

9.0(9) Caveats

9.0(8) Caveats/9.0(9) Modifications

This section describes possibly unexpected behavior by Release 9.0(8). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(8). For additional caveats applicable to Release 9.0(8), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(9).

Interfaces and Bridging [CSCdi09691]

ISO CLNS

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]

VINES

9.0(7) Caveats/9.0(8) Modifications

This section describes possibly unexpected behavior by Release 9.0(7). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(7). For additional caveats applicable to Release 9.0(7), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(8).

AppleTalk

When configuring an AppleTalk access group on an interface, the access-group command may allow or disallow traffic in violation of the list. A work around is to issue the interface sub-command **no apple route-cache**. [CSCdi12917]

When converting NBP BrRq packets into NBP FwdReq, the system does not preserve the original DDP source address. It, instead, uses the address of the outgoing interface. This can short-circuit **access-group** filtering. [CSCdi13287]

When NBP BrRq and NBP FwdReq packets are converted to NBP LkUps, the source address is not preserved. This can cause access-groups to inadvertently filter out the LkUps. The workaround is to disable access-groups. [CSCdi14245]

Devices that perform gleaning of MAC addresses from AppleTalk Phase 2 packets may experience connectivity problems. This problem can manifest itself as services on the local network appearing and disappearing in Mac Choosers. There is no workaround. An upgrade is necessary. [CSCdi14732]

Basic System Services

Certain debugging messages are unexpectedly displayed to the console regardless of the state of the **logging console** configuration command. [CSCdi12665]

Under conditions of high network or tacacs authentication server load, multiple responses can be received by the router or communication server. The multiple responses can be lost and cause the input queue to fill up on the interface the responses were received on. [CSCdi13626]

DECnet

The router fails to become a Decnet Designated Router on an FDDI interface when it is supposed to do so (it is the highest priority Decnet router on the FDDI ring). As a result, Decnet router hellos to end-nodes are not sent out on the FDDI ring and the end-nodes on the ring do not see a Designated Router. [CSCdi10442]

Cisco routers do not listen to the DECnet multicast address for level-2 only routers. This can create problems in situations where DEC routers are configured level-2 only. The workaround is to con-

figure those routers for both level-1 and level-2 routing. [CSCdi14521]

Interfaces and Bridging

In certain environments, use of the **source-bridge proxy-explorer** command may cause a router to reload, reporting a "Jump to Zero" error. [CSCdi12328]

IP Routing Protocols

Under extreme circumstances, if autonomous switching is enabled (that is, **ip route-cache cbus** is configured), the router will reload. [CSCdi12415]

ISO CLNS

Under conditions which are not yet well understood, when a CLNS NET is configured on a router using the command **clns router igrp areatag net nsap1**, and is then "undone" by the command **no clns router igrp areatag net nsap1**, and another NET is configured by the command **clns router igrp areatag net nsap2**, the system may reload. Caution is advised when adding and deleting CLNS NETs. [CSCdi09094]

IS-IS, when redistributing static routes, should not include the prefix in a level-2 LSP if the next-hop interface for the static route goes down. This is not a problem for ISO-IGRP. [CSCdi13023]

The configuration command **redistribute isis** is not properly written to non-volatile configuration memory. [CSCdi13154]

TN3270

Under certain rare circumstances, the communication server may hang running TN3270. [CSCdi13290]

VINES

If the router receives a redirect that lists itself as the next hop for a router, it will process the packet resulting in a circular routing table entry. This makes the destination listed in the redirect become unreachable from behind the router. [CSCdi12292]

Wide-Area Networking

The **dialer-list 10** command would cause the router to take an exception. This is because only dialer lists from 1 to 9 are allowed. [CSCdi11279]

XNS, Novell IPX, and Apollo Domain

If you configure a Novell IPX static RIP or static SAP entry using a host id which matches the host id being used by any cisco interface the static RIP or static SAP will be disallowed. The verification of host id should use the entire network.host-id address for a match instead of only the host-id. The 9.0 version of this bug only applies to static RIP configuration as static SAP configuration is not supported in the 9.0 software release. [CSCdi13332]

When responding to a RIP request from a NetWare 3.1x/4.x Server/Router the response is sent to an incorrect MAC address (0000.0000.0001) and therefor is never received. This will only happen on NetWare devices which use an internal network number, a response to normal NetWare Client is sent to the correct MAC address. [CSCdi13400]

In a topology where multiple equal cost routes exist to a destination and novell maximum-path is still at the default value of one a situation can happen such that an old route-cache entry exists pointing to a route that no longer exists. Using a non default value of novell maximum-path will avoid this issue, which will clear itself the next time the route cache changes, or when a clear novell cache is done. [CSCdi14410]

Novell routes are flushed whenever a "novell network xxx" command is issued against an interface, even if the network number is unchanged from it's previous value.

This is most often seen when a config file is uploaded using the ciscoworks config management feature.

The impact is novell routing to some destinations will stop for up to 1 minute while the novell route tables are rebuilt. SPX sessions, which have relatively short timeout values, may be dropped. [CSCdi14444]

9.0(6) Caveats/9.0(7) Modifications

This section describes possibly unexpected behavior by Release 9.0(6). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(6). For additional caveats applicable to Release 9.0(6), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(7).

AppleTalk

AppleTalk GMZ (GetMyZone) packets received on a nonextended interface are not handled properly by the system and get held in the small buffer pool. Evidence for this problem would include a slow depletion on the available system memory (as shown by **show memory**) and a continuous rise in the small buffer "total" count (as shown by **show buffers**). A GMZ request on a nonextended interface is an undefined call and should be ignored, but some AppleTalk-based network management packages use these packets to determine network configuration. [CSCdi10715]

Occasionally, a newly configured MACIP server in a running router will not begin operation. Instead, it will hang in state "initial." This problem will only occur in routers that have been running for more than three weeks. The workaround is to configure MACIP prior in the first three weeks of operation, or to restart the router and reconfigure MACIP. This problem will not occur in routers that have a continuously running MACIP server. [CSCdi10771]

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]

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]

When **debug apple-events** and **debug apple-routing** are enabled, state changes for routes are reported with incorrect cable ranges. There is no system impact. To get an accurate picture of the state of a route, use **show appletalk route**. [CSCdi11558]

During a **write terminal** or a **show configuration**, trailing blank space in a zone name is not visible, although present. There is no system impact. [CSCdi11847]

Basic System Services

Changing the logging level via the **logging console** global configuration command does not limit the display of logging messages to the console. The workaround is to log in via a virtual terminal and control the logging of messages with the **logging monitor** global configuration command. [CSCdi11676]

DECnet

A router running with IV/V conversion enabled converts any Phase IV hellos it receives and adds it to the Phase V adjacency database. The format of this entry in the Phase V database is recorded

as Phase IV. If a corresponding Phase V hello comes in (that is, the other router is also running Phase V), it should overwrite the entry in the Phase V adjacency database that was always forwarding to the final destination instead of the next hop. A IV adjacency is stored in the V adjacency database as noted above. This information is also entered into the V routing table, so that it is propagated through the OSI cloud. The caveat results in the router not updating this route, so the route would go into holddown and ultimately go away. Therefore, Phase IV ES information never stays long enough in the V routing table. [CSCdi11174]

All data structures used by the 'show' routines are locked to avoid a premature 'free' of that chunk of memory. This chunk of memory is freed at the end of the 'show' routine. There were code paths where the 'show' function could 'return', without freeing the locked chunk of memory. This could lead to memory leaks. [CSCdi11545]

EXEC and Configuration Parser

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]

IBM Connectivity

The token ring interface was sending ring status messages to the lan manager when it was in the "DOWN" state. The status messages are valid only after the interface has begun the insertion process. [CSCdi10364]

Interfaces and Bridging

If you run **setup** from enabled mode and configure a Token Ring interface that was previously shut down and had no ring speed, the configuration fails because **setup** configures **no shutdown** before **ring-speed 16**. [CSCdi09032]

Spurious entries may appear in the bridge table (**show bridge**) when the MAC address of an interface changes (for example, in reconfiguring an interface with a different DECnet address). This can be corrected with **clear bridge n**, where *n* is the bridge group identifier. [CSCdi09802]

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]

There is a window in which commands to the interface get dropped. The fix is to protect against interrupts when issuing commands. In this case, the system drops the command to throttle the interface. When the system later tries to unthrottle the interface, it can get passed random pointer values to the interfaces shared memory.

Also, store the throttle count in idb and display in **show controller**. [CSCdi11046]

The bridge is not forwarding broadcast packets over a bridge circuit group. The packets propagate on both serial links but are blocked at the second serial interface on the other end. The show span command will display that the second interface is in the blocking state. [CSCdi11811]

IP Routing Protocols

If a router receives IP packet fragments which are broadcasts, or addressed to the router and the fragments arrive more quickly than they can be reassembled, large amounts of processor memory can be consumed. [CSCdi10903]

When RIP is turned on a connected route and RIP routes are redistributed into OSPF, the connected route is not redistributed into OSPF routes after **shutdown** and **no shutdown** is executed on that interface. [CSCdi10957]

With OSPF, **distribute-list OUT access-list #** with the interface specification does not work. In fact, providing interface on **distribute-list OUT** does not mean anything in OSPF. In later versions

of 9.1 release, users are not allowed to specify interface. [CSCdi11048]

The router will not accept a partial command for **ip route-cache** because of the addition of a new command **ip route-cache-same-interface**. [CSCdi11171]

Under certain unknown but infrequent conditions the memory pool may get corrupted causing the router to reload because of lack of memory. A show memory command prior to reload will show substantially less than expected memory in the Total Bytes column. [CSCdi11392]

Routes learned via core EGP are redistributed into BGP with an AS path of zero (0) rather than the AS of the remote peer. [CSCdi11575]

OSPF sets the forwarding address when the next hop is through an unnumbered serial interface. This points to an address in the external Link State Adjacency (LSA) that may or may not be in the OSPF domain. OSPF sets the forwarding address when the next hop is through an unnumbered serial interface. This points to an address in the external Link State Adjacency (LSA) that may or may not be in the OSPF domain. [CSCdi11583]

ISO CLNS

ISIS Hellos possibly appear to be causing an input queue to wedge. Removing ISIS from the interface and reloading clears the problem. [CSCdi10948]

CLNS pings do not work to a Token Ring interface. There is no workaround. [CSCdi11265]

The formula for metric calculation was not correct; in particular, setting K4 to zero and K5 to 1 would make the denominator of the expression to be zero, causing a 'zero error divide'. [CSCdi11705]

TN3270

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]

VINES

If a user waits at the --More-- prompt in the middle of **show vines neighbor** or **show vines route** output for any period of time, it is possible that the router will reload when the output is continued. This will only happen if the neighbor entry about to be displayed is deleted before the user continues. This is very unlikely to happen in normal usage of the router. [CSCdi10788]

In their Release 5.50, Banyan changed the way that a client determines the name of its routing server. This fix changes the router to support that new method as well as the old method. [CSCdi11384]

If a network is set up such that two or more routers are connected to a LAN segment containing a server, and the router interfaces connected to that LAN segment have been configured as serverless, then it is likely that there will be a broadcast storm. The workaround is to correctly configure the routers by removing the serverless specification on the interfaces that have servers connected to them. [CSCdi11991]

If **debug vines packet** debugging is turned on and the router is pinged by a 9.21 router, it prints messages about invalid packet formats. The router correctly returns the packets, but it should not print error messages. These messages occur because of over zealous checking in earlier software releases. 9.21 has a complete IPC implementation, and thus has nonzero values where 9.0 and 9.1 would enter zeros. [CSCdi12228]

Wide-Area Networking

Incoming SMDS ARPs are not entered into the SMDS ARP table. This is only evident in test situations where the interface is looped. There is no regular operational impact. [CSCdi10269]

XNS, Novell IPX, and Apollo Domain

The SAP Flash updates that result from adding a static SAP to a router are not filtered according to any assigned SAP filter list. SAP poison packets, hop count 16, are not filtered according to the configured SAP filter access list on the outgoing interface. Static SAP entries are Flash-announced to the world at the wrong hop count. When the correct hop count is sent in the periodic updates, it will cause neighbor routers to think the topology has changed and to place the service into hold down, timeout, and flash an advertisement of hops equal 16 before advertising the correct hop count. [CSCdi10834]

If a Novell SAP update is received that has more than the normal seven services per frame advertised and all those services are new, there is a strong possibility that memory will be corrupted. [CSCdi12108]

The optional behavior of the **rip-check** command installed as of CSCdi09056 has now become the default. To turn off the rip-check handling of RIP requests use the **no novell rip-check** command. Two new counters have been added to the Show Novell Traffic display: SAP format errors and RIP format errors. Should these counters be incrementing on a router, it might be prudent to investigate which client is sending malformed RIP requests by turning on **debug novell-rip-event**, information will then be displayed about the next one of these packets, which arrives along with other RIP events (which may or may not be interesting). Note: turning on debugging may cause unwanted overhead on the router. Use of an analyzer may also be warranted. [CSCdi12244]

9.0(5) Caveats/9.0(6) Modifications

This section describes possibly unexpected behavior by Release 9.0(5). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(5). For additional caveats applicable to Release 9.0(5), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(6).

AppleTalk

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

The router may remember old, deconfigured AppleTalk networks as directly connected for reconfigured AppleTalk interfaces in rare circumstances if all of the following are true:

1) The router has an existing cable-range or address associated with it. 2) The router is reconfigured with a new cable-range or address while it is not administratively shut down.

This problem does not exist in releases 9.1 and above of the router software. The workaround in release 9.0 is to ensure that one of the above two conditions are not met by either administratively shutting the interface down with the **shutdown** command, or removing the existing AppleTalk address with either the **no appletalk address** or **no appletalk cable-range** command. [CSCdi09635]

Under certain conditions, the configuration interface subcommand **multiring all** or **multiring appletalk** will prevent the router from being able to acquire an AppleTalk node address, thereby preventing the interface from becoming active as a routing node. You can detect this condition using the command **debug apple-arp**, which shows the router attempting to probe for an address indefinitely, incrementing the requested node address at each cycle. To circumvent this condition, remove the **multiring** command from the afflicted interface. (Multiring is necessary only if AppleTalk traffic will be source-routed from the adjacent Token Ring network to remote Token Ring networks.) If multiring is necessary, a temporary workaround is to disable multiring only during the AppleTalk ARP process. Once the interface has become operational for AppleTalk, multiring can again be applied to the interface. However, if the interface should restart for any reason, AppleTalk will again be disabled, so this should be considered an emergency workaround.

only. [CSCdi09753]

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]

NBP registered entities within the router on phase1 interfaces can not be seen by the directly connected devices on that interface nor a reply is sent for the nbp requests in nbp debugging. End result is neighbors may not see these entities even though they are registered in the router.No work around [CSCdi09819]

A MacIP server configuration may get deleted if the IP interface that it is associated with does not become operational immediately. The workaround is to ensure that the interface is up and then re-configure the MacIP server commands. [CSCdi09824]

A pending ZIP garbage collection request may not be fulfilled as expected. This can occur whenever a route and its associated zone is deleted. There is no user visible impact. [CSCdi10254]

All inactive zones may not be freed during ZIP garbage collection; **show appletalk zone** will display zones without any networks. This can occur when a large number of routes and their associated zones are deleted. There is no router impact. [CSCdi10279]

no appletalk cable-range and **no appletalk address** do not properly release assigned zone(s). As a result, the zone(s) may not be properly cleared during ZIP garbage collection and may show up as orphaned zones in **show appletalk zone**. There is no significant router impact. To ensure proper cleanup of zones, the user should issue a **no appletalk zone** command before issuing either of the two previous commands. [CSCdi10297]

Partially qualified AppleTalk addresses of the form, 0.X, are unexpectedly inserted into the AARP cache on all nonextended interfaces. Since the entries are not valid, they will shortly age out. No user intervention is required. [CSCdi10426]

AARP response debugging messages prints incorrect return addresses when **debug apple-arp** is enabled. [CSCdi10439]

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]

The AppleTalk address of the dissenting router is incorrectly reported as 0.0 when a network number conflict is discovered during port startup of an extended interface. There is no system impact. [CSCdi10839]

Basic System Services

FDDI valid transmission time will display 4294965054 in some cases. [CSCdi08923]

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

A condition can occur in which the available free memory becomes fragmented, and there is insufficient contiguous memory for certain processes to occur. For example executing the command 'write terminal' may result in the error message "Not enough memory, try again later".

A system reload may correct this condition. [CSCdi09382]

DECnet

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]

Any FDDI attached DECNET Phase IV end-nodes will have an OSI adjacency entry with a mul-

ticast SNPA. This occurs only when DECNET conversion is enabled on the router. [CSCdi09956]
DECnet was sending L2 updates starting with area number zero. This upset DEC routers, since valid DECnet areas start from 1. [CSCdi09981]

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]

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

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 its unique address, which is a problem in the above scenario. [CSCdi07105]

There is currently a problem (in both Software Release 9.0 and 9.1) of setting access filters on source-route bridge networks based on SNAP type codes. [CSCdi09010]

When a Cisco router with multiple Token Ring interfaces runs with the DECnet protocol, there are duplicate Token Ring MAC addresses on the bridge network because the Cisco implementation of DECnet modifies all the Token Ring interface MAC addresses to the same address. The IBM LNM protocol does not allow multiple stations with the same MAC address to exist on the bridge network. All the LNM functionality that relates to the duplicate MAC addresses, such as path test, station, profile, and link with bridge, will not perform normally.

A configuration command was added to allow the Cisco router's LNM module to accept link requests from the adapter that is not closer to the LNM station ring. In a normal case, the LNM station links with the adapter of a bridge that is closer to the LNM ring and expects to receive an error if an LNM station tries to link with the other end of a bridge. This addition allows a Cisco router to stay linked with LNM station and to report problems. However, other LNM station-related functionality is still not acting properly.

The following is the procedure to configure the router and LNM station:

1. Define the Cisco router as a bridge on LNM station. Use the burn-in address and the virtual interface address.
2. Issue the **lnm duplicate-address** global command on the Cisco router to turn on the option. [CSCdi09396]

The RSRB state machine goes to a null state when one of the peers of the WAN peers is power cycled. The workaround is to reset both routers. [CSCdi09767]

A slow memory leak occurred when LNM configured. (leading to router reload in more than a week). [CSCdi09881]

Interfaces and Bridging

The **transmitter-delay microseconds** command does not show up when issuing a **write terminal** or a **show config** command on ethernet or token ring interfaces. For this reason, the command must be issued at each *reload* for it to take effect. Serial interfaces function as described in the manual and don't exhibit the same failure. [CSCdi08710]

Static routes with destination gateways routed to via an interface that goes down (or is shutdown) are not always removed from the main routing table. [CSCdi09374]

If multiple NETs are configured on an interface, the router should send out multiple IS Hellos out

that interface, one for every NET configured (for that interface). [CSCdi09414]

Configuring local acknowledgment on only one side of a peer set may result in a system reload. [CSCdi09536]

OSPF doesn't listen to multicasts on an old Type 2 ethernet card. OSPF does not listen to multicasts on an old Type 2 Ethernet card. [CSCdi09553]

In the function for dealing with ring status messages, there is a test for the state DOWN that declares the interface UP, the assumption being that the interface does not issue ring status messages until it is inserted onto the ring. This is a breach of the keepalive process and preempts an attempt to put the ring into state TESTING. The offending code has been removed. [CSCdi09742]

In pre-9.0(5.4) environments, IP fast switching is not allowed on the same interface. This becomes desirable in a scenario like with three routers connected to the same Frame Relay network, where router A has DLCI to B and router B has DLCI to C. There is no DLCI between A and C, so traffic between A and C would have to go through B.

A new IP subinterface command has been defined to allow IP fast switching on the same interface:

int s 0 ip route-cache-same-interface

IP fast switching on the same interface and ICMP redirects are incompatible. Therefore, when the user enters the **ip route-cache-same-interface** command, ICMP redirects are never sent on the specific interface. If the user enters the command **ip redirect**, ICMP redirects are sent and the IP fast-switching cache is not updated with new entries if the output and input interface are the same. IP fast switching between serial interfaces does not work properly on low-end products in 9.0(5.3) and previous environments. This fix includes changes to the IP fast switching code to properly handle the frame header when switching between serial lines. [CSCdi09761]

If an interface flaps, or if an IP routing protocol is removed from the configuration, then the "gateway of last resort" will be lost. [CSCdi09903]

the Lan Net Manager "frame forward" used to verify an SRB route was causing a call to the function "send_trace_report()" with parameters in reverse order. This caused an attempt to jump to a null vector, thus "jump to zero error." The patches not only fix the function call, but also puts in paranoid code to check for invalid pointers. [CSCdi09980]

BGP routes aged out of IP forwarding table. [CSCdi09983]

OSPF default hello interval for non-broadcast interface is not set to 30 seconds as documented. Instead, it is set to 10 seconds. Interface subcommand **ip ospf hello-interval number-of-seconds** can be used to specify this interval. [CSCdi10027]

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]

In the case where there are excessive token-to-mother interrupts, the system should call `str_reset()` instead of `str_soft_reset()` so that the interface transitions correctly. [CSCdi10116]

On Cisco routers with Token Ring interfaces, enabling OSPF using the following commands may cause the router to execute an immediate system reload:

router ospf *ospf-process-id*

network *address wildcard-mask* **area** *area-id*

There is no workaround. Users wishing to use OSPF under these circumstances are advised to call the Cisco TAC for more information. [CSCdi10488]

The R16M will accept the configuration command **ring-speed 4/16** even though its ring speed can only be changed by a jumper. The interface display will show the ring speed from the **configuration** command. However, the ring will continue to operate at the correct (jumped) speed. The fix is to reject an attempt to change the ring speed on interfaces that are hardware configurable only.

[CSCdi10617]

IP Routing Protocols

It was observed once that a router was continuously looping running SPF, which resulted in locking the router. Not enough data is available on this bug. [CSCdi08089]

OSPF fails to install an external route which it receives in external link state advertisement in some circumstance. The mechanism that causes this to happen is still unknown. The way to work around when this happens is to cause the shortest path calculation to run again by issuing a **clear ip route *** command. [CSCdi09149]

OSPF installs a wrong next hop for a route that is advertised in AS external advertisement. This happens when there is more than one AS external advertisement to the same destination. [CSCdi09213]

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]

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]

The OSPF summary lock timer is created as a continuous timer where it should be a one-shot timer. If this timer is set once, it will try to come back even when it is not supposed to. [CSCdi09684]

If a new BGP neighbor is configured after the router has been operational for 24 days, BGP will not attempt to start the session. The workaround is to manually start the session with **clear ip bgp**. [CSCdi09732]

When an interface whose IP address is used as router ID by an OSPF router is shut down, the router mistakenly regenerates a router LSA with the old router ID that consequently fails to be deleted after an acknowledgement is received. This causes it to be continuously retransmitted.

Note that this does not prevent the router from performing the normal operation. The router changes its router ID and reforms adjacencies with its neighbors with the new router ID correctly.

This caveat is introduced in 9.1(3.1) and 9.0(4.2). [CSCdi09931]

If the IP encapsulation for an interface is changed from the default value of 'ARPA', fastswitching of other protocols may be disabled, particularly Novell. Packets will instead be process-switched. [CSCdi10046]

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]

The **show ip ospf database** can cause the system to reload when the link state advertisement is removed from the OSPF database after the command has been issued. [CSCdi10228]

This is a dynamic configuration problem. If you issue an **area range** command while the router is in operation, the router will not remove the summary LSAs that fall into that range. The workaround is that after completing the configuration, do a **write memory** and remove the OSPF process. Then configure the process again from memory. [CSCdi10293]

In OSPF routers, the **area area-id stub** command causes the router to lose neighbors. [CSCdi10295]

A router which is configured as an area border router in OSPF domain fails to generate a summary network link state advertisement into the backbone area for a network in non-backbone area that is configured as an interface's secondary address. [CSCdi10302]

A spurious OSPF HELLO packet may be sent during a system reload. The HELLO may not contain proper router ID information and can confuse other vendors' OSPF implementations, resulting in improper startup. [CSCdi10516]

The router executes a system reload if there is a virtual link configured and the interface whose IP address is used as router ID is shut down. The workaround is not to shut down the interface whose IP address is used as the OSPF router ID. [CSCdi10555]

When two routers are connected by an unnumbered serial link, OSPF does not calculate the routes properly. The workaround is to number the unnumbered link. [CSCdi10563]

The area route summarization command **area range xxxx xxxx** accepts 0.0.0.0 as the summary address even though this address might cause routing loops. You should not add 0.0.0.0 as the range address. [CSCdi10627]

ISO CLNS

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

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]

IS-IS will send level-1 LSPs over a point-to-point link to a level-2 adjacency. The router on the other end discards the packet, and the sending side continually retransmits these LSPs. [CSCdi09335]

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]

Duplicate adjacencies are formed (both system ID and SNPA are the same) when CLNS cluster aliasing is enabled on an interface. This happens for ISO-IGRP and DECnet Phase IV systems. This does not happen for IS-IS and OSI end-system adjacencies. [CSCdi09525]

There are four obscure cases where IS-IS does not purge its own LSPs. The effect is LSPs harmlessly stay in the database longer than necessary. [CSCdi09526]

IS-IS sends point-to-point IIHs out HSSI interfaces that are 1 byte larger than the allowable MTU. This results in a %TOOBIG.... error message. The adjacency still forms. [CSCdi09538]

If IS-IS areas are configured in neighboring routers such that they are not in the correct order, a level-1 adjacency will not form. This only occurs in multihomed areas over point-to-point links. [CSCdi09555]

Interface static routes with no SNPA specified will not be deleted from the configuration file. They are deleted from the routing table. [CSCdi09579]

The Cisco router will create an adjacency with an end-system that has advertised an invalid NSAP format in its ESHes. [CSCdi09670]

If an OSI end-system advertises an NSAP address that exceeds the legal length (20 octets), the Cisco router will accept and process the NSAP and will build an adjacency. [CSCdi09672]

If there exists a Phase IV end node directly connected to a router, and IS-IS is enabled where the router is designated router, the Phase IV end system is not inserted into the Level 1 routing table and therefore is not reachable. This is a problem for end systems that are both Phase IV and Phase V. [CSCdi09678]

IS-IS does not free the memory used for any LSP when the Lifetime expires and it is deleted from the link state database. This event does not occur very often. [CSCdi09759]

There are rare occurrences that the system may reload when a "show isis database detail" command is issued when the link state database contents is changing. [CSCdi09805]

The NSAP lookup routine goes through the entire hash table even when a matching entry has been found. [CSCdi09915]

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]

Under certain circumstances, CLNS over x.25 would not work because the encapsulation of the CLNS packets would fail. The call to the x.25 address look-up routine did not zero out the structure containing the x.121 address correctly, thus causing the look-up routine to fail; this error was propagated up to the encapsulation routine which also would fail. [CSCdi09966]

If a static CLNS route to the zero-length prefix ("default") is configured, it will not be written correctly to NVRAM.

The workaround is to install a small number of static routes of length one instead. [CSCdi09997]

If there is a neighboring IS on a LAN, and a router is configured to run IS-IS on the interface, the router does not advertise the IS as an ES link in the pseudo-node LSP.

This fix allows ISes that do not run IS-IS to be reachable via the IS-IS running domain. [CSCdi10002]

When there exists static routes in which the next-hop interface is no longer reachable, and ISO-IGRP is redistributing static routes, it will continue to do so if the interface goes down or the next-hop goes unreachable. [CSCdi10060]

When deconfiguring an ISO-IGRP routing process, static prefix routes learned by that process are not deleted from the routing table. These routes stay in the table indefinitely. A system reload is the only cure for the problem. [CSCdi10406]

There are situations when two routers running IS-IS are brought up on a serial interface and all the LSPs are not flooded to each other. [CSCdi10532]

If static adjacencies are configured before the IS-IS routing process is configured, the adjacencies are not inserted into the nonpseudo node LSP. This is a race condition and does not happen very often. [CSCdi10587]

The router will reject IS-IS packets when more than one SNPA with the same address is present in the CLNS neighbor table. This can be determined with the EXEC command **debug clns-routing**. [CSCdi10931]

The router may reload if more than one IS-IS PSNP/CSNP is sent. [CSCdi10939]

When the CLNS route cache is enabled (default) and a DECnet Phase IV adjacency has been established, it is possible for the Phase IV-Phase V conversion routine to forward Phase V packets (CLNS) to the Phase IV end system. This would result in a loss of connectivity when the Phase IV end system is attempting to connect to a Phase V host. Turning off the CLNS route cache via the interface subcommand **no clns route-cache** will act as a workaround, but may negatively impact

performance. [CSCdi10980]

LAT

A protocol translation to a LAT printer will routinely lose the end of the job. [CSCdi08256]

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]

If two users attempt a TACACS login or SLIP address request at the same time, the password one user types in can be sent with both authentication requests, causing authentication failures. This is due to the use of a static buffer. The problem will be fixed by using dynamic storage. If two users attempt a TACACS login or SLIP address request at the same time, the password one user types in can be sent with both authentication requests, causing authentication failures. This is due to the use of a static buffer. The problem will be fixed by using dynamic storage. If two users attempt a TACACS login or SLIP address request at the same time, the password one user types in can be sent with both authentication requests, causing authentication failures. This is due to the use of a static buffer. The problem will be fixed by using dynamic storage. [CSCdi10479]

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]

When a router is reloaded or powercycled, IP fast switching will not be enabled by default for an unnumbered serial interface. The workaround is to either explicitly enable fast switching on the unnumbered interface using the **ip route-cache** command or to assign an IP address to the interface, remove the IP address, and then make it unnumbered. This needs to be done every time the router is reloaded or power-cycled. [CSCdi09493]

A TCP connection that has transmitted a very large amount of data (on the order of 2 billion bytes) can remove packets from the retransmission queue prematurely, causing the connection to unexpectedly close due to a retransmission timeout, even though the network path is working correctly. This can affect router functions like remote source route bridging, which can transmit large amounts of data over a long period of time. [CSCdi09764]

The ability to debug TCP-based remote source-route bridging, X.25 switching, and SDLC tunneling, is inadequate. The new commands **debug ip-tcp-driver** and **debug ip-tcp-driver-pak** would be very useful. [CSCdi10382]

TN3270

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

Under certain circumstances, use of TN3270 may cause the Communication Server to hang. [CSCdi09987]

Wide-Area Networking

The protocol translator resets X.25 switches by sending pad control packets with the Q-bit set immediately after full data packets which have the M-bit set. [CSCdi09201]

An advertent bug was introduced in 9.0(4.5) that broke the ARP mechanism for SMDS. All ARPs, over an SMDS link, were being discarded preventing routing of IP traffic over SMDS. [CSCdi09781]

The protocol translator will now respond with a pad error message to X.3 parameter sets that are outside the range of recommendations X.3. [CSCdi10042]

With remote X.25 switching, and routing to a destination that cannot route the call, the originating router pauses indefinitely. [CSCdi10051]

XNS, Novell IPX, and Apollo Domain

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] 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]

The SHOW NOVELL/XNS ROUTE display is missing the count of learned routes in the header of the display. [CSCdi09923]

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]

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]

9.0(4) Caveats/9.0(5) Modifications

This section describes possibly unexpected behavior by Release 9.0(4). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(4). For additional caveats applicable to Release 9.0(4), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(5).

AppleTalk

The computed total in the summary line of the **show appletalk zones** command is not the same as the number of zone names shown in the output of the command. This is cosmetic and does not affect routing operation. [CSCdi06993]

The AppleTalk name lookup cache may not always be reflected in the output of various **show apple** commands. This affects the output of these **show** commands only; it does not affect any core router functionality. [CSCdi07775]

A **clear interface** command does not clear the IPtalk port. Use the configuration command **no apple iptalk** instead. [CSCdi07778]

AppleTalk zone multicasts such as NBP lookups are unexpectedly ignored on FDDI interfaces. [CSCdi09424]

DECnet

DECnet fast-switching on the Cisco 7000 works fine on Ethernet when the encapsulation is ARPA. It does not seem to work for ISO1/SNAP encapsulations. A (crude) workaround to the ISO1/SNAP ethernet encap switching problem is to: 1) enable the default encap (ARPA) 2) set your switching mode 3) enable the desired encap (ISO1 or SNAP)

Fast-switching is supported only for ARPA encapsulation (for Ethernet). The problem is that the code does not change the decnet 'fast-switch' flag to FALSE when the encaps flag is changed to anything other than ARPA. Likewise, the code does not set the flag to TRUE when the encaps is changed back to ARPA (and DECnet fast-switching is turned on). [CSCdi08415]

Turning on fast switching on an interface should be disallowed if that interface does not support fast switching, or in the case of serial interfaces, if the encapsulation does not support fast switching. [CSCdi08806]

The DECnet fast-switching code will not process an extended ACL if no standard ACL is present. To be consistent with the slow-switched case, the check for the presence of a standard ACL should be removed so that a list consisting of only extended ACEs will be processed. [CSCdi08875]

The DECnet fast-switching code will not process an extended ACL if no standard ACL is present. To be consistent with the slow-switched case, the check for the presence of a standard ACL should be removed so that a list consisting of only extended ACEs will be processed.

DECnet fast-switching on a Cisco 7000 works fine on Ethernet when the encapsulation is ARPA. It does not seem to work for ISO1/SNAP encapsulations. A (crude) workaround to the ISO1/SNAP Ethernet encapsulation switching problem is to enable the default encapsulation (ARPA), set your switching mode, and enable the desired encapsulation (ISO1 or SNAP). Fast switching is supported only for ARPA encapsulation (for Ethernet). The problem is that the code does not change the DECnet fast-switch flag to FALSE when the encapsulation flag is changed to anything other than ARPA. Likewise, the code does not set the flag to TRUE when the encapsulation is changed back to ARPA (and DECnet fast-switching is turned on). [CSCdi08879]

If a DECnet Phase V end node sends both Phase IV hellos and ESHs, the Cisco router continually changes the adjacency type stored in the OSI adjacency database. Therefore, packets are sometimes converted and sometimes not. The correct behavior is to set the adjacency type to Phase V and use this adjacency. Phase IV packets should then always be converted. Phase V packets should not. [CSCdi09235]

EXEC and Configuration Parser

The **lapb hold-queue** interface subcommand is not properly stored in the router's configuration memory. [CSCdi08957]

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]

Misconfiguration of the router with peers that do not exist or are powered down can cause the box to lose all memory. [CSCdi09041]

Ignore Format 3 type 4 XIDs for SDLLC connections. This is sent by NCP when VTAM is brought down. Once VTAM is brought down, there is no point for the router to initiate connection. [CSCdi09211]

An SNA frame that was sent and replied to with RNR, is never resent and causes the FEP to eventually drop the session by sending the DM. [CSCdi09218]

Interfaces and Bridging

The system reloads when OSPF processes the link state advertisement retransmission list. This happens right after the system starts up. [CSCdi04617]

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

A message appears when a LAN Network Manager trace frame is accepted by the router and forwarded onto other interfaces on that router, causing a minor memory leak because the wrong packet inside the router is trying to be freed. There is no workaround. [CSCdi07950]

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

The Chaos, PUP, and Hello routing protocols do not properly expire old routing entries, leading to a memory leak, race conditions, crashes, and incorrect routing decisions. [CSCdi08881]

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]

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]

There is a messaging scheme whereby the Token Ring interface board can send status information to the system. There was no protection against a runaway board dominating the system with interrupts. The fix is to watch for excessive amounts of interrupts over a short period and reset the board if necessary. [CSCdi09022]

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]

When TCP sessions go across a serial line that has TCP header compression enabled, those sessions may fail. This failure occurs if IP packets are fast switched from the serial interface that has the TCP header compression active to any other interface. Users may work around this problem by disabling IP fast switching in systems that have the TCP header compression active. This disabling need be done only for interfaces other than the serial interface doing TCP header compression. This is because TCP header compression automatically disables fast switching of IP packets on the serial interface. The configuration interface subcommand **no ip route-cache** can be used to disable

IP fast switching on an interface. [CSCdi09069]

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]

There was a condition whereby the Token Ring chipset would become the Ring Parameter Server but the LAN Network Manager could not discover this fact and so would not respond to requests by other stations to insert onto the ring. [CSCdi09108]

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]

Systems with X.25 map entries for more than one remote AppleTalk host on a single interface are subject to reloads. [CSCdi09328]

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

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

During a designated router election process, a router that used to be a designated router but just lost the election fails to choose itself as backup designated router when it should. The correct behavior is to choose a router with the highest router priority among the rest, excluding the router that declared itself as designated router. [CSCdi08732]

OSPF generates a Seq Number Mismatch event after receiving a duplicate database description packet after it moved into state Full and it was a slave during database synchronization. The correct behavior is to simply discard it, up until Dead Interval time since transition into state Full. After that period, it will generate a Seq Number Mismatch event. [CSCdi08829]

When configuring a router with the **redistribute static metric-type 1** router subcommand for OSPF router, the metric-type 1 argument is correctly set for redistributed routes but it is not recorded in configuration file as indicated by **write terminal** command. This can cause the router to use the default metric-type of 2 if the incorrect configuration file is written to either file or memory, then reloaded back to the router. [CSCdi08870]

BGP does not accept advertisements of network 0.0.0.0. [CSCdi08880]

The system reloads after loading configuration file with the **distribute-list access-list-number out** router subcommand for an OSPF router. This only happen when loading configuration file from TFTP server. Configuring From the console will not cause a reload. [CSCdi08956]

If a BGP router learns a route via IBGP and it has an EBGp neighbor as the next hop, and it then advertises the same route to the EBGp neighbor, the resulting next hop will be the EBGp neighbor

itself. This will cause the BGP session to disconnect. [CSCdi08963]

An OSPF packet is sent with IP-TTL 1 on a virtual link. This can cause the packet to be discarded when it is crossing the transit area. The IP-TTL for packet to virtual link is now set to 255. [CSCdi09000]

There are some cases when OSPF processes an incoming summary link state advertisement, the system will reload. This problem occurs under heavy OSPF load conditions. [CSCdi09090]

OSPF removes the wrong instance of link state advertisement from link state retransmission list after receiving a link state acknowledgment. This happens in a rare circumstance when the acknowledgment is for an older instance of link state advertisement. [CSCdi09189]

OSPF module miscalculates whether two link state advertisements are the same instance. [CSCdi09190]

OSPF module miscalculates whether two link state requests are for the same link state advertisement instance. [CSCdi09338]

Debugging messages showed by the OSPF module during the designated router election process shows a wrong router ID. [CSCdi09411]

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 redistributing core egp into another protocol, the command **redistribute egp 0** is written out as **redistribute egp** which is an invalid command. This only happens if the EGP AS is 0. [CSCdi09524]

ISO CLNS

CLNS static routes will not be written to NVRAM when a routing protocol has learned the same route and has better administrative distance. The correct behavior is for static routes to be written to NVRAM. [CSCdi05767]

If there are multiple options present in an IS-IS hello packet, there are cases that the area address is not extracted and stored in the adjacency database. This occurs when the router on the other end of a serial link advertises both an IP address and an area address. This does not occur between two Cisco routers, since integrated IS-IS is not supported until a later release. [CSCdi09048]

When using IS-IS as the OSI routing protocol, any static routes that are configured are not entered into the Level 1 IS-IS routing table. As a result, route table lookups on the static address fail. The IS-IS code will add a route to the routing table if the route is IS-IS or ES-IS derived; it should also add the route if the route is a static one. [CSCdi09053]

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

When redistributing ISO-IGRP routes into IS-IS, there are cases where some routes do not get redistributed. This occurs when the number of ISO-IGRP prefix routes causes more than one IS-IS Level 2 LSP to be generated. The routes that overflow the first LSP do not get generated. [CSCdi09144]

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

Certain OSI packets are transmitted with a nonzero padding byte. The padding byte is inserted between the OSI protocol data unit and the relevant data link layer encapsulation. Since padding bytes are supposed to contain the value zero, this condition can confuse equipment from other ven-

dors and cause interoperability problems. This caveat does not occur between Cisco equipment because Cisco equipment always expects 1-byte padding and does not specifically check the value of this padding byte. [CSCdi09238]

There are situations where IS-IS will delete the wrong link in an LSP. This results in either duplicate entries or corrupt LSPs. [CSCdi09466]

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]

Run from ROM software (igs-kr or igs-bprx) may not properly advertise LAT services defined in the communication server or protocol translator. [CSCdi08837]

Local Services

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]

If a router or comm server is configured with a username having an encrypted password of invalid format, it is possible that the unit will reload when someone tries to log in using that username. The only way to get an encrypted password is for the cisco unit to create it; users should not enter: username myname password 7 mypassword Since "mypassword" is not a valid format for a type 7 encrypted password. [CSCdi08805]

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]

Terminal Service

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 output** 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 output** command. [CSCdi08625]

VINES

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]

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

When receiving a Novell RIP request packet with an invalid length, it is possible that a system data structure could become corrupted. In extreme cases this could cause the system to crash at some later time. The **novell rip-check** command will turn on special processing of these invalid RIP request packets, avoiding the undesirable behavior. [CSCdi09056]

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

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]

9.0(3) Caveats/9.0(4) Modifications

This section describes possibly unexpected behavior by Release 9.0(3). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(3). For additional caveats applicable to Release 9.0(3), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(4).

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]

This is a cosmetic fix. When a router receives a nbp lookup to its net# and the zone, it finds all the matching entities in that entire zone and sends the nbp reply. Correct behaviour is to just respond with the entities registered on that interface only. Routers do not have to reply with all entities in that zone this is a NBP protocol function to generate the lookups to all the networks that correspond to this zone. This redundancy is eliminated [CSCdi08746]

When the **appletalk permit-partial-zones** command is enabled, the **appletalk distribute-list access-list out** and **appletalk getzonelist-filter access-list** commands unexpectedly permit all networks and zones in RTMP updates and GetZoneList replies when used with access lists that contain no zone information (that is, network number restrictions/permissions only). [CSCdi08819]

Basic System Services

show process memory can be inaccurate due to incorrect accounting of deallocated memory. [CSCdi07586]

Output of "show process" outputs octal numbers, even if service decimal-tty is set (the default on the protocol translator and 500-CS). [CSCdi08240]

The router may experience a software error when the command **show memory free** is executed, and the command must pause for output at any time in displaying the results of the command. The workaround is to ensure that the output does not pause by using the command **terminal length 0** before issuing the **show memory free** command. [CSCdi08368]

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]

A packet going from one DECnet host to another on the same LAN should not be subject to access control checks. Making these packets go through the access control check serves no useful purpose since end systems can easily discover that they are on the same LAN and bypass the router altogether. This makes any access control set-up useless for such packets. Also, the result of this is that two end systems on the same LAN cannot talk to each other if they end up using the router to "discover" each other for the first time. [CSCdi08121]

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]

When a DECnet extended access list is configured with a destination address, the code ignores the destination/mask information in the ACL. If a match was found in the connect part of the ACE, it would return TRUE, that is, grant access, regardless of the destination/mask information. For example, **access-list 300 permit 1.400 0.0 1.999 0.0 eq any** should allow only packets from 1.400 to 1.999 to go through. The observed behavior was all packets would go through, regardless of destination. The fix is to just check that the source address/mask (and destination/mask, if applicable) specified in the access list matches the corresponding values in the incoming packet. [CSCdi08760]

With DECnet access lists, the destination address/mask is ignored, regardless of what is in the connect part of the ACL. If the connect part of the ACL matches, access is granted, regardless of the destination address/mask. [CSCdi08818]

EXEC and Configuration Parser

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

IBM Connectivity

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]

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

Initializing Token Ring causes existing LNM links to be dropped. There is no further information available concerning this problem. [CSCdi07235]

The **show sdlc** should be removed from **show ?** display. There is no further information available concerning this problem. [CSCdi07287]

If the ring-group parameter for the **sdlc traddr configuration command** is configured before defining the ring-group (by issuing the **source-bridge ring-group** configuration command) it could cause the router to crash.

Now, the **sdlc traddr** command will not be accepted, if the ring-group parameter specified is not already defined by the **source-bridge ring-group** command. [CSCdi07317]

Repeated disconnections of the router could cause the router to hang. This was especially seen with Lan Network Manager sessions.

The problem was that multiple llc2 control blocks would get allocated pointing to the same session. [CSCdi08350]

Interfaces and Bridging

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

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

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]

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]

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]

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

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]

If IS-IS is not configured to redistribute static routes but is configured to redistribute ISO-IGRP routes, in some cases the ISO-IGRP routes are not propagated. [CSCdi08231]

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]

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]

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]

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 RIF structures are now initialized before use. It is possible that a previous use of a RIF structure had entries that could affect operations when the RIF entry is reused a second time for a different purpose. This has caused problems of pings being unsuccessful, unable to reach SRB hosts, and so on. Initializing an entry will clear out all previous usage and start afresh. [CSCdi08790]

Novell SNAP 8137 and Novell ISO1 E0E0 on both TR and FDDI are currently classified as the same type. Incoming FDDI SNAP 8137 packets will not be fast switched. [CSCdi08820]

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]

IP Routing Protocols

In certain obscure circumstances and configurations, internal BGP paths which are not yet synchronized can be preferred over external BGP paths. This can cause instability in both BGP and in the IGP. [CSCdi08113]

When a subnet is known via OSPF and is redistributed into some other protocols (for example, BGP or another OSPF) and the route to the subnet is removed, the other protocol may remove that entire network from its routing table. [CSCdi08129]

In a very large networks, it is possible for fragmentation to occur on OSPF packets. This can cause problems with routers that do not do proper reassembly. [CSCdi08210]

Duplicate AS path regular expressions are not ignored with the consequence that they will show up more than once in the list if a box is configured with the same set of ACL's more than once. Duplicate AS path regular expressions are not ignored with the consequence that they will show up more than once in the list if a box is configured with the same set of ACL's more than once. [CSCdi08228]

Whenever inconsistent metrics are assigned to a router interface, it is possible to run into this bug. The result of this bug is that the route entries in the IP routing table will sometimes drop the interface or will have a wrong interface. The workaround is to have consistent metrics in the network. [CSCdi08297]

When a route boots from ROM, it ignores OSPF configuration in NVRAM. After booting, enter the command **config mem**. [CSCdi08409]

If a summary LSA is regenerated within 5 seconds, the flooding of the LSA may not happen resulting in inconsistent database. The fix will be available in a future release. [CSCdi08463]

When a link is flapping continuously, it is possible to run SPF calculations after each topology change, resulting in locking the router. There is no workaround. This will be fixed in a future release. [CSCdi08600]

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; other-

wise, 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 **no clns enable** command does not check to see whether or not a dynamic protocol is active on an interface before disabling CLNS on the interface. [CSCdi07413]

The MTU of CLNS is always set to be three less than the IP MTU on the same interface. This works for Ethernet/802.3, but is incorrect for other media. This bug could cause CLNS to attempt to generate fragments larger than can be reasonably sent on an interface, resulting in packet loss, although this is unlikely to happen in practice. [CSCdi07875]

The **show clns route** command will display unused next-hop addresses when one of the equal-cost routes goes down. [CSCdi08262]

If the **isis metric value** interface subcommand is entered and the IS-IS process is not created (no previous **router isis** command), the system may crash. [CSCdi08434]

If a Cisco router is an IS-IS designated router on a multiaccess network, it will transmit LSP entries in CSNP packets with a negative lifetime. This is only a problem if a receiver uses the lifetime information, and Cisco routers do not. This was found while doing interoperability testing with IBM's IS-IS implementation. [CSCdi08435]

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]

This problem only occurs when you run an ISO-IGRP routing process where you enable Level 2 only routing for all interfaces for the processes routing domain. For example:

```
router iso-igrp 39 net 39.0001.0000.0c00.ffff.00 int e 0 clns router iso-igrp 39 level 2 int e 1 clns
router iso-igrp 39 level 2
```

ISO-IGRP routes are created, ISO-IGRP adjacencies are not. Routes may not go away. [CSCdi08745]

If you enter the command **no router iso-igrp 39** all prefix routes created by this process will not be removed from the CLNS prefix routing table. A workaround is to do a **clear clns routes**. Also, if you enter **router iso-igrp 39** and **distance 90**. Prefix routes that are created by this process are not assigned a distance of 90. A workaround is to enter the **clear clns routes** command. The next

updates received will build routes with a distance of 90. [CSCdi08755]

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]
```

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]

If the subnet mask is changed after a system has been operational, the new subnet mask will not be reflected in the IP routing table. A workaround is to reload the system after changing the subnet mask. [CSCdi05915]

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]

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 car-

riage 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 processing 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 implementation 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]

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]

Wide-Area Networking

Once enabled, disabling X.25 routing with the **no x25 routing** command does not disable X.25 call forwarding. [CSCdi06840]

When bridging over a frame relay network that supports a multicast facility, the router learns the multicast DLCI as the source rather than the DLCI of the sending router. This causes all bridged packets to be sent on the multicast channel rather than being sent directly to the destination DLCI. This does not impact operation when using the frame relay map bridge command.

When bridging over a frame relay network that supports a multicast facility, the router learns the multicast DLCI as the source rather than the DLCI of the sending router. This causes all bridged packets to be sent on the multicast channel rather than being sent directly to the destination DLCI. This does not impact operation when using the frame relay map bridge command. [CSCdi08280]

Back-to-back routers cannot communicate using frame relay. [CSCdi08408]

The change made for CSCdi08408 had the side effect of clearing dynamically learned frame relay maps every 10 seconds. This behavior is seen when the frame relay network supports a multicast network and dynamic address resolution is being used. In addition, the status display from the show frame-relay map command indicates all DLCIs are "added" rather than "defined" as expected. [CSCdi08482]

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]

XRemote

Invalid Xremote packets from an X-terminal may cause the cisco Communication Server to reload. [CSCdi08110]

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

9.0(2) Caveats/9.0(3) Modifications

This section describes possibly unexpected behavior by Release 9.0(2). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(2). For additional caveats applicable to Release 9.0(2), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(3).

AppleTalk

Customer sites which have AppleTalk networks which are incorrectly configured would experience a gradual loss of available free memory in the router.

This problem would be exacerbated by the existence of a non-extended (ie, Phase I) route in a fully extended AppleTalk internet. A fully extended internet is one which does not meet the Phase I to Phase II transition criteria which are:

- no cable may be configured with more than one zone name.
- no cable may have a 'wide' cable range, eg 100-101, as opposed to the compatibility cable range of 100-100.

Customer sites that experience a high amount of appletalk route instability (which can be determined with the configuration command 'appletalk event-logging') may find that the router loses memory at a faster rate. In situations where a large number of routes are lost at the same time, the resulting loss of free memory may occur quickly. [CSCdi05619]

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]

Due to a logic error in the IPTalk functionality of AppleTalk, packets sent to a Unix(tm) host running CAP (Columbia AppleTalk Package) would be receive a double encapsulation of 6 bytes of AppleTalk LLAP information instead of 3 bytes. This would cause the packets to be rejected by CAP or if they were accepted, to be rejected by the CAP applications such as 'AUFS' and 'lwsrv'. [CSCdi05850]

The AppleTalk executive command **show appletalk neighbor** *net.node* will not show neighbor entries which are associated with an interface which has either lost line protocol or has been administratively shut down. These neighbor entries will be seen in the list of neighbors produced by

the **show appletalk neighbors** command without any argument, but cannot be specified in the detailed **show appletalk neighbor** *net.node* version of the same command. [CSCdi06089]

A spurious "couldn't register" traceback message may occasionally be seen on an interface that is being enabled for AppleTalk support if the interface resets at the same time the support is being started. This traceback message is harmless and can be ignored. [CSCdi06171]

The AppleTalk nbptest facility does not correctly represent 8-bit characters in output resulting from the use of 8-bit characters in NBP entity names or zone names. As this is a diagnostic facility, there is no impact upon routing functionality. There is no workaround for this problem at this time. This error will be fixed in a future release. [CSCdi06266]

The AppleTalk "show appletalk zones" command may show zones which are pending for deletion. The AppleTalk implementation will not free unused zone names from the zone list immediately; users executing the "show appletalk zones" command may indeed see zone names without a corresponding network number. This is nominal.

No workaround is needed.

In a future release, the "show appletalk zones" command will not list those zones for which deletion is pending. [CSCdi06269]

If an AppleTalk IPtalk interface was configured to use a UDP port other than the default (768), the configuration command "appletalk iptalk-baseport" was not written to NVRAM or the configuration file.

A workaround for routers which boot their configurations from a TFTP server would be to manually edit the config file to add the command "appletalk iptalk-baseport n" to the configuration file after the interface configurations. [CSCdi06297]

When the global configuration command **appletalk ipatalk-baseport** would be used to change the default IPtalk base UDP port, there was no way to restore the default setting with a no **appletalk iptalk-baseport** command.

A workaround to this error would be to use the command **appletalk iptalk-baseport 768**. [CSCdi06901]

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 second 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 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]

When the **debug apple-arp** and **debug apple-error** commands were enabled and the ARP cache ager was activated (eg, as a result of the four-hour timer expiring) the router would hang when the AppleTalk cache invalidation code would attempt to print out an informational message that the AppleTalk cache was being cleared.

The workaround to this problem is to use the **debug apple-arp** command to debug specific problems and to not leave it enabled any longer than necessary. y [CSCdi07102]

When **debug apple-nbp** was enabled, useful information about the processing of NBP Lookups and Replies was not generated. [CSCdi07172]

AppleTalk implementations on some other vendor's equipment can generate incorrectly addressed packets that could cause Cisco routers to retransmit the packet out the same interface from which it was received. This unexpected behavior can only occur on wide extended-cable configurations (those interfaces that are configured with a cable-range that does not have the same starting and ending network number.) [CSCdi07345]

Loading interface-specific MacIP configuration commands across the network will cause MacIP to fail when the interface is already configured for AppleTalk and the configuration file read from the network also contains the same configuration commands that reenables AppleTalk on that interface. The workaround for this problem is either to not load in MacIP commands across the network, by saving the commands locally in NVRAM on the router, or to not specify the interface-specific AppleTalk commands in the host configuration file being read across the network. [CSCdi07353]

When a router is configured with an AppleTalk zone name that begins or ends with a special 8-bit graphics character, NBP lookup queries made to this zone in the router will cause the router to reload. The workaround is to not configure zone names that begin or end in 8-bit characters. A crash can also occur when an NBP search is performed with graphics characters at the beginning or end of the type field. For example, a server with a trademark symbol at the end of the server-type name will cause the router to crash if it is installed on a zone connected to the router. The workaround is to move the server to a zone not assigned to the router, so that lookup requests for this type of service will not be directed at the router. [CSCdi07672]

No user impact.

When the router loses connectivity with a neighbor, the output of SHOW APPLE NEIGHBOR reports that the neighbor is overdue instead of down. It should only report that the neighbor is overdue after the update interval but still within the valid interval. [CSCdi07752]

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]

Users are starting to find that they are running out of memory (seen by a "buffer overflow" message) on their csc3 processors when netbooting 9.1 images on top of 9.1 roms. Due to Cisco's ever expanding software base, the csc3 does not hold enough memory to directly do this.

As a work-around, users should try one of the following:

1) consider buying a csc4 which has 16Meg. 2) use the secondary bootstrap by putting a jumper in bit9 on their csc3. 3) netboot a compressed image. This will run just a fast once uncompressed. [CSCdi05751]

Attempting a LAT connection to a line configured with an extended access list (access list of 100 or greater) will cause an error message to be generated and the connection to fail. [CSCdi05928]

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]

Memory intensive applications can fragment memory. [CSCdi07275]

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]

OSPF may stop working after 49 days. The work around is to deconfigure, then reconfigure, OSPF. [CSCdi07671]

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

Under some conditions the **show decnet route** command may cause the router to reload. [CSCdi05272]

The default 'max address' value in DECnet was defined to be 255, instead of 1023. This means that if the router was not explicitly configured to accept node numbers larger than 255, it would not talk to any DECnet host whose node number was greater than 255.

The workaround was to just set the 'max address' value to 1023. The fix is to set the default to the maximum value i.e. 1023, so that no explicit configuration is necessary. [CSCdi05380]

When DECNET Phase IV to Phase V conversion is enabled, but DECNET Phase IV routing is not enabled on an interface on which both Phase IV DECNET hellos and Phase V/CLNP ESHs are being received from the same DECNET node, no working DECNET adjacency is formed with that node. The node in question will appear in the DECNET routing table alternately as a Phase IV and a Phase V neighbor, but no traffic will be exchanged with it.

This may be worked around by making sure to enable Phase IV DECNET routing on all interfaces on which Phase IV hellos may be received. [CSCdi05885]

For systems running releases before 9.0(2.1), a memory leak occurs when an OSI route lookup fails and DECNET conversion is enabled in the router. [CSCdi06837]

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]

When a router is configured to perform DECnet conversion and IS-IS is the routing protocol enabled, OSI adjacencies created by DECnet IV are not inserted in IS-IS LSPs. This causes loss of connectivity of Phase IV end-systems through IS-IS clouds. [CSCdi07850]

DECnet Phase IV end-systems do not get propagated through an IS-IS routing domain. [CSCdi07938]

EXEC and Configuration Parser

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]

When the system is configured for ANSI Annex D LMI on a frame relay interface and the user writes the configuration to non-volatile memory, the system generates an unneeded command defining the LMI DLCI. When the nonvolatile memory is read, the system complains about the unneeded command. There is no impact on system operation. [CSCdi07735]

IBM Connectivity

In early versions of the bridging software IEEE BPDUs weren't always well formed. That is TCN

BPDU's would not get transmitted properly (like not at all). [CSCdi05981]

show interface serial displays the hardware mac address of the token ring board instead of the virtual mac address. The correct behavior is for the address configured with the command **sdllc trad-dr** to be displayed. [CSCdi06061]

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 responding to LLC XID request frames, the router responds with an XID response which is truncated by three bytes. The behavior is present in all versions of the router which support LLC. Most LLC implementations will still interoperate with the router truncating the frame response, and the impact is a minimal amount of extra traffic to establish an LLC2 connection. Once the connection is made, there is no impact at all. The truncated XID response frame is rejected by the originator who is requesting a connection. The frame rejection (FRMR frame) causes the router to trigger a connection request frame (SABME frame) and the connection continues as specified by normal LLC2 protocols.

The router should respond with XID of correct length. Cisco expects to resolve this behavior in a future release of 8.3. [CSCdi06733]

IBM OS/2 2.0 generates an incorrect response to received SABME frames with the Poll bit set to zero. OS/2 responds with a UA frame with the Poll bit set to one. The LLC2 standard requires that an outstanding SABME be acknowledged by a UA frame which has its Final bit set to the state of the Poll bit of the SABME, therefore, the LLC2 component in the Cisco router ignores the acknowledgment and the connection is never made. Therefore, attempts to link to the LAN Network Manager component of the Cisco router with LAN Manager or LAN Network Manager running under OS/2 2.0 will fail, and it will be impossible to manage or monitor the IBM features of the router from an OS/2 2.0-based management platform. Two workarounds exist. The first is to manage the router with management software running under older versions of OS/2 2.0, or to manage the router with NetCentral or some other management agent other than OS/2 2.0, that does not exhibit this bug in LLC2 behavior. In a future release, the Cisco router will send a SABME with Poll bit set to one when requesting a connection to IBM Network Management SAPs so the response, which has the Final bit set in any case, will be acknowledged, and the bug in OS/2 2.0 will be avoided. [CSCdi07429]

When a source-route bridge uses a locally administered MAC address, LAN Manager 1.3 and LAN Network Manager 1.0 require that bridge to respond to a TEST frame with a RIF field even when both the bridge and the LAN Manager station are on the same ring. Cisco routers do not respond in this manner. As such, the LAN manager station closes its adaptor when attempting to link to a Cisco source-route bridge. Users can use the universally administered MAC addresses on the Cisco routers to work around this problem. [CSCdi07598]

llc2 operation parameters always use the default value for the rsrb virtual interface. They can not be changed. [CSCdi07600]

IBM OS/2 2.0 generates an incorrect response to received SABME frames with the Poll bit set to zero. OS/2 responds with a UA frame with the Poll bit set to one. The LLC2 standard requires that an outstanding SABME be acknowledged by a UA frame which has its Final bit set to the state of the Poll bit of the SABME, therefore, the LLC2 component in the Cisco router ignores the acknowledgment and the connection is never made. Therefore, attempts to link to the LAN Network Manager component of the Cisco router with LAN Manager or LAN Network Manager running under OS/2 2.0 will fail, and it will be impossible to manage or monitor the IBM features of the

router from an OS/2 2.0-based management platform, Two workarounds exist. The first is to manage the router with management software running under older versions of OS/2 2.0, or to manage the router with NetCentral or some other management agent other than OS/2 2.0, that does not exhibit this bug in LLC2 behavior. In a future release, the Cisco router will send a SABME with Poll bit set to one when requesting a connection to IBM Network Management SAPs so the response, which has the Final bit set in any case, will be acknowledged, and the bug in OS/2 2.0 will be avoided. [CSCdi08704]

Interfaces and Bridging

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]

ES-IS and IS-IS do not use ISO 10589 multicast addresses for 802.5. [CSCdi05093]

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]

test interfaces command is not working. [CSCdi05977]

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

The router can reload in some cases if a BGP call collision occurs.

Generating large IGRP update can cause performance problems.

BGP update messages could be incorrectly generated.

If unknown BGP attributes are received, memory can be lost.

If the router is low on memory, BGP can corrupt its internal data structures.

BGP can interact with the environmental monitor card, corrupting BGP data structures. [CSCdi06106]

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]

Bridged packets received on an FDDI interface are flooded in FDDI-encapsulated form onto Frame Relay links in the bridge group. As a result, MAC addresses across the FDDI are not learned across the FR link. This is a problem only for bridge groups of more than 2 interfaces. [CSCdi07130]

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]

Other broken implementations of EGP can incorrectly send a subnet address in the POLL message. This message will be ignored. [CSCdi07271]

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]

copy to/erase of the MC+ fails: _str_mc_write_short There is no further information available concerning this problem. [CSCdi07314]

Bridged packets flooded from a FDDI interface will have their trailing byte corrupted when the packet length is 3 bytes more than any multiple of 4. [CSCdi07401]

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 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]

IP Routing Protocols

For area border routers, connected to the backbone only by a virtual link, the **show ip ospf** will indicate an area count one greater than expected and the record for the backbone will appear twice. [CSCdi04917]

When multiple ASBRs in a network generate functionally equivalent AS external advertisements, then the router advertisement with the higher OSPF router ID is used. If the LSA with the higher ID router is MAXAGED due to loss of route, the correct behavior is for the lower ID router to regenerate the LSA if it has the route. This regeneration does not work if only some of the routes are lost and not all. [CSCdi05681]

Executing the command **show ip ospf database** may cause a system reload to occur.

[CSCdi05692]

If a system is running OSPF, and two of its interfaces are connected to the same physical subnet, it will attempt to form an adjacency with itself, resulting in excessive network traffic and CPU loading. [CSCdi05898]

When OSPF is configured to redistribute RIP routes and default with RIP metric received, OSPF default route metric does not change with RIP default route metric. [CSCdi06010]

Issuing the command **show ip route** may cause a reload to occur. [CSCdi06011]

If a BGP connection is started over a network that has secondary addresses, the underlying TCP connection will always use the primary address. [CSCdi06866]

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

Even when a summary LSA is Maxaged, it is still used in calculating the route resulting in invalid route.

Work around:

Clear the ip routing Table

"clear ip ro *" [CSCdi07335]

If BGP is configured with the **no synchronization** command, and a route is learned via internal BGP from a BGP peer which is not adjacent, BGP may subsequently advertise the route with a sub-optimal next hop. In this case, the next hop will point to the BGP speaker itself which will forward traffic correctly. This caveat only affects forwarding efficiency. This caveat was introduced in 9.0(2). [CSCdi07531]

After a topology change occurs, OSPF waits for 8 seconds before running SPF on the database in order to protect the router from link flapping. This delay is not absolutely necessary and will be removed in a future release. [CSCdi07562]

When there are multiple External LSAs for the default route (0.0.0.0) in OSPF domain, there is a possibility for the default route to disappear from the IP routing table. There is no workaround. The fix for this will be available in a future release. [CSCdi07576]

When an IP address used as the OSPF router id is removed from the router and used in a different router, there is the possibility of two OSPF routers using same router id, thereby causing an inconsistent OSPF database. The workaround is to remove the OSPF configuration and reenter it when changing OSPF router id's. [CSCdi07602]

If a router is configured for BGP and no BGP neighbors are active, either because they are not configured or because their BGP connections have not been established, then excessive memory utilization will occur. The workaround is to configure at least one BGP peer and to insure that it becomes active. This caveat was introduced in 9.0(2). [CSCdi07626]

In obscure circumstances which are not fully understood, a problem with a TCP connection underlying a BGP connection may cause the router to reload. [CSCdi07637]

Fails to redistribute Connectd nets into OSPF There is no further information available concerning this problem. [CSCdi07885]

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]

ISO CLNS

If IS-IS is configured on an interface and an ISH is received, an IIH is sent. This results in an increased frequency of sending IIH's. The correct behavior is for an IIH to be sent only when establishing new adjacencies or adjacencies in Init state. [CSCdi05098]

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

A system reload may occur if DECnet conversion is enabled, two ISO-IGRP processes are redistributing each others routes and the system receives a DECnet packet that it cannot route. [CSCdi05883]

DECnet Phase IV to Phase V conversion does not work when two ISO-IGRP processes are redistributing each other's routes. [CSCdi06087]

When running CLNS, a router would send out IS hellos even when there was no NET configured for the interface. The fix is to check that there is an NET configured before sending out the IS hellos. [CSCdi06104]

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]

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]

If you are running pre 9.0(2.1), there may be buffer loss problems when running CLNS over Frame Relay. [CSCdi07183]

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]

If there exists any IS-IS routers in a network that originates LSPs with an LSP number of non-zero, the destinations in that LSP will not be inserted into the routing table. The only workaround is for LSPID's of the form **xxxx.xxxx.xxxx.yy-zz**, **zz** must be **00**. [CSCdi07491]

If a ISO-IGRP route exists where a valid next-hop is used but there is no adjacency entry for the next-hop, the system may crash.

This situation is very rare. If a ISO-IGRP route exists where a valid next-hop is used but there is no adjacency entry for the next-hop, the system may crash.

This situation is very rare. [CSCdi07502]

The router may crash in very rare circumstances. If a single LSP exists in the AVL tree and it is deleted, it may dereference a NULL pointer. [CSCdi07683]

The **show clns routes** command can, under some circumstances, cause the router to reload. [CSCdi07710]

If CLNS is configured on an interface (either by the "clns enable" or "clns router iso-igrp ..." command), and IS-IS is not configured on the interface, received IS-IS packets will consume packets and not return buffers to the system. If CLNS is configured on an interface (either by the "clns enable" or "clns router iso-igrp ..." command), and IS-IS is not configured on the interface, received IS-IS packets will consume packets and not return buffers to the system. [CSCdi07758]

If an IS-IS link resides in a non-zero LSP number, it will be displayed as appearing in LSP number 0 as well. This creates the illusion that it is advertised more than once. This causes no problems with connectivity. [CSCdi07827]

The system may reload when the last IS-IS LSP is deleted from the link state database. This may happen when you are deconfiguring IS-IS from the system. [CSCdi07846]

A DECNET created adjacency in the OSI adjacency database will have a data-link address that is multicast. This results in multicast packets transmitted for each packet sent to the host. [CSCdi07939]

LAT

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

TCP/IP Host-Mode Services

For the IGS platform, IP crc errors may occur when packets are sent using tcp header compression over a serial line. [CSCdi04783]

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]
```

This is a random crash which happens due to accessing a deleted route. This crash is seen when an interface is shutdown and try to boot. [CSCdi06658]

The config command "no ip routing" only deletes the first of the defined static routes from the configuration, when in fact all of them should be deleted. [CSCdi07190]

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]

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
```

where 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

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]

TCP header compression over X.25 does not work in the initial release of 9.0(1). [CSCdi03839]

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]

If a Virtual Circuit is established in order to forward a packet, the packet may not be forwarded immediately on receipt of the CALL CONFIRM. [CSCdi07560]

The system does not reset the sequence counters used for the LMI keepalive information element when the LMI type is changed. This behavior occurs if the LMI type is changed (from Cisco to ANSI Annex D or vice versa) after the system has been in operation for some period of time. This behavior has no impact on operation but does not conform to the detail of the specification. [CSCdi07649]

XNS, Novell IPX, and Apollo Domain

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]

Novell RIP updates were sometimes sent with more than the maximum of 50 routes in a single

packet. Always enforce a limit of 50 networks in any single Novell RIP packet. [CSCdi06999]

9.0(1) Caveats/9.0(2) Modifications

This section describes possibly unexpected behavior by Release 9.0(1). Unless otherwise noted, these caveats apply to all 9.0 releases up to and including 9.0(1). For additional caveats applicable to Release 9.0(1), see the caveats sections for newer 9.0 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.0(2).

AppleTalk

When the system replies to AppleTalk ZIP GetNetInfo and GetLocalZones requests which originated on networks other than the network about which they request information (that is, requests which have been routed from distant networks), the router sets the source DDP address in the reply incorrectly, using a broadcast address. Furthermore, when a request is addressed to one of the system's interfaces, but received on a different interface, the information returned is taken from the receiving interface instead of from the interface to which the request is originated. In addition, the address of the receiving interface is always used as the DDP-layer source address of the reply packet; correct behavior would be to use the address to which the request was originally sent.

This behavior has no impact on ordinary AppleTalk applications, but may confuse some network management software. [CSCdi04809]

Serious errors reported by the AppleTalk code are not logged to the console or to any syslog server, regardless of the configured logging levels. [CSCdi04812]

If AppleTalk IPtalk is enabled, then disabled, a block of system memory will be "lost" and rendered unusable until the system is rebooted. The block involved is not large, and the behavior rarely if ever has operational impact. [CSCdi04821]

For Appletalk networks with macip service enabled, macip traffic will not be forwarded thru the router. This problem will only affect those networks where the macip packets must traverse the router acting as a macip server to reach another macip server. [CSCdi04909]

IPtalk does not function. [CSCdi04932]

When an AppleTalk ARP reply is received on a Token Ring interface, the sanity check that prevents entering multicast MAC addresses into the ARP table is done incorrectly; the least-significant bit of the first octet of the address is checked instead of the most-significant. This may result in the system accepting invalid AppleTalk ARP replies, or, usually more seriously, in its ignoring valid ones. This can be worked around by reconfiguring other nodes to use Token Ring MAC addresses which do not have the least significant bits set in their first octets. [CSCdi05137]

This problem will prevent cisco routers from properly interoperating with other AppleTalk implementations on FDDI media which do implement the March 8, 1991 Preliminary Proposal. [CSCdi05298]

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. [CSCdi05437]

This problem will prevent cisco routers from properly interoperating with other AppleTalk implementations on FDDI media which do implement the March 8, 1991 Preliminary Proposal. [CSCdi05464]

When a router running a software release supporting FDDITalk, ie, 9.0(2), is on the same FDDI ring as a router running a pre-FDDITalk release, the AppleTalk protocol will be restarted on the FDDI interface every time a pre-FDDITalk RTMP broadcast is heard on the ring.

Workaround: upgrade all routers running AppleTalk on the FDDI ring to 9.0(2) or do not upgrade any of the routers. [CSCdi06129]

Basic System Services

Permit the deletion and/or redefinition of NetBIOS access-lists used by access-expressions. Also, change the access-expression parameters "input" and "output" to "in" and "out" for consistency. [CSCdi04815]

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]

The protocol translation option of the IGS router software fails to properly initialize the allowed transport outputs to include X.25 PAD service. This will result in messages of the form "% pad connections not permitted from this terminal" when a user attempts to create a PAD connection. Outgoing PAD connections configured via **translate** commands will operate correctly. A workaround is to configure the virtual terminals on the IGS to explicitly include the PAD capability by using the command **transport output lat pad telnet**. [CSCdi05115]

On very heavily loaded systems, the CPU utilization percentages given by the **show process** and **show cpu** commands and the interface utilization percentages given by the **show interface** command may fail to decay properly or may be displayed as impossible values. [CSCdi05168]

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]

Enabling debugging for the OSPF protocol may result in a loss of neighbors. This is caused by the logging process running at too high a priority. Note, logging messages may now be delayed due to this change in behavior. [CSCdi05202]

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]

Under unusual circumstance when an SNMP packet is received some memory will be lost, over time this could use up all system memory. Two things must be true for this to happen; a bad community name is in the snmp request resulting in an authentication trap, and the snmp request must have over 14 variables in it. [CSCdi06309]

DECnet

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. [CSCdi04755]

If the DECnet Phase IV to Phase V conversion prefix greater than 11 octets is configured then the router may reload. [CSCdi05376]

The split horizon rule is applied to DECNET Phase IV routing updates sent on Frame Relay interfaces; no information on routes learned through a Frame Relay interface is included in outgoing updates transmitted through that interface. If each of two remote routers is connected by a PVC to the local router, but no PVC connects the remote routers themselves themselves, the two remote routers will be unable to communicate with one another via the local router. This problem may be avoided by providing full-mesh PVC connectivity among all the routers on a Frame Relay network. [CSCdi05827]

EXEC and Configuration Parser

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 **appletalk event-logging** is not removed from the configuration when AppleTalk routing is removed. Attempts to remove it after AppleTalk routing has been removed will fail because no AppleTalk commands can be executed when AppleTalk routing is not running. The workaround is to turn off AppleTalk event logging before disabling AppleTalk routing. [CSCdi04793]

Setup does not exit automatically when modem disconnect is detected. At this point the user must type control c to exit from setup. [CSCdi04940]

The command **show flash** is not currently supported on terminal servers and protocol translators. [CSCdi05506]

IBM Connectivity

When a LAN Manager query is received for a nonexistent station attached to a virtual ring results in the issuance of "SYS-2-SHARE" error messages. This behavior has no operational impact other than the issuance of the messages. [CSCdi04342]

Older IBM documentation used the values 516, 1470, 2052, 4472, 8144, 11454, and 17800 as the possible largest frame values for token ring. The current IBM documentation uses the values 516, 1500, 2052, 4472, 8144, 11407, and 17800.

The 9.0 software was changed to reflect these new values. However this makes it incompatible with 8.X releases since the 9.0 sw running with 8.X configurations will drop remote-peer configuration commands when largest frame sizes were configured as either 1470 or 11454 bytes.

To workaroud this problem change the 8.X configuration when using 9.0 software for remote peer commands to the new values if 1470 or 11454 were previously used as largest frame sizes. [CSCdi05036]

Lan Network Manager will not work with the 9.0(1) software release when the router is more than two hops away from the machine running Lan Network Manager. Therefore, the PC running Lan Network Manager should be located on a ring that is directly attached to a Cisco router. [CSCdi05073]

Lan Network Manager will not work with the 9.0(1) software release when the router is more than two hops away from the machine running Lan Network Manager. Therefore, the PC running Lan Network Manager should be located on a ring that is directly attached to a Cisco router. [CSCdi05105]

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]

Running Lan Manager with the older Netronix Token Ring card, if ring beaconing occurs the router may reload. [CSCdi05258]

If an LLC2 session is lost by the router sending a disconnect frame in an SDLLC LLC2/SDLC pair, the SDLLC state will not be reset to "disconnect". Therefore, the background task that tries to re-establish sessions for SDLLC sessions that use the **sdlc partner** command will never attempt to try the reestablishment, as the SDLLC state was still "connect". Effectively, that makes SDLLC useless for that link until the router is reloaded.

To determine if the router is exhibiting this incorrect behavior, exam the output of **show interface** command for the SDLLC interface. It will show an SDLC state of "disconnect" but an SDLLC state of "connect" for the affected SDLC address. To clear this state the router must be reloaded. [CSCdi05335]

&&IGNORE

No customer impact, merely removes a nuisance error message. [CSCdi05517]

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]

A system reload will occur if a Lan Manager workstation attempts to link to a router containing a single token ring port. [CSCdi05699]

LNM passwords must be exactly 8 characters in length. The correct behavior would allow for the password to be 6 to 8 characters in length. [CSCdi05892]

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 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]

Interfaces and Bridging

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]

For the OSPF protocol if a **redistribute ospf n metric n** command is issued. Then there is no way to remove that command. If a **no redistribute ospf** command followed by a **redistribute ospf n** is issued then all parameters are restored to their default values, not just the ones configured. [CSCdi04393]

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]

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]

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]

Packets received over the Ultraset 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]

In environments with very large BGP updates (as in NSFnet regional networks, which pass data for the entire IP catenet), it is possible for the BGP process to consume all the buffers in the system, and still be unable to continue because of insufficient additional buffers being available. This is manifested as a stoppage of all process-level network activity. [CSCdi04872]

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]

For the OSPF protocol, a administrative distance change on a routing process does not affect existing routes. [CSCdi04920]

If any of the following IP-only specific router subcommands are issued then the CLNS/IS-IS data structures become corrupted. **default-information, default-metric, distribute-list, metric, neighbor area, network, offset-list, passive-interface, timers, variance** This causes a system reload when if any if the the following CLNS/IS-IS command sequences are later issued: **router isis timers, basic** or **router iso-igrp, variance**. [CSCdi04936]

An ultranet interface configured for bridging accepts it's 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]

The **dialer fast-idle** command ignores parameters. [CSCdi05002]

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]

The **no ip routing** command does not stop IP routing processes. [CSCdi05157]

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]

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]

Specifying **ring-speed 4** actually results in **ring-speed 16** and vice versa. [CSCdi05224]

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]

Administrative distance is not properly respected for network 0.0.0.0. A new route with lesser distance but higher metric than the one in the routing table will be ignored. [CSCdi05399]

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]

When under high loads, the CSC-2R may enter the "initializing" state. If initialization is delayed because of excessive input load, the system may eventually time out and shut down its interface to the ring. This will result in an interface reset. [CSCdi05446]

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 EGP is used to generate a route for network 0.0.0.0, a route to 0.0.0.0 can be accepted from another EGP speaker, causing a default routing loop. [CSCdi05459]

Using local SRB under heavy load may cause SETFUNFAIL errors and a high rate of resets for both the CSC-2R and CSC-R16 Token Ring interfaces. [CSCdi05499]

The **sdlc-largest-frame** command for sdllc is written incorrectly in nv ram. This results in a parser error when an attempt is made to execute the command. The work-around is to enter the command manually each time the router is reload. Please note that this only applies if the value for sdlc-largest-frame has been changed from the default. [CSCdi05655]

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

No IP flash routing updates are sent (with any routing protocol) when an interface is administratively shut down. This may result in connected routers being slow to react to the loss of the newly shut-down path. [CSCdi05794]

On GS routers with no FDDI interfaces and no SBE16 token ring interfaces, bridging of maximum size packets will corrupt a couple of bytes of heap memory. This corruption will cause the 'show memory' command to display less free memory than expected and terminate prematurely with SMASHED BLOCK error messages. [CSCdi06229]

After a system has been operational for 24 days, the IGRP, RIP, HELLO and CHAOS routing processes will stop sending updates. The cessation will occur if the routing process has been running the entire time the system has been operational or if the process was manually started any time after system start up.

There is a workaround for IGRP. Assuming the user is not using non-default values for the IGRP timers, simply use the following router subcommand:

timers basic 90 270 280 630 1

The only value that helps the workaround case is setting the fifth parameter equal to one. The other values do not affect the problem and should be set according to the users wishes. The above example is the normal case.

A workaround does not exist for RIP, HELLO and CHAOS. [CSCdi06310]

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]

When using the **no neighbor** command to delete configured neighbors on non-broadcast multi-access networks such as X.25 for OSPF protocol the following error message is generated and no action is taken, OSPF: no interface specified. Work around is to remove the ospf process and recreate it. [CSCdi05291]

For the OSPF protocol, given that the designated router for a particular net has been disabled, the backup router is then promoted to designated router. This is correct behavior. However no backup router is selected to replace the backup router promoted to router. Although this behavior is incorrect, impact is minor resulting in slight increased traffic. [CSCdi05309]

When the **ospf neighbor** command issued with a local interface address, this causes a poll to be sent to itself. [CSCdi05586]

When an igmp route is overwritten with OSPF route, the old igmp route redistributed into OSPF domain is not flushed. [CSCdi05605]

ISO CLNS

When a router running ISO IS-IS and supporting level 1 routing is reconfigured to support only level 2 routing, it does not send notification that its level 1 links are no longer available, nor does it explicitly resign as level 1 designated router on subnets for which it has been elected such. [CSCdi04171]

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 relearned. [CSCdi04753]

ES-IS intermediate system hellos are sent on every CLNS interface, even when router (or the interface in question) is configured only for level 2 operation. [CSCdi04866]

CLNS dynamic routing protocols used over SMDS networks do not properly capture the SMDS source addresses from which updates were received. This makes it impossible to use dynamic routing protocols for CLNS over SMDS. This may be worked around by the use of static routing. [CSCdi04891]

After the **no router isis** command is issued, the **show timers** displays two sets of timers. [CSCdi04892]

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]

For the IS-IS protocol, the prefix route selector is not included when selecting a NET to delete. This may result in some NETs being incorrectly removed from the routing table. [CSCdi04934]

For the IS-IS protocol, the retransmission timer over serial lines is incorrectly set to 5000 seconds rather than 5. This causes unreliable delivery of IS-IS Link State Packets over serial links for the first 5000 seconds after the link is configured to run IS-IS. This may result in serious problems such

as routing loops. [CSCdi04943]

The IS-IS protocol will incorrectly advertise serial link adjacencies created by the ISO-IGRP or ES-IS protocol. This can produce unpredictable routing behaviour. [CSCdi04944]

Static neighbors are not added to ISO-IGRP level-1 routing table if they are entered before ISO-IGRP is enabled on an interface. [CSCdi04976]

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]

When more than one adjacency with the same system id exists in the adjacency database, and one of these adjacencies is deleted, the associated link in the LSP is not deleted. [CSCdi05067]

The **show clns route** command may cause the router to reload. [CSCdi05111]

If ISO IS-IS is used in environments with large numbers of link states (on the order of 50 or more), system reloads or other unexpected behavior may occur. [CSCdi05151]

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 **no redistribute static** command does not work for ISO-IGRP. [CSCdi05284]

The **show clns redirect** command may cause the router to pause indefinitely. [CSCdi05367]

In configurations that deploy DECNET IV to V conversion where the OSI backbone runs IS-IS, the DECNET Phase IV created adjacencies in the CLNS adjacency database are not inserted in the fastswitching cache. This causes slow switching to occur for these systems. [CSCdi05477]

ES-IS supplies NET in RD PDU for redirects to end-systems. The correct behavior is for this to occur only if redirecting the node to an intermediate-system. [CSCdi05674]

LAT

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

Certain LAT error messages do not give sufficient data to actually tell what is wrong. In particular, the "% Reach limit of struct" message didn't give any indication of which struct was involved. [CSCdi05178]

Local Services

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]

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]

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]

Under rare circumstances, it is believed to be possible for a proxy ARP reply to be processed in-

correctly, resulting in incorrect entries in the ARP table. These entries will give valid MAC addresses for incorrect IP addresses. This behavior has never actually been observed in the field, and should occur only when the interface on which the original proxy ARP reply is received undergoes an up-to-down state transition immediately after the packet arrives. [CSCdi05169]

Race condition can occur where packets are not transmitted because they get overwritten before being transferred to the transmit interface. [CSCdi05678]

TN3270

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 kermi, 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

Login strings do not work properly. If a connection is made to a host for which a login string has been defined, the login string is not sent, and a "bad login string" message is issued on the system console. There is no workaround. [CSCdi05791]

Uncategorized Items

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]

Wide-Area Networking

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]

The **show interface** and **show X25 vc** commands did not indicate when the window at packet level(x25) and/or frame level(lapb) was closed.

The **show interface** and **show X25 vc** commands have been modified to display this message "Window is closed" For the **show x25 vc** command the above message indicates the VC is packet is level flow controlled and the window is closed.

For the **show interface** command the above message indicates the interface is frame level flow controlled and the window is close. [CSCdi04981]

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]

Attempting to issue a **clear x25-vc** command to remove idle X.25 SVCs may cause the router to reload. [CSCdi05037]

Issuing the command **no dialer fast-idle** incorrectly resets the **dialer idle-timeout** instead of the **dialer fast-idle timeout**. [CSCdi05041]

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 one X25 facility is configured, and the **x25 rpoa wan** command is one of those facilities, then disabling the rpoa facility may cause the router to reload. [CSCdi05219]

When configured for ANSI ANNEX D frame relay, the router incorrectly uses dlci 1023. This causes the line protocol to be declared down. The correct behavior is to use dlci 0.

The workaround is to disable keepalives on a particular interface. [CSCdi05280]

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]

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]

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]

If two **no dialer commands** are issued in a row, there is a high probability that the router may reload. [CSCdi05594]

On the IGS platform only, transparent bridging over Frame Relay does not work. [CSCdi08495] [CSCdi05664]

XNS, Novell IPX, and Apollo Domain

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 is possible for the router to reload. This occurs infrequently as packets corrupted in this manner are fairly rare. [CSCdi04921]

When a change in the XNS, Novell IPX, or Apollo Domain routing table triggers a flash routing update, information about the networks whose status has not changed is included, while information about the networks whose status has changed is omitted. This is exactly opposite the correct behavior. Because information about routing changes is not propagated correctly in flash updates, routing convergence may be slower than would otherwise be expected. In addition, in large networks with many unstable links, flash update traffic may consume enough bandwidth and/or router CPU resources to have strong adverse effects on network performance. [CSCdi04959]

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]

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]

XRemote

XDM will not allow a user to abort a session being set up (with the ^x sequence) once a host has been selected. This can cause the session to hang if the TCP connection to actually start the session is never made. [CSCdi05184]