



Release Notes for the Catalyst 3550 Multilayer Switch Cisco IOS Release 12.1(9)EA1

April 26, 2002

The Cisco IOS Release 12.1(9)EA1 runs on all Catalyst 3550 multilayer switches.

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

- If you are installing a new switch, refer to the IOS release label on the rear panel of your switch.
- If your switch is on, use the **show version** privileged EXEC command. See the [“Determining the Software Version and Feature Set”](#) section on page 7.
- If you are upgrading to a new release, refer to the software upgrade filename for the IOS version.

For the complete list of Catalyst 3550 switch documentation, see the [“Related Documentation”](#) section on page 32.

This IOS 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 IOS releases become available, they will be posted to Cisco.com (previously Cisco Connection Online [CCO]) in the Cisco IOS software area.

Contents

This information is in the release notes:

- [“System Requirements”](#) section on page 2
- [“Downloading Software”](#) section on page 6
- [“Installation Notes”](#) section on page 9
- [“New Features”](#) section on page 14
- [“Limitations and Restrictions”](#) section on page 16
- [“Important Notes”](#) section on page 22



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

Copyright © 2002. Cisco Systems, Inc. All rights reserved.

- [“Open Caveats” section on page 23](#)
- [“Resolved Caveats” section on page 29](#)
- [“Documentation Updates” section on page 31](#)
- [“Related Documentation” section on page 32](#)
- [“Obtaining Documentation” section on page 32](#)
- [“Obtaining Technical Assistance” section on page 33](#)

System Requirements

These are the system requirements for this IOS release:

- [“Hardware Supported” section on page 2](#)
- [“Software Compatibility” section on page 3](#)

Hardware Supported

[Table 1](#) lists the hardware supported by this IOS release.

Table 1 *Supported Hardware*

Switch	Description
Catalyst 3550-12T	10 Gigabit Ethernet 10/100/1000BASE-T ports and 2 Gigabit Interface Converter (GBIC)-based Gigabit Ethernet slots
Catalyst 3550-12G	10 GBIC-based Gigabit Ethernet slots and 2 10/100/1000BASE-T ports
Catalyst 3550-24	24 autosensing 10/100 Ethernet ports and 2 GBIC-based Gigabit Ethernet slots
Catalyst 3550-48	48 autosensing 10/100 Ethernet ports and 2 GBIC-based Gigabit Ethernet slots
Catalyst 3550-24-DC	24 autosensing 10/100 Ethernet ports, 2 GBIC-based Gigabit Ethernet slots, and an on-board direct-current (DC) power converter
GBIC modules	<ul style="list-style-type: none"> • 1000BASE-SX GBIC • 1000BASE-LX/LH GBIC • 1000BASE-ZX GBIC • 1000BASE-T GBIC • GigaStack GBIC • Course Wave Division Multiplexer (CWDM) fiber-optic GBIC
Redundant power system	Cisco RPS 300 Redundant Power System

Software Compatibility

These are the software compatibility requirements for this IOS release:

- “Recommended Platform Configuration for Web-Based Management” section on page 3
- “Operating System and Browser Support” section on page 3
- “Installing the Required Plug-In” section on page 4
- “Creating Clusters with Different Releases of IOS Software” section on page 5

Recommended Platform Configuration for Web-Based Management

Table 2 lists the recommended platforms for Web-based management.

Table 2 Recommended Platform Configuration for Web-Based Management

OS	Processor Speed	DRAM	Number of Colors	Resolution	Font Size
Windows NT 4.0 ¹	Pentium 300 MHz	128 MB	65,536	1024 x 768	Small
Solaris 2.5.1	SPARC 333 MHz	128 MB	Most colors for applications	—	Small (3)

1. Service Pack 3 or higher is required.

The minimum PC requirement is a Pentium processor running at 233 MHz with 64 MB of DRAM. The minimum UNIX workstation requirement is a Sun Ultra 1 running at 143 MHz with 64 MB of DRAM.

For information about supported operating systems, see the next section.

Operating System and Browser Support

You can access the web-based interfaces by using the operating systems and browsers listed in Table 3. The switch checks the browser version when starting a session to ensure that the browser is supported. If the browser is not supported, the switch displays an error message, and the session does not start.

Table 3 Supported Operating Systems and Browsers

Operating System	Minimum Service Pack or Patch	Netscape Communicator ¹	Microsoft Internet Explorer ²
Windows 95	Service Pack 1	4.61, 4.7x	4.01a, 5.0, 5.5
Windows 98	Second Edition	4.61, 4.7x	4.01a, 5.0, 5.5
Windows NT 4.0	Service Pack 3 or later	4.61, 4.7x	4.01a, 5.0, 5.5
Windows 2000	None	4.61, 4.7x	4.01a, 5.0, 5.5
Solaris 2.5.1 or later	Sun-recommended patch cluster for the OS and Motif library patch 103461-24	4.61, 4.7x	Not supported

1. Netscape Communicator versions 4.60 and 6.0 are not supported.

2. Service Pack 1 or higher is required for Internet Explorer 5.5.

**Note**

If your browser is Internet Explorer and you receive an error message stating that the page might not display correctly because your security settings prohibit running activeX controls, this might mean that your security settings are set too high. To lower security settings, go to **Tools > Internet Options**, and select the **Security** tab. Select the indicated **Zone**, and move the **Security Level for this Zone** slider from **High to Medium** (the default).

**Note**

In Cluster Management displays, Internet Explorer versions 4.01 and 5.0 might not display edge devices that are not connected to the command switch. Other functionality is similar to that of Netscape Communicator.

Installing the Required Plug-In

A Java plug-in is required for the browser to access and run the Java-based Cluster Management Suite (CMS). Download and install the plug-in before you start CMS. Each platform, Windows and Solaris, supports three plug-in versions. For information on the supported plug-ins, see the “[Windows 2000, Windows 95, Windows 98, and Windows NT 4.0 Platforms](#)” section on page 5 and the “[Solaris Platforms](#)” section on page 5.

You can download the recommended plug-ins from this URL:
<http://www.cisco.com/cgi-bin/tablebuild.pl/java>

**Note**

Uninstall older versions of the Java plug-ins before installing the Java plug-in.

If the Java applet does not initialize after you have installed the plug-in, open the Java Plug-in Control Panel (**Start > Programs > Java Plug-in Control Panel**), and verify these settings:

In the Proxies tab, verify that **Use browser settings** is checked and that no proxies are enabled.

**Note**

If you are running an Internet virus checker on Windows 2000 and the plug-in takes a long time to load, you can speed up CMS operation by disabling the virus checker filter option or download option or both.

On McAfee VirusScan, from the Start menu, to disable the VirusScan Internet Filter option, the Download Scan option, or both, select **Start > Programs > Network Associates > Virus Scan Console > Configure**.

or

From the taskbar, right-click the Virus Shield icon and in the Quick Enable menu, disable the options by deselecting **Internet Filter** or **Download Scan**.

Windows 2000, Windows 95, Windows 98, and Windows NT 4.0 Platforms

These Java plug-ins are supported on the Windows platform:

- Java plug-in 1.3.1
- Java plug-in 1.3.0
- Java plug-in 1.2.2_05

You can download these plug-ins from this URL:

<http://www.cisco.com/pcgi-bin/tablebuild.pl/java>



Note

If you start CMS without having installed the required Java plug-in, the browser automatically detects this. If you are using a supported Internet Explorer browser, it automatically downloads and installs the Java plug-in 1.3.0 (default). If you are using a supported Netscape browser, the browser displays a Cisco.com page that contains the Java plug-in and installation instructions. If you are using Windows 2000, Netscape Communicator might not detect the missing Java plug-in.

Solaris Platforms



Caution

These Java plug-ins are supported on the Solaris platform:

To avoid performance and compatibility issues, do not use Java plug-ins later than Java plug-in 1.3.1.

- Java plug-in 1.2.2_07
- Java plug-in 1.3.0
- Java plug-in 1.3.1

You can download these plug-ins and instructions from this URL:

<http://www.cisco.com/pcgi-bin/tablebuild.pl/java>

To install the Java plug-in, follow the instructions in the README_FIRST.txt file.

Creating Clusters with Different Releases of IOS Software

When a cluster consists of a mixture of other Catalyst switches, we strongly recommend using only the Catalyst 3550 switches as the command and standby command switches. When the command switch is a Catalyst 3550 switch, all standby command switches must also be Catalyst 3550 switches. The Catalyst 3550 switch that has the latest software should be the command switch. If the command switch is a Catalyst 3550 Gigabit Ethernet switch and the standby command switch is a Catalyst 3550 Fast Ethernet switch, command switch port speeds are reduced if the standby command switch takes over.

If your cluster has Catalyst 2950, Catalyst 2900 XL, and Catalyst 3500 XL switches, the Catalyst 2950 switch (with the latest software release) should be the command switch. The Catalyst 2950 switch that has the latest software should be the command switch.

If your switch cluster has Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, and Catalyst 3500 XL switches, either the Catalyst 2900 XL or Catalyst 3500 XL (whichever has the latest software release) should be the command switch.

Table 4 lists the cluster capabilities and software versions for the switches.

Table 4 *Switch Software and Cluster Capability*

Switch	IOS Release	Cluster Capability
Catalyst 3550	Release 12.1(4)EA1 or later	Member or command switch
Catalyst 3500 XL	Release 12.0(5.1)XU or later	Member or command switch
Catalyst 2950	Release 12.0(5.2)WC(1) or later	Member or command switch
Catalyst 2900 XL (8-MB switches)	Release 12.0(5.1)XU or later	Member or command switch
Catalyst 2900 XL (4-MB switches)	Release 11.2(8.5)SA6 (recommended)	Member switch only
Catalyst 1900 and 2820	Release 9.00(-A or -EN) or later	Member switch only

Some versions of the Catalyst 2900 XL software do not support clustering and if you have a cluster with switches that are running different versions of IOS software, software features added on the latest release might not be reflected on switches running the older versions. For example, if you start Visual Switch Manager (VSM) on a Catalyst 2900 XL switch running Release 11.2(8)SA6, the windows and functionality can be different from a switch running Release 12.0(5)WC(1) or later.



Note

The CMS is not forward-compatible, which means that if a member switch is running a software version that is newer than the release running on the command switch, the new features are not available on the member switch. If the member switch is a new device supported by a software release that is later than the software release on the command switch, the command switch cannot recognize the member switch and it is displayed as an unknown device in the Front Panel view. You cannot configure any parameters or generate a report through CMS for that member; instead, you must launch the Device Manager application to perform configuration and obtain reports for that member.

Downloading Software

These are the procedures for downloading software:

- [“Determining the Software Version and Feature Set” section on page 7](#)
- [“Which Files to Use” section on page 7](#)
- [“Upgrading a Switch by Using CMS” section on page 7](#)
- [“Upgrading a Switch by Using the CLI” section on page 8](#)



Note

Before downloading software, read this section for important information.

Determining the Software Version and Feature Set

The IOS image is stored as a *.bin* file in a directory that is named with the IOS release. A subdirectory contains the HTML files needed for web management. 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. In the display, check the line that begins with *System image file is*. It shows the directory name in Flash memory where the image is stored. A couple of lines below the image name, you see *Running Layer 2/3 Switching Image* if you are running the enhanced multilayer software image, or *Running Layer 2 Switching Image Only* if you are running the standard multilayer software image.



Note

Although the **show version** output always shows the software image running on the switch (Layer 2 or Layer 2/3), the model name shown at the end of this display is the factory configuration (SMI or EMI) and does not change if you upgrade the software image.

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.

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 both the IOS image file and the HTML files (needed for the CMS). You must use the combined *.tar* file to upgrade the switch through the CMS.

The *.tar* file is an archive file from which you can extract files by using the **tar** command. You also use the *.tar* file to upgrade the system by using the **archive download-sw** privileged EXEC command.

Table 5 lists the software file names for this IOS release.

Table 5 Cisco IOS Software Files for Catalyst 3550 Switches

Filename	Description
c3550-i9q3l2-tar.121-9.EA1.tar	IOS image file and HTML files This image, the standard multilayer software image (SMI), has Layer 2+ features only.
c3550-i5q3l2-tar.121-9.EA1.tar	IOS image file and HTML files This image, the enhanced multilayer software image (EMI), has both Layer 2 and Layer 3 features.



Note

All Catalyst 3550 Gigabit Ethernet switches ship with the enhanced multilayer software image (EMI) installed. This image is an orderable upgrade for Catalyst 3550 Fast Ethernet switches with the standard multilayer software image (SMI) pre-installed.

Upgrading a Switch by Using CMS

You can upgrade switch software by using CMS. From the menu bar, select **Administration > Software Upgrade**. For detailed instructions, click **Help**.

Upgrading a Switch by Using the CLI

This procedure is for copying the combined .tar file to the Catalyst 3550 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.

To download software, and if necessary, the TFTP server application, follow these steps:

-
- Step 1** Use [Table 5 on page 7](#) to identify the file that you want to download.
- Step 2** Download the software image file.
- 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 SMI and EMI files, select **Download Cisco Catalyst 3550 software**.

- Step 3** Download the Cisco TFTP server from the URL link from Step 2, if necessary. The information on this page describes how to download and configure the TFTP server.
- Step 4** Copy the image to the appropriate TFTP directory on the workstation, and make sure the TFTP server is properly configured.
- For more information, refer to Appendix B in the *Catalyst 3550 Multilayer Switch Software Configuration Guide*.
- Step 5** Log in to the switch through the console port or a Telnet session.
- Step 6** 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 7** 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 172.20.129.10 and to overwrite the image on the switch:

```
Switch# archive download-sw /overwrite tftp://172.20.129.10/c3550-i5q312-tar.121-9.EA1.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.

Upgrading with a Non-Default System MTU Setting

If the switch was running Release 12.1(8)EA1c or earlier and you had used the **system mtu** global configuration command to configure a non-default system MTU size on your switch, follow these steps to upgrade your switch to 12.1(9)EA1:

-
- Step 1** Upgrade the IOS software to 12.1(9)EA1.
- Step 2** If a system MTU size of greater than 2000 is configured on the Catalyst 3550-12T or Catalyst 3550-12G, use the **system mtu** global configuration command to set it to the maximum supported MTU size.
-  **Note** The maximum allowable system MTU for Catalyst 3550 Gigabit Ethernet switches is 2000 bytes; the maximum system MTU for Fast Ethernet switches is 1546 bytes.
-
- Step 3** Save the running configuration by entering the **copy running-config startup-config** privileged EXEC command.
- Step 4** Reload the switch using the new IOS software.
- Step 5** When the switch comes back up with 12.1(9)EA1, reload the switch a second time by using the **reload** privileged EXEC command so that the **system mtu** command takes effect.
-

Recovering from Software Failure

In the software fails, you can reload the software. For detailed recovery procedures, refer to the “Troubleshooting” chapter in the *Catalyst 3550 Multilayer Switch Software Configuration Guide*.

Installation Notes

You can assign IP information to your switch by using the setup program, the Dynamic Host Configuration Protocol (DHCP)-based autoconfiguration (refer to the *Catalyst 3550 Multilayer Switch Software Configuration Guide*), or by manually assigning an IP address (refer to the *Catalyst 3550 Multilayer Switch Software Configuration Guide*).

These are the installation procedures:

- [“Setting Up the Catalyst 3550 Initial Configuration” section on page 10](#)
- [“Configuring Browsers and Accessing CMS” section on page 12](#)

Setting Up the Catalyst 3550 Initial Configuration

The first time that you access the switch, it runs a setup program that prompts you for an IP address and other configuration information necessary for the switch to communicate with the local routers and the Internet. This information is also required if you plan to use the CMS to configure and manage the switch.



Note

If the switch will be a cluster member managed through the IP address of the command switch, it is not necessary to assign IP information or a password. If you are configuring the switch as a standalone switch or as a command switch, you must assign IP information.

Follow these steps to create an initial configuration for the switch:

Step 1 Enter **Yes** at the first two prompts.

```
Would you like to enter the initial configuration dialog? [yes/no]: yes
```

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '['].

```
Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system.
```

```
Would you like to enter basic management setup? [yes/no]: yes
```

Step 2 Enter a host name for the switch, and press **Return**.

On a command switch, the host name is limited to 28 characters; on a member switch to 31 characters. Do not use *-n*, where *n* is a number, as the last character in a host name for any switch.

```
Enter host name [Switch]: host_name
```

Step 3 Enter a secret password, and press **Return**.

The password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, allows spaces, but ignores leading spaces.

```
Enter enable secret: secret_password
```

Step 4 Enter an enable password, and press **Return**.

```
Enter enable password: enable_password
```

Step 5 Enter a virtual terminal (Telnet) password, and press **Return**.

The password can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

```
Enter virtual terminal password: terminal-password
```

Step 6 (Optional) Configure Simple Network Management Protocol (SNMP) by responding to the prompts.

Step 7 Enter the interface name (physical interface or VLAN name) of the interface that connects to the management network, and press **Return**. For this release, always use **vlan 1** as that interface.

```
Enter interface name used to connect to the
management network from the above interface summary: vlan 1
```

- Step 8** Configure the interface by entering the switch IP address and subnet mask and pressing **Return**:

```
Configuring interface vlan 1:
Configure IP on this interface? [yes]: yes
IP address for this interface: 10.4.120.106
Subnet mask for this interface [255.0.0.0]: 255.255.255.0
```

- Step 9** Enter **Y** to configure the switch as the cluster command switch. Enter **N** to configure it as a member switch or as a standalone switch.

If you enter **N**, the switch appears as a candidate switch in the CMS. In this case, the message in [Step 10](#) is not displayed.

```
Would you like to enable as a cluster command switch? [yes/no]: yes
```

- Step 10** Assign a name to the cluster, and press **Return**.

```
Enter cluster name: cluster_name
```

The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.

The initial configuration appears:

The following configuration command script was created:

```
hostname host-name
enable secret 5 $1$LiBw$0XclwyT.PXPkuhFwqyhVi0
enable password enable-password
line vty 0 15
password terminal-password
snmp-server community public
!
no ip routing
!
interface vlan 1
no shutdown
ip address 10.4.120.106 255.255.255.0

interface GigabitEthernet0/1
no ip address
!
interface GigabitEthernet0/2
no ip address
!
...<output abbreviated>
!
interface GigabitEthernet0/12
no ip address

cluster enable cluster-name
!
end
```

Step 11 These choices are displayed:

```
[0] Go to the IOS command prompt without saving this config.
```

```
[1] Return back to the setup without saving this config.
```

```
[2] Save this configuration to nvram and exit.
```

```
Enter your selection [2]:2
```

Make your selection, and press **Return**.

After you complete the setup program, the switch can run the created default configuration. If you want to change this configuration or want to perform other management tasks, use one of these tools:

- Command-line interface (CLI)
- Cluster Management Suite (CMS) from your browser

Configuring Browsers and Accessing CMS

For the browser to use CMS, a Java plug-in is required, as described in the [“Installing the Required Plug-In” section on page 4](#). After you have assigned an IP address to the switch and installed the plug-in, you can access the switch from your browser and use the CMS to configure other switches. To use the web-based tools, see the [“Software Compatibility” section on page 3](#) to set up the appropriate browser options.

These are the installation procedures:

- [“Configuring Netscape Communicator \(All Versions\)” section on page 12](#)
- [“Configuring Microsoft Internet Explorer \(4.01\)” section on page 13](#)
- [“Configuring Microsoft Internet Explorer \(5.0\)” section on page 13](#)
- [“Displaying the CMS Access Page” section on page 14](#)

Configuring Netscape Communicator (All Versions)

Follow these steps to configure Netscape Communicator:

- Step 1** Start Netscape Communicator.
 - Step 2** From the menu bar, select **Edit > Preferences**.
 - Step 3** In the Preferences window, click **Advanced**.
 - Step 4** Check the **Enable Java**, **Enable JavaScript**, and **Enable Style Sheets** check boxes.
 - Step 5** From the menu bar, select **Edit > Preferences**.
 - Step 6** In the Preferences window, click **Advanced Cache**, and select **Every time**.
 - Step 7** Click **OK** to return to the browser Home page.
-

Configuring Microsoft Internet Explorer (4.01)

Follow these steps to configure Microsoft Internet Explorer 4.01:

-
- Step 1** Start Internet Explorer.
 - Step 2** From the menu bar, select **View > Internet Options**.
 - Step 3** In the Internet Options window, click the **Advanced** tab.
 - a. Scroll through the list of options until you see Java VM. Check the **Java logging enabled** and **Java JIT compiler enabled** check boxes.
 - b. Click **Apply**.
 - Step 4** In the Internet Options window, click the **General** tab.
 - a. In the Temporary Internet Files section, click **Settings**.
 - b. In the Settings window, select **Every visit to the page**, and click **OK**.
-

Configuring Microsoft Internet Explorer (5.0)



Note

During the installation of this browser, make sure to check the **Install Minimal or Customize Your Browser** check box. In the Component Options window in the Internet Explorer 5 section, make sure to check the **Microsoft Virtual Machine** check box to display applets written in Java.

Follow these steps to configure Microsoft Internet Explorer 5.0:

-
- Step 1** Start Internet Explorer.
 - Step 2** From the menu bar, select **Tools > Internet Options**.
 - Step 3** In the Internet Options window, click the **Advanced** tab.
 - a. Scroll through the list of options until you see Java VM. Check the **Java logging enabled** and **JIT compiler for virtual machine enabled** check boxes.
 - b. Click **Apply**.
 - Step 4** In the Internet Options window, click the **General** tab.
 - a. In the Temporary Internet Files section, click **Settings**.
 - b. In the Settings window, select **Every visit to the page**, and click **OK**.
-

If you are using Microsoft Internet Explorer 5.0 to make configuration changes to the switch, note that this browser does not automatically reflect the latest configuration changes. Make sure that you click **Refresh** for every configuration change.

Displaying the CMS Access Page

After the browser is configured, display the CMS access page:

-
- Step 1** Enter the switch IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Internet Explorer), and press **Return**.
- Step 2** Enter your username and password when prompted. The password provides level 15 access.



Note The browser always prompts for username and password. If no username is configured on your switch, you only need to enter the enable password in the appropriate field.

The Cisco Systems Access page appears. For more information on setting passwords and privilege levels, refer to the *Catalyst 3550 Multilayer Switch Software Configuration Guide*.

- Step 3** Click **Web Console** to launch the CMS applet.
- If you access CMS from a standalone or a cluster-member switch, Device Manager appears.
-

New Features

These are the new supported hardware and the new software features provided in IOS Release 12.1(9)EA1:

- [“New Hardware Features” section on page 14](#)
- [“New Software Features” section on page 14](#)

New Hardware Features

For a list of supported hardware, see the [“Hardware Supported” section on page 2](#).

New Software Features

Cisco IOS Release 12.1(9)EA1 contains these new features or enhancements:

- Cisco Intelligence Engine 2100 (IE2100) Series Cisco Networking Services (CNS) embedded agents for automating switch management, configuration storage, and delivery
- Voice VLAN for creating subnets for voice traffic from Cisco IP Phones
- IEEE 802.1Q tunneling so that customers with users at remote sites across a service provider network can keep VLANs segregated from those of other customers
- Layer 2 protocol tunneling to ensure that customers across a service provider network have complete STP, CDP, and VTP information about all users
- IEEE 802.1S Multiple Spanning Tree Protocol (MSTP) for grouping VLANs into a spanning-tree instance and for providing multiple forwarding paths for data traffic and load balancing
- IEEE 802.1W Rapid Spanning Tree Protocol (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the forwarding state

- DHCP relay agent information (option 82) for subscriber identification and IP address management
- Port security aging to set the aging time for secure addresses on a port
- Support for VLAN IDs in the full 1 to 4094 range allowed by the IEEE 802.1Q standard, and introduction of a new configuration mode (config-VLAN)
- Support for standard and extended IP access control lists (ACLs) and MAC extended ACLs for defining security policies on inbound Layer 2 interfaces (port ACLs)
- Order-dependent ACL merge (ODM)

The software uses an algorithm that produces an order-dependent set of merged ternary content addressable memory (TCAM) entries while processing ACLs used for security filtering and for some kinds of quality of service (QoS) classification. This algorithm does not result in any significant operational differences for configurations using the old algorithm that fit into the hardware. However, many configurations that previously did not fit now fit in the hardware because of the new merge algorithm. This algorithm is enabled by default and is not configurable.

The Catalyst 3550 hardware has only one input security lookup, one QoS lookup, and one output security lookup for each frame. The result of each of these lookups must yield the net result of all the configured features of that type. For example, if both a VLAN map and an input router ACL are configured, the packet is forwarded normally (and possibly routed) only if both the VLAN map and the router ACL allow forwarding. With ODM, the access list entries (ACEs) from the separate ACLs in the VLAN map and the router ACL are merged into a single unified set of entries that are loaded into the TCAM. A single TCAM lookup returns the net result of matching both the VLAN map and the input router ACL.

The ODM algorithm works by intersecting the ACLs, ACE by ACE, in order, until it reaches an ACE with a result that determines the final disposition of any frame that matches that series of ACEs. Two ACEs are said to intersect if a single packet could match both ACEs. The intersection of two ACEs is an ACE that matches all the packets that could match both original ACEs and only matches those packets. When a final disposition is reached or when the last ACL in the last feature is reached, any remaining ACLs are ignored, and a TCAM entry is generated. The entry describes the intersection of all the ACEs traversed up to that point and the final result of matching those ACEs.

After the merge is completed, the final result is order-dependent because it preserves the order dependence of the ACEs in the original ACLs.

- CMS support for these features:
 - ACLs on switch ports
 - Voice VLANs
 - Extended VLAN IDs (VLAN IDs 1006 to 4094)
 - CWDM GBIC modules

Limitations and Restrictions

You should review this section before you begin working with the switches. 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.

These are the limitations and restrictions:

- [“IOS Limitations and Restrictions” section on page 16](#)
- [“Cluster Limitations and Restrictions” section on page 20](#)
- [“Cluster Management Suite Limitations and Restrictions” section on page 21](#)

IOS Limitations and Restrictions

These limitations apply to IOS configuration:

- Storm control or traffic suppression (configured by using the **storm-control { broadcast | multicast | unicast } level level [.level]** interface configuration command) is supported only on physical interfaces; it is not supported on EtherChannel port channels even though you can enter these commands through the CLI.
- The Cisco RPS 300 Redundant Power System supports the Catalyst 3550 multilayer switch and provides redundancy for up to six connected devices until one of these devices requires backup power. If a connected device has a power failure, the RPS immediately begins supplying power to that device and sends status information to other connected devices that it is no longer available as a backup power source. As described in the device documentation, when the RPS LED is amber, the RPS is connected but down. However, this might merely mean that the RPS is in standby mode. Press the **Standby/Active** button on the RPS to put it into active mode. You can view RPS status through the CLI by using the **show rps** privileged EXEC command. For more information, refer to the *RPS 300 Hardware Installation Guide*.
- You can connect the switch to a PC by using the switch console port and the supplied rollover cable and the DB-9 adapter. You need to provide a RJ-45-to-DB-25 female DTE adapter if you want to connect the switch console port to a terminal. You can order a kit (part number ACS-DSBUASYN=) with this RJ-45-to-DB-25 female DTE adapter from Cisco.
- Modifying a multicast boundary access list does not prevent packets from being forwarded by any multicast routes that were in existence before the access list was modified if the packets arriving on the input interface do not violate the boundary. However, no new multicast routes that violate the updated version of the multicast boundary access list are learned, and any multicast routes that are in violation of the updated access list are not relearned if they age out.
After updating a multicast boundary, the workaround is to use the **clear ip mroute** privileged EXEC command to delete any existing multicast routes that violate the updated boundary. (CSCdr79083)
- When an IP packet with a cyclic redundancy check (CRC) error is received, the per-packet per-DSCP counter (for DSCP 0) is incremented. Normal networks should not have packets with CRC errors. (CSCdr85898)
- The **mac-address** interface configuration command does not properly assign a MAC address to an interface. This command is not supported on Catalyst 3550 switches. (CSCds11328)

- If you configure the DHCP server to allocate addresses from a pool to the switch, two devices on the network might have the same IP address. Pooled addresses are temporarily allocated to a device and are returned to the pool when not in use. If you save the configuration file after the switch receives such an address, the pooled address is saved, and the switch does not attempt to access the DHCP server after a reboot to receive a new IP address. As a result, two devices might have the same IP address.

The workaround is to make sure that you configure the DHCP server with reserved leases that are bound to each switch by the switch hardware address. (CSCds55220)

- The **show ip mroute count** privileged EXEC command might display incorrect packet counts. In certain transient states (for example, when a multicast stream is forwarded only to the CPU during the route-learning process and the CPU is programming this route into the hardware), a multicast stream packet count might be counted twice. Do not trust the counter during this transient state. (CSCds61396)
- When changing the link speed of a Gigabit Ethernet port from 1000 Mbps to 100 Mbps, there is a slight chance that the port will stop transmitting packets. If this occurs, shut down the port, and re-enable it by using the **shutdown** and **no shutdown** interface configuration commands. (CSCds84279)
- In IP multicast routing and fallback bridging, certain hardware features are used to replicate packets for the different VLANs of an outgoing trunk port. If the incoming speed is line rate, the outgoing interface cannot duplicate that speed (because of the replication of the packets). As a result, certain replicated packets are dropped. (CSCdt06418)
- When you use the **no interface port-channel** global configuration command to remove an EtherChannel group, the ports in the port group change to the administratively down state.
When you remove an EtherChannel group, enter the **no shutdown** interface configuration command on the interfaces that belonged to the port group to bring them back on line. (CSCdt10825)
- In the output displayed after a **show interface interface-id** privileged EXEC command, the *output buffer failures* field shows the number of packets lost before replication, whereas the *packets output* field shows the successful transmitted packets after replication. To determine actual discarded frames, multiply the output buffer failures by the number of VLANs on which the multicast data is replicated. (CSCdt26928)
- Internet Group Management Protocol (IGMP) packets classified by quality of service (QoS) to map the Differentiated Service Code Point (DSCP) value and the class of service (CoS) value in a QoS policy map might only modify the DSCP property and leave the CoS value at zero. (CSCdt27705)
- If you assign both tail-drop threshold percentages to 100 percent by using the **wrr-queue threshold** interface configuration command and display QoS information for this interface by using the **show mls qos interface statistics** privileged command, the drop-count statistics are always zero even if the thresholds were exceeded. To display the total number of discarded packets, use the **show controllers ethernet-controllers interface-id** privileged EXEC command. In the display, the number of discarded frames includes the frames that were dropped when the tail-drop thresholds were exceeded. (CSCdt29703)
- Open Shortest Path First (OSPF) path costs and Interior Gateway Routing Protocol (IGRP) metrics are incorrect for switch virtual interface (SVI) ports. You can manually configure the bandwidth of the SVI by using the **bandwidth** interface configuration command. Changing the bandwidth of the interface changes the routing metric for the routes when the SVI is used as an outgoing interface. (CSCdt29806)
- On the Catalyst 3550, coldStart and warmStart traps are not consistently sent. (CSCdt33779)
- Remote Monitoring (RMON) collection functions on physical interfaces, but it is not supported on EtherChannels and SVIs. (CSCdt36101)

- Multicast router information is displayed in the **show ip igmp snooping mrouter** privileged EXEC command when IGMP snooping is disabled. Multicast VLAN Registration (MVR) and IGMP snooping use the same commands to display multicast router information. In this case, MVR is enabled, and IGMP snooping is disabled. (CSCdt48002)
- When a VLAN interface has been disabled and restarted multiple times by using the **shutdown** and **no shutdown** interface configuration commands, the interface might not restart following a **no shutdown** command. To restart the interface, re-enter a **shutdown** and **no shutdown** command sequence. (CSCdt54435)
- When you configure the **ip pim spt-threshold infinity** interface configuration command, you want all sources for the specified group to use the shared tree and not use the source tree. However, the switch does not automatically start to use the shared tree. No connectivity problem occurs, but the switch continues to use the shortest path tree for multicast group entries already installed in the multicast routing table. You can enter the **clear ip mroute *** privileged EXEC command to force the change to the shared tree. (CSCdt60412)
- If the number of multicast routes configured on the switch is greater than the switch can support, it might run out of available memory, which can cause it to reboot. This is a limitation in the platform-independent code.

The workaround is to not configure the switch to operate with more than the maximum number of supported multicast routes. You can use the **show sdm prefer** and **show sdm prefer routing** privileged EXEC commands to view approximate maximum configuration guidelines for the current SDM template and the routing template. ((CSCdt63354))

- Configuring too many multicast groups might result in an extremely low memory condition and cause the software control data structure to go out of sync, causing unpredictable forwarding behavior. The memory resources can only be recovered by issuing the **clear ip mroute** privileged EXEC command. To prevent this situation, do not configure more than the recommended multicast routes on the switch. (CSCdt63480)
- The **dec** keyword is not supported in the **bridge bridge-group protocol** global configuration command. If two Catalyst 3550 switches are connected to each other through an interface that is configured for IP routing and fallback bridging, and the bridge group is configured with the **bridge bridge-group protocol dec** command, both switches act as if they were the spanning tree root. Therefore, spanning-tree loops might be undetected. (CSCdt63589)
- When you configure an EtherChannel between a Catalyst 3550 and a Catalyst 1900 switch, some of Catalyst 3550 links in the EtherChannel might go down, but one link in the channel remains up, and connectivity is maintained.

The workaround is to disable the Port Aggregation Protocol (PAgP) on both devices by using the **channel-group channel-group-number mode on** interface configuration command. PAgP negotiation between these two devices is not reliable. (CSCdt78727)

- When the switch is operating with equal-cost routes and it is required to learn more unicast routes than it can support, the CPU might run out of memory, and the switch might fail.

The workaround is to remain within the documented recommended and supported limits. (CSCdt79172)

- The behavior of a software ACL with QoS is different from a hardware ACL with QoS. On the Catalyst 3550 switch, when the QoS hardware rewrites the DSCP of a packet, the rewriting of this field happens before software running on the CPU examines the packet, and the CPU sees only the new value and not the original DSCP value.

When the security hardware ACL matches a packet on input, the match uses the original DSCP value. For output security ACLs, the security ACL hardware should match against the final, possibly changed, DSCP value as set by the QoS hardware. Under some circumstances, a match to a security ACL in hardware prevents the QoS hardware from rewriting the DSCP and causes the CPU to use the original DSCP.

If a security ACL is applied in software (because the ACL did not fit into hardware, and packets were sent to the CPU for examination), the match probably uses the new DSCP value as determined by the QoS hardware, regardless of whether the ACL is applied at the input or at the output. When packets are logged by the ACL, this problem can also affect whether or not a match is logged by the CPU even if the ACL fits into hardware and the permit or deny filtering was completed in hardware.

To avoid these issues, whenever the switch rewrites the DSCP of any packet to a value different from the original DSCP, security ACLs should not test against DSCP values in any of their access control elements (ACEs), regardless of whether the ACL is being applied to an IP access group or to a VLAN map. This restriction does not apply to ACLs used in QoS class maps.

If the switch is not configured to rewrite the DSCP value of any packet, it is safe to match against DSCP in ACLs used for IP access groups or for VLAN maps because the DSCP does not change as the packet is processed by the switch.

The DSCP field of an IP packet encompasses the two fields that were originally designated precedence and TOS (type of service). Statements relating to DSCP apply equally to either IP precedence or IP TOS. (CSCdt94355)

- Disabling autonegotiation on a GBIC interface by using the **speed nonegotiate** interface configuration command might cause the interface to show that the physical link is up, even when it is not connected. (CSCdv29722)
- If you configure a trunk port for Dynamic Trunking Protocol (DTP) nonegotiate mode and change the encapsulation type from ISL to 802.1Q by using the **switchport trunk encapsulation** interface configuration command, the port becomes an access port and is no longer trunking. (CSCdv46715)
- On earlier versions of Catalyst 3550-24 switches, if a 10/100BASE-TX port on the switch is connected to a Catalyst 2820 or Catalyst 1900 switch through an ISL trunk at 100 Mbps, bidirectional communication cannot be established. The Catalyst 2820 or Catalyst 1900 switch identifies the Catalyst 3550-24 switch as a CDP neighbor, but the Catalyst 3550-24 switch does not recognize the Catalyst 2820 or Catalyst 1900 switch. On these switches, you should not use ISL trunks between the Catalyst 3550-24 and a Catalyst 2820 or Catalyst 1900 switch. Configure the link as an access link instead of a trunk link.

This problem has been fixed in hardware on Catalyst 3550-24 switches with motherboard assembly number 73-5700-08 or later. To determine the board level on your switch, enter the **show version** privileged EXEC. Motherboard information appears toward the end of the output display. (CSCdv68158)

- When IGMP filtering is enabled and you use the **ip igmp profile** global configuration command to create an IGMP filter, reserved multicast addresses cannot be filtered. Because IGMP filtering uses only Layer 3 addresses to filter IGMP reports and due to mapping between Layer 3 multicast addresses and Ethernet multicast addresses, reserved groups (224.0.0.x) are always allowed through the switch. In addition, aliased groups can leak through the switch. For example, if a user is allowed to receive reports from group 225.1.2.3, but not from group 230.1.2.3, aliasing will cause the user to receive reports from 230.1.2.3. Aliasing of reserved addresses means that all groups of the form y.0.0.x are allowed through. (CSCdv73626)

If you use the **ip igmp max-groups** interface configuration command to set the maximum number of IGMP groups for an interface to 0, the port still receives group reports from reserved multicast groups (224.0.0.x) and their Layer 2 aliases (y.0.0.x). (CSCdv79832)

- The switch might reload when it is executing the **no snmp-server host** global configuration command. This is a rare condition that can happen if SNMP traps or informs are enabled and the SNMP agent attempts to send a trap to the host just as it is being removed from the configuration and if the IP address of the host (or the gateway to reach the host) has not been resolved by Address Resolution Protocol (ARP).

The workaround is to ensure that the target host or the next-hop gateway to that host is in the ARP cache (for example, by issuing a **ping** command) before removing it from the SNMP configuration. Alternatively, disable all SNMP traps and informs before removing any hosts from the SNMP configuration. (CSCdw44266)

- When you access CISCO-STACK-MIB portTable, the mapping might be off by one from the mapping given by the switch. The objects in this table are indexed by two numbers: portModuleIndex and portIndex. The allowable values for portModuleIndex are 1 through 16. Because 0 is not an allowable value, the value 1 represents module 0.

The workaround is to use the value 1 to represent module 0. (CSCdw71848)

- If a port on the Catalyst 3550 switch that is running the Multiple Spanning Tree Protocol (MSTP) is connected to another switch that belongs to a different multiple spanning tree (MST) region, the Catalyst 3550 port is not recognized as a boundary port when you start the protocol migration process by using the **clear spanning-tree detected-protocols interface interface-id** privileged EXEC command. This problem occurs only on the root bridge, and when the root bridge is cleared, the boundary ports are not shown because the designated ports do not receive any bridge protocol data units (BPDUs) unless a topology change occurs. This is the intended behavior.

The workaround is to configure the Catalyst 3550 switch for PVST by using the **spanning-tree mode pvst** global configuration command bridge, and then change it to MSTP by using the **spanning-tree mode mst** global configuration command. (CSCdx10808)

- If you apply an ACL to an interface that has a QoS policy map attached and the ACL is configured so that the packet should be forwarded by the CPU or if the configured ACL cannot fit into the TCAM, all packets received from this interface are forwarded to the CPU. Because traffic forwarded to the CPU cannot be policed by the policer configured on the interface, this traffic is not accurately rate-limited to the configured police rate.

The workaround, when QoS rate limiting is configured on an interface, is to configure applied ACLs so that packets are not forwarded by the CPU or reduce the number of ACEs in the ACL so that it can fit into the TCAM. (CSCdx30485)

- Catalyst 3550 switches do not take into account the Preamble and Inter Frame Gap (IFG) when rate limiting traffic, which could result in a slightly inaccurate policing rate on a long burst of small-sized frames, where the ratio of the Preamble and IFG to frame size is more significant. This should not be an issue in an environment where the frames are a mix of different sizes.

Cluster Limitations and Restrictions

These limitations apply to cluster configuration:

- When there is a transition from the cluster active command switch to the standby command switch, Catalyst 1900, Catalyst 2820, and Catalyst 2900 4-MB switches that are cluster members might lose their cluster configuration. You must manually add these switches back to the cluster. (CSCds32517, CSCds44529, CSCds55711, CSCds55787, CSCdt70872)

- When a Catalyst 2900 XL or Catalyst 3500 XL cluster command switch is connected to a Catalyst 3550 switch, the command switch does not find any cluster candidates beyond the Catalyst 3550 switch if it is not a member of the cluster. You must add the Catalyst 3550 switch to the cluster. You can then see any cluster candidates connected to it. (CSCdt09918)
- When clustering is enabled, do not configure SNMP community strings of more than 59 bytes, or clustering SNMP might not work correctly. (CSCdt39616)
- If both the active command-switch and the standby command switch fail at the same time, the cluster is not automatically recreated. Even if there is a third passive command switch, it might not recreate all cluster members because it might not have all the latest cluster configuration information. You must manually recreate the cluster if both the active and standby command switches simultaneously fail. (CSCdt43501)

Cluster Management Suite Limitations and Restrictions

These limitations apply to CMS configuration:

- Host names and Domain Name System (DNS) server names that contain commas on a cluster command switch, member switch, or candidate switch can cause CMS to behave unexpectedly. You can avoid this instability in the interface by not using commas in host names or DNS names. Do not enter commas when also entering multiple DNS names in the IP Configuration tab of the IP Management window in CMS.
- ACEs that contain the **host** keyword precede all other ACEs in standard ACLs. You can reposition the ACEs in a standard ACL with one restriction: No ACE with the **any** keyword or a wildcard mask can precede an ACE with the **host** keyword.
- CMS performance degrades if the Topology View is open for several hours on a Solaris machine. The cause might be a memory leak.
The workaround is to close the browser, reopen it, and launch CMS again. (CSCds29230)
- If you are printing a Topology View or Front Panel View that contains many devices and are running Solaris 2.6 with JDK1.2.2, you might get an *Out of Memory* error message.
The workaround is to close the browser, re-open it, and launch CMS again. Before you perform any other task, bring up the view that you want to print, and click **Print** in the **CMS** menu. (CSCds80920)
- If a PC running CMS has low memory and CMS is running continuously for two to three days, the PC runs out of memory.
The workaround is to relaunch CMS. (CSCdv88724)
- When a VLAN or a range of VLANs is already configured and you specify VLAN filter for a SPAN session, the current configuration for that session is overwritten with the new entry. Although the CLI appends new entries after the existing ones, CMS recreates the whole session, overwrites the current entry, and provides only a single VLAN filter per entry.
The workaround is to use the CLI; it is the only method for specifying multiple VLANs for filtering in a SPAN session. (CSCdw93904)

Important Notes

These are the important notes related to this IOS release:

- [“IOS Notes” section on page 22](#)
- [“Cluster Notes” section on page 22](#)
- [“Cluster Management Suite Notes” section on page 23](#)

IOS Notes

These notes apply to IOS configuration:

- If you configure a port ACL on a physical interface on a switch that has VLAN maps or input router ACLs configured, or if you configure a VLAN map or input router ACL on a switch that has port ACLs configured, a *CONFLICT* message is generated but the configuration is accepted. The port ACL action has priority on that port over actions in a router ACL or VLAN map applied to the VLAN to which the port belongs.

The result is that packets received on that physical port will be permitted or denied based on the port ACL action without regard to any permit or deny statements in any router ACL or VLAN map, while packets received on other physical ports in the VLAN will still be permitted or denied based on any router ACLs or VLAN maps applied to the VLAN. If the port ACL is applied to a trunk port, it overrides any other input ACLs applied to all VLANs on the trunk port.

- The default system MTU for traffic on the Catalyst 3550 switch is 1500 bytes. The 802.1Q tunneling feature increases the frame size by 4 bytes. Therefore, when you configure 802.1Q tunneling, you must configure all switches in the 802.1Q network to be able to process maximum frames by increasing the switch system MTU size to at least 1504 bytes. You configure the system MTU size by using the **system mtu** global configuration command.
- Beginning with IOS release 12.1(8)EA1, to configure traffic suppression (previously configured by using the **switchport broadcast**, **switchport multicast**, and **switchport unicast** interface configuration commands), you use the **storm-control {broadcast | multicast | unicast} level level [.level]** interface configuration commands. For more information about these commands, refer to the *Catalyst 3550 Multilayer Switch Command Reference*.
- When you are configuring a cascaded stack of Catalyst 3550 switches by using the GigaStack GBIC and want to include more than one VLAN in the stack, be sure to configure all the GigaStack GBIC interfaces as trunk ports by using the **switchport mode trunk** interface configuration command and to use the same encapsulation method by using the **switchport encapsulation {isl | dot1q}** interface configuration command. For more information about these commands, refer to the *Catalyst 3550 Multilayer Switch Command Reference*.
- If the 1000BASE-T GBIC (WS-G5482) is not securely inserted, the switch might fail to recognize it or might display an incorrect media type following a **show interface** privileged EXEC command entry. If this happens, remove and reinsert the GBIC.

Cluster Notes

This note applies to cluster configuration:

The **cluster setup** privileged EXEC command and the **standby mac-address** interface configuration command have been removed from the CLI and the documentation because they did not function correctly.

Cluster Management Suite Notes

These notes apply to CMS configuration:

- If you use CMS on Windows 2000, it might not apply configuration changes if the enable password is changed from the CLI during your CMS session. You have to restart CMS and enter the new password when prompted. Platforms other than Windows 2000 prompt you for the new enable password when it is changed.
- CMS does not display QoS classes that are created through the CLI if these classes have multiple match statements. When using CMS, you cannot create classes that match more than one match statement. CMS does not display policies that have such classes.
- If you use Internet Explorer Version 5.5 and select a URL with a nonstandard port at the end of the address (for example, *www.add.com:84*), you must enter *http://* as the URL prefix. Otherwise, you cannot launch CMS.
- Within an ACL, you can change the sequence of ACEs that have the **host** keyword. However, because such ACEs are independent of each other, the change has no effect on the way the ACL filters traffic.
- If you use the Netscape browser to view the CMS GUI and you resize the browser window while CMS is initializing, CMS does not resize to fit the window.

Resize the browser window again when CMS is not busy.

- CMS does not start if the temporary directory on your computer runs out of memory. This problem can occur because of a bug in the 1.2.2 version of the Java plug-in. The plug-in creates temporary files in the directory whenever it runs CMS, and the directory eventually runs out of plug-in space.

The workaround is to remove all the *jar_cache*.tmp* files from the temporary directory. The path to the directory is different for different operating systems:

```
Solaris: /var/tmp
Windows NT and Windows 2000: \TEMP
Windows 95 and 98: \Windows\Temp
```

Open Caveats

These are the open caveats with possible unexpected activity in this IOS release:

- [“Open IOS Caveats” section on page 24](#)
- [“Open Cluster Caveats” section on page 28](#)
- [“Open Cluster Management Suite Caveats” section on page 28](#)

Open IOS Caveats

These are the severity 3 IOS configuration caveats:

- CSCdt48614

When applied to routed ports, configurations for these two keywords are not properly retained after a reboot:

- the **tunnel** keyword relating to IP multicast routing, such as the **tunnel number**, **tunnel source ip-address**, and **tunnel destination ip-address** interface configuration commands
- the HSRP **track** keyword in the **standby group-number track type number [interface-priority]** interface configuration commands.

There is no workaround.

- CSCdv66568

After you change connections between GigaStack ports, the link might not be established, and LEDs on the GigaStack GBIC might continue to blink for more than 2 minutes. (It is normal for the LEDs to blink for a short time.)

The workaround is to disconnect and reconnect one of the links connected to the GigaStack GBIC with the continuous blinking LEDs.

- CSCdv79737

If a stack contains both Catalyst 3550 switches and Catalyst 3500 XL or Catalyst 2900 XL switches, cross-stack UplinkFast does not function if the management VLAN on the Catalyst 3500 XL or Catalyst 2900 XL switches is changed to other than VLAN 1 (the default).

The workaround is to make sure that the management VLAN of all Catalyst 3500 XL or 2900 XL switches in the stack is set to VLAN 1.

- CSCdw27519

Multicast data might be temporarily lost when a link comes up in a redundant network and causes the reverse path forwarding (RPF) to change. This only occurs when there are multiple paths between the rendezvous point (RP) and the multicast source. If the RP loses link state on the incoming interface, it quickly fails-over to a different interface. However, if the original interface comes up again, data may be lost for about 1 minute because PIM hello packets are being dropped by the RP while the interface is coming up. If link is not lost (for example, the RP port is connected to a hub, or the path is interrupted elsewhere), the data loss does not occur.

- CSCdx05364

When you use the **ip-access group** or **mac-access group** interface configuration command to apply a port ACL to a physical Layer 2 interface that is a member of an EtherChannel, the command is accepted even though switch does not support port ACLs on an EtherChannel or an interface belonging to an EtherChannel.

There is no workaround.

- CSCdx00558

When you enter the **clear adjacency** privileged EXEC command to clear the adjacency table and the **clear ip route *** privileged EXEC command to remove all routing table entries, a SYS-3-CPUHOG error message might appear.

The workaround is to selectively clear the routing table by using the **clear ip route network [mask]** privileged command. When the specific route is cleared, the adjacencies of that route are also cleared.

- CSCdx06694

If you configure the Catalyst 3550 switch with multiple SVIs and an IP address and the VLAN Membership Policy Server (VMPS) server does not have routes configured to reach all the subnets on the switch, the VMPS might not assign VLANs to the switch dynamic-access ports. This is because the switch randomly selects one of the SVI IP addresses (rather than the one configured on the VLAN interface that is used to reach the VMPS) when it sends VLAN Query Protocol (VQP) requests to the VMPS server. The VQP server responds with the VLAN assignment only if it knows how to reach the SVI IP address used in the VQP request.

The workaround is to configure the VMPS server so that it has routes to reach all subnets on the switch.

- CSCdx07308

If two separate HSRP groups are misconfigured with the same standby IP address, an ARP storm occurs, affected interfaces have poor or no connections, and warnings appear on the console (typically *%IP-4-DUPADDR* and *%SYS-2-MALLOCFAIL*).

The workaround is to avoid misconfiguring two HSRP groups with the same standby IP address. Configure different addresses for the two groups.

- CSCdx17189

If the number of VLANs configured on a switch is close to the maximum (1005) and you have configured a large number of trunk ports with an allowed VLAN list, when you use the **no switchport trunk allowed vlan** interface range command to remove the allowed list for all trunk ports on the switch, you might see a *SYS-3-CPUHOG* message.

The workaround is to not use the **interface range** command to remove the allowed list for all trunk ports. Instead, you should enter the individual commands in interface configuration mode for each trunk port.

- CSCdx19540

If the switch fails for any reason while you are exiting VLAN configuration mode (accessed by entering the **vlan database** privileged EXEC command), there is a slight chance that the VLAN database might get corrupted. After resetting from the switch, you might see these messages on the console:

```
%SW_VLAN-4-VTP_INVALID_DATABASE_DATA: VLAN manager received bad data of type device
type: value 0 from vtp database
```

```
$SW_VLAN-3-VTP_PROTOCOL_ERROR: VTP protocol code internal error
```

The workaround is to use the **delete flash:vlan.dat** privileged EXEC command to delete the corrupted VLAN database. Then reload the switch by using the **reload** privileged EXEC command.

- CSCdx20106

When 1000 VLANs and more than 40 trunk ports are configured, and the spanning-tree mode changes from MSTP to PVST or vice versa, this message appears on the console:

```
%ETHCNTR-3-RA_ALLOC_ERROR: RAM Access write pool I/O memory allocation failure
```

There is no workaround. However, we recommend that you reload the switch by using the **reload** privileged EXEC command. To avoid this problem, configure the system with fewer VLANs and fewer trunk ports, or use the **switchport trunk allowed vlan** interface configuration command to reduce the number of active VLANs on each trunk port.

- CSCdx20421

When you attempt an SNMP GetNext operation to retrieve the value of the `cIgmppFilterEditSpinLock` object in the CISCO-IGMP-FILTER-MIB, the object ID (OID) returned by the switch is not an IGMP Filter Editor Group object but an OID outside the IGMP Filter Editor Group.

There is no workaround.

- CSCdx24363

When you add an entry that checks Transmission Control Protocol (TCP) flags to an access list that is being used for QoS classification, the system might report that the hardware limitation has been reached for the policy map. This can occur when the policy map already contains several other access list entries (ACEs) that check different TCP flags or that check TCP or User Datagram Protocol (UDP) port numbers using an operation different from *equal* (such as *not equal*, *less than*, *greater than*, or *range*). When the hardware limitation is reached, the **service-policy input** interface configuration command is removed from the running configuration of the interface. The problem occurs because checks for TCP flags and TCP/UDP port numbers using operators other than *equal* share hardware resources and there is a limit to the number of checks supported within a single policy map.

Similar limits affect port ACLs, VLAN maps, and router ACLs. Port ACLs and policy maps share the same pool of resources, so the resource usage of a policy map added to the resource usage of a port ACL applied to the same interface determine if the hardware resources have been exceeded. VLAN maps and router ACLs that share the same VLAN label also share a single pool of hardware resources, separate from the one shared by policy maps and port ACLs.

These are possible workarounds:

- Rearrange the order of classes within the policy map, the order of entries in the individual ACLs used in the policy map, and the order of entries within an IP port ACL applied to the interface so that checks for TCP flags are made as early as possible within the policy map. You can also rearrange the order of individual ACLs within a VLAN map and the order of entries that make up a security ACL.
- Add an extra entry to the beginning of the ACL to check for the same TCP flags that are checked later in the ACL. Creating the entry is easy to do if the first entry of the ACL matches the TCP protocol because you can create the new entry by duplicating the first entry and adding the check for the TCP flags to the duplicate.
- Reduce the number of combinations of TCP flags that are being tested.
- If all other workarounds fail, avoid combining the established keyword or any other check against the TCP flags with *greater than*, *less than*, *not equal*, or *range* checks within the policy map and the port ACL configured on the interface or within the VLAN maps and router ACLs that share the same VLAN label.



Note

To determine which VLAN label is assigned to a VLAN or interface, use the **show fm vlan *vlan-id*** or **show fm interface *interface-id*** privileged EXEC command. Then use the **show fm vlan-label *label-id*** privileged EXEC command to determine which set of VLAN maps or router ACLs share the label.

- CSCdx29360

If you create multiple loopback interfaces by using the **loopback interface 0** global configuration command, you cannot delete the loopback interfaces by using the **no interface loopback 0** command.

The workaround to delete the loopback interfaces is to reload the switch by using the **reload** privileged EXEC command.

- CSCdx37772

When you change the spanning-tree mode from MST to PVST and the number of VLANs is greater than 128, the traffic might not be forwarded on VLANs for which the spanning-tree instance is not created (the maximum number of spanning-tree instance is 128).

The workaround is to enter the **shutdown** and the **no shutdown** interface configuration commands on the port for which no addresses are learned.

- CSCdx39914

In some rare cases, you might see one of these messages:

```
Assertion failed:(used_before == used_after), file ../src-vegas/vqatm.c, line
or
Assertion failed:(used_before + num_entries == used_after), file
../src-vegas/vqatm.c, line
```

You might receive one of these messages after a **reload** privileged EXEC command or when you are configuring or changing any of these features:

- access lists
- VLAN maps
- MAC access groups
- policy maps
- IP addresses
- secondary IP addresses
- IP unreachables
- IP redirects
- IP multicast boundaries

If the message is seen during configuration, the workaround is to follow these steps:

1. Remove or detach all access groups, policy maps, VLAN maps, and multicast boundaries.
2. Attach the access groups, policy maps, VLAN maps, multicast boundaries in a different order.
3. If the error message is not seen again, the operation was successful.

Otherwise, repeat Steps 1 and 2 again using an order different than the one in Step 2.



Note In some cases, it might be difficult or impossible to discover an order that will eliminate the error message.

If a message is seen during a reload, you must edit the config.text file externally and then download the edited config.text file to the switch. Editing the config.text file involves changing the order in which the configurations for different interfaces are listed. You might also need to reorder **vlan filter** global configuration commands in relation to each other and in relation to the configuration of one or more of the interfaces.

For information about how to edit the config.txt file, refer to “Appendix B, Working with the IOS File System, Configuration Files, and Software Images” in the *Catalyst 3550 Multilayer Switch Software Configuration Guide*.

Open Cluster Caveats

These are the severity 3 cluster configuration caveats:

- CSCdw38507

When you enter a remote command on a Catalyst 3550 member switch from a command switch that is not a Catalyst 3550 switch, if the command generates a lot of output and the output is paused and restarted, communication between the two switches might halt.

The workaround is to follow the documented cluster configuration guidelines, which recommend using a Catalyst 3550 switch as the command switch in mixed-model clusters.

- CSCdw91356

When Catalyst 1900, 2820, or 2900 XL 4 MB series switches are participating in a cluster and the active command switch fails and then is restored after the standby command switch has updated the members, it is possible for these legacy switches to miss the restoration and retain the standby command switch's MAC address.

The workaround is to manually reset the command switch MAC address on each member switch.

Open Cluster Management Suite Caveats

These are the severity 3 CMS configuration caveats:

- CSCdu79932

If you try to enable Port Fast on an interface that does not accept it—a trunk port, for example—no message warns you that Port Fast was not enabled.

There is no workaround.

- CSCdv35455

If you select multiple FastEthernet ports on a Catalyst 3550 switch, the speed of 1000 Mbps is shown as an option in the Modify Port Settings window. Ignore this speed option.

- CSCdv57881

You cannot modify the multicast groups that are shown in the IGMP Snooping window.

The workaround is to delete the group that you want to modify and then recreate it with the change that you want.

- CSCdw01109

In the CMS QoS Policies window, the Attach tab does not show any egress policy information. Even when some interfaces have an egress QoS policy associated with them, the policy does not appear in the Egress Policy column of the table of attached QoS policies. Note that attachment and detachment do work correctly, but the results cannot be viewed on CMS.

There is no workaround.

- CSCdx13380

If both a port ACL and a VLAN map are already configured on the Catalyst 3550 switch and you try to attach a port ACL through the CMS Security Wizard, you can use the Security Wizard to attach the port ACL to a switch port. The attached port ACL conflicts with the existing VLAN map and this is not a allowed configuration.

The workaround is to verify that both a port ACL and a VLAN map are not configured on the switch before using the Security Wizard.

Resolved Caveats

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

- [“Resolved IOS Caveats” section on page 29](#)
- [“Resolved Cluster Caveats” section on page 31](#)
- [“Resolved Cluster Management Suite Caveats” section on page 31](#)

Resolved IOS Caveats

These IOS caveats were resolved in this release:

- CSCdt51254
If you try to attach an ACL that uses the **log** keyword to a class-map, an error is displayed, but a **match none** statement is no longer added to the class map.
- CSCdu87797
If Catalyst 3550 GBIC ports containing GigaStack GBICs are configured as routed ports with EtherChannel group assignments, they are no longer allowed to join the channel group. The Catalyst 3550 switch does not support GigaStack GBICs as EtherChannel group members.
- CSCdv10257
You can now set the vmVlanType object in CISCO-VLAN-MEMBERSHIP-MIB and convert a port from static to dynamic access by using SNMP.
- CSCdv10276
The switch now generates the vmVmmpsChange trap in CISCO-VLAN-MEMBERSHIP-MIB, so you can use SNMP to monitor VMPS change events.
- CSCdv23503
The **standby mac-address** interface configuration command has been removed from the CLI.
- CSCdv28246
When you use the **system mtu** global configuration command to increase the system MTU to more than 1500 bytes and to reload the switch, large frames that are sent to the CPU are no longer truncated.
- CSCdv59364
If you use the setup program for initial configuration and set an IP address for a specific port (making it a router port and not a switch port), the setup program now correctly writes the **no switchport** interface configuration command to the configuration file, and manual reconfiguration of the port is not required.
- CSCdv70296
If a port is put into self-looped state, Spanning Tree Protocol (STP) no longer remains in the blocking state after the self-looped condition is removed.
- CSCdw07801
After you have changed the system maximum transmission unit (MTU) size by using the **system mtu** global configuration command, the new MTU value is applied to all interfaces on the switch, and the **show interfaces** privileged EXEC command now correctly displays the new MTU value.

- CSCdw17713
When you configure a switch port to block flooded traffic by using the **switchport block** interface configuration command, and then configure the port as a SPAN destination port by using the **monitor session session_number destination interface interface_id** global configuration command, the block settings no longer remain enabled during the SPAN session and cause traffic sent to unknown addresses to be blocked instead of forwarded to the destination port.
- CSCdw18077
When you connect two Catalyst 3550 switches through a GBIC port and configure both with the **speed nonegotiate** interface configuration command, the port now has a linkup status after you enter the **shutdown** and **no shutdown** interface configuration commands.
- CSCdw27053
When you use the **sdm prefer** global configuration command to select a template to be used for Switch Database Management (sdm) resource allocation, the configuration is stored in nonvolatile RAM, and the new template takes effect after the next switch reload. You can use the **show sdm prefer** privileged EXEC command to see which template is configured.
- CSCdw40980
The **copy running-config startup-config** privileged EXEC command no longer fails if the running configuration of the Catalyst 3550 exceeds 32 K.
- CSCdw51665
When the switch is configured with login authentication with local usernames, memory use no longer increases with each login to the switch.
- CSCdw52807
SNMP no longer returns errors or zero values for 64-bit counters in the IF-MIB.
- CSCdw57401
Packets denied by an output access list or a VLAN map on an output VLAN are no longer forwarded to the CPU for processing rather than being dropped by hardware resources in the Catalyst 3550 if all these conditions are true:
 - The destination is a directly connected host.
 - The ACL denying the packet is an output ACL (input ACLs should not have this problem).
 - None (or no significant amount) of the traffic to this host is being permitted, so no ARP entry is created.
 - The destination host is not originating any traffic that needs to go through the router (if it were, the host would employ ARP for the router's MAC address, which should cause the software to create an ARP entry).
- CSCdw57545
When a GBIC port that has a GigaStack GBIC module installed is administratively down, you can now bring up the link by using the **no shutdown** interface configuration command.
- CSCdw63139
When a Catalyst 3550 switch has been changed from the HSRP active router into the standby router, the switch now responds to a trace route through it.
- CSCdw84685
When a corrupted STP packet has a maximum age between 0 and 1 second, the switch no longer ages the bridge protocol data unit (BDPU). This causes a new spanning-tree root to be elected.

- CSCdx16941

When you enter the **show mls qos aggregate-policer** privileged EXEC command, the switch no longer fails if you create a policy map that uses an undefined aggregate policer, then delete the policy map, and then try to show or delete the aggregate policer.

Resolved Cluster Caveats

This cluster caveat was resolved in this release:

- CSCdw50973

The **cluster setup** privileged EXEC command has been removed from the CLI.

Resolved Cluster Management Suite Caveats

These CMS caveats were resolved in this release:

- CSCdw63797

In the Port Statistics window, the statistics appear on the correct ports.

- CSCdw63890

When an IP phone is connected to a customer premise equipment (CPE) device that is connected to a Catalyst 2900-LRE-XL switch and the Topology View is used to view the network, it now correctly shows the IP phone connectivity.

- CSCdw92499

You can now use CMS to create a VLAN map without a sequence number.

Documentation Updates

You can access all Catalyst 3550 documentation at this URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/index.htm>

This section provides updates to the product documentation. These changes will be included in the next version of the documentation.

Modifications

These are modifications or corrections for the *Catalyst 3550 Multilayer Switch Software Configuration Guide*:

- Encrypted Secure Shell (SSH), described in “Chapter 7, Administering the Switch,” is not available in this release.

- In the “Configuring Port Security” section in Chapter 19, this statement is on page 19-8:
“It is a security violation when
 - A station whose MAC address is configured as a secure MAC address on another secure port attempts to access the interface.”

The statement should be:

“It is a security violation when an address learned or configured on one secure interface is seen on another secure interface in the same VLAN.”

Related Documentation

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

<http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/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 “Ordering Documentation” section on page 33.

- *Catalyst 3550 Multilayer Switch Software Configuration Guide* (order number DOC-7811194=)
- *Catalyst 3550 Multilayer Switch Command Reference* (order number DOC-7811195=)
- *Catalyst 3550 Multilayer Switch System Message Guide* (order number DOC-7811196=)
- Cluster Management Suite (CMS) online help (available only from the switch CMS software)
- *Catalyst 3550 Multilayer Switch Hardware Installation Guide* (order number DOC-7811358=)
- *1000BASE-T Gigabit Interface Converter Installation Note* (not orderable but is available on Cisco.com)
- *Catalyst GigaStack Gigabit Interface Converter Hardware Installation Guide* (order number DOC-786460=)

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following URL:

<http://www.cisco.com>

Translated documentation is available at the following URL:

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

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco product documentation from the Networking Products MarketPlace:
http://www.cisco.com/cgi-bin/order/order_root.pl
- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:
<http://www.cisco.com/go/subscription>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on the World Wide Web, you can send us your comments by completing the online survey. When you display the document listing for this platform, click **Give Us Your Feedback**.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Cisco Systems, Inc.
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

<http://www.cisco.com>

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

<http://www.cisco.com/tac>

All customers, partners, and resellers who have a valid Cisco services contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to the following URL to register:

<http://www.cisco.com/register/>

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

<http://www.cisco.com/tac/caseopen>

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer will automatically open a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to the following URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled; for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). In addition, please have available your service agreement number and your product serial number.

This document is to be used with the documentation listed in the “[Related Documentation](#)” section.

CCIP, the Cisco *Powered Network* mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, Internet Quotient, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ Logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to

Copyright © 2002, Cisco Systems, Inc.
All rights reserved.

