

8.3(9) Caveats

This section describes possibly unexpected behavior by Release 8.3(9). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(9).

AppleTalk

The conversion of special characters to upper case for use in zone name comparisons is incorrect. This may result in incorrect responses to ZIP queries for zone names containing such characters. The workaround 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 non-discovery 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 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; i.e "MktgLaser" appears in "ZoneA" and "ZoneB", although it is defined to reside in "ZoneA" only.

This is known to occur with Apple LaserWriter IIG's and pre-1.5 version Dayna EtherPrints. [CSCdi04951]

In large Appletalk internets, an 8.3 router which is missing a zone name may experience unusually high CPU utilization as displayed by 'show process'. The process AT MAINT will take an increasingly larger interval each time it runs (as shown in the 'usecs' column of the 'show process' display). The workaround is to correct the missing zone condition. This problem is corrected in release 9.0. [CSCdi09487]

AppleTalk packets cannot be routed from X.25 interfaces directly onto either FDDI or Token Ring interfaces. This capability is provided in all router software releases above and including 9.0. [CSCdi09630]

The command **show apple adjacent-routes** in the online help list for **show apple** should really read as **show apple adjacent**. [CSCdi10469]

When an interface is configured for nonextended AppleTalk, it will unexpectedly try to bring itself up after an AppleTalk address is assigned but before a zone is specified. This leads to improper port startup. This can be avoided by specifying the zone first and the AppleTalk address second. [CSCdi11516]

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** (or **write** with no argument), **write erase**, **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" display. This may be worked around by explicitly configuring buffer management parameters using the **buffers command**. [CSCdi04062]

If a router gets its nvram erased or corrupted, and a reload occurs, an administrator must be

present to reconfigure the box. Instead, it would be good to turn on **service-config** automatically in such a circumstance so that the box can restart on its own. [CSCdi07521]

CBUS tokens do not populate the memory location that SNMP looks in for the data. Therefore the token ring upstream neighbor is always reported as 0000.0000.0000 [CSCdi13489]

DECnet

DECnet does not accept HELLOs from nodes with node-address greater than the max-address configuration parameter if the HELLO contains an area address which is different from the router's area address. [CSCdi13136]

IBM Connectivity

When a bridge number of a token ring interface is changed using Lan Network Manager, the output of SHOW SOURCE still displays the old bridge number. [CSCdi14351]

Interfaces and Bridging

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]

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]

When an IP IGRP update is created for a major network which is subnetted and directly connected to the router, but the update in question is being sent through an interface which 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]

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

An exceedingly rare race condition with IGRP can cause the router to reload. IGRP must simultaneously learn a new route while the routing table is being cleared. [CSCdi07276]

No ip broadcast-address should return to default. There is no further information available concerning this problem. [CSCdi07563]

In bridge tables with large numbers of entries or more than one bridge group, dynamic station entries may appear with an S in the Age field. These entries will not be properly aged or relearned, which may result in a station being unreachable from a bridge should the spanning tree reform. These entries may be removed manually using the **no bridge group address MAC-address** command. This action will allow the entry to be relearned. These entries can be removed from the bridge table as a whole only by reconfiguring the affected bridge group. [CSCdi08239]

If secondary addresses are used, in some circumstances IGRP can duplicate network routes in secondary advertisements. Normal operation is unimpaired, but excessive bandwidth is used. This is fixed in a future major release. [CSCdi08511]

When IP extended access lists are used and the extended access list has not yet been defined, some usages result in all packets being denied. Other usages result in all packets being permitted. [CSCdi08718]

If a **neighbor** command is used with IGRP, RIP, or Hello, and the neighbor is not in the major net as the primary address of the outbound interface, the routing update will be sent with an incorrect source address. This can result in incorrect routing at the neighbor. [CSCdi08770]

The **mac-address address** command does not work on serial interfaces, even though a serial interface may want to use this, for example, when a communication server is originating LAT packets on the serial link. [CSCdi09015]

If a very large FDDI SMT frame is received or sent by the router, and **debug fddi-smt** is configured, the debug output for that frame may be corrupted. [CSCdi09114]

IP accounting is not supported for UltraNet interfaces. Incorrect data is entered into the accounting table. The fix is to disable IP accounting on UltraNet interfaces. Future releases will prevent this unsupported configuration from being set up. [CSCdi10595]

Value in "Counters last cleared" field sometimes shows as "*****" after an extended period of time. The same field also sometimes remains stuck at 0:00:00. [CSCdi11305]

A few more statistics show up when you execute a **show controller ethernet interface #** when LANCE is used. [CSCdi13355]

IP Routing Protocols

show traffic will display certain fields as negative numbers once the values wrap into the sign bit. [CSCdi06979]

A race condition in the **show ip cache** command can cause the router to reload. This caveat cannot be completely fixed in 8.2 and 8.3. [CSCdi07900]

The system can refuse to allow the user to remove static ARP entries that were specified by the user, with the error message "Can't unset interface address." The system is wrongly confusing the user supplied ARP entry with the system generated ARP entries for its local network interfaces. The correct behavior is to allow the user to remove any ARP entries they added to the ARP table. This can happen when the user explicitly specifies an ARP entry for the local IP address of an interface on which ARP is not running. [CSCdi08523]

If an interface is configured with the **ip unnumbered** and **no ip split-horizon** commands, no routing updates will be received on that interface. [CSCdi08717]

Static routes that point to destinations reached via a route that has expired are not removed from the routing table. [CSCdi09564]

If a route is known to a network or subnet and a secondary address is configured on a down interface, and the secondary address matches the network or subnet in the routing table, the route will be replaced. The result is a connected route to a down interface. [CSCdi09845]

On routers with multiple interfaces with the same MAC address (for example with DECnet enabled) and multiple connections to the same network segment (to use multiple encapsulations for example), a condition exists where one or more of the interfaces may accept IP packets even though IP is not defined for that interface (i.e no ip address). A workaround for this problem is to explicitly disable the fast-switching cache on the interfaces routing ip with the **no ip route-cache** interface subcommand. [CSCdi10647]

OSPF neighbor config lost when the interface is shutdown. [CSCdi13549]

Configure doesn't save changes in " ip ospf hello-interval" for non-broadcast networks. The change is seen when doing "show ip ospf int s x" but not when doing "write term". [CSCdi15717]

ISO CLNS

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

This problem applies only when doing ISO-IGRP inter-domain routing over links that split horizon is not performed. This includes X.25 PDNs, Frame Relay and SMDS networks. Prefix route advertisements will count to infinity over these networks when a prefix goes unreachable. CSC [CSCdi07379]

ES-IS cache entries for a disabled interface are not flushed when the interface is disabled. This means that packets destined to systems that were formerly reachable through that interface may be lost until the cache entries time out (maximum of five minutes). [CSCdi08490]

CLNS packets that are slow switched will always have their checksums calculated from scratch, even when the incoming packet has checksums turned off. This has no operational impact, other than slowing down packet forwarding and receipt if the original packet did not have checksums enabled. [CSCdi08567]

If there are any CLNS discard routes configured and they are redistributed into ISO-IGRP, they will not be advertised. The workaround is to configure a fictitious static route so it can be redistributed. [CSCdi09917]

After an uptime of nearly 25 days the IS-IS level 2 LSP may stop being sent, causing the IS-IS routing entry to disappear in the neighbour router. This is likely to happen if a router has only one level 2 adjacency. [CSCdi13482]

LAT

A LAT protocol translation session can fail to be destroyed properly under some circumstance when output is still in progress as the connection is closed by the remote LAT host. [CSCdi07506]

Protocol Translation

When "login" is specified as an option in a **translate** command, the user will end up being queried for his password before any telnet option negotiation has been done. Frequently, this will result in echoing the password and double-echoing the username. [CSCdi04686]

TCP/IP Host-Mode Services

[\$\$Ignore\$\$] [CSCdi04220]

A router may experience large processing demands for a TCP connection on closure if the TCP protocol exchange for the close is unduly delayed. This was detected and traced in connection with Cisco's X.25-over-TCP implementation where X.25 caused the connection to linger in a half-closed state. The X.25 behavior was reported and fixed as bug report CSCdi05031. [CSCdi05515]

While routing IP, if two ARP-style interfaces have the same IP address and one of those interfaces is shut down, the wrong MAC address could get entered into the ARP table. The workaround is to remove the duplicate IP address from the shutdown interface with the **no ip address** interface sub-command. [CSCdi07036]

In some netbooting configurations, a client may have multiple interfaces that it could use to traffic data back and forth to the server while it is netbooting. The first thing a client will do if the server is not on the same physical wire as one of its interfaces is broadcast a request for a proxy ARP to get to the server. This is asking a neighbor to help it traffic to the server. Once a neighbor responds, data will be forwarded to the server. In some cases, a second neighbor might step in and tell the client that it will act as the proxy ARP. When this happens, the client gets confused as his original path to the server has now changed. It is more common that two or more parallel IEEE media between the client and its only neighbor will also cause this to happen. This will most likely cause an error similar to the following:

```
Booting gs3-k.91 from 223.255.254.254: !O.OO.O..... [timed out] [CSCdi07727]
```

TN3270

TN3270 sessions may pause indefinitely when attempting to connect to Unisys mainframes. [CSCdi09547]

Uncategorized Items

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 workaround is to manually reduce the interface MTU to 1500. [CSCdi01406]

VINES

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

Traceroute did not function because routers were not sending notification that a packet had been dropped. That notification has been re-instated, and has been modified so that only dropped packets addressed to the ECHO port will elicit a notification. [CSCdi14256]

Wide-Area Networking

When a switch is re-configured to use a different DLCI to reach the same end address, the router doesn't flush the "deleted" map entry and attempt to learn a new mapping. [CSCdi03757]

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]

The frame relay encapsulation code doesn't 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 workaround is to configure the router with the correct DLCIs. [CSCdi05481]

The **x25 pvc bridge number** interface command is not properly stored in the router's configuration memory. [CSCdi06683]

Under unusual circumstances, a RESET of a virtual circuit (VC) may not properly discard all in-transit data. This may cause an additional RESET of the VC to occur. [CSCdi07811]

The Cisco X.25 implementation allows both modulo 8 and modulo 128 virtual circuits to coexist on the same interface; this is nonstandard. [CSCdi07812]

The X.25 idle timer previously applied to SVCs that were switched (via the **x25 route** command) or nonswitched on an interface. Now only nonswitched SVCs are subject to the X.25 idle timer. [CSCdi09927]

Due to a parsing error, the interface subcommand **frame-relay lmi-type ANNEX D** is not accepted. This occurs even when the system is reading a configuration file written by the software, as from non-volatile memory. A workaround is to load a configuration file at boot time containing the alternative form, **frame-relay lmi-type ansi** which is accepted. [CSCdi15175]

XNS, Novell IPX, and Apollo Domain

Adjusting the Novell output-sap-delay to a large number, for instance 200ms, may cause an increase in input queue drops. A workaround for this would be to use a smaller number for the delay, and/or increase the size of the input queue. Novell recommends a SAP inter-packet delay of 55ms. [CSCdi07338]

If a Novell network number is assigned to an interface that is administratively shut down and the router has a valid alternative route to that same network in its routing table, poison SAPs will be routed to that network. A result of this possibly unexpected behavior is that it will sometimes appear that the router is violating split-horizon and sending poison SAPs back out the interface they arrived on. Regular periodic SAP updates do not display this behavior. The workaround is to remove Novell network numbers from interfaces that are administratively shut down. [CSCdi07425]

A race condition in the **show novell cache** command can cause the router to reload. [CSCdi09163]

If a Novell access-list has a host id portion but not a mask the mask portion may take an incorrect value and cause the access-list to fail. A workaround is to always specify a mask even when none

should be required.

```
access-list 800 permit 1.0000.0000.0001 0000.0000.0000 2.0000.0c00.acbe 0000.0000.0000  
instead of
```

```
access-list 800 permit 1.0000.0000.0001 2.0000.0c00.acbe [CSCdi17980]
```

8.3(8) Caveats/8.3(9) Modifications

This section describes possibly unexpected behavior by Release 8.3(8). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(8). For additional caveats applicable to Release 8.3(8), 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(9).

AppleTalk

AppleTalk addresses of the form 0.X, where X is any valid node number, are erroneously entered into the fast-switching cache. This may possibly affect systems with more than one operational nonextended interface. [CSCdi10802]

Zone names that begin with one or more leading blank spaces are not properly stored in the configuration memory. This may lead to zone conflicts when the system is rebooted; the parser will consume all leading white space when parsing the zone name. To prevent such a situation, zone names with leading blank spaces should not be used. The correct system behavior would be to store the first leading blank space as the sequence :20 using the special colon notation. [CSCdi11052]

A ZIP GetMyZone reply is sent in response to a ZIP GetLocalZones request on nonextended interfaces. This is an unexpected response on Macintoshes running AppleTalk v58. The correct behavior is to respond with a GetLocalZones reply. [CSCdi11248]

Basic System Services

An Ethernet frame padded to meet the minimum required length sent from an Ethernet across a serial line to another Ethernet the 6 byte padding is dropped . So it resulted in a runt with 58 byte length seen on the destination Ethernet across the serial line. This caveat occurs only when the system bridges Ethernet frames with length field (802.3, ISO, and SNAP) with more than two interfaces in the same bridge group.

This is configuration sensitive when using 9.1 or 9.0 s/w on one router and 8.3 on the other router. [CSCdi10148]

EXEC and Configuration Parser

The parser sometimes claims that incomplete command names are not unique. [CSCdi10554]

The **enable password** can contain a maximum of 80 characters. An individual had defined a 60 character password, but was only able to enter 49 characters at the enable prompt, and was therefore not able to enter enable mode. [CSCdi10832]

Interfaces and Bridging

When bridging over circuit groups all broadcast traffic was forwarded over 1 of the line in the circuit group, instead of being load-balanced. [CSCdi10071]

IP Routing Protocols

Upon receipt of IP directed broadcast packets, the system erroneously attempts to resolve the directed broadcast address via HP Probe address resolution broadcasts. This occurs if the directed broadcast is destined for a directly connected interface, and that interface is configured for **arp probe**. The system then also correctly forwards the directed broadcast as a data link layer broadcast (if not disabled via the configuration command **no ip directed-broadcast**). The system should be sending the directed broadcast as a (data link layer) broadcast out the directly connected interface, but should not be attempting to perform address resolution on the IP directed broadcast address.

[CSCdi09627]

TN3270

Under certain circumstances, use of TN3270 may cause the Communication Server to hang.

[CSCdi09987]

TN3270 may return modified data fields to the host in the incorrect order. This is primarily manifested in applications complaining of invalid data in fields that do indeed have the correct data.

[CSCdi10344]

Under some circumstances, a terminal server running TN3270 may display the message: %SYS-3-BADPARAM: Function memNSchr. This is cosmetic and can be ignored. [CSCdi10773]

Wide-Area Networking

XNS, Novell IPX, and Apollo Domain

The Cisco IPX **ping** command was limited to a maximum of 1500 bytes. This patch increases the **ping** maximum to 4096 bytes for segments which supports that size. [CSCdi10130]

8.3(7) Caveats/8.3(8) Modifications

This section describes possibly unexpected behavior by Release 8.3(7). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(7). For additional caveats applicable to Release 8.3(7), 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(8).

AppleTalk

Serial interfaces configured with discovery mode never become operational. [CSCdi09532]

This does not affect the user . Software clean up. [CSCdi09787]

AppleTalk Packets cannot be fast switched between MEC Ethernet controllers and HSSI serial controllers when the Ethernet interface is running Phase I AppleTalk, and the HSSI interface is running Phase II AppleTalk. [CSCdi09818]

Basic System Services

A terminal line configured for flow control will not successfully time out (due to a "session-timeout" configuration) if XOFF is selected for the line at the time of the timeout. [CSCdi09310]

DECnet

Under some condtions the **show decnet route** command may cause the router to reload.

This has been fixed in 9.x . [CSCdi07664]

DECnet should look at the MAX AREA parameter and not advertise reachability to any areas greater than this parameter. Likewise, it should not advertise reachability to a node that is greater than the MAX NODE parameter. It should also not accept hellos from such nodes. [CSCdi09716]

EXEC and Configuration Parser

When netbooting the operating system and config file, the **ip split-horizon** statement disappears. This is because **ip split-horizon** is located after the **encapsulation frame-relay** statement, [CSCdi08462]

IBM Connectivity

The problem was simply that the system did not learn the Burned In Address of the token ring adapter card until after the interface inserted onto the ring. If the interface was shutdown when the router was booted and the router was configured for bridging, the virtual ring address would be configured with the address 4000.0000.0000 ... clearly invalid.

This happened because the virtual ring uses the Burned In Address of the adapter, logically 'OR'ed with the '4' to obtain it's unique address, which is a problem in the above scenerio. [CSCdi07105]

Interfaces and Bridging

The system normally disallows multiple interfaces to be configured with IP addresses on the same subnet. Such IP address overlap should be allowed when it occurs between a transmit only interface and its associated receive interface, as designated by the **transmit-interface** interface subcommand. [CSCdi09300]

Due to interactions between the bridging code and driver code, the spanning-tree state is handled incorrectly. In pre-9.1, this shows up most readily on serial lines. If a serial line is shut and then no-shut, the port goes into blocking and then stays there. Similarly, if you have an Ethernet port and you pull the cable out, the port will go down. But if you wait for a minute or so (give the spanning-tree protocol time to recompute) and then plug the cable back in, you will see the port go into Forwarding immediately. This can cause temporary network meltdowns. [CSCdi09535]

IP Routing Protocols

Source-routed IP packets that are supposed to be discarded by the system are sometimes not. Such packets are being packet switched when the local system does not appear as the next hop in the source route. These packets should never be packet switched when the user has entered the **no ip source-route** configuration command. This unexpected behavior can pose a security problem for sites that use this command to restrict access. [CSCdi09517]

When initiating a TFTP read request, the system can generate TFTP packets with invalid UDP checksums. This only happens when the request is transmitted out an unnumbered interface. If the TFTP server has UDP checksumming enabled, TFTP read requests via the unnumbered interface will fail. Turning off UDP checksumming at the TFTP server or restricting TFTP reads to numbered interfaces avoids this problem. [CSCdi09577]

If the system tries to clear pending output on an inbound telnet connection that is in the process of closing, it is possible that the system may reload in some cases. [CSCdi10087]

ISO CLNS

A redirect sent out over an X25 interface does not get encapsulated and CLNS returns a failure. [CSCdi04417]

CLNS fast switching over a serial interface with HDLC encapsulation falls back to slow switching. [CSCdi09172]

When CLNS receives a packet that needs to be fragmented, but the 'segmentation permitted' bit in the packet is off, it should send back an error packet (ERPDU) indicating this situation. [CSCdi09413]

Local Services

If an attempt is made to either write a read-only object or read a write-only object, the wrong error code is returned. [CSCdi09714]

TCP/IP Host-Mode Services

When a TCP connection has a closed window, packets containing valid ACKs 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 especially a problem in the initial stages of a connection, when we send the SYN-ACK with a 0 window. If the ACK to our SYN contains data also, we will not process that ACK, and the connection never gets to ESTABLISHED state. [CSCdi05962]

TN3270

Communication servers under certain circumstances may drop TN3270 connections and eventually reload. [CSCdi09197]

VINES

This problem only occurs when a client is initially powered on, and the first login attempt results

in a forced password change. The user will not be able to change his password and will not be able to log in. The workaround is to have another user log in and log out at that client, then the affected user will be able to log in and change his password. [CSCdi09467]

XNS, Novell IPX, and Apollo Domain

This patch fixes an interoperability issue between the cisco Novell IPX routing fast switching implementation between release 9.1 and 8.3 or 9.0 software releases before either 8.3(7.2) or 9.0(5.1).

Note: 8.2 has the same problem as 8.3 and 9.0, but no fix will be generated for that release.

In the 9.1 release fast switching was enhanced to allow communication to FDDI and Serial end hosts. Before 9.1, the router did not fast switch Novell frames to a Novell FDDI end host, but would always process switch them instead, so communication between actual end hosts was always effective.

The older release Novell fast switching code wrote packets sent to next-hop remote routers on FDDI and Serial links with extra padding bytes, in such a way that it guaranteed that Novell frames output on Ethernet interfaces by the remote router would always have at least 64 bytes of data (plus 4 bytes of checksum).

The 9.1 fast switching code generates correctly formatted frames on FDDI and Serial interfaces. However, the older releases of software will misinterpret these frames when fast switching, and generate output frames on Ethernet that, while valid frames, are smaller than 64 bytes.

Some versions of PC Ethernet drivers seem to require a 64 byte minimum frame size (plus 4 bytes of checksum). As such, if they are in a setting where a 9.1 and previous release router are running in series, they will not be able to accept the smaller frames.

This patch allows 8.3 and 9.0 to operate correctly with both correctly formatted input frames from release 9.1, or incorrectly formatted input frames from previous releases, on both FDDI or Serial.

Note 1: The problem in 8.3 and 9.0 can be worked around by turning off fast switching on the 9.1 router's FDDI or Serial interface.

Note 2: This patch will also fix problems where 8.3 or 9.0 cannot correctly forward frames sent by a PC FDDI end host onto an Ethernet. [CSCdi09754]

Novell, XNS, and Apollo maximum-path 0 is accepted and displayed by the system, but the default maximum-paths is 1. If a user types a maximum path of 0, make this return to the default setting of 1. [CSCdi09955]

8.3(6) Caveats/8.3(7) Modifications

This section describes possibly unexpected behavior by Release 8.3(6). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(6). For additional caveats applicable to Release 8.3(6), 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(7).

AppleTalk

An error has been found in the AppleTalk fastswitching functionality which results in invalid AppleTalk packets being generated in the case of a packet being received on a cBus FDDI interface running extended AppleTalk and being destined for a non-extended Ethernet MEC interface.

This error can be worked around by disabling the AppleTalk route cache on either the MEC Ethernet interface or the FDDI interface. [CSCdi08211]

Basic System Services

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]

Output of "show process" outputs octal numbers, even if service decimal-tty is set (the default on

the protocol translator and 500-CS). [CSCdi08240]

entering multiple **logging buffered** commands without an intervening **no logging buffered** command can cause meaningless output to be included in the output of the **show logging** command. [CSCdi08459]

System images from the 8.3, 9.0, 9.1, and 9.14 releases could not be successfully netbooted on IGS boxes with 8.2 EPROMs. The ROM monitor in the 8.2 EPROMs did not support some functions that the newer releases use. The system image should protect itself by error checking the return code from all ROM monitor calls. [CSCdi08521]

An "event-dismiss" error message can be encountered when debug output is being output on the console while running a bootstrap system image; for example, igs-rxboot, xx-rxboot, csc3-boot, and so on.

```
(boot)ROUTER#debug tokenring %SYS-2-INTSCHED: event dismiss at level 4 -Process = "Exec", level= 4, pid= 11 -Traceback= A87C A8D6 1418C 9422 9EB2 15FA 304D8 70DEC %SYS-2-INTSCHED: event dismiss at level 4 -Process = "Exec", level= 4, pid= 11 -Traceback= A87C A8D6 1418C 9422 9 [CSCdi08533]
```

DECnet

DECnet address translation fails on IGS platform routers in the cases where both interfaces are not fast switched and one of the interfaces is capable of being fast switched. The workaround is to configure both interfaces for DECnet fast switching. Since this is not possible for all interfaces and encapsulations, such as Token Ring, X.25, and Frame Relay interfaces, some configurations cannot support ATG on IGS platform routers. [CSCdi07652]

EXEC and Configuration Parser

The **debug ?** command does not show serial options if only serial interface type is HSSI. [CSCdi07674]

IBM Connectivity

When routing IP in conjunction with bridging, HP Probe packets will be bridged rather than received by the router. Cisco Systems expects to resolve this problem in a future release. [CSCdi07039]

Interfaces and Bridging

Output drops double counted when output holdq is full. There is no further information available concerning this problem. [CSCdi07195]

The **debug broadcast** command does not work on FDDI broadcast packets unless the hidden **debug fddi-event** command is enabled. [CSCdi08137]

When removing a remote Source Route Bridge peer the unit may crash. [CSCdi08152]

Static IP routes can fail to be removed from the routing table when an unnumbered interface goes down. This can result in host or network routes pointing to a down interface to continue to be advertised via routing protocols. When the interface goes down, the router should remove the static route from the routing table for as long as the interface remains down. Until fixed, static IP routes should not be used with unnumbered interfaces. [CSCdi08180]

When reconfiguring the priority on an interface used for transparent bridging, we delay reconfiguring the port until we receive the following BPDU message. This can cause a significant delay in the convergence of the spanning tree. This caveat is present in all previous releases. The port is now reconfigured as soon as the configuration command is executed. [CSCdi08296]

When use process PCM and dual-homing connection, if the user issues a **cmt disconnect** command to a standby port the CPU utilization will go very high. [CSCdi08427]

When a communication server line is configured for modem control and with a session time-out,

the session time-out will not be honored if the line is running in SLIP mode. [CSCdi08562]

On the IGS, Cisco 3000, and Cisco 4000 serial network interfaces, we check the status of DCD before we assert DTR. This means that loopback interfaces that connect our output DTR signal to our input DCD signal will not work, because DCD will never be asserted. We should assert DTR before checking for DCD. [CSCdi08612]

When an IP packet with IP options is received on a fast-switching interface, the system sometimes fails to decrement the IP TTL before forwarding the packet. This is most noticeable when a "traceroute" program is being used with source-routing options, and causes the system to sometimes fail to show up as an intermediate hop in the traceroute output. [CSCdi08699]

If an unnumbered interface is shut down, it is periodically removed from the IP routing table. This causes unnecessary routing table activity and can introduce other detrimental side effects. [CSCdi08715]

The **clear counter** [*type unit*] command always clears the counters regardless of the user's response to confirmation. [CSCdi08774]

MCI/SCI will become unusable when the MTU is 4K or above because there is only one buffer for the receive side. We recommend that MTU should be less than 4.5K. [CSCdi08842]

When using a protocol translator or communication server without IP routing enabled, ARPs for IP aliases or addresses associated with translate commands may not be answered correctly. As a workaround, turn on IP routing to restore the expected behavior. [CSCdi08981]

When a system is attempting to TFTP boot, it may not know a route to the TFTP server. If the system has multiple interfaces by which it might contact the TFTP server, it can fail to continue to use the interface on which the TFTP transfer was just established. The result is that the TFTP boot attempt fails. The system should remember by means of its ARP table the interface to use to reach the TFTP server. Configuring the system's NVRAM so that it can only reach the server by one interface at boot time avoids this problem. [CSCdi09068]

A Cisco router sends VINES routing updates as spanning tree explorers whereas a VINES server sends routing updates as all-routes explorers. The Cisco implementation provides lower explorer impact upon the network, whereas the Banyan implementation finds the shortest path between any two nodes. The fix for this behavior allows choosing between spanning tree explorers and all-routes explorers on a per protocol basis. This is done via an extension to the **multiring** command. The new command syntax is [**no**] **multiring** {*protocol* | **all**} [**all-routes** | **spanning**]. The trailing **all-routes** and **spanning** keywords specify the explorer type to be used. The default is to use spanning tree explorers. [CSCdi09091]

IP Routing Protocols

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

The system does not properly process RARP response packets received where these packets are responses for requests not initiated by the system. The system allows such packets to remain in the input queue, resulting in two user visible problems. First, the network interface input queue can fill up with RARP response packets, causing all subsequent packets destined for the system to be dropped. Second, the system fails to bridge these RARP response packets. The correct behavior is to bridge such packets in the case where the system is configured to bridge RARP packets; otherwise, it should ignore these packets. [CSCdi08651]

The **distribute-list** command sometimes makes access list changes even when a parsing error is detected and an error message is printed. The software continues processing this command even though an error has been detected. Because of this aspect of the implementation, the system will

treat a **distribute-list** command that specifies a nonexistent interface as if no interface has been specified, thus unexpectedly applying the access list to all interfaces. If the user receives parser errors in response to their **distribute-list** configuration commands, it is recommended that they verify that the system has not wrongly interpreted their commands by examining the distribute-list commands reported by **write terminal**. [CSCdi08668]

ISO CLNS

The "better SNPA" field in an ES-IS redirect is always sent in native bit order. This can cause OSI End Systems on Token Ring networks to be unable to reach some destinations when more than one router is present on the token ring. [CSCdi07200]

The encapsulation type for CLNS is sometimes displayed incorrectly when a **show clns interface** command is entered. This is a cosmetic defect only. [CSCdi08467]

CLNS fast switching does not properly fragment packets. Packets received on FDDI that are larger than 1497 octets will not be forwarded properly over serial and 802.3 interfaces. This is not typically a problem, since CLNS packets are seldom this large. The workaround is to disable CLNS fast switching on the FDDI interface with the **no clns route-cache** command. [CSCdi08494]

If the CLNS trace facility is used to trace a path that goes through another Cisco router on the same LAN, the second of the three trace packets may not work. This has no operational impact, other than causing a 3-second delay in the execution of the trace. [CSCdi08653]

CLNP packets received by a router with a lifetime field of zero will be forwarded (with a lifetime of 255) if slow-switched. This has no operational impact whatsoever unless a host is emitting packets with a lifetime of zero. [CSCdi08654]

When an invalid ER PDU is received, we should just discard it without sending an ER PDU in response. [CSCdi09139]

LAT

When a passthrough connection is made to another LAT system over an existing LAT connection, the break key does not return the terminal server to the correct mode. [CSCdi07815]

Under certain conditions, the LAT disconnect sequence may cause the Communication server to reload. [CSCdi08636]

Local Services

Any attempt to query an unimplemented SNMP MIB variable will cause the system to return the snmpEnableAuthenTraps variable. The correct behavior is to indicate that the variable requested is not available, and this will be corrected in a future release. [CSCdi04806]

A box with TR crashed with the following:

```
IP-3-Desthost:src=200.2.3.1 dst=0.0.0.0 Null desthost Process="SNMP Server",level=0,pid=28  
Traceback=23628 23364 2500e 26a14 269ae 26c00 391da 81bbc [CSCdi05629]
```

If **enable use-tacacs** is configured on either a communication server or protocol translator without defining a **tacacs-server host**, then any username/password combination will allow any user to enable. [CSCdi08070]

On routers without NVRAM, part of the sequence used to determine IP addresses is to send a BOOTP request. The replies to these requests are being ignored. [CSCdi08893]

TCP/IP Host-Mode Services

TCP connections can exhibit long pauses in data delivery if the cisco is attempting to send data faster than the foreign host can use that data. This happens most often in cases of protocol translation, sdlc tunneling, remote source route bridging, and X.25 switching. TCP connections can exhibit long pauses in data delivery if the cisco is attempting to send data faster than the foreign host can use that data. This happens most often in cases of protocol translation, sdlc tunneling, remote

source route bridging, and X.25 switching. [CSCdi07964]

TN3270

Keymaps are not currently parsed correctly. Each keymap consists of the name of the keymap, the terminal types to which it applies, and the various mappings. When parsing the terminal types, only the first one is read correctly. The result is that the keymap will only be selected when the user's terminal type matches either the name of the keymap or the first terminal type in the keymap. This will be fixed by changing the software to correctly parse the terminal types in the keymap. [CSCdi05677]

The login-string configuration command is not correctly implemented for tn3270 connections. As currently implemented, it merely sends the ASCII text of the login-string to the host at the other end of the connection. This is fine for Telnet and Rlogin connections, but for tn3270 connections, the login-string must be passed through the tn3270 input path.

The problem will be fixed by passing the login-string through the tn3270 input path on tn3270 connections. Additionally, two special escape characters have been added, %t for tab, and %m for carriage return. In order to place a tab in a login-string, one will enter %t. Likewise, one will use %m at the end of the login-string to achieve a carriage return, as normal telnet processesing would send an undesirable line feed after the carriage return. [CSCdi08252]

Clear to end of line is currently done by writing spaces. This is very slow and can be painful on low-speed dialup lines.

It will be fixed by using two attributes in the ttycap, ms: and cx:. If both attributes are in the terminal's ttycap cisco's tn3270 implentation will use the clear to end of line command rather than sending spaces to the terminal. This will be the default behavior. Note that this may not be appropriate when a terminal is in underline mode. Removing the cx: attribute from the termcap will cause cisco's tn3270 to clear to the end of line by sending spaces. [CSCdi08441]

Terminal Service

When tn3270 has a buffer of data to send which is exactly the same size as the packet that it is sending it in, the packet is sent without the TCP PUSH flag set. Some host implementations will not act on the data unless the TCP PUSH is set. Connections to these hosts can pause for the session timeout period. This will be fixed by having all tn3270 packets sent with the push flag set. When TN3270 has a buffer of data to send that is exactly the same size as the packet that it is sending it in, the packet is sent without the TCP PUSH flag set. Some host implementations will not act on the data unless the TCP PUSH is set. Connections to these hosts can pause for the session timeout period. This will be fixed by having all TN3270 packets sent with the push flag set. [CSCdi08034]

If a line is configured with **session-timeout n output**, the "output" part of the command will remain in effect even if a new **session-timeout n** command is given (without "output" specified). A workaround is to turn off the "output" part explicitly with a **no session-timeout 0 ouput** command. If a line is configured with the **session-timeout n output** command, the **output** option will remain in effect even if a new **session-timeout n** command is given (without **output** specified). A workaround is to turn off the **output** option explicitly with a **no session-timeout 0 ouput** command. [CSCdi08625]

VINES

A recent Vines bug is causing Vines clients to send broadcast streettalk packets. Because the Cisco router floods streettalk broadcasts, this can cause congestion in wide area links. The change to the router code is to only flood streettalk broadcasts sent from a server. [CSCdi08277]

If a VINES SPP packet is addressed directly to a router, it will discard the packet twice producing a "Buffer in list" error message. This error is very unlikely, and is also harmless. [CSCdi08362]

The Cisco router now accepts and process VINES redirects from other servers. Prior to this fix, redirect messages were ignored. This patch also fixes some minor problems generating redirect messages. [CSCdi09088]

A Cisco router may occasionally send an ICP error message with an error code of zero. Receipt of this message can cause a Banyan server to drop some or all communications links passing through the router. [CSCdi09175]

If a station is removed from an interface that uses one type of encapsulation and is added to another interface that uses a different encapsulation before the neighbor entry expires, communication to the station will never be reestablished. [CSCdi09294]

There is a condition where some serverless networks will have extreme difficulty logging in to any server. This is caused by a packet sent by the router not being understood by a VINES server. The workaround to this problem is to shorten the name of the Cisco router to be 15 characters or fewer. [CSCdi09372]

XNS, Novell IPX, and Apollo Domain

The **ping** command will display incorrect round-trip times for 32-, 33-, or 34-byte Novell IPX or XNS packets. Use larger sizes when sending IPX or XNS echoes from the router to obtain more accurate round-trip times. [CSCdi07529]

On media other than 802.x, the **show xns int** command will display the wrong encapsulation type if the default encapsulation has been changed. For example, on an SMDS interface **show xns int** will display "XNS encapsulation is HDLC." We should only display XNS encapsulation types for 802.x media. [CSCdi07929]

When a Cisco unit has a large number of the same type of interfaces, the **show novell cache** or **show xns cache** commands will display the interface limited to nine characters, which allows only Ethernet1 to be displayed when it is in fact Ethernet11. The initial 9.1 release changed this to ten characters, which corrects Ethernet names, but Token Ring will have a similar problem unless the length is eleven characters. [CSCdi08236]

When a Cisco router generates an XNS error response packet, it is sent out with a source address equal to the original source of the packet that caused the error response. The source address should be that of the router itself. [CSCdi08377]

In certain topologies, fast switch looping of (Novell) multicast packets can occur when received on an interface that is active, but not configured for Novell. This is now corrected. [CSCdi08722]

Certain Novell packets may be received and processed by the local interface when they have been sent by a misconfigured client, server, or router. For example, a SAP Get Nearest File Server packet sent on network 0xA1 from a host whose network number has been misconfigured as 0xA2. These misconfigured packets should be ignored and counted as bad packets. In the Show Novell Traffic display the packets pitched counter should be incremented when we receive one of these packets. [CSCdi09178]

XRemote

Xremote debugging messages may appear even if debugging is disabled. Also, some of the debug messages are incorrect. [CSCdi08259]

8.3(5) Caveats/8.3(6) Modifications

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). For additional caveats applicable to Release 8.3(5), 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(6).

AppleTalk

In 8.3(3), and 9.0(1), a non-extended 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]

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

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

This message may 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]' function.

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

This problem manifests itself in two observable ways in the Appletalk component of the router software. The first is that once the router has been up for more than 24 and one-half days, clearing, resetting or reconfiguring an AppleTalk interface will cause the interface in question to attain a status of "Restart port pending" which will not change, no matter how the interface is configured or cleared.

The seconds manifestation of this problem is cosmetic in nature. Times which are expressed as an interval of time, particularly in the output of the command **show appletalk neighbor will show neighbor "up times" of "never" after the router has been up for 24 and one-half days or longer.**

The only workaround for this problem is to reload the router every three weeks. [CSCdi06929]

The **show appletalk command is does not accept the "talk" portion of the keyword "appletalk." This is not a serious problem, as it is easily worked around by using the keyword "apple" in exec commands.** [CSCdi06988]

Basic System Services

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 workaround is to reload the router after a reconfiguration that has invalidated the address the router was using to source syslog messages. [CSCdi04906]

If a user connected via TELNET to a router leaves the **show process** display at the "--more--" prompt, and the virtual terminal session idle timer expires, a system reload may occur. [CSCdi05633]

The **srp output-address-list list** command is mistakenly applied to the source MAC address and not to the destination MAC address. [CSCdi06347]

get_pak_size missing support for huge buffers There is no further information available concerning this problem. [CSCdi07091]

On the 8.3, 9.0, and 9.1 releases, the Ethernet and serial interfaces on the IGS use larger buffers than is required if a Token Ring is configured in the system. This wastes shared (buffer) memory. On the 9.1 release, the Cisco 4000 also uses larger buffers than is required if a Token Ring Network Interface Module (NIM) is configured in the system. [CSCdi07369]

Configuring a location string longer than 69 characters can cause the system to reload. After configuring, the system prints out a message saying that the system was configured from and gives the location. If the location is greater than 69 characters in length, it can cause a system reload. The

correct behavior would be to truncate the location string and will be implemented in a future release. [CSCdi07834]

DECnet

A Cisco router running DECnet IV with conversion enabled does not ignore Phase IV hellos sent from a Phase V router. As such, the router will try to set up a Phase IV adjacency with the Phase V router, while the Phase V router ignores the Phase IV hellos that the Cisco router sends. In effect, this causes the adjacency to be one-way, and will show up in the Cisco router's DECnet IV routing table as initializing. [CSCdi07393]

We were not ignoring IV hellos sent by a router running V (cisco or DEC). This created problems when a DEC V router was adjacent to a cisco router, because we were accepting the DEC's IV hellos while the DEC router was rejecting our IV hellos. The result was a half-baked IV adjacency. Bug 7393 added code to ignore IV hellos from a V router when we were running OSI, IV and had conversion turned on. This fixed the original problem, but it resulted in an interesting side effect: we were now refusing IV hellos from cisco routers as well and this caused a DECnet IV network to get partitioned when there were cisco routers running with IV, OSI and conversion on. [CSCdi08164]

EXEC and Configuration Parser

The **no service ipname** command is never written to nonvolatile memory on STS-10 or STS-10x terminal servers. As a result, IEN116 name server lookup will be re-enabled when a system is re-loaded. [CSCdi05273]

The "setup" command does not allow CLNS station IDs containing a zero to be entered if an ID other than the default was desired. Possible workarounds include using the default station ID supplied, or using a station ID that does not contain a zero. [CSCdi06665]

IBM Connectivity

In early versions of the bridging software IEEE BPDUs weren't always well formed. That is TCN BPDUs would not get transmitted properly (like not at all). [CSCdi05981]

During process-level bridging, the non-flood bridge forwarding code does not check to make sure that it does not output a packet on the interface upon which it arrived. The behavior has been present in all versions of the router supporting process-level bridging.

Normal transparent bridging does not notice this, as it runs fast switched and the the check is correctly applied in the fast switching code. However, bridging that runs at the process level (SR/TLB, bridging with Priority Output, and bridging over X.25 or Frame Relay) runs into this problem. Symptoms of this problem are seen in packets that are duplicated on the receiving interface. The correct behavior is that packets should not be retransmitted on receiving interface.

The impact is on certain protocols that are sensitive to packet duplication and that may not function properly. Process-level bridging performance will degrade. There are no known workarounds. Cisco expects to resolve this behavior in a future release of 8.3. [CSCdi06609]

Router issues a %SYS-2-INTSCHED message and traceback when operating with **debug rif** enabled. The behavior has been present in all versions of the code supporting process-level bridging. After the command has been issued, the router may begin to display the message. The length of time depends upon how much traffic is presented to the router. Higher levels of traffic cause the problem to appear sooner. Once the condition has been triggered, the router continually sends error message and traceback information.

The impact is a potential performance for process level activities. The workaround is to not use the **debug rif** command. The behavior has been present in all versions of the router supporting rif caching. [CSCdi06634]

It is possible for a RIF entry to be updated by a received frame at the same time it is being used to generate a frame. In this case there is a possibility that a frame with a circular RIF will be generated. [CSCdi06673]

When doing pure bridging some forms of communication with the router/bridge using IP wouldn't work correctly. [CSCdi06687]

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]

Shutting down interfaces that are members of a bridge group and are in a forwarding state, and then bringing them back up may result in forwarding loops in the spanning tree. These loops will manifest themselves in saturated traffic levels on the interfaces and excessive CPU utilization.

Systems in this situation typically must be reloaded to recover normal operation. [CSCdi05010]

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 a CSC/4 processor with an ethernet MCI, keepalives won't 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 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 then "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 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** followed by the **keepalive n**.

Note: The only action above 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. [CSCdi05172]

Configuring "ip route 0.0.0.0 Null 0" will result in the route showing up multiple times in the routing table. [CSCdi05754]

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

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

show rif could cause a router to crash when the rif cache is getting updated. This fix resolves the problem. [CSCdi06016]

Router has problems netbooting when there are multiple paths to the remote tftp server. [CSCdi06088]

When bridging is enabled, SNAP encapsulated packets will be bridged even when the relevant routing protocols are enabled. Bridge filters may be used to constrain the propagation of this traffic by SAP, but no solution is available for receiving or routing these packets. [CSCdi06109]

RIP, HELLO, and IGRP advertisements being broadcast on unnumbered serial links will not advertise the major network number of the associated numbered interface. [CSCdi06205]

CSCdi05488 caused the router to not send complete RIP updates to explicitly configured RIP neighbors. [CSCdi06285]

The router software decrements the reset counter after some internally generated interface resets,

e.g. after the "mac-address" command has been issued. There is no check to see if the reset counter is zero before decrementing it, thus it is possible to decrement the counter to a negative value. Because the value is always displayed as an unsigned positive number, it shows up as a number near 4294967295. [CSCdi06490]

It is possible for the router to reload in the **show controller token** command. This can only happen if a CSC-R16 or CSC-R16M token ring card is in the reset state. [CSCdi06681]

In a Spanning Tree environment for bridging some transitions from Forwarding to Blocking wouldn't work correctly. This could result in inconsistent Spanning Tree state with possible network outages resulting. [CSCdi06689]

If split horizon is disabled and the interface is numbered, the router should not accept IGRP, RIP, or HELLO routing updates from other hosts on that interface but not on the subnets configured on that interface. [CSCdi06885]

If a SMT frame comes in on the FDDI the wrong thing happened and we would lose buffers. [CSCdi07080]

Multicast Fddi packets that did not have a UI (0x03) control field would not get bridged at all. [CSCdi07107]

In a bridge environment ARP entries can be heard for a given node on either a FDDI or an Ethernet. If the node is on FDDI we should keep it there but due to a bug we will hear it on Ethernet later and force it to change which causes communications to not take place. [CSCdi07139]

When configured to encapsulate vines packets with a snap header, the router currently uses the header AAAA.0300.0000.0BAD. This fix changes the code to use the proper header of AAAA.0300.0000.80C4. [CSCdi07196]

In a pure bridged environment (ie IP is being bridged rather than routed), under different topologies other nodes would sometimes not be able to communicate directly to the Cisco Router/Bridge. This includes SNMP and Telnet traffic. This makes the Router/Bridge effectively unmanageable. [CSCdi07417]

When routing IP in conjunction with transparent bridging, 802.3 SNAP encapsulated IP will be bridged rather than routed. Cisco Systems expects to resolve this problem in a future release. [CSCdi07495]

In a bridged environment there were a number of bugs that would cause various failures. This included not garbage collecting bridge table entries at the proper time as well as some corner cases in the Spanning Tree transitions. [CSCdi07532]

In very rare circumstances, EGP can cause a router to reload if another process attempts to clear the IP routing table while an EGP update is being processed. [CSCdi07587]

When there is a single fiber break or the neighbor station sends constant halt line state(HLS), system CPU utilization will reach 100%. [CSCdi07682]

When the Cisco router receives a IEEE 802.2 TEST and XID frame that contains both a RIF field which indicates that the frame should traverse the Cisco router, and a destination address which indicates the frame should terminate at the Cisco router, the Cisco router chooses to terminate the frame and reply to it, if needed. This is not in compliance with a strict definition of source-route bridging. This is a minor problem that has little, if any, actual functional impact in most source-bridged networks. [CSCdi07722]

If extended access lists are used on an MCI, SCI or cBus interface, the IP route cache is enabled, and also the 'established' keyword is used, it can be improperly evaluated. This can permit packets which should be filtered and exclude packets which should be permitted. This behavior was first introduced in 8.2. [CSCdi07901]

A bridge configured with **no bridge acquire** will continue to flood and forward packets for other than statically configured MAC addresses. In some cases, bridge filters may be used instead to achieve the desired pattern of traffic containment. [CSCdi07934]

Regarding multibus timeouts and RESETFAIL errors:

Please note the linkage between the following system versions and the sbemon & strmon token ring firmware versions:

```
FIRMWARESYSTEM 8.3SYSTEM 9.0SYSTEM 9.1 =====  
===== sbemon 3.28.3(5.14)9.0(3.1)9.1(1.4) strmon 1.2 N/A9.0(3)9.1(1.4)
```

It is the firmware that is linked to the system versions and will cause a crash if earlier systems are used.

CUSTOMER ENGINEERING:

The older versions of the firmware will work as before with all systems. This means that if a customer has token rings with older firmware that are working fine then there is no need to upgrade them.

On the otherhand, if the system is crashing with a multi-bus timeout (the stack trace should indicate the crash happened in madge_input), or the token ring adapter board dies with a RESETFAIL error from the system, then a firmware upgrade with the proper system software will solve the problem. This is a valid solution as of November, 1992, prior to the above system and firmware releases.

MANUFACTURING:

As I understand it, manufacturing does not have a process for issuing different versions of firmware and they prefer to use the latest firmware as their "default" image, but this would cause the older systems to crash.

A partial solution is to release the new firmware along with, or after, the above system versions. The problem, then, is if a customer buys a new token ring adapter (or upgrades an existing one): they **must** also have the right system. Currently, the way we deal with this problem of incompatible hardware is to wait until the bug generates a call from the field and then we check the version numbers, recommending the proper software upgrade. This is how we do it with the cbus microcode. With this in mind, strmon 1.2 and sbemon 3.2 probably won't be released until January, 1993. [CSCdi08087]

When the system is bridging IP, ARPs originated by the system cause an error message to be generated. This behavior is seen only with packets originated by the system and impacts the use of IP for management of a bridge with a frame relay interface. [CSCdi08293]

Under certain circumstances a pure IP bridge (**no ip routing**) wouldn't be able to communicate with other IP hosts in the presence of topology changes. [CSCdi08349]

IP Routing Protocols

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

ISO CLNS

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

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

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

In some CLNS displays, when an X.121 address is displayed, an 8 digit is printed as 0 and a 9 digit is printed as 1. [CSCdi06308]

When an NSAP address with length of 0 is present in a CLNS packet, the fast switching routines corrupt memory and causes the system to reload. [CSCdi06370]

When a CLNS area is deleted, the process associated with the area's domain is deleted, even if other areas exist in the domain. In effect, this will leave orphan areas. [CSCdi06666]

This fix is an addition to bug 6666, and is specific to 8.3 only. [CSCdi07104]

The system does not properly fragment CLNP packets in some cases. If the packet length is slightly larger than the MTU of the outgoing subnetwork, the packet may either be sent as-is (oversized), or it may have a short final fragment (the ISO 8473 standard requires all fragments to carry at least eight octets of data).

This may cause packets to be undeliverable if the receiving End System enforces the final fragment size requirement, or if the packet is sent with a size greater than the subnetwork MTU. [CSCdi07646]

LAT

LAT status returns "already queued" on first queue attempt There is no further information available concerning this problem. [CSCdi07893]

Local Services

sysLocation is read-only. As a workaround, the location can be set with the **snmp-server location** configuration command. [CSCdi07909]

SNMP GetNext will return incorrect responses for certain queries. [CSCdi08044]

TCP/IP Host-Mode Services

If an interface is shut down and assigned an IP address, then the router should ignore that interface when trying to determine if it is on the same subnet as various other IP addresses. This affects various calculations, notably BGP NEXT_HOP calculations. [CSCdi05356]

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]

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

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 non broadcast media such as X.25 and some frame relay networks. [CSCdi06001]

In some cases we are sending tftp ACK responses after an out of order packet has been received by a client while netbooting. If the server is busy, this is quite a possible event. Sending a second ACK response causes the client and server to get into an argument over what packet to send, and in many topologies it will fail. Common cases look like:

```
!!!!!!O.....[timeout]
```

```
!!!!!!OOOOOOOOOO!OOOOOOOOOO!OOOOOOOOOO!OOOO...[timeout]
```

```
!!!!!!!O..... [timeout] [CSCdi06319]
```

Terminal Service

Under certain circumstances, a reload may occur when switching between multiple TCP connections or using the state machine feature of the communications server. [CSCdi06884]

Uncategorized Items

cisco's implementation of IPTalk is intended to allow Unix(tm) hosts running CAP (Columbia AppleTalk Package) using the non-native AppleTalk encapsulation (ie, AppleTalk encapsulation inside IP datagrams) to communicate with an existing AppleTalk network. The cisco implementation

of IPtalk does not currently provide for router-to-router IP encapsulation tunnels. [CSCdi05452] 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]

VINES

Problem Statement: 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.

Resolution: When a server discovery broadcasts is forwarded onto the network containing the nearest server, it will be forwarded as a MAC layer broadcast. This means that all servers on that physical network will see and respond to this frame, instead of one single server.

There is also a change to the output of "show vines route" so you can easily see which vines server is considered the 'nearest' vines server. The new output is:

4 routes, next update 77 seconds Codes: R - RTP derived, C - connected, S - static

RN Net 0027AF9A [2] via 0027AF9A:1, 10 sec, 0 uses, Ethernet0 C Net 30004355 is this router's network, 0 uses R Net 002ABFAA [2] via 002ABFAA:1, 10 sec, 0 uses, Ethernet0 R Net 3000FB06 [1] via 3000FB06:1, 8 sec, 0 uses, Fddi0

were the letter 'N' indicates that this server is the nearest server, and it is on the local network. The letter 'n' is used to indicate that this server is considered the nearest server, but it is not on the local network. [CSCdi02868]

It is possible, but not probable, that you can crash the router while running the command "show vines route". If you issue this command and let the display sit at the "--More--" prompt until the last route displayed expires from the routing table, the router will crash when you hit the space bar to continue. This DDTS fixes this problem. [CSCdi05330]

This problem occurs when a server is moved from one physical cable segment to another, and both cable segments are connected to a router. The router must expire the neighbor entry for the old cable before it can learn a new entry for the new cable. During this period, as it receives routing updates on the new interface, it continues to process them even though they do not match the current neighbor entry for the server. [CSCdi06994]

Provide the ability to disable split horizoning of vines routing updates. This is needed to build a vines networks over a non-broadcast media, such as frame relay, when there is not complete connectivity between all nodes in the network. [CSCdi07034]

Provide quicker learning of alternate route when an interface goes down. [CSCdi07037]

When operating in serverless mode, some customers need the ability to flood a received broadcast to all other interfaces instead of choosing the best interface and sending the frame. This bug fix adds this capability and the supporting code so it may be configured. The new command is "vines serverless broadcast". [CSCdi07599]

Wide-Area Networking

An interface input queue may fill up and not recover if an X.25 provider in the RNR state receives and discards an I Frame and then violates the LAPB protocol by exiting from the RNR state with an RR instead of an REJ frame. The symptom is that the serial interface pauses indefinitely and ceases transmission. [CSCdi05957]

The X.25 PAD code will return a list of ALL X.3 parameters if we received an x.29 "read request" message with more than one parameter requested. This is improper, and will cause some X.25 implementations to clear the connection. [CSCdi06432]

The error message and traceback:

```
%X25-3-INTIMEQ Interface [chars], LCN [dec] already in timer queue, new time [dec]
```

is used as a diagnostic aid; although an unexpected condition was detected and reported, the operation of the router and the X.25 protocol are not affected. If this message is produced, contact Cisco Systems; include the text and traceback of this message as well as the information from the **show version** command. [CSCdi07238]

XNS, Novell IPX, and Apollo Domain

When we miss a SAP update we mark the entry as poisoned but if a subsequent SAP update is received we never remove it from the poisoned state so the SAP entry will always time out, even if only one update was missed. This problem has always existed but another patch added recently (CSCdi05359) has now exacerbated this previously unnoticed bug. [CSCdi06315]

Correct usage of Novell/XNS/Apollo transportControl (hop count) field, read/increment only hop count bits, discard packet when 16th router reached (hop count = 15, *not* 16), preserve reserved bits as packet transits router, minimize impact on Novell fast switching code when reserved bits are 0 (the normal case). [CSCdi06340]

XRemote

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 would be rejected as "martians" preventing nodes from acquiring their initial configuration when connected to a new network.

The workaround to this problem is to have at least one router on the cable which is not running version 8.3(4). [CSCdi06137]

Interfaces and Bridging

8.3(3) Caveats/8.3(4) Modifications

This section describes possibly unexpected behavior by Release 8.3(3). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(3). For additional caveats applicable to Release 8.3(3), 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(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 message. [CSCdi03924]

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 re-configure. [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]

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]

When an AppleTalk ARP reply is received on a Token Ring interface, the check that prevents entering multicast MAC addresses into the ARP table is done incorrectly; the least-significant bit (LSB) of the first octet of the address is checked instead of the most-significant bit (MSB). Configuring DECnet on a Token-Ring interface will force the MAC address of that interface to change to the DECnet standard, and the LSB becomes non-zero.

This may result in the system accepting invalid AppleTalk ARP replies, or, usually more seriously, in its ignoring valid ones. If the protocol allows, this can be worked around by reconfiguring other nodes to use Token Ring MAC addresses which do not have the least significant bits set in their first octets. [CSCdi05167]

This bug would affect the ability of a non-extended 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 other router than a Shiva FastPath on the cable, the interface will start routing as expected. [CSCdi05440]

Basic System Services

CPT now supports DDN X.25 [CSCdi04635]

The number of PT sessions supported has been increased from 32 to 64. [CSCdi04820]

The "show conf" 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 a user does a write memory command, it will save this config to the NVRAM. This will cause them to stay permanently in your config even in future releases. A user must use the NO commands for each line to clear the extra messages. [CSCdi04904]

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

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

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 workaround 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 workaround for problems with modems sending junk characters during various types of speed negotiation. The command is disabled by default. [CSCdi04852]

IBM Connectivity

SRB proxy explorer does not work. [CSCdi04671]

The **no bridge n address** command does not work properly. Although the specified entry is re-

moved, 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 re-computed 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]

Under certain conditions on the token ring interface (generally high traffic or noisy media), a message similar to:

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

may appear, indicating that the token ring interface was unable to reset itself. [CSCdi05644]

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]

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 address 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]

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]

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

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]

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

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

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

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]

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]

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

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

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]

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

If two interfaces have the same IP address and one of them is shut down, the other interface will not respond to an IP ping. [CSCdi04913]

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

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]

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]

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

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]

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

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

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

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]

If an interface enabled for multiring is reset, either by user action or by keepalives, then the router may issue "Bad enqueue" messages. The format of the message is "%SYS-2-LINKED: Bad en-

queue of 26BFE8 in queue 1E5450 -Process= "Net Background", ipl= 4, pid= 9 -Traceback= 7442 323F8 2EFF2 13ABA 10FF6 2434". [CSCdi05570]

ISO CLNS

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

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

When a prefix route goes unreachable an update is sent out with an infinity metric. Routers that receive the update that are not using the originator of the update as the current next-hop, will send a flash update about the destination. This causes unnecessary excess use of bandwidth and can lead to meltdown conditions. [CSCdi04845]

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

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]

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

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 workaround 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]

The routines 'iso_igrp_sendhello' and 'iso_igrp_sendupdate' always use the first NET that was configured for the router as the source NET. This would cause problems when hellos or updates were sent out an interface that was configured with a different NET. [CSCdi05542]

LAT

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

Certain LAT implementations generate messages with invalid (non-zero) contents of reserved fields. The cisco implementation, adhering to the spec, rejected such invalid messages. This causes problems communicating with some LAT implementations. [CSCdi04803]

Enabling **debug lat-packet** may cause a system reload to occur. [CSCdi05100]

Local Services

tacacs last-resort succeed does not work on lines configured for dynamic assignment of SLIP addresses. This problem is fixed in release 8.3(5), 9.0(1), and later releases. [CSCdi02330]

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

Setting the SNMP tsMsgInterval 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 0 for everything except login and slip address). This should only happen for responses to login requests. [CSCdi05175]

Protocol Translation

If multiple translate commands from tcp to lat are defined with the same tcp address, but different ports, the software will ignore all but the last of these, and all tcp connections will get translated to the lat host specified in the last translate command entered. [CSCdi04702]

TCP/IP Host-Mode Services

If a router is configured with a unnumbered serial interface and the serial interface is down, the corresponding numbered interface will not respond to IP pings. [CSCdi04236]

IP accounting reports the length of fastswitched IP packets incorrectly. [CSCdi04472]

If a FIN arrives out of order (eg because of a lost packet), then 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]

Under some obscure conditions (TCP connection receives a RST packet while the connection is closing and we 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 the condition 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]

TN3270

A user cannot put a character in the top left corner of a CICS request for transaction. A workaround is to hit the tab key, which will move the user over one character to the right. [CSCdi04643]

Transparent mode is not supported. Applications that depend on the passthru function of this mode will not work correctly. Some applications known to use this mode are kermit, SAS graphics stuff, and a serial printing application called TPRINT. [CSCdi04645]

This problem is caused by the IBM host sending a SET BUFFER ADDRESS command for a 132 column terminal. The IBM 3278-2 terminal (and cisco's implementation of tn3270) does not support 132 columns. In releases prior to 8.3(4), sending a SET BUFFER ADDRESS command that was out of range could cause the terminal server to pause indefinitely. [CSCdi05323]

Terminal Service

A system reload may occur if a modem line gets in a stuck TCP state and the user hangs up. [CSCdi04585]

Autobaud does not work to 19.2k with **databits 7**. [CSCdi04657]

If there are two or more LAT sessions on a terminal line, one of which has been suspended, the session timeout set by the **session-timeout** will not affect any of the LAT sessions. [CSCdi05480]

Uncategorized Items

Incoming connections fail to return to default settings once the session is terminated. [CSCdi04522]

When multiple connections come very quickly for the same port a race condition can occur which will cause a system reload. [CSCdi04569]

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]

VINES

On systems with token ring interfaces not configured for multiring, arp will fail if an arp request with a rif is received. [CSCdi04274]

Wide-Area Networking

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]

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]

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

The X.25 switch code doesn't 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]

Appletalk phase I fails to route over serial links configured for SMDS encapsulation. [CSCdi04914]

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, then the router may reload. [CSCdi05031]

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

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]

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]

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

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]

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]

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 workaround is to administratively shut the interface and then reinitialize it. [CSCdi05483]

XNS, Novell IPX, and Apollo Domain

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]

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** 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]

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

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 0x7ffffff. [CSCdi05048]

When an interface is shutdown, 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 shutdown, 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 3 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 a 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]

Novell SAP advertisements between parallel routers may loop when a server/service is down, until the hop count reaches 16 on all routers in parallel. The SAP loop may not subside until 3 routers * 60 secs (SAP interval) * 16 hop or 48 Minutes for three routers in parallel. [CSCdi05359]

XRemote

The restriction that limited XRemote X clients to 31 has been removed. The new maximum is 127. [CSCdi04672]

8.3(2) Caveats/8.3(3) Modifications

This section describes possibly unexpected behavior by Release 8.3(2). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(2). For additional caveats applicable to Release 8.3(2), 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(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 which is a multiple of 64 + 1 will fail. To see if this problem exists in a particular configuration, exam the output of the **show configuration** command. If there is a duplicate entry the list is broken. A possible workaround 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 insure that all packets will be accepted by those end-nodes. [CSCdi04611]

Basic System Services

The terminal server uses 7-bit comparison to identify escape, flow-control, and dispatch characters in the user data stream. This can result in problems when 8-bit character sets are in use, since ordinary user characters may be misinterpreted as control characters. The correct behavior is to give the user control over the type of comparison used. [CSCdi03972]

DECNET static mapping addresses didn't show up properly with the **show smds map** command. The addresses were incorrectly displayed as zero. CLNS didn't work properly with SMDS encapsulation. Appletalk didn't work correctly with SMDS encapsulation. [CSCdi04322]

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]

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

EXEC and Configuration Parser

The "setup" dialog for terminal servers support defaults to "yes" for the "XRemote font servers" response. The default should be "no". [CSCdi03928]

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]

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

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

IBM Connectivity

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

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 10 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]

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]

The **early-token-release** command requires the presence of SBEMON version 3.1 or STRMON version 1.1. [CSCdi07069]

Interfaces and Bridging

There is no way to disable 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 workaround 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 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]

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 workaround is to remove all knowledge of the network before changing the static route. [CSCdi03863]

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]

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]

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]

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

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 may 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]

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]

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

ISO CLNS

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 workaround is to configure CLNS consistently at both ends of each serial line. [CSCdi02905]

When Decnet Phase V (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]

Clns prefix routes which 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 re-learned. [CSCdi04583]

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]

Local Services

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]

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

TCP/IP Host-Mode Services

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

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

Uncategorized Items

If the router is assigned a CLNS NET using the **clns net command, and that NET is then removed using no clns net , 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]

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]

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

If you power-cycle one peer of an hdcl rsrp 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]

Wide-Area Networking

Under some conditions the router may reload when the **show x25 status** command is typed 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]

Under heavy load, LAPB could mis-handle 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]

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]

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

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]

XNS, Novell IPX, and Apollo Domain

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. 9.0 Cisco routers will not answer such echo requests. This means the the Novell ping command will work from 9.0 to any 8.3/8.2. It will not work from versions previous to 8.2(8)/8.3(3) to 9.0. [CSCdi03913]

1500 bytes is the largest IPX packet size currently supported. 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 6 times the update interval. A more reasonable value is 3-4 times the interval. [CSCdi04238]

In a network with equal cost multiple paths, the router may hear advertisements for the same service thru 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]

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]

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 as 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]

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

XRemote

XRemote font loading does not send a file data reply until additional data are sent by the remote terminal. This causes delays in font loading. [CSCdi02032]

8.3(1) Caveats/8.3(2) Modifications

This section describes possibly unexpected behavior by Release 8.3(1). Unless otherwise noted, these caveats apply to all 8.3 releases up to and including 8.3(1). For additional caveats applicable to Release 8.3(1), 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(2).

AppleTalk

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 **no appletalk distribute-list n in|out**. 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 workaround, the command may be added to a remote configuration file using a text editor. [CSCdi02792]

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 workaround, 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 **appletalk proxy-nbp** commands 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]

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

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

Basic System Services

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]

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 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]

If a terminal line has saved defaults (configuration items changed by the user which will be changed back to their normal values when the user logs off), and the "transport" configuration command is used on the line, the free memory header for the line will be corrupted. This will usually result in a system reload. [CSCdi02917]

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

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

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 which might fall into a load-sharing set, or on no interfaces which 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 Phase V to 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 Phase V to 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 Phase IV source area number to create the 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]

EXEC and Configuration Parser

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

IBM Connectivity

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

It is possible to define a new "stun schema" which 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]

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 func-

tionality, and has minimal impact on performance. [CSCdi03173]

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 workaround is to manually reload the router. [CSCdi03636]

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 as 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]

Interfaces and Bridging

When IP routing updates are sent through interfaces which have secondary addresses which 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]

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]

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 inability to communicate across serial lines between routers which are neither bridging nor routing IP. [CSCdi02692]

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

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]

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 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]

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

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

An IGS-PT without the bridging option cannot send LAT traffic over serial lines. [CSCdi02902]
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 which 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]

The show ip redirects 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 no ip routing , 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]

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]

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

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]

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]

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 acces-class". The system will not parse the files correctly when they are read back in. [CSCdi03630]

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

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

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]

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

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]

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 have 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]

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]

ISO 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 **clns enable**. **The workaround is to use the newer clns router static syntax. [CSCdi03258]**

LAT

Enabling flow control on a lat vty does not propagate the state change out the incoming connection. This can result in incorrect flow control being set on a connection. [CSCdi02409]

If the LAT engine retransmits an unacknowledged "balanced mode" message, it never sets the RRF flag, meaning that it doesn't solicit the remote node for permission to send a message with data. As a result, no data passes from the slave node to the master node until the master node sends another message. This can result in pauses for up to the keep-alive timer interval. [CSCdi02922]

Reverse LAT connections from hosts to terminal server rotaries do not work reliably. "Service in use" messages will often be returned for requests whose target lines are actually available. [CSCdi03766]

Local Services

If no domain name has been set using **ip domain-name**, 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]
TCP/IP Host-Mode Services

Transit packets from which the router has stripped IP security options are output malformed. The workaround 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 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 non-optimal in most user environments. The correct behaviour 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 acknowledgement, 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]

TN3270

A colon is missing from the default ttycap entry for VT100 terminals. This results in inefficient cursor motion during TN3270 sessions and in an incorrect value for the long-form terminal type string, which is rarely used. [CSCdi03744]

Terminal Service

If a line has been configured with **modem callout** or **modem inout**, and an outgoing TCP connection is made to that line, no data will be read from it until some data have been sent. A workaround is to use some other form of modem control; modem cts-required is a candidate for most applications. [CSCdi02704]

The system will sometimes fail to drop DTR in response to loss of CTS on a line configured with **modem cts-required**. [CSCdi02707]

The autobaud code does not correctly determine line speed for lines using baud rates outside the standard set. Although this is difficult to correct in general, the code will be modified to work with 14,400 baud, which is expected to become a common speed with the wide deployment of V.32bis dialup modems. [CSCdi03002]

Uncategorized Items

no ip-forward-protocol udp 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]

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]

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

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]

When a line's permanent configuration calls for hardware flow control, and the user temporarily reconfigures the line for software flow control, the flow control options will not be correctly reset when the line is freed. [CSCdi03795]

VINES

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

Wide-Area Networking

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 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]

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 a 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 IPX, and Apollo Domain

Some third-party Novell applications issue SAP updates listing services with network numbers of zero. The system readvertizes 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]

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 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 **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]

If a Novell SAP packet which 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]

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]