

New and Changed Information

The following sections list the new hardware and software features supported by the Cisco IOS Release 12.0.

New Hardware Features in Cisco IOS Release 12.0(24)

Cisco 3640A Router

The Cisco 3640A router extends the life of the Cisco 3640 router, providing identical features, functionality, and performance. The Cisco 3640A router will support the same Cisco IOS feature sets as the Cisco 3640 router, but requires a different minimum version of Cisco IOS software.

New Software Features in Cisco IOS Release 12.0(1)

The following new software features are supported by Cisco IOS Release 12.0.

AAA Authorization and Accounting Feature

In earlier Cisco IOS releases, only named authentication method lists were supported under Cisco's Authentication, Authorization, and Accounting (AAA) network security services. Now, AAA has been extended to support both authorization and accounting named method lists. Named method lists for authorization and accounting function the same way as those for authentication. Named method lists for AAA authorization and accounting allow you to define different methods for authorization and accounting and apply those methods on a per-interface or per-line basis.

AAA Scalability

The Authentication, Authorization and Accounting (AAA) Scalability feature enables you to configure and monitor the number of background processes allocated by the PPP manager in a network access server (NAS) to deal with AAA authentication and authorization requests. In previous Cisco IOS releases, only one background process was allocated to handle all AAA requests for PPP. This meant that parallelism in AAA servers could not be fully exploited. The AAA Scalability feature enables you to configure the number of processes used to handle AAA requests for PPP, thus increasing the number of users that can be simultaneously authenticated or authorized.

The AAA Scalability feature provides an increase in the number of parallel authentication and authorization requests the NAS can forward to the AAA server.

AAA Support for MS-CHAP

Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) is the Microsoft version of CHAP and is an extension to RFC 1994. Like the standard version of CHAP, MS-CHAP is used for PPP authentication; in this case, authentication occurs between a PC using Microsoft Windows NT or Microsoft Windows 95 and a Cisco router or access server acting as a network access server (NAS).

MS-CHAP differs from the standard CHAP as follows:

- MS-CHAP is enabled by negotiating CHAP Algorithm 0x80 in LCP option 3, Authentication Protocol.
- The MS-CHAP Response packet is in a format designed to be compatible with Microsoft Windows NT 3.5 and 3.51, Microsoft Windows 95, and Microsoft LAN Manager 2.x. This format does not require the authenticator to store a clear or reversibly encrypted password.
- MS-CHAP provides an authenticator-controlled authentication retry mechanism.
- MS-CHAP provides an authenticator-controlled change password mechanism.
- MS-CHAP defines a set a “reason-for-failure” codes returned in the Failure packet message field.

Depending on the security protocols you have implemented, PPP authentication using MS-CHAP can be used with or without Authentication, Authorization and Accounting (AAA) security services. If you have enabled AAA, PPP authentication using MS-CHAP can be used in conjunction with both TACACS+ and RADIUS. Two new vendor-specific RADIUS attributes (IETF Attribute 26) were added to enable RADIUS to support MS-CHAP.

Additional Vendor-Proprietary RADIUS Attributes

Remote Authentication Dial-In User Server (RADIUS) is an access server authentication, authorization, and accounting protocol originally developed by Livingston, Inc. Although an Internet Engineering Task Force (IETF) draft standard for RADIUS specifies a method for communicating vendor-proprietary information between the network access server and the RADIUS server, some vendors have extended the RADIUS attribute set in a unique way. In this release, Cisco IOS software introduces support for additional vendor-proprietary RADIUS attributes.

For a complete list of supported IETF and vendor-proprietary RADIUS attributes, refer to the “RADIUS Attributes” appendix in the Cisco IOS Release 11.3 *Security Configuration Guide*.

Users who have implemented security solutions using a vendor-proprietary implementation of RADIUS can now integrate Cisco access routers into their networks more easily.

Airline Product Set (ALPS)

The Airline Product Set (ALPS) feature is a tunneling mechanism that transports airline protocol data across a Cisco router-based TCP/IP network to an X.25-attached mainframe. This feature provides connectivity between agent set control units (ASCUs) and a mainframe host that runs the airline reservation system database.

The basic ALPS topology is composed of three major components that provides the end-to-end transportation of airline protocol traffic across the network: the Airline Control (ALC) protocol, the TCP-based transport protocol, and the AX.25/EMTOX access to the mainframe.

The ALPS feature is integrated in the Cisco IOS software and allows airlines to replace their existing hardware and software with Cisco routers. For customers who already use Cisco routers, this feature allows them to consolidate networking overhead and functionality.

Always-On/Dynamic ISDN

Always On/Dynamic ISDN (AO/DI) is an on-demand service that is designed to optimize the use of an existing Integrated Services Digital Network (ISDN) signaling channel (D channel) to transport X.25 traffic. The X.25 D channel call is placed from the subscriber to the packet data service provider. The use of PPP allows protocols to be encapsulated within the X.25 logical circuit carried by the D channel. The bearer channels (B channels) use the Multilink protocol without the standard Q.922 and X.25

encapsulations, and invoke additional bandwidth as needed. Optionally, the Bandwidth Allocation Control Protocol (BACP) and the Bandwidth Allocation Protocol (BAP) can be used to negotiate bandwidth allocation as required.

AO/DI takes full advantage of existing packet handlers at the central office by using an existing D channel to transport the X.25 traffic. The link associated with the X.25 D channel packet connection is used as the primary link of the Multilink bundle. The D channel is a connectionless, packet-oriented link between the Customer Premise Equipment (CPE) and the central office. Since the D channel is always available, it is possible to in turn offer “always available” services. On-demand functionality is achieved by using the B channels to temporarily boost data throughput and are disconnected after use.

ATM E.164 Auto Conversion

The ATM E164 Auto Conversion feature allows a Cisco router to set up ATM switched virtual circuit (SVC) connections based on E.164 addresses.

E.164 is an International Telecommunications Union (ITU) specification for the ISDN international telephone numbering plan, which has traditionally only been used in telephone networks. The Asynchronous Transfer Mode (ATM) Forum has defined three different 20-byte ATM End System Address (AESA) formats, along with the native E.164 format, for use in ATM networks. One of these 20-byte formats is the embedded E.164 AESA (E164_AESA) format. The ATM E164 Auto Conversion feature allows networks that operate based on ATM addressing formats to interconnect with networks based on E.164 addressing formats.

The ATM E164 Auto Conversion feature requires components from addressing, routing, and signaling to perform properly.

Automated Double Authentication

The automated double authentication feature enhances the existing double authentication feature.

Previously, with the existing double authentication feature, a second level of user authentication is achieved when the user Telnets to the network access server or router and enters a username and password. Now, with automated double authentication, the user does not have to Telnet anywhere but instead responds to a dialog box that requests a username and password or PIN. For information about the existing double authentication feature, refer to the “Configuring Authentication” chapter of the *Cisco IOS Release 11.3 Security Configuration Guide*.

This feature has all the security benefits of double authentication, but provides a simpler, more user-friendly interface for remote users. Users are no longer required to Telnet to a remote device; they can simply respond to on-screen dialogs.

The remote user hosts must be running a companion client application. As of the first publication of this document, the only client application software available is the Glacier Bay application server software for PCs.

Automatic Protection Switching of Packet-over-SONET Circuits

The automatic protection switching (APS) feature is supported on Cisco 7500 series routers. This feature allows switch over of packet-over-SONET (POS) circuits and is often required when connecting SONET equipment to telco equipment. APS refers to the mechanism of bringing a “protect” POS interface into the SONET network as the “working” POS interface on a circuit from the intervening SONET equipment.

The protection mechanism used for this feature is “1+1, Bidirectional, nonrevertive” as described in the Bellcore publication TR-TSY-000253, SONET Transport Systems; Common Generic Criteria, Section 5.3. In the 1+1 architecture, there is one working interface (circuit) and one protect interface, and the same payload from the transmitting end is sent to both the receiving ends. The receiving end decides which interface to use. The line overhead (LOH) bytes (K1 and K2) in the SONET frame indicate both status and action.

The protect interface is configured with the IP address of the router that has the working interface. The APS Protect Group Protocol, which runs on top of UDP, provides communication between the process controlling the working interface and the process controlling the protect interface.

Using this protocol, POS interfaces can be switched due to a router failure, degradation or loss of channel signal, or manual intervention. In bidirectional mode, the receive and transmit channels are switched as a pair. In unidirectional mode, the transmit and receive channels are switched independently. For example, if the receive channel on the working interface has a loss of channel signal, both the receive and transmit channels are switched.

In addition to the new Cisco IOS commands added for the APS feature, the POS interface configuration commands **pos threshold** and **pos report** have been added to support user configuration of the bit error rate (BER) thresholds and reporting of SONET alarms.

Bridging Code Rework

In Cisco IOS Release 12.0, the bridging code has been reworked to modularize the dependencies between Cisco IOS VLAN, the L2 path, and the IEEE 802.1d Spanning Tree.

Certification Authority Interoperability

Certification Authority (CA) interoperability is provided in support of the IP Security (IPSec) standard. CA interoperability permits Cisco IOS devices and CAs to communicate so that your Cisco IOS device can obtain and use digital certificates from the CA. Although IPSec can be implemented in your network without the use of a CA, using a CA provides manageability and scalability for IPSec. Without CA interoperability, Cisco IOS devices could not use CAs when deploying IPSec. CAs provide a manageable, scalable solution for IPSec networks.

CIP Core Dump Support

A class of global configuration commands exist in the Cisco IOS software that allow you to output a core dump from a router or a processor card in the router when the Cisco IOS software halts unexpectedly. These are the **exception** global configuration commands. The **exception** commands were hidden from the user configuration interface until Cisco IOS Software Release 11.2(8).

The **exception slot** global configuration command lets you configure a CIP to output a core dump when the Cisco IOS software running on the CIP halts unexpectedly.

The **exception slot** command uses FTP to transfer data from the CIP to a host system.



Note

The output obtained by the **exception** command can be interpreted by a qualified Cisco technical support person.

Cisco AS5300 Additional Functionality

The following additional software support is provided for the Cisco AS5300 access server:

- Modem pooling
- E1 R2 signaling
- Fast Ethernet 110BaseT interface support

Cisco Database Connection

The Cisco Database Connection feature enables Cisco routers to implement IBM's distributed relational database architecture (DRDA) level 3 over TCP/IP. The Cisco router with Database Connection exists in the TCP/IP network, and clients use the Database Connection IP address and port on the router to connect to the IBM host system that exists in the SNA network.

When Database Connection is configured on a router, client-based Open Database Connectivity (ODBC) applications can connect to IBM's family of IBM D2 relational databases which include:

- DB2 for OS/390 (MVS)
- DB2 for Virtual Machine (VM)
- DB2 for Virtual Storage Extended (VSE) (SQL/DS)
- DB2 for OS/400
- DB2 Universal Server (AIX, HP-UX, UNIX, Solaris, Windows NT, Windows 95, OS/2, SCO OpenServer)

The router with Database Connection converts DRDA packets over TCP/IP to DRDA packets over APPC (LU 6.2) and then routes them to DB2 databases. Database Connection runs as a TCP/IP daemon on the router, accepting DRDA client connections over TCP/IP. When a client connects to the database on an IBM mainframe host, Database Connection allocates an APPC conversation over SNA to an IBM server, and acts as a gateway between DRDA over TCP/IP and DRDA over APPC.

When configured on a router, the Database Connection feature enables desktop applications to access data in remote databases located on IBM hosts. Database Connection receives database access messages from the client over a TCP/IP link. Database Connection converts the messages to SNA and transmits them to the host using APPC services provided by the Cisco IOS APPN software.

The Database Connection feature offers the following benefits:

- Leverages existing TCP/IP network—Because Database Connection converts DRDA packets over TCP/IP to DRDA packets over APPC (LU 6.2) you can leverage TCP/IP in your enterprise.
- Improves manageability—You can manage enterprise-wide access to DB2 from a single centralized location within the data center.
- Maximizes computer resources—Database Connection takes advantage of distributed processing and standard communication protocols and reduces CPU utilization on the host.
- Eliminates special software—Database Connection works on your router without the need for any special software on the IBM host. It allows ODBC clients to connect to IBM's DB2 relational databases using TCP/IP without the need for communication packages on the desktop. Database Connection supports ODBC on client systems, allowing you to use applications of your choice that are enabled with ODBC. Some examples of applications that utilize ODBC are Microsoft Excel, Microsoft Access, Lotus 1-2-3, Visual Basic, Visual C++, and PowerBuilder.
- Increases data speed access—Cisco routers enable high-speed connections to DB2 on hosts, and these connections are faster than native host TCP/IP.

Cisco Express Forwarding

Cisco express forwarding (CEF) is advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions.

Although you can use CEF in any part of a network, it is designed for high-performance, highly resilient Layer 3 IP backbone switching.

CEF offers these benefits:

- Improved performance—CEF is less CPU-intensive than fast or optimum switching route caching. More CPU processing power can be dedicated to Layer 3 services such as quality of service (QoS) and encryption.
- Scalability—CEF offers full switching capacity at each line card when dCEF mode is active.
- Resilience—CEF offers unprecedented level of switching consistency and stability in large dynamic networks. In dynamic networks, fast switching cache entries are frequently invalidated due to routing changes. These changes can cause traffic to be process switched using the routing table, rather than fast switched using the route cache. Because the FIB lookup table contains all known routes that exist in the routing table, it eliminates route cache maintenance and the fast switch/process switch forwarding scenario. CEF can switch traffic more efficiently than typical demand caching schemes.

Cisco IOS File System

The Cisco IOS File System (IFS) feature provides a single interface to all file systems the router uses, including:

- Flash memory file systems
- Network file systems (TFTP, rcp, and FTP)
- Any other endpoint for reading or writing data (such as NVRAM, the running configuration, ROM, raw system memory, system bundled microcode, Xmodem, Flash load helper log, modems, and BRI MUX interfaces).

IFS provides the following benefits:

- File viewing and classification—With IFS, all files can be viewed and classified (image, text file, and so forth), including files on remote servers. For example, you may want to determine the size and type of an image on a remote server before you copy it to ensure that it is a valid image. You can also view a configuration file on a remote server to verify that it is the correct configuration file before you load the file on the router.
- Platform-independent commands—With IFS, the file system user interface is no longer platform specific. Commands have the same syntax, regardless of which platform is used. Thus, you can use the same commands for all of your routers. However, not all commands are supported on all platforms and file systems. Because different types of file systems support different operations, certain commands are not available for all file systems. Platforms will support commands for the file systems they use.
- Minimal prompting for commands—IFS minimizes the required prompting for many commands, such as the copy command. You can enter all of the required information in the command line, rather than having to provide information when the system prompts you for it. For example, if you want to copy a file to an FTP server, you can specify the specific location on the router of the source file,

the specific location of the destination file on the FTP server, and the username and password to use when connecting to the FTP server, all on a single line. However, if you wish to have the router prompt you for the needed information, you can still enter the minimum form of the command.

Depending on the current configuration of the file prompt command and the type of command you entered, the router may prompt you for confirmation, even if you have provided all the information in the command. In these cases, the default value will be the value entered in the command. Press Return to confirm the values.

- Directory navigation and creation—With IFS, you can move around to different directories and list the files in a directory. On newer platforms, you can create subdirectories in Flash memory or on a disk.
- URLs—The new file system interface uses Uniform Resource Locators (URLs) to specify the location of a file. URLs are commonly used to specify files or locations on the World Wide Web. However, on Cisco routers, they can now be used to specify the location of files on the router or remote file servers.
- Network files—You can specify FTP, rcp, and TFTP files on the network server. The location can be an IP address or a host name. The username variable, if specified, overrides the username specified by the **ip rcmd remote-username** or **ip ftp username** commands. The password overrides the password specified by the **ip ftp password** command.

The file path (directory and filename) is specified relative to the directory used for file transfers. For example, on UNIX file servers, TFTP pathnames start in the /tftpboot directory, and rcp and FTP paths start in the home directory associated with the username.

- Local files—For local files (files on the router), the syntax is basically the same as the syntax previously used on high-end platforms; however, you can now specify directories.
- URL prefix—The URL prefix specifies the file system. The list of available file systems differs by platform and operation. Refer to your product documentation or use the **show file systems** command to determine which prefixes are available on your platform.

CLAW IP Packing

Cisco IOS implements the Common Link Access to Workstation (CLAW) Channel Protocol to transport data between the mainframe and the Cisco CIP in TCP/IP environments. The CLAW packing feature is an enhancement to the CLAW protocol support which enables the transport of multiple IP packets in a single channel operation.

The CLAW packing feature significantly increases throughput performance between a mainframe and a Cisco CIP.

Currently, IBM's TCPIP stack does not support the CLAW packing feature. However, the original implementation of the CLAW IP datagram support will continue to work with IBM's stack, even concurrently with the CLAW packing feature.

Committed Access Rate

The Committed Access Rate (CAR) feature performs the following functions:

- Limits the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.
- Classifies packets by setting the IP precedence or QoS group.

CAR can be used to rate-limit traffic based on certain matching criteria, such as incoming interface, IP precedence, QoS group, or IP access list criteria. CAR provides configurable actions, such as transmit, drop, set precedence, when traffic conforms to or exceeds the rate limit.

Conditionally Triggered Debugging

The Conditionally Triggered Debugging feature limits debugging messages based on their related interface or subinterface. When this feature is enabled, the router generates debugging messages for packets entering or leaving the router on a specified interface. However, the router does not generate debugging output for packets entering or leaving through a different interface. This feature allows you to focus debugging output on the problematic interface or interfaces.

You can specify the interfaces explicitly. For example, you may only want to see debugging messages for one interface or subinterface. You can also turn on debugging for all interfaces that meet specified conditions, such as a particular username, calling party number, or called party number. If you specify multiple conditions, the interface must meet at least one of the conditions.

This feature controls the output from the following protocol-specific **debug** commands:

- **debug aaa {accounting | authorization | authentication}**
- **debug dialer {events | packets}**
- **debug isdn {q921 | q931}**
- **debug modem {oob | trace}**
- **debug ppp {all | authentication | chap | error | negotiation | multilink events | packet}**

While this feature limits the output of the above commands, it does not automatically enable the generation of debugging output from these commands. Debugging messages are generated only if the protocol-specific debug command is enabled.

This feature is useful on dial access servers, which have a large number of ports. Normally, the router will generate debugging messages for every interface, resulting in a large number of messages. The large number of messages consumes system resources. For example, if the messages are displayed on the console, the router must take the time to send every message to the console. Similarly, if you are sending messages to a syslog server, the large number of generated output consumes network resources.

In addition, the large number of messages can make it difficult to find the specific information you need. Messages may scroll by on the console too quickly, or the logging buffer may wrap around before you are able to examine the contents of the buffer.

By limiting the debugging messages, you can receive messages related to only the ports you want to troubleshoot. This limiting decreases the number of generated messages, thus consuming fewer resources and making it easier for you to find the information you want.

Context-Based Access Control Feature

The Cisco IOS Firewall feature set combines existing Cisco IOS firewall technology and the new context-based access control (CBAC) feature. When you configure the Cisco IOS Firewall feature set on your Cisco router, you turn your router into an effective, robust firewall.

The Cisco IOS Firewall feature set is designed to prevent unauthorized external individuals from gaining access to your internal network, and to block attacks on your network, while at the same time allowing authorized users to access network resources.

You can use the Cisco IOS Firewall feature set to configure your Cisco IOS router as:

- An Internet firewall or part of an Internet firewall

- A firewall between groups in your internal network
- A firewall providing secure connections to or from branch offices
- A firewall between your company's network and your company's partners' networks

The Cisco IOS Firewall feature set provides the following benefits:

- Protects internal networks from intrusion
- Monitors traffic through network perimeters
- Enables network commerce via the World Wide Web

Dialer Watch

Dialer Watch is a backup feature that integrates dial backup with routing capabilities. Prior dial backup implementations used the following conditions to trigger backup:

- Interesting packets were defined at central and remote routers using Dial on Demand routing (DDR).
- Connection loss occurred on a primary interface using a back up interface with floating static routes.

Distributed Cisco Express Forwarding Netflow

The dCEF NetFlow feature allows the VIP and VIP2 to NetFlow switch packets and perform NetFlow data export similar to CEF Flow switching on the RSP. Each VIP maintains its own independent flow cache and can generate its own export packets containing statistics on expired flows. Use the **show ip cache flow** command to display per-protocol statistic summaries for packets switch by the RSP and VIP. The **show ip cache flow** command shows only flow details for RSP switched flows.

Distributed Weighted Fair Queuing

Flow-based weighted fair queuing (WFQ) controls the ratio of transmission bandwidth allocation among different traffic flows during periods of congestion. Class-based WFQ allocates transmission bandwidth among different traffic flows or QoS groups during periods of congestion.

The Distributed Weighted Fair Queuing (DWFQ) feature uses the VIP rather than the RSP to perform the queuing; therefore, it requires a Cisco 7500 series router or Cisco 7000 series router with RSP7000.

Flow-based WFQ provides the following benefits:

- All flows are allocated equal bandwidth.
- Well-behaved hosts are protected from badly behaved hosts.
- WFQ provides absolute allocation of bandwidth in unequal amounts based on traffic requirements.
- DWFQ uses the VIP to provide a faster implementation of WFQ than the RSP implementation.

Distributed Weighted Random Early Detection

Random Early Detection (RED) is a congestion avoidance mechanism that takes advantage of TCP's congestion control mechanism. By randomly dropping packets prior to periods of high congestion, RED tells the packet source to decrease its transmission rate. Assuming the packet source is using TCP, it will decrease its transmission rate until all the packets reach their destination, indicating that the congestion is cleared.

Weighted RED (WRED) generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. Thus, higher priority traffic is delivered with a higher probability than lower priority traffic.

However, you can also configure WRED to ignore IP precedence when making drop decisions so that non-weighted RED behavior is achieved.

WRED is useful on any output interface where you expect to have congestion. However, WRED is usually used in the core routers of a network, rather than the edge. Edge routers assign IP precedences to packets as they enter the network. WRED uses these precedences to determine how it treats different types of traffic.

The Distributed WRED (DWRED) feature uses the VIP rather than the RSP to perform the queuing; therefore, it requires a Cisco 7500 series router or Cisco 7000 series router with RSP7000.

DRP Server Agent Enhancements

The DRP Server Agent enhancements are as follows:

- DistributedDirector can use BGP Multi-Exit Discriminators in traffic redirection decisions.
- The DRP Server can measure client-to-server link latency (round trip time) for use in traffic redirection decisions.

E1 R2 Country Support and Modem Management Enhancements

R2 signaling is an international signaling standard that is common to channelized E1 networks. However, there is no single signaling standard for R2. The ITU-T Q.400-Q.490 recommendation defines R2, but a number of countries and geographic regions implement R2 in entirely different ways. Cisco Systems addresses this challenge by supporting many localized implementations of R2 signaling in its Cisco IOS software.

Cisco System's E1 R2 signaling default is ITU, which supports the following countries: Denmark, Finland, Germany, Russia (ITU variant), Hong Kong (ITU variant), and South Africa (ITU variant). The expression "ITU variant" means there are multiple R2 signaling types in the specified country, but Cisco supports the ITU variant.

Cisco Systems also supports specific local variants of E1 R2 signaling in the following regions, countries, and corporations:

- | | |
|---------------------------------------------------------------|------------------------------------|
| • Argentina | • Malaysia |
| • Australia | • New Zealand |
| • Brazil | • Paraguay |
| • China | • Peru |
| • Columbia | • Philippines |
| • Costa Rica | • Saudi Arabia |
| • East Europe (includes Croatia, Russia, and Slovak Republic) | • Singapore |
| • Ecuador ITU | • South Africa (Panafstel variant) |
| • Ecuador LME | • Telmex corporation (Mexico) |
| • Greece | • Telnor corporation (Mexico) |
| • Guatemala | • Thailand |

- Hong Kong (uses the China variant)
- Indonesia
- Israel
- Korea
- Uruguay
- Venezuela
- Vietnam

**Note**

Only MICA modems support R2 functionality. Microcom modems do not support R2.

Benefits include:

- R2 custom localization—R2 signalling is supported for a wide range of countries and geographical regions. Cisco is continually supporting new countries.
- Broader deployment of dial access services—The flexibility of a high-density access server can be deployed in E1 networks.

Encryption over Frame Relay

You can now use any type of encapsulation with IP encryption, except as follows: If you have a second-generation Versatile Interface Processor (VIP2) with a serial interface, encryption will not work for traffic on the serial interface unless you use the Point-to-Point Protocol (PPP), High-Level Data Link Control (HDLC) protocol, or Frame Relay protocol. For example, you cannot use encryption if you have X.25 or SMDS configured for the serial interface of a VIP2.

Enhanced ATM VC Configuration and Management

The Enhanced ATM VC Configuration and Management feature set includes new and enhanced capabilities that allow you to create and manage ATM PVCs and SVCs with more ease and improved integrity. This feature set includes the following five subfeatures:

- **New VC Configuration**—The New VC Configuration subfeature allows you to create ATM permanent virtual circuits (PVCs), switched virtual circuits (SVCs), static maps, and associated virtual circuit (VC) parameters more easily and with fewer errors using new ATM commands in new VC command modes.
- **VC Integrity Management**—The VC Integrity Management subfeature allows you to manage your ATM PVCs and SVCs so that your router receives immediate notification of when these VCs go down in your network. Upon notification, protocols can reroute packets and prevent unpredictable and relatively long timeout periods.
- **PVC Discovery**—The PVC Discovery subfeature allows you to enable your router to automatically assign (or discover) PVCs on an ATM interface or subinterface using information from an attached adjacent switch.
- **Multiprotocol Inverse ARP**—The Multiprotocol Inverse ARP subfeature allows you to enable a dynamic protocol mapping between an ATM PVC and a network address by configuring Inverse Address Resolution Protocol (Inverse ARP) on ATM PVCs running IP or IPX.
- **Rate Queue Tolerance**—The Rate Queue Tolerance subfeature allows you to configure a range of peak rates on a single rate queue, thereby improving ATM rate queue usage.

Use the Enhanced ATM VC Configuration and Management feature set to simplify and expedite PVC and SVC configurations and improve the management of PVC and SVC integrity. The benefits of this feature set include:

- Simplified ATM PVC, SVC, and static map configuration.
- VC management that detects connections and disconnections of PVCs and SVCs immediately so that packets are rerouted upon notification.
- Automatic assignment (or discovery) of ATM PVCs on an ATM interface or subinterface using information from an attached adjacent switch.
- Dynamic protocol mapping between a PVC and a network address so that you no longer have to manually configure an ATM static map.
- Improved rate queue usage when you configure a range of peak rates on a single rate queue.

Expanded IP Access Lists

This feature expands the extended IP access list range as follows:

- 1-99—IP standard access list
- 100-199—IP extended access list
- 1300-1999—IP standard access list (expanded range)
- 2000-2699—IP extended access list (expanded range)

Fast EtherChannel

The Fast EtherChannel feature allows multiple Fast Ethernet point-to-point links to be bundled into one logical link to provide bidirectional bandwidth of up to 800 Mbps. Fast EtherChannel can be configured between Cisco 7500 series routers and Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI) or between a Cisco 7500 series router or a Cisco 7000 series router with the RSP7000 and RSP7000CI and a Catalyst 5000 switch.

FDDI Frames Per Token Limit

This feature allows an interface to transmit multiple frames per token, instead of only a single frame at a time. Users can specify the maximum number of frames to be transmitted with each token capture. Throughput is significantly increased, which benefits customers with heavy or very bursty traffic.

FEIP2-DSW Second-Generation Fast Ethernet Interface Processors

The FEIP2-DSW second-generation Fast Ethernet Interface Processor is a replacement for the FEIP2-2TX and FEIP2-2FX, which are available on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

Second-Generation Fast Ethernet Interface Processors—The second-generation Fast Ethernet Interface Processors (FEIP2-2TX and FEIP2-2FX) are available on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

Frame Relay-ATM Interworking on the Cisco MC3810

The Cisco MC3810 supports the FRF.5 Frame Relay-ATM Interworking function, which enables Frame Relay voice or data traffic to be encapsulated in ATM cells. You can transport Frame Relay traffic over an ATM cloud via a virtual interface within the Cisco MC3810. Using the encapsulation process, you can migrate from Frame Relay to ATM, or you can tunnel Frame Relay traffic across an ATM backbone to a second Cisco MC3810 or other Frame Relay device, and then extract the ATM traffic back to Frame Relay. The Frame Relay traffic is encapsulated in the ATM data cells.

FTP Server

The FTP Server feature enables a Cisco IOS device to act as an FTP server. This feature was first introduced in Cisco IOS Release 11.3 AA for the Cisco 3640 platform.

HSRP MAC Refresh Interval

When Hot Standby Router Protocol (HSRP) runs over FDDI, you can change the interval at which a packet is sent to refresh the MAC cache on learning bridges or switches. HSRP hello packets use the burned-in address instead of the MAC virtual address. Refresh packets keep the switch's or learning bridge's MAC cache current.

You can change the refresh interval on FDDI rings to a longer or shorter interval, thereby using bandwidth more efficiently. You can prevent the sending of any MAC refresh packets if you don't need them (if you have FDDI but do not have a learning bridge or switch).



Note

This feature applies to HSRP running over FDDI only. You do not need to configure the MAC refresh interval if you have the **standby use-bia** command configured.

Internet Key Exchange Protocol

The Internet Key Exchange (IKE) Protocol is a key management protocol standard which is used in conjunction with the IPSec standard. IPSec is an IP security feature that provides robust authentication and encryption of IP packets.

IPSec can be configured without IKE, but IKE enhances IPSec by providing additional features, flexibility, and ease of configuration for the IPSec standard.

IKE is a hybrid protocol that implements the Oakley key exchange and Skeme key exchange inside the ISAKMP framework. (ISAKMP, Oakley, and Skeme are security protocols implemented by IKE.)

IP Directed Broadcast Changes

By default, IP directed broadcasts are no longer forwarded; they are dropped. However, you can enable IP directed broadcasts, and optionally specify an access list. Prior to Cisco IOS Release 12.0, IP directed broadcasts were forwarded by default. By dropping IP directed broadcasts, routers are less susceptible to denial-of-service attacks.

IP Host Backup Feature

The IP Host Backup feature permits a mainframe operating system to be moved from one mainframe to another without requiring a change to the router configuration at the time of the move.

Multiple mainframes can be connected to a single Channel Interface Processor (CIP) by means of an ESCON director. Often, these mainframes run using the Multiple Image Facility (MIF), which permits the physical machine to be divided into multiple logical partitions (LPARs). By defining an unused partition on another mainframe, a user can move the operating system from a failed mainframe or mainframe partition to the unused partition. By having multiple paths to each device, the move is accomplished without changing the mainframe software. This function also permits moving an IP stack between multiple operating system images.

On the CIP, each IP connection is treated as a physical device. The CIP does not support multiple paths to a single IP connection (or device). Prior to IP Host Backup, the router configuration had to be changed whenever the mainframe operating system was moved from one mainframe or LPAR to another.

**Note**

IP Host Backup does not provide single system image or automatic failover to a waiting backup application. Host operator action on the mainframe is required in these instances.

IPSec Network Security

IPSec is a framework of open standards developed by the Internet Engineering Task Force (IETF).

IPSec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec devices (“peers”) such as Cisco routers.

IPSec provides the following network security services:

- Privacy—IPSec can encrypt packets before transmitting them across a network.
- Integrity—IPSec authenticates packets at the destination peer to ensure that the data has not been altered during transmission.
- Authentication—IPSec peers authenticate the source of all IPSec-protected packets.
- Anti-replay protection—IPSec prevents capture and replay of packets, and helps protect against denial-of-service attacks.

With IPSec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as virtual private networks (VPNs), extranets, and remote user access.

IPSec services are similar to those provided by Cisco Encryption Technology, a proprietary security solution introduced in Cisco IOS Software Release 11.2; however, IPSec provides a more robust security solution, and is standards-based.

IPX Infrastructure Enhancements

IPX infrastructure enhancements include a new trace route capability for IPX to troubleshoot typical network connectivity problems. IPX infrastructure enhancements also include improvements to the following two features:

- IPX ping now has the diagnostics feature. This feature addresses diagnostic-related issues by accepting and processing unicast or broadcast diagnostic packets. It can ping other stations using the diagnostic packets and display the configuration information in the response packet.

**Note**

IPX ping already includes Cisco and Novell features.

- IPX Enhanced IGRP incremental SAP now has split horizon, which instructs the routers not to advertise the SAP to the same interface from where that SAP is received. It also allows you to display the Enhanced IGRP neighbor server table using regular expressions.

IPX infrastructure enhancements offer the following benefits:

- IPX trace route probes all the intermediate IPX routers and servers traversed along the path to the final destination; it measures round trip delays and displays results.
- IPX Enhanced IGRP incremental SAP split horizon eliminates the return of the update packets to the same interface from where that SAP is received.
- IPX ping default feature addresses diagnostic request and response packets.

The following restrictions apply to the IPX infrastructures enhancements:

- To use the IPX trace route feature, ensure that all intermediate routers respond to the socket number (0x874E) and process trace route requests; the target node processes trace route requests and as a last resort also processes diagnostic requests destined to the diagnostic socket (0x456).
- When the IPX ping default is set to diagnostic, you are not allowed to ping to broadcast address because Cisco routers do not have an application to collect, analyze, and account all diagnostic responses from numerous nodes. Cisco routers need Cisco IOS Release 12.0 software or later to use the IPX ping default feature.
- IPX Enhanced IGRP split horizon is off for WAN interfaces and on for LAN interfaces. The global default stays off. The interface setting takes precedence if the interface setting is modified or when both the global and interface settings are unmodified. The global setting is used only when the global setting is modified and the interface setting is unmodified.

MAC Address and Precedence Accounting

The MAC address accounting feature provides accounting information for IP traffic based on the source and destination MAC address on LAN interfaces. For example, with this feature you can determine how much traffic is destined for various peers at the network access points. This feature is currently supported on Ethernet and FDDI interfaces. The precedence accounting feature provides accounting information for IP traffic based on the precedence on any interface. This feature calculates the total packet and byte counts for an interface that receives or sends IP packets and sorts the results based on IP precedence. This feature is supported on all interfaces and subinterfaces and supports CEF, dCEF, flow, and optimum switching.

Modem Pooling for the Cisco AS5200

Cisco IOS Release 12.0 provides the modem pooling feature to the Cisco AS5200 access server. Modem pooling allows service providers to define, select, and use separate pools of modems within a single access server or router to provide different dial-in services. Modem allocation is based on the dialed number identification service (DNIS) and a predetermined number of modem ports based on DNIS.

There are a number of applications for using the call set up information, including DNIS/ANI, processing incoming call requests with CallerID, and selecting services to setup “automatically” for specified calls. These uses generally fall into two categories, those requiring allocation of a specific number of modems for a specific service, and those requiring allocation of specific physical modems.

Some wholesale service providers need to allocate a minimum (guaranteed) number of ports per customer and provide some level of extra (overflow) ports. Some service providers use different dial-in numbers for different wholesale customer service offerings. This is one way of differentiating between services or customers for port allocation.

MPPC—MS PPP Compression

Microsoft Point-to-Point Compression (MPPC) is a scheme used to compress Point-to-Point Protocol (PPP) packets between Cisco and Microsoft client devices. The MPPC algorithm is designed to optimize bandwidth utilization in order to support multiple simultaneous connections.

The MPPC algorithm uses a Lempel-Ziv (LZ) based algorithm with a continuous history buffer, called a dictionary.

MS Callback

The MS Callback feature provides client-server callback services for Microsoft Windows 95 and Microsoft Windows NT clients. MS Callback supports the Microsoft Callback Control Protocol (MSCB). MSCB is Microsoft's proprietary protocol that is used by Windows 95 and Windows NT clients. MS Callback supports negotiated PPP Link Control Protocol (LCP) extensions initiated and agreed upon by the Microsoft client. MS Callback is added to existing PPP Callback functionality. Therefore, if you configure your Cisco access server to perform PPP Callback using Cisco IOS Release 12.0(1) or later, MS Callback is automatically available.

Multicast Distributed Switching

Prior to multicast distributed switching (MDS), IP multicast traffic was always switched at the Route Processor (RP) in the Route Switch Processor (RSP)-based platforms. With this release, IP multicast traffic can be distributed switched on RSP-based platforms with VIPs.

Switching multicast traffic at the RP had disadvantages:

- The load on the RP increased. This affected important route updates and calculations (for BGP, among others) and could stall the router if the multicast load was significant.
- The net multicast performance was limited to what a single RP could switch.
- MDS solves these problems by performing distributed switching of multicast packets received at the line cards (VIPs in the case of RSP, and line cards in the case of GSR). The line card is the interface card that houses the VIPs (in the case of RSP) and the GSR line card (in the case of GSR). MDS is accomplished using a forwarding data structure called a Multicast Forwarding Information Base (MFIB), which is a subset of the routing table. A copy of MFIB runs on each line card and is always kept up to date with the RP MFIB table.

In the case of RSP, packets received on non-VIP IPs are switched by the RP. MDS can work in conjunction with Cisco Express Forwarding (CEF), unicast distributed fast switching (DFS), or flow switching. The multicast switching load is kept off the RP, improving the performance of the router.

Multihop Virtual Private Dialup Network (VPDN)

Multihop Virtual Private Dialup Network (VPDN) allows packets of the same file that are received by two different Home Gateways from a remote client to be recombined successfully in the first Home Gateway contacted.

Multihop VPDN solves the problem of packets generated from the same file being unable to traverse two VPDNs in route to the first router contacted. Typically, packets from the same file can be received by two different Home Gateways when the file is big enough that it requires multiple calls to be made to send the entire file.

Multilayer Switching

Multilayer Switching (MLS) provides high-performance Layer 3 switching for the Catalyst 5000 series LAN switches. MLS switches IP data packets between subnets using advanced application specific integrated circuit (ASIC) switching hardware. Standard routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Routing Information Protocol (RIP), and Intermediate System-to-Intermediate System (IS-IS), are used for route determination.

The Route Switch Module (RSM) performs route processing and central configuration and control for the Catalyst 5000 series switch. Routing services can also be provided by an externally attached router.

MLS also provides traffic statistics as part of its switching function. These statistics are used for identifying traffic characteristics for administration, planning, and troubleshooting. MLS uses NetFlow Data Export (NDE) to export the flow statistics.

Multiple ISDN Switch Types

The Multiple ISDN Switch Types feature allows you to configure more than one ISDN switch type per router. You can apply an ISDN switch type on a per interface basis, thus extending the existing global ISDN switch-type command to the interface level. This allows Basic Rate Interfaces (BRI) and Primary Rate Interfaces (PRI) to run simultaneously on platforms that support both interface types.

Multiprotocol over ATM

MPOA enables the fast routing of internetwork-layer packets across a nonbroadcast multi-access (NBMA) network. MPOA replaces multi-hop routing with point-to-point routing using a direct virtual channel connection (VCC) between ingress and egress edge devices or hosts. An ingress edge device or host is defined as the point at which an inbound flow enters the MPOA system; an egress edge device or host is defined as the point at which an outbound flow exits the MPOA system.

The following components are required for an MPOA network:

- MPOA Client (MPC)
- MPOA Server (MPS)
- Catalyst 5000 series ATM module
- LAN Emulation (LANE)

An MPC identifies packets sent to an MPS, establishes a shortcut VCC to the egress MPC, and then routes these packets directly over the shortcut VCC. An MPC can be a router or a Catalyst 5000 series ATM module. An MPS can be a router or a Catalyst 5000 series Route Switch Module/Versatile Interface Processor 2 (RSM/VIP2) with an ATM interface.



Note

Since the RSM/VIP2 can also be used as a router, all references to router in this document refer to both a router and the RSM/VIP2 with an ATM interface.

MPOA provides the following benefits:

- Eliminates multiple router hops between the source and the destination points of the ATM cloud by establishing shortcuts for IP packets and other protocol packets.
- Frees the router for other tasks by reducing IP traffic.
- Provides backward compatibility as an ATM network by building upon LANE, and can be implemented using both MPOA and LANE-only devices.

For MPOA to work properly, a LANE client must have an ELAN ID for all ELANs represented by the LANE client.

Named Method Lists for AAA Authentication and Accounting

In earlier Cisco IOS releases, only named authentication method lists were supported under Cisco's Authentication, Authorization, and Accounting (AAA) network security services. With Cisco IOS Release 11.3(3)T, AAA was extended to support both authorization and accounting named method lists. Named method lists for authorization and accounting function the same way as those for authentication.

Named method lists for AAA authorization and accounting allow you to define different methods for authorization and accounting and apply those methods on a per-interface or per-line basis.

National ISDN Switch Type

National ISDN Switch Types for Basic Rate and Primary Rate Interfaces introduces changes to ISDN switch types for Primary Rate Interfaces (PRI) and Basic Rate Interfaces (BRI) as follows:

- Adds a new switch type for PRI interfaces (`isdn switch-type primary-ni`).
- Changes the BRI `basic-ni1` switch type to `basic-ni` (`isdn switch-type basic-ni`).
- Removes the ISDN `vn2` switch type (`isdn switch-type vn2`) used in France. The existing `vn3` switch type (`isdn switch-type vn3`) supports French `vn2` switches.
- Removes the ISDN `basic-nwnet3` switch type (`isdn switch-type basic-nwnet3`) used in Norway. The `basic-net3` switch type (`isdn switch-type basic-net3`) supports Norway NET3 switches.
- Removes the ISDN `basic-nznet3` switch type (`isdn switch-type basic-nznet3`) used by New Zealand NET3 switches. The ISDN `basic-net3` switch type (`isdn switch-type basic-net3`) supports New Zealand NET3 switches.
- Adds the ability to configure outgoing PRI B channel selection for the T1 controller in ascending order (channel 1 to channel 23) or descending order (channel 23 to channel 1). Previously, the router selected a B channel for outgoing calls from the highest free channel in descending order. The E1 controller channel selection for ascending order is channel 1 to 31, and 31 to 1 for descending order.



Note

The command parser will still accept the following switch types: `basic-nwnet3`, `vn2`, and `basic-net3`; however, when viewing the NVRAM configuration using either the `show running configuration` or `write terminal` command, the `basic-net3` or `vn3` switch types are displayed respectively.

National ISDN Switch Types for Basic Rate and Primary Rate Interfaces provides the following benefits:

- Unlike previous custom implementations, such as `basic-5ess`, `basic-dms100`, `primary-5ess`, and `primary-dms100`, the National ISDN specification is designed to be switch independent. This increases flexibility in adapting to evolving standards and future enhancements.
- The ability to select PRI B channel order election for outgoing calls allows extended flexibility and compatibility with a variety of ISDN switch type service implementations. Additionally, this ability reduces ISDN switch misconfigurations, which can delay initial service activation.

NetFlow Switching Enhancements

The new `ip flow-cache active-timeout` configuration command lets you specify the timeout period for the NetFlow cache.

New and Changed Show Commands for the Cisco 2600 Series Routers

This release introduces new and changed **show** commands for the Cisco 2600 series routers.

The Cisco 2600 series routers is a new family of cost-effective, modular access routers designed to enable customers to easily adopt future technologies and scale to accommodate network expansion, thereby protecting technology investments. The Cisco 2600 series is a key component of Cisco's data/voice/video integration strategy, enabling corporate customers to consolidate data, voice, and video traffic to reduce costs, deploy new business applications, and improve network performance.

NFAS with D-channel Backup

The DMS100 and NI2 switch types have been added to the existing Non-Facility Associated Signaling (NFAS) with D Channel Backup feature. ISDN NFAS allows a single D channel to control multiple PRI interfaces. A backup D channel can be configured for use when the primary NFAS D channel fails. Once the channelized T1 controllers are configured for ISDN PRI, only the NFAS primary D channel must be configured; its configuration is distributed to all the members of the associated NFAS group.

NHRP Enhancements

There are two enhancements to Next Hop Resolution Protocol (NHRP) when it is running with BGP over ATM media:

- NHRP now works on Cisco Express Forwarding (CEF) platforms.
- On such platforms, you can now configure NHRP to initiate switched virtual circuits (SVCs) once a configured traffic rate is reached. Similarly, SVCs can be torn down when traffic falls to another configured rate.

A third enhancement can be applied to NHRP on any platform:

- When an interface is placed in NHRP server-only mode, you have the option to specify non-caching. In this case, NHRP does not store information in the NHRP cache, such as NHRP responses that could be used again. To save memory, the non-caching option is generally used on a router located between two other routers.

Cisco's implementation of NHRP now supports IETF's draft version 11 of "NBMA Next Hop Resolution Protocol (NHRP)."

NPE-200 Network Processing Engine

This release of the Cisco IOS software introduces the NPE-200 for Cisco 7200 series routers. The network processing engine maintains and executes the system management functions for Cisco 7200 series routers. The network processing engine also shares the system memory and environmental monitoring function with the I/O controller. The NPE-200 has an R5000 microprocessor that operates at an internal clock speed of 200 MHz, 4 MB of SRAM, and erasable programmable read-only memory (EPROM) for storing sufficient code for booting the Cisco IOS software.

OSPF LSA Group Pacing

The OSPF LSA group pacing feature allows the router to group together OSPF link state advertisements (LSAs) and pace the refreshing, checksumming, and aging functions. The group pacing results in more efficient use of the router.

Prior to the LSA group pacing feature, the Cisco IOS software would perform refreshing on a single timer, and checksumming and aging on another timer. In the case of refreshing, for example, the software would scan the whole database every 30 minutes, refreshing every LSA the router generated, no matter how old it was.

This problem is now solved by each LSA having its own timer. Again using the example of refreshing, each LSA gets refreshed when it is 30 minutes old, independent of other LSAs. So CPU is used only when necessary. However, LSAs being refreshed at frequent, random intervals would require many packets for the few refreshed LSAs the router must send out. That would be inefficient use of bandwidth. Therefore, the router delays the LSA refresh function for an interval of time instead of performing it when the individual timers are reached. The accumulated LSAs constitute a group, which is then refreshed and sent out in one packet or more. Thus, the refresh packets are paced, as are the checksumming and aging. The pacing interval is configurable; it defaults to 4 minutes, which is randomized to further avoid synchronization.

The router groups together OSPF LSAs and paces the refreshing, checksumming, and aging functions so that sudden hits on CPU usage and network resources are avoided. This feature is most beneficial to large OSPF networks.

OSPF Point to Multipoint

OSPF has two new features related to point-to-multipoint networks. One feature applies to broadcast networks; the other feature applies to nonbroadcast networks.

- On point-to-multipoint broadcast networks, there is no need to specify neighbors. However, you can specify neighbors with the `neighbor` command, in which case you should specify a cost to that neighbor.
- On point to multipoint nonbroadcast networks, you now use the `neighbor` command to identify neighbors. Assigning a cost to a neighbor is optional.

Before this feature, some OSPF point-to-multipoint protocol traffic was treated as multicast traffic. Therefore, the `neighbor` command was not needed for point-to-multipoint interfaces because multicast took care of the traffic. Hellos, updates and acknowledgments were sent using multicast. In particular, multicast hellos discovered all neighbors dynamically. However, some customers were using point-to-multipoint on nonbroadcast media (such as classic IP over ATM), so their routers could not dynamically discover their neighbors. This feature allows the **neighbor** command to be used on point-to-multipoint interfaces.

On any point-to-multipoint interface (broadcast or not), the Cisco IOS software assumed the cost to each neighbor was equal. The cost was configured with the `ip ospf cost` command. In reality, the bandwidth to each neighbor is different, so the cost should be different. With this feature, you can configure a separate cost to each neighbor. This feature applies to point-to-multipoint interfaces only.

This feature allows you to configure neighbors on point-to-multipoint interfaces and assign a cost to each neighbor. These capabilities allow the router to dynamically discover neighbors over nonbroadcast media and to prefer some routes over others by assigning different costs to neighbors.

PowerQuicc

This release provides baseline platform support for a new processor, the MPC860 PowerQUICC. It is used by the Cisco MC3810 and Cisco 2600 routers.

Protocol-Independent Multicast (PIM) Version 2

Protocol-Independent Multicast (PIM) Version 2 includes the following improvements over PIM Version 1:

- A single, active rendezvous point (RP) exists per multicast group, with multiple backup RPs. This compares to multiple active RPs for the same group in PIM Version 1.
- A bootstrap router (BSR) provides a fault-tolerant, automated RP discovery and distribution mechanism. Thus, routers dynamically learn the group-to-RP mappings.
- Sparse mode and dense mode are properties of a group, as opposed to an interface. We strongly recommend sparse-dense mode, as opposed to either sparse mode or dense mode only.
- PIM Join and Prune messages have more flexible encodings for multiple address families.
- A more flexible Hello packet format replaces the Query packet to encode current and future capability options.
- Register messages to an RP indicate whether they were sent by a border router or a designated router.
- PIM packets are no longer inside IGMP packets; they are stand-alone packets.

POS Command Enhancement

The new interface command **pos scramble-atm** enables SONET payload scrambling on a POS interface. SONET payload scrambling applies a self-synchronous scrambler ($x^{43}+1$) to the Synchronous Payload Envelope (SPE) of the interface to ensure sufficient bit transition density.

Quality of Service Policy Propagation via Border Gateway Protocol

The Quality of Service (QoS) policy propagation via Border Gateway Protocol (BGP) feature allows you to classify packets based on access lists, BGP community lists, and BGP autonomous system (AS) paths. The supported classification policies include Internet Protocol (IP) precedence setting and the ability to tag the packet with a QoS class identifier internal to the router (available in a future maintenance release of the software). After a packet has been classified, you can use other QoS features such as Committed Access Rate (CAR) and Weighted Random Early Detection (WRED) to specify and enforce business policies to fit your business model.

The QoS policy propagation via BGP feature has the following enhancements:

- QoS group ID—You can set an internal QoS group ID that can be used later to perform rate-limiting or weighted fair queuing based on the QoS group ID. In the previous release you could only set up to eight IP precedence level to classify packets. By setting the QoS group ID in addition to the IP precedence, you can now have more than eight classes on which to perform rate-limiting or weighted fair queuing.
- Source and destination address lookup—You can specify whether the IP precedence level or QoS group ID used is obtained from the source (input) address or destination (output) address entry in the route table. In the previous release you could only use the destination address. You can now specifying the input or output address.

BGP policy propagation provides the following benefits:

- Allows you to classify packets using access lists, community lists, and AS paths.
- Leverages BGP to distribute QoS policy to remote routers in your network.
- Allows ingress routers to prioritize incoming and outgoing traffic.
- Allows you to classify packets based on IP precedence or QoS group ID.

R1 Modified Signaling for the Cisco AS5200 and AS5300 Access Servers

Enabling R1 modified signaling allows a Cisco AS5200 or Cisco AS5300 universal access server to talk to central office trunks that also use R1 modified signaling. R1 signaling is an international signaling standard that is common to channelized T1/E1 networks; however, Cisco only has made this feature available in Taiwan. You can configure a channelized T1/E1 interface to support different types of R1 modified signaling, which is used in older analog telephone networks.

This type of signaling is not the same as ITU R1 signaling; it is R1 signaling modified for Taiwan specifically.



Note

In the future, R1 Modified Signaling will be supported by the Cisco AS5800 access server, and will be available in Turkey as well as Taiwan.

Cisco now supports native R1 signaling on both E1 and T1 interfaces. This version of software supports R1 signaling customized for Taiwan only. This feature allows enterprises and service providers to fully interoperate with the installed Taiwanese telecommunications standards, providing interoperability in addition to the vast array of Cisco's IOS troubleshooting and diagnostic capability. This will provide customers with a seamless, single-box, solution for their Taiwan signaling requirements.

RIF Pass-through in DLSw+

By default, DLSw+ terminates the RIF for Token Ring, terminates the LLC for all media types and forwards data only across a WAN with DLSw+ and TCP/IP headers. The RIF is a field in source-route bridged frames that indicates the SRB path the frame should take when traversing a Token Ring network. In the case of an explorer packet, the RIF is a field of the source-route bridged frame that indicates the SRB path that the SRB explorer has traversed so far. The RIF is limited to seven hop counts by the IBM standards. Because DLSw+ terminates the RIF at the virtual ring, the network's scalability increases because the hop count of the packet starts over, and the packet can traverse seven additional hops. Also, RIF termination simplifies network design because ring numbers no longer have to be unique throughout an entire enterprise.

RJ-45 Interface Support

Cisco 7200 series routers support a new I/O-controller with an RJ-45 interface. The optional Fast Ethernet port is configurable for use at 100 Mbps full-duplex or half-duplex operation (half duplex is the default). The Fast Ethernet port is equipped with either a single MII receptacle or an MII receptacle and an RJ-45 receptacle.

To support this new feature, the **media-type** interface command has been modified. The **media-type** interface command now supports two options:

- **100basex**—Specifies an RJ-45 100BaseX physical connection
- **mii**—Specifies a media-independent interface

When using the I/O controller that is equipped with an MII receptacle and an RJ-45 receptacle, only one receptacle can be configured for use at a time.

SNMP Inform Request

The SNMP Inform Requests feature allows routers to send inform requests to SNMP managers.

Routers can send notifications to SNMP managers when particular events occur. For example, an agent router might send a message to a manager when the agent router experiences an error condition.

SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response PDU. If the manager does not receive an inform request, it does not send a response. If the sender never receives a response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination.

Because they are more reliable, informs consume more resources in the router and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform may be retried several times. The retries increase traffic and contribute to a higher overhead on the network. Thus, traps and inform requests provide a trade-off between reliability and resources.

SNMP Manager Feature

The SNMP Manager feature allows a router to serve as an SNMP manager. As an SNMP manager, the router can send SNMP requests to agents and receive SNMP responses and notifications from agents. When the SNMP manager process is enabled, the router can query other SNMP agents and process incoming SNMP traps.

Security Considerations

Most network security policies assume that routers will be accepting SNMP requests, sending SNMP responses, and sending SNMP notifications. With the SNMP manager functionality enabled, the router may also be sending SNMP requests, receiving SNMP responses, and receiving SNMP notifications. Your security policy implementation may need to be updated prior to enabling this feature.

SNMP requests are typically sent to UDP port 161. SNMP responses are typically sent from UDP port 161. SNMP notifications are typically sent to UDP port 162.

SNMP Sessions

Sessions are created when the SNMP manager in the router sends SNMP requests, such as inform requests, to a host or receives SNMP notifications from a host. One session is created for each destination host. If there is no further communication between the router and host within the session timeout period, the session will be deleted.

The router tracks statistics, such as the average round-trip time required to reach the host, for each session. Using the statistics for a session, the SNMP manager in the router can set reasonable timeout periods for future requests, such as informs, for that host. If the session is deleted, all statistics are lost. If another session with the same host is later created, the request timeout value for replies will return to the default value.

Sessions consume memory. A reasonable session timeout value should be large enough that regularly used sessions are not prematurely deleted, yet small enough such that irregularly used, or one-shot sessions, are purged expeditiously.

Standard IP Access List Logging

The Cisco IOS software can now provide logging messages about packets permitted or denied by a standard IP access list. That is, any packet that matches the access list will cause an informational logging message about the packet to be sent to the console. The level of messages logged to the console is controlled by the logging console command. This capability was previously only available in extended IP access lists.

The first packet that triggers the access list causes a logging message right away, and subsequent packets are collected over 5-minute intervals before they are displayed or logged. The logging message includes the access list number, whether the packet was permitted or denied, the source IP address of the packet, and the number of packets from that source permitted or denied in the prior 5-minute interval.

You can monitor how many packets are being permitted or denied by a particular access list, including the source address of each packet.

System Controller FTP Server

The FTP Server feature configures a router to act as an FTP server. FTP clients can copy files to and from certain directories on the router. In addition, the router can perform many other standard FTP server functions.

When the router receives a request for an FTP connection, the FTP server process is started. The FTP server prompts for a username and password. After you supply a valid username and password, you can enter various commands.

The FTP server allows you to retrieve files, such as syslog files, from the disk file system on the router. Not all FTP commands are supported by this FTP server implementation.

System Controller Health Monitor

The Health Monitor feature monitors key performance attributes of the shelves managed by the system controller. The Health Monitor feature continually polls its managed shelves to obtain the information stored in the Health Monitor MIB. Management stations collect information for all the shelves from the system controller rather than by polling each shelf individually.

In addition, you can configure specific performance thresholds for all managed shelves through simple commands on the system controller. The system controller uses SNMP to automatically configure the following on each managed shelf:

- Expressions in the EXPRESSION-MIB to calculate the attributes
- RMON alarms to poll the attributes at specific intervals
- RMON events to send traps to the system controller when an attribute exceeds its specified threshold

When threshold traps are received by the system controller, they are converted to Health Monitor traps and sent to trap destinations configured in the system controller.

The Health Monitor feature provides the following benefits:

- Simplified configuration of SNMP-based monitoring functions. Entering a few commands on the system controller configures all of the managed shelves to send traps.
- Management systems poll only the system controller to get Health Monitor MIB data. The management systems do not have to poll the individual shelves. Thus, this feature reduces network traffic and system resources used by management systems.

System Controller Performance Data Collection

The Performance Data Collection feature allows a system controller to collect and store SNMP MIB data from its managed router and dial shelves. The system controller then serves as a central point for network management data collection.

The system controller collects the raw data from the managed shelves periodically, saves the data, and provides a single access point for a central network management application. The data can then be uploaded to a network management station using FTP or TFTP.

Performance data is stored on a disk local to the system controller. A new file is created each time the system controller collects data from a shelf.

The Performance Data Collection feature provides the following benefits:

- Remote network management stations can get performance data from one place, as a single file transferred via FTP or TFTP. This benefit reduces network traffic and resources in the management station because the station does not have to poll each individual shelf.
- The bulk transfer method of collecting data generates less traffic on the network than collecting the same amount of information using SNMP requests. The bulk transfer method also impacts the managed shelves less than SNMP polling.

System Controller Shelf Discovery and Autoconfiguration

The Shelf Discovery and Autoconfiguration feature allows a system controller to automatically discover new shelves and properly configure them to interact with the system controller. The system controller communicates with its managed shelves through the Shelf Discovery Protocol (SDP), which runs on top of UDP.

The Shelf Discovery and Autoconfiguration feature provides the following benefits:

- Control of multiple platforms from one location
- Easier method of configuring commands on all shelves
- Consolidated list of managed shelves and all interfaces on managed shelves

System Controller Syslog Server

The Syslog Disk Logging feature allows you to collect, store, and retrieve all managed shelf syslog messages through the system controller. The system controller receives syslog messages from managed shelves and stores these messages in subfiles on its disk.

In addition, this feature provides an enhanced method of viewing messages in the logging history table. Messages can be displayed based on host IP address, time received, and order received.

The Syslog Disk Logging feature provides the following benefits:

- The system controller provides one storage and retrieval location for syslog messages from multiple hosts on the network.
- You can display syslog messages based on time, hostname, or order received.
- Subfiles can store a large number of messages.
- Messages are preserved across system reloads. Without this feature, messages are stored in syslog history tables, which are lost when the system reboots.

System Controller Virtual Console

The Virtual Console feature allows you to access dial and router shelves connected to a system controller. During a system controller session, you can connect to a router or dial shelf at the same privilege level as the current system controller session.

By entering one command, you can Telnet directly to a shelf, provide a username and password, and then go to the same privilege level as the system controller.

The Virtual Console feature allows you to connect to all managed shelves through one session and switch between sessions easily. You do not have to reenable privileged EXEC mode every time you switch to another shelf session. This feature is useful when you need to do quick tasks on different managed shelves.

Tag Switching

Tag Switching combines the performance and capabilities of Layer 2 (data link layer) switching with the proven scalability of Layer 3 (network layer) routing. It enables service providers to meet challenges brought about by explosive growth, and provides the opportunity for differentiated services without sacrificing existing infrastructure. The Tag Switching architecture is flexible, allowing data to be transferred over any combination of Layer 2 technologies. Support is offered for all Layer 3 protocols, and scaling is possible beyond anything offered in today's networks.

Tag Switching can efficiently enable the delivery of IP services over an ATM switched network. It supports the creation of different routes between a source and a destination on a purely router-based Internet backbone. Service providers who use Tag Switching can save money and increase revenue and productivity.

Token Ring ISL Feature

The Token Ring Inter-Switch Link (TRISL) feature is a Cisco protocol for interconnecting multiple routers and switches and maintaining VLAN information as traffic goes between routers and switches. The TRISL feature provides a method to transport native Token Ring frames from multiple VLANs across a 100 MB Fast Ethernet link.

Cisco 7500 series, or Cisco 7200 series routers installed with any one of the following port adapters support the transmission of Token Ring frames from multiple VLANs across a 100 Mbps Fast Ethernet link:

- 2-port Fast Ethernet/ISL 100BaseTX
- 2-port Fast Ethernet/ISL 100BaseFX
- 1-port Fast Ethernet 100BaseTX
- 1-port Fast Ethernet 100BaseFX

The 2-port Fast Ethernet/ISL port adapters support frame sizes up to 17800 bytes, and the 1-port Fast Ethernet port adapters support a frame size of 1500 bytes.

TRISL provides a bridging technology between switches and routers that will transport traffic from both types of LANs.

The TRISL feature provides the following new functions for routers and switches:

- Inter-VLAN routing for fast switched IP and IPX routing protocols between Token Ring or Ethernet LANs for traffic with or without RIF.

- Source-route bridging (SRB), transparent bridging, source route translational (SRT), and source route translational (SR/TLB) bridging between ISL or TRISL VLANs and interfaces that are enabled for bridging. These interfaces can include Token Ring, Ethernet, FDDI, TR-LANE, Ethernet-LANE, and any other media with encapsulations that support transparent bridging.

TRISL uses a Fast Ethernet interface to provide connectivity between routers and switches or between switches and extends the VLAN capabilities of a switch by tagging the standard Token Ring frame with the necessary VLAN information.

For detailed information on how Token Ring switching is implemented on switches, refer to the *Catalyst Token Ring Switching Implementation Guide*.

Video Support on the Cisco MC3810

The Cisco MC3810 supports video traffic within a data stream in two ways:

- Video in Pass-Through Mode—Using this method, video traffic received from a video CODEC connected to a universal I/O serial port can be transported on a dedicated timeslot between systems using the TDM functionality of the T1/E1 trunk.
- Video over ATM AAL1—A serial stream from a video CODEC connected to a Cisco MC3810 on serial port 0 or 1 can be converted to ATM and transported across an ATM network using AAL1 Circuit Emulation Services (CES) encapsulation.



Note

Only V.35 cable is supported for video traffic over serial port 0 or 1.

VIP Enhancements

New privileged EXEC commands provide more information about the Versatile Interface Processor (VIP). The command **show controllers logging** displays logging information about a VIP. The command **show controllers tech-support** displays general information about a VIP when reporting a problem. The command **show controllers align** shows NULL pointer dereferences and misaligned accesses for a VIP.

Voice over ATM on the Cisco MC3810

Voice over ATM enables a Cisco MC3810 to carry live voice traffic (for example, telephone calls and faxes) over an ATM network. The Cisco MC3810 supports compressed Voice over ATM on virtual interface ATM 0 only. Voice over ATM uses AAL5 encapsulation, which is designed to accommodate bursty traffic.



Note

The Cisco MC3810 does not support ATM SVCs.



Note

When using ATM on the Cisco MC3810, the channel group, TDM group, and Channel Associated Signaling (CAS) functionality is not available on the Multiflex Trunk (MFT) because ATM uses all T1/E1 time slots.

Voice over Frame Relay on the Cisco MC3810

Voice over Frame Relay enables a Cisco MC3810 concentrator to carry live voice traffic (for example, telephone calls and faxes) over a Frame Relay network. Voice over Frame Relay on the Cisco MC3810 is supported on serial ports 0 or 1, and on the T1/E1 trunk.

**Note**

The Cisco MC3810 does not support Frame Relay SVCs.

Voice over HDLC on the Cisco MC3810

Voice over HDLC enables a Cisco MC3810 concentrator to carry live voice traffic (for example, telephone calls and faxes) back-to-back to a second Cisco MC3810. Voice over HDLC on the Cisco MC3810 is supported on serial ports 0 or 1, and the T1/E1 trunk.

Voice over IP

Voice over IP enables a Cisco router to carry live voice traffic (for example, telephone calls and faxes) over an IP network.

VPDN MIB

The Virtual Private Dialup Network (VPDN) Management Information Base (MIB) feature is intended to support all the tables and objects defined in the Cisco VPDN Management MIB for VPDN user sessions. VPDN system-wide information is available. This includes active VPDN tunnels, active user sessions in active VPDN tunnels, and failure history information, per username.

The VPDN MIB feature provides generic logging output for VPDN information, such as Layer 2 Forwarding Protocol (L2F). The syslog messages are generated to inform authentication or authorization errors, resource issues, and time-out events.

The VPDN MIB feature offers a mechanism to track failures of user calls in a VPDN system allowing SNMP retrieval of user call failure information, on a per user basis. The VPDN Syslog Facility feature offers real-time access to VPDN fault information.

Web Cache Control Protocol

The Web Cache Control Protocol feature transparently redirects HTTP requests from the intended server to a Cisco Cache Engine. When the Cisco Cache Engine receives the request, it attempts to service the request from its own cache. If the requested information is not present, the Cisco Cache Engine then makes a request to the web server to get the required information. After receiving the required information from the web server, the Cisco Cache Engine passes the information back to the client and possibly caches it to fill future requests.

xDSL Bridging

The x Digital Subscriber Line bridge support feature enables you to configure a router for intelligent bridge flooding for x digital subscriber line and other bridge applications.

New Hardware Features in Cisco IOS Release 12.0(1)

The following new hardware features are supported by the Cisco IOS Release 12.0.

1-Port ATM-25 Network Module for the Cisco 3600 Series

The 1-port asynchronous transfer mode (ATM-25) network module provides full 25.6-Mbps connectivity to an external asynchronous digital subscriber line (ADSL) modem or ATM switch for Cisco Series 3600 routers. This network module provides ATM traffic shaping for use with ADSL uplink speeds and protocol support for both permanent virtual circuit (PVC) and switched virtual circuits (SVC) environments. This network module provides full support for multiprotocol encapsulation over ATM Adaptive Layer 5 (RFC 1483), classic IP over ATM encapsulation (RFC 1577), and ATM User Network interface signaling.

1-Port HSSI Network Module

The Cisco 3600 series 1-port high-speed serial interface (HSSI) network module provides full-duplex connectivity at Synchronous Optical Network (SONET) OC-1/STS-1 (51.840 Mhz), T3 (44.736 MHz), and E3 (34.368 MHz) rates in conformance with the EIA/TIA-612 and EIA/TIA-613 specifications. The actual rate of the interface depends on the external data service unit (DSU) and the type of service to which it is connected. This 1-port HSSI network module can reach speeds of up to 52 Mbps in unidirectional traffic with 1,548-byte packets and 4,250 packets per second.

The 1-port HSSI network module provides the following benefits:

- Supports speeds up to 52 Mbps
- Supports a range of connectivity options: ATM, Frame Relay, PPP, and SMDS
- Supports EIA/TIA-612 and EIA/TIA-613 specifications at T3, E3, SONET OC1/STS-1 and NXT1 substrates

Channel Port Adapter

The Channel Port Adapter (CPA) expands the value of Cisco's Channel Interface Processor (CIP) solution. The CIP continues to be the industry's premier high-performance mainframe channel connect solution. The CPA extends the CIP architecture to customers requiring mid-range mainframe channel connectivity.



Note

The Cisco Mainframe Channel Connection (CMCC) product family includes the CIP on the Cisco 7500 series routers and the CPA on the Cisco 7200 series routers.

The CPA is a standard single-width port adapter supporting ESCON or parallel (also known as bus and tag) channel interfaces to IBM and IBM-compatible mainframes.

The CPA comes with a default of 16 MB of RAM. Customers may upgrade to 32 MB of RAM before the CPA is shipped from Cisco or as a field upgrade after the CPA has been installed.

The only difference between channel software applications (IP Datagram, Cisco SNA, TCP/IP Offload, TN3270 Server, and CMPC) running on the CIP and the CPA is performance. The CIP will typically have higher performance and capacity than the CPA because the CIP has more memory (128 MB of RAM compared to 32 MB for the CPA) and a faster internal bus.

Each CPA provides a single channel interface for Cisco 7200 series routers. In some situations, this eliminates the need for a separate front-end processor. The ESCON CPA contains a single ESCON I/O connector.

**Note**

The Cisco 7200 series router supports online insertion and removal (OIR), which allows you to install or remove port adapters while the system is operating.

The key benefits of the Cisco CPA in a Cisco 7200 series router are as follows:

- Cost-effective—Both the CPA and Cisco 7200 series routers provide industry-leading price performance.
- Extends current product offering to mid-range—Offers a midrange alternative to the high-performance, high-end Cisco CIP and Cisco 7500 series router solution.
- Simplified migration path—CPA and CIP microcode support the same features and applications, enabling seamless migration for network expansion.
- Flexibility of Cisco 7200 platform available to channel—Extends capabilities of the popular router platform for distributed mainframe networks in branch offices, redundant channel connectivity, and LAN-WAN/mainframe traffic consolidation.

The CPA is a high-speed port adapter. (A Fast Ethernet port adapter is an example of another type of high-speed port adapter.) A single Cisco 7200 series router can support up to three high-speed port adapters.

Cisco 7200

The Cisco 7202 is part of the Cisco 7200 series routers, which consists of the 2-slot Cisco 7202, 4-slot Cisco 7204, and the 6-slot Cisco 7206. The Cisco 7202 supports multiprotocol, multimedia routing and bridging with a wide variety of protocols and any combination of Ethernet, Fast Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode (ATM), Integrated Services Digital Network (ISDN), and serial media.

Network interfaces reside on port adapters that provide the connection between the router's Peripheral Component Interconnect (PCI) buses and external networks. The Cisco 7202 has two slots (slot 1 and slot 2) for the port adapters, one slot for an Input/Output (I/O) controller, and one slot for a network processing engine. You can place the port adapters in either of the two available slots.

**Note**

The Cisco 7202 does not support a mixture of AC- and DC-input power.

The Cisco 7202 provides the following features:

- Online insertion and removal (OIR)—Allows you to add, replace, or remove port adapters without interrupting the system or entering any console commands.
- Dual hot-swappable, load-sharing power supplies—Provide system power redundancy; if one power supply or power source fails, the other power supply maintains system power without interruption. Also, when one power supply is powered off and removed from the router, the second power supply immediately takes over the router's power requirements without interrupting normal operation of the router.
- Environmental monitoring and reporting functions—Allow you to maintain normal system operation by resolving adverse environmental conditions before any loss of operation.

- Downloadable software—Allows you to load new images into Flash memory remotely, without having to physically access the Cisco 7202 router, for fast, reliable upgrades.

Cisco uBR7246 Universal Broadband Router Features

Cisco uBR7246 universal broadband features enable the Cisco uBR7246 universal broadband router to communicate with a hybrid fiber coax (HFC) cable network via a Cisco MC11 cable modem card. Cisco MC11 cable modem cards allow you to connect cable modems on the HFC network to a Cisco uBR7246 in a Community Antenna Television (CATV) headend facility. The modem card provides the interface between the Cisco uBR7246 protocol control information (PCI) bus and the radio frequency (RF) signal on the HFC network.

The MC11 cable modem cards consist of the following components:

- One downstream cable F-connector port—The downstream port supports quadrature amplitude modulation (QAM) speeds of 64-QAM, or 6 bits per symbol. A symbol is the basic unit of modulation in CATV systems.
- One upstream cable F-connector port—The upstream port supports quadrature phase shift keying (QPSK) modulation, or 2 bits per symbol.
- Cable Media Access Control (MAC) unit—The cable MAC frames and encrypts the downstream signal for RF transmission and passes the signal to the downstream physical layer (PHY). Reverses the signal framing and encryption on the upstream signal from the upstream PHY.
- Downstream PHY unit—The downstream PHY generates a modulated, intermediate frequency (IF) output signal at a frequency of 44 MHz and passes the IF signal to an external IF to RF upconverter installed in the downstream path.
- Upstream PHY unit—The upstream PHY receives the modulated, upstream signal at a frequency of 5 MHz to 40 MHz and passes the signal to the cable MAC to remove the framing and encryption formats.
- Spectrum manager—The spectrum manager continuously monitors the noise in unused upstream channels. If the signal-to-noise ratio reaches an unacceptable level on a particular channel, the spectrum manager will automatically assign a new upstream channel to the cable modem using that channel. This feature is referred to as frequency agility.

E1-G.703/G.704 Serial Port Adapter

The E1-G.703/G.704 serial port adapters (PA-4E1G-120 and PA-4E1G-75) are available on Cisco 7500 series routers, Cisco 7200 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

JT2 6.3-MHz Serial Port Adapter

The JT2 6.3-MHz serial port adapter (PA-2JT2) is available on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

PA-12E/2FE Ethernet Switch 10BaseT and 100BaseTX Port Adapters

The PA-12E/2FE Ethernet switch 10BaseT and 100BaseTX port adapters are available on Cisco 7200 series routers. The PA-12E/2FE port adapter provides up to twelve 10-Mbps and two 10/100-Mbps switched Ethernet (10BaseT) and Fast Ethernet (100BaseTX) interfaces for an aggregate bandwidth of 435 Mbps, full-duplex.

PA-4R-D TR Port Adapter

The Dedicated Token Ring port adapter (PA-4R-DTR) is available on Cisco 7500 series routers, Cisco 7200 series routers, and Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

The PA-4R-DTR provides up to four IBM Token Ring or IEEE 802.5 Token Ring interfaces. Each Token Ring interface can be set for 4 Mbps or 16 Mbps half-duplex or full-duplex operation and can operate as a standard Token Ring station or as a concentrator port. The default for all interfaces is Token Ring station mode with half-duplex 16-Mbps operation. The PA-4R-DTR connects over Type 1 lobe or Type 3 lobe cables, with each interface providing an RJ-45 receptacle.

PA-A2 ATM-CES Port Adapter and Enhancement

The PA-A2 ATM-CES port adapters (PA-A2-4T1C-OC3SM, PA-A2-4T1C-T3ATM, PA-A2-4E1XC-OC3SM, PA-A2-4E1XC-E3ATM, PA-A2-4E1YC-OC3SM, and PA-A2-4E1YC-E3ATM) are available on Cisco 7200 series routers.

The ATM-CES port adapters now support the following new features:

- Available Bit Rate (ABR)—The ABR service category is specified in the ATM Forum Traffic Management Specification Version 4.0.
- Virtual Path Shaping—A virtual path (VP) is a logical association or bundle of virtual circuits (VCs).

In addition, all traffic shaping features available with the **atm pvc** interface command (*peak average burst*) are supported, and you can now configure the number of transmit channels for the interface with the **atm tx-channels** interface configuration command.

PA-A3-T3, PA-A3-E3, PA-A3-OC3MM, PA-A3-OC3SMI, PA-A3-OC3SMI Port Adapters

The enhanced ATM port adapter is a new generation of single-wide, single-port ATM port adapters available on Cisco 7200 series routers, Cisco 7500 series routers, and Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

The PA-A3 port adapters include five hardware versions that support the following standards-based physical interfaces:

- DS3
- E3
- OC-3c/STM1 multimode
- OC-3c/STM-1 single-mode intermediate reach
- OC-3c/STM-1 single-mode long reach

The PA-A3 port adapters support all Cisco IOS features and ATM-specific features available in Release 11.1(18)CC except for the available bit rate (ABR) permanent virtual circuit (PVC) features. ABR will be supported in a future maintenance release of Cisco IOS Release 11.1 CC.

The enhanced ATM port adapter supports the following features:

- Up to 4,096 total connections (open VCs)
- Up to 1024 segmentation and reassemblies (SARs) on the VIP2, up to 800 SARs on the NPE-200, and up to 200 SARs on the NPE-150
- ATM adaptation layer 5 (AAL5) for data traffic
- Physical interface: DS3 and E3 electrical coax (UNI3.x), (G.707), and (G.708); SONET/SDH optical fiber (OC-3 or STM-1)
- Traffic shaping on a per-VC basis
- Non-real time variable bit rate (NRT VBR) and unspecified bit rate (UBR) quality of service (QoS)
- Operation, Administration, and Maintenance (OAM) cells
- Online insertion and removal (OIR) on Cisco 7200 series routers

For more information on the PA-A3 port adapters, refer to the PA-A3 Enhanced ATM Port Adapter Installation and Configuration publication that accompanies the hardware.

PA-CT3/4T1 Channelized T3 Dual-Wide Port Adapter

The channelized T3 dual-wide port adapter (PA-CT3/4T1) is now available on Cisco 7200 series routers.

PA-E3 and PA-2E3 Serial Port Adapters

The PA-E3 and PA-2E3 serial port adapters are available on Cisco 7200 series routers, on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI). For information on interoperability guidelines for E3 serial port adapter DSUs, refer to the *E3 Serial Port Adapter Installation and Configuration* publication that ships with the product.

PA-H and PA-2H Revision B High-Speed Serial Interface Port Adapters

The PA-H Rev. B HSSI (High-Speed Serial Interface) port adapter is available on Cisco 7200 series routers, on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI). Although the PA-H was introduced in Cisco IOS Release 11.1(6)CA, the minimum Cisco IOS Release required by the PA-H is Release 11.1(12)CA or later, or 11.2(7)P or later. For more information on the PA-H and PA-2H port adapters, refer to the *Field Notice: HSSI Port Adapters* publication.

PA-T3 and PA-2T3 Serial Port Adapter

The PA-T3 and PA-2T3 serial port adapters are available on Cisco 7200 series routers, on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI). For information on interoperability guidelines for T3 serial port adapter DSUs, refer to the *T3 Serial Port Adapter Installation and Configuration* publication that ships with the product.

Switched 56K Digital Dial-in Over Channelized T1 and Robbed Bit Signaling

Internet service providers can provide switched 56-kbps access to their customers using a Cisco AS5300 or Cisco AS5200.

Switched 56K digital dial-in enables many services for ISPs. When using traditional ISDN PRI, the access server uses the bearer capability to determine the type of service. However when providing switched 56K over a CT1 RBS connection, the DS0s in the access server can be configured to provide either modem or 56-kbps data service. The dial-in user can access a 56-kbps data connection using either an ISDN BRI connection or a 2- or 4-wire switched 56-kbps connection. The telco to which the access server connects must configure its switches to route 56-kbps data calls and voice (modem) calls to the appropriate DS0.

Likewise, an enterprise can provide switched 56-kbps digital dial-in services to its full time telecommuters or small remote offices using ISDN PRI or a CT1 RBS connection.

Additional benefits:

- Enables ISDN BRI clients to connect to a Cisco AS5300 or Cisco AS5200 over switched 56K and T1 CAS.
- Provides switched 56K dial-in services over T1 CAS to remote clients that do not have access to ISDN BRI. For example, a remote PC making digital calls over a 2- or 4-wire switched 56-kbps connection and a CSU.

T1 CAS Support for the Cisco 3640 Digital Modem Network Module

The Digital Modem Network Module for the Cisco 3640 router is a high-density digital network module containing 6, 12, 18, 24, or 30 digital (MICA) modems. These modems, along with the T1 (or E1) port module, provide a direct digital connection to an Integrated Services Digital Network (ISDN) Primary Rate Interface (PRI) channel. The T1 CAS feature enables these network modules to support voice call transmission using channelized T1 lines (CT1) with channel associated signaling (CAS).

CAS is a form of signaling used on a T1 line. With CAS, a signaling element is dedicated to each channel in the T1 frame.

This type of signaling is sometimes called Robbed Bit Signaling (RBS) because a bit is taken out (or robbed) from the user's data stream to provide signaling information to and from the switch. The T1 CAS feature enables the modems on the Digital Network Modem Module to receive and transmit incoming and outgoing call signaling (such as on-hook and off-hook) through each T1 controller that is configured for a channelized T1 line.

Depending on the modem license you purchase with your Cisco 3640, the modems on the Digital Modem Network Module are either manageable or not manageable by Cisco IOS software commands. If the license you purchase includes this modem management capability, you can use the modem management commands to gather call and performance statistics at any time, even if there is an active call on the modem.

The Digital Modem Network Module for the Cisco 3640 provides the following benefits:

- The T1 CAS feature enables you to use T1 CAS signaling with the Digital Modem Network Module for the Cisco 3640 router.
- Enables you to support a mix of digital (ISDN) and basic telephone service analog modem calls over a single digital network interface.
- Modem management commands enable you to gather call and performance statistics.
- Supports 56-kbps modem connections via the K56 Flex and V.90 standards when the portware for these standards becomes available.

T1 CSU WIC for the Cisco 3600 and Cisco 1600 Series

The Cisco T1 data service unit/channel service unit (DSU/CSU) WAN interface card is an integrated, managed, T1 or fractional T1 WAN interface card. It provides nonchannelized data rates of 1 to 24 X 64 kbps or 1 to 24 X 56-kbps and follows ANSI T1.403 and AT&T Publication 62411 standards.

The Cisco DSU/CSU WAN T1 interface management features include the following:

- You can remotely configure the interface using Telnet and the Cisco IOS command line interface (CLI).
- For monitoring purposes, the router and DSU/CSU are manageable as a single Simple Network Management Protocol (SNMP) entity, using CiscoWorks or CiscoView. DSU/CSU statistics are accessed from the CLI.
- The SNMP agent supports the standard Management Information Base II (MIB II), Cisco integrated DSU/CSU MIB, and T1 MIB (RFC 1406).
- Loopbacks (including a manual button for a network line loopback) and bit error rate tester (BERT) tests are provided for troubleshooting.
- Test patterns, alarm counters, and performance reports are accessible using the CLI.
- The module has carrier detect, loopback, and alarm LEDs.

VIP2-50

The VIP2-50 is available on Cisco 7500 series routers, and on Cisco 7000 series routers with the 7000 Series Route Switch Processor (RSP7000) and 7000 Series Chassis Interface (RSP7000CI).

MIBs

Current MIBs

To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL:

<http://tools.cisco.com/ITDIT/MIBS/servlet/index>

If Cisco MIB Locator does not support the MIB information that you need, you can also obtain a list of supported MIBs and download MIBs from the Cisco MIBs page at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

To access Cisco MIB Locator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions found at this URL:

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

Deprecated and Replacement MIBs

Old Cisco MIBs will be replaced in a future release. Currently, OLD-CISCO-*MIBs are being converted into more scalable MIBs without affecting existing Cisco IOS products or network management system (NMS) applications. You can update from deprecated MIBs to the replacement MIBs as shown in Table 6.

Deprecated MIB	Replacement
OLD-CISCO-APPLETALK-MIB	RFC1243-MIB
OLD-CISCO-CHASSIS-MIB	ENTITY-MIB
OLD-CISCO-CPUK-MIB	In development
OLD-CISCO-DECNET-MIB	In development
OLD-CISCO-ENV-MIB	CISCO-ENVMON-MIB
OLD-CISCO-FLASH-MIB	CISCO-FLASH-MIB
OLD-CISCO-INTERFACES-MIB	IF-MIB CISCO-QUEUE-MIB
OLD-CISCO-IP-MIB	In development
OLD-CISCO-MEMORY-MIB	CISCO-MEMORY-POOL-MIB
OLD-CISCO-NOVELL-MIB	NOVELL-IPX-MIB
OLD-CISCO-SYS-MIB	Compilation of OLD*MIBs
OLD-CISCO-SYSTEM-MIB	CISCO-CONFIG-COPY-MIB
OLD-CISCO-TCP-MIB	CISCO-TCP-MIB
OLD-CISCO-TS-MIB	In development
OLD-CISCO-VINES-MIB	CISCO-VINES-MIB
OLD-CISCO-XNS-MIB	In development

Important Notes

This section contains important information about the use of your Cisco IOS Release 12.0 software.

Field Notices and Bulletins

For general information about the types of documents listed in this section, refer to the following document:

http://www.cisco.com/warp/customer/cc/general/bulletin/software/general/1654_pp.htm

- **Field Notices**—Cisco recommends that you view the field notices for this release to see if your software or hardware platforms are affected. If you have an account on Cisco.com, you can find field notices at http://www.cisco.com/warp/customer/tech_tips/index/fn.html. If you do not have a Cisco.com login account, you can find field notices at http://www.cisco.com/warp/public/tech_tips/index/fn.html.

- Product Bulletins—If you have an account on Cisco.com, you can find product bulletins at <http://www.cisco.com/warp/customer/cc/general/bulletin/index.shtml>. If you do not have a Cisco.com login account, you can find product bulletins at <http://www.cisco.com/warp/public/cc/general/bulletin/iosw/index.shtml>.
- *What's Hot for IOS Releases: Cisco IOS 12.0*—*What's Hot for IOS Releases: Cisco IOS 12.0* provides information about caveats that are related to deferred software images for Cisco IOS Release 12.0. If you have an account on Cisco.com, you can access *What's Hot for IOS Releases: Cisco IOS 12.0* at <http://www.cisco.com/kobayashi/sw-center/sw-ios.shtml> or by logging in and selecting **Software Center: Cisco IOS Software: What's Hot for IOS Releases: Cisco IOS 12.0**.
- *What's New for IOS*—*What's New for IOS* lists recently posted Cisco IOS software releases and software releases that have been removed from Cisco.com. If you have an account on Cisco.com you can access *What's New for IOS* at <http://www.cisco.com/kobayashi/sw-center/sw-ios.shtml> or by logging in and selecting **Software Center: Cisco IOS Software: What's New for IOS**.

Cisco IOS Release 12.0(23)—Caveat CSCdx40124

Deferral of Cisco 1600 Images

Four images in Cisco IOS Release 12.0(20) through Cisco IOS Release 12.0(23) were deferred because of a severe defect. This defect has been assigned Cisco caveat ID CSCdx40124. This caveat affects the following images:

- c1600-y-mz
- c1600-bnr2y-mz
- c1600-osy56i-mz
- c1600-bnor2sy56i-mz

With caveat CSCdx40124, a Cisco 1600 router may run out of processor memory even when the total memory (Processor + IO) is sufficient for the normal working of the router.

The software solution for these deferred images is Cisco IOS Release 12.0(24).

In order to increase network availability, Cisco recommends that you upgrade affected Cisco IOS images with the suggested replacement software images. Cisco will discontinue manufacturing shipment of affected Cisco IOS images. Any pending order will be substituted by the replacement software images.



Note

Please be aware that failure to upgrade the affected Cisco IOS images may result in network downtime.

The terms and conditions that governed your rights and obligations and those of Cisco, with respect to the deferred images will apply to the replacement images.

Cisco IOS Release 12.0(22)—Caveat CSCdx40124

Deferral of Cisco 1600 Images

Four images in Cisco IOS Release 12.0(20) through Cisco IOS Release 12.0(23) were deferred because of a severe defect. This defect has been assigned Cisco caveat ID CSCdx40124. This caveat affects the following images:

- c1600-y-mz
- c1600-bnr2y-mz
- c1600-osy56i-mz
- c1600-bnor2sy56i-mz

With caveat CSCdx40124, a Cisco 1600 router may run out of processor memory even when the total memory (Processor + IO) is sufficient for the normal working of the router.

The software solution for these deferred images is Cisco IOS Release 12.0(24).

In order to increase network availability, Cisco recommends that you upgrade affected Cisco IOS images with the suggested replacement software images. Cisco will discontinue manufacturing shipment of affected Cisco IOS images. Any pending order will be substituted by the replacement software images.



Note

Please be aware that failure to upgrade the affected Cisco IOS images may result in network downtime.

The terms and conditions that governed your rights and obligations and those of Cisco, with respect to the deferred images will apply to the replacement images.

Cisco IOS Release 12.0(21)—Caveat CSCdx40124

Deferral of Cisco 1600 Images

Four images in Cisco IOS Release 12.0(20) through Cisco IOS Release 12.0(23) were deferred because of a severe defect. This defect has been assigned Cisco caveat ID CSCdx40124. This caveat affects the following images:

- c1600-y-mz
- c1600-bnr2y-mz
- c1600-osy56i-mz
- c1600-bnor2sy56i-mz

With caveat CSCdx40124, a Cisco 1600 router may run out of processor memory even when the total memory (Processor + IO) is sufficient for the normal working of the router.

The software solution for these deferred images is Cisco IOS Release 12.0(24).

In order to increase network availability, Cisco recommends that you upgrade affected Cisco IOS images with the suggested replacement software images. Cisco will discontinue manufacturing shipment of affected Cisco IOS images. Any pending order will be substituted by the replacement software images.



Note

Please be aware that failure to upgrade the affected Cisco IOS images may result in network downtime.

The terms and conditions that governed your rights and obligations and those of Cisco, with respect to the deferred images will apply to the replacement images.

Cisco IOS Release 12.0(20)—Caveat CSCdx40124

Deferral of Cisco 1600 Images

Four images in Cisco IOS Release 12.0(20) through Cisco IOS Release 12.0(23) were deferred because of a severe defect. This defect has been assigned Cisco caveat ID CSCdx40124. This caveat affects the following images:

- c1600-y-mz
- c1600-bnr2y-mz
- c1600-osy56i-mz
- c1600-bnor2sy56i-mz

With caveat CSCdx40124, a Cisco 1600 router may run out of processor memory even when the total memory (Processor + IO) is sufficient for the normal working of the router.

The software solution for these deferred images is Cisco IOS Release 12.0(24).

In order to increase network availability, Cisco recommends that you upgrade affected Cisco IOS images with the suggested replacement software images. Cisco will discontinue manufacturing shipment of affected Cisco IOS images. Any pending order will be substituted by the replacement software images.

**Note**

Please be aware that failure to upgrade the affected Cisco IOS images may result in network downtime.

The terms and conditions that governed your rights and obligations and those of Cisco, with respect to the deferred images will apply to the replacement images.

Cisco IOS Release 12.0(17)—Caveat CSCdu10569

Deferral of AS5300 Boot Image

The c5300-boot-mz image has been deferred in Cisco IOS Release 12.0(17) because of a severe defect. This defect has been assigned Cisco Caveat ID CSCdu10569. The software solution for this defect is the c5300-boot-mz image in Cisco IOS Release 12.0(4)T1.

In order to increase network availability, Cisco recommends that you upgrade affected Cisco IOS images with the suggested replacement software images. Cisco will discontinue manufacturing shipment of affected Cisco IOS images. Any pending order will be substituted by the replacement software images.

**Caution**

Please be aware that failure to upgrade the affected Cisco IOS images may result in network downtime.

The terms and conditions that governed your rights and obligations and those of Cisco, with respect to the deferred images will apply to the replacement images.

Cisco IOS Release 12.0(14)—Caveat CSCdr91706 and IOS HTTP Vulnerability

A defect in multiple releases of Cisco IOS software will cause a Cisco router or switch to halt and reload if the IOS HTTP service is enabled, browsing to `http://router-ip/anytext/?` is attempted, and the enable password is supplied when requested. This defect can be exploited to produce a denial of service (DoS) attack.

The vulnerability, identified as Cisco bug ID CSCdr91706, affects virtually all mainstream Cisco routers and switches running Cisco IOS software releases 12.0 through 12.1, inclusive. This is not the same defect as CSCdr36952.

The vulnerability has been corrected and Cisco is making fixed releases available for free to replace all affected IOS releases. Customers are urged to upgrade to releases that are not vulnerable to this defect as shown in detail below.

This vulnerability can only be exploited if the enable password is known or not set.

You are strongly encouraged to read the complete advisory, which is available at <http://www.cisco.com/warp/public/707/ioshttpserverquery-pub.shtml>.

Cisco IOS Release 12.0(6a) Fixes Caveat CSCdm82691

Cisco IOS Software Release 12.0(6) was deferred due to a severe defect. This defect has been assigned Cisco Caveat ID CSCdm82691. This caveat affects only the Cisco 7000/7500 RSP platform images (image names beginning with `rsp-`) and Catalyst 5000 RSM platform images (image names beginning with `c5rsm-`).

A Cisco router running Cisco IOS Release 12.0(6) RSP code with a VIP based channelized E1, T1, 4T, or 8T port adapter, may display the following error messages because of the changes from CSCdm69227:

```
00:33:29: %RSP-3-BADBUFHDR: Invalid buffer ptr, address 58000000
-Traceback= 6025DD38 60266EB8 60237798
00:33:29: %RSP-3-INVRTN: Invalid return queue, next=0x61323238,
  hwidb=0x617A0A80, type=0x0
  queue_ptr=0x130, bufhdr_offset=0x0, id=0, bcast_id=0
  bufhdr 58007420: 00000000 00000128 01280000 00000000
-Traceback= 60266E74 6023779800
```

It was determined that Caveat CSCdm82691 was significant enough to merit a software rebuild.

This rebuild includes the caveat fix and is numbered 12.0(6a).

Cisco IOS Syslog Failure

Certain versions of Cisco IOS software may fail or hang when they receive invalid User Datagram Protocol (UDP) packets sent to their syslog ports (port 514). At least one commonly-used Internet scanning tool generates packets which can cause such problems. This fact has been published on public Internet mailing lists, which are widely read both by security professionals and by security crackers. This information should be considered in the public domain.

Attackers can cause Cisco IOS devices to repeatedly fail and reload, resulting in a completely disabled Cisco IOS device that will need to be reconfigured by its administrator. Some Cisco IOS devices have been observed to hang instead of failing when attacked. These devices do not recover until they are manually restarted by reset or power cycling. An administrator must personally visit an attacked, hung

device to restart it, even if the attacker is no longer actively sending any traffic. Some devices have failed without providing stack traces; some devices may indicate that they were “restarted by power-on”, even when that is not the case.

Customers should assume that any potential attacker is likely to know that the existence of this vulnerability and the ways to exploit it. An attacker can use tools available to the public on the Internet. An attacker does not need to write any software to exploit the vulnerability. Minimal skill is required. No special equipment is required.

Despite Cisco’s specifically inviting such reports, Cisco has received no actual reports of malicious exploitation of this vulnerability.

This vulnerability notice was posted on Cisco’s World Wide Web site:

<http://www.cisco.com/warp/public/770/iossyslog-pub.shtml>

This information was also sent to the following e-mail and Usenet news recipients:

- cust-security-announce@cisco.com
- bugtraq@netspace.org
- first-teams@first.org (includes CERT/CC)
- first-info@first.org
- cisco@spot.colorado.edu
- comp.dcom.sys.cisco
- nanog@merit.edu

Affected Devices and Software Versions

Vulnerable devices and software versions are specified in Table 5 of *Software Versions and Fixes*.

Affected versions include 11.3AA, 11.3DB, and all 12.0 versions (including 12.0 mainline, 12.0S, 12.0T, and any other regular released version whose number starts with 12.0), up to the repaired releases listed in Table 5. Cisco is correcting the vulnerability in certain special releases and will correct it in future maintenance and interim releases. See *Software Versions and Fixes* for details. Cisco intends to provide fixes for all affected IOS variants.

No particular configuration is needed to make a Cisco IOS device vulnerable. It is possible to filter out attack traffic using access lists. See *Workarounds* for techniques. However, except at Internet firewalls, the appropriate filters are not common in customer configurations. Carefully evaluate your configuration before assuming that any filtering you have already configured protects you against this attack.

The most commonly used or asked-about products are listed below. If you are unsure whether your device is running Cisco IOS software, log in to the device and issue the **show version** command. Cisco IOS software will identify itself simply as “IOS” or “Internetwork Operating System Software”. Other Cisco devices will not have the **show version** command, or they will identify themselves differently in their output. The most common Cisco devices that run Cisco IOS software include the following:

- Cisco routers in the AGS/MGS/CGS/AGS+, IGS, RSM, 800, ubr900, 1000, 2500, 2600, 3000, 3600, 3800, 4000, 4500, 4700, AS5200, AS5300, AS5800, 6400, 7000, 7200 (including the ubr7200), 7500, and 12000 series,
- Most recent versions of the LS1010 ATM switch,
- Some versions of the Catalyst 2900XL LAN switch,
- Cisco DistributedDirector.

Affected software versions, which are relatively new, are not necessarily available on every device listed above.

If you are not running Cisco IOS software, you are not affected by this vulnerability. The following Cisco devices are *not* affected:

- 700 dialup routers (750, 760, and 770 series) are not affected.
- Catalyst 1900, 2800, 2900, 3000, and 5000 LAN switches are not affected, except for some versions of the Catalyst 2900XL. However, optional router modules running Cisco IOS software in switch backplanes, such as the RSM module for the Catalyst 5000 and 5500, are affected.
- WAN switching products in the IGX and BPX lines are not affected.
- MGX (formerly known as the AXIS shelf) is not affected.
- Host-based software is not affected.
- Cisco PIX Firewall is not affected.
- Cisco LocalDirector is not affected.
- Cisco Cache Engine is not affected.

This vulnerability has been assigned Cisco bug ID CSCdk77426.

Solution

Cisco offers free software updates to correct this vulnerability for all affected customers, regardless of their contract status. However, because this vulnerability information has been disseminated by third parties, Cisco has released this notice before updates are available for all software versions. Table 5 gives Cisco's projected fix dates.

Make sure your hardware had adequate RAM to support the new software before installing it. Amount of RAM is seldom a problem when you upgrade within a major release (say, from 11.2(11)P to 11.2(17)P), but it is often a factor when you upgrade between major releases (say, from 11.2 P to 11.3 T).

Because fixes will be made available for all affected releases, this vulnerability will rarely, if ever, require an upgrade to a new major release. Cisco recommends very careful planning for any upgrade between major releases. Make certain no known bugs will prevent the new software from working properly in your environment.

Further upgrade planning assistance is available on the World Wide Web at:

<http://www.cisco.com>

Customers with service contracts should obtain new software through their regular update channels (generally via Cisco's World Wide Web site). They may upgrade to any software release, but they must remain within the boundaries of the feature sets they have purchased.

Customers without service contracts may upgrade to obtain only the bug fixes; they are not offered upgrades to versions newer than required to resolve the defects. In general, these customers will be restricted to upgrading within a single row of Table 5 below, except when no upgrade within the same row is available in a timely manner. Obtain updates by contacting one of the following Cisco Technical Assistance Centers (TACs):

- +1 800 553 2447 (toll-free from within North America)
- +1 408 526 7209 (toll call from anywhere in the world)
- tac@cisco.com

Give the URL of this notice (<http://www.cisco.com/warp/public/770/iossyslog-pub.shtml>) as evidence of your entitlement to a free update. Free updates for non-contract customers must be requested through the TAC. Please do not contact either “psirt@cisco.com” or “security-alert@cisco.com” for software updates.

Workarounds

You can work around this vulnerability by preventing any affected Cisco IOS device from receiving or processing UDP datagrams addressed to its port 514. This can be done either using packet filtering on surrounding devices, or by using input access list filtering on the affected IOS device itself.

If you use an input access list, that list should be applied to all interfaces to which attackers may be able to send datagrams. Interfaces include not only physical LAN and WAN interfaces, but virtual subinterfaces of those physical interfaces, as well as virtual interfaces and/or interface templates corresponding to GRE, L2TP, L2F, and other tunneling protocols.

The input access list must block traffic destined for UDP port 514 at any of the Cisco IOS device’s own IP addresses, as well as at any broadcast or multicast addresses on which the Cisco IOS device may be listening. Be sure to block both old-style “all-zeros” broadcasts and new-style “all-ones” broadcasts. It is not necessary to block traffic being forwarded to other hosts; only traffic actually addressed to the Cisco IOS device is of interest.

No single input access list works in all configurations. Know the effect of your access list in your specific configuration before activating it.

The following example shows a possible access list for a three-interface router, along with the configuration commands needed to apply the list. The example assumes input filtering is not needed, other than as a workaround for this vulnerability.

```
! Deny all multicasts, and all unspecified-net broadcasts, to port 514
access-list 101 deny udp any 224.0.0.0 31.255.255.255 eq 514
! Deny old-style unspecified-net broadcasts
access-list 101 deny udp any host 0.0.0.0 eq 514
! Deny network-specific broadcasts. This example assumes that all of
! the local interfaces are on the class B network 172.16.0.0, subnetted
! everywhere with mask 255.255.255.0. This will differ from network
! to network. Note that we block both new-style and old-style broadcasts.
access-list 101 deny udp any 172.16.0.255 0.0.255.0 eq 514
access-list 101 deny udp any 172.16.0.0 0.0.255.0 eq 514
! Deny packets sent to the addresses of our own network interfaces.
access-list 101 deny udp any host 172.16.1.1 eq 514
access-list 101 deny udp any host 172.16.2.1 eq 514
access-list 101 deny udp any host 172.16.3.3 eq 514
! Permit all other traffic (default would be to deny)
access-list 101 permit ip any any

! Apply the access list to the input side of each interface
interface ethernet 0
ip address 172.16.1.1 255.255.255.0
ip access-group 101 in

interface ethernet 2
ip address 172.16.2.1 255.255.255.0
ip access-group 101 in

interface ethernet 3
ip address 172.16.3.3 255.255.255.0
ip access-group 101 in
```

Listing all possible addresses—especially all possible broadcast addresses—to which attack packets might be sent is complicated. If you do not need to forward any legitimate syslog traffic received on an interface, you can block all syslog traffic arriving on that interface. Remember that blocking will affect traffic routed through the Cisco IOS device as well as traffic destined to the device; if the IOS device is expected to forward syslog packets, you will have to do the detailed filtering. Because input access lists impact system performance, install them with caution, especially on systems running very near their capacity.

Software Versions and Fixes

Many Cisco software images have been or will be specially reissued to correct this vulnerability. For example, regular released version 12.0(2) is vulnerable, as are interim versions 12.0(2.1) through 12.0(2.3). The first fixed interim version of 12.0 mainline software is 12.0(2.4). However, a special release, 12.0(2a), contains only the fix for this vulnerability and does not include any other bug fixes from later 12.0 interim releases.

If you are running 12.0(2), and want to fix this problem without risking possible instability presented by installing the 12.0(2.4) interim release, you can upgrade to 12.0(2a). Release 12.0(2a) is a “code branch” from the 12.0(2) base, which will merge back into the 12.0 mainline at 12.0(2.4).

Special releases, like 12.0(2a), are one-time, spot fixes, and they will not be maintained. Thus, the upgrade path from 12.0(2a) is to 12.0(3).

Table 5 specifies information about affected and repaired software versions.



Note

All dates within this table are subject to change.

Table 5 Affected and Repaired Software Versions

Cisco IOS Major Release	Description	Special Fix ¹	First Fixed Interim Release ²	Fixed Maintenance Release ³
Unaffected Releases				
11.2 and earlier—all variants	Unaffected early releases (no syslog server)	Unaffected	Unaffected	Unaffected
11.3, 11.3T, 11.3DA, 11.3MA, 11.3NA, 11.3WA, 11.3(2)XA	11.3 releases without syslog servers	Unaffected	Unaffected	Unaffected
Releases based on 11.3				
11.3AA	11.3 early deployment for AS58xx	11.3(7)AA2, 8-JAN-1999 ⁴	11.3(7.2)AA	11.3(8)AA, 15-FEB-1999
11.3DB	11.3 for Cisco NRP routing blade in Cisco 6400 xDSL DSLAM			11.3(7)DB2, 18-JAN-1999
Releases based on 12.0				
12.0	12.0 Mainline	12.0(2a), 8-JAN-1999	12.0(2.4)	12.0(3), 1-FEB-1999
12.0T	12.0 new technology early deployment	12.0(2a)T1, 11-JAN-1999	12.0(2.4)T	12.0(3)T, 15-FEB-1999
12.0S	ISP support; 7200, RSP, GSR		12.0(2.3)S, 27-DEC-1998	12.0(2)S ⁵ , 18-JAN-1999
12.0DB	12.0 for Cisco 6400 universal access concentrator node switch processor (lab use)			12.0(2)DB, 18-JAN-1999
12.0(1)W	12.0 for Catalyst 8500 and LS1010	12.0(1)W5(5a) and 12.0(1a)W5(5b) (LS1010 platform only)	12.0(1)W5(5.15)	12.0(1)W5(6) (platform support for Catalyst 8540M will be in 12.0(1)W5(7))
12.0(0.6)W5	One-time early deployment for CH-OC12 module in Catalyst 8500 series switches.	Unaffected; one-time release	Unaffected	Unaffected; general upgrade path is via 12.0(1)W5 releases.
12.0(1)XA3	Short-life release; merged to 12/0T at 12.0(2)T	Obsolete	Merged	Upgrade to 12.0(2a)T1 and/or to 12.0(3)T.
12.0(1)XB	Short-life release for Cisco 800 series; merged to 12.0T and 12.0(3)T	12.0(1)XB1	Merged	Upgrade to 12.0(3)T.

Table 5 *Affected and Repaired Software Versions (continued)*

Cisco IOS Major Release	Description	Special Fix¹	First Fixed Interim Release²	Fixed Maintenance Release³
12.0(2)XC	Short-life release for new features in Cisco 2600, Cisco 3600, ubr7200, ubr900 series; merged to 12.0T at 12.0(3)T.	12.0(2)XC1, 7-JAN-1999	Merged	Upgrade to 12.0(3)T
12.0(2)XD	Short-life release for ISDN voice features; merged to 12.0T at 12.0(3)T.	12.0(2)XD1, 18-JAN-1999	Merged	Upgrade to 12.0(3)T
12.0(1)XE	Short-life release	12.0(2)XE, 18-JAN-1999	Merged	Upgrade to 12.0(3)T

1. A special fix is a one-time release that provides the most stable immediate upgrade path.
2. Interim releases are tested less rigorously than regular, maintenance releases; interim releases may contain serious bugs.
3. Fixed maintenance releases are on a long-term upgrade path. Other long-term upgrade paths also exist.
4. All dates in this table are estimates, subject to change.
5. This entry is not a misprint. The 12.0(2.3)S interim release is available before the 12.0(2)S regular release in which the vulnerability is fixed.

Field Notices

Cisco Systems recognizes that to remain a leader in the internetworking industry, we must provide our customers with reliable, available and serviceable products.

To continually fulfill these goals, we provide our customers information necessary to their successful use of our products. This includes notification of any critical problems with Cisco products or technologies. These notifications, called “Field Notices,” include problem descriptions, safety or security issues, workarounds, and corrective actions necessary to ensure overall customer satisfaction and Cisco's highest overall product quality in the internetworking industry.

Field notices can be accessed on Cisco.com at the following address:
<http://www.cisco.com/warp/public/770/index.shtml>.

Using the Cisco MC3810 with the PSTN

This section describes important notes regarding use of the Cisco MC3810 with the Public Switched Telephone Network (PSTN).

Connections to a PSTN

Care should be exercised when connecting switched voice ports on the Cisco MC3810 directly to the PSTN because improper configurations can expose the corporate network to telephone fraud.

Switched Access from the PSTN

The Cisco MC3810 has the capability to connect a user from the PSTN directly to the corporate wide-area telephone network. As a phone switch, the Cisco MC3810 can be configured to switch the user to any location in that network, even remote locations that are connected again to another PSTN. However, the Cisco MC3810 does not provide any mechanism to restrict where users can call after they are connected. Without proper network design, this condition could result in the unauthorized use of the corporate network for making calls at the corporation's expense. To prevent this from occurring, Cisco does not recommend connecting a switched voice interface on the Cisco MC3810 directly to the PSTN. Instead, it should be connected to a PBX that implements a security scheme that prevents unauthorized use.

Non-Switched Calls

The same opportunity for illicit use does not exist for non-switched call types such as pass-through connections. Pass-through calls create a path to only a single location specified by the network administrator. For example, a pass-through connection might be used to pass a trunk from a PBX to the PSTN. In this case, the trunk on the PBX will always pass straight through the Cisco MC3810 to the PSTN. As a result, the necessary security is provided by the PBX.

Caveats for Cisco IOS Release 12.0

For a list of the software caveats that apply to Release 12.0, refer to *Caveats for Cisco IOS Release 12.0*. This caveats document is updated for every maintenance release and is located on Cisco.com and the Documentation CD-ROM. Caveats describe unexpected behavior or severe defects in Cisco IOS software releases. Severity 1 caveats are more serious; Severity 2 caveats are less serious.

**Note**

If you have an account with Cisco.com, you can also use the Bug Toolkit to find select caveats of any severity. To reach the Bug Toolkit, **log in** to Cisco.com and click **Service & Support: Software Center: Cisco IOS Software: BUG TOOLKIT**. Another option is to go to http://www.cisco.com/cgi-bin/Support/Bugtool/launch_bugtool.pl.
