DOC-GSR-8.3E 78-0843-05
August 13, 1992

CISCO SYSTEMS

Router Products Release Note for Software Release 8.3

This release note describes the features, modifications, and caveats for Software Release 8.3, including 8.3(1) through 8.3(5).

Complete documentation for Release 8.3 is contained in the Cisco Systems publication *Router Products and Configuration Reference* dated October 1991. A list and description of the current software versions available from Cisco Systems is included in the "Preface" to this document.

Preface

This preface contains version numbers of the router software currently available and Cisco Systems' service and support provisions. It also provides a list of documents that supplement this release note and procedures for obtaining online documents from Cisco via the File Transfer Protocol (FTP).

Current Software Version Levels

The table that follows describes the current software versions for Cisco router and TRouter™ software. Refer to these version numbers when ordering software updates.

Version	System	Description	ROMs
8.3(1-5)	GS3	CSC/3 Gateway Server Sets:	
		GS3-F	8
		GS3-BF	8
		GS3-FX	8
		GS3-BFX	8
8.3(1-5)	GS2	CSC/2 Gateway Server Sets:	
		GS2-R	8
		GS2-BR	8
		GS2-RX	8
		GS2-BRX	8
8.3(1-5)	TR3	CSC/3 TRouter Sets:	
		TR3-X	8
8.3(1-5)	TR2	CSC/2 TRouter Sets:	
		TR2-RX	8
8.3(1-5)	IGS	IGS Server Sets:	
		IGS-R	8
		IGS-BR	8
		IGS-RX	8
		IGS-BPRX	8
		IGS-BRX	8

Letter Key:
B—Bridging software
F—Standard system software with cBus complex
P—Protocol translation software
R—Standard system software which executes out of ROM
X—Standard and Commercial/DDN X.25 software

The software images that are run on a CSC/2 processor have been expanding as new feature code has been added for each software release. Some of these images (GS2-R, GS2-BR, GS2-RX, GS2-BRX, TR2-RX) have now exceeded the 1 MB-ROM capacity of the current CSC/2 processor boards. As a result, these images will be shipped on 2 MB-ROMs as of Software Release 8.3. This change requires an accompanying PAL change to support the addressing of the added megabyte. The new PAL (Cisco Part Number 17-0987-01) that Cisco provides for this support is compatible only with 2MB ROMs and will not support the use of 1 MB-ROMs.

The procedures for updating your system with the latest software version, including procedures for EPROM replacement, are contained in the Cisco Systems publication *Modular Products Hardware Installation and Reference*, October 1991.

Reference Documents

This release note supplements the following documents:

Router Products Configuration and Reference, October 1991
IGS Hardware Installation and Reference, October 1991
Modular Products Hardware Installation and Reference, October 1991
North American Customer Services Product Guide, February 1990

Warranty Information

All Cisco Systems products are covered under a limited factory warranty. This warranty covers defects in the hardware, software, or firmware. Refer to the Cisco Systems *Customer Services Product Guide* for more information on Cisco's warranty policy, or contact Customer Service at 1-800-553-NETS or (415) 326-1941.

Note: Warranty and other service agreements may differ for international customers. Contact your closest Cisco regional representative for more information.

Maintenance Agreements

Cisco Systems offers a Comprehensive Hardware Maintenance Agreement throughout North America that includes on-site remedial services, software support, a 24-hour emergency hot line, overnight parts replacement, and an escalation procedure. Cisco also offers software, maintenance, and advanced replacement services under a SMARTnet agreement for customers who desire those services. Noncontract maintenance services are provided at current time-and-materials rates. For more information, contact Customer Service at 1-800-553-NETS or (415) 326-1941.

Customer Support

Cisco's maintenance strategy is based upon customer-initiated service requests to the Cisco Systems Technical Assistance Center (TAC). The TAC coordinates all customer services, including hardware and software telephone technical support, onsite service requirements, and module exchange and repair.

The TAC is available Monday through Friday from 5:00 a.m. to 6:00 p.m. Pacific Coast time (excluding company holidays) at the numbers that follow. If you must return your Cisco equipment for repair or replacement, contact the TAC or a Cisco regional representative for more information.

Hardware and software support specialists who help diagnose and solve customer problems will be able to isolate and solve your problem much faster if you are prepared with the information they need (see the TAC escalation procedures page shipped with this product). When you call the TAC, have the following information ready:

- Chassis serial number
- Maintenance contract number
- Software version and hardware configuration

You can display your software version level and your hardware configuration by using the **show version** command.

Technical Assistance (TAC):

1-800-553-2447Fax:(415) 903-8787
(415) 688-8209 E-mail:tac@cisco.com

Sales, Orders, Questions, and Comments:

1-800-553-NETS (6387)Fax:(415) 903-8080
(415) 903-7208E-mail:csrep@cisco.com

Ordering Additional Cisco Publications

To order additional copies of a Cisco manual, contact your sales representative. (The Customer Order Number for each manual is located at the bottom of the title page.) Customer Service can provide you with the name of your sales representative if necessary.

1-800-553-NETS (6387)
E-mail: customer-service@cisco.com

Obtaining Cisco Technical Information Electronically

Cisco provides a directory of documents that you can access electronically using File Transfer Protocol (FTP). The directory includes such publications as product release notes, descriptions of Management Information Bases (MIBs), commonly used Requests for Comments (RFCs), and technical notes. The directory does not include electronic versions of Cisco technical manuals.

To obtain these technical documents, proceed as follows:

Step 1: At your server prompt, use the **ftp** command to connect to address *ftp.cisco.com.*

```
% ftp ftp.cisco.com
```

When you connect to the directory, you are greeted with an informational banner:

```
Connected to dirt.cisco.com.
220 dirt FTP server (Version 5.51.28 Mon Jan 13 17:51:58 PST 1992)
ready.
```

This is followed by a login prompt.

Step 2: Enter the word **anonymous** as your login name:

```
Name (ftp.cisco.com:cindy): anonymous
The system responds with this message:
331 Guest login ok, send ident as password.
Password:
```

Step 3: Enter your login name at the `Password:` prompt.

The following message and `ftp>` prompt appear:

```
230 Guest login ok, access restrictions apply.
ftp>
```

Step 4: To obtain a list of available files, enter **get README** at the `ftp>` prompt:

```
ftp> get README
200 PORT command successful.
150 Opening ASCII mode data connection for README (10093 bytes).
226 Transfer complete.
local: README remote: README
10307 bytes received in 0.17 seconds (59 Kbytes/s)
```

Step 5: Enter the **get** command and the full file name for each file you require.

Step 6: To exit FTP, use the **quit** command.

```
ftp> quit
221 Goodbye.
```

Note: In the FTP directory, the **ls** command does not accept wildcards; therefore, you cannot use this command to obtain a list of available files. To obtain a list of available files, you must use the README file.

Obtaining Information from Other Sources

This section describes how to obtain RFCs and technical standards.

Obtaining RFCs

Information about the Internet suite of protocols is contained in documents called *Requests for Comments*, or *RFCs*. These documents are maintained by Government Systems, Inc. (GSI). You can request copies by contacting GSI directly, or you can use the TCP/IP File Transfer Protocol (FTP) to obtain an electronic copy.

Contacting GSI

You can contact GSI through mail, by telephone, or through electronic mail.

Government Systems, Incorporated
Attn: Network Information Center
14200 Park Meadow Drive, Suite 200
Chantilly, Virginia 22021

1-800-365-3642
(703) 802-4535
(703) 802-8376 (FAX)

NIC@NIC.DDN.MIL
Network address: 192.112.36.5
Root domain server: 192.112.36.4

Obtaining an Electronic Copy

To obtain an electronic copy of an RFC via FTP, complete the following steps:

Step 1: At your server prompt, use the **ftp** command to connect to address *nic.ddn.mil*:

```
% ftp nic.ddn.mil
```

The following display appears, followed by a login prompt:

```
Connected to nic.ddn.mil.
220-*****Welcome to the Network Information Center*****
    *****Login with username "anonymous" and password "guest"
    *****You may change directories to the following:
      ddn-news          - DDN Management Bulletins
      domain            - Root Domain Zone Files
      ien               - Internet Engineering Notes
      iesg              - IETF Steering Group
      ietf              - Internet Engineering Task Force
      internet-drafts   - Internet Drafts
      netinfo           - NIC Information Files
      netprog           - Guest Software (ex. whois.c)
      protocols         - TCP-IP & OSI Documents
      rfc               - RFC Repository
      scc               - DDN Security Bulletins
220 And more.
```

Step 2: At the login prompt, enter the word **anonymous** as your login name:

```
Name (nic.ddn.mil:cindy): anonymous
The NIC responds with this message:
331 Guest login ok, send "guest" as password.
Password:
```

Step 3: Enter the word **guest** at the `Password:` prompt. The following message and `ftp>` prompt appear:

```
230 Guest login ok, access restrictions apply.
ftp>
```

Step 4: Use the **cd** command to change directories.

The following example illustrates how to change the RFC directory and obtain RFC 1158:

```
ftp> cd rfc
250 CWD command successful.
ftp> get rfc1158.txt
```

Step 5: To exit the FTP facility, enter the **quit** command at the `ftp>` prompt.

Release 8.3(1) Features and Enhancements

This section describes the major functions introduced in Release 8.3(1) of the router software.

Note: For descriptions of all commands, refer to the Command Summaries at the end of each chapter or section of the *Router Products and Configuration Reference* publication. The "Index" and "Table of Contents" indicate the appropriate chapters.

New Hardware Features

Release 8.3(1) introduces support for the High-Speed Serial-Port Communications Interface (HSCI) complex, which provides a connection for two new interfaces: HSA, which provides connection for the High-Speed Serial Interface (HSSI) specification, and ULA, which provides connection to UltraNet supercomputer environments.

The HSA interface provides a single, full-duplex synchronous serial connection capable of transmitting and receiving data at up to 52 Mbps. The HSSI specification is a de facto industry standard providing connectivity to DS3, E3, frame relay at DS3, and other high-speed wide-area services through a DSU or line termination unit.

The ULA interface product, which is available exclusively from Ultra Network Technologies, provides a fiber or coax interface to supercomputer environments through an Ultra Network Technologies hub at rates of up to 125 Mbps.

Note: Use of these devices requires installation of new microcode levels. Refer to Table 1 in the "Microcode Revisions" section for information on the mandatory upgrades required for these new Release 8.3 features.

System and Interface Configuration Features

Release 8.3(1) includes the following enhancements and changes to its interface configuration capabilities.

- Support has been added for the following new UltraNet interface command, as listed on page 6-43 of the *Router Products Configuration and Reference* manual:

  ultranet address *ultranet-mac-address*

- Support has been added for the new HSSI interface command, as described on page 7-13 of the *Router Products Configuration and Reference*. Enhancements to the **loopback** command now allow testing of Cisco's UltraNet hardware, testing of the HSSI interface at the applique, the DTE side of the DSU, the line side of the DSU, and the remote DSU. The format of this command follows:

[no] loopback {applique | dte | line | remote}

- Several changes have been made to the **setup** facility, including the addition of new prompts to allow users to leave the System Configuration Dialog and to configure the DEC MOP server feature. Several minor modifications have also been made to the onscreen prompts within the configuration command script.
- The following **service** commands have changed format for Release 8.3(1):

Old Format	New Format
service domain	**ip domain-lookup**
service ipname	**ip ipname-lookup**
service subnet-zero	**ip subnet-zero**

Note: The old formats for the service commands are accepted in configuration input, but the output of the **write terminal** or **show config** commands will display the new forms.

- The **tcp-keepalives** {*in* | *out*} keyword has been added to the list of **service** commands.
- Extensions have been added to the **banner** command to display a message-of-the-day banner and a banner upon opening an EXEC process or an incoming message.
- Addition of the **[no] exec-banner** line command allows users to enable or disable banner commands.
- Support for configurable buffer sizes has been added through the new **buffers** global command, as listed on page 4-4 of the *Router Products Configuration and Reference* manual.
- A new **description** interface subcommand has been implemented to add a description of an interface to a configuration file.
- Additional flags have been added to the **transmitter-delay** command to allow configuration for the IGS router and the new HSSI hardware.
- The **priority-list** *list* interface configuration command sets up priority queuing on a specified interface. Optional keywords for this command allow fine-tuning of the packet count and enable traffic priority to be assigned by access list and Ethernet type code access list number, as well as by origin or destination to FTP or UDP ports.
- The **error-threshold** command now provides a means to configure the frequency at which the error recount will be set as listed on page 7-11 of the *Router Products Configuration and Reference.*
- The **mtu** command has been added to allow adjustment of the default maximum packet size.

New Routing Configuration Features

Software Release 8.3(1) includes enhancements to Cisco's routing configuration capabilities. These modifications are described in the *Router Products Configuration and Reference manuals.*

Internet Protocol (IP) Routing

The following enhancements have been made to Cisco's IP routing implementation:

- IP autonomous switching support has been added to provide faster packet processing for AGS+ systems by allowing the cBus complex to switch packets independently, without interrupting the system processor.

  The new command follows:

  [no] ip route-cache [cbus]

Note: Customers who want to use autonomous switching must upgrade the microcode on their MEC, FDDI, and CCTL cards. (The "Microcode Revisions" section provides more information about this.) Because a significant number of components needs to be replaced on these boards, we recommend that users take advantage of Cisco's advance board replacement service to implement these upgrades. For more details on this service, contact Customer Service at 800-553-NETS (6387).

- For ICMP echo requests, support has been added to set the DF ("Don't Fragment") bit in the IP header and to report ICMP unreachables with code equal to "Fragmentation needed but DF set." For information on these capabilities, refer to the "IP Ping Command" section of the *Router Products Configuration and Reference* manual.
- IP header compression support has been implemented in accordance with RFC 1144. This feature allows compression of TCP/IP packets along HDLC serial links, and can be executed with the following command:

  [no] ip tcp header-compression [passive]

- Support has been added for multiple IP helper addresses per interface, executed through the command **ip helper command address** *address.*
- IP Path MTU Discovery (associated with the **ip mtu** command, as documented on page 13-18 of the *Router Products Configuration and Reference* manual) has been added to allow dynamic discovery of the maximum transmission unit of an internet path, according to RFC 1191.
- Support has been added for HP Probe Proxy to allow a route to respond to HP Probe Proxy Name requests. Users can enable this feature through the interface configuration subcommand:

  ip probe proxy

- EGP route time-out periods have been modified to occur in the correct intervals.

AppleTalk Routing

Cisco's AppleTalk protocol implementation has undergone the following modifications:

- Changes have been made to the command syntax to replace references to Phase 1 and Phase 2 with the corresponding references: nonextended (for Phase 1) and extended (for Phase 2). These changes are documented throughout the *Router Products Configuration and Reference* manual.
- The following global configuration commands have been added:

  [no] apple strict-rtmp

  [no] apple send-rtmp

  [no] apple proxy-nbp

  appletalk iptalk-baseport

  appletalk timers (documented on page 10-15 of the *Router Products Configuration and Reference*)
- The following interface subcommands have been added:

  appletalk iptalk (documented on page 10-13 of the *Router Products Configuration and Reference*)

  [no] apple distribute-list {in | out}

  appletalk discovery (documented on page 10-13 of the *Router Products Configuration and Reference*)
- The following **clear** commands have been added:

  clear apple neighbors

  clear apple route

ISO CLNS Routing

Release 8.3 includes the following modifications to Cisco's ISO CLNS protocol implementation:

- The maximum number of ISO IGRP routing processes has been increased to ten (previously, the limit was six).
- Interfaces which are running ISO IGRP can now be restricted to sending routing updates for level 2 only. This feature can be defined through the following command:

  clns router igrp ***tag*** **level2**

Apollo Routing

- Cisco's Apollo protocol implementation now supports access lists which can be referenced by name through use of the command **apollo access-list**.

Banyan VINES Routing

Release 8.3(1) includes the following changes to Cisco's Banyan VINES protocol implementation:

- The **vines serverless** command has been added to allow configuration of a network without a server.

IBM Connectivity Features

There are several new additions to Cisco's support for IBM connectivity environments with Release 8.3(1). (For information on source-route bridging, refer to the "Bridging" part of the reference manual.)

SDLC Transport (Serial Tunnel)

Software Release 8.3(1) introduces support for Serial Tunnel (STUN) functionality, also known as SDLC Transport, for encapsulating SDLC-framed traffic into IP packets and routing them over any IP-supported media through use of the TCP transport mechanism.

- Support is offered for the following STUN global commands:

  [no] stun peer-name *ip-address*

  [no] stun poll-interval *milliseconds*

  [no] stun primary-pass-through *seconds*

  [no] stun protocol-group *group-number protocol*

  [no] stun schema *name* **offset** *constant-offset length address-length* **format** *format-keyword*

- Support is offered for the following STUN interface subcommands:

  encapsulation stun

  [no] stun group *group-number*

  [no] stun proxy-poll address *address* **modulus** *modulus* **{primary | secondary}**

  [no] stun proxy-poll address *address* **discovery**

  [no] stun route all tcp *ip-address*

  [no] stun route all interface serial *interface-number*

  [no] stun route all interface serial *interface-number* **direct**

  [no] stun route address *address-number* **tcp** *ip-address*

  [no] stun route address *address-number* **interface serial** *interface-number*

  [no] stun route address *address-number* **interface serial** *interface-number* **direct**

NetBIOS

- Support for NETBIOS Access Filters has been added to control packets transmitted across a Token Ring bridge using the NETBIOS interface. Cisco has implemented two types of filters: one for source and destination station names and one for arbitrary byte patterns in the packet itself. The new commands follow:

  [no] netbios access-list host *name* **{permit | deny}** *pattern*

  [no] netbios input-access-filter-host *name*

  [no] netbios output-access-filter-host *name*

  [no] netbios access-list bytes *name* **{permit | deny}** *offset pattern*

  [no] netbios input-access-filter bytes *name*

  [no] netbios output-access-filter-bytes *name*

WAN Features

With Software Release 8.3(1), Cisco introduces new and enhanced support for WANs.

- Frame relay is now supported as an encapsulation for the routing of IP, DECnet, AppleTalk, XNS, Novell, VINES, and ISO CLNS protocols, and for transparent bridging through use of the command **frame relay map**. SDLC Transport (also known as serial tunneling, or STUN) and source-route bridging (SRB) are also supported over frame relay by virtue of their encapsulation within TCP/IP.
- Dial backup functionality has been added to provide protection against WAN downtime by allowing configuration of a backup serial line via a circuit-switched connection. Support has been added for the following Dial Backup commands:

  [no] backup delay {*enable-delay* **|never**} {*disable-delay* **|never**}

  [no] backup interface *interface-name*

  [no] backup load {*enable-threshold* **|never**} {*disable-load* **|never**}

- Support is provided for Switched Multimegabit Data Service (SMDS), a packet-switched WAN service provided by the Regional Bell Operating Companies (RBOCs) and other telephone service carriers. Cisco provides an SMDS interface at T1 rates. The following SMDS interface subcommands, as listed on page 8-53 of the *Router Products Configuration and Reference,* have been added:

  encapsulation smds

  [no] smds address *smds-address*

  [no] smds att-mode

  [no] smds enable-arp

  [no] smds multicast *protocol-type smds-group-address*

  [no] smds static-map *protocol-type protocol-address smds-address*

Network Management Features

SNMP

Cisco's SNMP support has undergone several changes.

- Support has been added for SNMP MIB II (RFC 1157). Instructions for the retrieval of the MIB II document are included at the end of this section on "SNMP."
- The global configuration command **snmp-server system-shutdown** has been added.

Miscellaneous Enhancements

- A new **clear counters** EXEC command has been added to clear interface counters.

Bridging Features

Cisco has added new bridging support with Release 8.3(1), including enhancements to transparent bridging and Source-Route Bridging (SRB).

Transparent Bridging

Release 8.3(1) includes the following changes to Cisco's software support for transparent bridging:

- For the IEEE spanning tree, multiple spanning-tree domains are now supported. Domains are given a value from 1 to 10 and are specified with the following command:

  bridge *group* **domain** *domain-number*

- Support has been added to filter LAT frames to allow the selective inclusion or exclusion of LAT multicast service announcements on a per-interface basis. The new commands follow:

  Global:

  bridge *group* **lat-service-filtering**

  Interface:

  bridge-group *number* **input-lat-service-deny** *grouplist*

  bridge-group *number* **input-lat-service-permit** *grouplist*

  bridge-group *number* **output-lat-service-deny** *grouplist*

  bridge-group *number* **output-lat-service-permit** *grouplist*

- The **show span** command has been enhanced to display LAT group code filtering.

- Transparent bridging software has been modified to allow bridging of packets in X.25 frames. The **x25 map** command has been modified to allow this capability.
- Transparent bridging software now supports bridging of packets over frame relay networks. This feature works on networks that support a multicast facility as well as those that do not support multicasts. The **frame-relay map** interface subcommand has been modified to allow this capability.

Source-Route Bridging

Following is a list of the Release 8.3(1) changes to Cisco's Source-Route Bridging software.

- The **multiring** command has been extended to enable collection and use of routing information fields (RIFs) for all protocols. The software now allows per-protocol specification on a given interface to use multiring protocols. The protocols supported within this feature are Apollo Domain, AppleTalk, ISO CLNS, DECnet, IP, Novell IPX, Banyan VINES, and XNS.
- A new command has been added to limit the size of the backup queue for remote source-route bridging. This command controls the number of packets that can wait for transmission to a remote ring without being discarded.

  [no] source-bridge tcp-queue-max *number*

  The **show source-bridge** EXEC command now displays the queue length.
- The command **source-bridge remote-peer** may now be completed with an optional new keyword, **lf** *size*. This keyword allows the maximum-size frame to be sent to the remote peer to be specified.

Documentation Enhancements

As of Software Release 8.3(1), Cisco's documentation set has a new format that includes the following changes:

- A comprehensive error message appendix has been added that includes error messages for all Cisco products.
- Manuals have been reorganized to support specific tasks. Each new software manual provides sections on using the **setup** command facility, using the system, configuring the system, system management, and protocol-specific configuration. The new manuals also provide summaries of relevant commands with each chapter.

Obsolete Commands and Capabilities

This section lists the commands and capabilities of the Cisco router software which are no longer supported as of Release 8.3.

- Cisco's support for the Banyan VINES protocol no longer includes the **vines propagate** command.

- Cisco's support for the AppleTalk protocol no longer includes the **show apple detailed** command.
- Cisco's support for the AppleTalk protocol no longer includes the **show apple lock** command.
- Cisco's support for frame relay no longer includes the command, **frame-relay dlci-bits**
- Cisco's support for SNMP no longer includes X.25 virtual circuits clear traps.
- The **show priority** command is no longer supported in Release 8.3.

Release 8.3(2) Features and Enhancements

Enable and Console Passwords and the SNMP Community String

With Software Release 8.3(2), the software no longer allows the enable password or the console password to be used as the community string for SNMP.

New Routing Configuration Features in Release 8.3(2)

Software Release 8.3(2) includes enhancements to Cisco's routing configuration capabilities. These modifications are described in the *Router Products Configuration and Reference* manual.

Internet Protocol (IP) Routing

The following enhancements have been made to Cisco's IP routing implementation:

- Cisco has changed the defaults for commands that configure address resolution using proxy ARP and probe ARP. The defaults in Release 8.3(2) are as follows:

  no ip proxy arp

  no arp probe

- Cisco has improved support for HP Probe Proxy, first introduced in Release 8.3(1), to allow a route to respond to HP Probe Proxy Name requests. Users enable this feature through the interface configuration subcommand:

  ip probe proxy

Changes in HP Probe Proxy Behavior

Because the defaults are now the **no arp probe** and **no ip proxy arp** commands, you must specifically configure the **arp probe** command on all interfaces that will support HP probe proxy. Other changes include the following:

- Unsolicited probe replies are now cached. This improves user response time when starting sessions and helps eliminate unnecessary probeVNA exchanges.
- ARP probe is now supported over both Ethernet and IEEE encapsulation. Ethernet encapsulation is used whenever possible.
- DTC probe still needs to be bridged. The router notices the difference between DTC probes and other probes and will not answer DTC probes.
- The only limitation to the proxy table is the amount of memory in the router and the amount of NVRAM, if the proxy table is stored there. Alternatively, the proxy table can be netbooted after the router reloads.

Release 8.3(5) Features and Enhancements

IGS/TR Functionality Supported

Software Release 8.3(5) adds the ability to perform remote source-route bridging over X.25 on the IGS/TR platform.

The operational ring speed for the ISG Token Ring connector is set with the following configuration command:

ring-speed *speed*

The argument *speed* is **4** for 4 Mbps, or **16** for 16 Mbps.

Additional User Notes

Token Ring Restarts

If the system receives an indication of a cabling problem from a CSC-R16 Token Ring interface, that interface is placed in a reset state. The system does not attempt to restart the interface. To restart the interface, correct the cabling problem and use the **clear interface** command to reset it.

The system functions in this manner because periodic attempts to restart the Token Ring interface have drastic effects on the stability of routing tables, and sometimes on the stability of Token Ring networks themselves.

Netboot Restrictions

It should be noted that netbooting over X.25 or frame relay is not allowed to a broadcast address. You must specify the address of a server host to successfully netboot the system files. Use an off-net map entry of the destination. This means that you cannot simply have an X.25 or frame relay map entry for the next hop router. You need a map entry (use the **x25 map** or **frame-relay map** commands) for the host from which you will boot, even if that host is not on a directly-connected network.

X.25 Example

The **x25 map** command is used to map an IP address into an X.121 address. There *must* be an **x25 map** command which matches the IP address given on the **boot system** command line. In order to netboot over X.25, the address of the system from which to netboot *must* be given explicitly, and an **x25 map** entry must exist for that site, as the following example illustrates.

```
boot system gs3-bfx.83-2.0 131.108.13.111
!
interface Serial 1
ip address 131.108.126.200 255.255.255.0
encapsulation X25-DCE
x25 address 10004
x25 map IP 131.108.13.111 10002 BROADCAST
lapb n1 12040
clockrate 56000
```

Frame Relay Example

If file *gs3-bfx* is to be booted from a host with IP address *131.108.126.2*, the following would need to be in the configuration:

```
boot system gs3-bfx 131.108.126.2
!
interface Serial 0
encapsulation frame-relay
frame-relay map IP 131.108.126.2 100 broadcast
```

SMDS Interoperability

Routers running version 8.3(1) cannot interoperate over SMDS with routers running 8.3(2) or later software.

Subnetting the Same IP Address Across X.25

In order to configure load sharing across multiple serial lines, the entries for all the adjacent interface IP addresses need to be included in the **x25 map** command for each serial interface.

As an example, two routers, Cisco A and Cisco B, each with two serial interfaces, would require the following configuration files to allow subnetting the same IP address:

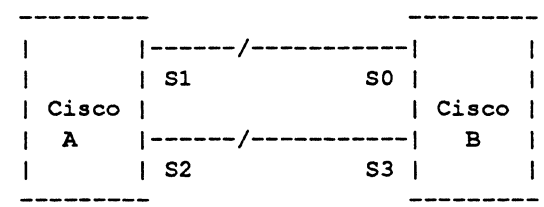

Cisco A

```
interface serial 1
ip 131.108.170.1 255.255.255.255.0
x25 address 11
x25 map ip  131.108.170.3 13
x25 map ip  131.108.170.4 13

interface serial 2
ip 131.108.170.2 255.255.255.255.0
x25 address 12
x25 map ip  131.108.170.4 14
x25 map ip  131.108.170.3 14
```

Cisco B

```
interface serial 0
ip 131.108.170.3 255.255.255.255.0
x25 address 13
x25 map ip 131.108.170.1 11
x25 map ip 131.108.170.2 11

interface serial 3
ip 131.108.170.4 255.255.255.255.0
x25 address 14
x25 map ip 131.108.170.2 12
x25 map ip 131.108.170.1 12
```

AppleTalk over FDDI

There is an interaction between AppleTalk and FDDI when mixing system version 8.2(4) (only) and later versions. If a router has both MCI Ethernet and FDDI interfaces while running 8.2(4), runts may be generated by the MCI interfaces when packets are sent from a router running a version later than 8.2(4) across the FDDI to the 8.2(4) router and forwarded via the MCI Ethernet interface(s).

The solution is to upgrade all 8.2(4) routers with FDDI and MCI Ethernet interfaces before upgrading any other routers on the FDDI ring in question. A work around is to disable AppleTalk fast switching on the MCI Ethernet interfaces of the affected 8.2(4) routers. This is done with the command **no apple route-cache**.

Novell IPX Packet Sizes

With release 8.3(2), Cisco's Novell IPX implementation supports packet sizes of more than 576 bytes on media that are capable of carrying packets of that size. Until recently, no Novell end node would send or accept a packet larger than this size; this has changed in newer Novell software. The software now accepts Novell IPX packets up to the maximum size allowed on the media. [CSCdi04193]

XNS Ungerman-Bass

The "Configuring Ungermann-Bass Net/One XNS" section (page 19-7 in the *Router Products Configuration and Reference*) states that netbooting does not work in the current Cisco software; however, there is a work around for this restriction. Since Ungermann-Bass devices have the ability to boot from another protocol that cannot be routed, you can use bridging to pick up the network identifier (ID). The steps for this work around follow:

Step 1: In order for the NIU to netboot correctly across a Cisco router, you need to bridge 0x7000 (etype 7000) to 0x7005 inclusively.

Step 2: The download server (DLS) on the NetDirector has to run with the **-Bytes** option on. This option causes the NIUs to receive their XNS network ID as it is configured in their LC file. The default startup for a download server causes an NIU to use the same XNS network ID as the download server.

Step 3: To locate service names *not* on the same segment, the users must type ***service_name** instead of just **service_name**, so that the NIU will do an *all-xns-net-broadcast* instead of just a *local-xns-net broadcast* (XNS type 4 destined to *-1.ffff.ffff.ffff*).

Release 8.3 Caveats

The following items identify unexpected behavior of the Cisco router software which is not included in the 8.3 release documentation.

We have introduced an internal reporting system for tracking modifications and caveats. For your reference, identification numbers follow the description of the caveat or modification.

Internet Protocol (IP)

IP network 0.0.0.0 is always a candidate default network. The asterisk marking a default candidate is not shown when this network is displayed in response to the **show ip routes** command. This problem has no operational impact. [CSCdi02203]

Use of the **ip security strip** command to strip IP security options causes packet corruption. [CSCdi02286]

The IP maximum transmission unit (MTU) for an interface is not adjusted upward when the physical MTU is changed to a larger-than-default value. [CSCdi02684]

If IP routing is disabled, IP packets are not correctly encapsulated for serial lines on which bridging is not enabled. Correct behavior for the software would be to encapsulate IP packets as bridged Ethernet packets across bridged networks, and in native HDLC form for nonbridged connections. [CSCdi02692]

CLNS

When a CLNS NET is removed from a router's configuration, the router may continue indefinitely sending intermediate system hellos using that NET. The work around is to reload the router when the NETs are changed. [CSCdi02578]

The **clns configuration-time** command has no effect. [CSCdi02852]

Fast switching of CLNS over FDDI does not function correctly. By default, this feature is disabled. Users should not enable it with 8.3(1) software. [CSCdi01839]

AppleTalk

The command **no apple distribute-list** *number* is treated identically to the command **apple distribute-list** *number*. [CSCdi02729]

It is possible for zones to be retained in the AppleTalk Zone Information Table after the networks to which they correspond have been aged out from the routing table. [CSCdi02791]

The **apple proxy-nbp** command is never written to nonvolatile RAM or to configuration files stored on the TFTP servers, which prevents this command from being displayed with a **write term** or **show config** command. Routers must be reconfigured with the **apple proxy-nbp** command at each reboot. Users may be impacted by this problem if there are no other routers in the AppleTalk internet configured with an equivalent command. [CSCdi02792]

SNMP

The **snmp-server host** command issues no error message when it rejects a host on the basis of an invalid format. [CSCdi02757]

Apollo Domain

The command **show access** does not list any Apollo access lists which may have been defined. [CSCdi02864]

HSSI

When switching packets from an HSSI to a slower-speed interface (such as an MCI), the HSSI may lose keepalives, which can cause the interface to go down. This condition occurs when a heavy load is presented to an output interface. There are two work arounds for this problem: either set the HSSI MTU to 1500 bytes or turn off keepalives. [CSCdi02738]

Release 8.3(2) Caveats

This section describes possibly unexpected behavior by release 8.3(2).

AppleTalk

The conversion of special characters to uppercase for use in zone name comparisons is incorrect. This may result in incorrect responses to ZIP queries for zone names containing such characters. The work around is to use only alphanumeric characters in zone names. [CSCdi02637]

It is possible for the 802.2 length field to be set incorrectly on AppleTalk packets fast switched from Ethernet/802.3 media to FDDI media. [CSCdi02653]

If AppleTalk routing is configured for nondiscovery mode on a Token Ring interface connected to a network whose zone-name configuration does not match that of the router, the interface will still be used for AppleTalk routing. The correct behavior would be to shut down AppleTalk routing on the interface in question until the configuration problems had been resolved. [CSCdi03451]

The AppleTalk fast-switching cache is not always invalidated when a new route to a destination is discovered. This may result in failure to deliver traffic to one or more nodes until the cache entry times out. [CSCdi04016]

CLNS

CLNS packets are not fast-switched correctly onto FDDI media. CLNS fast switching should be disabled on all FDDI interfaces. [CSCdi01839]

If the router is assigned a CLNS NET using the **clns net** command, and that NET is then removed using the **no clns net** command, the router will continue to send intermediate system hello messages claiming the removed NET. Note that the **clns net** command is seldom used and is supported primarily for historical reasons. [CSCdi02578]

DECnet

When the router converts DECnet Phase V packets into DECnet Phase IV packets, occasional packets are malformed. [CSCdi03717]

When a DECnet Phase IV packet is converted to a CLNS packet, the size of the clns packet buffer is computed incorrectly, causing overflow when converting large packets. This overflow may result in occasional malformed packets or in system reloads. [CSCdi03963]

DECnet Phase V (CLNS) packets whose destination NSAPs have selector fields which do not correspond to NSP are not converted to Phase IV. [CSCdi04103]

Frame Relay

If a router is receiving routing updates (for any protocol) over a frame relay multicast DLCI, it will learn routes via the frame relay interface, even if the individual data DLCIs associated with remote hosts or routers are defunct. This can result in failure to route around some frame relay failures. [CSCdi02499]

IP Routing Protocols

There is no way to disable the application of the split horizon rule for IP routing. IP routes are not advertised over the interfaces through which they were learned. On a frame relay or SMDS network that is not connected in a full mesh where secondary IP addresses are in use, some routers will never exchange sufficient routing information, resulting in a partitioned network. The work around is to configure frame relay and SMDS networks such that all routers connected to them can communicate directly. This problem will be resolved in Software Release 8.3(3) by the new interface configuration command **no ip split-horizon**. The improved code will disable split horizon by default on frame relay and SMDS interfaces. [CSCdi03430]

If a route learned from EGP in the local autonomous system is redistributed into BGP, and the route is to be sent to another internal BGP peer, that peer will refuse the BGP connection. [CSCdi03853]

When the next hop for a static route which is being redistributed into BGP is changed, the redistributed BGP route does not change. The work around is to remove all knowledge of the network before changing the static route. [CSCdi03863]

When a default route is being learned from RIP, and there is more than one candidate default router with the same metric, the route chosen will oscillate among the candidates. Correct behavior is to choose one default route and use it until there is a real reason to change. [CSCdi04137]

Whenever it appears in the IP routing table, network 0.0.0.0 is a candidate default route. This is not always reflected by the **show ip route** command display. [CSCdi02203]

IBM Connectivity

It is possible for frames being source-route bridged between CSC-R16 interfaces to be reordered over serial links. [CSCdi03110]

Interfaces and Bridging

Under rare conditions, it is possible for a race in the code for the **show ip arp** command to result in system reloads. This command should be used with care. [CSCdi02706]

ARP packets sent on FDDI sometimes use hardware type codes other than the Ethernet code. RFC 1188 calls for ARP on FDDI always to use the Ethernet code. [CSCdi04119]

It is possible for use of the **cbus-buffers** command on busy networks to cause system reloads at the time the command is processed because of a race condition. Failures are extremely rare. [CSCdi04033]

If the system is started with an HSSI interface configured for SMDS, and the encapsulation for that interface is later changed to HDLC, the interface MTU value is not reset, even though the cBus buffers are reapportioned as though the MTU had been reduced to 1500. The work around is to manually reduce the interface MTU to 1500. [CSCdi01406]

When the bandwidth parameter for an interface is changed while that interface is running the Spanning Tree protocol, the interface's path cost is not recalculated to reflect the change, even if the path cost was originally computed from the previous bandwidth setting. This results in the spontaneous appearance of a **path-cost** command in configuration files written after the change, since the path cost no longer reflects the default that would be calculated from the new bandwidth setting. The path cost may be manually set to match the cost that would have been calculated from the new bandwidth. [CSCdi03807]

Basic System Services

Changing IGS serial interface MTU values, or enabling the SMDS encapsulation on IGS serial interfaces, may result in miscalculation of the new buffer quotas. This damage manifests itself as the appearance of incorrect or negative values for buffer quotas in the show buffers command display. This may be worked around by explicitly configuring buffer management parameters using the **buffers** command. [CSCdi04062]

EXEC and Configuration Parser

If the user issues multiple **configure** commands, specifying configuration from the network, only the first dialog will default to the correct TFTP server. Subsequent dialogs will default to broadcast TFTP. [CSCdi04128]

Local Services

SNMP's dnAreaTable does not provide access to the last entry in the system's actual DECNET routing table. [CSCdi03933]

The ifMTU variable reflects the configured IP-specific MTU for the interface. It should reflect the configured overall/physical MTU. [CSCdi04022]

Under rare circumstances, sending of SNMP *tty* enterprise traps may result in router reloads. [CSCdi04138]

TCP/IP Host-mode Services

When TCP receives an acknowledgment for a retransmitted segment, and further segments remain to be retransmitted, the additional segments may not be sent immediately. The condition, which will resolve itself within a few seconds, has minimal operational impact. [CSCdi01517]

X.25

If an X.25 virtual circuit is established over a TCP connection, and the X.25 end nodes in question have different default packet and/or window size parameters, and neither of the end nodes includes facilities specifying these parameters at connection establishment, the end nodes may be given different ideas of the packet and window sizes in effect for the connection, without any translation of packet or window sizes being done by the router. This may be worked around by configuring default packet and window sizes consistently, or by forcing the end nodes to include appropriate facilities in their connection establishment packets. [CSCdi01624]

When X.25 switching is enabled, X.29 calls to subaddresses of the system's main X.25 address will not be accepted and forwarded to rotaries as documented. [CSCdi03285]

The router does not support X.25 clear request packets which have facilities or call user data attached. These packets are neither accepted on connections terminating at the router nor forwarded by the X.25 switching code. [CSCdi04048]

A number of races exist in the X.25 code. These may result in the issuance of spurious traceback messages, or, rarely, in system reloads. Problems will be observed most often on busy X.25 links connected to busy routers. [CSCdi04049]

XNS/Novell/Apollo

Novell echo request packets are sent with an echo reply type code instead of an echo request code. 9.0 Cisco routers will not answer such echo requests. [CSCdi03913]

Release 8.3(3) Caveats

This section describes possibly unexpected behavior by release 8.3(3).

AppleTalk

If the initial address given for an AppleTalk interface does not agree with that interface's cable range, the port may be driven into a continuous reset state. The correct behavior is to reject the attempt to configure an invalid address and issue an error messsage. [CSCdi03924]

Interfaces connected to end-nodes using AppleTalk for VMS, prior to version 3.01, should have the AppleTalk fast-switching cache disabled to insure that all packets will be accepted by those end-nodes. [CSCdi04696]

AppleTalk does not correctly track changes to the encapsulation type set on a serial interface. To work around this problem, clear the AppleTalk configuration on the interface and reconfigure. [CSCdi04609]

The informational level message, *AT-6-ADDRUSED*, will display gibberish numbers for the AppleTalk address in use for the interface in question. [CSCdi04706]

The system will permit configuration of AppleTalk cable ranges on serial interfaces with SMDS and frame relay encapsulations. In fact extended networks are not supported for such interfaces. [CSCdi04771]

CLNS

CLNS packets are not sent on Token Ring media. CLNS is not usable over Token Ring networks. [CSCdi04498]

CLNS packets are not fast switched correctly onto FDDI media. CLNS fast switching should be disabled on all FDDI interfaces. [CSCdi01839]

Under some circumstances, prefix routes will be identified as permanent but unreachable. As a result, if that prefix route is the best match for a particular destination, packets may be dropped. The only way to clear this condition is to reload the router. [CSCdi03030]

The command **clns hold-time** does not work. Although the value is set, it is not used when generating IS hellos. The default is used instead. [CSCdi04388]

IP Routing Protocols

The **offset-list** command does not work properly. Even though a route is included in the access list to have its metric adjusted, no adjustment is made. The default metric is used instead. [CSCdi04312]

When IP traffic is being fast switched on an IGS, and IP accounting is enabled, it is possible for system reloads to occur. The work around is to disable either IP accounting or IP fast switching. [CSCdi04467]

When the router is configured with two interfaces onto a network, and the first interface fails, EGP sessions will still use this address as the source address of their packets. This creates a "black hole" with a loss of connectivity. [CSCdi04549]

After an interface fails, all serial routes are momentarily removed from the IP routing table. Note that this is self-healing, since the routes are then put back in the table. This will cause some routing instability. [CSCdi04579]

EGP per-protocol access lists are broken. For outbound updates, access lists are not applied, thus no filtering is done on these updates. [CSCdi04794]

Whenever it appears in the IP routing table, network 0.0.0.0 is a candidate default route. This is not always reflected by the **show ip route** display. [CSCdi02203]

Interfaces and Bridging

Under rare conditions, it is possible for a race in the code for the **show ip arp** command to result in system reloads. This command should be used with care. [CSCdi02706]

When multiple IP helper addresses are defined, broadcast packets going out the first interface in the list could be sent with bad checksums. [CSCdi04326]

AppleTalk does not work over frame relay in 8.3(2). [CSCdi04547]

The router may deliver RSRB and STUN packets out of order when using raw (or direct) serial encapsulation. Some network applications cannot tolerate receiving packets out of order. [CSCdi04775]

The **no priority-group** command does not accept a number argument. For instance, **no priority-group 10** would incorrectly generate an error. [CSCdi04527]

Attempts to send AppleTalk broadcasts on a frame relay network causes the router to pause indefinitely. This problem occurs on a frame relay network that does not support multicast and has three or more nodes running AppleTalk. [CSCdi04767]

If a router is receiving routing updates (for any protocol) over a frame relay multicast DLCI, it will learn routes via the frame relay interface, even if the individual data DLCIs associated with remote hosts or routers are defunct. This can result in failure to route around some frame relay failures. [CSCdi02499]

If the system is started with an HSSI interface configured for SMDS, and the encapsulation for that interface is later changed to HDLC, the interface MTU value is not reset, even though the cBus buffers are reapportioned as though the MTU had been reduced to 1500. The work around is to manually reduce the interface MTU to 1500. [CSCdi01406]

LAT

LAT break sequences sent by connected hosts are not always honored until the host has sent the next data character. [CSCdi03935]

Basic System Services

Changing IGS serial interface MTU values, or enabling the SMDS encapsulation on IGS serial interfaces, may result in miscalculation of the new buffer quotas. This damage manifests itself as the appearance of incorrect or negative values for buffer quotas in the **show buffers** display. The work around is to explicitly configure buffer management parameters using the **buffers** command. [CSCdi04062]

EXEC and Configuration Parser

If during **setup** user input is delayed, a possible timeout will occur. The router will then loop indefinitely requesting user input; however, no input will be accepted. At this point the router would have to be rebooted to clear the condition. [CSCdi04427]

Local Services

SNMP's dnAreaTable does not provide access to the last entry in the system's actual DECnet routing table. [CSCdi03933]

Under circumstances which are not well understood, badly formed tty traps are output when the SNMP table becomes corrupted. [CSCdi04744]

If extended TACACS is enabled, under certain rare conditions involving retransmissions, corrupted memory could cause the router to reload. [CSCdi04165]

SRT Bridging

SRB proxy explorer does not work. [CSCdi04671]

The **no bridge** *n address* command does not work properly. Although the specified entry is removed, the configuration is modified so that bridge n address commands for stations that were not previously modified are introduced. [CSCdi04700]

TCP/IP Host-Mode Services

When TCP receives an acknowledgment for a retransmitted segment, and further segments remain to be retransmitted, the additional segments may not be sent immediately. The condition, which will resolve itself within a few seconds, has minimal operational impact. [CSCdi01517]

IP accounting reports the length of fast-switched IP packets incorrectly. [CSCdi04472]

If a FIN arrives out of order (for example, because of a lost packet), the connection (now in CLOSEWAIT state) will no longer accept the missing packets in between, leaving the connection permanently paused. [CSCdi04615]

When a router has been up more than approximately 25 days, TCP connections to VTYs may take 4 to 6 minutes to be removed after they have been closed. [CSCdi04738]

VINES

On systems with Token Ring interfaces not configured for multiring, ARP will fail if an ARP request with a RIF is received. [CSCdi04274]

X.25

If X25 encapsulation fails, buffers may be lost. This manifests as a slow loss of memory. [CSCdi04449]

Under some conditions the router may reload when the **show x25 vc** command is entered. [CSCdi04481]

If an X.25 virtual circuit is established over a TCP connection, and the X.25 end nodes in question have different default packet and/or window size parameters, and neither of the end nodes includes facilities specifying these parameters at connection establishment, the end nodes may be given different ideas of the packet and window sizes in effect for the connection, without any translation of packet or window sizes being done by the router. This may be worked around by configuring default packet and window sizes consistently or by forcing the end nodes to include appropriate facilities in their connection establishment packets. [CSCdi01624]

The router does not support X.25 clear request packets which have facilities or call user data attached. These packets are neither accepted on connections terminating at the router nor forwarded by the X.25 switching code. [CSCdi04048]

Under some conditions the router may reload when the **show x25 map** command is typed. [CSCdi04536]

XNS/Novell/Apollo

When a router with Novell IPX routing is being booted over the network, it is possible for received IPX traffic to fill internal buffers without being processed. Buffer starvation may prevent the router from completing its boot process. [CSCdi02722]

Release 8.3(4) Caveats

This section describes possibly unexpected behavior by release 8.3(4).

AppleTalk

In large AppleTalk networks with large Phase 1 components, networks numbers that would normally age out of routing tables may persist indefinitely. This is due in part to the lack of split-horizon processing in Phase 1 environments and changes made to the RTMP aging process in 8.3.

One possible work around is to apply access-lists to block the invalid network numbers from being propagated using the **appletalk distribute-list [in | out]** command.

Upgrading to Phase 2, extended operation on all networks, also corrects the problem. [CSCdi05913]

The conversion of special characters to uppercase for use in zone name comparisons is incorrect. This may result in incorrect responses to ZIP queries for zone names containing such characters. The work around is to use only alphanumeric characters in zone names. [CSCdi02637]

It is possible for the 802.2 length field to be set incorrectly on AppleTalk packets fast switched from Ethernet/802.3 media to FDDI media. [CSCdi02653]

If AppleTalk routing is configured for nondiscovery mode on a Token Ring interface connected to a network whose zone-name configuration does not match that of the router, the interface will still be used for AppleTalk routing. The correct behavior would be to shut down AppleTalk routing on the interface in question until the configuration problems had been resolved. [CSCdi03451]

Interfaces connected to end nodes using AppleTalk for VMS prior to version 3.01 should have the AppleTalk fast switching cache disabled to ensure that all packets will be accepted by those end nodes. [CSCdi04696]

On extended AppleTalk networks with multiple defined zone names, some devices may appear in more than one zone when viewed from a Macintosh that lies across the router; for example, "MktgLaser" appears in "ZoneA" and "ZoneB," although it is defined to reside in "ZoneA" only. This is known to occur with Apple LaserWriter IIg printers and pre-1.5 version Dayna EtherPrints. [CSCdi04951]

The command **appletalk event-logging** is not working. [CSCdi05694]

Bridging

TCP/IP ARP replies are sometimes bridged when both transparent bridging and IP routing are enabled. The conditions under which this occurs are not yet fully understood. [CSCdi05156]

With Novell routing enabled, packets of type AFAF are not bridged. [CSCdi05201]

CLNS

Disabling CLNS routing may cause a system reload to occur. [CSCdi05019]

It is possible for the **show clns routes** command to cause system reloads when issued after the **clear clns route** command. The conditions under which this may occur are not yet fully understood. [CSCdi05143]

CLNS packets are not fast switched correctly onto FDDI media. CLNS fast switching should be disabled on all FDDI interfaces. [CSCdi01839]

Issuing the command **clear clns route** may cause a system reload to occur. [CSCdi05343]

Under some circumstances, the **clns routing** command may show up twice in the routers configuration file. This does not affect the operation of CLNS routing. [CSCdi05196]

NSAP masks are not included in RDPDUs [CSCdi05049]

CLNS does not support both static and dynamic routing simultaneously within a router. [CSCdi05893]

Frame Relay

The frame relay encapsulation code does not correctly check the status of a DLCI. The result is that packets can be sent on a DLCI which the frame relay switch has indicated as deleted via the LMI messages. This problem shows up if a router is misconfigured such that a mismatch exists between the router's DLCI and those defined in the frame relay switch. The work around is to configure the router with the correct DLCIs. [CSCdi05481]

If a router is receiving routing updates (for any protocol) over a frame relay multicast DLCI, it will learn routes via the frame relay interface, even if the individual data DLCIs associated with remote hosts or routers are defunct. This can result in failure to route around some frame relay failures. [CSCdi02499]

IP Routing Protocols

When an IP IGRP update is created for a major network that is subnetted and directly connected to the router, but the update in question is being sent through an interface that does not lie in the network in question, the metric chosen for the major net may not be the best of the metrics to any of its subnets. Furthermore, if the connection to the major net in question is through a secondary address, the network will not be included in the IGRP update at all. [CSCdi02859]

When a router is configured with two interfaces onto an IP network, if the first interface fails, EGP sessions will still use this address as the source address of their packets. This creates a "black hole" with a loss of connectivity. [CSCdi04549]

The **distribute-list out** command is not honored for BGP routing processes. [CSCdi04825]

BGP will accept a NEXT_HOP path attribute that is the router's own address. [CSCdi04961]

Whenever it appears in the IP routing table, network 0.0.0.0 is a candidate default route. This is not always reflected by the **show ip route** command display. [CSCdi02203]

ARP requests generated on FDDI by systems that are bridging IP are sent using the common FDDI SNAP encapsulation. Other systems on the FDDI ring will not bridge these packets onto Ethernets that may be connected to them, and ARP table entries will therefore never be learned for systems on those Ethernets. The correct behavior is to use the Ethernet-over-FDDI encapsulated bridging format for ARP packets generated on FDDI by units bridging IP. [CSCdi05482]

When using the **domain-list** command, the software may fail to properly update domain cache entries that have been timed out. [CSCdi03896]

If RIP is run across an unnumbered link, and the associated numbered interface has a nondefault broadcast address, the RIP updates on the unnumbered links will have an incorrect checksum generated. The work around is to use the default broadcast address on the associated numbered interface. [CSCdi04838]

In the presence of CSCdi05031, large amounts of processing time may be consumed by the TCP driver. The high TCP usage may be experienced upon closure of an X.25/TCP tunnel between two routers, one of which is experiencing CSCdi05031. [CSCdi05515]

Interfaces and Bridging

Under rare conditions, it is possible for a race in the code for the **show ip arp** command to result in system reloads. This command should be used with care. [CSCdi02706]

If IP autonomous switching is enabled on one FDDI interface and not enabled on another, it is possible for incorrect IP ARP table entries to be created for nodes directly connected on the FDDI rings. These entries can be recognized in the **show ip arp** display by their encapsulation fields, which are shown as ARPA. The ARPA encapsulation is not supported for FDDI, except when IP is being encapsulation-bridged. It is not possible to communicate with or through the nodes with incorrect ARP table entries. The condition can be avoided by configuring autonomous switching consistently on all FDDI interfaces in the system. [CSCdi05835]

The system may periodically display bad share count messages when running SRB and RSRB. [CSCdi04971]

On a CSC/4 processor with an Ethernet MCI, keepalives will not bring back an Ethernet interface that is down (transceiver cable disconnect, cable unterminated, and so on).

For an Ethernet with keepalives enabled, a keepalive packet is sent at every keepalive interval. In this scenario, if a user were to disconnect the transceiver cable to the Ethernet and three keepalives were sent but not received, "line protocol" would go down, and the interface would be unusable, as expected. If the user was to then reconnect the transceiver cable, the correct behavior would be for the keepalives to bring the interface back up within the keepalive period. This does not happen with the CSC/4 processor. The interface will remain down despite attempts to lengthen the keepalive period, generate more keepalives, or clear the Ethernet interface with the **clear interface** command.

The work around is to toggle the keepalives for that particular ethernet interface using the **no keepalive** command followed by the **keepalive** *n* command. [CSCdi05172]

Note: The only action that is *required* for the interface to come back up is to turn off keepalives. Turning them back on is optional, but doing so will correctly turn off "line protocol" if the line goes down in the future.

Under certain conditions on the Token Ring interface (generally high traffic or noisy media), a message similar to the following may appear, indicating that the Token Ring interface was unable to reset itself. [CSCdi05644]:

```
%TR-3-RESETFAIL: Unit 0, reset failed, error code 00007F32.
-Traceback= 97F84 97CFA 970A2 96FBE 9C5E8 12766 37F8 1D1E
```

If the system is started with an HSSI interface configured for SMDS, and the encapsulation for that interface is later changed to HDLC, the interface MTU value is not reset, even though the cBus buffers are reapportioned as though the MTU had been reduced to 1500. The work around is to manually reduce the interface MTU to 1500. [CSCdi01406]

The router will reload if the interface subcommand **bandwidth** is set to zero. [CSCdi05964]

Basic System Services

It is possible for system reloads to occur when the nonvolatile configuration memory is manipulated from more than one terminal session. Only one terminal at a time should do commands from the set **show config, write memory, write** with no argument, **write erase**, or **config** from memory]. [CSCdi03856]

Changing IGS serial interface MTU values or enabling the SMDS encapsulation on IGS serial interfaces, may result in miscalculation of the new buffer quotas. This damage manifests itself as the appearance of incorrect or negative values for buffer quotas in the **show buffers** command display. This can be worked around by explicitly configuring buffer management parameters using the **buffers big permanent 5** command. [CSCdi04062]

The router does not change the source address it uses for syslog messages after the address is no longer valid. The correct behavior is for a new address to be selected. A work around is to reload the router after a reconfiguration that has invalidated the address the router was using to source syslog messages. [CSCdi04906]

The **stopbits 1.5** command is never written to nonvolatile RAM or to remote network configuration files, even for lines that have been configured using it. [CSCdi05124]

Local Services

SNMP's dnAreaTable does not provide access to the last entry in the system's actual DECNET routing table. [CSCdi03933]

SRT Bridging

The spanning-tree path cost is not set correctly for interfaces using frame relay encapsulation. The work around is to manually specify the path cost with the command **bridge-group path-cost**. [CSCdi05593]

TCP/IP Host-mode Services

The **service tcp-keepalive** command only applies to terminal ports and VTYs. [CSCdi05905]

UDP echo requests are only responded to correctly for the first request received. Subsequent responses will be sent to the initial requesting address regardless of who issues the request. The correct behavior is for the response to be sent to the address making the request. [CSCdi05721]

UDP port filtering is only done on packets arriving with a media broadcast indication. Consequently, the udp port filtering mechanism **ip forward protocol udp** is ignored when receiving packets from nonbroadcast media such as X.25 and some frame relay networks. [CSCdi06001]

VINES

Server discovery broadcasts received on interfaces configured with **vines serverless** are always forwarded to the nearest server listed in the routing table. The nearness of the server in question is calculated from the router's point of view, rather than from the point of view of the client. This behavior may cause overloading of the nearest server while other servers are left underutilized. [CSCdi02868]

The router always chooses the last entry in the neighbor's table when responding to a client request. The correct behavior is to respond with the first entry in the table. [CSCdi05000]

X.25

If an X.25 virtual circuit is established over a TCP connection, and the X.25 end nodes in question have different default packet and/or window size parameters, and neither of the end nodes includes facilities specifying these parameters at connection establishment, the end nodes may be given different ideas of the packet and window sizes in effect for the connection, without any translation of packet or window sizes being done by the router. This may be worked around by configuring default packet and window sizes consistently, or by forcing the end nodes to include appropriate facilities in their connection-establishment packets. [CSCdi01624]

The router does not support X.25 clear request packets that have facilities or call user data attached. These packets are neither accepted on connections terminating at the router nor forwarded by the X.25 switching code. [CSCdi04048]

Release 8.3(5) Caveats

This section describes possibly unexpected behavior by release 8.3(5). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(5).

AppleTalk

Entering the command **appletalk event-logging** returns a spurious message:

```
% One of "probe" or "request"
```

This message can be ignored. [CSCdi05694]

In large AppleTalk networks with large Phase 1 components, networks numbers that would normally age out of routing tables may persist indefinitely. This is due in part to the lack of split-horizon processing in Phase 1 environments and changes made to the RTMP aging process in 8.3.

One possible work around is to apply access lists to block the invalid network numbers from being propagated using the **appletalk distribute-list [in | out]** command. Upgrading to Phase 2 extended operation on all networks also corrects the problem. [CSCdi05913]

The conversion of special characters to uppercase for use in zone name comparisons is incorrect. This may result in incorrect responses to ZIP queries for zone names containing such characters. The work around is to use only alphanumeric characters in zone names. [CSCdi02637]

It is possible for the 802.2 length field to be set incorrectly on AppleTalk packets that are fast-switched from Ethernet/802.3 media to FDDI media. [CSCdi02653]

If AppleTalk routing is configured for nondiscovery mode on a Token Ring interface connected to a network whose zone name configuration does not match that of the router, the interface will still be used for AppleTalk routing. The correct behavior would be to shut down AppleTalk routing on the interface in question until the configuration problems have been resolved. [CSCdi03451]

Interfaces connected to end nodes using AppleTalk for VMS prior to version 3.01 should have the AppleTalk fast switching cache disabled to insure that all packets will be accepted by those end-nodes. [CSCdi04696]

On extended AppleTalk networks with multiple defined zone names, some devices may appear in more than one zone when viewed from a Macintosh that lies across the router; for example, "MktgLaser" appears in "ZoneA" and "ZoneB," although it is defined to reside in "ZoneA" only. This is known to occur with Apple LaserWriter IIg printers and pre-1.5 version Dayna EtherPrints. [CSCdi04951]

Cisco's AppleTalk implementation over HDLC-encapsulated serial lines padded all packets to even byte boundaries for performance reasons in MCI microcode releases before release 1.9. However, when routing a packet that arrived in the router from a HDLC-encapsulated serial line and that was routed out an extended 802 interface (for example, 802.3 Ethernet, Token Ring, or FDDI), the extra byte appended to the end of the HDLC frame would be included in the packet sent out the 802 interface. Some implementations of AppleTalk (particularly AppleTalk for VMS 3.01) would then reject a packet which had traversed a HDLC-serial-to-extended 802 path, claiming that the packet did not conform to the LLC frame specification. Previously, this problem could be worked around by disabling fast switching on the Ethernet , Token Ring, or FDDI interface connected to the VAX running AppleTalk for VMS. This fix, in conjunction with MCI microcode release 1.9, allows fast switching to be enabled on interfaces connected to VAXes running AppleTalk for VMS in networks where AppleTalk traffic transits HDLC-serial links before arriving at the VAX. [CSCdi05439]

Contrary to the documentation for 8.3 and 9.0, the following commands have not been implemented [CSCdi05597]:

clear appletalk routes
clear appletalk zones
clear appletalk neighbors

In software releases 8.3(3) and 9.0(1), a nonextended interface can become operational in spite of the fact that an adjacent and active neighbor has a different configuration. Although the interface becomes operational, connectivity through any routes controlled by that neighbor is lost. [CSCdi05642]

After the router has been operational for approximately two weeks, elapsed time fields in AppleTalk **show** EXEC commands will read as never instead of an elapsed time. This is due to wrapping of the router's internal timekeeping calculations. The indication of a never-elapsed time has no effect on functioning of the router. Approximately two weeks after the elapsed-time fields read as *never*, they will once again indicate correct elapsed time values. The indication of elapsed time in routers which have been operational for more than two weeks will be addressed in a future release of the software. [CSCdi06268]

Bridging

TCP/IP ARP replies are sometimes bridged when both transparent bridging and IP routing are enabled. The conditions under which this occurs are not yet fully understood. [CSCdi05156]

With Novell routing enabled, packets of type AFAF are not bridged. [CSCdi05201]

ISO CLNS

Disabling CLNS routing can cause a system reload to occur. [CSCdi05019]

CLNS packets are not fast switched correctly onto FDDI media. CLNS fast switching should be disabled on all FDDI interfaces. [CSCdi01839]

Issuing the command **clear clns route** can cause a system reload to occur. [CSCdi05343]

Forwarding a converted DECnet Phase IV packet causes a DECnet PhaseV redirect. For example, a CLNS packet is received on an interface. It is converted into a DECnet Phase IV packet, which is then sent back out the interface, and an ES-IS redirect PDU is erroneously sent. [CSCdi06121]

Under some circumstances, the **clns routing** command may show up twice in the routers configuration file. This does not affect the operation of CLNS routing. [CSCdi05196]

NSAP masks are not included in RD PDUs [CSCdi05049]

CLNS does not support both static and dynamic routing simultaneously within a router. [CSCdi05893]

Frame Relay

The frame relay encapsulation code does not correctly check the status of a DLCI. The result is that packets can be sent on a DLCI which the frame relay switch has indicated as deleted in the LMI messages. This problem shows up when a mismatch exists between the router's DLCI and those defined in the frame relay switch. The work around is to configure the router with the correct DLCIs. [CSCdi05481]

If a router is receiving routing updates (for any protocol) over a frame relay multicast DLCI, it will learn routes via the frame relay interface, even if the individual data DLCIs associated with remote hosts or routers are defunct. This can result in failure to route around some frame relay failures. [CSCdi02499]

IP and IP Routing Protocols

When using the **domain-list** command, the software may fail to properly update domain cache entries that have been timed out. [CSCdi03896]

When an IP IGRP update is created for a major network that is subnetted and directly connected to the router, but the update in question is being sent through an interface that does not lie in the network in question, the metric chosen for the major network may not be the best one for any of its subnets. Furthermore, if the connection to the major network in question is through a secondary address, the network will not be included in the IGRP update at all. [CSCdi02859]

If, when a router is configured with two interfaces onto an IP network, the first interface fails, EGP sessions will still use this address as the source address for their packets. This creates a "black hole" with a loss of connectivity. [CSCdi04549]

The **distribute-list out** command is not honored for BGP routing processes. [CSCdi04825]

BGP will accept a NEXT_HOP path attribute that is the router's own address. [CSCdi04961]

If routers using secondary addresses are inconsistent about the primary address, routing updates are not generated correctly. [CSCdi05942]

Whenever it appears in the IP routing table, network 0.0.0.0 is a candidate default route. This is not always reflected by the **show ip route** command display. [CSCdi02203]

If RIP is run across an unnumbered link, and the associated numbered interface has a nondefault broadcast address, then the RIP updates on the unnumbered links will have an incorrect checksum generated. The work around is to use the default broadcast address on the associated numbered interface. [CSCdi04838]

Interfaces and Bridging

Under rare conditions, it is possible for a race in the code for the **show ip arp** command to result in system reloads. This command should be used with care. [CSCdi02706]

The system may periodically display bad share count messages when running SRB and RSRB. [CSCdi04971]

Keepalives will not bring back an Ethernet interface that is down (transceiver cable disconnected, cable unterminated, and so on) on a CSC/4 processor with an Ethernet MCI. For an Ethernet with keepalives enabled, a keepalive packet is sent every keepalive interval. In this scenario, if a user disconnects the transceiver cable to the Ethernet, and three keepalives were sent but not received, "line protocol" would go down, and the interface would be unusable, as expected. If the user then reconnects the transceiver cable, the correct behavior would be for the keepalives to bring the interface back up within the keepalive period. This does not happen with the CSC/4 processor. The interface remains down despite attempts to lengthen the keepalive period, generate more keepalives, or attempt to clear the ethernet interface with the **clear interface** command. The work-around is to toggle the keepalives for that particular ethernet interface using the **no keepalive** command followed by the **keepalive** *n* command. [CSCdi05172]

Note: The only action that is *required* for the interface to come back up is to turn off keepalives. Turning them back on is optional, but doing this will correctly turn off "line protocol" if the line goes down in the future.

ARP requests generated on FDDI by systems which are bridging IP are sent using the common FDDI SNAP encapsulation. Other systems on the FDDI ring will not bridge these packets onto Ethernets which may be connected to them, and ARP table entries will therefore never be learned for systems on those Ethernets. The correct behavior is to use the Ethernet-over-FDDI encapsulated bridging format for ARP packets generated on FDDI by units bridging IP. [CSCdi05482]

Under conditions such as high traffic or noisy media on the Token Ring interface, a message similar to the following may appear, indicating that the Token Ring interface was unable to reset itself: [CSCdi05644]

```
%TR-3-RESETFAIL: Unit 0, reset failed, error code 00007F32.
-Traceback= 97F84 97CFA 970A2 96FBE 9C5E8 12766 37F8 1D1E
```

If the system is started with an HSSI interface configured for SMDS, and the encapsulation for that interface is later changed to HDLC, the interface MTU value is not reset, even though the cBus buffers are reapportioned as though the MTU had been reduced to 1500. The work around is to manually reduce the interface MTU to 1500. [CSCdi01406]

The router will reload if the interface subcommand bandwidth is set to zero. [CSCdi05964]

Basic System Services

It is possible for system reloads to occur when the nonvolatile configuration memory is manipulated from more than one terminal session. Only one terminal at a time should do commands from the set **show config, write memory, write** with no argument, **write erase**, or **config** from memory]. [CSCdi03856]

Changing IGS serial interface MTU values or enabling the SMDS encapsulation on IGS serial interfaces, may result in miscalculation of the new buffer quotas. This damage manifests itself as the appearance of incorrect or negative values for buffer quotas in the **show buffers** command display. This may be worked around by explicitly configuring buffer management parameters using the **buffers** command. [CSCdi04062]

The router does not change the source address it uses for syslog messages after the address is no longer valid. The correct behavior is for a new address to be selected. A work around is to reload the router after a reconfiguration that has invalidated the address the router was using to source syslog messages. [CSCdi04906]

The **stopbits 1.5** command is never written to nonvolatile RAM or to remote network configuration files, even for lines which have been configured using it. [CSCdi05124]

Local Services

SNMP's dnAreaTable does not provide access to the last entry in the system's actual DECnet routing table. [CSCdi03933]

SRT Bridging

The spanning-tree path cost is not set correctly for interfaces using frame relay encapsulation. The work around is to manually specify the path cost with the command **bridge-group path-cost**. [CSCdi05593]

TCP/IP Host-Mode Services

The **service tcp-keepalive** command only applies to terminal ports and VTYs. [CSCdi05905]

UDP echo requests are only responded to correctly for the first request received. Subsequent responses will be sent to the initial requesting address regardless of who issues the request. The correct behavior is for the response to be sent to the address making the request. [CSCdi05721]

When a TCP connection has a closed window, packets containing valid acknowledgment (ACK) packets are discarded if they also contain any data since the data is outside of the window. The correct behavior is to continue to process the ACKs for segments with reasonable ACK values. This is a problem in the initial stages of a connection, when the server sends the synchronous-acknowledgment (SYN-ACK) packet with a 0 window. If the ACK to the server's SYN also contains data, the server will not process that ACK, and the connection never gets to ESTABLISHED state. [CSCdi05962]

UDP port filtering is only done on packets arriving with a media broadcast indication. Consequently, the UDP port-filtering mechanism command **ip forward-protocol udp** is ignored when receiving packets from nonbroadcast media such as X.25 and some frame relay networks. [CSCdi06001]

In the presence of CSCdi05031, large amounts of processing time can be consumed by the TCP driver. The high TCP usage may be experienced upon closure of an X.25/TCP tunnel between two routers, one of which is experiencing CSCdi05031.

VINES

Server discovery broadcasts received on interfaces configured with the **vines serverless** command are always forwarded to the nearest server listed in the routing table. The nearness of the server in question is calculated from the router's point of view, rather than from the point of view of the client. This behavior may cause overloading of the nearest server, while other servers are left underutilized. [CSCdi02868]

The router always chooses the last entry in the neighbor's table when responding to a client request. The correct behavior is to respond with the first entry in the table. [CSCdi05000]

X.25

An interface input queue may fill up and not recover if an X.25 provider violates the LAPB protocol by exiting from the RNR state with an RR frame instead of a REJ frame. This can cause the serial interface to pause indefinitely and cease transmission. [CSCdi05957]

If an X.25 virtual circuit is established over a TCP connection, and the X.25 end nodes in question have different default packet and/or window size parameters, and neither of the end nodes includes facilities specifying these parameters at connection establishment, the end nodes can be given different ideas of the packet and window sizes in effect for the connection, without any translation of packet or window sizes being done by the router. This can be worked around by configuring default packet and window sizes consistently, or by forcing the end nodes to include appropriate facilities in their connection establishment packets. [CSCdi01624]

The router does not support X.25 clear request packets which have facilities or call user data attached. These packets are neither accepted on connections terminating at the router nor forwarded by the X.25 switching code. [CSCdi04048]

8.3(4) Caveats/8.3(5) Modifications

This section describes possibly unexpected behavior by release 8.3(4). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(4). For additional caveats applicable to release 8.3(4), see the caveats sections for newer 8.3 releases. The caveats for newer releases precede this section. All the caveats listed in this section are resolved in release 8.3(5).

AppleTalk

AARP packets from nodes in the startup range are rejected as "martians" preventing nodes from acquiring their initial configuration when connected to a new network. The work around is to have at least one router on the cable that is not running version 8.3(4). [CSCdi06137]

IP Routing Protocols

After a system has been operational for 24 days, the IGRP, RIP, HELLO and CHAOS routing processes stop sending updates. The cessation occurs when the routing process has been running the entire time the system has been operational or when the process is manually started any time after system start up. There is a work around for IGRP. Assuming nondefault values for the IGRP timers, use the following router subcommand:

timers basic 90 270 280 630 1

The only value that helps the work around case is setting the fifth parameter equal to one. The other values do not affect the problem and should be set according to the users' needs. The above example is the normal case. A work around does not exist for RIP, HELLO and CHAOS. [CSCdi06310]

All the caveats listed in this section are resolved in release 8.3(4).

AppleTalk

If the initial address given for an AppleTalk interface does not agree with that interface's cable range, the port may be driven into a continuous reset state. The correct behavior is to reject the attempt to configure an invalid address and issue an error messsage. [CSCdi03924]

In certain unusual circumstances, the router can fail to acquire zone information from neighbors for valid routes. This results in partial loss of connectivity. Turning off AppleTalk, and/or restarting the router may act as a work around. [CSCdi04999]

AppleTalk does not correctly track changes to the encapsulation type set on an serial interface. To work around this problem, clear the AppleTalk configuration on the interface and reconfigure. [CSCdi04609]

The informational level message, AT-6-ADDRUSED, will display gibberish numbers for the AppleTalk address in use for the interface in question. [CSCdi04706]

The system will permit configuration of AppleTalk cable ranges on serial interfaces with SMDS and frame relay encapsulations. In fact, extended networks are not supported for such interfaces. [CSCdi04771]

When an AppleTalk ARP reply is received on a Token Ring interface, the sanity check that prevents entering multicast MAC addresses into the ARP table is done incorrectly; the least-significant bit of the first octet of the address is checked instead of the most-significant. This may result in the system accepting invalid AppleTalk ARP replies, or, usually more seriously, in its ignoring valid ones. This can be worked around by reconfiguring other nodes to use Token Ring MAC addresses that do not have the least significant bits set in their first octets. [CSCdi05167]

This bug would affect the ability of a nonextended AppleTalk interface in discovery mode to start when there is only a Shiva FastPath on the cable to perform the function of seed router. If there is already some router other than a Shiva FastPath on the cable, the interface will start routing as expected. [CSCdi05440]

A system reload of a router may occur under very rare circumstances while performing a **show apple arp** command as a result of an ARP table entry being removed while the **show apple arp** command was traversing the ARP table. [CSCdi05232]

Bridging

If a bridge group containing three or more interfaces is established, and if any of the interfaces in that bridge group is an X.25 or frame relay serial link, random data may be sent in place of the correct data for bridged frames being flooded over that link. This manifests itself both in incorrect delivery of traffic and in the appearance of incorrect MAC addresses in the bridging database of the bridge(s) at the other end of the X.25 or frame relay link. [CSCdi05027]

Broadcast 802.2/802.3 packets with DSAP/SSAP pairs of FE/FE (usually CLNS packets) are not bridged. This behavior is present in release 8.3(3), but not in release 8.3(1). [CSCdi05009]

CLNS

CLNS packets are not sent on Token Ring media. CLNS is not usable over Token Ring networks. [CSCdi04498]

Routers performing DECnet Phase V/CLNS to DECnet Phase IV conversion may rapidly run out of system memory. [CSCdi05021]

The command **clns hold-time** does not work. Although the value is set, it is not used when generating IS hellos. The default is used instead. [CSCdi04388]

Frame Relay

When a frame relay interface transitions from up to down and vice versa, the system variables are updated but no SNMP trap is generated. This is incorrect behavior. The correct behavior is to generate the SNMP trap. [CSCdi05198]

There are instances where the frame relay initialization does not clear the loopback flag. An interface will incorrectly report that it is in loopback if the interface is in loopback mode with HDLC encapsulation, then reconfigured for frame relay encapsulation without shutting down the interface. The work around is to administratively shut down the interface and then reinitialize it. [CSCdi05483]

AppleTalk does not work over frame relay in 8.3(2) and 8.2(3). [CSCdi04547]

If the **frame relay map** command is issued before the **encapsulation frame relay** command, then no action is taken. This is the correct behavior. So although no action is taken no error message is generated. Not generating an error message in this case was incorrect; an error message is now generated. [CSCdi04576]

IP Routing Protocols

IP fast switching continues to use a default route for a network even after receiving a valid route for that network. [CSCdi04804]

If a network broadcast address and a default subnet are configured, the Cisco router will erroneously route a network broadcast to the default subnet. This can lead to routing table instabilities. A work around is to specify the broadcast address of 255.255.255.255. [CSCdi05052]

When IP traffic is being fast switched on an IGS, and IP accounting is enabled, it is possible for system reloads to occur. This can be worked around by disabling either IP accounting or IP fast switching. [CSCdi04467]

After an interface fails, all serial routes are momentarily removed from the IP routing table. Note that this is self-healing because the routes are then put back in the table. This will cause some routing instability. [CSCdi04579]

EGP per-protocol access lists are broken. For outbound updates, access lists are not applied; thus no filtering is done on these updates. [CSCdi04794]

Attempts to create IP static interface routes through interfaces that do not have IP addresses assigned will fail. [CSCdi04898]

If IP accounting is disabled, or if the IP accounting database is cleared or checkpointed while a **show ip accounting [checkpoint]** command is being issued, a system reload may occur. [CSCdi05159]

The way EGP-handled routes are aged out is incorrect in the case where the router drops the route, and the neighbor stays up. The incorrect behavior is to use a multiple of invalid time. The correct behavior is to subtract invalid time from flush time and use that value as a multiple to age the routes. [CSCdi05170]

An IP accounting filter disables fast switching for packets that do not match the filter. [CSCdi05299]

If the command **no ip split-horizon** is enabled on an interface with secondary addresses, RIP updates are only issued for those secondary addresses on a different major network number from the primary. The correct behavior is for a RIP update to be sent out for each secondary address. [CSCdi05448]

ICMP Information requests do not cause entries to be made in the ARP table. Instead, an ARP request is broadcast before sending the ICMP reply. This can cause problems with devices that need to learn the subnet portion of their IP addresses from the ICMP Reply. [CSCdi04328]

If an IP address is removed from an interface using the **no ip address**, all routes using that interface are deleted from the IP routing table. This is sometimes unnecessary when there is an additional path to the target. [CSCdi04396]

If the next hop router specified for a static route goes down, ISO-IGRP incorrectly sends out a flash update with a noninfinity metric for that static route. [CSCdi04927]

ISO-IGRP flash-update storms occur when there are parallel adjacencies on interfaces with different ISO-IGRP metrics. The storm occurs for prefix routes only. A work around is to make the metrics the same on the interfaces. This is accomplished by setting the bandwidth and the delay to be the same on each interface involved. [CSCdi05235]

HP Probe

Under some circumstances, primarily involving a nonzero hold queue on an ethernet interface, the use of the HP Probe feature may cause the router to lose memory. [CSCdi05186]

Older HP probe clients (notably old versions of OfficeShare) require support for the "where is gateway" packet. This feature is not supported. [CSCdi04667]

Interfaces and Bridging

When multiple IP helper addresses are defined, broadcast packets going out the first interface in the list could be sent with bad checksums. [CSCdi04326]

The router may deliver RSRB and STUN packets out of order when using raw (or direct) serial encapsulation. Some network applications cannot tolerate receiving packets out of order. [CSCdi04775]

An Ultranet interface configured for bridging accepts its own broadcasts. This can cause the bridging table to become corrupted. [CSCdi04954]

The **no priority-group** command does not accept a number argument. For instance, the command **no priority-group 10** would incorrectly generate an error. [CSCdi04527]

Attempts to send AppleTalk broadcasts on an frame relay network causes the router to pause indefinitely. This problem occurs on a frame relay network that does not support multicast and has three or more nodes running AppleTalk. [CSCdi04767]

Packets received over the Ultranet interface that are within seven bytes of maximum size will be incorrectly counted as giants. [CSCdi04817]

No ARP cache entry is made for the system's own IP address on an Ultranet interface. This results in the system being unable to "talk to itself" using IP over that interface. [CSCdi04828]

When an IP packet with options and a time-to-live field of one is received on a fast-switching interface, the packet is erroneously treated as having an IP header checksum error. This is most noticeable when a **traceroute** program is being used with source-routing options. [CSCdi04830]

The router allows Bridging Circuit Groups to be configured on interfaces supporting frame relay and X.25. This functionality is not supported for frame relay and X.25. The correct behavior is for the router to not allow Bridging Circuit Groups to be configured on interfaces supporting frame relay and X.25. [CSCdi04998]

When issuing the command **show interface token 0**, the bia is displayed as *0000.0000.0000*. The correct behavior is for the actual burned-in address of the board to be displayed. [CSCdi05404]

If an interface enabled for multiring is reset, either by user action or by keepalives, the router may issue "Bad enqueue" messages. The format of the message follows [CSCdi05570]:

```
%SYS-2-LINKED: Bad enqueue of 26BFE8 in queue 1E5450 -Process= "Net
Background", ipl= 4, pid= 9 -Traceback= 7442 323F8 2EFF2 13ABA 10FF6 2434
```

Very high average output rates can result in overflows in the computation of the five-minute data rates in the **show interface** command display. This manifests itself as the appearance of nonsensically large values. [CSCdi04665]

Initiating a LAT translation session with transparent bridging enabled causes a system reload to occur. [CSCdi05229]

Basic System Services

The **show conf** command displays the following buffer numbers :

buffers small min-free 20

buffers middle min-free 10

buffers big min-free 5

Extra lines of default buffers clutter the NVRAM listing. If the **write memory** command is executed, it will save this configuration to the NVRAM. This will cause the lines to stay permanently in the configuration, even in future releases. The no variations of the commands must be entered in order for each line to clear the extra messages. [CSCdi04904]

CLNS hosts do not increment the line count correctly in the **show host** command display. Consequently, the command does not respect the settings in the **term length** *n* command. [CSCdi05083]

EXEC and Configuration Parser

If during setup user input is delayed, a possible timeout will occur. The router will then loop indefinitely requesting user input. However no input will be accepted. At this point, the router would have to be reloaded to clear the condition. [CSCdi04427]

When setup is used to configure a router, the **router igrp** command is removed from the configuration file on reload. The work around is to modify the configuration file by hand and add back the missing command. [CSCdi04641]

The command **service exec-wait**, which causes the EXEC process to wait if there is input pending on a modem line, has been implemented. This command is intended as a work around for problems with modems sending junk characters during various types of speed negotiation. The command is disabled by default. [CSCdi04852]

Local Services

The **tacacs last-resort succeed** command does not work on lines configured for dynamic assignment of SLIP addresses. [CSCdi02330]

Under circumstances that are not well understood, badly formed tty traps are output when the SNMP table becomes corrupted. [CSCdi04744]

Setting the SNMP tsMsgIntervaltim variable to zero prevents any issuance of the message. The correct behavior is for the message to be issued at intervals decided by the system itself. [CSCdi04860]

Any authenticated extended TACACS request will change the user's access class. If the field is set in the packet, the TACACS server supplied leaves it set to zero for everything except the login and SLIP address. This should only happen for responses to login requests. [CSCdi05175]

SMDS

AppleTalk Phase I fails to route over serial links configured for SMDS encapsulation. [CSCdi04914]

The OUI fields of outgoing SMDS packets may contain random data. This can interfere with communication to nodes that do very strict packet checking. The correct behavior is to zero these fields. [CSCdi05119]

SRT Bridging

SRB proxy explorer does not work. [CSCdi04671]

The **no bridge** *n address* command does not work properly. Although the specified entry is removed, the configuration is modified so that **bridge** *n address* commands for stations that were not previously modified are introduced. [CSCdi04700]

Path costs for Spanning Tree Protocol not recomputed when enabling DEC spanning tree protocol. A potential side effect of this is that interfaces configured for bridging after the **bridge** *n* **proto dec** command has been issued may have different path costs than those configured before the command. [CSCdi05251]

TCP/IP Host-Mode Services

IP accounting reports the length of fast-switched IP packets incorrectly. [CSCdi04472]

If a FIN arrives out of order (for example, because of a lost packet), the connection (now in the CLOSEWAIT state) will no longer accept the missing packets in between, leaving the connection permanently paused. [CSCdi04615]

When a router has been up more than approximately 25 days, TCP connections to VTYs may take 4 to 6 minutes to be removed after they have been closed. [CSCdi04738]

Under some obscure conditions (TCP connection receives a RST packet while the connection is closing, and you are waiting for data to go to the terminal), TCP does not release all buffers. Eventually this causes the interface input queue to fill up. The router must be reloaded in order to clear up this condition. This problem is not so serious because it occurs infrequently. [CSCdi04957]

The success rate for the **ping** command may incorrectly report a low success if ping is run for a very long time. The counter containing the successful ping count overflows. [CSCdi05163]

VINES

On systems with Token Ring interfaces not configured for multiring, ARP will fail if an ARP request with a RIF is received. [CSCdi04274]

X.25

With X25 TCP enabled, if data continues to be sent to a TCP connection in the CLOSEWAIT state after the X25 connection has been removed, the router may reload. [CSCdi05031]

If X25 encapsulation fails, buffers may be lost. This manifests as a slow loss of memory. [CSCdi04449]

Under some conditions, the router may reload when the **show x25 vc** command is typed. [CSCdi04481]

If more than 22 parameter/value pairs are entered in an **x29 profile** command, memory will become corrupted, leading to a possible system failure. [CSCdi05307]

A number of races exist in the X.25 code. These may result in the issuance of spurious trace back messages, or, rarely, in system reloads. Problems will be observed most often on busy X.25 links connected to busy routers. [CSCdi04049]

Under some conditions the router may reload when the **show x25 map** command is entered. [CSCdi04536]

The X.25 switch code does not properly handle forwarding of a RESET packet, causing it to be returned on the line instead of forwarded over the TCP connection. [CSCdi04663]

When an X.25 PAD connection receives an INDICATION OF BREAK packet, that indication is not forwarded into the data stream of any possible outgoing connection. [CSCdi04908]

X.25 virtual circuits over which no data have ever been sent are not closed when the configured idle time has passed. If any traffic whatsoever is sent over a virtual circuit, the idle timer will be applied thereafter. [CSCdi05123]

The **no x25 facility throughput** command does not work. There is no way to remove this facility. [CSCdi05217]

Additional calls cannot be made if all available VCs are open, and the first VC is busy, even if the remaining VCs are idle. The correct behavior is to check all VCs and not just the first one on the list. [CSCdi05374]

XNS/Novell/Apollo

If a Novell packet is corrupted such that the checksum field is not *0xFFFF*, it is possible for the router to reload. This occurs infrequently as packets corrupted in this manner are fairly rare. [CSCdi04921]

When a router with Novell IPX routing is being booted over the network, it is possible for received IPX traffic to fill internal buffers without being processed. Buffer starvation may prevent the router from completing its boot process. [CSCdi02722]

XNS routes that have been filtered out by **xns output-network-filter** command are still being advertised with a hop count of 16 (inaccessible). The correct behavior is for these networks not to be included in the routing update. [CSCdi03844]

When an interface is shut down, only the connected route to that network is removed from the routing table. All other Novell routes that were learned via that interface remain until they are timed out. [CSCdi05087]

When an interface is shut down, the Novell static routes associated with that interface will age out of the routing table. The correct behavior is for static routes not to age out. [CSCdi05090]

When Novell routing is disabled on an interface, the Novell routes learned via that interface are not deleted from the table. These routes must time out for three minutes. The correct behavior is for the routes to be flushed from the table when Novell routing is disabled. [CSCdi05144]

For the Novell protocol, the router is too restrictive when deciding which packets to forward in a mixed-media environment. If a packet is sourced from a station on a Token Ring with the address *0100.xxxx.xxxx*, that the packet will not make it past the second router in the path to the destination. The reason is that while *0100* is not multicast on TR, when the packet then is sent on an Ethernet to another router, it becomes sourced from a multicast address and is thrown away. The same would hold true for a source address of *8000.xxxx.xxxx* on Ethernet arriving at a router via a Token Ring interface. [CSCdi05177]

When IPX extended access lists (lists numbered 900 through 999) are written to nonvolatile memory, explicitly specified port numbers are written using syntax that the configuration parser will not accept correctly. This has the effect of forcing all explicit port numbers to 0 when the configuration is reread. [CSCdi01836]

For the Novell protocol, the **ping** command round-trip time may be calculated incorrectly for ping packets with a data size of exactly 32 bytes. The numbers will be out of the range of possibility. [CSCdi04937]

XNS ping packets with a data size of 32 bytes may produce incorrect round trip times. The numbers will be unreasonably large. [CSCdi04984]

The command **show novell route net** will display the entire Novell routing table for Novell network numbers greater than 0x7fffffff. [CSCdi05048]

8.3(2) Caveats/8.3(3) Modifications

All the caveats listed in this section are resolved in release 8.3(3).

AppleTalk

AppleTalk fast-switching cache entries are not always invalidated when the metric associated with a route changes. This may result in misdelivery of some packets. [CSCdi04098]

The AppleTalk background process was erroneously changed to low priority. On a very busy router routes start aging out, even though updates were received in time. [CSCdi04191]

A problem exists with AppleTalk access lists. The problem is visible when a network entry hashes to the same value as the all-nets entry. Any network number that is a multiple of 64 + 1 will fail. To see if this problem exists in a particular configuration, examine the output of the **show configuration** command. If there is a duplicate entry, the list is broken. A possible work around is use a different network number that does not hash to the all-nets entry. [CSCdi04201]

When parallel paths exist, the AppleTalk fast-switching cache is invalidated too frequently. This has a negative impact on performance. [CSCdi04280]

Interfaces connected to end-nodes using AppleTalk for VMS, prior to version 3.01, should have the AppleTalk fast-switching cache disabled to ensure that all packets will be accepted by those end-nodes. [CSCdi04611]

CLNS

CLNS prefix routes that are advertised more than four hops away may not be retained in the routing table. Also, convergence for prefix routes is very slow: when they go away, it may take a long time for them to be removed; when they come back, it may take a long time for them to be relearned. [CSCdi04583]

When Decnet PhaseV (CLNS) packets are being converted to DECNET Phase IV, and CLNS fast switching is enabled for the output interface, all but the first packet for a given Phase IV destination will be dropped. This can be worked around by disabling CLNS fast switching on the output interface. [CSCdi03931]

If the **broadcast** command flag is set for a CLNS neighbor connected by an X.25 network, that neighbor will be sent IGRP updates even if it is an end system. This has no real operational impact, but is an unnecessary use of bandwidth. [CSCdi04416]

The route for an area will not be removed after that area is deleted. In addition the router will continue to use that NET after an area is gone. [CSCdi04680]

The router does not recognize CLNS packets as such unless CLNS routing is enabled. When CLNS packets are received over an HDLC serial line by a router without CLNS routing enabled, that router will log an "Unknown HDLC" message for each packet. The work around is to configure CLNS consistently at both ends of each serial line. [CSCdi02905]

If the router is assigned a CLNS NET using the **clns net** command, and that NET is then removed using the **no clns net** command, the router will continue to send intermediate system hello messages claiming the removed NET. Note that the **clns net** command is seldom used and is supported primarily for historical reasons. [CSCdi02578]

DECnet

When the router converts DECnet PhaseV packets into DECnet Phase IV packets, occasional packets are malformed. [CSCdi03717]

When a DECnet Phase IV packet is converted to a CLNS packet, the size of the CLNS packet buffer is computed incorrectly, causing overflow when converting large packets. This overflow may result in occasional malformed packets or in system reloads. [CSCdi03963]

DECnet PhaseV (CLNS) packets whose destination NSAPs have selector fields that do not correspond to NSP are not converted to Phase IV. [CSCdi04103]

DECnet Phase IV NCP commands directed to a DECnet Phase IV router across a DECnet PhaseV backbone do not pass through the DECnet PhaseV backbone correctly. This means that NCP commands cannot be executed across a DECnet Phase V backbone. When fixed, reachability still will be limited to routers no more than one hop away. [CSCdi04719]

IP Routing Protocols

Under some conditions the router may reload when the **show ip route** command is entered. [CSCdi04132]

It is possible for Cisco HDLC packets to be sent on interfaces configured for X.25, frame relay, or SMDS during router initialization. The actual sending of the packets has no known negative operational impact, but may result in illegal packet reports from frame relay switches. Sending of HDLC packets through X.25 interfaces, however, violates internal assumptions of the router software and may result in system reloads during initialization on X.25 networks. [CSCdi04462]

There is no way to disable application of the split horizon rule for IP routing. IP routes are not advertized over the interfaces through which they were learned. On a frame relay or SMDS network that is not connected in a full mesh where secondary IP addresses are in use, some routers will never exchange sufficient routing information, resulting in a partitioned network. The work around is to configure frame relay and SMDS networks such that all routers connected to them can communicate directly. This problem is resolved by the new interface configuration command **no ip split-horizon**. The improved code will disable split horizon by default on frame relay and SMDS interfaces. [CSCdi03430]

If a route learned from EGP in the local autonomous system is redistributed into BGP, and the route is to be sent to another internal BGP peer, that peer will refuse the BGP connection. [CSCdi03853]

When the next hop for a static route which is being redistributed into BGP is changed, the redistributed BGP route does not change. The work around is to remove all knowledge of the network before changing the static route. [CSCdi03863]

BGP next hop updates can be transmitted out the wrong interface. Insufficient checking of next hop information allows incorrect data to be entered into the routing table. [CSCdi04055]

When a default route is being learned from RIP, and there is more than one candidate default router with the same metric, the route chosen will oscillate among the candidates. Correct behavior is to choose one default route and use it until there is a real reason to change. [CSCdi04137]

If the system is directly connected to a subnetted major IP network, with its address on that network being one of its secondary addresses, and no default subnet exists for the major network in question, but the router does have a default route for general use, packets for unknown subnets may be forwarded through the main default route, which may send them outside of their major network entirely. This can be worked around by making one of the router's IP addresses on the major network in question a primary address. [CSCdi04215]

It is not possible to add a static interface route to null 0. [CSCdi04270]

IBM Connectivity

It is possible for frames being source-route bridged between CSC-R16 interfaces to be reordered. [CSCdi03110]

The Serial Tunnel **route** command will not parse changes to existing entries. Instead of overwriting the old, it will incorrectly add the new entry alongside the old. [CSCdi04310]

Interfaces and Bridging

ARP packets sent on FDDI sometimes use hardware type codes other than the Ethernet code. RFC 1188 calls for ARP on FDDI always to use the Ethernet code. [CSCdi04119]

If an error is made while configuring the encapsulation method, the encapsulation will incorrectly be set to NULL. This will be display as encapsulation unknown. [CSCdi03593]

It is possible for use of the **cbus-buffers** command on busy networks to cause system reloads at the time the command is processed. This is caused by a race condition, and failures are extremely rare. [CSCdi04033]

MAC level address access lists for SRB do not work. [CSCdi04559]

If you power-cycle one peer of an HDLC RSRB connection in 8.3(2), it will occasionally fail to re-establish the session. In this state, if you power cycle the other side, or if you remove then reinstate the remote-peer statement on the router that was cycled, it will re-establish the session. [CSCdi04508]

When the bandwidth parameter for an interface is changed while that interface is running the Spanning Tree protocol, the interface's path cost is not recalculated to reflect the change, even if the path cost was originally computed from the previous bandwidth setting. This results in the spontaneous appearance of a **path-cost** command in configuration files written after the change, because the path cost no longer reflects the default that would be calculated from the new bandwidth setting. The path cost may manually be set to match the cost that would have been calculated from the new bandwidth. [CSCdi03807]

For the IGS platform, bridge packets to multicast addresses using static bridge table entries do not work correctly. Packets were not getting forwarded to the multicast targets and the router was dropping them. This results in a loss of connectivity. [CSCdi04141]

It is possible for Cisco HDLC packets to be sent on interfaces configured for X.25, frame relay, or SMDS during router initialization. The actual sending of the packets has no known negative operational impact, but may result in illegal packet reports from frame relay switches. Sending of HDLC packets through X.25 interfaces, however, violates internal assumptions of the router software and may result in system reloads during initialization on X.25 networks. [CSCdi04462]

Basic System Services

DECnet static mapping addresses did not show up properly with the **show smds map** command. The addresses were incorrectly displayed as zero. CLNS did not work properly with SMDS encapsulation. AppleTalk did not work correctly with SMDS encapsulation. [CSCdi04322]

EXEC and Configuration Parser

Under some conditions the router may reload when the **show users** command is entered. [CSCdi04339]

If the user issues multiple **configure** commands, specifying configuration from the network, only the first dialog will default to the correct TFTP server. Subsequent dialogs will default to broadcast TFTP. [CSCdi04128]

Depending on the different types of error correction enabled (V.42, MNP, none) at the two modem sides, junk input characters may be passed to the terminal server as one modem attempts to negotiate a type of error correction that the other modem does not support. As a result, these junk characters are passed as input to the password prompt, and generally fail the login and disconnect the modem. [CSCdi04261]

If a **clear line** *n* command is issued for a line that has no process associated with it (for instance a SLIP line), the command will fail, and the line will not be cleared. [CSCdi04530]

Local Services

If extended TACACS is enabled, under certain rare conditions involving retransmissions, corrupted memory could cause the router to reload. [CSCdi04165]

The tsMsgTmpBanner and tsMsgSend variables can be neither read nor written. [CSCdi03894]

The ifMTU variable reflects the configured IP-specific MTU for the interface. It should reflect the configured overall/physical MTU. [CSCdi04022]

Under rare circumstances, sending of SNMP tty enterprise traps may result in router reloads. [CSCdi04138]

SRT Bridging

The default spanning tree path-cost value chosen for an interface is always computed according to the algorithm for IEEE spanning tree, even if the DEC spanning tree protocol is in use on that interface. This results in a default cost a factor of ten higher than that used by other DEC-compatible bridges for comparable media. This can be worked around by manually configuring a cost for each interface. [CSCdi04211]

TCP/IP Host-Mode Services

Computation of UDP checksums for packets whose UDP length fields have been corrupted may cause system reloads. [CSCdi03433]

TFTP over parallel links does not always behave correctly. [CSCdi01274]

X.25

Under heavy load, LAPB could mishandle the N(R) field in outgoing I-frames after receipt of a REJ frame. This caused the other end of the link to issue a FRMR frame to reset the link level, which has the side effect of clearing any X.25 virtual circuits going over the link. [CSCdi03558]

Under some conditions the router may reload when the **show x25 status** command is entered with X.25 debugging enabled. [CSCdi00832]

If X.25 switching is enabled, X.29 calls subaddresses of the system's main X.25 address will not be accepted and forwarded to rotaries as documented. [CSCdi03285]

In a SABM collision, it was possible for LAPB to get confused about its state. The link did come up, but only after a prolonged and unusual exchange of frames. [CSCdi03559]

A number of races exist in the X.25 code. These may result in the issuance of spurious traceback messages, or, rarely, in system reloads. Problems will be observed most often on busy X.25 links connected to busy routers. [CSCdi04948]

X.29 access lists are not checked for outgoing X.29 connections [CSCdi03891]

XNS/Novell/Apollo

Novell echo request packets from some versions of the system software previous to 9.0 are sent with an echo reply type code instead of an echo request code. Cisco 9.0 routers will not answer such echo requests. This means the Novell **ping** command will work from 9.0 to any 8.3/8.2 software version. It will not work from versions prior to 8.2(8)/8.3(3) to 9.0. [CSCdi03913]

SAP service entries will expire every three timeout intervals. This produces very unstable SAP tables causing poor performance. This problem was introduced in 8.3(2). [CSCdi04720]

The largest IPX packet size currently supported is 1500 bytes. This is not a problem except in networks utilizing Novell's BIGPACK.NLM. The correct behavior is to allow IPX packets up to the size of the interface MTU. [CSCdi04193]

The hold-down time used for Novell and XNS routes is six times the update interval. A more reasonable value is three to four times the interval. [CSCdi04238]

In a network with equal cost multiple paths, the router may hear advertisements for the same service through two interfaces. The advertisement coming from the second interface is accepted without verifying that it is from the same source as the entry in the SAP table. This prevents the SAP entry from aging out when the path thru the first entry no longer exists. This behavior can lead to some server/clients being isolated from the rest of the network. [CSCdi04327]

The router does not respond correctly to a Novell SAP get server request when the server type requested was *-1* (all services). This is not a very serious problem because very few applications use this function. [CSCdi04649]

Novell broadcasts with the destination network zero were not forwarded even when a helper address was present. Applications that depend on broadcasts to network zero being forwarded across the network will not work properly. [CSCdi04658]

For non-NetBIOS Novell service, flooding the helper address of *-1.ffff.ffff.ffff* is used when forwarding flooded traffic. *-1.ffff.ffff.ffff* translates to *ffffffff.ffff.ffff.ffff* when forwarded. Some Novell servers do not recognize the *ffffffff.ffff.ffff.ffff* broadcast address, and the flooded packet is ignored. The correct behavior is for the local net number to be used when flooding the packet. [CSCdi04494]

Novell access list checks are not applied to NetBIOS when flooding is enabled. The correct behavior is for NetBIOS traffic to be subject to the access list checks and not flooded by default. [CSCdi04496]

8.3(1) Caveats/8.3(2) Modifications

All the caveats listed in this section are resolved in release 8.3(2).

AppleTalk

When a route is deleted from the AppleTalk routing table, there is a possibility of corruption of the table data structure. This corruption most often results in system reloads shortly thereafter. This problem is most often observed in very unstable networks. There is no direct work around, but the frequency of failures can be reduced by correcting flapping lines and other sources of instability. [CSCdi03060]

If more than one **appletalk proxy-nbp** command is issued for the same network number, the system will pause indefinitely. This can be avoided by not issuing the **appletalk proxy-nbp** command for networks which have already been specified in such commands. [CSCdi03061]

It is not possible to configure an SMDS or frame relay network as an AppleTalk network. [CSCdi03106]

The data length fields of 802.3 packets containing AppleTalk data are sometimes set incorrectly. Some implementations will ignore such packets or count them as errors. Connections with such implementations through Cisco routers may fail either consistently or sporadically. [CSCdi03377]

It is possible, but rare, for corruption of system data structures to take place during gleaning of node MAC addresses from AppleTalk transit traffic. Such corruption may result in system reloads and/or in the issuance of SYS-2-SMASHED messages. [CSCdi03397]

A race condition between the AppleTalk routing and memory management processes may occasionally result in system reloads. [CSCdi03720]

Filters applied to AppleTalk routing updates using the **appletalk distribute-list** command are not applied to responses to ZIP GetZoneList queries. This may result in clients receiving information about zones and networks they cannot actually reach, which may in turn result in services being offered in user menus when the services are not in fact available. [CSCdi02688]

It is not possible to delete routing filters using the **no appletalk distribute-list** *n* **in | out** command. You can remove an AppleTalk routing filter by disabling AppleTalk and reconfiguring it from scratch. [CSCdi02729]

The **appletalk nbp-proxy** global configuration command is never written to NVRAM or to remote configuration files. As a work around, the command can be added to a remote configuration file using a text editor. [CSCdi02792]

Under some circumstances, the **show apple** command may display the number of busy nodes as negative. [CSCdi03659]

CLNS

Different functional addresses are used for ES-IS in different versions of the standard for CLNS over Token Ring networks; not all of these addresses are supported. The router is unable to exchange ES-IS frames with nodes using functional addresses other than the ones it knows. The correct behavior for Cisco is to support all the functional addresses actually in use on installed networks. [CSCdi02903]

It is possible under some circumstances for IGRP, ISO IGRP, and IS-IS processes to overflow their process stacks when their associated routing protocols are used over X.25 networks. This can result in system reloads. [CSCdi03124]

Intermediate system hellos are never sent on interfaces configured with the **clns enable** command. The work around is to use the newer **clns router static** command syntax. [CSCdi03258]

DECnet

If there is more than one possible path to a DECnet destination, and if DECnet fast switching is disabled for the output interface(s) associated with one or more of the paths while being enabled on the interface(s) associated with the other(s), an error in the internal traffic allocation logic may cause traffic to avoid one of the paths completely. This can be worked around by enabling DECnet fast switching either on all interfaces that might fall into a load-sharing set or on no interfaces that might fall into that set. Cisco recommends consistent use of fast-switching options on load-shared interfaces regardless of the presence of this caveat. [CSCdi02689]

It is possible for incorrect values to be placed in the selector fields of NSP packets being converted from DECnet Phase IV to DECnet Phase V. [CSCdi03109]

When converting packets from DECnet Phase V to DECnet Phase IV, the algorithm for determining if the selector field is a valid NSP value is wrong. As a result, some packets which have valid NSP values will not be converted from DECnet Phase V to DECnet Phase IV. [CSCdi03145]

The DECnet Phase IV destination area number is used to form the source area number of the output packet when a packet is being converted from DECnet Phase IV to DECnet Phase V. The correct behavior is to use the DECnet Phase IV source area number to create the DECnet Phase V source area number. [CSCdi03562]

It is possible for Cisco's MOP server to send MOP console carrier packets with lengths greater than 256. Some MOP products (including the DECServer 90L), do not accept packets this long. [CSCdi03667]

IP and IP Routing Protocols

It is possible for a race between the code for BGP and the code for other IP routing protocols to result in system reloads. [CSCdi02834]

When IP routing updates are sent through interfaces that have secondary addresses that lie in different major networks than their primary addresses, the split horizon rule is not applied to information about the secondary networks. The operational impact of this behavior is minimal, and it can be worked around entirely by the use of output routing filters. [CSCdi01355]

Routers that are heavily loaded and that are sending traffic into congested X.25 networks may issue the SYS-2-INTSCHED messages. These messages may appear in such numbers as to make the router's console unusable. Routers that are running dynamic routing protocols and injecting large routing updates into X.25 networks are especially vulnerable to this failure. The work around is to reduce network congestion. [CSCdi02772]

When an IP RIP update is sent from a secondary IP address, no more than one packet of data is sent, regardless of the actual amount of routing data eligible for inclusion. In addition, updates never contain data regarding major networks other than the network in which the secondary address lies, nor do they contain default route data. [CSCdi02857]

The typical size of EGP packets on the MILNET has become too large for the internal buffers used to process such packets. The router may ignore EGP packets received from the MILNET. [CSCdi02898]

The **show ip redirect** and **show ip aliases** commands do not exist on routers. When IP routing is enabled, these commands do not provide useful information, but when IP routing has been disabled with the **no ip routing** command, their output may be of interest. [CSCdi02980]

When an IP RIP update containing exactly one maximum-sized packet's worth of entries is generated, it is followed by a RIP packet containing no entries. Such packets are illegal and may cause error reports to be issued by third-party equipment. [CSCdi03059]

IP routes that use an interface are not deleted immediately when the **no ip address** command is given for that interface. The work around is to remove the routes manually using the **clear ip route** EXEC command. [CSCdi03319]

The error message returned by BGP when a peer system attempts to open a connection using a version number of 3 or higher requests the use of an illegal protocol version instead of the use of version 2. This results in incorrect version negotiation with third-party equipment. [CSCdi03358]

If a static IP route is configured via a gateway which is not directly reachable, and an alternate route exists to that gateway, the configured gateway's address will be overwritten in the routing table and in saved configurations with that of the first hop router in the alternate path. [CSCdi03419]

IP RIP updates are not sent from secondary addresses when the secondary major networks are not subnetted. [CSCdi03638]

RIP default routes will never replace static routes to net 0.0.0.0 in the IP routing table, regardless of the administrative distances assigned. [CSCdi03701]

Attempts to change the autonomous system number associated with an EGP neighbor always fail. This can be worked around by reconfiguring, then reloading, the router. [CSCdi03702]

If memory is exhausted, the router may fail to properly process **network** commands, without giving any indication to the user that the commands have failed. [CSCdi02816]

If a dynamic ARP reply is received for an IP address for which a static ARP table entry has been configured, the static entry will be overwritten by the dynamic information. Correct behavior would be to ignore ARP replies for addresses with static ARP entries. [CSCdi00118]

Because of a race between the code for printing the IP routing table and the code that actually maintains that table, it is possible for use of the **show ip routes** command to result in system reloads. This is especially likely in unstable networks. The **show ip routes** command should be used with care. [CSCdi03277]

The **no ip-forward-protocol udp** command does not reinitialize the UDP forwarding table to the default before disabling UDP forwarding. A later **ip forward-protocol udp** command causes earlier port enable/disables to become active again. [CSCdi03261]

IBM Connectivity

It is not possible to use SDLC tunneling in a system with DECnet routing enabled (or vice versa). [CSCdi03170]

In certain corner cases, SDLC proxy polling can cause an extra RR to be sent from the primary host, causing the secondary to resend its first I-frame in a series twice. This does not affect functionality, and has minimal impact on performance. [CSCdi03173]

A Serial Tunnel (STUN) TCP connection to a remote Cisco router could hang in the rare circumstance of the TCP connection being aborted by one side of the connection at the precise moment that the other end of the connection was just finishing reading previously sent data from the side closing the connection. In practice, this rarely occurs because TCP connections that abort due to an error usually do so after a long idle period in the traffic flow between the two TCP peers. [CSCdi03648]

It is possible to define a new STUN schema that has the same name as an existing predefined STUN type (such as SDLC). When such a new definition is made, it overrides the existing predefined definition type requiring a reload of the router to restore the accessibility of the predefined version. [CSCdi03066]

Interfaces and Bridging

Frame relay DLCI numbers are not learned properly for the MAC addresses of nodes across frame relay networks. This results in excessive frame relay multicasting of bridged traffic. [CSCdi03103]

The D15 mode of SMDS is not supported. [CSCdi03660]

It is possible for IGS routers to choose spanning tree bridge identifiers that are not based on their actual Ethernet/802.3 addresses. Furthermore, these identifiers are chosen from a relatively small number of possibilities, and often will overlap. This may cause disruption of spanning trees. [CSCdi03703]

The IGS serial interface cannot receive any frames until after it has itself sent at least one frame. This generally has minimal operational impact except for SDLC tunneling. If the IGS is connected to an SDLC primary device, it must wait for a poll from the primary before sending any data. Since the IGS cannot receive the poll until some data has been sent, the line is never activated. This can be worked around by changing the line encapsulation to HDLC for a brief period when the system is first brought up. [CSCdi03820]

It is possible for interface-related counter values returned by SNMP to decrease between successive samples when they are expected to increase monotonically. The conditions under which this occurs are not yet well understood. [CSCdi02452]

When an IGS router is bridging Ethernet traffic onto a congested HDLC serial line, some packets may be corrupted. The corruption will consist of the insertion of extra data bytes before the destination MAC address. This will result in undesired traffic on the remote Ethernet and in erroneous bridging cache entries on the remote router. [CSCdi02563]

It is theoretically possible for garbage messages to be issued when certain types of CSC-R16 failures occur. These failures have never been observed with released Cisco software. [CSCdi02618]

If an interface's MTU is adjusted upward, the IP and CLNS MTUs for that interface are not adjusted to match. The correct behavior is to adjust the IP and CLNS MTUs unless they have been explicitly configured to be different from the interface MTU. [CSCdi02684]

IGS routers will not bridge DEC RBMS (Remote Bridge Management System) frames. [CSCdi02872]

Type 2 (Interlan) CSC-E Ethernet interfaces may experience rare output hangs. Type 2 interfaces were eliminated from Cisco's product line several years ago and are not supported with CSC/3 processors. [CSCdi02927]

The router does not respond to HP probe packets that use Ethernet (ARPA) encapsulation. The router does not properly bridge HP probe name requests and replies to and from HP DTC devices. The router does not listen to HP probe unsolicited replies, resulting in poor performance. The router does not generate HP probe VNA requests in Ethernet encapsulation. Due to the additional overhead, the interface configuration command **no arp probe** is now the default. [CSCdi02949]

Entries may occasionally be dropped from the frame relay DLCI map for an interface. This occurs when new entries are added, and is more likely when large numbers of map entries exist. [CSCdi03355]

cBus buffer sizes for UltraNet interfaces are sometimes set to too small values. This may result in inability to receive or transmit maximum-sized UltraNet datagrams. [CSCdi03438]

The system will allow configuration of priority queueing for LAPB interfaces. This should not be done; configuring priority queueing on a LAPB interface will result in LAPB protocol errors. [CSCdi03500]

The **slip access-class** configuration command is written to nonvolatile memory and to remote configuration files as **slip acess-class**. The system will not parse the files correctly when they are read back in. [CSCdi03630]

If an asynchronous connection is lost while a SLIP packet is being transmitted over the line, the packet buffer for that packet will not be returned to the free buffer pool. In addition, the packet will remain permanently charged against the input queue quota for the interface on which it arrived. Over very long periods, these conditions can have the cumulative effect of shutting down a terminal server and/or its network interface. This can often be worked around by remedying conditions that lead to unexpected modem line drops and/or by occasionally reloading the terminal server. [CSCdi03785]

A **frame-relay local-dlci** command will be written to NVRAM or to a network configuration file even if the configured local DLCI is the default. This is harmless. [CSCdi03846]

If IP routing is disabled, and an IP packet is sent out of a serial line, the packet is sent as a bridged packet, even if bridging is not enabled. This can lead to an inability to communicate across serial lines between routers which are neither bridging nor routing IP. [CSCdi02692]

It is possible for the caching of Token Ring RIFs to cause router reloads. This is especially likely in busy networks. This limitation can sometimes be worked around by disabling multiring mode on Token Ring interfaces. [CSCdi03298]

Basic System Services

The maximum number of "middle" buffers that can be allocated in an IGS is lower than the number many applications require to operate comfortably. [CSCdi02961]

There is no way to see the internal state of the environmental monitor card from the system command interpreter. The **show envm** command will remedy this. [CSCdi02761]

The **show interface** command display does not mention the fact that the interface counters have never been cleared if they have not been, but it does mention when they were cleared if they have been. [CSCdi02882]

It is possible for use of the **show host** command while the host-name cache is being updated to result in system reloads. The **show host** command should be used with care. [CSCdi02918]

The **clear line** command has no effect on lines configured for SLIP. [CSCdi03372]

EXEC and Configuration Parser

The **arp** interface configuration command does not work on STS-10X terminal servers. [CSCdi02979]

Local Services

If no domain name has been set using the **ip domain-name** command, the value returned for the SNMP sysName variable will be invalid. [CSCdi03250]

The fact that the system enable password is always accepted as a read-write SNMP community string creates a security hole. Correct behavior is to require the user to explicitly configure any community strings to be used. [CSCdi03418]

Incorrect data are returned for the ifPhysAddress MIB variable on FDDI interfaces. [CSCdi03568]

When multiple **boot host** commands are specified, there is no failover from the primary server to the secondary server(s). [CSCdi03290]

SRT Bridging

A race exists between the transparent bridging code for learning MAC address locations from unicast packets and that for learning them from broadcasts or multicasts. In busy networks with many nodes, this race may cause corruption of internal bridging data structures. This corruption causes the router to cease functioning without reloading; the only work around is to manually reload the router. [CSCdi03636]

On boot up, on the IGS platform, bridging does not work over HDLC. Clearing the serial line should restore functionality. [CSCdi02959]

TCP/IP Host-Mode Services

Transit packets from which the router has stripped IP security options are output malformed. The work around is to disable stripping of security options. [CSCdi02286]

Overly optimistic assumptions are made about path latency when an incoming TCP connection is accepted. This may result in over-eager retransmission during the early life of the connection. [CSCdi03099]

When a TCP segment is acknowledged, the software does not reset the time for retransmission based on the original transmission time of the following segment (if one is queued), but does the first retransmission of the following segment at the time it would have retransmitted the acknowledged segment. This can cause many extra retransmissions when the time between packet sends is close to the calculated initial round-trip time. [CSCdi03136]

HP Probe is on by default. This has been determined to be nonoptimal in most user environments. The correct behavior is for this to be off by default. [CSCdi03597]

If a TFTP transfer is in progress, and the system receives a retransmission or other packet while expecting an acknowledgment, the transfer will be aborted completely. This can generally be worked around by retrying transfers or configuring the system to retry automatic transfers. Operational impact is usually minor. [CSCdi03810]

VINES

Banyan VINES did not work properly over frame relay. [CSCdi03100]

X.25

When an AppleTalk broadcast packet (usually a routing update) is replicated for transmission via multiple virtual circuits on an X.25 interface, all copies but the first are corrupted. This means that it is essentially impossible to use AppleTalk over X.25 with more than one remote router on the X.25 network. [CSCdi03122]

Clearing X.25 virtual circuits with the **clear x25-vc** command may result in system reloads, especially when many circuits are being established and cleared by other means. The **clear x25-vc** command should be used with caution in busy environments. [CSCdi01622]

When transparent bridging is being used over X.25 links, it is possible for a race condition to cause system reloads or other unexpected, apparently nondeterministic behavior. [CSCdi03178]

If the NVC option is changed for an interface, this change is not properly executed. It may be applied to another unrelated X25 interface. [CSCdi03790]

System software cannot be booted over X.25 links. [CSCdi03811]

XNS/Novell/Apollo

Attempts to reduce the maximum number of parallel paths available to XNS, Novell, or Apollo traffic (using the **xns maximum-paths**, **novell maximum-paths**, or **apollo maximum-paths** command) will result in a router reload. To reduce the maximum number of available paths, disable routing for the protocol in question entirely, and reconfigure that protocol from scratch. [CSCdi02775]

Replies to Novell RIP requests are sometimes sent with destination network numbers of zero. The correct behavior is to use a destination reflecting the network number actually used on the cable. Some Novell applications rely on the correct behavior, and will not learn their network numbers properly if it is not followed. [CSCdi02779]

If a Novell SAP packet that would ordinarily cause the sending of a flash update is received, but output SAP filters prevent the sending of the actual flash update, a buffer will be lost. In unstable networks, the cumulative effect of such lost buffers will be the complete depletion of the router's memory pool. In addition, if a flash update would ordinarily be sent, but the interface through which the update would be sent is not up, a "SYS-2-INLIST" message will be issued. This latter behavior is harmless, but often results in flurries of "SYS-2-INLIST" messages being issued at startup, especially on routers with Token Ring interfaces. [CSCdi02876]

Some third-party Novell applications issue SAP updates listing services with network numbers of zero. The system readvertises these services on its other networks with the original zero network numbers. The correct behavior is to rewrite zero network numbers to the network number of the network on which the update was received. [CSCdi01348]

If multiple flash updates are sent in response to a Novell SAP packet, the hop count(s) in each flash update sent will be one greater than the hop counts(s) in the previous one. The correct behavior would be to have all flash update hop counts the same, and one greater than the value in the original input packet. [CSCdi02571]

If an XNS error report packet is received, but cannot be forwarded because no route to its destination is known, the buffer holding that packet will not be returned to the free pool. In unusual environments and/or over very long uptimes, this can result in router failure. [CSCdi02863]

The command **no apollo access-group** *x* is not interpreted correctly. The only way to remove an Apollo access list from an interface is to shut Apollo routing down entirely and reconfigure it from scratch. [CSCdi03133]

When Novell or XNS RIP updates are sent, networks which are denied by routing filters are mentioned in the updates, but with hop counts of 16 (RIP's "infinity" hop count). While this does not produce any routing problems Cisco is aware of, it is an inefficient use of bandwidth. Correct behavior would be not to mention the filtered networks at all. [CSCdi03517]

The command **show access-lists** does not display access lists defined for the Apollo Domain routing protocol. The correct behavior is to display the contents of all access lists. [CSCdi02864]

New Microcode Version Levels

No new microcode was released specifically for Software Release 8.3(2-4); however, new microcode was made available with the 8.3(1) system software release for several of the AGS+ cards. Table 1 outlines the minimum microcode versions required for key new features available with 8.3(1). In addition, new microcode was released with Software Release 9.0 that provides bug fixes and improved functionality for 8.3 customers. Table 2 summarizes the most recent microcode versions available to 8.3 users. Subsequent sections describe the features, modifications, and caveats for the various microcode versions.

Note: In Table 1, the upgrade requirements apply to *all* boards listed for particular features. The upgrades are mandatory for users who choose to implement these new features.

Table 1 Minimum Microcode Upgrade Requirements for New 8.3 Features

Feature	cBus	MEC 5.0*	MEC 5.1*	FDDI	Miscellaneous
IP Autonomous Switching	2.0	1.7	2.2	1.0	
HSSI, UltraNet	2.0	1.7	2.2	1.0	HSCI 1.0
SMDS and Frame Relay					SCI 1.2 or MCI 1.7

* MEC microcode versions depend on the hardware version. Refer to the section "MEC Microcode Revisions."

Note: These are the minimum levels required for the listed features. For the recommended and most recent versions, see Table 2.

Microcode Interoperability Summary

Table 2 shows the compatibility between microcode versions (including the latest ones, released with 8.3(1)) and various system software versions. Refer to these numbers when ordering microcode upgrades.

Table 2 New Microcode Compatibility

System Software	SCI	MCI	CSC-R	CSC-R16	ENVM	CCTL	MEC Revs* 5.0	MEC Revs* 5.1	FDDI	HSCI
8.2(4–5)	1.2	1.8	2.2	2.0	1.1	2.0	1.7	2.2	1.0	N/A
8.2(6-8)	1.2	1.8	2.2	3.0	1.1	2.0	1.7	2.2	1.0	N/A
8.3(1-5)	1.3	1.9	2.2	3.0	2.0	2.0	1.7	2.3	1.0	1.0

For users who choose not to upgrade to the new microcode levels, Table 3 shows the latest microcode versions recommended for each system software version.

Table 3 Compatibility of Older Microcode Levels

System Software	SCI	MCI	CSC-R	CSC-R16	ENVM	CCTL	MEC Revs* 5.0	MEC Revs* 5.1	FDDI	HSCI
8.1(25) & 8.2(1–3)	1.1	1.8	2.2	N/A	1.1	1.0	1.6	2.1	128.43	N/A
8.2(4–5)	1.1	1.8	2.2	2.0	1.1	1.0	1.6	2.1	128.45	N/A
8.2(6-8)	1.1	1.8	2.2	3.0	1.1	1.0	1.6	2.1	128.45	N/A
8.3(1-5)	1.1	1.8	2.2	3.0	1.1	1.0	1.6	2.1	128.45	N/A

*MEC microcode versions depend on the hardware version (see the section "MEC Microcode Revisions").

cBus Controller Microcode Revisions

This section describes the microcode revisions to the Cisco proprietary cBus controller card (CSC-CCTL).

cBus Microcode Version 1.0

Microcode version 1.0 was the first officially released version of cBus microcode.

cBus Microcode Version 2.0

Microcode version 2.0 for the cBus was introduced with Software Release 8.3(1) to support IP autonomous switching and the new HSSI interface.Version 2.0 requires the replacement of all nine registered EPROMs. For more information on the upgrade procedure from version 1.0 to 2.0, refer to the Cisco publication *Upgrading the cBus Controller Card Microcode Version 1.0 to Version 2.0* (Part Number 78-0856).

Modifications

- Support for IP Autonomous Switching has been added.
- Support for the HSCI interface controller has been added.

System Software and Microcode Prerequisites

- Requires Release 8.3(1) for autonomous switching or HSCI support.
- Refer to Tables 1 through 3 for information on the microcode levels required for other boards used in conjunction with the CSC-CCTL.

CSC-R16 Microcode Revisions

This section describes the microcode revisions to the CSC-R16 Token Ring interface card.

CSC-R16 Microcode Version 2.0

This was the first officially released version of CSC-R16/SBEMON microcode. It was released in June 1991.

Caveats

If cabling problems exist, the router prints the following error message and does not attempt to restart the interface:

```
%TR-3-WIREFAULT: Unit 1, wire fault: check the lobe cable MAU
connection
```

This occurs to prevent flapping of routes in the various network layer protocols. This behavior is new as of release 8.2(6). An explicit **clear interface** command will cause the router to attempt to restart the interface.

System Software Prerequisites

- Interoperates with Release 8.2(4) and later system software.

CSC-R16 Microcode Version 3.0

This version of CSC-R16 microcode contains Madge microcode and requires the replacement of the two SBEMON ROMs on the CSC-R16 card. For more information on the upgrade procedure from version 2.0 to 3.0, refer to the Cisco publication *Upgrading the CSC-R16 Token Ring Interface Card to Microcode Version 3.0* (Part Number 78-0848).

Modifications

- Incorporates Madge microcode for significantly improved performance.

Caveats

Continued extreme load conditions can result in faulty error counts due to priority lockout of the SBEMON firmware. If cabling problems exist, the router prints the following error message and does not attempt to restart the interface:

```
%TR-3-WIREFAULT: Unit 1, wire fault: check the lobe cable MAU
connection
```

This occurs to prevent flapping of routes in the various network layer protocols. An explicit **clear interface** command will cause the router to attempt to restart the interface.

System Software Prerequisites

- Requires system software 8.2(6) or later.

Environmental Monitor (ENVM) Microcode Revisions

This section describes the microcode revisions to the environmental monitor (ENVM) card.

ENVM Microcode Version 1.1

Version 1.1 was the first officially released version of ENVM ("ECMON") microcode.

ENVM Microcode Version 2.0

Released with SR 9.0 in April 1992, ECMON 2.0 provides additional accuracy and functionality to the ENVM card. For information on the upgrade procedure from version 1.1 to 2.0, refer to *Installing and Upgrading ENVM Cards* (Part Number 78-0899-01).

Modifications

- Gives improved accuracy in voltage, temperature and air flow sensing, warning messages and shutdown
- Provides a log of previous shutdown history upon subsequent boot
- Provides access to environmental parameters via SNMP queries

System Software Prerequisites

- Requires System Software 9.0 or later for SNMP query features

FDDI Microcode Revisions

This section describes the microcode revisions to the Cisco Fiber Distributed Data Interface controller card (CSC-FCI).

Note: The revision numbering scheme for the FDDI microcode has been amended to align with the numbering scheme for Cisco's other microcode versions; thus, FDDI microcode version 1.0 is the latest revision.

FDDI Microcode Version 128.43

FDDI Microcode Version 128.43 was the first officially released version of FDDI microcode. It was released in May 1990.

Caveats

- Routers can hang in high-traffic environments where the FDDI ring is extremely unstable (continual transitions). Customer perceived problems follow:
  - — The router locks up, and a power cycle is the only way to correct the problem.
  - — Performance degrades over time.
  - — High instances of output hangs and cBus controller restarts occur.
  - — Symptoms may include the following error message (traceback information, which is release version-dependent, needs to be examined).

```
May 21 15:24:23 145.1.89.254 19: %SYS-2-GETBUF: Bad getbuffer, bytes= 16654
-Process="Virtual Exec", level=4, pid= 41 -Traceback=5014 4510E 66BE 9E380
9B408 A07B8 A1 028 A0100
```

- Performance can degrade significantly due to excessive processor utilization if a MIC connector is pulled from one of the PHY ports. (Both PHY ports erroneously contend to be active.) The problem can be remedied with the **cmt disconnect phy {A | B}** command when the problem is observed.

FDDI Microcode Version 128.45

Released in June 1991, this version requires the replacement of all eight EPROMs.

Modifications

- Fixes potential hang condition in high-traffic environments.

Caveats

- Performance can degrade significantly due to excessive processor utilization if a MIC connector is pulled from one of the PHY ports. (Both PHY ports erroneously contend to be active.) The problem can be remedied with the **cmt disconnect phy {A | B}** command when the problem is observed.

System Software Prerequisites

- Does not work with Software Releases 8.2(1), 8.2(2), or 8.2(3).

FDDI Microcode Version 1.0

Released in October 1991, FDDI microcode version 1.0 requires the replacement of all eight registered EPROMs. For more information on the upgrade procedure from version 128.43 or version 128.45 to version 1.0, refer to the Cisco publication *Upgrading the CSC-FCI Card* (Part Number 78-0857).

Modifications

- Fixes the excessive processor utilization problem when a MIC connector is pulled from one of the PHY ports.
- Performance is improved in excessively bursty environments (such as environments with frequent NFS timeouts).
- Support has been added to interoperate with most single mode and multimode converters.

System Software and Microcode Dependencies

- For information on the microcode levels required for other boards used in conjunction with the CSC-FCI, refer to Tables 1 through 3.
- For FDDI microcode version 1.0, Software Release 8.2(4) or later is required for support of the three modifications listed above.
- When upgrading the FDDI microcode to version 1.0, the microcode on the cBus card (CSC-CCTL) must also be upgraded to version 2.0 if it has not already been done. FDDI version 1.0 will not interoperate with CSC-CCTL version 1.0.

Note on System Limitation

A condition can arise in Software Release 8.3 in some fully loaded AGS+ configurations, specifically if all four cBus slots are used and there is at least one MEC, one FDDI, and one HSCI interface controller card, and the fourth slot contains an MEC or FDDI interface controller, where the transmit and receive memory buffers are oversubscribed. This is generally due to the default MTU size being set to a size greater than 1500 bytes for the HSSI or Ultranet interface. This condition can result in excessive numbers of dropped packets on the FDDI interface and can cause the FDDI interface to hang and produce the following error message:

```
Interface fddi 0 output hung, restarting cBus 0 controller -
mci_output()
```

The interface can be reset by entering the following command:

```
Router>clear interface fddi
```

The condition can be avoided by manually setting the HSCI MTU size to the default Ethernet size of 1500 by entering the following interface subcommand in configuration mode:

```
mtu 1500
```

This command will effectively alleviate buffer starvation.

HSCI Microcode Revisions

This section describes the microcode revisions to the Cisco High-speed Serial Communications Interface (HSCI), which supports both the High-speed Serial Interface (HSSI) and UltraNet interface specifications.

HSCI Microcode Version 1.0

HSCI microcode version 1.0 is the first officially released version of HSCI microcode.

System Software and Microcode Dependencies

- HSCI microcode version 1.0 requires Software Release 8.3(1) or later.
- For information on the microcode levels required for other boards used in conjunction with the HSCI, refer to Tables 1 through 3.

MCI Microcode Revision Summary

The following section describes the various revisions of microcode for the Multiport Communications Interface card (CSC-MCI).

Note: Unless otherwise stated, the system software prerequisites (earliest level required) can be found in Tables 2 and 3 in the previous section, "Microcode Interoperability Summary."

MCI Microcode Version 1.4

Microcode Version 1.4 was the first officially released version of MCI microcode. It was released in the summer of 1988.

Caveats

- Ethernet might hang with keepalives turned on in high collision environments.
- Serial line reports DCD transitions when the interface is disabled. This causes the excessive error rate shutdown feature to fail when DCD flapping occurs.
- AppleTalk fast switching is not supported.

MCI Microcode Version 1.5

Released in March 1989, this version requires the replacement of all 13 registered EPROMs.

Modifications

- Support for non-volatile memory card (CSC-MC and CSC-MT).
- Ethernets no longer hang with keepalives turned on in high collision environments.
- Serial DCD transition interrupts no longer occur when the interface is disabled.
- Supports additional protocols for use with fast switching. The protocol-specific fast-switching support is specified in the "Features" section of the appropriate software release note.

Caveats

- Serial interface reports CRC errors as encapsulation failure.
- An Ethernet interface, under certain traffic patterns, can receive an errored frame as often as once in 3000 to 4000 frames.

MCI Microcode Version 1.6

Released in July 1990, this version only requires the replacement of two of the registered EPROMs if MCI microcode version 1.5 is currently installed. For more information on the upgrade procedure, refer to the Cisco publication *Upgrading MCI Cards to Version 1.6* (Part Number 78-0722).

Modifications

- Serial interface reports of CRC errors as encapsulation failure has been corrected.

Caveats

- Serial performance with X.25 at rates of 2 Mbps or higher can be impacted by very short (2- to 6-byte) packets.

- In extremely rare cases, packets received on Ethernet can have the last byte or word corrupted due to extended dribble bit errors from certain other vendors' equipment.

System Software Prerequisites

- Requires system code 8.1(14) or later.

MCI Microcode Version 1.7

Released in July 1990, this version requires replacement of all 13 of the registered EPROMs.

Modifications

- In MCI microcode version 1.5, an Ethernet interface, under certain traffic patterns, could receive an errored frame as good once in 3000 to 4000 frames. This has been corrected.
- Supports additional protocols for use with fast switching. The protocol-specific fast-switching support is specified in the "Features" section of the appropriate system software release note.
- Ethernet enhanced to completely receive frames which have the framing or dribble bit error.
- It has been shown that the SEEQ Ethernet controller used on the MCI reports framing or dribble errors with certain other vendors' Ethernet interfaces. This behavior has been most prevalent in 10Base-T environments or where Ethernet hub or concentrator equipment is used. With MCI microcode version 1.7, in combination with Software Release version 8.1(19) or later, if only framing errors are reported, then frames which have the framing or dribble bit error will be received as good.
- In combination with the newer revision (R68561AP) of the Rockwell MPCC integrated circuit, MCI microcode Version 1.7 now supports the sending of zero in the first byte. This capability is required for SMDS and frame relay.

Caveats

- An Ethernet interface could experience repeated resets in an environment of late collisions in fiber or broadband Ethernets and, if keepalives are turned off, then appear to hang.
- Serial performance with X.25 at rates of 2 Mbps or higher can be impacted by very short (2- to 6-byte) packets.
- In extremely rare cases, packets received on Ethernet can have the last byte or word corrupted due to extended dribble bit errors from certain other vendors' equipment.

System Software Prerequisites

- Requires system code 8.1(19) or later.

MCI Microcode Version 1.8

Released in January 1991, this version only requires the replacement of two of the registered EPROMs if MCI microcode version 1.7 is currently installed. For more information on the upgrade procedure from version 1.5 to 1.8, refer to the Cisco publication *Upgrading MCI Microcode Version 1.5 to Version 1.8* (Part Number 78-0728). For more information on the upgrade procedure from version 1.7 to 1.8, refer to the Cisco publication *Upgrading MCI Microcode Version 1.7 to Version 1.8* (Part Number 78-0766).

Modifications

- Previously an Ethernet interface could experience repeated resets in an environment of late collisions in fiber or broadband Ethernets and, if keepalives were turned off, then appear to hang. This condition has been corrected.

Caveats

- Serial performance with X.25 at rates of 2 Mbps or higher can be impacted by very short (2- to 6-byte) packets.
- In extremely rare cases, packets received on Ethernet can have the last byte or word corrupted due to extended dribble bit errors from certain other vendors' equipment.

System Software Prerequisites

- Interoperates with Release 8.1(25) and later system software.

MCI Microcode Version 1.9

Released with SR 9.0 in April 1992, this version requires replacement of all 13 registered EPROMs. For more information on the upgrade procedure from older versions of MCI microcode, refer to *Upgrading MCI Cards from Microcode Version 1.5 to 1.9* (Part Number 78-0728-01), or *Upgrading MCI Cards from Microcode Version 1.7/1.8 to 1.9* (Part Number 78-0766-01).

Modifications

- Improved serial performance with X.25 at rates of 2 Mbps or higher.

- Previously, in extremely rare cases, packets received on Ethernet could have the last byte or word corrupted due to extended dribble bit errors from certain other vendors' equipment. This has been corrected.

System Software Prerequisites

- Interoperates with Release 8.1(25) and later system software.

MEC Microcode Revisions

This section describes the microcode revisions to the Cisco Multiple Ethernet Controller (MEC) interface card.

The proper version of microcode for use with an MEC card depends on the hardware revision of the card. To determine revision levels and microcode versions in use on a given system, use the **show controller cbus** command, or examine the label on the card edge.

The following limitations apply:

Table 4 MEC Hardware/Microcode Dependencies

Controller Type	Revision Level	Compatible Microcode
5.0 A*	1.0	1.1, 1.4
5.0 C,D	1.0	1.6, 1.7**
5.1 E,F	3.0	2.1, 2.2, 2.3***

*MEC controller type 5.0, revision A, is incompatible with the new MEC microcode. Customers using this revision should upgrade.
**Version 1.7 is provided for customers with cards of earlier revisions who would like to use Autonomous Switching.
*** MEC microcode version 2.2 is required for autonomous switching.

Note: An MEC controller type 5.0 with 1.6 microcode is identical in function to an MEC controller type 5.1 running 2.1 microcode. Similarly, an MEC controller type 5.0 with microcode version 1.7 is identical in function to an MEC controller type 5.1 running 2.2 microcode.

MEC Microcode Version 1.1

MEC Microcode Version 1.1 was the first officially released version of MEC microcode.

Caveats

- Under heavy traffic loads, the cBus controller memory can be corrupted, resulting in a gradual decrease in performance. In extreme cases, an interface can hang.

MEC Microcode Version 1.4

This version requires the replacement of three registered EPROMs if MEC microcode version 1.1 is currently installed. Refer to Table 4 at the beginning of the "MEC Microcode Revision" section for hardware revision prerequisites associated with this version of microcode.

Modifications

- MEC microcode version 1.4 fixes the caveat in MEC version 1.1.

MEC Microcode Version 1.6

MEC Microcode Version 1.6 was released in February 1991. This version requires the replacement of all 13 registered EPROMs. For hardware revision prerequisites associated with this version of microcode, refer to Table 4 at the beginning of the "MEC Microcode Revision" section.

Modifications

- MEC microcode version 1.6 supports the hardware upgrade from 5.0, Revision A to 5.0, Revision C.

Caveats

- MEC Microcode Version 1.6 does not properly support packets starting on odd-byte boundaries. This problem impacts AppleTalk fast-switching with some other vendors' software.

MEC Microcode Version 1.7

Released with the 8.3(1) system software in October 1991, this version requires the replacement of all 13 registered EPROMs. For more information on the upgrade procedure from version 1.6 to 1.7, refer to the Cisco publication "Upgrading MEC Microcode Version 1.6/2.1 to Version 1.7/2.2" (78-0858). For hardware revision prerequisites associated with this version of microcode, refer to Table 4 at the beginning of the "MEC Microcode Revision" section.

Modifications

- Support has been added for autonomous switching.
- As a performance improvement for AppleTalk fast switching with Software Release 8.3(1), support has been added for handling packets which start on odd-byte boundaries.

Caveats

- MEC microcode version 1.7 does not properly support packets starting on odd-byte boundaries when used with system software versions earlier than 8.2(6). This problem impacts AppleTalk fast switching with some other vendors' software.

MEC Microcode Version 2.1

MEC microcode Version 2.1 was released in February 1991 with the MEC controller type 5.1 card. For hardware revision prerequisites associated with this version of microcode, refer to Table 4 at the beginning of the "MEC Microcode Revision" section.

Modifications

- MEC microcode version 2.1 supports the hardware upgrade from controller type 5.0 to revision 5.1.

Caveats

- MEC Microcode Version 2.1 does not properly support packets starting on odd-byte boundaries. This problem impacts AppleTalk fast switching with some other vendors' software.

MEC Microcode Version 2.2

Released with the 8.3(1) system software in October 1991, MEC microcode version 2.2 requires the replacement of all 13 registered EPROMs. For more information on the upgrade procedure from version 2.1 to 2.2, refer to the Cisco publication *Upgrading MEC Microcode Version 1.6/2.1 to Version 1.7/2.2* (Part Number 78-0858). Refer to Table 4 at the beginning of the "MEC Microcode Revision" section for hardware revision prerequisites associated with this version of microcode.

Modifications

- Support has been added for autonomous switching.
- As a performance improvement for AppleTalk fast switching with Software Release 8.3(1), support has been added for handling packets which start on odd-byte boundaries.

Caveats

- The MEC microcode version 2.2 does not properly support packets that start on odd-byte boundaries when used with system software versions earlier than 8.2(6). This problem impacts AppleTalk fast switching with some other vendors' software.
- In extremely high collision environments, the transmitter may hang. In rare instances, if the output interface queues are full (due to continued extreme collision rates), the interface may not recover, even when issuing a **clear interface** command. This condition requires a system reset.

MEC Microcode Version 2.3

Released with SR 9.0 in April 1992, MEC microcode version 2.3 requires replacement of all 13 registered EPROMs. For more information on the upgrade procedure from version 2.1 to 2.3 and from version 2.2 to version 2.3, refer to *Upgrading MEC Cards from Microcode Version 2.2 to 2.3* (Part Number 78-0934-01). Refer to Table 3 for hardware revision prerequisites associated with this version of microcode.

Modifications

- Fixes the transmitter hang problem that could occur in extremely high collision environments.

Caveats

- The MEC microcode version 2.3 does not properly support packets starting on odd-byte boundaries when used with system software versions earlier than 8.2(6). This problem impacts AppleTalk fast switching with some other vendors' software.

SCI Microcode Revision Summary

This section describes the microcode revisions to the Serial-port Communications interface card (CSC-SCI).

SCI Microcode Version 1.0

This was the first officially released version of SCI microcode. It was released in January 1989.

Caveats

- Serial interface reports CRC errors as encapsulation failures. This problem is identical to the serial line bug in MCI version 1.5.
- Serial performance with X.25 at rates of 2 Mbps or higher can be impacted by very short (2- to 6-byte) packets.

SCI Microcode Version 1.1

Released simultaneously with MCI Version 1.6, this version requires the replacement of two of the registered EPROMs if SCI microcode version 1.0 is currently installed. For more information on the upgrade procedure from version 1.0 to 1.1, refer to the Cisco publication *Upgrading SCI Microcode Version 1.0 to Version 1.1* (Part Number 78-0733).

Modifications

- Serial interface reports of CRC errors as encapsulation failures has been corrected.

Caveats

- Serial performance with X.25 at rates of 2 Mbps or higher can be impacted by very short (2- to 6-byte) packets.

SCI Microcode Version 1.2

Released with 8.3(1) system software, SCI microcode version 1.2 requires the replacement of all 12 registered EPROMs. For more information on the upgrade procedure from version 1.0 to 1.2, refer to the Cisco publication *Upgrading SCI Microcode*