Release Notes for Catalyst 2948G-L3 and Catalyst 4908G-L3 for Cisco IOS Release 12.0(25)W5(27)

March 21, 2005

Current Release:
12.0(25)W5(27c)

Previous release: 12.0(25)W5(27b), 12.0(25)W5(27a), 12.0(25)W5(27), 12.0(18)W5(22b), 12.0(18)W5(22a), 12.0(14)W5(20), 12.0(10)W5(18e), 12.0(7)W5(15d)

This publication describes the current Catalyst 2948G-L3 and Catalyst 4908G-L3 switches software features and caveats for Cisco IOS Release12.0(25)W5(27c).

Contents

This publication contains the following sections:

- Introduction, page2
- Version and Part Number, page7
- System Requirements, page2
- New Features and Changed Information, page7
- Limitations and Restrictions, page10
- Caveats, page11
- Error Messages, page23
- Related Documentation, page24
- Service and Support, page24
- Obtaining Documentation, page25
- Obtaining Technical Assistance, page26

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

Introduction

The Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers are high-performance Layer 3 switch routers that share the same software image. The Catalyst 2948G-L3 is a multiprotocol 10/100/1000 Ethernet switch router. The Catalyst 4908G-L3 is a multiprotocol Gigabit Ethernet switch router.

A Layer 3 switch router performs the following three major functions:

- Packet switching
- Route processing
- Intelligent network services

Compared to other routers, Layer 3 switch routers process more packets faster by using application-specific integrated circuit (ASIC) hardware instead of microprocessor-based engines. Layer 3 switch routers also improve network performance with two software functions, route processing and intelligent network services.

System Requirements

This section describes the system requirements for Release 12.0(25)W5(27c) and includes the following sections:

- Memory Defaults, page3
- Hardware Supported, page3
- Features Supported, page3
- Features Not Supported, page6
- Determining the Software Version, page6
- Version and Part Number, page7

Memory Defaults

Table1lists the default Flash and DRAM memory defaults for the Catalyst 2948G-L3 and Catalyst4908G-L3 switch routers.

Table1 Default Memory by Platform

Layer 3 Switch Router	Flash Memory	DRAM
Catalyst 2948G-L3	16 MB	64 MB
Catalyst 4908G-L3	16 MB	64 MB

Hardware Supported

Table2lists the interfaces that the Catalyst 2948G-L3 and the Catalyst 4908G-L3 switch routers support.

Table2 Interfaces Supported by Platform

Layer 3 Switch Routers	Interface Types	No. of Ports
Catalyst 2948G-L3	10/100 Mbps Fast Ethernet—UTP	48
	1 Gbps Gigabit Ethernet	2
Catalyst 4908G-L3	1 Gbps Gigabit Ethernet	8

Features Supported

Table3lists the software features of the Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers.

Table3 Feature Set for the Catalyst 2948G-L3 and Catalyst 4908G-L3 Switch Routers

Layer 1 Features
10/100BASE-TX half duplex and full duplex (Catalyst 2948G-L3 only)
1000BASE-SX,-LX, and long haul (-LX/LH, -ZX) full duplex
Layer 2 Bridging Features
Layer 2 transparent bridging
Layer 2 MAC[1] learning, aging, and switching by hardware
Spanning Tree Protocol (IEEE 802.1D) per bridge group
Maximum of 16 active bridge groups supported
Up to 4000 MAC addresses
Integrated routing and bridging (IRB)
VLAN[2] features
ISL[3]-based VLAN trunking
IEEE 802.1Q-based VLAN trunking

Table3 Feature Set for the Catalyst 2948G-L3 and Catalyst 4908G-L3 Switch Routers (continued)

Layer 3 Routing, Switching, and Forwarding
IP, IPX, and IP multicast routing and switching between Ethernet ports
Constrained multicast flooding (CMF)
QoS-based forwarding based on IP precedence
Load balancing among equal cost paths, based on source and destination IP and IPX[4] addresses
CEF[5] load balancing on Gigabit Ethernet ports using tunnel and universal load balancing algorithms
Layer 2 entries, IP routing, IP multicast routing, and Novell IPX routing share the 2 4 KB CAM[6] on the Catalyst2948G-L3 switch router and the 32 KB CAM on the Catalyst 4908G-L3switch router
Up to 18,000 IP routes
Up to 20,000 IP host entries
Up to 20,000 IPX routes
Up to 20,000 IPX host entries
Up to 12,000 IP multicast route entries
Access Control Lists (Gigabit Ethernet ports)
RADIUS[7] server support
IP uplink redirect
Supported Routing Protocols
RIP[8] and RIP II
IGRP[9]
EIGRP[10]
OSPF[11]
IPX RIP and EIGRP
PIM[12]—sparse and dense mode
Secondary addressing
Static routes
Classless interdomain routing (CIDR)
Local Proxy ARP[13]
BGP (Border Gateway Protocol)
Bundling of up to four FEC ports
Fast EtherChannel (FEC) Features (Catalyst 2948G-L3 only)
Load balancing among equal cost paths, based on source and destination IP and IPX[14] addresses
CEF[15] load balancing on Gigabit Ethernet ports using tunnel and universal load balancing algorithms
Load sharing for bridged traffic based on MAC address
ISL on the FEC

Table3 Feature Set for the Catalyst 2948G-L3 and Catalyst 4908G-L3 Switch Routers (continued)

Fast EtherChannel (FEC) Features (Catalyst 2948G-L3 only) (continued)
IRB on the FEC
IEEE 802.1Q trunking on the FEC
Up to 16 active FEC port channels
Note The Catalyst 4908G-L3 switch router does not have Fast Ethernet interfaces, which can be assigned to an EtherChannel.
Gigabit EtherChannel (GEC) Features
Bundling of the two Gigabit Ethernet ports on the Catalyst 2948G-L3 switch router and up to four active GEC port channels on the Catalyst 4908G-L3 switch router
Load balancing among equal cost paths, based on source and destination IP and IPX[16] addresses
CEF[17] load balancing on Gigabit Ethernet ports using tunnel and universal load balancing algorithms
Load sharing for bridge traffic based on MAC address
ISL suppported on the external GEC
IRB on the GEC
IEEE 802.1Q trunking on the GEC
One active GEC[18] port channel on the Catalyst 2948G-L3 switch router and up to four active GEC port channels on the Catalyst 4908G-L3 switch router
Additional Protocols and Features
Bootstrap Protocol (BOOTP)
Cisco Discovery Protocol (CDP) support on Ethernet ports
Cisco Group Management Protocol (CGMP) server support
DHCP[19] relay
HSRP[20] over 10/100 Ethernet, Gigabit Ethernet, FEC, GEC, and Bridge Virtual Interface(BVI)
ICMP[21]
IGMP[22]
SAP and IPX SAP[23] filtering
IRB[24] routing mode support
SNMP[25]
Per-port QoS[26] rate-limiting and shaping

1. MAC = Media Access Control
2. VLAN = Virtual LAN
3. ISL = Inter-Switch Link
4. IPX = Internet Packet Exchange
5. CEF = Cisco Express Forwarding
6. CAM = content addressable memory
7. RADIUS=Remote Authentication Dial-in User Service
8. RIP = Routing Information Protocol
9. IGRP = Interior Gateway Routing Protocol

10. EIGRP = Enhanced Interior Gateway Routing Protocol
11. OSPF = Open Shortest Path First
12. PIM = Protocol Independent Multicast
13. ARP = Address Resolution Protocol
14. IPX = Internet Packet Exchange
15. CEF = Cisco Express Forwarding
16. IPX = Internet Packet Exchange
17. CEF = Cisco Express Forwarding
18. GEC = Gigabit EtherChannel
19. DHCP = Dynamic Host Configuration Protocol
20. HSRP = Hot Standby Router Protocol
21. ICMP = Internet Control Message Protocol
22. IGMP = Internet Group Management Protocol
23. IPX SAP = Internet Packet Exchange Service Advertisement Protocol
24. IRB = Integrated Routing and Bridging Protocol
25. SNMP = Simple Network Management Protocol
26. QoS = Quality of Service

Features Not Supported

Table4lists some of the features not supported on the Catalyst2948G-L3 and the Catalyst 4908G-L3 switch routers.

Table4 Features Not Supported on the Catalyst 2948G-L3 and the Catalyst 4908G-L3 Switch Routers

Features Not Supported
Layer 2 source MAC address filtering with standard Access Control List (ACL)
User Datagram Protocol (UDP) turbo flooding
Port-based snooping (SPAN)
DEC spanning tree
Per packet load balancing
AppleTalk 1 and 2 routing
AppleTalk Routing Table Maintenance Protocol (RTMP)
AppleTalk Update-based Routing Protocol (AURP)
CGMP over BVI
Policy Base Routing (PBR)
Generic Routing Encapsulation (GRE)
PIM over tunnel interfaces

Determining the Software Version

The Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers share the same software version. To determine the version of the Cisco IOS software currently running on your switch router, log in to the switch router and enter the **show version** EXEC command.

Version and Part Number

Table5lists the features and license numbers for each platform.

Table5 Features and License Numbers by Platform

Platform	Features Included	License Number
Catalyst 2948G-L3	OSPF, IGRP, EIGRP	FR2948GL3-IP
Catalyst 2948G-L3	IPX	FR2948GL3-IPX
Catalyst 4908G-L3	IPX	FR4908GL3-IPX

New Features and Changed Information

This section lists the new features available in each release.

Features in Release 12.0(25)W5(27c)

There are no new features in Cisco IOS Release 12.0(25)W5(27c).

Features in Release 12.0(25)W5(27b)

There are no new features in Cisco IOS Release 12.0(25)W5(27b).

Features in Release 12.0(25)W5(27a)

There are no new features in Cisco IOS Release 12.0(25)W5(27a).

Features in Release 12.0(25)W5(27)

There are no new features in Cisco IOS Release 12.0(25)W5(27).

Features in Release 12.0(18)W5(22b)

There are no new features in Cisco IOS Release 12.0(18)W5(22b).

The 12.0(18)W5(22b) release contains important fixes. If you are currently running 12.0(18)W5(22a) or any earlier release you should migrate to the 12.0(18)W5(22b) release.

New Features in 12.0(18)W5(22a)

Software release 12.0(18)W5(22a) supports the following new features:

- Local Proxy ARP, page8
- RADIUS Server, page8
- CEF Load Balancing on Gigabit Ethernet Ports, page8

Local Proxy ARP

The local proxy ARP feature allows the route processor to respond to ARP requests for IP addresses within a subnet where normally no routing is required. With the local proxy ARP feature enabled, the route processor responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the switch to which they are connected.

The local proxy ARP feature is disabled by default. Use the **ip local-proxy-arp** interface configuration command to enable the local proxy ARP feature on an interface. Use the **no ip local-proxy-arp** interface configuration command to disable the feature. ICMP redirects are disabled on interfaces where the local proxy ARP feature is enabled.

To use the local proxy ARP feature, the IP proxy ARP feature must be enabled. The IP proxy ARP feature is enabled by default. Refer to the *Cisco IOS Release 12.0 Network Protocols Configuration Guide Part 1* , "IP Addressing and Services," "Configuring IP Addressing," "Configure Address Resolution Methods," at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt2/1cipadr.htm

RADIUS Server

The RADIUS feature is a distributed client/server system that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco routers and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

RADIUS is a fully open protocol, distributed in source code format, that can be modified to work with any security system currently available on the market.

Cisco supports RADIUS under its AAA security paradigm. RADIUS can be used with other AAA security protocols, such as TACACS+, Kerberos, or local username lookup. RADIUS is supported on all Cisco platforms. Refer to the *Cisco IOS Release 12.0 Security Configuration Guide*, "Security Server Protocols," "Configuring RADIUS," at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_c/scprt2/scrad.htm

CEF Load Balancing on Gigabit Ethernet Ports

CEF load balancing is based on a combination of source and destination packet information; it allows you to optimize resources by distributing traffic over multiple paths for transferring data to a destination.

You can configure CEF load balancing on a per-destination basis. Load distortions may occur across multiple switch routers when the same CEF load balancing algorithm is used on every switch router. You can resolve these distortions by selecting a specific CEF load balancing algorithm based on your network environment.

New Features in 12.0(14)W5(20)

No new features were added to the 12.0(14)W5(20) release.

New Features in 12.0(10)W5(18e)

The Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers support BGP.

Layer 3 switching software on the Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers features ACLs for Gigabit Ethernet ports. The primary function of an ACL is to provide network control and security, allowing you to filter packet flow into or out of the switch router interface.

Layer 3-switching software also supports IEEE 802.1Q bridging over Fast Ethernet, Gigabit Ethernet, FEC, GEC and BVI. The switch router can be deployed in environments with the 802.1Q trunking protocol and can bridge between ISL and 802.1Q stations.

The Catalyst 2948G-L3 switch router supports IP uplink redirect. IP uplink redirect enables traffic between Fast Ethernet interfaces to be switched through the Gigabit Ethernet interface.

The Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers support per-port QoS traffic conditioning features, such as rate-limiting and shaping.

New Features in 12.0(7)W5(15d)

The Catalyst 2948G-L3 switch router can recover a system image using Xmodem and Ymodem protocols.

Note This feature is new only to the Catalyst 2948G-L3 switch router. The Catalyst 4908G-L3 switch router has this function as part of its basic software features.

New Features in 12.0(7)WX5(15a)

Layer 3 switching software on the Catalyst 2948G-L3 and the Catalyst 4908G-L3 switch router features the switching database manager (SDM). SDM resides on the central processor; its primary function is to maintain the Layer3 switching database in ternary content addressable memory (TCAM). SDM maintains the address entries contained in TCAM in an appropriate order. SDM manages TCAM space by partitioning protocol-specific switching information into multiple regions.

The key benefit of SDM in Layer 3 switching is its ability to configure the size of the protocol regions in TCAM. SDM enables exact-match and longest-match address searches, which result in high-speed forwarding.

For additional information on SDM, refer to the *Catalyst 2948G-L3 and Catalyst 4908G-L3 Software Feature and Configuration Guide.*

Limitations and Restrictions

The following limitations and restrictions apply to the Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers:

- HSRP performance drops when two switch routers are configured with BVI. When HSRP over BVI is configured on both the active and standby HSRP switch routers, and the HSRP-routed packets pass through the standby switch router to reach the active router, the performance of traffic traversing this path may degraded more slowly. You can resolve this problem by using the **standby use-bia** command to configure HSRP to use the burnt in address (BIA) MAC address instead of the HSRP virtual MAC address.
- In the absence of any egress traffic on the master port (one of the following ports: f2, f6, f10, f12, f18, f22, f26, f30, f34, f38, f42, or f46) of the 10/100 Fast Ethernet interface, the Ethernet ports that are in the same EPIF as the master port learn MAC entries slower than their usual rate. For instance, if port f2 has no egress traffic, ports f1, f3, and f4 learn MAC entries at a slower rate than usual.
- CDP will fail on a Gigabit port when trunking is enabled. The switch router will not send CDP packets on a trunk port connected to a Catalyst 4000 Family switch when CDP packets are coming on a VLAN for which a subinterface is not configured. To receive CDP packets, configure a dummy VLAN subinterface on the trunk port connected to the Catalyst 4000 family switch.
- Under extreme conditions, MAC_learn IPC might be lost. A host move under high traffic conditions might cause a MAC entry in the Cisco IOS bridging table to be lost. Routing over BVI might cause loss of connectivity. You can resolve this problem by entering the **clear bridge** command
- A root tree pointer may become invalid for an existing subinterface. This happens very infrequently, when a large configuration is copied to the running configuration under heavy traffic loads. Use the **clear bridge** command to remove the invalid root tree pointer.
- The CLI command **no qos switching** is not supported. Use the **qos mapping precedence** *value* **wrr-weight** *weight* command to configure the same WRRweight for all the precedence values globally, using the CLI.
- When the **no negotiation auto** command is used on a Gigabit port, the link status of that port appears up regardless of the presence of a cable or GBIC on that port.
- If the interface encapsulation is changed to ISL or IEEE 802.1Q on a particular port while there is traffic on the interface, runts and input error counters might increase. However, after the link is stable and normal operation resumes, these counters should not continue to increase.
- An invalid value is returned for SNMP requests for the CiscoFlashDeviceCard MIB object.
- Catalyst 2948G-L3 and Catalyst 4908G-L3 switch routers do not block SNAP encapsulated ARP packets, even though there is switching support only for ARPA encapsulated IP packets. Because of this ARP entries for non-supported IP encapsulations can make it to the ARP table.
- Without IEEE 802.1Q or ISL encapsulation on a subinterface, inbound IPX ACLs do not take effect.
- When spanning tree is disabled in a bridge group, dynamically learned MAC entries will not be deleted immediately from the CAM. If the interface on which the MAC entries were learned goes down the entries will be aged out and removed.

Caveats

Caveats describe unexpected behavior in Cisco IOS software releases. Severity 1 caveats are the most serious caveats; severity 2 caveats are less serious. Caveats listed as open in a prior release are carried forward to the next release as either open or resolved.

For information on caveats in Cisco IOS Release 12.0, see "Caveats for Cisco IOS Release 12.0," which lists severity 1 and 2 caveats for Release 12.0 on Cisco.com and the Documentation CD-ROM.

Note Caveats about Fast Ethernet interfaces do not apply to the Catalyst 4908G-L3 switch router, which has only Gigabit Ethernet interfaces.

Open Caveats in Release 12.0(25)W5(27c)

This section describes the open caveats in Cisco IOS Release 12.0(25)W5(27c):

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

Resolved Caveats in Release 12.0(25)W5(27c)

This section describes the resolved caveats in Cisco IOS Release 12.0(25)W5(27c):

- A specifically crafted Transmission Control Protocol (TCP) connection to a telnet or reverse telnet port of a Cisco device running Internetwork Operating System (IOS) may block further telnet, reverse telnet, Remote Shell (RSH), Secure Shell (SSH), and in some cases Hypertext Transport Protocol (HTTP) access to the Cisco device. Telnet, reverse telnet, RSH and SSH sessions established prior to exploitation are not affected. All other device services will operate normally.

  Conditions: User initiated specially crafted TCP connection to a telnet or reverse telnet port results in blocking further telnet sessions. Whereas, services such as packet forwarding, routing protocols and all other communication to and through the device remains unaffected.

  Workaround: The detail advisory is available at: http://www.cisco.com/warp/public/707/cisco-sa-20040827-telnet.shtml (CSCef46191)

- TCP connections may be vulnerable to spoofed ICMP packets. A spoofed ICMP packet may cause the TCP connection to use a very low segment size for 10 minutes at a time.

  This symptom is observed when TCP connections are configured for PMTU discovery. Note that PMTU discovery is disabled by default on a router.

  Workaround: Disable PMTU discovery. (CSCed78149)

- A Cisco device running IOS and enabled for the Border Gateway Protocol (BGP) is vulnerable to a Denial of Service (DoS) attack from a malformed BGP packet. Only devices with the command bgp log-neighbor-changes configured are vulnerable. The BGP protocol is not enabled by default, and must be configured in order to accept traffic from an explicitly defined peer. Unless the malicious traffic appears to be sourced from a configured, trusted peer, it would be difficult to inject a malformed packet.

  Cisco has made free software available to address this problem.

  This issue is tracked by CERT/CC VU#689326.

  This advisory will be posted at

  http://www.cisco.com/warp/public/707/cisco-sa-20050126-bgp.shtml

  (CSCee67450)

- A document that describes how the Internet Control Message Protocol (ICMP)could be used to perform a number of Denial of Service (DoS) attacks against the Transmission Control Protocol (TCP) has been made publicly available. This document has been published through the Internet Engineering Task Force (IETF) Internet Draft process, and is entitled “ICMP Attacks Against TCP” (draft-gont-tcpm-icmp-attacks-03.txt).

  These attacks, which only affect sessions terminating or originating on a device itself, can be of three types:

  – 1. Attacks that use ICMP “hard” error messages

  – 2. Attacks that use ICMP “fragmentation needed and Don't Fragment (DF) bit set” messages, also known as Path Maximum Transmission Unit Discovery (PMTUD) attacks

  – 3. Attacks that use ICMP “source quench” messages

  Successful attacks may cause connection resets or reduction of throughput in existing connections, depending on the attack type.

  Multiple Cisco products are affected by the attacks described in this Internet draft.

  Cisco has made free software available to address these vulnerabilities. In some cases there are workarounds available to mitigate the effects of the vulnerability.

  This advisory is posted at

  http://www.cisco.com/warp/public/707/cisco-sa-20050412-icmp.shtml.

  The disclosure of these vulnerabilities is being coordinated by the National Infrastructure Security Coordination Centre (NISCC), based in the United Kingdom. NISCC is working with multiple vendors whose products are potentially affected. Its posting can be found at:

  http://www.niscc.gov.uk/niscc/docs/re-20050412-00303.pdf?lang=en.

  (CSCef60659)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Open Caveats in Release 12.0(25)W5(27b)

This section describes the open caveats in Cisco IOS Release 12.0(25)W5(27b):

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Resolved Caveats in Release 12.0(25)W5(27b)

This section describes the resolved caveats in Cisco IOS Release 12.0(25)W5(27b):

- A vulnerability in the Transmission Control Protocol (TCP) specification (RFC793) has been discovered by an external researcher. The successful exploitation enables an adversary to reset any established TCP connection in a much shorter time than was previously discussed publicly. Depending on the application, the connection may get automatically re-established. In other cases, a user will have to repeat the action (for example, open a new Telnet or SSH session). Depending upon the attacked protocol, a successful attack may have additional consequences beyond terminated connection which must be considered. This attack vector is only applicable to the sessions which are terminating on a device (such as a router, switch, or computer) and not to the sessions that are only passing through the device (for example, transit traffic that is being routed by a router). In addition, this attack vector does not directly compromise data integrity or confidentiality.

  All Cisco products which contain TCP stack are susceptible to this vulnerability.

  This advisory is available at:

  http://www.cisco.com/warp/public/707/cisco-sa-20040420-tcp-ios.shtml, and it describes this vulnerability as it applies to Cisco products that run Cisco IOS® software.

  A companion advisory that describes this vulnerability for products that do not run Cisco IOS software is available at:

  http://www.cisco.com/warp/public/707/cisco-sa-20040420-tcp-nonios.shtml. (CSCed27956 and CSCed38527)

- A Cisco device running IOS and enabled for the Border Gateway Protocol (BGP) is vulnerable to a Denial of Service (DOS) attack from a malformed BGP packet. The BGP protocol is not enabled by default, and must be configured in order to accept traffic from an explicitly defined peer. Unless the malicious traffic appears to be sourced from a configured, trusted peer, it would be difficult to inject a malformed packet. BGP MD5 is a valid workaround for this problem.

  Cisco has made free software available to address this problem. For more details, refer to this advisory, available at http://www.cisco.com/warp/public/707/cisco-sa-20040616-bgp.shtml. (CSCdu53656 and CSCea28131)

Open Caveats in Release 12.0(25)W5(27)

This section describes the open caveats in Cisco IOS Release 12.0(25)W5(27):

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Resolved Caveats in Release 12.0(25)W5(27)

This section describes the resolved caveats in Cisco IOS Release 12.0(25)W5(27):

- Appletalk will stop working and the router will stop responding to “GETNETINFO” requests that happen during a Macintosh clients’ bootup when you go to code a cat4232-in-mz.120-25.W5.27.bin image.

  Workaround: Downgrade to an earlier image. (CSCeb70373)

- Cisco routers and switches running Cisco IOS software and configured to process Internet Protocol version 4 (IPv4) packets are vulnerable to a Denial of Service (DoS) attack. A rare sequence of crafted IPv4 packets sent directly to the device may cause the input interface to stop processing traffic once the input queue is full. No authentication is required to process the inbound packet. Processing of IPv4 packets is enabled by default. Devices running only IP version 6 (IPv6) are not affected. A workaround is available.

  Cisco has made software available, free of charge, to correct the problem.

  This advisory is available at

  http://www.cisco.com/warp/public/707/cisco-sa-20030717-blocked.shtml. (CSCea02355)

Open Caveats in Release 12.0(18)W5(22b)

This section describes the open caveats in Cisco IOS Release 12.0(18)W5(22b):

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3. The maximum equal hop paths is set to 2. The IOS routing table shows two destination paths (N1 and N2) in the IPX routing table. When interface I2 is shut down, the IOS routing table should show N1 and N3 as two equal hop paths (because all three paths are equal hop paths.); however, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload.

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Resolved Caveats in Release 12.0(18)W5(22b)

This section describes the resolved caveats in Cisco IOS Release 12.0(18)W5(22b):

- An error can occur with management protocol processing. Go to the following URL for further information:

  http://www.cisco.com/cgi-bin/bugtool/onebug.pl?bugid=CSCdw65903

Open Caveats in Release 12.0(18)W5(22a)

This section describes the open caveats in Cisco IOS Release 12.0(18)W5(22a):

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Resolved Caveats in Release 12.0(18)W5(22a)

This section describes the resolved caveats in Cisco IOS Release 12.0(18)W5(22a):

- An ARP packet received by the router that has the router's own interface address, but with a different MAC address, can overwrite the router's own MAC address in the ARP table, causing that interface to stop sending and receiving traffic. This attack is successful only against interfaces on the Ethernet segment that is local to the attacking host.

  Workaround: Hardcode the interface's ARP table entry by using the **arp** *ip-address hardware-address type* [*alias*] command. This entry will remain in the ARP table until you enter the **clear arp** command.

  Refer to the advisory at the following URL:

  http://www.cisco.com/warp/public/707/IOS-arp-overwrite-vuln-pub.shtml

  This vulnerability does not apply to switches running Catalyst operating system software, only to switches running Cisco IOS software. (CSCdu81936)

- A CPU HOG condition occurs on the switch router after you enter the **no ipx router eigrp** command for routes learned through IEEE 802.1Q encapsulation on the Gigabit Ethernet port. After approximately 15 seconds, the console prompt returns. (CSCdp37972)
- ARP requests intended for any local IP address on the Catalyst 2948G-L3 switch router will be consumed even if the interface is configured only for bridging. The ARP requests will not be flooded inside the bridge group.

  Workaround: Configure static ARPs on the device. (CSCdt43641)

Open Caveats in Release 12.0(14)W5(20)

This section describes the open caveats in Cisco IOS Release 12.0(14)W5(20):

- A CPU HOG condition occurs on the switch router after you enter the **no ipx router eigrp** command for routes learned through IEEE 802.1Q encapsulation on the Gigabit Ethernet port. After approximately 15 seconds, the console prompt returns. (CSCdp37972)
- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- ARP requests intended for any local IP address on the Catalyst 2948G-L3 switch router will be consumed, even if the interface is configured only for bridging. The ARP requests are consumed they will not be flooded inside the bridge group.

  Workaround: Configure static ARPs on the device which requests it. (CSCdt43641)

Resolved Caveats in Release 12.0(14)W5(20)

This section describes the resolved caveats in Cisco IOS Release 12.0(14)W5(20):

- Enabling port-qos features on an interface with the B4 version of XPIF leads to high-performance drops.

  Workaround: A special image that includes a fix is available. Contact TAC for a copy of the image. (CSCdt08870)

- Occasionally when you apply port-qos features, such as rate-limiting and shaping, an interface does not take packet length into account, making the output bit rate inaccurate.

  Workaround: This problem is fixed in Cisco 12.0(10)W05(17.113), which is an integrated release not available on Cisco.com. This IOS image will be provided on a case-by-case basis for customers requiring port-qos features. The next maintenance release will incorporate this fix. (CSCds82323)

- When accessed through SNMP, the QoS mapping table lists an entry with the wrong precedence index value of 4. This value must be in a range from 0 to 3. (CSCdr24893)

- A Border Gateway Protocol (BGP) UPDATE contains Network Layer Reachability Information (NLRI) and attributes that describe the path to the destination. Each path attribute is a type, length, value (TLV) object.

  The type is a two-octet field that includes the attribute flags and the type code. The fourth high-order bit (bit 3) of the attribute flags is the Extended Length bit. It defines whether the attribute length is one octet (if set to 0) or two octets (if set to 1). The extended length bit is used only if the length of the attribute value is greater than 255 octets.

  The AS_PATH (type code 2) is represented by a series of TLVs (or path segments). The path segment type indicates whether the content is an AS_SET or AS_SEQUENCE. The path segment length indicates the number of autonomous systems (ASes) in the segment. The path segment value contains the list of ASes (each AS is represented by two octets).

  The total length of the attribute depends on the number of path segments and the number of ASes in them. For example, if the AS_PATH contains only an AS_SEQUENCE, then the maximum number of ASes (without having to use the extended length bit) is 126 [= (255-2)/2]. If the UPDATE is propagated across an AS boundary, then the local Abstract Syntax Notation (ASN) must be appended and the extended length bit used.

  The caveat was caused by the mishandling of the operation during which the length of the attribute was truncated to only one octet. Because of the internal operation of the code, the receiving border router would not be affected, but its iBGP peers would detect the mismatch and issue a NOTIFICATION message (update malformed) to reset their session.

  The average maximum AS_PATH length in the Internet is between 15 and 20 ASes, so there is no need to use the extended length. The failure was discovered because of a malfunction in the BGP implementation of another vendor. There is no workaround.

  [Part of the text was taken from rfc 1771.] (CSCdr54230)

- Cisco Security Advisory:

  Cisco IOS Software TCP Initial Sequence Number Randomization Improvements

  Revision 1.0: INTERIM

  For Public Release 2001 February 27 20:00 US/Eastern (UTC+0500)

  Summary

  Cisco IOS software contains a flaw that permits the successful prediction of TCP Initial Sequence Numbers.

  This vulnerability is present in all released versions of Cisco IOS software running on Cisco routers and switches. It only affects the security of TCP connections that originate or terminate on the affected Cisco device itself; it does not apply to TCP traffic forwarded through the affected device in transit between two other hosts.

  To remove the vulnerability, Cisco is offering free software upgrades for all affected platforms. The defect is described in DDTS record CSCds04747.

  Workarounds are available that limit or deny successful exploitation of the vulnerability by filtering traffic containing forged IP source addresses at the perimeter of a network or directly on individual devices.

  This notice will be posted at http://www.cisco.com/warp/public/707/ios-tcp-isn-random-pub.shtml. (CSCds04747)
- When BGP sessions get reset, currently, with lob neighbor-changes, the even is errlogged. However, to find out the reasons as to why there was a reset, one has to turn on the debugs. This fix will automatically errlog the NOTIFICATION message when the sessions are reset. This feature will be turned on by the same log neighbor-changes knob. (CSCdr54231

Open Caveats in Release 12.0(10)W5(18e)

This section describes the open caveats in Cisco IOS Release 12.0(10)W5(18e):

- Enabling port-qos features on an interface with the B4 version of XPIF leads to high performance drops.

  Workaround: A special image that includes a fix is available. Contact TAC for a copy of the image. (CSCdt08870)
- Occasionally when you apply port-qos features, such as rate-limiting and shaping, an interface does not take packet length into account, making the output bit rate inaccurate.

  Workaround: This problem is fixed in Cisco 12.0(10)W05(17.113), which is an integrated release not available on Cisco.com. This IOS image will be provided on a case-by-case basis for customers requiring port-qos features. The next maintenance release will incorporate this fix. (CSCds82323)
- When accessed through SNMP, the QoS mapping table lists an entry with the wrong precedence index value of 4. This value must be in a range from 0 to 3. (CSCdr24893)
- A CPU HOG condition occurs on the switch router after you enter the **no ipx router eigrp** command for routes learned through IEEE 802.1Q encapsulation on the Gigabit Ethernet port. After approximately 15 seconds the console prompt returns. (CSCdp37972)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

Resolved Caveats in Release 12.0(10)W5(18e)

This section describes the resolved caveats in Cisco IOS Release 12.0(10)W5(18e):

- Address resolution protocol (ARP) packets are consumed and flooded by IOS even though IP routing is turned off globally. (CSCdr39535)
- When IP multicast routing is enabled on a BVI, nonReverse Path Forwarding (RPF) IP multicast data packets are discarded. Because of this some hosts in the same bridge domain might be disconnected. (CSCdr35926)
- When IP multicast routing is not enabled over BVI and BVI is configured to route IP, IP multicast data packets are discarded. (CSCdr35855)
- Under heavy traffic conditions, remote MAC entries are not learned consistently on all FEC members. This causes traffic to flood in one direction.

  Workaround: Use the **clear bridge** command to resync the remote entries on all ports. (CSCdp53223)

- When BVI is configured, the following message appears when the switch router is reloaded:

  ```
  %CEF hwidb not found for BVI1
  ```

  This is a harmless message. (CSCdp51343)

- The total number of IPX networks that can be added to the BVI interfaces is restricted to 32. (CSCdp49222)
- Packets are switched out on the native VLAN, leading to routing by the CPU (with BVI). Untagged packets coming in on the 802.1Q native VLAN are not processed by the microcode. Instead they are given to the CPU, and the CPU does the processing. This means that high CPU utilization will be seen if untagged packets are received at a high rate on the native VLAN subinterfaces. (CSCdp33630)

Note Generally, only management data, transmitted at a very low rate, would be seen on the native VLAN, because it is mainly used for network management purposes.

Open Caveats in Release 12.0(7)W5(15d)

This section describes the open caveats in Cisco IOS Release 12.0(7)W5(15d):

- Address resolution protocol (ARP) packets are consumed and flooded by IOS even though IP routing is turned off globally. (CSCdr39535)
- When IP multicast routing is enabled on a BVI, nonReverse Path Forwarding (RPF) IP multicast data packets are discarded. Because of this some hosts in the same bridge domain might be disconnected. (CSCdr35926)
- When IP multicast routing is not enabled over BVI and BVI is configured to route IP, IP multicast data packets are discarded. (CSCdr35855)
- When accessed through SNMP, the QoS mapping table lists an entry with the wrong precedence index value of 4. This value must be in a range from 0 to 3. (CSCdr24893)
- Under heavy traffic conditions, remote MAC entries are not learned consistently on all FEC members. This causes traffic to flood in one direction.

  Workaround: Use the **clear bridge** command to resync the remote entries on all ports. (CSCdp53223)
- When BVI is configured, the following message appears when the switch router is reloaded:

  ```
  %CEF hwidb not found for BVI1
  ```

  This is a harmless message. (CSCdp51343)
- The total number of IPX networks that can be added to the BVI interfaces is restricted to 32. (CSCdp49222)
- A CPU HOG condition occurs on the switch router after you enter the **no ipx router eigrp** command for routes learned through IEEE 802.1Q encapsulation on the Gigabit Ethernet port. After approximately 15 seconds the console prompt returns. (CSCdp37972)
- Packets are switched out on the native VLAN, leading to routing by the CPU (with BVI). Untagged packets coming in on the 802.1Q native VLAN are not processed by the microcode. Instead they are given to the CPU, and the CPU does the processing. This means that high CPU utilization will be seen if untagged packets are received at a high rate on the native VLAN subinterfaces. (CSCdp33630)

Note In general, only management data transmitted at a very low rate would be seen on the native VLAN, because it is mainly used for network management purposes.

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload.

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Resolved Caveats in Release 12.0(7)W5(15d)

This section describes the resolved caveats in Cisco IOS Release 12.0(7)W5(15d):

- When the interface encapsulation of a port channel subinterface is changed from ISL to 802.1Q, or vice versa, and if this subinterface is the only member of a bridge-group, the console might not respond. However, the system runs normally, and the user can access the console using Telnet as long as there is IP connectivity to the system. (CSCdp56448)

  Workaround: Before attempting an encapsulation change on the port channel subinterface, do one of the following:

  – Remove the bridge-group globally, by entering the **no bridge** *bridge_id* command.

  – Add an additional member to the bridge-group.

  – Remove all members of the port channel.

- The 64-byte packets transmitted by the Gigabit Ethernet ports are not updated in the interface statistics. This does not impact performance. (CSCdp27003)
- When traffic to the CPU from many interfaces is very high, the CPU might have memory allocation failures for the packet buffers. (CSCdp54671)
- The front panel link LED for Fast Ethernet ports is green for the 10/100 Mbps link speed. The link LED is not amber for 10 Mbps, as mentioned in the *Catalyst 2948G-L3 Hardware Installation Guide*. (CSCdp43594)
- With constrained multicast flooding (CMF), when a bridge group member joins and leaves a multicast group, the multicast traffic is still forwarded out of that bridge member.

  Workaround: Use the **clear bridge multicast** command to stop forwarding traffic out of the bridge member. (CSCdp42656)

Open Caveats in Release 12.0(7)Wx5(15a)

This section describes the open caveats in Cisco IOS Release 12.0(7)Wx5(15a):

Note Caveats about Fast Ethernet interfaces do not apply to the Catalyst 4908G-L3 switch router, which has only Gigabit Ethernet interfaces.

- Address resolution protocol (ARP) packets are consumed and flooded by IOS even though IP routing is turned off globally. (CSCdr39535)
- When IP multicast routing is enabled on a BVI, nonReverse Path Forwarding (RPF) IP multicast data packets are discarded. Because of this some hosts in the same bridge domain might be disconnected. (CSCdr35926)

- When IP multicast routing is not enabled over BVI and BVI is configured to route IP, IP multicast data packets are discarded. (CSCdr35855)
- When accessed through SNMP, the QoS mapping table lists an entry with the wrong precedence index value of 4. This value must be in a range from 0 to 3. (CSCdr24893)
- When traffic to the CPU from many interfaces is very high, the CPU might have memory allocation failures for the packet buffers. (CSCdp54671)
- Under heavy traffic conditions, remote MAC entries are not learned consistently on all FEC members. This causes traffic to flood in one direction.

  Workaround: Use the **clear bridge** command to resync the remote entries on all ports. (CSCdp53223)
- When the interface encapsulation of a port channel subinterface is changed from ISL to 802.1Q, or vice versa, and if this subinterface is the only member of a bridge-group, the console might not respond. However, the system runs normally, and the user can access the console using Telnet as long as there is IP connectivity to the system. (CSCdp56448)

  Workaround: Before attempting an encapsulation change on the port channel subinterface, do one of the following:

  – Remove the bridge-group globally, by entering the **no bridge** *bridge_id* command.

  – Add an additional member to the bridge-group.

  – Remove all members of the port channel.
- When BVI is configured, the following message appears when the switch router is reloaded:

  ```
  %CEF hwidb not found for BVI1
  ```

  This is a harmless message. (CSCdp51343)
- The total number of IPX networks that can be added to the BVI interfaces is restricted to 32. (CSCdp49222)
- A CPU HOG condition occurs on the switch router after you enter the **no ipx router eigrp** command for routes learned through IEEE 802.1Q encapsulation on the Gigabit Ethernet port. After approximately 15 seconds the console prompt returns. (CSCdp37972)
- The front panel link LED for Fast Ethernet ports is green for the 10/100 Mbps link speed. The link LED is not amber for 10 Mbps, as mentioned in the *Catalyst 2948G-L3 Hardware Installation Guide*. (CSCdp43594)
- With constrained multicast flooding (CMF), when a bridge group member joins and leaves a multicast group, the multicast traffic is still forwarded out of that bridge member.

  Workaround: Use the **clear bridge multicast** command to stop forwarding traffic out of the bridge member. (CSCdp42656)
- Packets are switched out on the native VLAN, leading to routing by the CPU (with BVI). Untagged packets coming in on the 802.1Q native VLAN are not processed by the microcode. Instead they are given to the CPU, and the CPU does the processing. This means that high CPU utilization will be seen if untagged packets are received at a high rate on the native VLAN subinterfaces. (CSCdp33630)

Note In general, only management data transmitted at a very low rate would be seen on the native VLAN, because it is mainly used for network management purposes.

- Cisco IOS does not update the IPX routing table when more than two equal hop paths are available and one of them is shut down. For example, a switch with three interfaces (I1, I2, and I3) might have an IPX network configured on each interface, as N1, N2, and N3, respectively. A remote IPX network (R) is accessible through N1, N2, and N3 interfaces. The maximum equal hop paths is set to 2. The IOS routing table will show two destination paths (N1 and N2) in the IPX routing table, the interface I2 will shut down. Because all the three paths are equal hop, the IOS routing table should show N1 and N3 as two equal hop paths. However, the routing table shows only N1 as the destination path.

  Workaround: Enter the **clear ipx route** command. The routing table will show N1 and N3 as the destination next hop paths. (CSCdp13515)

- The 64-byte packets transmitted by the Gigabit Ethernet ports are not updated in the interface statistics. This does not impact performance. (CSCdp27003)

- In configurations with a large number of bridge groups and bridge group members, you might see the following traceback message during reload:

  ```
  00:00:38:%SYS-3-CPUHOG: Task ran for 3084 msec (437/1), process = CDP Protocol, PC =
  6015BD40.
  ```

  Workaround: Disable CDP on all interfaces or reduce the number of interfaces on which CDP is configured until the problem no longer occurs. (CSCdr97028)

Resolved Caveats in Release 12.0(7)Wx5(15a)

There were no resolved caveats for Release 12.0(7)Wx5(15a).

Error Messages

This section describes Catalyst 2948G-L3 and Catalyst 4908G-L3 switch router error messages.

Note Error messages about Fast Ethernet interfaces do not apply to the Catalyst 4908G-L3 switch router, which has only Gigabit Ethernet interfaces.

- When you use the **sdm size** command in an attempt to add more IP multicast routes than the size previously configured, the following error message is displayed.

  For example, if the number of IP multicast routes added to SDM exceeds the size configured with the **sdm size** command, the following error message is displayed:

  ```
  %LSS-1-SDM: IP Multicast, Region reached limit Cannot accept more entries
  ```

  Explanation Cannot accept more routes in SDM.

  Recommended Action To rectify this situation, increase the protocol size using the **sdm size configuration** command and reload the switch router.

- When you attempt to add more routes than the Fast Ethernet CAM size allows, the following error message is displayed:

  ```
  7:28:06:%LSS-4-INTERFACE:(Interface FastEthernet2) CAM reached limit. Cannot
  accept more route entries
  ```

  Explanation Cannot accept more routes in CAM.

  Recommended Action None.

Related Documentation

This section describes the documentation available for the Catalyst 2948G-L3 and Catalyst4908G-L3 switch routers. Both printed manuals and electronic documents are available.

The most current documentation is available on Cisco.com and the Documentation CD-ROM. These electronic documents might contain updates and modifications made after the hard-copy documents were printed.

Use these release notes with the following documents:

- *Catalyst 2948G-L3 Hardware Installation Guide*
- *Catalyst 2948G-L3 and Catalyst 4908G-L3 Software Feature and Configuration Guide*
- *Catalyst 4908G-L3 Hardware Installation Guide*
- For information about MIBs, refer to:
  http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Service and Support

For service and support for a product purchased from a reseller, contact the reseller. Resellers offer a wide variety of Ciscoservice and support programs, which are described in the section “Service and Support” in the information packet that was shipped with your product.

Note If you purchased your product from a reseller, you can access Cisco.com as a guest. Cisco.com is CiscoSystems’ primary real-time support channel. Your reseller offers programs that include direct access to Cisco.com services.

For service and support for a product purchased directly from Cisco, use Cisco.com.

Software Configuration Tips on the Cisco TAC Home Page

For helpful tips on configuring Cisco products, follow this path on Cisco.com:

Service & Support: Technical Assistance Center

“Software Technical Tips” are popular tips and hints gathered from Cisco's Technical Assistance Center (TAC). Most of these documents are also available from the TAC's Fax-on-Demand service. To access Fax-on-Demand and receive documents at your fax machine, call 888-50-CISCO (888-502-4726). From international areas, call 650-556-8409.

In addition to “Software Technical Tips,” the following sections are on the Technical Documents page:

- Cisco Product Catalog—MultiNet & Cisco Suite 100, Network Management, Cisco IOS Software Bulletins, CiscoPro Configurations.
- Field Notices—Notification of critical issues regarding Cisco products. These include problem descriptions, safety or security issues, and hardware defects.
- Hardware Technical Tips—Technical tips related to specific hardware platforms.
- Hot Tips—Popular tips and hints for a range of product suites, gathered from Cisco's Technical Assistance Center (TAC).
- Internetworking Technical Tips—Tips for using and deploying Cisco IOS software features and services.
- Sample Configurations—Actual configuration examples complete with topology and annotations.
- Special Collections—Other helpful documents: Frequently Asked Questions, Security Advisories, References & RFCs, Case Studies, and the CiscoPro Documentation CD-ROM.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation at this URL:

http://www.cisco.com/univercd/home/home.htm

You can access the Cisco website at this URL:

http://www.cisco.com

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpck/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:

  http://www.cisco.com/en/US/partner/ordering/index.shtml
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408526-7208 or, elsewhere in North America, by calling 800553-NETS (6387).

Documentation Feedback

You can send comments about technical documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, Cisco Technical Support provides 24-hour-a-day, award-winning technical assistance. The Cisco Technical Support Website on Cisco.com features extensive online support resources. In addition, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not hold a valid Cisco service contract, contact your reseller.

Cisco Technical Support Website

The Cisco Technical Support Website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, 365 days a year at this URL:

http://www.cisco.com/techsupport

Access to all tools on the Cisco Technical Support Website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

http://tools.cisco.com/RPF/register/register.do

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool automatically provides recommended solutions. If your issue is not resolved using the recommended resources, your service request will be assigned to a Cisco TAC engineer. The TAC Service Request Tool is located at this URL:

http://www.cisco.com/techsupport/servicerequest

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco TAC engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)
EMEA: +32 2 704 55 55
USA: 1 800 553 2447

For a complete list of Cisco TAC contacts, go to this URL:

http://www.cisco.com/techsupport/contacts

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—Your network is "down," or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

  http://www.cisco.com/go/marketplace/

- The Cisco *Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the Cisco Product Catalog at this URL:

  http://cisco.com/univercd/cc/td/doc/pcat/

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

  http://www.ciscopress.com

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

  http://www.cisco.com/packet

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

  http://www.cisco.com/go/iqmagazine

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

  http://www.cisco.com/ipj

- World-class networking training is available from Cisco. You can view current offerings at thisURL:

  http://www.cisco.com/en/US/learning/index.html

This document is to be used in conjunction with the documents listed in the "Related Documentation" section.