cisco IOS Release 12.2(50)SG, supplicants authorized on PVLAN in multi-auth host mode are not moved to an Unauthorized state when the PVLAN is removed.

This problem occurs only when a port is configured with PVLAN and 802.1X multi-auth.

Workaround: Shut down then reopen the interface. (CSCsr58573)
- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)
- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)
- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

 - Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
 - Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)
- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed , do one of the following:

 - Enter any commands for that port.
 - Insert an SFP+ in that port.
 - Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the `snmp mib target list vrf` command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- The switch does not accept the `snmp mib target list vrf` command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbours.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tracking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)

- The switch may reload after destroying the `expExpressionTable` row via SNMP when you enable the **debug management expression evaluator** command.

Workaround: Remove the above debug command. (CSCsu67323)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- A router may crash when a privilege-level 15 user logs on with the `callback` or `callback-dialstring` attribute.

This problem is seen on all Catalyst 4500 or 4900 chassis running Cisco IOS Release 12.2.(50)SG. The problem occurs when the following conditions are present:

- The router is configured with AAA authentication and authorization.
- The AAA server runs CiscoSecure ACS 2.4.
- The `callback` or `callback-dialstring` attribute is configured on the AAA server for the user.

Workarounds: Do not configure the `callback` or `callback-dialstring` attribute for the user. If you use the `callback-dialstring` attribute in the TACACS+ profile, ensure that the NULL value is not configured. (CSCei62358)

- When you attempt an ISSU upgrade or downgrade between Cisco IOS Release 12.2(50)SG and 12.2(44)SG or 12.2(46)SG, the switch displays a traceback.

Workaround: None. (CSCsw32519)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.

- d. Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102  Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- Entering the channel-group x mode or channel-protocol followed by lacp or pagp command on an fa1 management interface causes the active supervisor engine to reload.

Port-channel functionality is not supported on the management interface.

This is a configuration error.

Workaround: None. (CSCsv91302)

- On a Catalyst 4900M switch, when you use a WS-X4908-10GE card with CVR-X2-SFP twin gig converters, the giga ports do not link up to the peer device that sends a remote fault. The **show int status | inc gi x/y** command indicates notconnect.

Similar behavior is observed with Supervisor Engine 6-E uplinks and the WS-X4706-10GE line card.

This behavior is seen in Cisco IOS Releases 12.2(50)SG thru 12.2(50)SG3 when the peer device sends a remote fault.

Workaround: Disable auto negotiation at both ends.

(CSCta02425)

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20

IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

Resolved Caveats in Cisco IOS Release 12.2(50)SG1

This section lists the resolved caveats in Release 12.2(50)SG1:

- When port security is configured on a port connected to a host via an IP phone and the host is disconnected, the host's MAC address is not removed from the port security MAC address table even if the IP phone and switch support the CDP 2nd port disconnect TLV feature.

Workaround: To remove the host's MAC address from the port security MAC address table, unconfigure and reconfigure port security on the port. (CSCsr74097)

Open Caveats in Cisco IOS Release 12.2(50)SG

This section lists the open caveats in Cisco IOS Release 12.2(50)SG:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.

- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.
On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submenu. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan** *vlan* command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map** *name*, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCsl39767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name(device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPs, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootup, the router port is created first.

(CSCsq63051)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- When you configure **ip source binding** statically on an interface, and then remove linecard on which the interface resides, the entries are not removed from the running config.

Workaround: Before removing a linecard, delete the statically configured **ip source binding** entries on any of the interfaces on the line-card. (CSCsv54529)

- If you configure OFM on an Etherchannel (with at least two interfaces), when you shut or remove the first member that joined the channel, the CFM neighbor is lost.

Workaround: Clear the errors with the **clear ethernet cfm errors** command in EXEC mode. (CSCsv43819)

- Duplicate serial number error messages are reported on switching One X Convertor with SFP+, SFP+, X2 to another port, the inserted port enters a faulty status.

This problem impacts X2, OneX converters, and SFP+ on the Supervisor Engine 6-E, and linecards.

Workaround: Remove and reinsert the One X Convertor with SFP+ , SFP+ alone, or X2 after some perceivable delay. (CSCsu43461)

- Ping does not execute prior to a posture validation.

Workaround: Reapply the identity policy on the interface with the **permit icmp** command. (CSCsu03507)

- On a Catalyst 4500 switch running 12.2(50)SG, when the access VLAN is deleted and then restored on a port configurd with 802.1x multi-auth, authorized 802.1X clients cannot pass traffic because the spanning tree remains in a Disabled state after the access VLAN is restored.

This problem occurs when an 802.1X client is authorized on a multi-auth port. After the access VLAN is deleted, then restored, the client is reauthorized but the spanning tree state of the access VLAN remains Disabled.

Workaround: Shut down then reopen the interface.

(CSCso50921)

- On a switch running Cisco IOS Release 12.2(50)SG, supplicants authorized on PVLAN in multi-auth host mode are not moved to an Unauthorized state when the PVLAN is removed.

This problem occurs only when a port is configured with PVLAN and 802.1X multi-auth.

Workaround: Shut down then reopen the interface. (CSCsr58573)

- When the switch port configured with 802.1X Multi-Domain Authentication (MDA) and Guest VLAN is connected to a non-802.1X supplicant PC through a hub, the port falls back to guest VLAN. Subsequently, it is stuck in the guest VLAN and ignores all EAPOL traffic from another 802.1X supplicant PC connected to the hub.

Workaround: None. (CSCsu42775)

- VTP databases do not propagate through promiscuous trunk ports. If only promiscuous trunks are configured, users will not see the VLAN updates on the other switches in the VTP domain.

Workaround: For VTP database propagation, configure ISL/dot1q trunk port. (CSCsu43445)

- Egress traffic may not be allowed when 802.1X is configured as a Unidirectional Controlled Port.

Workarounds: Do one of the following:

- Enter **spanning-tree portfast** then **authentication control-direction in** on a 802.1X port.
- Enter **shut** then **no shut** on a 802.1X port.

(CSCsv05205)

- When you remove an SFP+ from a OneX converter in a X2 slot, it takes roughly 45 seconds for the system to recognize this. Any commands during this time will indicate that the SFP+ is still present. Reinserting the SFP+ in another port or inserting another SFP+ in the same port can result in Duplicate Seeprom error message.

Workaround: When a log message appears indicating that the SFP+ has been removed, do one of the following:

- Enter any commands for that port.
- Insert an SFP+ in that port.
- Reinsert the removed SFP+ in any other port.

(CSCsv90044)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- The switch does not accept the snmp mib target list vrf command. This CLI is rejected even if the vrf is present in the DUT.

Workaround: None. (CSCsr95941)

- When a PVLAN isolated port is connected to a router serving as a multicast source, and you enable igmp snooping, the routers connected to the isolated ports display as PIM neighbours.

Workaround: Do one of the following:

- Do not attach routers to PVLAN isolated ports.
- Disable igmp snooping (either globally or on the VLAN).
- Do not use a router connected to PVLAN isolated port as a multicast source.

(CSCsu39009)

- When you delete and recreate an interface, the tracking process is unable to track its state track.

Workaround: Reconfigure tracking on the newly created interface. (CSCsr66876)

- The switch may reload after destroying the expExpressionTable row via SNMP when you enable the **debug management expression evaluator** command.

Workaround: Remove the above debug command. (CSCsu67323)

- IP Router Option may not work with IGMP version 2.

Workaround: None. (CSCsv42869)

- A router may crash when a privilege-level 15 user logs on with the callback or callback-dialstring attribute.

This problem is seen on all Catalyst 4500 or 4900 chassis running CiscoIOS Release 12.2.(50)SG. The problem occurs when the following conditions are present:

- The router is configured with AAA authentication and authorization.
- The AAA server runs CiscoSecure ACS 2.4.
- The callback or callback-dialstring attribute is configured on the AAA server for the user.

Workarounds: Do not configure the callback or callback-dialstring attribute for the user. If you use the callback-dialstring attribute in the TACACS+ profile, ensure that the NULL value is not configured. (CSCei62358)

- When you attempt an ISSU upgrade or downgrade between Cisco IOS Release 12.2(50)SG and 12.2(44)SG or 12.2(46)SG, the switch displays a traceback.

Workaround: None. (CSCsw32519)

- If VLAN Load Balancing is progressing, and you reconfigure VLAN Load Balancing to reflect different blocking ports, manual preemption does not occur.

Workaround: To reconfigure VLAN Load Balancing with a different configuration, do the following:

- Reconfigure the VLAN Load Balancing configuration on the desired REP ports.
- Shut** any one REP port in the segment to cause a failure in that segment.
- No-shut** that port to restore normal REP topology with one ALT port.
- Invoke manual preemption on a primary edge port to obtain VLAN Load Balancing with the new configuration.

(CSCsv69853)

- After posture validation succeeds, the following benign traceback messages may appear after you unconfigure the global RADIUS and IP device tracking commands:

```
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.101 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
%SM-4-BADEVENT: Event 'eouAAAAuthor' is invalid for the current state 'eou_abort':
eou_auth 4.1.0.102 Traceback= 101D9A88 10B76BB0 10B76FE0 10B7A114 10B7A340 1066A678
106617F8
```

This applies to classic or E-series Catalyst 4500 supervisor engines running Cisco IOS Release 12.2(50)SG

Workaround: None. (CSCsw14005)

- Entering the channel-group x mode or channel-protocol followed by lacp or pagp command on an fa1 management interface causes the active supervisor engine to reload.

Port-channel functionality is not supported on the management interface.

This is a configuration error.

Workaround: None. (CSCsv91302)

- On a Catalyst 4900M switch, when you use a WS-X4908-10GE card with CVR-X2-SFP twin gig converters, the giga ports do not link up to the peer device that sends a remote fault. The **show int status | inc gi x/y** command indicates notconnect.

Similar behavior is observed with Supervisor Engine 6-E uplinks and the WS-X4706-10GE line card.

- On a Catalyst 4900M switch, when you use a WS-X4908-10GE card with CVR-X2-SFP twin gig converters, the giga ports do not link up to the peer device that sends a remote fault. The **show int status | inc gi x/y** command indicates notconnect.

Similar behavior is observed with Supervisor Engine 6-E uplinks and the WS-X4706-10GE line card.

This behavior is seen in Cisco IOS Releases 12.2(50)SG thru 12.2(50)SG3 when the peer device sends a remote fault.

Workaround: Disable auto negotiation at both ends.

(CSCta02425)

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

- When using dynamic policy installation for a client or host that is authenticated on a secure port, the traffic from the client is not permitted even though the **permit ip any any** command is specified as the dynamic policy for the client.

This occurs only if the following conditions are satisfied:

- Multi-host mode configured on the port with the **authentication host-mode multi-host** command.
- Default ACL (the IP access-list) configured on the interface specifies **deny ip any any**.
- Dynamic policy authorization for the client specifies **permit ip any any**.

Workaround: None.

CSCsz63739

Resolved Caveats in Cisco IOS Release 12.2(50)SG

This section lists the resolved caveats in Release 12.2(50)SG:

- With CFM, if the VLAN associated with the service instance or MEP is allocated after the Inward Facing MEP (IFM) is configured on an interface whose status is **down**, the IFM CC status remains **inactive** in the output of the **show ethernet CFM maintenance local** command. Also, the CFM remote neighbor is not seen.

This behavior is only seen when VLAN is allocated after the IFM is configured.

Workaround: Unconfigure with the **no ethernet cfm mep level mpid vlan** command, then reconfigure the IFM with the **ethernet cfm mep level mpid vlan** command on the port after the VLAN is allocated. Verify that the C-Status of the IFM is Active with the **show ethernet cfm maintenance-points local** command. (CSCsm85460)

- Occasionally, if a PC continues to send traffic behind an 802.1X capable phone that is plugged into a port configured with MDA (Multi-Domain Authentication), MAB (MAC Authentication Bypass) and port security, a 802.1X security violation is triggered if the port observes traffic from the PC before the phone is fully authorized on the port.

Workaround: Authenticate the phone before plugging a PC behind the phone. (CSCsq92724)

- After CFM is disabled globally and then a switch is reloaded with the CFM configuration in place, and after reload when cfm is enabled globally, the cfm meps are being inactive, which results in loss of cfm neighbors.

Workarounds: Do one of the following:

- Reapply the cfm configuration; at a minimum, remove and re-add the MEPs configured on all the interfaces of the switch.
- Deallocate cfm service VLANs and reallocate them.

(CSCsq90598)

- When policer or shape or shape values are specified in terms of percentage of link bandwidth on a policy and the interface on which it is attached is forced to a specific speed with the **speed 10/100/1000** command, the applied policer or shape or shape value might correspond to the new forced speed.

Service policy has to be configured with percentage police or shape or share values and the link speed is forced to a specific values. For example

```
Policy-map p1
  class-map c1
    police rate percent 10
```

Workaround: Either use the **speed auto 10/100/1000** command or the absolute policer, shape or shape values rather than percentage values. For example,

```
Policy-map p1
  class-map c1
    police rate 10 mbps
```

(CSCsk56877)

- When the trusted boundary feature is enabled on an interface, there is no command to check the current operating state.

Workaround: None. You cannot explicitly check the trusted boundary state. However, you can indirectly determine this state:

The trusted boundary feature ensures whether the packet's COS/DSCP value will be trusted or not. When the interface is not in a trusted state, the COS/DSCP fields are forced to zero on a received packet.

A QoS policy exists on that interface that uses that COS/DSCP value for classification. Therefore, if the packet classification is based on the packet value, you can infer that the interface is in a trusted state. (CSCsh72408)

- Cisco IOS software contains a vulnerability in multiple features that could allow an attacker to cause a denial of service (DoS) condition on the affected device. A sequence of specially crafted TCP packets can cause the vulnerable device to reload.

Cisco has released free software updates that address this vulnerability.

Several mitigation strategies are outlined in the workarounds section of this advisory.

This advisory is posted at

<http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20090325-tcp>

CSCsr29468

- Symptoms: Several features within Cisco IOS software are affected by a crafted UDP packet vulnerability. If any of the affected features are enabled, a successful attack will result in a blocked input queue on the inbound interface. Only crafted UDP packets destined for the device could result in the interface being blocked, transit traffic will not block the interface.

Cisco has released free software updates that address this vulnerability.

Workarounds that mitigate this vulnerability are available in the workarounds section of the advisory. This advisory is posted at the following link:

<http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20090325-udp>

CSCsk64158

- Symptoms: SSLVPN service stops accepting any new SSLVPN connections.

Conditions: A device configured for SSLVPN may stop accepting any new SSLVPN connections, due to a vulnerability in the processing of new TCP connections for SSLVPN services. If “debug ip tcp transactions” is enabled and this vulnerability is triggered, debug messages with connection queue limit reached will be observed. This vulnerability is documented in two separate Cisco bug IDs, both of which are required for a full fix: CSCso04657 and CSCsg00102.

CSCso04657

- If a redundant switch is in SSO mode or during an ISSU upgrade/downgrade, and the standby supervisor is running IOS software release 12.2(44)SG or 12.2(46)SG, when you enter the **auto qos voip trust** command on an interface with an attached service-policy, the standby supervisor engine reboots.

Workaround: Remove all service-policies from the interface before entering the **auto qos voip trust** command.

CSCsq37471

- Multiple Cisco products are affected by denial of service (DoS) vulnerabilities that manipulate the state of Transmission Control Protocol (TCP) connections. By manipulating the state of a TCP connection, an attacker could force the TCP connection to remain in a long-lived state, possibly indefinitely. If enough TCP connections are forced into a long-lived or indefinite state, resources on a system under attack may be consumed, preventing new TCP connections from being accepted. In some cases, a system reboot may be necessary to recover normal system operation. To exploit these vulnerabilities, an attacker must be able to complete a TCP three-way handshake with a vulnerable system.

In addition to these vulnerabilities, Cisco Nexus 5000 devices contain a TCP DoS vulnerability that may result in a system crash. This additional vulnerability was found as a result of testing the TCP state manipulation vulnerabilities.

Cisco has released free software updates for download from the Cisco website that address these vulnerabilities. Workarounds that mitigate these vulnerabilities are available.

This advisory is posted at

<http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20090908-tcp24>

CSCsv04836

Open Caveats in Cisco IOS Release 12.2(46)SG

This section lists the open caveats in Cisco IOS Release 12.2(46)SG:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.

Workarounds: None.

The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)

- When policer or shape or shape values are specified in terms of percentage of link bandwidth on a policy and the interface on which it is attached is forced to a specific speed with the **speed 10/100/1000** command, the applied policer or shape or shape value might correspond to the new forced speed.

Service policy has to be configured with percentage police or shape or share values and the link speed is forced to a specific values. For example

```
Policy-map p1
  class-map c1
    police rate percent 10
```

Workaround: Either use the **speed auto 10/100/1000** command or the absolute policer, shape or shape values rather than percentage values. For example,

```
Policy-map p1
  class-map c1
    police rate 10 mbps
```

(CSCsk56877)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps `system-cpp-dhcp-cs`, `system-cpp-dhcp-sc`, and `system-cpp-dhcp-ss` may not be effective.

Workaround: None. (CSCsk67395)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- When the trusted boundary feature is enabled on an interface, there is no command to check the current operating state.

Workaround: None. You cannot explicitly check the trusted boundary state. However, you can indirectly determine this state:

The trusted boundary feature ensures whether the packet's COS/DSCP value will be trusted or not. When the interface is not in a trusted state, the COS/DSCP fields are forced to zero on a received packet.

A QoS policy exists on that interface that uses that COS/DSCP value for classification. Therefore, if the packet classification is based on the packet value, you can infer that the interface is in a trusted state. (CSCsh72408)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submode. Then, apply the new class-map with the updated changes.

CSCsk70826)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsi71036)

- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.

Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)

- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.

Workaround: None. (CSCs139767)

- IGMP snooping entries are active even after disabling IGMP snooping globally and per VLAN.

Workarounds: Disable IGMP snooping on all the relevant VLANs before disabling it globally.

- IPv6 MLD entries are active even if an IPv6 MLD related configuration does not exist.

Workaround: Unconfigure all generic QoS policies from the system. (CSCsq84853)

- IPv6 entries are active in the CAM; the CPU receives IPv6 packets.

Workaround: Unconfigure any generic QoS policies from the system. The QoS policies with the **match any** attribute cause IPv6 entries to become active. If the switch is a pure Layer 2 device, remove the generic protocol family attributes and narrow it to the protocol family.

(CSCsq84796)

- Initially, REP configured with VLAN Load Balancing (VLB) works correctly. When you issue a force-switchover on the switch, that has a port acting as the secondary ALT port, a loop is induced in the topology.

Workaround: Enter shut, then no-shut on any REP port (of the same segment in which VLB is configured) in the topology. (CSCsq75342)

- In Cisco IOS Release 12.2(46)SG, if flexlink is applied to a pair of etherchannels, then flexlink config may not be applied after a reboot, if the backup EtherChannel is defined after the flexlink configuration.

Workaround: Define the backup etherchannel before applying flexlink command. (CSCsq13477)

- In Cisco IOS Release 12.2(46)SG, if an etherchannel is a member of a flexlink pair, then static MAC addresses configured on the EtherChannel are not moved to the alternate port when the EtherChannel fails (flexlink failure)

Workaround: None. (CSCsq99468)

- Performing a default interface operation on an interface with auto-QoS enabled results in an error message and the loss of the auto-QoS configuration. For example, the following sequence of operation results in a loss of the configuration:

```
config-if# auto qos voip cisco-phone
config# default interface interface-name
```

Workaround: Replace the **default interface** command with the following:

```
config# interface interface-number
config-if# switchport
```

(CSCsq47116)

- The IPv6 ICMP neighbor state changes from **REACH** to **STALE** after 15 secs of inactivity on the link.

Workaround: Ping the global and link local addresses of the neighbor to ascertain and reinstate reachability. (CSCsq77181)

- IPv6 EIGRP routes are not learned through the port channel.

Workaround: Unconfigure the port channel and the associated physical port, and reconfigure them. (CSCsq74229)

- When a CFM Inward Facing MEP(IFM) is configured on a VLAN that is not allocated on a switch port that is DOWN, the **show ethernet cfm maintenance-points local** command displays the IFM CC Status as **Inactive**. Then, you allocate the VLAN, the CC-status remains **Inactive**.

You only see this symptom if you did not allocate a VLAN before you configure the IFM, then at a later time allocate the same VLAN.

Workaround: Unconfigure, then reconfigure the IFM on the port.

- With CFM, if the VLAN associated with the service instance or MEP is allocated after the Inward Facing MEP (IFM) is configured on an interface whose status is **down**, the IFM CC status remains **inactive** in the output of the **show ethernet CFM maintenance local** command. Also, the CFM remote neighbor is not seen.

This behavior is only seen when VLAN is allocated after the IFM is configured.

Workaround: Unconfigure with the **no ethernet cfm mep level mpid vlan** command, then reconfigure the IFM with the **ethernet cfm mep level mpid vlan** command on the port after the VLAN is allocated. Verify that the C-Status of the IFM is Active with the **show ethernet cfm maintenance-points local** command. (CSCsm85460)

- Occasionally, if a PC continues to send traffic behind an 802.1X capable phone that is plugged into a port configured with MDA (Multi-Domain Authentication), MAB (MAC Authentication Bypass) and port security, a 802.1X security violation is triggered if the port observes traffic from the PC before the phone is fully authorized on the port.

Workaround: Authenticate the phone before plugging a PC behind the phone. (CSCsq92724)

- Ordinarily, the output of a CFM Traceroute from a MEP normally lists down the next hop name(device/host name) for each hop till the other MEP. When CFM over EtherChannel exists between the two MEPS, CFM Traceroute issued from a MEP does not show the next hop name.

Workaround: None. (CSCso50659)

- An IP unnumbered configuration is lost after a reload.

Workarounds: Do one of the following:

- After a reload, copy the startup-config to the running-config.
- Use a loopback interface as the target of the **ip unnumbered** command
- Change the CLI configuration such that during bootstrap, the router port is created first.

(CSCsq63051)

- After CFM is disabled globally and then a switch is reloaded with the CFM configuration in place, and after reload when cfm is enabled globally, the cfm meps are being inactive, which results in loss of cfm neighbors.

Workarounds: Do one of the following:

- Reapply the cfm configuration; at a minimum, remove and re-add the MEPs configured on all the interfaces of the switch.
- Deallocate cfm service VLANs and reallocate them.

(CSCsq90598)

- In SSO mode, when a port-channel is created, deleted, and re-created on an active supervisor engine with the same channel-number, the standby port-channel state goes out of sync. After a switch over, the following message displays:

```
%PM-4-PORT_INCONSISTENT: STANDBY:Port is inconsistent:
```

Workaround: When the port channel starts to flap, enter **shut** and **no shut** on the port channel. After the first switchover and after deleting the portchannel, create a new channel. (CSCsr00333)

- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.

The following conditions may cause a RACL to malfunction:

- ACL are applied on the output direction of the interface.
- IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).

Here are two examples of such non-functioning RACL:

```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

```
IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```

Workaround: None.

CSCtc13297

Resolved Caveats in Cisco IOS Release 12.2(46)SG

This section lists the resolved caveats in Release 12.2(46)SG:

- When a service-policy is removed from a physical port that is member of an ether channel, a LACP or PAGP protocol-based ether channel goes down. The port-channel members get bundled back in but remain in *suspended* state due to failure to exchange the protocol packets with the other end.

Workarounds: Before removing the service policy from a ether channel member, remove it from the channel. Then, return it to the channel. (CSCsk70568)

- When using *bandwidth percentage* actions in a queuing policy-map, the actual bandwidth share differs from that of the configured policy-map.

In a queuing QoS policy, there can be zero or more queuing classes that have an explicit, user specified, bandwidth share specified. There can be zero or more queuing classes that do not have such user specified bandwidth share. The system takes the unallocated bandwidth share and allocates it equally among the latter set of classes.

When using percentage-based bandwidth allocation, if the share comes to less than 1%, the queues corresponding to those classes do not get updated in hardware with the new bandwidth share. These queues get more than the expected share of bandwidth.

Workarounds: Ensure that the unallocated bandwidth percentage is at least equal to the number of queues that do not have the explicit **bandwidth percentage** command. This should include the default as well as priority queues. (CSCsk77757)

- Not all combinations of features can be simultaneously supported by the hardware. When such a feature combination is configured, packets will be processed in software and a log message indicating this will be generated:

```
%C4K_HWACLMAN-4-ACLHWLABELERR: Path (in :50, 1006) label allocation failure:
SignatureInconsistent - packets will be handled in software, QoS is disabled.
```

One feature combination that can trigger this problem is the attempt to combine a QoS policy that matches on cos bits with IPv6 ACL configuration that matches on IPv6 source addresses that partially mask in the lower 48 bits of the address. (IPv6 subnets in the /81 to /127 range will also trigger this behavior if IPv6 multicast routing is enabled.)

Workaround: Do not configure feature combinations that conflict. Currently the above conflict between QoS policies matching on COS bits and IPv6 configuration with partial masking of the lower 48 bits of the source address is the only known conflicting feature combination. If matching on COS bits is required by the QoS policy, architect the IPv6 network using /80 subnets or larger. (CSCsk79791)

- When a service policy on a port-channel member port is modified, traffic may be dropped for some of the classes.

Workaround: Do the following:

1. Un-configure the interface(s) on which this policy-map is attached from the portchannel.
2. Modify the policy-map.
3. Configure the interface(s) in the portchannel.

(CSCsk77119)

- When two switches are connected back-to-back via two or more links and when a packet is locally-originated, the source IP address may not correspond to the IP address of the outgoing interface. A switch receiving such a packet with unicast RPF feature enabled might drop the incoming packet.

Workaround: None. (CSCsh99124)

- In policy map, if a queuing class with the **bandwidth remaining percent** <> command sits before a priority queuing class configuration, the **bandwidth remaining percent** <> command action is applied on reload.

Workaround: Re-apply the policy-map. (CSCsk75793)

- A port can be either a member of a portchannel or have auto-QoS applied to it, but not both. The two are mutually exclusive features.

Currently, if it is applied to a port that is already a member of a portchannel, the application is rejected with an error message. However, the reverse is not true. If auto-QoS is applied first and then the port joins a portchannel, the command is accepted.

The following example using port g2/1 shows the type of usage that should be avoided:

```
conf t
int g2/1
auto qos voice trust
channel-group 10 mode auto
```

This example applies auto-QoS on a port (g2/1) and subsequently makes the port a member of portchannel (10).

Workaround: Do not make a port with auto-QoS enabled a member of a portchannel. (CSCsi95018)

- If *exceed burst* is not explicitly configured for a dual rate policer, the **show policy-map** command displays “0” as the burst value.
Workaround: Enter the **show policy-map interface** command. (CSCsj44237)
- When a queuing policy is attached to a trunk port configured with a per-port per-VLAN QoS policy, the port-level queuing policy is processed as part of a per-VLAN policy and is rejected on bootup. Queuing policy is supported on a physical interface in the output direction only.
Workaround: After bootup, reattach a queuing policy on a physical interface. (CSCsk87548)
- When you delete a port-channel with a per-port per-VLAN QoS policy, the switch crashes.
Workaround: Before deleting the port-channel, do the following:
 1. Remove any per-port per-VLAN QoS policies, if any.
 2. Remove the VLAN configuration on the port-channel with the **no vlan-range** command. (CSCsk91916)
- The cbQosPoliceCfgTable mib object is *not* populated by the **police bps byte** command.
Workaround: None. (CSCsk45940)
- On rare occasions, a Catalyst 4900M switch may undergo restart if ARP requests are sent to all ports on the switch and “debug ip arp” is enabled.
Workaround: None. (CSCs126706)
- Storm control may not work as expected on Tengig ports 1/1 and 1/3.
Workaround: None. (CSCs137599)

Open Caveats in Cisco IOS Release 12.2(40)XO

This section lists the open caveats in Cisco IOS Release 12.2(40)XO:

- Software qos does not match a .1Q packet properly for applying the desired qos actions.
Workarounds: None.
The support to handle .1Q packets for software QoS lookup unavailable in the Cisco IOS Release 12.2(40)SG release. (CSCsk66449)
- When a service-policy is removed from a physical port that is member of an ether channel, a LACP or PAGP protocol-based ether channel goes down. The port-channel members get bundled back in but remain in *suspended* state due to failure to exchange the protocol packets with the other end.
Workarounds: Before removing the service policy from a ether channel member, remove it from the channel. Then, return it to the channel. (CSCsk70568)
- When using *bandwidth percentage* actions in a queuing policy-map, the actual bandwidth share differs from that of the configured policy-map.
In a queuing QoS policy, there can be zero or more queuing classes that have an explicit, user specified, bandwidth share specified. There can be zero or more queuing classes that do not have such user specified bandwidth share. The system takes the unallocated bandwidth share and allocates it equally among the latter set of classes.
When using percentage-based bandwidth allocation, if the share comes to less than 1%, the queues corresponding to those classes do not get updated in hardware with the new bandwidth share. These queues get more than the expected share of bandwidth.

Workarounds: Ensure that the unallocated bandwidth percentage is at least equal to the number of queues that do not have the explicit **bandwidth percentage** command. This should include the default as well as priority queues. (CSCsk77757)

- Not all combinations of features can be simultaneously supported by the hardware. When such a feature combination is configured, packets will be processed in software and a log message indicating this will be generated:

```
%C4K_HWACLMAN-4-ACLHWLABELERR: Path (in :50, 1006) label allocation failure:
SignatureInconsistent - packets will be handled in software, QoS is disabled.
```

One feature combination that can trigger this problem is the attempt to combine a QoS policy that matches on cos bits with IPv6 ACL configuration that matches on IPv6 source addresses that partially mask in the lower 48 bits of the address. (IPv6 subnets in the /81 to /127 range will also trigger this behavior if IPv6 multicast routing is enabled.)

Workaround: Do not configure feature combinations that conflict. Currently the above conflict between QoS policies matching on COS bits and IPv6 configuration with partial masking of the lower 48 bits of the source address is the only known conflicting feature combination. If matching on COS bits is required by the QoS policy, architect the IPv6 network using /80 subnets or larger. (CSCsk79791)

- When policer or shape or shape values are specified in terms of percentage of link bandwidth on a policy and the interface on which it is attached is forced to a specific speed with the **speed 10/100/1000** command, the applied policer or shape or shape value might correspond to the new forced speed.

Service policy has to be configured with percentage police or shape or share values and the link speed is forced to a specific values. For example

```
Policy-map p1
  class-map c1
    police rate percent 10
```

Workaround: Either use the **speed auto 10/100/1000** command or the absolute policer, shape or shape values rather than percentage values. For example,

```
Policy-map p1
  class-map c1
    police rate 10 mbps
```

(CSCsk56877)

- Under some conditions, one or more flows continue to be dropped because of DBL even after DBL has been removed from the service-policy.

When an output service-policy is attached to an interface and if the policy is configured to apply DBL on a queue, the flows that are enqueued to the queue are subjected to the DBL algorithm. If one or more flows are classified as *belligerent* (flows do not back-off in response to drops because of congestion in the queue), those flows continue to be classified as belligerent even when DBL is disabled on that queue.

For this condition to persist, the transmit queues in question must remain congested for a long period of time and that congestion must be caused by flows that remain belligerent.

Workaround: Provided the queue in question is non-default (queuing actions are not configured in the class-default class of the policy-map), detach and re-attach the service-policy.

If this happens on the default queue, modifying and resetting some queuing parameters like bandwidth/shape fixes the issue. (CSCsk62457)

- When an Catalyst 4900M switch encounters either a fan tray failure or a supervisor critical temperature, the chassis shuts off. The output of the **show crashdump** command will *not* indicate the cause of the power-down.

Workarounds: Use the **show log** command to determine the cause of the power-down.

- If the log has *LogGalInsufficientFansDetected* messages, the cause was a fan-tray failure.
- If the log has *LogRkiosModuleShutdownTemp* messages, the cause was that the supervisor critical temperature exceeded the failure threshold.

(CSCsk48632)

- When a service policy on a port-channel member port is modified, traffic may be dropped for some of the classes.

Workaround: Do the following:

1. Un-configure the interface(s) on which this policy-map is attached from the portchannel.
2. Modify the policy-map.
3. Configure the interface(s) in the portchannel.

(CSCsk77119)

- When two switches are connected back-to-back via two or more links and when a packet is locally-originated, the source IP address may not correspond to the IP address of the outgoing interface. A switch receiving such a packet with unicast RPF feature enabled might drop the incoming packet.

Workaround: None. (CSCsh99124)

- A Catalyst 4900M switch will support a maximum of 32 MTU values system wide.

On a Catalyst 4900M running Cisco IOS Release 12.2(40)SG, all MTU values configured on a line card are set to default when the module is reset. Furthermore, MTU values are not retained for modules that are physically moved.

Workaround: None. (CSCsk52542)

Workaround: Reinsert the X2. (CSCsk43618)

- On rare occasions, if you use an X2 SR transceiver on a WS-X4706-10GE running Cisco IOS Release 12.2(40)SG, you will observe CRC errors after a reload or power cycle when you insert the card or the X2.

Workaround: Reinsert the X2. (CSCsk43618)

- Control plane policing applied to DHCP traffic as identified by the system class-maps system-cpp-dhcp-cs, system-cpp-dhcp-sc, and system-cpp-dhcp-ss may not be effective.

Workaround: None. CSCsk67395)

- In policy map, if a queuing class with the **bandwidth remaining percent <>** command sits before a priority queuing class configuration, the **bandwidth remaining percent <>** command action is applied on reload.

Workaround: Re-apply the policy-map. (CSCsk75793)

- When the CPU transmits .1X packet on an interface that has an egress qos policy attached, the packet is not matched and exits without any QoS marking actions.

When a packet is sent to the CPU it may get sent out on some other interface. If so, the original COS value for a .1X packet cannot be matched by software QoS (as per CSCsk66449). The packet is transmitted with whatever COS value it was generated with (7, for the MLDv1 packets described here).

Workaround: None.

Part of the root cause of this problem is captured through CSCsk66449, which indicates that the software QoS cannot match against a .1X packet. (CSCsk72544)

- When the trusted boundary feature is enabled on an interface, there is no command to check the current operating state.

Workaround: None. You cannot explicitly check the trusted boundary state. However, you can indirectly determine this state:

The trusted boundary feature ensures whether the packet's COS/DSCP value will be trusted or not. When the interface is not in a trusted state, the COS/DSCP fields are forced to zero on a received packet.

A QoS policy exists on that interface that uses that COS/DSCP value for classification. Therefore, if the packet classification is based on the packet value, you can infer that the interface is in a trusted state. (CSCsh72408)

- A port can be either a member of a portchannel or have auto-QoS applied to it, but not both. The two are mutually exclusive features.

Currently, if is applied to a port that is already a member of a portchannel, the application is rejected with an error message. However, the reverse is not true. If auto-QoS is applied first and then the port joins a portchannel, the command is accepted.

The following example using port g2/1 shows the type of usage that should be avoided:

```
conf t
int g2/1
auto qos voice trust
channel-group 10 mode auto
```

This example applies auto-QoS on a port (g2/1) and subsequently makes the port a member of portchannel (10).

Workaround: Do not make a port with auto-QoS enabled a member of a portchannel. (CSCsi95018)

- If a class-map is configured with **exceed-action drop**, re-configuring the same class-map with **exceed-action transmit** causes class-map configurations to conflict for the same class-map.

Workaround: If you plan to change a class-map action, such as **exceed-action**, you need to remove the class-map with the **no class c1** command under policy-map submenu. Then, apply the new class-map with the updated changes.

CSCsk70826)

- Policing actions are not applied if they appear at the child level of a two-level hierarchical policy-map.

The switch supports two-level hierarchical policy-maps. Policing actions can be present at only one of the two levels (parent or child). If they are present at the child level, they are not applied.

Workaround: None. (CSCsl0631)

- Applying a policy to a VLAN that has been allocated to a routed port causes the internal VLAN to be policed.

Workaround: Avoid creating a VLAN that has been allocated internally to a routed port. (CSCsh60244)

- If *exceed burst* is not explicitly configured for a dual rate policer, the **show policy-map** command displays "0" as the burst value.

Workaround: Enter the **show policy-map interface** command. (CSCsj44237)

- If *burst* is not explicitly configured for a single rate policer, the **show policy-map command** displays an incorrect burst value.
Workaround: Enter the **show policy-map interface** command. (CSCsi71036)
- When a queuing policy is attached to a trunk port configured with a per-port per-VLAN QoS policy, the port-level queuing policy is processed as part of a per-VLAN policy and is rejected on bootup. Queuing policy is supported on a physical interface in the output direction only.
Workaround: After bootup, reattach a queuing policy on a physical interface. (CSCsk87548)
- When you delete a port-channel with a per-port per-VLAN QoS policy, the switch crashes.
Workaround: Before deleting the port-channel, do the following:
 1. Remove any per-port per-VLAN QoS policies, if any.
 2. Remove the VLAN configuration on the port-channel with the **no vlan-range** command. (CSCsk91916)
- The `cbQosPoliceCfgTable` mib object is *not* populated by the **police bps byte** command.
Workaround: None. (CSCsk45940)
- When you enter the **show policy-map vlan *vlan*** command, unconditional marking actions that are configured on the VLAN are not shown.
Workaround: None. However, if you enter the **show policy-map *name***, the unconditional marking actions are displayed. (CSCsi94144)
- You observe a .05% loss on WS-X4908-10GE when sending traffic at 99% of the port capacity.
Workaround: None. (CSCsI39767)
- On rare occasions, a Catalyst 4900M switch may undergo restart if ARP requests are sent to all ports on the switch and “debug ip arp” is enabled.
Workaround: None. (CSCsI26706)
- Storm control may not work as expected on Tengig ports 1/1 and 1/3.
Workaround: None. (CSCsI37599)
- Output IPv6 ACLs with Ace to match on the ICMP option fail on a switch.
The following conditions may cause a RAACL to malfunction:
 - ACL are applied on the output direction of the interface.
 - IPv6 ACL contain Ace to match on the ICMP option fields (ICMP Type or ICMP Code).
 Here are two examples of such non-functioning RAACL:


```
IPv6 access list a1
  permit icmp any any nd-ns sequence 10
  deny ipv6 any any sequence 20

IPv6 access list a2
  permit icmp 2020::/96 any nd-ns sequence 10
  deny ipv6 any any sequence 20
```


Workaround: None.
CSCtc13297

Resolved Caveats in Cisco IOS Release 12.2(40)XO

This section lists the resolved caveats in Release 12.2(40)XO:

- None

Troubleshooting

These sections provide troubleshooting guidelines for the Catalyst 4900M series switch running IOS supervisor engines:

- [Netbooting from the ROMMON](#), page 143
- [Troubleshooting at the System Level](#), page 144
- [Troubleshooting Modules](#), page 144
- [Troubleshooting MIBs](#), page 144

Netbooting from the ROMMON

Netbooting using a boot loader image is not supported. Instead, use one of the following options to boot an image:

1. Boot from a CompactFlash card by entering the following command:

```
rommon 1> boot slot0:<bootable_image>
```

2. Use ROMMON TFTP boot.

The ROMMON TFTP boot is very similar to the BOOTLDR TFTP boot, except that:

- the BOOTLDR variable should *not* be set
- the TFTP server must be accessible from the Ethernet management port on the supervisor engine.

To boot from ROMMON, perform the following tasks while in ROMMON mode:

- a. Ensure that the Ethernet management port is physically connected to the network.
- b. Verify that bootloader environment is not set by entering the **unset bootldr** command.
- c. Set IP address of the Ethernet management port on the supervisor engine by entering the following command: **set interface fa1 ip_address <ip_mask**

For example, to set the supervisor engine Ethernet port with an IP address 172.16.1.5 and IP mask 255.255.255.0, enter the following command:

```
rommon 2> set interface fa1 172.16.1.5 255.255.255.0
```

- d. Set default gateway for the Ethernet management port on the supervisor engine by entering the following command: **set ip route default gateway_ip_address**. The default gateway should be directly connected to the supervisor engine Ethernet management port subnet.
- e. Ping the TFTP server to ensure that there is connectivity to the server from the Ethernet management port on the supervisor engine by entering the following command: **ping <tftp_server_ip_address>**.

- f. Once the ping is successful, boot the image from the TFTP server by entering the following command: **boot tftp://tftp_server_ip_address/<image_path_and_file_name**

For example, to boot the image name **cat4500-ipbase-mz** located on the TFTP server 172.16.1.8, enter the following command:

```
rommon 3> boot tftp://172.16.1.8/tftpboot/cat4500-ipbase-mz
```

Troubleshooting at the System Level

This section contains troubleshooting guidelines for system-level problems:

- When the system is booting and running power-on diagnostics, do not reset the switch.
- Ensure that you do not mix the serial and Ethernet cables. The Ethernet Management port is inoperative. An Ethernet cable plugged into the Ethernet port is active only in ROMMON mode.

Troubleshooting Modules

This section contains troubleshooting guidelines for the Catalyst 4900M series switch:

- Whenever you connect an interface that has duplex set to autonegotiate to an end station or another networking device, ensure that the other device is configured for autonegotiation as well. If the other device is not set to autonegotiate, the port set to autonegotiate will remain in half-duplex mode, which can cause a duplex mismatch resulting in packet loss, late collisions, and line errors on the link.

Troubleshooting MIBs

For general information on MIBs, RMON groups, and traps, refer to the Cisco public MIB directory (<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>). For information on the specific MIBs supported by the Catalyst 4900M series switches, refer to the Catalyst 4000 MIB Support List located at <ftp://ftp.cisco.com/pub/mibs/supportlists/cat4000/cat4000-supportlist.html>.

Related Documentation

Although their Release Notes are unique, the 4 platforms (Catalyst 4500, Catalyst 4900, Catalyst ME 4900, and Catalyst 4900M) use the same *Software Configuration Guide*, *Command Reference Guide*, and *System Message Guide*. Refer to the following home pages for additional information:

- Catalyst 4500 Series Switch Documentation Home
<http://www.cisco.com/go/cat4500/docs>
- Catalyst 4900 Series Switch Documentation Home
<http://www.cisco.com/go/cat4900/docs>
- Cisco ME 4900 Series Ethernet Switches Documentation Home
http://www.cisco.com/en/US/products/ps7009/tsd_products_support_series_home.html

Hardware Documents

Installation guides and notes including specifications and relevant safety information are available at the following URLs:

- *Catalyst 4500 Series Switches Installation Guide*
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/hardware/installation/guide/78-14409-08/4500inst.html>
- *Catalyst 4500 E-series Switches Installation Guide*
<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/hardware/catalyst4500e/installation/guide/Eseries.html>
- For information about individual switching modules and supervisors, refer to the *Catalyst 4500 Series Module Installation Guide* at:
http://www.cisco.com/en/US/products/hw/switches/ps4324/prod_installation_guides_list.html
- *Regulatory Compliance and Safety Information for the Catalyst 4500 Series Switches*
http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/hardware/regulatory/compliance/78_13233.html
- Installation notes for specific supervisor engines or for accessory hardware are available at:
http://www.cisco.com/en/US/products/hw/switches/ps4324/prod_installation_guides_list.html
- Catalyst 4900 and 4900M hardware installation information is available at:
http://www.cisco.com/en/US/products/ps6021/prod_installation_guides_list.html
- Cisco ME 4900 Series Ethernet Switches installation information is available at:
http://www.cisco.com/en/US/products/ps7009/prod_installation_guides_list.html

Software Documentation

Software release notes, configuration guides, command references, and system message guides are available at the following URLs:

- Catalyst 4500 release notes are available at:
http://www.cisco.com/en/US/products/hw/switches/ps4324/prod_release_notes_list.html
- Catalyst 4900 release notes are available at:
http://www.cisco.com/en/US/products/ps6021/prod_release_notes_list.html
- Cisco ME4900 4900 Series Ethernet Switch release notes are available at:
http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/release/note/OL_11511.html

Software documents for the Catalyst 4500 Classic, Catalyst 4500 E-Series, Catalyst 4900, and Cisco ME 4900 Series Ethernet Switches are available at the following URLs:

- *Catalyst 4500 Series Software Configuration Guide*
http://www.cisco.com/en/US/products/hw/switches/ps4324/products_installation_and_configuration_guides_list.html
- *Catalyst 4500 Series Software Command Reference*
http://www.cisco.com/en/US/products/hw/switches/ps4324/prod_command_reference_list.html

- *Catalyst 4500 Series Software System Message Guide*
http://www.cisco.com/en/US/products/hw/switches/ps4324/products_system_message_guides_list.html

Cisco IOS Documentation

Platform-independent Cisco IOS documentation may also apply to the Catalyst 4500 and 4900 switches. These documents are available at the following URLs:

- Cisco IOS configuration guides, Release 12.x
http://www.cisco.com/en/US/products/ps6350/products_installation_and_configuration_guides_list.html
- Cisco IOS command references, Release 12.x
http://www.cisco.com/en/US/products/ps6350/prod_command_reference_list.html
You can also use the Command Lookup Tool at:
<http://tools.cisco.com/Support/CLILookup/clkSearchAction.do>
- Cisco IOS system messages, version 12.x
http://www.cisco.com/en/US/products/ps6350/products_system_message_guides_list.html
You can also use the Error Message Decoder tool at:
<http://www.cisco.com/cgi-bin/Support/Errordecoder/index.cgi>
- For information about MIBs, refer to:
<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

Notices

The following notices pertain to this software license.

OpenSSL/Open SSL Project

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

License Issues

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License:

Copyright © 1998-2007 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: “This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)”.
4. The names “OpenSSL Toolkit” and “OpenSSL Project” must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called “OpenSSL” nor may “OpenSSL” appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:
 “This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)”.

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT “AS IS” AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (ey@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License:

Copyright © 1995-1998 Eric Young (ey@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (ey@cryptsoft.com).

The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young’s, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

“This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)”.

The word ‘cryptographic’ can be left out if the routines from the library being used are not cryptography-related.

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: “This product includes software written by Tim Hudson (tjh@cryptsoft.com)”.

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License].

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What’s New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What’s New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS Version 2.0.

This document is to be used in conjunction with the documents listed in the “[Related Documentation](#)” section.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

*Release Notes for the Catalyst 4900 Series Switch, Cisco IOS Release 12.2(53)SG
Copyright © 2008-2009, Cisco Systems, Inc. All rights reserved.*

