



Release Notes for the Catalyst 3750 Metro Switch, Cisco IOS Release 12.1(14)AX2

Revised April 2005

The Cisco IOS Release 12.1(14)AX2 runs on all Catalyst 3750 Metro switches.

These release notes include important information about this Cisco IOS release and any limitations, restrictions, and caveats that apply to it. Verify that these release notes are correct for your switch:

- If you are installing a new switch, refer to the Cisco IOS release label on the rear panel of your switch.
- If your switch is on, use the **show version** privileged EXEC command. See the [“Finding the Software Version and Feature Set”](#) section on page 3.
- If you are upgrading to a new release, refer to the software upgrade filename for the software version.

For the complete list of switch documentation, see the [“Related Documentation”](#) section on page 29.

You can download the switch software from these sites:

- <http://www.cisco.com/kobayashi/sw-center/sw-lan.shtml>
(for registered Cisco.com users with a login password)
- <http://www.cisco.com/public/sw-center/sw-lan.shtml>
(for nonregistered Cisco.com users)

This software release is part of a special release of Cisco IOS software that is not released on the same 8-week maintenance cycle that is used for other platforms. As maintenance releases and future software releases become available, they will be posted to Cisco.com (previously Cisco Connection Online [CCO]) in the Cisco IOS software area.



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

Copyright © 2004 Cisco Systems, Inc. All rights reserved.

Contents

This information is in the release notes:

- [“Hardware Supported” section on page 2](#)
- [“Downloading Software” section on page 3](#)
- [“Installation Notes” section on page 5](#)
- [“New Features” section on page 5](#)
- [“Minimum Cisco IOS Release for Major Features” section on page 6](#)
- [“Limitations and Restrictions” section on page 6](#)
- [“Open Caveats” section on page 10](#)
- [“Resolved Caveats” section on page 16](#)
- [“Documentation Updates” section on page 22](#)
- [“Related Documentation” section on page 29](#)
- [“Obtaining Documentation” section on page 29](#)
- [“Obtaining Technical Assistance” section on page 30](#)
- [“Obtaining Additional Publications and Information” section on page 31](#)

Hardware Supported

Table 1 lists the supported hardware and the minimum Cisco IOS release required.

Table 1 Supported Hardware

Switch	Description	Supported by Minimum Cisco IOS Release
Catalyst 3750 Metro 24-AC switch	24 10/100 Ethernet ports, 2 1000X standard SFP ¹ module slots, 2 1000X ES ² SFP slots, and field-replaceable AC power supply	Cisco IOS Release 12.1(14)AX
Catalyst 3750 Metro 24-DC switch	24 10/100 Ethernet ports, 2 1000X standard SFP module slots, 2 1000X ES SFP slots, and field-replaceable DC power supply	Cisco IOS Release 12.1(14)AX
SFP modules	1000BASE-T, 1000BASE-SX, and 1000BASE-LX	Cisco IOS Release 12.1(14)AX
	1000BASE-ZX and CWDM ³	Cisco IOS Release 12.1(14)AX1

1. SFP = small form-factor pluggable
2. ES = enhanced services
3. CWDM = coarse wavelength-division multiplexer

Downloading Software

These are the procedures for downloading software:

- “Finding the Software Version and Feature Set” section on page 3
- “Deciding Which Files to Use” section on page 3
- “Upgrading a Switch by Using the CLI” section on page 4
- “Recovering from a Software Failure” section on page 5



Note

Before downloading software, read this section for important information.

Finding the Software Version and Feature Set

The Cisco IOS image is stored as a bin file in a directory that is named with the Cisco IOS release. The image is stored on the system board flash device (flash:).

You can use the **show version** privileged EXEC command to see the software version that is running on your switch.

You can also use the **dir filesystem:** privileged EXEC command to see the directory names of other software images that you might have stored in flash memory.

Deciding Which Files to Use

The upgrade procedures in these release notes describe how to perform the upgrade by using a combined tar file. This file contains the Cisco IOS image file. To upgrade the switch through the command-line interface (CLI), use the tar file and the **archive download-sw** privileged EXEC command.

[Table 2](#) lists the software filename for this software release.

Table 2 Cisco IOS Software Image Files for Catalyst 3750 Metro Switches

Filename	Description
c3750me-i5-tar.121-14.AX2.tar	Cisco IOS image tar file. This image has Layer 2+ and Layer 3 features.
c3750me-i5k2-tar.121-14.AX2.tar	Cisco IOS crypto image tar file. This image has the Kerberos, SSH, Layer 2+, and Layer 3 features.

Upgrading a Switch by Using the CLI

This procedure is for copying the tar file to the switch. You copy the file to the switch from a TFTP server and extract the files. You can download an image file and replace or keep the current image.

Download the software from Cisco.com to your management station by following these steps:

-
- Step 1** Use [Table 2 on page 3](#) to identify the file that you want to download.
- Step 2** Download the software image file from Cisco.com.
- If you have a SmartNet support contract, go to this URL and log in to download the appropriate files:
<http://www.cisco.com/kobayashi/sw-center/sw-lan.shtml>
 - If you do not have a SmartNet contract, go to this URL and follow the instructions to register on Cisco.com and download the appropriate files:
<http://www.cisco.com/public/sw-center/sw-lan.shtml>

To download the files, click the link for your switch platform, and then follow the links on the page to select the correct tar image file.

- Step 3** Copy the image to the appropriate TFTP directory on the workstation, and make sure that the TFTP server is properly configured.
- For more information, refer to Appendix B in the software configuration guide for this release.
- Step 4** Log in to the switch through the console port or a Telnet session.
- Step 5** Check your VLAN 1 configuration by using the **show interfaces vlan 1** privileged EXEC command, and verify that VLAN 1 is part of the same network as the TFTP server. (Check the *Internet address is* line near the top of the display.)
- Step 6** Download the image file from the TFTP server to the switch. If you are installing the same version of software that is currently on the switch, overwrite the current image by using this privileged EXEC command:

```
archive download-sw /overwrite /reload tftp:[[//location]/directory]/image-name.tar
```

The **/overwrite** option overwrites the software image in flash memory with the downloaded one.

The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved.

For *//location*, specify the IP address of the TFTP server.

For */directory/image-name.tar*, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

This example shows how to download an image from a TFTP server at 198.30.20.19 and to overwrite the image on the switch:

```
Switch# archive download-sw /overwrite tftp://198.30.20.19/c3750me-i5-tar.121-14.AX2.tar
```

You can also download the image file from the TFTP server to the switch and keep the current image by replacing the **/overwrite** option with the **/leave-old-sw** option.

Recovering from a Software Failure

Switch software can be corrupted during an upgrade, by downloading the wrong file to the switch, and by deleting the image file. In all of these cases, the switch does not pass the power-on self-test (POST), and there is no connectivity. You can use the Xmodem protocol to recover from these failures.

For detailed recovery procedures, refer to the “Troubleshooting” chapter in the software configuration guide for this release.

Installation Notes

You can assign IP information to your switch by using these methods:

- The Express Setup program (Refer to the *Catalyst 3750 Metro Switch Hardware Installation Guide*.)
- The CLI-based setup program (Refer to the *Catalyst 3750 Metro Switch Hardware Installation Guide*.)
- The DHCP-based autoconfiguration (Refer to the *Catalyst 3750 Metro Switch Software Configuration Guide*.)
- Manually assigning an IP address (Refer to the *Catalyst 3750 Metro Switch Software Configuration Guide*.)

New Features

These are the new supported hardware and the new software features provided this release:

- [“New Hardware Features” section on page 5](#)
- [“New Software Features” section on page 5](#)

New Hardware Features

There are no new hardware features in this release.

New Software Features

This release contains these new switch features or enhancements:

- Quality of service (QoS) egress priority queue on a port. Shaped round robin (SRR) services the priority queue until it is empty before servicing the other three queues. You enable the egress priority queue by using the **priority-queue out** interface configuration command.
- QoS Differentiated Services Code Point (DSCP) transparency allows the user to prevent the switch from rewriting the DSCP field in the user IP packet when QoS is enabled.

For information about these features, see the [“Documentation Updates” section on page 22](#).

Minimum Cisco IOS Release for Major Features

Table 3 lists the minimum software release required to support the major features on the Catalyst 3750 Metro switch.


Note

Features not included in the table are available in all releases. You can see a list of features from the first release at this URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750m/12114ax/3750msg/swintro.htm#68019>

Table 3 Catalyst 3750 Metro Switch Features and the Minimum Cisco IOS Release Required

Feature	Minimum Cisco IOS Release Required	For more information,
QoS egress priority queue	12.1(14)AX2	See the “Documentation Updates” section on page 22.
QoS DSCP transparency	12.1(14)AX2	See the “Documentation Updates” section on page 22.
Point-to-point Layer 2 protocol tunneling	12.1(14)AX1	See the “Documentation Updates” section on page 22.

Limitations and Restrictions

You should review this section before you begin working with the switch. These are known limitations that will not be fixed, and there is not always a workaround. Some features might not work as documented, and some features could be affected by recent changes to the switch hardware or software.

Cisco IOS Limitations and Restrictions

These limitations apply to Cisco IOS configuration:

- The switch does not support tunnel interfaces for unicast routed traffic. Only Distance Vector Multicast Routing Protocol (DVMRP) tunnel interfaces are supported for multicast routing.
- Nonreverse-path forwarded (RPF) IP multicast traffic to a group that is bridged in a VLAN is leaked onto a trunk port in the VLAN even if the port is not a member of the VLAN group, but it is a member in some other VLAN group. Unnecessary traffic is sent on the trunk port and needlessly reduces the bandwidth of the port. There is no workaround because non-RPF traffic is continuous in certain topologies. As long as the trunk port is a member on a trunk port in at least one VLAN, this problem for the non-RPF traffic occurs. (CSCdu25219)
- If a bridge group contains a VLAN that has a static MAC address configured, all non-IP traffic in the bridge group with this MAC address destination is sent to all ports in the bridge group. The workaround is to remove the VLAN from the bridge group or to remove the static MAC address from the VLAN. (CSCdw81955)

- If the number of multicast routes and Internet Group Management Protocol (IGMP) groups are more than the maximum number in the Switch Database Management (SDM) template shown with the **show sdm prefer** global configuration command, the traffic received on unknown groups is flooded in the received VLAN even though the **show ip igmp snooping multicast-table** privileged EXEC command output shows otherwise. The workaround is to reduce the number of multicast routes and IGMP snooping groups to less than the maximum supported value. (CSCdy09008)
- An egress SPAN copy of routed unicast traffic might show an incorrect destination MAC address on both local and remote SPAN sessions. This limitation does not apply to bridged packets. The workaround for local SPAN is to use the replicate option. There is no workaround for a remote SPAN session. This is a hardware limitation. (CSCdy72835)
- Egress SPAN routed packets (both unicast and multicast) show the incorrect source MAC address. For remote SPAN packets, the source MAC address should be the MAC address of the egress VLAN, but instead the packet shows the MAC address of the remote SPAN (RSPAN) VLAN. For local SPAN packets with native encapsulation on the destination port, the packet shows the MAC address of VLAN 1. This problem does not appear with local SPAN when the encapsulation replicate option is used. This limitation does not apply to bridged packets. The workaround is to use the **encapsulate replicate** keywords in the **monitor session** global configuration command. Otherwise, there is no workaround. This is a hardware limitation. (CSCdy81521)
- IGMP filtering is applied to packets that are forwarded through hardware. It is not applied to packets that are forwarded through software. Hence, with multicast routing enabled, the first few packets are sent from a port even when IGMP filtering is set to deny those groups on that port. There is no workaround. (CSCdy82818)
- A static IP address might be removed when the previously acquired DHCP IP address lease expires. This problem occurs under these conditions:
 - When the switch is booted without a configuration (no config.text file in flash memory).
 - When the switch is connected to a DHCP server that is configured to give an address to it (the dynamic IP address is assigned to VLAN 1).
 - When an IP address is configured on VLAN 1 before the dynamic address lease assigned to VLAN 1 expires.
 The workaround is to reconfigure the static IP address. (CSCCea71176)
- The switch treats frames received with mixed encapsulation (802.1Q and Inter-Switch Link [ISL]) as frames with FCS errors, increments the error counters, and causes the port LED to blink amber. This happens when an ISL-unaware device receives an ISL-encapsulated packet and forwards the frame to an 802.1Q trunk interface. There is no workaround. (CSCdz33708)
- IP-option traffic is sometimes leaked unnecessarily on a trunk port. Suppose the trunk port in question is a member of an IP multicast group in VLAN X, but it is not a member in VLAN Y. In VLAN Y, there is another port that has membership in the group, and VLAN Y is the output interface for the multicast route entry corresponding to the group. IP-options traffic received on an input interface VLAN (other than VLAN Y) is unnecessarily sent on the trunk port in VLAN Y because the trunk port is forwarding in VLAN Y (even though the port has no group membership in VLAN Y). There is no workaround. (CSCdz42909)

- Known unicast (secured addresses) are flooded within a bridge group under this condition: If secure addresses are learned or configured on a port and the VLAN on this port is part of a bridge group, non-IP traffic destined to the secure addresses is flooded within the bridge group. The workaround is to disable fallback bridging. To remove an interface from a bridge group and to remove the bridge group, use the **no bridge-group** *bridge-group* interface configuration command. Another workaround is to disable port security on all ports in all VLANs participating in fallback bridging by using the **no switchport port-security** interface configuration command. (CSCdz80499)
- When you use the **ip access-group** interface configuration command with a router access control list (ACL) to deny access to a group in a VLAN, multicast data to the group that is received in the VLAN is always flooded in the VLAN regardless of IGMP group membership in the VLAN. This provides access to directly connected clients, if any, in the VLAN. The workaround is to not apply a router ACL configured to deny access to a VLAN interface. Apply the security through other means; for example, apply VLAN maps to the VLAN instead of using a router ACL for the group. (CSCdz86110)
- SNAP-encapsulated IP packets are dropped without an error message being reported at the interface. The switch does not support SNAP-encapsulated IP packets. There is no workaround. (CSCdz89142)
- The switch does not create an adjacency table entry when the Address Resolution Protocol (ARP) timeout value is 15 seconds and the ARP request times out.
The workaround is to set an ARP timeout value higher than 120 seconds. (CSCea21674)
- A route map that contains an ACL with a DSCP clause cannot be applied to a Layer 3 interface. The switch rejects this configuration and issues an error message that shows that the route map is unsupported. There is no workaround. (CSCea52915)
- During periods of very high traffic and when two RSPAN source sessions are configured, the VLAN ID of packets in one RSPAN session might overwrite the VLAN ID of the other RSPAN session. If this occurs, packets intended for one RSPAN VLAN are incorrectly sent to the other RSPAN VLAN. This problem does not affect RSPAN destination sessions. The workaround is to configure only one RSPAN source session. (CSCea72326)
- Some switch queues are disabled if the buffer size or threshold level is set too low with the **mls qos queue-set output** global configuration command. The ratio of buffer size to threshold level should be greater than ten to avoid disabling the queue. The workaround is to choose compatible buffer sizes and threshold levels. (CSCea76893)
- When an IP phone is connected to the switch, the port VLAN ID (PVID) and the voice VLAN ID (VVID) both learn its MAC address. However, after dynamic MAC addresses are deleted, only VVID relearns the IP phone MAC address. MAC addresses are deleted manually or automatically for a topology change or when port security or an 802.1x feature is enabled or disabled. There is no workaround. (CSCea80105)
- After changing the access VLAN on a port that has 802.1x enabled, the IP phone address is removed. Because learning is restricted on 802.1x capable ports, it takes approximately 30 seconds before the address is relearned. There is no workaround. (CSCea85312)
- The egress-SPAN data rate might degrade when fallback bridging or multicast routing is enabled. The amount of degradation depends on the processor loading. Typically, the switch can egress-SPAN at up to 40,000 packets per second (64-byte packets). As long as the total traffic being monitored is below this limit, there is no degradation. However, if the traffic being monitored exceeds the limit, only a portion of the source stream is spanned. When this occurs, this console message appears: `Decreased egress SPAN rate.`

In all cases, normal traffic is not affected; the degradation limits only how much of the original source stream can be egress-spanned. If fallback bridging and multicast routing are disabled, egress-SPAN is not degraded. There is no workaround. If possible, disable fallback bridging and multicast routing. If possible, use ingress-SPAN to observe the same traffic. (CSCeb01216)

- Some IGMP report and query packets with IP options might not be ingress-spanned. Packets that are susceptible to this problem are IGMP packets containing 4 bytes of IP options (IP header length of 24). An example of such packets would be IGMP reports and queries having the router alert IP option. Ingress-spanning of such packets is not accurate and can vary with traffic rate. Typically, very few or none of these packets are spanned. There is no workaround. (CSCeb23352)
- If the number of VLANs times the number of trunk ports exceeds the recommended limit of 13000, the switch can halt. The workaround is to reduce the number of VLANs or trunks. (CSCeb31087)
- Port-based Ethernet over Multiprotocol Label Switching (EoMPLS) sessions do not function if the incoming port is configured as an Inter-Switch Link (ISL) trunk. The workaround is to configure the incoming ports as an 802.1Q trunk or as an access port. (CSCeb44014)
- When traffic with different class of service (CoS) values is sent into a 802.1Q tunnel, only the CoS 0 statistics increment in the **show mls qos interface** user EXEC command display. There is no workaround. (CSCeb75230)
- The display for the **show mpls ldp neighbor ipaddr-of-neighbor detail** user EXEC command always shows the targeted hello holdtime value as *infinite*. The workaround is to use the **show mpls ldp parameter** user EXEC command to see the configured value. (CSCeb76775)
- The **bandwidth** interface configuration command is not supported at the interface level, but it appears in the CLI. There is no workaround. (CSCeb80223)
- The **random-detect** interface configuration command is not supported at the interface level, but it appears in the CLI. There is no workaround. (CSCeb80300)
- The display for the **show policy-map interface** user EXEC command shows zeros for the counters associated with class-map match criteria. There is no workaround. (CSCec08205)
- When MPLS is enabled, traceroute is not supported. There is no workaround. (CSCec13655)
- The **priority** policy-map class configuration command cannot be configured for the default traffic class in a policy map. The workaround is to configure explicit matches for traffic that requires priority treatment. (CSCec38901)
- VLAN mappings can be configured on a per-interface basis. A different set of mappings can be configured on each ES interface. The per-interface VLAN mappings remain in effect even when the ES ports are bundled in an EtherChannel. For example, if you map Gigabit Ethernet 1/1/1 to VLAN 20 through VLAN 50 and Gigabit Ethernet 1/1/2 to VLAN 20 through VLAN 70, traffic on VLAN 20 leaving the switch through the ES port bundle should be load balanced across the individual ES interfaces. However, some of that traffic is incorrectly translated to VLAN 50, and some is incorrectly translated to VLAN 70. The workaround is to configure identical VLAN mappings on both ES ports if they are going to be bundled into an EtherChannel. (CSCec49520)
- Modifying a QoS class within a very large service policy that is attached to an ES port can cause high CPU utilization and an unresponsive CLI for an excessive period of time. The workaround is to detach the service policy from the port while making the modifications and then to re-attach the service policy. (CSCec75945)
- When packets are queued for egress on an ES port due to the application of a QoS service policy, they consume packet buffer memory on the switch. If many queues are simultaneously congested and are unable to drain, packet loss can occur in either direction (ingress or egress) due to the lack of buffer memory. If this becomes a problem, you can change switch behavior by using the **queue-limit** policy-map class configuration command at the class level to set shorter queue depths. Each shaper has an associated buffer queue with a default depth of 128 packets.

For example:

```
Switch(config)# policy-map cos2-policy
Switch(config-pmap)# class cos2
Switch(config-pmap-c)# bandwidth 50000
Switch(config-pmap-c)# queue-limit 32
```

The point at which buffer memory is exhausted depends on the number of queues, the sizes of the queued packets, and whether or not the traffic pattern being sent to the switch allows the queues to drain at all. (CSCed83886)

- The Catalyst 3750 Metro switch does not learn its own MAC address on Layer 2 interfaces. For example: Ports 1/0/1 and 1/0/2 belong to VLAN x, port 1/0/3 is a Layer 3 port with an IP address that belongs to the subnet of VLAN x, and ports 1/0/2 and 1/0/3 are connected. In this case, a host connected to port 1/0/1 cannot ping port 1/0/3. The switch does not update the CAM table and does not use the MAC address of port 1/0/3 in the CAM table for port 1/0/2.

The workaround is to statically configure the MAC address of port 1/0/3 in the CAM table of the switch bound to port 1/0/2 by using the **mac address-table static mac-addr vlan vlan-id interface fastethernet1/0/2** global configuration command. (CSCee87864)

Open Caveats

These are the open Cisco IOS configuration caveats:

- CSCdz30046

When multicast VLAN registration (MVR) groups are added or deleted, the receiver port that joined the groups after the addition still receives traffic even after the group is deleted. The correct behavior is that MVR data traffic to the group should stop flowing to the receiver port immediately after the **no mvr group ip-address** global configuration command is entered.

The workaround is to disable MVR by using the **no mvr** global configuration command and then to re-enable it by using the **mvr** command. Add and delete the groups that have problems by using the **mvr group ip-address** and the **no mvr group ip-address** global configuration commands.

- CSCea90131

Under these conditions, the switch might report a false security violation after an 802.1x supplicant is authenticated and assigned a new VLAN by the RADIUS server:

- 802.1x, port security, and voice VLAN are configured on a port.
- The maximum number of secure addresses has been learned on the port before it is authenticated.
- The VLAN assigned by the RADIUS server is different than the access VLAN configured on the port.

This problem does not prevent traffic from being forwarded to the 802.1x client, but the **show port-security** privileged EXEC command output might show that the port is *SecureDown* when it is actually *SecureUp* and forwarding traffic correctly.

The workaround is to restart the interfaces that appear to be out of sync by using the **shutdown** and then the **no shutdown** interface configuration commands.

- CSCeb10032
The switch might not be able to pass Vine (Advanced Research Projects Agency) ARPA frames over bridge groups.
The workaround is to use Subnetwork Access Protocol (SNAP) frames.
- CSCeb14406
DVMRP does not correctly forward packets.
There is no workaround.
- CSCeb29898
After starting up a switch that has more than 300 VLANs and the maximum number of static Etherchannel groups (12), all interfaces that are part of an Etherchannel might stay down. This occurs because the remote switch detects an Etherchannel misconfiguration and disables its ports. This problem can occur in either per-VLAN-spanning-tree plus (PVST+) or rapid-PVST+ mode.
The workaround is to restart the EtherChannel ports or to configure automatic recovery:
 - Use the **shutdown** and **no shutdown** interface configuration commands on the remote switch to restart all err-disabled interfaces.
 - Use the **errdisable recovery cause channel-misconfig** global configuration command to enable automatic link recovery on the remote switch, and use the **errdisable recovery interval** global configuration command to configure a short recovery interval.
- CSCeb35422
On a voice VLAN port with both 802.1x and port security enabled, dynamic secure addresses might not get deleted when the port is changed from multihost mode to single-host mode. This means that addresses learned in the multihost mode are still allowed after changing to single-host mode. This problem occurs under these conditions:
 - The port is in authorized state.
 - The port learns the MAC addresses of multiple hosts.
 - VLAN assignment is not enabled for the authorized host.
 The workaround is to disable and then re-enable port security on the port.
- CSCeb42949
The switch does not work with the User Registration Tool (URT). The PC attempting to connect to the network can login successfully, but is not allowed to pass traffic after the port is moved to the user VLAN. The MAC address for that device shows BLOCKED.
There is no workaround.
- CSCeb54159
If an interface on the switch is mapped to queue-set 2, and you disable and then re-enable multilayer QoS globally by using the **mls qos** global configuration command, the interface is no longer mapped to the correct egress queue-set.
The workaround is to reconfigure the interface queue-set by using the **no queue-set** interface configuration command followed by the **queue-set 2** interface configuration command.
- CSCeb56226
If an 802.1x port is configured for forced-unauthorized port control mode and voice VLAN, after you remove the voice VLAN and disable 802.1x on the port, the port no longer passes traffic.
The workaround is to restart the port by using the **shutdown** and then the **no shutdown** interface configuration commands.

- CSCec01607

When you configure a policy-rate policer and attach it to an interface, the cdQosPoliceCfg table shows two identical policers when only one has been configured. The cdQosPoliceStats table also contains status for two policers.

The workaround is to ignore the second policer.
- CSCec13135

The cbQosREDCfg, cbQosREDClassCfg and cbQosREDClassStats tables in CISCO-CLASS-BASED-QOS-MIB are unsupported.

The workaround is to use the **random-detect** policy-map class configuration command to configure Weighted Random Early Detection (WRED) and the **show policy-map interface** user EXEC command to monitor WRED.
- CSCec19825

When the receive rate is 100 Mbps and the sample interval (historyControlInterval) is more than 45 seconds, the calculation of the SNMP etherHistoryUtilization report is incorrect and shows a much lower utilization than expected. This also occurs at lower receive rates if the interval is long.

The workaround is to set the historyControlInterval to less than 30 seconds. Longer intervals can cause the counters to overflow if the traffic is intense.
- CSCec52524

When a switch boots without a configuration file, you are prompted for a *Yes* or *No* response to the `Continue with configuration dialog? [yes/no]` question. Even after you enter a response, pressing the Mode button for up to 2 seconds puts the switch in Express Setup mode and erases any configuration information that has been entered.

The workaround is to not enter Express Setup mode while working through the initial configuration dialog after configuration information has been entered.
- CSCec57743

The **no setup express** global configuration command does not appear in the CLI menu when you enter the **no ?** command.

The workaround is to enter the **no setup express** global configuration command even though it does not appear in the help.
- CSCec66730

The CLI allows an interface configured for 802.1x with the **dot1x port-control auto** interface configuration command to be added to a port-channel group, even though the features are mutually exclusive.

There is no workaround. Do not configure 802.1x on a port channel.
- CSCec70857

If you change the priority queue settings for ingress priority queue 2 by using the **mls qos srr-queue input priority-queue 2 bandwidth** global configuration command, the command is accepted correctly. However, the configuration resulting from the **show running-config** privileged EXEC command contains an extra **input**, for example, **mls qos srr-queue input priority-queue input 2 bandwidth**. This causes subsequent problems if the command is saved and if the switch is reloaded.

The workaround is to follow these steps:

 1. Copy the flash:config.text file to a PC or terminal.
 2. Edit the config.text file, and remove the extra **input** keyword.

3. Copy the file to the switch.
4. Reload the switch.

- CSCec71041

When automatic QoS (auto-QoS) is configured on an interface and that interface is changed from routed mode to switched mode or switched mode to routed mode, the trust policies displayed by the **show running-config** user EXEC command and the **show mls qos interface** user EXEC command are incorrect for the new interface type.

The workaround is to disable auto-QoS on the interface by using the **no auto qos voip [cisco-phone | trust]** interface configuration command, to change the interface to routed or switched mode, and then to reconfigure auto-QoS by using the **auto qos voip {cisco-phone | trust}** interface configuration command.

- CSCec72190

Changing the STP mode from PVST to MST (by using the **spanning-tree mode mst** global configuration command) or from MST to PVST (by using the **spanning-tree mode pvst** global configuration command) causes the LEDs for Layer 3 interfaces to turn amber, even though the ports are up.

The workaround is to use the **shutdown** and then the **no shutdown** interface configuration commands on each Layer 3 interface to force the LEDs back into sync.

- CSCec72935

The port bandwidth limit displayed by **show mls qos interface** user EXEC command is 100 minus the configured value. For example, if the configured value is 70, the display shows 30.

There is no workaround. This is a display issue only; the configured settings work correctly.

- CSCec73580

When the **switchport voice vlan {vlan-id | dot1p | none | untagged}** interface configuration command and the **spanning-tree bpduguard {disable | enable}** interface configuration command are configured together on an interface connected to another Catalyst 3750 Metro switch, the interface does not become err-disabled when BPDUs are detected (which is the expected action). If the voice VLAN configuration is removed, the BPDU guard works as expected.

There is no workaround.

- CSCec84254

The switch does not link up with some media converters running at 100 Mbps. This problem occurs on 10/100BASE-T interfaces.

There is no workaround.

- CSCed16780

Layer 2 Protocol tunneling does not function reliably on EtherChannel interfaces.

There is no workaround.

- CSCed23767

When you enter the **switchport port-security aging time** interface configuration command, the CLI help shows that entering 0 disables the aging time. The help is incorrect; a value of 0 is an invalid option.

The workaround to disable the aging time is to use the **no switchport port-security aging time** interface configuration command.

- CSCed26316
When an ES port is configured as an ISL trunk port, sending jumbo frames (1500 to 9000 bytes) through the ES port causes the connected link to receive fragments and cyclic redundancy check (CRC) errors on a high percentage of the traffic. After a prolonged traffic run, the ES ports might stop forwarding traffic.
The workaround is to use only 802.1Q encapsulation on ES ports that might carry jumbo frames
- CSCed27873
When an ES port is configured as an ISL trunk port, the interface counters and the **show interface** user EXEC command display count ISL packets that are within the system MTU size (1500 bytes) as giant packets that were discarded because they exceeded the system MTU.
The workaround is to configure the ES port as an 802.1Q trunk port.
- CSCed29932
When you change the MPLS router ID by entering the **mpls ldp router-id [loopback value] force** global configuration command, the local router ID is changed, but the label distribution protocol (LDP) does not bind with the LDP neighbor until the loopback interface is shut down and brought back up.
The workaround is to enter a **shutdown** and then a **no shutdown** interface configuration command on the loopback interface after you change the MPLS router ID.
- CSCed30184
An EtherChannel configured for 802.1Q tunneling and either the Port Aggregation Protocol (PAgP) or the Link Aggregation Control Protocol (LACP) does not come up.
The workaround is to use **channel-group channel-group-number mode on** interface configuration command.
- CSCee29107
If a switch has more than one route to reach a remote destination over an MPLS cloud, it does not support an EoMPLS configuration to that remote destination. This applies to both port-mode and VLAN-mode EoMPLS sessions. Trying to configure an EoMPLS session to that destination causes a Layer 2 loop for the EoMPLS session. Any traffic to an unknown destination address or multicast or broadcast traffic enters a permanent loop if both the ES ports (Gigabit Ethernet 1/1/1 and Gigabit Ethernet 1/1/2) are carrying the remote adjacency, either as a switch virtual interface (SVI) or a routed port. The Layer 2 loop is formed due to the internal implementation of EoMPLS.
There is no workaround. This configuration is not supported.
- CSCee30284
The EoMPLS tunneled ports (both port-based and VLAN-based) on a switch do not tunnel Layer 2 protocol packets that are received from the customer switch to the remote (egress) provider edge switch. The Layer 2 protocols packets are either consumed by the switch or, if Layer 2 protocol tunneling is enabled, are passed along by the standard Layer 2 protocol tunneling mechanism (outside the EoMPLS tunnels).
There is no workaround.
- CSCee33525
If you configure hundreds of QoS class maps, the switch might display a SYS-3-CPUHOG message when you apply the service policy to an interface by using the **service-policy** interface configuration command or when you remove the service policy.
There is no workaround.

- CSCee87630

When an ES interface on a switch is configured as a (Layer 2) switch port and MPLS is configured on a VLAN interface, changing the MPLS MTU on the VLAN interface is not effective.

The workaround is to change the MPLS MTU on the ES interface by entering the **mpls mtu bytes** interface configuration command.



Note When choosing an MTU size, use a value other than the default value. Choosing 1546 as an MTU size resolves the problem, but after a reload, the problem re-appears because the default values are not showing up in the startup configuration.

- CSCed70488

When an MPLS label is added to traffic leaving an ES interface, any Layer 3 “set” operations (for rewriting the IP precedence or DSCP) that should apply to that traffic are not performed.

There is no workaround

- CSCee22721

In order to make Intermediate System-to-Intermediate System (IS-IS) function on an SVI interface, you must manually set the Connectionless Network Service (CLNS) MTU to 1497.

The workaround, when configuring IS-IS on an SVI, is to use the **clns mtu 1497** interface configuration command on the SVI to set the CLNS MTU to 1497.

- CSCef27079

The switch reloads when you perform either of these CLI sequences:

Sequence A:

1. Attach an existing policy, *policy_a*, to an interface on the switch.
2. Rename this policy.
3. Create a new policy-map named *policy_a*.
4. Delete the policy-map *policy_a*.

Sequence B:

1. Attach an existing policy, *policy_a*, to an interface on the switch.
2. Enable Auto-QoS on the interface by entering the **auto qos voip cisco-phone** interface configuration command.
3. Rename the policy map *AutoQoS-Police-SoftPhone*.
4. Disable Auto-QoS on the interface by entering the **no auto qos voip cisco-phone** interface configuration command.
5. Enable Auto-QoS again on the same interface.

The workaround is that before deleting a renamed policy or disabling Auto-QoS on an interface to which a policy map is attached, detach the renamed policy from the interface by using the **no service-policy input** or **no service-policy output** interface configuration command.

Resolved Caveats

These are the caveats that have been resolved in this release.

- “Caveats Resolved in Cisco IOS Release 12.1(14)AX2” section on page 16
- “Caveats Resolved in Cisco IOS Release 12.1(14)AX1” section on page 19
- “Caveats Resolved in Cisco IOS Release 12.1(14)AX” section on page 20

Caveats Resolved in Cisco IOS Release 12.1(14)AX2

These caveats are resolved in Cisco IOS Release 12.1(14)AX2:

- CSCdz32659

The Cisco Discovery Protocol (CDP) process no longer causes memory-allocation failed messages.

- CSCec25430

A connection to a faulty Cisco 7935 IP phone no longer might cause the switch to reload.

- CSCed78149

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>.

- CSCee24347

When IS-IS is configured on an SVI, adjacency with its neighbor is now established.

- CSCee43832

A switch configured as an MPLS VPN provider-edge device no longer incorrectly displays traceback failure messages.

- CSCee45346
In an EoMPLS environment, the default mapping of COS to MPLS experimental bits now correctly copies the COS bits into both the outer Interior Gateway Protocol (IGP) label and the inner virtual connection (VC) label.
- CSCee48187
The switch no longer reloads when you modify a large policy map that is attached to an ES port.
- CSCee48384
The **no mpls ip propagate-ttl** global configuration command now works correctly on the switch.
- CSCee54095
The switch no longer reloads when you configure a policy map.
- CSCee60873
The switch no longer reloads when you configure Border Gateway Protocol (BGP) authentication.
- CSCee64947
Configuring an IP address on either interface VLAN 27 or interface VLAN 28 no longer causes the switch to reload.
- CSCee72865
The switch no longer reloads after you enter a **show policy-map interface *interface-id* output class *class-name*** user EXEC command.
- CSCee73194
When you configure IS-IS on an SVI the **ping clns** privileged EXEC command now works correctly.
- CSCee80101
The switch no longer reloads when you enter the **spanning-tree bpdudfilter enable** interface configuration command on a VLAN interface.
- CSCee82516
When the switch is configured as a service provider-edge (PE) device, clients attached to a routed interface can now ping across the MPLS cloud when all VLANs are removed and added back on the uplink trunk.
- CSCee89222
When VPN routing and forwarding (VRF) is configured, if an external devices attempts to ping a loopback interface, the switch now correctly responds to the ping.
- CSCee90338
When PVST+ or MSTP is enabled and the spanning tree root VLAN ages out, a transient loop no longer occurs when the network reconverges.
- CSCee92228
When two Cisco-approved SFP modules are inserted in the switch at the same time, both SFPs are now recognized by the switch.
- CSCee96113
QoS interface statistics now correctly increment on ES ports.

- CSCef02264
Removing the priority action from a class within an egress service-policy on an ES interface by using the **no priority** policy-map class configuration command causes the service policy to become invalid; however, the switch now properly reverts to the last valid configuration.
- CSCef04425
CPU usage no longer increases to 99 per cent every 5 seconds when a large policy map is attached to an ES port.
- CSCef05460
The switch no longer shuts down if you enter a **duplex** command in interface-range configuration mode for a range of Fast Ethernet interfaces that belong to a port channel.
- CSCef12796
When MPLS is enabled, the switch now correctly adds the MPLS labels so that the packets are correctly processed by the next-hop routers.
- CSCef13260
When a DSCP mutation map is applied to an ES interface and then removed, the CoS value of the packets are now correct.
- CSCef15180
An EoMPLS connection no longer might stop working if an access port in the EoMPLS VLAN changes administrative state.
- CSCef20483
If a large quantity of traffic is policy-based routed on the switch, this no longer causes an increase in the CPU load.
- CSCef26094
When a switch is configured for MPLS VPN and the ARP entry of the next hop ages out or is manually cleared, packets coming in on a VRF interface are no longer sent out untagged. MPLS VPN site connectivity is not broken.
- CSCef26204
The switch no longer stops tagging packets to the default route learned on a VRF interface when it receives a route update (delete or learn) from one of its BGP neighbors for that VRF.
- CSCef31332
When a switch is used as a pure Layer 2 switch with MPLS disabled, and ES ports are used to connect to the provider-edge routers, this no longer breaks MPLS VPN connectivity.
- CSCef38562
QoS policy-maps now work correctly when you change the policer from a bit rate (**police cir** policy-map class configuration command) to a percentage (**police cir percent** policy-map class configuration command) or from a percentage to a bit rate.
- CSCin67568
The switch now correctly handles CDP packets with hostnames that are longer than 255 characters.

Caveats Resolved in Cisco IOS Release 12.1(14)AX1

These caveats are resolved in Cisco IOS Release 12.1(14)AX1:

- CSCeb78921
When MPLS traffic has been running long enough for the MPLS byte counters to roll over, entering the **show mpls forwarding-table** user EXEC command no longer results in a display with large negative bytes.
- CSCec46189
A switch running Routing Information Protocol (RIP) version 1 no longer discards RIP packets destined for directed broadcast addresses.
- CSCec55073
If you paste a configuration that has a large ACL into the running configuration of the switch, the console no longer halts, and the switch no longer reboots.
- CSCec57826
A switch configured to send DHCP unicast requests to a specified DHCP server no longer reloads when it tries to boot up.
- CSCec62437
Address Resolution Protocol (ARP) responses destined to the switch CPU that are received on a Layer 2 interface that has a MAC access list applied to it are no longer dropped, and ARP learning works correctly.
- CSCec69183
The switch no longer reloads when you add an aggregate policer to a new policy map. In previous releases, this happened if the aggregate policer had previously been used in a policy map that was attached to an interface, even if that policy map had been detached and removed from the interface.
- CSCec76671
The switch no longer times out the source and multicast group address (S, G) entries when passing low traffic (such as 3 packets per minute) from the source.
- CSCec84210
The switch now communicates correctly with media converters running at 10 Mbps.
- CSCec86127
The spanning-tree algorithm now blocks a Layer 2 loop if you change the native VLAN or a trunk port to a VLAN that you have not yet created.
- CSCec86621
When you enter the **ip default-network** global configuration command, all packets that match the default route are no longer sent to the CPU.
- CSCec87974
When using SNMP to poll a switch, the ifHCInOctets and ifHCOutOctets are no longer 0 for Gigabit EtherChannel interfaces.
- CSCed03214
The switch no longer pauses indefinitely when 1-byte frames are received.

- CSCed06621
The dynamic MAC address of the Hot Standby Router Protocol (HSRP) group is now relearned on the standby switch even if several interfaces have the same HSRP standby group.
- CSCed10210
The switch no longer allows Telnet sessions to the device from unauthorized hosts when you apply an access class to inbound vty lines.
- CSCed11323
If there are multiple aggregate policers configured on a switch and one of the policers is used in a policy map that has been applied to an interface, the switch no longer fails if you remove the aggregate policer without first detaching it from the policy map. In previous releases, this occurred when you first applied the command or after you saved the configuration and then reloaded the switch.
- CSCed36621
When the MPLS label range is set to the default, the switch no longer rejects MPLS traffic with a label greater than 8192.
- CSCed47290
MPLS now functions correctly on switch virtual interfaces (SVIs).
- CSCed71197
Converting a Layer 2 port to a Layer 3 port by using the **no switchport** interface command no longer causes a remote ACL that is applied to the switch SVI to be removed from the ternary content addressable memory (TCAM).
- CSCed75115
You can now set the CLNS MTU to more than 1497 on a routed interface.
- CSCee14600
Layer 2 protocol tunneling packets received through an ES interface are now forwarded correctly.
- CSCee23626
When the switch is running an Any Transport over MPLS (AToM) process and an LDP session starts or stops, the switch no longer experiences a slow memory leak or has to be reloaded.

Caveats Resolved in Cisco IOS Release 12.1(14)AX

These caveats were resolved in release 12.1(14)AX:

- CSCdu53656
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, please refer to this advisory, available at <http://www.cisco.com/warp/public/707/cisco-sa-20040616-bgp.shtml>.

- CSCea28131

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, please refer to this advisory, available at <http://www.cisco.com/warp/public/707/cisco-sa-20040616-bgp.shtml>.

- CSCed27956

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>.

- CSCed38527

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>.

Documentation Updates

With this release, the switch supports new QoS features (see the “[New Software Features](#)” section on [page 5](#)) that were not described in the software documentation. These are the documentation updates for this release.

Software Configuration Guide

These are the documentation updates for the *Catalyst 3750 Metro Switch Software Configuration Guide*.



Note

In Cisco IOS Release 12.1(14)AX1, the switch supported point-to-point Layer 2 protocol tunneling, which was not documented in the Cisco IOS Release 12.1(14)AX software documentation. This information is in the *Release Notes for the Catalyst 3750 Metro Switch, Cisco IOS Release 12.1(14) AX1* at this URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750m/12114ax/ol464602.htm#wp44273>

This information is part of Chapter 26, “Configuring QoS.” For the complete chapter (minus these updates), go to this URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750m/12114ax/3750msg/swqos.htm>

Configuring the Egress Priority Queue



Note

This is a new section, not previously in the “Configuring QoS” chapter.

Beginning in Cisco IOS Release 12.1(14)AX2, you can ensure that certain packets have priority over all others by queuing them in the egress priority queue on a port. SRR services this queue until it is empty and before servicing the other queues.

All four queues participate in the SRR unless the egress priority queue is enabled, in which case the first bandwidth weight is ignored and is not used in the ratio calculation. The priority queue is serviced until empty before the other queues are serviced. You enable the priority queue by using the **priority-queue out** interface configuration command.

Follow these guidelines when the egress priority queue is enabled on a port; otherwise, the egress queues are serviced based on their SRR weights:

- If the egress priority queue is enabled, it overrides the SRR shaped and shared weights for queue 1.
- If the egress priority queue is disabled and the SRR shaped and shared weights are configured, the shaped mode overrides the shared mode for queue 1, and SRR services this queue in shaped mode.
- If the egress priority queue is disabled and the SRR shaped weights are not configured, SRR services this queue in shared mode.

On an ES port, you can use LLQ (enabled with the **priority** policy-map class configuration command) and the egress priority queue (enabled with the **priority-queue out** interface configuration command). By using these two features, you can give priority to a class of traffic and avoid losing traffic when the switch is congested. In previous releases (before the egress priority queue was supported), you could put a traffic class into the strict-priority queue, but congestion at the egress queue-sets could result in the

dropping of that priority traffic. The **priority-queue out** interface configuration command enables you to prioritize the same traffic class at the egress queue-sets, ensuring that priority traffic reaches the hierarchical queues and is processed with priority.

Beginning in privileged EXEC mode, follow these steps to enable the egress priority queue. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify a port, and enter interface configuration mode.
Step 3	priority-queue out	Enable the egress priority queue, which is disabled by default. When you configure this command, the SRR weight and queue size ratios are affected because there is one fewer queue participating in SRR. This means that <i>weight1</i> in the srr-queue bandwidth shape or the srr-queue bandwidth share command is not used in the ratio calculation.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the egress priority queue, use the **no priority-queue out** interface configuration command.

This example shows how to enable the egress priority queue when the SRR weights are configured. The egress expedite queue overrides the configured SRR weights.

```
Switch(config)# interface gigabitethernet1/0/1
Switch(config-if)# srr-queue bandwidth shape 25 0 0 0
Switch(config-if)# srr-queue bandwidth share 30 20 25 25
Switch(config-if)# priority-queue out
Switch(config-if)# end
```

Enabling DSCP Transparency Mode



Note

This is a new section, not previously in the “Configuring QoS” chapter.

By default, in Cisco IOS Release 12.1(14)AX2 or later, the DSCP transparency feature is disabled, and the DSCP value of an incoming packet is modified based on the QoS configuration. If you enable the DSCP transparency feature, the switch does not change (rewrite) the DSCP field of an IP packet at the ingress and egress interfaces. During QoS processing, the switch modifies the CoS value of the packet based on the internal DSCP value and DSCP-to-CoS map.

In software releases earlier than Cisco IOS Release 12.1(14)AX2, if QoS is disabled, the DSCP value of the incoming IP packet is not modified (the default). If QoS is enabled and you configure the interface to trust DSCP, the switch does not modify the DSCP value. If you configure the interface to trust CoS, the switch modifies the DSCP value according to the CoS-to-DSCP map.

If you enable DSCP transparency, these automatic actions occur:

- Even though the QoS configuration, including the DSCP input and output queue threshold maps, affects the internal DSCP value of the packet, the DSCP value of the packet at the ingress and the egress is the same.



Note Any “match” statements configured within an output service policy applied to an ES interface will match the original DSCP value, rather than the internal representation.



Note If an action in an output service policy that is applied to an ES interface modifies the DSCP value, that modification takes place even when DSCP transparency is enabled.

- If you apply the DSCP-to-DSCP mutation map to an interface, the internal DSCP value changes according to the mutation map.

Beginning in privileged EXEC mode, follow these steps to enable DSCP transparency on a switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS globally.
Step 3	no mls qos rewrite ip dscp	Enable DSCP transparency. The switch is configured to not modify the DSCP field of the IP packet.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface [<i>interface-id</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To configure the switch to modify the DSCP value based on the trust setting or on an ACL by disabling DSCP transparency, use the **mls qos rewrite ip dscp** global configuration command.

If you disable QoS by using the **no mls qos** global configuration command, the CoS and DSCP values are not changed (the default QoS setting).

If you enter the **no mls qos rewrite ip dscp** global configuration command to enable DSCP transparency and then enter the **mls qos trust [cos | dscp]** interface configuration command, DSCP transparency is disabled.

Command Reference

In the *Catalyst 3750 Metro Switch Command Reference*, these commands are new or have been modified to support QoS functionality.

- [mls qos rewrite ip dscp](#), page 25
- [priority-queue](#), page 27

mls qos rewrite ip dscp



Note

This is a new command, not previously in the command reference.

Use the **mls qos rewrite ip dscp** global configuration command to configure the switch to change (rewrite) the Differentiated Services Code Point (DSCP) field of an incoming IP packet. Use the **no** form of this command to configure the switch to not modify the DSCP field of the packet.

mls qos rewrite ip dscp

no mls qos rewrite ip dscp

Syntax Description

This command has no arguments or keywords.

Defaults

DSCP transparency is disabled. The switch changes the DSCP field of the incoming IP packet.

Command Modes

Global configuration

Command History

Release	Modification
12.1(14)AX2	This command was introduced.

Usage Guidelines

Use the **mls qos rewrite ip dscp** command to disable the DSCP transparency feature. The DSCP value of an incoming packet is modified based on the QoS configuration, such the port trust setting, policing and marking, and the DSCP-to-DSCP mutation map.

If you disable DSCP transparency and then enable QoS globally, these automatic actions occur:

- The switch does not modify the DSCP field in the packet.
- If policing is enabled and the packet is out of profile, the switch changes the internal DSCP value according to the policed-DSCP map and uses the modified internal DSCP value for queueing.
- The switch uses the internal DSCP value of the packet to represent the priority of the traffic and to generate a class of service (CoS) value.

To enable the DSCP transparency feature, you *must* use the **no mls qos rewrite ip dscp** global configuration command. When DSCP transparency is enabled, the switch does not change (rewrite) the DSCP field of an IP packet at the ingress and egress interfaces. During QoS processing, the switch modifies the CoS value of the packet based on the internal DSCP value and DSCP-to-CoS map.

If you enable DSCP transparency, these automatic actions occur:

- Even though the QoS configuration, including the DSCP input and output queue threshold maps, affects the internal DSCP value of the packet, the DSCP value of the packet at egress is the same as the DSCP value at the ingress.
- If you also apply the DSCP-to-DSCP mutation map to an interface, the internal DSCP value changes according to the mutation map.

If you disable QoS, the CoS and DSCP values of the incoming IP packet do not change (the default).

Examples

This example shows how to enable DSCP transparency and configure the switch to not change the DSCP value of the incoming IP packet:

```
Switch(config)# mls qos
Switch(config)# no mls qos rewrite ip dscp
Switch(config)# end
```

This example shows how to disable DSCP transparency and configure the switch to change the DSCP value of the incoming IP packet:

```
Switch(config)# mls qos
Switch(config)# mls qos rewrite ip dscp
Switch(config)# end
```

You can verify your settings by entering the **show mls qos interface** *[interface-id]* privileged EXEC command.

Related Commands

Command	Description
mls qos	Enables QoS globally.
show mls qos interface	Displays QoS information at an interface level.

priority-queue



Note

This is a new command, not previously in the command reference.

Use the **priority-queue** interface configuration command to enable the egress priority queue on a port. Use the **no** form of this command to return to the default setting.

priority-queue out

no priority-queue out

Syntax Description

out	Enable the egress priority queue.
------------	-----------------------------------

Defaults

The egress priority queue is disabled.

Command Modes

Interface configuration

Command History

Release	Modification
12.1(14)AX2	This command was introduced.

Usage Guidelines

When you configure the **priority-queue out** command, the shaped round robin (SRR) weight ratios are affected because there is one fewer queue participating in SRR. This means that *weight1* in the **srr-queue bandwidth shape** or the **srr-queue bandwidth share** interface configuration command is ignored (not used in the ratio calculation). Before servicing the other queues, SRR services the priority queue until it is empty.

Follow these guidelines when the priority queue is enabled; otherwise, the egress queues are serviced based on their SRR weights:

- If the egress priority queue is enabled, it overrides the SRR shaped and shared weights for queue 1.
- If the egress priority queue is disabled and the SRR shaped and shared weights are configured, the shaped mode overrides the shared mode for queue 1, and SRR services this queue in shaped mode.
- If the egress priority queue is disabled and the SRR shaped weights are not configured, SRR services the queue in shared mode.

To avoid a loss of priority traffic on an enhanced-services (ES) port when the switch is congested, you can configure both the egress priority queueing and the strict-priority queueing (low-latency queueing [LLQ]) features. Note that when you map certain traffic classes to the egress priority queue, they are not automatically placed into the strict-priority queue. You must configure both the **priority-queue out** interface configuration and the **priority** policy-map class configuration commands.

Examples

This example shows how to enable the egress priority queue when the SRR weights are configured. The egress expedite queue overrides the configured SRR weights.

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# srr-queue bandwidth shape 25 0 0 0
Switch(config-if)# srr-queue bandwidth share 30 20 25 25
Switch(config-if)# priority-queue out
```

This example shows how to disable the egress priority queue after the SRR shaped and shared weights are configured. The shaped mode overrides the shared mode.

```
Switch(config)# interface gigabitethernet1/0/2
Switch(config-if)# srr-queue bandwidth shape 25 0 0 0
Switch(config-if)# srr-queue bandwidth share 30 20 25 25
Switch(config-if)# no priority-queue out
```

You can verify your settings by entering the **show mls qos interface *interface-id* queueing** or the **show running-config** privileged EXEC command.

Related Commands

Command	Description
show mls qos interface queueing	Displays the queueing strategy (SRR, priority queueing), the weights corresponding to the queues, and the CoS-to-egress-queue map.
srr-queue bandwidth shape	Assigns the shaped weights and enables bandwidth shaping on the four egress queues mapped to a port.
srr-queue bandwidth share	Assigns the shared weights and enables bandwidth sharing on the four egress queues mapped to a port.

Hardware Installation Guide

The Preface for the *Catalyst 3750 Metro Switch Hardware Installation Guide* does not include the translations for the Warning symbol and explanation (Statement 1071) or a change to the Warning statement about installation for short-circuit (overcurrent) protection (Statement 1005-Circuit Breaker) in Appendix E, “Translated Safety Warnings.”

This information is in the *Release Notes for the Catalyst 3750 Metro Switch, Cisco IOS Release 12.1(14)AX* at this URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750m/12114ax/ol464601.htm#35851>

Related Documentation

These documents provide information about the switch and are available from this Cisco.com site:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat3750m/index.htm>

You can order printed copies of documents with a DOC-xxxxxx= number from the Cisco.com sites and from the telephone numbers listed in the “Obtaining Documentation” section on page 29.

- *Catalyst 3750 Metro Switch Software Configuration Guide* (order number DOC-7815870=)
- *Catalyst 3750 Metro Switch Command Reference* (order number DOC-7815871=)
- *Catalyst 3750 Metro Switch System Message Guide* (order number DOC-7815872=)
- *Catalyst 3750 Metro Switch Hardware Installation Guide* (order number DOC-7815869=)
- *Cisco Small Form-Factor Pluggable Modules Installation Notes* (not orderable but available on Cisco.com)

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_inpk/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 408 526-7208 or, elsewhere in North America, by calling 800 553-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 this URL:
<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.

CCSP, the Cisco Square Bridge logo, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0406R)

Copyright © 2003 - 2005 Cisco Systems, Inc. All rights reserved..