9.21(8) Caveats

This section describes possibly unexpected behavior by Release 9.21(8). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(8).

Basic System Services

Under certain circumstances, the order of "ntp peer" and "ntp server" statements in the system configuration may change. There is no workaround to this problem. [CSCdi18741]

The **show version** command does not label the booted image correctly when it was booted from ROM. The command shows "System image file is unknown, booted via" instead of the expected "Running default software." [CSCdi23575]

The sending or reception of telnet sub-negotiation strings is not show by **debug telnet** or individual connection debugging. [CSCdi25822]

If no tacacs-server hosts are configured, login attempts on lines configured for tacacs login will be rejected. This is different from historical behavior. [CSCdi28420]

Communication Server

Under rare circumstances, a cisco-500 line configured for hardware flow control may become locked. This can prevent the next user of the line seeing any response back from the cisco-500. [CSCdi27906]

DECnet

A router receiving a MOP connection request through its serial port for one of its LAN port addresses responds with the LAN port's burnt-in address instead of the actual hardware address. If the requesting host uses the DECnet-style MAC address of the router in the request packet, the host will not recognize the response packet sent by the router because it sees a different address in the "source" field. This causes the requesting host to time out on the connect request. [CSCdi26991]

EXEC and Configuration Parser

Problem: CiscoWorks Autoinstall overides the neighbor device configuration file in the following manner: ip forward-protocol udp 285 -- Original configuration of the neighbor device ip forward-protocol udp -- Ciscoworks sticks this in the config file.

This effectively nullifies the first command.

Solution: Once the new device is up and running, the user could telnet to the neighbor device and manually get rid of the command "ip forward-protocol udp". [CSCdi22012]

IBM Connectivity

On a 7000 with SRB configuration, when a packet with path tracing(e.g: decnet pings) passes thourgh a CTR interface, the following message will be displayed on the router console:

%LINK-3-BADMACREG : Interface, non-existant MACADDR registry for link 0 -Process= "*Sched*", ipl = x -Traceback = hhhhhhhh hhhhhhhh hhhhhhhh hhhhhhhh hhhhhhhh hhhhhhhh [CSCdi16761]

When IPX routing is enabled on a Token Ring interface and there is a source-route bridge network behind the ring, a **multiring ipx all** command is used to cache the RIF in the router. During normal operation all is well. But when a station is moved from one ring to another ring (for example, from 0B8 to 0B1), the station cannot reach the server. Looking at the RIF cache on the AGS+, it is fine. However, analyzing the frames with a Sniffer, we can see the "create connection request" from the station with a good RIF field but the answer from the AGS+ shows the previous RIF (the RIF before the station was moved). The workaround is to disable the IPX route cache or to clear the IPX cache when a station is moved. This is a general problem with all routed protocols. The RIF code does not inform the routing protocols when an entry in the table changes. Therefore, the cache entries become invalid. [CSCdi17099]

In a router with multiple STUN interfaces, Local-ack can get out of sync if one of the STUN interface is reset. [CSCdi17806]

When running RSRB with local-ack, certain topologies may cause the input queue to fill up. The result is that the interface with the filled input queue will no longer pass traffic. [CSCdi22676]

Interfaces and Bridging

Bridging of appletalk, clns and decnet over SMDSW does not work. This is due to multicast datagrams dropped when trying to flood such datagrams received from the SMDS network (clns and decnet), and incorrect processing of bridged datagrams (appletalk). [CSCdi10490]

The 4000 FDDI interface may not be able to receive traffic smoothly under very heavy loads (greater than 14 kpps), in which case the incoming traffic will only be accepted in small bursts every 10 seconds, until the load falls below that critical threshold. [CSCdi10848]

On Cisco 7000 systems with a Serial Interface Processor (SIP) installed (Not the FSIP), removing or replacing an EIP, TRIP, or other interface cards may cause the router to execute a system reload. Replacing the SIP with an FSIP via the available free upgrade program should eliminate this problem. Call 1-800-553-NETS if you have not yet received the upgrade. [CSCdi13319]

During the process of initializing a token ring interface the keepalive process on other interfaces may stop. This condition may result in other interfaces on the system experiencing resets. [CSCdi13654]

Normally, the fddi interface has a keepalive value set at 10 seconds. When the 9.22 OS version is booted from the network, these keepalives change to not set in the NV-RAM. This behaviour has been modified in later versions for correct operation. [CSCdi14647]

While netbooting, a couple of error messages show up for FDDI interface, such as **CBUS-3-INITERR** with an error code of (0000) [CSCdi14648]

AGS+ systems containing CSC-C2FCIT FDDI ciscoBus Interface cards configured for transparent bridging will experience dropped packets when packet sizes are smaller than 20 bytes. Packets of 20 bytes or larger will be bridged properly. [CSCdi15617]

On systems configured to support the DEC spanning-tree protocol, keepalives from one media type are not being filtered out of traffic pertaining to other media types also configured for DEC spanning-tree. For example, protocol analyzer traces take [CSCdi15729]

On systems configured for autonomous bridging on ciscoBus II-resident interfaces, if the bridge-group is removed, and standard bridging is configured on those interfaces, ciscoBus II-type bridging will instead be configured for those interfaces. T [CSCdi16310]

When a configuration change occurs that causes the Token Ring interface to initialize, there is a delay between the time the command is entered and the initialization process begins. The Token Ring initialization also impedes other process-level functions. [CSCdi16454]

3204 serial on v.35 resets interface every 60 seconds. [CSCdi19239]

The AGS+ with a CBus I controller which executes code image cctl132-3 and FDDI controller micro-code image 2.1 generates a "CBus_init" error when AppleTalk routing is enabled. This error will not occur when AppleTalk routing is not enabled.Upgradin [CSCdi19940]

A 4000 router may restart from bus error **HD64570_MSCI_RESET** if 4T NIM is down rev. Upgrade the 4T to rev 3 to resolve. [CSCdi21975]

Texas Instruments has stopped production of the TMS380C16 and switched to the TMS380C26 Token Ring chip. The new chip disables the SRA (source router accelerator chip) when the TMS380C26 chip is in promiscous mode. This means that the Token Ring interface can no longer support both source-route bridging and transparent bridging on the same interface. Whenever transparent bridging is turned on, the source route bridging ceases to function. The TMS380C26

chip is used in Cisco 2500, Cisco 4000, and Cisco 4500 routers. [CSCdi22815]

On Async Interfaces running SLIP, the interface input error count is off by one. [CSCdi26085]

IP Routing Protocols

The output of the command **show ip protocol** in regards to OSPF can be confusing. The fields displayed do not apply to the OSPF protocol. This will be corrected in 9.21. [CSCdi10880]

To maintain syntacs common through out the display the designations referring to autonomous systems is replaced with designations which refer to OSPF process-id. For Example amethyst#sh ip ospf ? 1)[] routing process information 2)amethyst#sh ip ospf database OSPF Router with id (180.180.177.2) (Autonomous system 345) [CSCdi13868]

The system does not disallow the assignment of the same IP address to multiple X.25 interfaces on the same system. [CSCdi15734]

Router does not establish adjacencies with the other router on Point-to-point frame-relay subinterface. The work around is to specify 'ip ospf network broadcast' under each serial subinterface.get established. [CSCdi16443]

"show ip bgp summary" displays "0:00:00" as the current up/down time if the state has not changed in 4 weeks. [CSCdi17438]

Starting from 9.1, the intended default for **redistribute ospf** command is to redistribute internal route only into EGP, but it is not enforced. This fix solves the problem for 10.0 and later versions. [CSCdi23229]

Starting in 9.21, a route-map command **match route-type external** can be used to indicate the desire to redistribute external route, including BGP, EIGRP and OSPF type 1/2 external route. The problem is that there is no way to specify the redistribution of only OSPF type 1 and OSPF type 2 route seperately. This fix solves the problem for 10.0 and later version. Now, the users can use **match route-type external type-1** to specify OSPF type 1 route only and **match route-type external type-2** to specify OSPF type 2 route only. The **match route-type external** is still used to indicate all types of external route as above. [CSCdi23279]

Starting from 10.0, the **redistribute ospf** *1* route-map *map* command with **match route-type external** clause in the route-map for EGP will result in nothing being redistributed. The only way to specify redistributing external route is through the use of the **match external for the redistribute** command. The fix correct the problem for 10.0 and later versions. [CSCdi23291]

On an interface with secondary addresses the router replies to ARP requests with the primary address only. [CSCdi25069]

[CSCdi25947]

ISO CLNS

First of all, interface static routes are always up, regardless of the interface state. The fake adjacency created in this case should probably vary based on the interface state. Secondly, next-hop static routes do not get added/withdrawn from the IS-IS database if they point through an adjacency learned through IS-IS rather than ES-IS and the adjacency goes up or down. [CSCdi19594]

When a self-generated LSP is received from the net and it appears newer than router's own. If anything other than LSP fragment 1 is received this way, IS-IS attempts to regenerate fragment 1 with this new sequence number. This may cause the IS-IS to regenerate LSP with wrong sequence number. [CSCdi20806]

If the DR generates a new set of LSP fragments and if, as a result of there being less info to stuff into the LSP, there are less LSP fragments generated, the old high-numbered fragments will not be flushed. This may cause IS-IS not to flush DR LSPs. [CSCdi20807]

In 9.21 and later versions, DECnet ping fails over OSI cloud if the destination is part of the OSI

cloud itself. In more general term, the problem occur exactly when the router receive a CLNS packet with a selector 0x20 (the value denoting NSP) destinated to it. No workaround is available. However, since it is seldom to have data packet destinated to a router (as in the case for ping), the impact of the problem is small. This fix solves this problem for 10.0 and later versions. [CSCdi22736]

ISIS should not send ISH over non-point-to-point network if it is configured to route IP only using the **ip router isis** interface command. The system works correctly under normal condition. However, if the same interface is turned on for routing CLNS using either **clns enable** or **clns router isis**, the system will start sending ISH over the interface and it continue to do so even CLNS routing is then turned off for the interface. This generates additional traffic that is not necessary and there is no way to stop it. This fix solves the problem for 10.0 and later versions. [CSCdi22765]

Starting from 9.1, the **no clns router iso-igrp** command can clns-disable the interface even though some other level 2 iso-igrp's are still configured over the interface. This fix solves the problem for 10.0 and later versions. [CSCdi23008]

ISIS LSPs may be seen in the database that appear to be corrupted; typically they have options like "code:0 length:253". This is a cosmetic problem only. There is no workaround to this problem. [CSCdi24977]

Novell IPX, XNS, and Apollo Domain

Changing ipx network number encapsulation sap to the current ipx network number breaks it into two seperate commands. We should display the command in one line. This is a cosmetic bug and does not prevent ipx routing from functioning. The display in write term has been fixed. We display it on one line. This fix was done when we applied CSCdi15627 as of 009.021(000.134). [CSCdi15585]

TCP/IP Host-Mode Services

User rlogin connections are not shown during a **show tcp** command. [CSCdi16427]

TN3270

$IGNORE

TN3270 connections to some types of hosts will come up as line mode if the "return" key is pressed more than once between issuing the command and completing the TCP connection. A suggested workaround is to use a **login-string** containing a keyboard-locking pause of several seconds (**login-string myhost /%5w/**) [CSCdi19385]

A ttycap entry with two colons in a row or a colon at the end of one line and the beginning of the next will fail to be interpreted correctly. [CSCdi27822]

Wide-Area Networking

The Release 9.14 code merge added Token Ring and FDDI fast switching to Release 9.21; however, PPP was forgotten (not in 9.1, but already in 9.21). The PPP encapsulation is added to HDLC, Frame Relay, and SMDS as a valid encapsulation. [CSCdi12076]

When a line changes state, a DLCI which was created dynamically is cleared and drops any state that may have been configured after the DLCI was created. [CSCdi15512]

Under unusual circumstances X.25 addresses may be improperly encoded. [CSCdi16601]

With HDLC encapsulation on a serial interface we now compare the keepalive time with the configured value, and display an error message if the times don't match. [CSCdi16626]

The **dialer-list x LIST y** required for dialer interfaces to work. PPP looks for a "real" interface with appropriate PPP state. [CSCdi17733]

When dialing for the first time after a reload, the following message may be displayed when the **debug dialer** flag is on: Dialer1: No free dialer - starting fast idle timer. This has no effect on the operations of DDR. [CSCdi18014]

AppleTalk zones are not set when running over a Frame Relay point-to-point subinterface. This problem is caused by incorrect DLCI mapping by the router. [CSCdi18233]

Frame Relay Multicast with CLNS not working correctly. No other information available. [CSCdi21301]

Dial-on-demand PPP connections to any router sending an IPCP request with an IP address of 0.0.0.0, such as as Wellfleet router, do not work. The workaround is to have the non-Cisco router propose a valid IP address in its IPCP packet. [CSCdi22160]

The command **Show frame-relay map** incorrectly reports the LMI type ANSI as CISCO. This has no effect on the operation of the ANSI LMI. [CSCdi22669]

The **dialer load-threshold** command is changed so that the parameter is not a number between 1 and 255 any more, but a percentage between 1 and 100. Configuring 100 is equivalent to configuring 255 in previous versions. [CSCdi23561]

On async lines, the cisco's implementation of PPP doesn't ignore spurious unescaped bytes received on the line. [CSCdi23988]

When typing **no smds static-map ...**, and the map is not configured, an error message is printed with a bogus network address. [CSCdi24672]

Adding more than one **dialer-list l protocol p permit/deny** configuration commands with the same list number, protocol and permit/deny clause results in as many lines in the configuration file, instead of just one. [CSCdi24793]

Enclosure # 2. Trial Modification of the second record. [CSCdi24817]

To use decnet over DDR, static maps for **decnet-router-l1** are required on top of static maps for **decnet**. This extra configuration should not be required, only static maps for **decnet** should be necessary. [CSCdi24862]

In some instances, when a Frame Relay subinterface with an inactive DLCI has been administratively shut down by a user, it may exit the shutdown state and return to the active state even though the DLCI is still in an inactive state. [CSCdi25156]

There is no form of modem control that offers the capabilities of **modem cts-required** or **modem callout** that also allows simultaneous HW flow control. [CSCdi26270]

No release-note needed. [CSCdi27075]

Frame relay counters are not updated when fast switching. [CSCdi27509]

When a 4000 or a 4500 places or receives a call on a serial line, the line goes up, then down, and then up again, instead of going and staying up. The router then believes that the call is an incoming call. This happens only with the Hitachi HD64570 serial controller; the Mostek MK5025 works fine. [CSCdi27742]

When removing dialer maps from a BRI configuration, the router may reload. To work around this problem, shutdown the interface before removing a map. [CSCdi28180]

It is possible to fool the X.25 software into believing an an X.25 SVC is a PVC if the interface is rapdily shut down and brought back up. [CSCdi29850]

If, under rare and poorly understood circumstances, the router initiates an XOT connection, sends a Call packet, and the remote XOT host violates the protocol by returning a Call packet instead of a Call Confirm, the router will reload at some later time. [CSCdi30338]

9.21(7) Caveats/9.21(8) Modifications

This section describes possibly unexpected behavior by Release 9.21(7). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(7). For additional caveats applicable to Release 9.21(7), see the caveats sections for newer 9.21 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.21(8).

9.21(6) Caveats/9.21(7) Modifications

This section describes possibly unexpected behavior by Release 9.21(6). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(6). For additional caveats applicable to Release 9.21(6), see the caveats sections for newer 9.21 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.21(7).

AppleTalk

AppleTalk ports can get stuck in the restart state when system uptime is greater than 24.85 days. There is no workaround; you must reload the system. [CSCdi25482]

Global Appletalk ARP commands have a side effect of changing the router ID number for Apple-Talk Enhanced IGRP. There is no workaround. [CSCdi25786]

The exec commands **show appletalk eigrp events and show appletalk eigrp sia-events. These commands will display the past 500 AppleTalk/EIGRP events which have occured on a router.**

The global configuration command **appletalk rtmp gc-interval has been added to allow users to configure longer intervals between AppleTalk RTMP route garbage collections. Increasing the RTMP garbage collection interval can decrease the CPU load on a router running AppleTalk RTMP which is experiencing route flapping, but users should be aware that increasing the garbage collection intervals results in the AppleTalk RTMP process keeping process memory which could be returned to the system free list immediately for a longer time. [CSCdi25924]**

If AppleTalk is started after the router has been up for more than three weeks, RTMP updates will not be sent out of the router.

The workaround to this is to enable AppleTalk before the router has been up for three weeks, or to reboot the router before enabling AppleTalk. [CSCdi26137]

When an AppleTalk distribution list is defined with at least one zone entry and no network entries, it is necessary to include **access-list** *number***permit other-access** as part of the access-list. Otherwise, improper filtering of routing updates occurs. [CSCdi26233]

ZIP GetNetInfo Request packets that contain the zonename, '*', are not correctly handled. Some printers generate this type of request when starting. There is no workaround. [CSCdi26491]

Basic System Services

NTP "master" mode does not work; the system will fail to declare itself synchronized, and thus no other system will synchronize to it. There is no workaround. [CSCdi27148]

Communication Server

The communication server leaks memory if the **nohangup** keyword of the **username** global configuration command is used to define special username entries. [CSCdi25520]

Interfaces and Bridging

When TCP/IP routing is enabled along with transparent bridging on the same interface, some SNAP encapsulated TCP/IP packets with destinations on the same network segment may be bridged to other networks. [CSCdi23944]

Extended bridging access lists are not evaluated correctly for flooded packets when the access list is applied on output from an interface. [CSCdi24778]

When multiple FDDI cards are present in the router, the interfaces in the lower slot positions may lose their downstream neighbors. [CSCdi25764]

Hitachi based serial ports may not transmit under severe load, resulting from the under-run inter-

rupt not being properly enabled [CSCdi26209]

This bug was found in the priority packet path(Eh: keepalives, bpdus etc). Holdq_enqueue can fail also due to the lack of available Q elements, in addition to the normal case when the Q becomes full. In such a case, make sure if a tail is present before unqueueing it to accomodate the current priority packet. If there is no tail, just flag failure.

The crash occured since there was no check for the valid tail. [CSCdi26417]

IP Routing Protocols

a summary route is advertised on an unnumbered serial when EIGRP is configured on a single sub-netted major network. Use **no auto-summary** as a workaround. [CSCdi25562]

Novell IPX, XNS, and Apollo Domain

The IPX Enhanced IGRP **distribute-list** command allows standard access lists only (access lists whose numbers are 800 through 899) only. It should also allow extended access lists (numbers from 900 through 999). [CSCdi25895]

when using IPX SAP filter using wildcard character '*', the last character before '*' is ingored. For example, given the access-list access-list 1000 permit -1 0 SAN* Server SAM should be de-nied, but it is accepted. [CSCdi27294]

The ipx watchdog-spoof command is written to non volatile memory before the dialer commands are written, upon a reload the system will complain about DDR not being enabled and will not en-able watchdog spoofing. Instead of enforcing watchdog spoofing on dialer configured interfaces allow spoofing on all serial or dialer interfaces. [CSCdi27326]

TCP/IP Host-Mode Services

When the sequence number for a TCP connection grows so large that the right edge of the window rolls over to zero, the usable window size calculation fails to calculate the correct usable window size. [CSCdi27537]

VINES

The VINES RIF cache becomes corrupted when an end station does an all routes broadcast/non-broadcast return. The problem is that the router returns a corrupt RIF to the end station. [CSCdi23239]

When the **vines serverless broadcast** command is configured in a redundant topology and all oth-er router interfaces are configured with the **vines serverless** command, a broadcast storm results. [CSCdi25597]

Enabling the 'vines decimal-addresses' command should affect all printing of vines addresses. It does not currently affect the printing of access lists, meaning that access lists can not be written to NVram and read back. This also affects a couple of debugging statements. [CSCdi25843]

When fast switching VINES over a source-route bridged Token Ring network, the router does not build its fast-switching cache entries properly. This prevents communication with stations that are across a bridge from the router. The workaround is to disable fast switching on the Token Ring interface. [CSCdi26288]

When a neighboring system changes MAC addresses between RTP or SRTP routing updates, the system may unexpectedly halt. A neighboring system may change MAC addresses for any number of reasons: swapped interfaces, started up Decnet, or **mac-address** interface command was used. Issuing **clear vines neighbor** on the system can prevent the system halt. [CSCdi27038]

Wide-Area Networking

Dynamic PVC's can lose track of the subinterface they are assigned to. If you divide a frame-relay interface into subinterfaces then it is possible for the router to think the PVC is used on serial 0, although it should be assigned to interface 0.1. Deleting/adding the map entry puts the PVC back

into the right subinterface. This problem can occur if the frame-relay network reports the DLCI deleted and then adds it again. [CSCdi16677]

The X.25 software typically does not encode address or facility information in a Call Accepted/Call Connected packet, which some X.25 equipment rejects with a "packet too short" diagnostic (38). [CSCdi21201]

Locally switched X.25 Calls will generate a %X25-3-SPURD1 error after a long period of time (typically 3 or more hours after connecting). [CSCdi24989]

Routing by NSAP (for CMNS) doesn't work. [CSCdi25326]

The Calling or Called Address Extension facility is formatted improperly in the "debug x25" output. [CSCdi25529]

The Frame Relay broadcast queue might exhibit drops under high broadcast volume. There will be an increase in "buffer element" misses at the same time the drops happen. [CSCdi25707]

The X.25 interface parameter th has an upper limit of the configured input window size; this is too restrictive because SVCs can negotiate larger window sizes. [CSCdi26730]

9.21(5) Caveats/9.21(6) Modifications

This section describes possibly unexpected behavior by Release 9.21(5). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(5). For additional caveats applicable to Release 9.21(5), see the caveats sections for newer 9.21 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.21(6).

AppleTalk

The system may halt unexpectedly when the command **show appletalk cache** is used. The behavior may be avoided by disabling screen pausing by issuing a **terminal lenth 0**. [CSCdi22893]

Multi-packet Extended ZIP Replies for a single network number are not correctly handled by the system. Multi-packet responses occur whenever the zonelist is too large to fit into a single packet. If the network is directly connected, this problem can be avoided by using a non-discovery mode configuration. [CSCdi22970]

This patch adds the command **appletalk eigrp active-timer to the AppleTalk implementation of EIGRP.**

This command is used to adjust or disable the amount of timeout period for EIGRP routes which are in the "active" state. [CSCdi23198]

The system can take an inordinate amount of time acquiring the zonelists for AppleTalk routes. This is especially a problem in very large AppleTalk internets. There is no workaround. The entire zone table should eventually be acquired. [CSCdi23874]

This DDT will add the new command **appletalk rtmp jitter which will allow a user to introduce "jitter" into the RTMP update interval.**

Jitter is useful to avoid convergence of RTMP routine update events in a large network after a long period of time. In large networks, especially "spoke-and-hub" topologies, it is possible that RTMP routing update events will start to synchronize in such a way that rather than routing updates being sent at various points in time during a 10-second interval by various routers in the network, most of the routers in the network will send their RTMP routing updates at the same time as most other routers in the network. This results in heavy packet loads on routers which have many neighbors in the routing topology.

This condition can be avoided by configuring the routers which are most central in the network topology with a jitter of 20 to 30 percent; ie, RTMP updates will be sent anywhere from 7 to 10 seconds after the last routing update has been sent, rather than every 10 seconds. [CSCdi24959]

CSCdi25131 corrects a mis-printing of the AppleTalk composite route metric whereby very large route metrics would be printed as negative numbers.

CSCdi25131 also adds the following commands:

clear apple interface

This command will clear and reset the current state of an AppleTalk interface without resetting the underlying hardware interface. If you want to reset or restart the underlying interface or protocol, you should use the command **clear interface.**

clear apple route-cache

This command will clear the AppleTalk fast-switching cache. Previously, the command **clear apple arp was used to clear both the AppleTalk ARP and fast-switching caches simultaneously. [CSCdi25131]**

Basic System Services

Rlogin support was removed from software images not containing the Communication Server subsystem beginning in release 9.21, but **rlogin** still appears (and is parsed) as a valid, but non-functional, command. [CSCdi23256]

The communication server does not do Telnet if the **/line** or **/stream** switch is used with **rlogin**. The communication server also does not fall back to Telnet protocol if the destination port is anything other than RLOGIN in the **rlogin** command. [CSCdi24813]

The Terminal Server crashes if an IP address is specified for network connection to a remote host, at the CS prompt or in the **Connect IP-addr** command. The crash would occur if 'rlogin' is the preferred transport protocol for outbound network connections on the user's line. [CSCdi24931]

The extensions of **Show Buffers** is unnecessarily available to non-privileged users. [CSCdi24956]

Communication Server

In 9.21, the changes to the "host" IP module apparently cause us to try Proxy ARP for a destination BEFORE trying to use the default gateway. This worked correctly in 9.1 (used default-gateway first.) [CSCdi14799]

The use of Extended TACACS, in conjuction with **tacacs-server last-resort password** can cause the ethernet input queue to fill up and the box to pause indefinately if repsonses to the TACACS queries are received WHILE the last-resort password is being entered. [CSCdi18919]

Cisco-500 comm servers will reload if an IP packet whose desintation IP address has 90 as the second octet is received on the ethernet. [CSCdi23494]

A router or communication server can spontaneously reload while attempting to hang up a line configured with the **autohangup** command once the last network connection on the line is closed. This crash happens only if the last connection was resumed using the **resume** EXEC command. [CSCdi24025]

DECnet

A router that has been configured as a Level 1 router should not send out Level 2 routing updates. [CSCdi20884]

EXEC and Configuration Parser

The sub-command **set level level-1** issued after a **route-map** command is rejected as an ambiguous command. [CSCdi23770]

The interface subcommand **access-expression [in|out]** *expression* is written to configuration memory as a filter for both inbound and outbound packets. [CSCdi24000]

The old form of the **frame-relay lmi-type ansi** interface subcommand, **frame-relay lmi-type Annex D** is not accepted by the parser. This can cause an interface to use the incorrect form of LMI after an upgrade from a software version earlier than version 9.21. [CSCdi24881]

Even if configured with **exec-character-bits 8**, user typin to the exec prompt is echoed and inter-pretted with the 8th bit cleared. [CSCdi25148]

Interfaces and Bridging

On Cisco 2502 and Cisco 2504 routers, IP and IPX packets of length 920 to 1050 bytes being rout-ed from Token Ring to serial interfaces may be corrupted. The workaround is to disable fast switching on the serial interface. [CSCdi19480]

MTU larger than 18000 on the low-end reloads the box. [CSCdi19751]

The command **show interface** now displays the signal lead states in the same way as the 7000 se-ries router does. Supported platforms include routers with the Hitachi HD64570 based serial ports; including the 3X04, 2500, and 4T NIM's. [CSCdi23344]

The command **show controller serial X** on a 4000 with 4t NIM card does not display correct in-formation. [CSCdi23470]

Under high process-switched traffic load low-end platforms may build up the input queue counter. [CSCdi24497]

If a modem hangs up while the cisco is in the process of implementing a user's **slip** or **ppp** com-mand, the cisco can fail to properly reinitialize modem control on that line, causing the line to be unusable until cleared, and appear "dead" to subsequent dialin users. This is most likely to happen if an extended tacacs server is delayed in responding to the xslipaddr or xslipon messages. [CSCdi24676]

IP Routing Protocols

When load-balancing IP traffic over multiple equal-cost paths, the system's routing table might reach an inconsistent state, leading to a system reload. Before the inconsistent state is reached, the system must have three or four equal-cost paths for a particular route. A routing update must then be received that causes the system to replace those paths with fewer (but still more than one), better metric paths. This route must then become used for further locally generated traffic. This problem is most likely to be seen after an interface flap (that is, after an interface's line state goes from up to down to up again) in an environment where there are redundant, but not symmetric, intercon-nections between routers. The problem also seems more likely in FDDI environments, where in-terfaces flap before fully coming up. These flaps can result in multiple back-to-back routing table changes. [CSCdi20674]

Enabling the Hewlett-Packard IP Probe protocol via the **ip probe proxy** command does not cor-rectly enable the protocol. There is no workaround for this behavior. [CSCdi23909]

This bug is introduced in 9.21. It happens when tthe router has RIP running on interface that use the **ip unnumbered** *numbered-interface* command. If *numbered-interface* has the **ip broadcast** command configured, the peer router will not get the RIP update. The workaround is to remove the **ip broadcast** command from the *numbered-interface*.

This fix provides a complete solution for 9.21 and later versions so that the **ip broadcast** will not cause the problem again. [CSCdi24719]

ISO CLNS

The problem will occur for serial, non point-to-point interfaces in general.

A --------/----------B ------------------ C

When the 'remote' router (i.e. A) comes up, the 'middle' router (B) should send a routing update immediately, and this is not happening. The update gets sent after 15 minutes. [CSCdi17808]

When changing the encapsulation on a serial link from a point-to-point mode (such as HDLC or PPP) to a "cloud" mode (such as SMDS or Frame Relay), IS-IS routing fails to consider the inter-face as broadcast media. Because of this, the CSNP will not be exchanged and hence the database

will not be synchronized. To work around this problem, unconfigure and configure ISIS after changing the encapsulation type. [CSCdi23691]

When attempting to configure **clns is-neighbors** on subinterfaces, the neighbor address is not saved when written to configuration memory. The partial configuration produces error messages and are ignored upon reboot. [CSCdi23750]

Novell IPX, XNS, and Apollo Domain

The interface configuration command **ipx watchdog-spoof** fails to properly enable Novell watchdog timer spoofing. There is no workaround for this behavior. [CSCdi23324]

The **rsup-only** keyword of the **ipx sap-incremental** command cannot be used on subinterfaces. [CSCdi24492]

Protocol Translation

The password is echoed and the username is double echoed if the "login" option is specified in a **translate** command for TCP to PAD connections. Remote telnet option negotiation should happen before the login sequence. [CSCdi22864]

Removing a **translate** command from the configuration can cause other translations using the same inbound ip address to stop working. A workaround is to configure the remaining translations again, e.g. by doing **write memory** and **config memory** . [CSCdi23621]

In LAT to PAD (X25) translated sessions, a CTRL-S followed by the entry of any character can sometimes cause a continuous stream of empty LAT messages, causing a session disconnect. [CSCdi24491]

TCP/IP Host-Mode Services

Under rare circumstances, an opening TCP connection can get stuck in CLOSEWAIT state. This can also result in a STUN peer session getting stuck in an OPENING state at the same time. [CSCdi23455]

If the **printer** option is applied to a translation from an ip address also used for other non-printer translations, connections to those other translations may fail. A workaround is to ensure that **printer** is used on either all or none of the translations sharing a "from" IP address. [CSCdi23757]

VINES

The VINES routing code was building neighbor entries based upon the first RTP packet seen. If the first packet happens to be an SRTP packet, then a neighbor entry is built from invalid data. [9.21 doesn't understand the SRTP frame format.] This bad neighbor entry can cause a chain of events that leads to a router crash. The solution is to make 9.21 filter out all SRTP frames, making sure that bad neighbor entries never occur. [CSCdi22826]

The router responds to packets sent to a non-existent client, if the client address is based upon the router's address. [CSCdi23085]

It is possible, but very unlikely, for two router to both get out of sequence on an IPC connection. This has only been seen at cisco when developing new features. If this occurs, issue the commands 'show vines ipc' and 'clear vines ipc ', where is the number of the open IPC Connection as given by the show command. [CSCdi23169]

After a sub-interface has been deleted, it should no longer appear in 'show vines' commands. [CSCdi23225]

This is a preventative change to the router software. When VINES split horizon is disabled on an interface, the default for redirect messages is changed so that they will never be sent. (This is the same thing as issuing the command 'vines redirect-interval 0'.) This should prevent customers from having reachability problems in partial mesh topologies. [CSCdi23607]

The VINES 'send' command does not work with early versions of the cisco 9.21 or 10.0 releases.

[CSCdi24505]
Routing VINES over X.25 links might cause the router to unexpectedly reload. [CSCdi24728]
This ddts greatly enhances the usability of several vines debugging commands through access lists. The new form of these commands are:
debug vines packet [ ] debug vines route [ ] [verbose] debug vines table [ ]
Number is an optional argument, and is an access list in the range of 201 to 300. For the first two commands, if the access list is supplied it will be used to filter debugging based upon the source address in a packet. For the last commands,if the access list is supplied it will be used to filter debugging based upon address in the router's tables. The 'debug vines route' command now only displays the presence of routing messages. Use of the 'verbose' argument will also display the contents of routing updates. [CSCdi25004]
The redirect logic does not correctly delete non-optimal routes when it installs a new optimal route. This does not cause an operational problem as the non-optimal routes will never be used, and will age out of the routing table normally. [CSCdi25037]
VINES Crash in vines_best_path_from_delt in 9.21, 10.0 and 10.2. Crash is seen when continuously issuing a "show vines routing" command on one exec process while issuing "clear vines neighbor *" commands on a second exec process. Fixed in 9.21(5.4), 10.0(5.5) and 10.2(1.1). [CSCdi25310]
Banyan Support has asked for the ability to diable the enhancements added to cisco's VINES RTP support to reduce network overhead. The first enhancement is the split horizon of regular routing updates. The second enhancement is that immediate updates, sent to announce topology changes, contain only the information that has changed. In both of these cases, a Banyan server would transmit the full topology. [CSCdi25325]
Wide-Area Networking
When acting as an IP bridge but using IP for management, the system is not recognizing packets destined for one of its IP addresses when received via frame relay. [CSCdi17739]
PPP drops LCP options, from config request if rejected by peer. This is not always correct [CSCdi19434]
When using a BRI interface as a backup interface, and the backup is being done based on load, the BRI interface may be taken down prematurely, even though the load is still high. [CSCdi20472]
The LAPB MIB values reported for lapbFlowRejOutPkts and lapbFlowRejInPkts are reversed. [CSCdi22390]
Under rare circumstances, a cisco-500 line configured for hardware flow control may become idle while the cisco has the RTS signal dropped. This can prevent some modems from answering the next call. [CSCdi22708]
Cisco 2500 routers with a BRI interface will now pass the Layer 1 homologation tests required for NET3 in Europe. This is due to the ability of the Labs to now be able to execute these tests. These are corner case tests and will probably never been seen by customers or in the field. [CSCdi23409]
If a switched VC is connected without explicit flow control negotiation, an interface that is not configured with the correct maximum packet size values may exhibit mysterious problems without an obvious cause.
When a partial data packet is switched on such a VC (and the D-bit is zero), the M-bit is forced to zero; this behavior is required per the standards and GOSIP certification, but the VC's message boundaries are lost. [CSCdi23465]
If a a frame relay interface is divided in subinterfaces, the mapping between the PVCs and the interfaces can get out of sync if the frame relay network reports the DLCI deleted and then adds it

again. [CSCdi23706]

The communication server or router may be restarted due to an address error when PAP authentication has been configured on an async line. A PPP client sending an invalid PPP frame may cause this to occur. [CSCdi24013]

Broadcasts, such as routing updates, are not sent out "receive only" DDR interfaces that have neither a **dialer map** nor a **dialer string** configured. Workarounds are to configure a **dialer string**, or **neighbor** commands for the routing protocol in use. [CSCdi24060]

When using IPX over PPP, if the node number is NAK'ed, we continue to ask to negotiate it. [CSCdi24078]

The ISDN software sends an invalid disconnect cause code for Japan. [CSCdi24172]

SMDS DXI frames from an SDSU/SWITCH greater then 6 bytes in length causes an error. The router atempts to fast switch the packet rather then reply to the DXI heartbeat request. [CSCdi24214]

Memory leak in ISDN BRI. PRIM type buffers are lost on ISDN BRI interfaces under heavy activation/deactivation. [CSCdi24495]

Packets less than the minimum size are erroneously accepted on SMDS interfaces. [CSCdi24560]

Broadcast queue drops the of packets that are less than the queue length. [CSCdi24700]

The lower limit for the bytes per minute transmission parameter used with **frame-relay broadcast-queue** is too high. This value was set to a bit rate rather than a byte rate. [CSCdi24940]

9.21(4) Caveats/9.21(5) Modifications

This section describes possibly unexpected behavior by Release 9.21(4). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(4). For additional caveats applicable to Release 9.21(4), see the caveats sections for newer 9.21 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.21(5).

AppleTalk

On an AppleTalk network with no routers, it is possible to configure a communication server with arbitrary network numbers and zones. Other AppleTalk devices on the network will pick these up and function correctly; however, ARAP connections made to the communication server will be unable to do NBP lookups on the network. The workaround is to use the **appletalk zone *** command to configure a network address in the start-up range. [CSCdi19202]

When system uptime exceeds approximately 24.45 days, AppleTalk interfaces can unexpectedly hang during restarts and never become operational. The only workaround is to reload the system. [CSCdi20052]

When entering more than one zone for an Appletalk-enabled interface, zones after the first are shown in the order opposite from what was entered. This has no functional impact on router operation whatsoever, however it may make comparisons of configuration files more difficult.

For example, if the following commands were entered:

interface ethernet 0 appletalk zone a appletalk zone b appletalk zone c appletalk zone d

...the following would be output on 'write terminal':

interface ethernet 0 appletalk zone a appletalk zone d appletalk zone c appletalk zone b

The user ordering will be preserved in an upcoming maintenance release of 9.21 and 10.0. [CSCdi20804]

When the command **no appletalk permit-partial-zones** is enabled, multi-zoned cable-ranges that have one or more restricted zones may be unexpectedly omitted from RTMP routing updates as a result of access control provided by the **appletalk distribute-list** facility. [CSCdi20945]

The NBP name cache is not cleared when **no appletalk name-lookup-interval** is issued. [CSCdi21020]

Attempting to send AppleTalk ping packets from a Cayman tunnel interface can crash the system. Because Cayman tunnels do not have any AppleTalk addresses, ping packets from such interfaces are not supported. [CSCdi21278]

GetZoneList replies inadvertently use one packet for each zone. While this does not affect system functionality, it is an inefficient use of network bandwidth. [CSCdi21291]

The **debug appletalk iptalk** debugging command is not correctly parsed. As a result debugging information for IPTalk cannot be displayed. [CSCdi21872]

When entering access lists, zones after the first are stored in reversed order of entry. This has no functional impact but may make comparisons with configuration files difficult. [CSCdi21970]

Illegal sized DDP packets can be generated unexpectedly. This behavior can occur during MacIP encapsulation of IP packets within DDP. The workaround for MacIP is to reduce the IP MTU to less than 586 bytes. [CSCdi22479]

Basic System Services

It seems that static routes are being reported as ipRouteIfIndex = 0 :

KEY=1.0.0.0 ipRouteDest=1.0.0.0 ipRouteIfIndex=0 ipRouteMetric1=0 ipRouteMetric2=-1 ipRouteMetric3=-1 ipRouteMetric4=-1 ipRouteNextHop=198.92.60.3 ipRouteType=direct ip-RouteProto=local ipRouteAge=18 ipRouteMask=255.0.0.0

Since the type of connection is direct, there is no way to see that this in fact is a static route.

RFC 1213 does not provide a mechanism to identify static routes. Rather than not report them, cisco's implementation reports them as ifIndex=0. [CSCdi17967]

The "show ntp associations" and "show ntp status" commands are only available to the user in enabled state.

There is no workaround to this problem. [CSCdi21281]

Under some circumstances, NTP may not use a full-accuracy timestamp. This can cause drift errors in the range of small numbers of milliseconds, and may contribute to problems experienced with the public domain "ntpdate" application.

There is no workaround to this problem. [CSCdi21744]

Communication Server

The **debug modem** command reports transitions on the DSR signal of the cisco-500 as being on the RING signal. The **debug modem** command reports transitions on the DSR signal of the 500-CS as being on the RING signal. [CSCdi05130]

An access class learned via extended TACACS is not applied to SLIP connections. [CSCdi18108]

DECnet

DECnet does not recognize the new L2 Hello multicast. This would cause disconnectivity in the networks which include new DECnet area router hellos. [CSCdi21862]

Configuration commands involving the DECnet Level 2 multicast are not consistent. Some of them expect the string "decnet_router_l1" (and "_l2"), while others expect to see "decnet_router-l1" (or "-l2"). [CSCdi22635]

EXEC and Configuration Parser

The context sensitive help for the options to the **tacacs-server notify** and **tacacs-server authenticate** commands is somewhat misleading. [CSCdi20670]

CMNS map commands defined for a non-serial interface will not be accepted on a reload because the "cmns enable" configuration command is placed after the maps. [CSCdi20811]

Priority queueing or custom queueing is not supported on X25 or lapb or multi-lapb, but the system

do not prevent this mis-configuration to be entered. This fix solves the problem for 9.21 and later by enforing the restriction. [CSCdi22009]

IBM Connectivity

When applying NetBIOS access lists with **rsrb remote-peer** access list statements on a system with active SRB traffic, the router may reload due to a bus error. The fix changes the system code so that it handles these conditions in a more graceful manner. [CSCdi18993]

Interfaces and Bridging

Formatting of token ring output from **show controller** is not correct. [CSCdi20105]

In rare circumstances a 4000 with a 2R NPM reloads with an address error. [CSCdi21361]

IP Routing Protocols

Invalid packets can be accepted by the router resulting in an exception dump. [CSCdi18314]

The system allows the static class D route, which is invalid, to be entered through the **ip route** command. This will lead to routing table corruption in which multiple entries for the same class D route is resulted. Remove the invalid static class D route recover the routing table. This fix solves the problem for 9.21 and later version by preventing Class D route from being accepted. [CSCdi20100]

This bug is introduced in 9.21. When using **redistribute ospf metric-type** *type-value*, the redistribution will be done correct using the *type-value* specified. But the **metric-type** *type-value* pair will never appear in NVGEN when doing the **write** command. So this piece of information will be lost after rebooting. There is no workaround in this case except reentering the command.

The fix take solve this problem in 9.21 and on. [CSCdi20758]

In OSPF, when a neighbor goes down, a host route for that neighbor is incorrectly added. A possible workaround is to trigger the rebuild of OSPF router link state advertisement by changing the interface metric or by rebooting. [CSCdi21103]

If a majornet route is DEXTERIOR and part of a subnetted net, then the DEXTERIOR bit is not set in the dummy ndb. This may cause show ip route to fail to show the route as exterior even though it is marked as such. IGRP fails to advertise it as exterior. [CSCdi21943]

When **ip as-path access-list** command is configured with list number 199, this command will not appear in NVGEN or **show access-list** command even though it exists. So the access-list is lost after reboot. This fix solves the problem for 9.21 and later version. [CSCdi22465]

ISO CLNS

The **default-information orignate** command for ISIS does not work. The parser does not accept the **orignate** keyword under ISIS at all. There is no workaround for this problem. This bug is introduced in 9.21.

The fix solves this problem. In addition, the **level-1**, **level-1-2** and **level-2** keywords for ISIS is removed. And the **metric-type** is no longer available under ISIS. It is supposed that the user will use **routemap** keyword to configure default route information.

The fix also has effect on OSPF. In 9.21, the **always** keyword cannot be entered with the **metric** keyword together. This problem is corrected by this fix too. [CSCdi19238]

Fast-Switching fails if the padding option is on for clns packets. The router would drop from fast-switching to process switching. This fix solves the problem for 9.21 and latter. [CSCdi20346]

ISIS may advertise passive interfaces when they are down. The system does not generate new LSP to reflect the interface state change. [CSCdi20407]

The system may crash while doing debug isis spf-events. [CSCdi21131]

Starting in 9.1, it is possible for the CLNS Input process to hog the CPU under heavy clns traffic. No workaround is available. This fix solves the problem for 9.1 and on. [CSCdi21418]

ISIS add/drops prefixes from LSP based on state of next hop. This happens When the nexthop of a static route points to adjacency learned via ISIS. The static in the LSP is not withdrawn when the ISIS adjacency goes down. [CSCdi21501]

LSP is corrupted in ISIS-IP process, if there are more than 21 internal reachability informations. [CSCdi21885]

The **setup** command does not prompt per-interface to enable CLNS when CLNS is enabled globally (but is not running prior to running setup). This fix solves the problem for 9.21 and later version. [CSCdi22244]

Novell IPX, XNS, and Apollo Domain

On larger IPX networks Show IPX Server may display a %SYS-3-CPUHOG message. This means the system has determined that for an interval of time the show process has used too much cpu time and is just an informational message. [CSCdi21050]

The original default of the **ipx gns-response-delay** command was 500 ms. This value fixes an issue in NetWare 2.x with dual-connected servers in parallel with a router NetWare 2.x was the most common release. NetWare 3.x and later do not have the same issue, and a nonzero GNS response delay may cause problems in certain situations. The default of the **ipx gns-response-delay** command has been changed to 0. [CSCdi22285]

Protocol Translation

A pre-mature termination of LAT TCP translate sessions could prevent removing the translate commands from the configuration file. Current active users will contain a non-zero count in "show translate" command. [CSCdi21414]

TCP/IP Host-Mode Services

Under certain conditions, failed protocol translation connections between TCP and either X.25 or LAT using the **translate** command's *printer* option can cause a reload of the Protocol Translator. More specifically, the problem occurs when the software detects an error on the incoming TCP connection after the outgoing X.25 or LAT connection has been set up. [CSCdi22217]

UDP broadcasts can be flooded even if TTL checks fail. [CSCdi22568]

VINES

VINES access lists for ARP, ICP, and RTP packets are not displayed properly. [CSCdi21060]

If the router is maintaining a large table of VINES neighbors, and many of those neighbors were learned via RTP redirects, the router will appear to "pause" once every 90 seconds for a couple of seconds. [CSCdi21257]

If a broadcast message is received from a VINES station that does not currently exist in the routers tables, the router may crash. [CSCdi22326]

If the MAC addresses of the router have been forced to something other than their original values, VINES may erroneously complain that there is a potentially duplicated VINES address. [CSCdi22490]

Wide-Area Networking

CLNS and APPLETALK do not work over LAPB. [CSCdi17431]

A system reload occurs intermittently when configuring or removing an existing translate command. [CSCdi17791]

This bug prevents fastswitching from a serial port running IETF frame relay to a token ring. [CSCdi18522]

Priority queuing is not supported on a BRI interface. [CSCdi18843]

A serial interface configured as a dialer rotary group member will not output traffic if the encapsulation on the dialer interface has not changed after the serial interface has been configured. The

workaround is to change the encapsulation on the dialer interface, for instance from HDLC to PPP, or vice-versa, and then back to what is required by the application. [CSCdi20201]

Under some circumstances, the modem control software of the cisco will change the state of the DTR output without issuing an appropriate **debug modem** debugging message. [CSCdi20564]

Protocols other than IP and CLNS over IETF Frame Relay will not interoperate with RFC 1294-compliant devices when the pad byte in the frame header is of size zero. [CSCdi20942]

The "local" global option on incoming x25 to TCP translate commands used in conjunction with the "profile" option, allows ECHO Telnet protocol negotiation NOT to be translated. Echoing of character is performed by the remote PAD. [CSCdi21087]

When using the help function for the frame-relay broadcast_queue command, the parameters are not displayed correctly. The help command displayed only the queue size and transmission rate. The third parameter, maximum number of packets sent in a second, was not displayed. However, typing in the command with all three parameters works as expected. [CSCdi21089]

When issuing the clear counter command the sent and dropped counts for the frame relay broadcast queue are not reset. [CSCdi21094]

[CSCdi21124]

Improved the negotiation time of PPP encapsulated ISDN call setups over the BRI interface. [CSCdi21126]

When removing a dialer from a serial interface, fastswitching is not always reenabled, causing subsequent traffic to be processed switched. [CSCdi21177]

There are times when the 'show smds map' command would enter an infinite loop when the right combination of STATIC map entries exist in the configuration. This would only happen if 2 protocol addresses would HASH into the same location. [CSCdi21239]

When X.25-over-TCP (XOT) sends a Call Confirm that modifies one of the two proposed flow control facilities (window sizes or maximum packet sizes), the values may be set to 0, which is illegal. [CSCdi21602]

When the system is using Frame Relay maps that were created using Inverse ARP, these maps should be dropped when a DLCI becomes inactive or is deleted. In addition, if the DLCI used by a box at the far end changes, the map entry should be updated. The second scenario might occur when Frame Relay is being accessed using dial-up service and the far-end system makes two calls in rapid succession. [CSCdi21870]

$IGNORE [CSCdi21902]

Cisco routers with an ISDN BRI interface may not accept calls that have a Progress information element in the incoming messages. This may cause a correct call to be rejected.

This may be seen with debug isdn-q931 showing a RELEASE or RELEASE COMPLETE message sent with a Cause value of 0xE4, "Invalid IE contents". [CSCdi22621]

Cisco routers with an ISDN BRI interface may not accept incoming 56 kb calls when using basic-net3 switchtype.

Also, show the speed of the outgoing call when debug isdn-event is enabled. [CSCdi22624]

Dialer rotary groups do not work, they are unable to dial out a phone number. [CSCdi22715]

XRemote

9.21(3) Caveats/9.21(4) Modifications

This section describes possibly unexpected behavior by Release 9.21(3). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(3). For additional caveats applicable to Release 9.21(3), see the caveats sections for newer 9.21 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.21(4).

AppleTalk

Executing the show apple interface command may cause the system to restart itself. This happened on interfaces configured with many zones. [CSCdi18875]

Access list-based priority queuing for AppleTalk was incorrectly based on source network numbers instead of destination network numbers. [CSCdi18935]

The command **appletalk route-redistribution** is now assumed to be the default when AppleTalk EIGRP routing is enabled. To be consistent with the handling of other commands in the router, we do not write the default setting of configuration commands into NVRAM. [CSCdi19394]

If the router's configuration contained the global configuration command **appletalk lookup-type** , the router would crash with a bus error. [CSCdi19463]

The appletalk cable-range switches between good and bad state when it has been deleted and a network is then added. [CSCdi19712]

Apple HSSI port has problem aquiring zones even though both parties are seed ports. This is most likely to happen for HSSI interfaces configured for no apple protocol rtmp or apple protocol eigrp. [CSCdi19901]

IPTalk is broken. [CSCdi19995]

A number of ZIP packets are not fully validated before they are processed. There is no significant system impact. [CSCdi20397]

Basic System Services

Under rare circumstances, the **clear line** command fails to clear the process running on that line. A **show process** command shows that the process on that line has an inappropriate and rapidly increasing number in the "invoked" column. [CSCdi16063]

If two systems are configured with "ntp peer" statements pointing at one another, one system will lose its configuration statement after NTP starts up. The only workaround is to use "ntp server" instead of "ntp peer", or to configure only one of the two systems (since it isn't really necessary to configure both of them). [CSCdi16597]

If a SAP update packet is received with an invalid length, much larger than the data actually contained in the packet, the system may reload. It is also possible, but unlikely, that invalid server entries may appear in the **show ipx server** table. When these packets are received, they should be counted as SAP format errors and the counter displayed by the **show ipx traffic** command should increment. [CSCdi19010]

If the configuration command "no ntp broadcast" is entered for a particular interface, NTP broadcast will be turned off on all interfaces. The workaround is to re-configure NTP broadcast on those interfaces on which it is still desired. [CSCdi19103]

If the system is restarted before summer time went into effect, and a "show version" command is entered after summer time has gone into effect, the restart date and time will be displayed in terms of summer time instead of standard time. [CSCdi19108]

When using the command **tacacs authenticate slip**, only users who have already logged in via TACACS will be authenticated. Users not logged in will be allowed to pass without a challenge. The new behavior is to allow the administrator to require users not logged in via TACACS to log in before being allowed to run slip (or any other protocol requiring authorization).

The new behavior only applies when using the optional **always** keyword is added, e.g., **tacacs-server authenticate slip always** [CSCdi19246]

A new error message has been added to the system. It has the form: %SYS-3-CPUHOG: Task ran for nnn msec, Process = xxx, PC = yyyyy If you encounter this message, please contact customer

support. [CSCdi19294]

Communication Server

The line subcommands "parity" ,"stopbits" and "databits" are accepted for an AUX port but are never displayed or saved to memory. [CSCdi18794]

The Communications Server **rlogin** command does not correctly parse an IP address as an argument, resulting in the error message "% Invalid IP address". The work-around is to use a host name. [CSCdi20178]

TACACS usernames that contain an '@' will have the '@' and the data after the '@' stripped due to the new TACACS directed-request feature.

For most customers this is not a problem, but some customers parse the '@' in their tacacs server and take special actions based on the data after the '@'. These customers should use the global configuration command

"no tacacs-server directed-request"

to cause the cisco product to pass the entire user input to the tacacs server. [CSCdi20394]

DECnet

When updating DECnet access-lists, the new entries are improperly appended to the access-list. This behaviour deviates from the standard access-list behaviour. [CSCdi12660]

In the **decnet conversion - prefix - , when the prefix is longer than the regular size, the router reloads. The correct behaviour is to give an error message indicating the prefix length is too long. [CSCdi19151]**

If decnet conversion is enabled but the **clns router isis** is not set on the pt-to-pt interface from which decnet adjacency is came from, the information of this ES decnet adjacency will not be put inside the L1 LSP even the **clns router isis** command is set later.

This fix take care about it and make sure decnet adjacency information will be put into the L1 LSP immediately whenever **clns router isis** is set on the interface. [CSCdi19190]

DECnet packets received on the FDDI interface of a Cisco 4000 router are always sent to the system processor for processing, ignoring the setting of the **decnet route-cache** interface attribute. This caveat was introduced in Release 9.21(1). [CSCdi19689]

When DECnet is configured over a tunnel interface, the hellos appear to be sent out on the interface, but the adjacencies do not come up. [CSCdi19902]

While converting from DECnet Phase IV to Phase V (and vice versa), the router holds back a converted packet once in a while and sends it out when some other event happens (for example, routing update and keepalives). This sporadic delay in packet transmission results in degradation of end-to-end DECnet performance. [CSCdi20151]

A problem due to an incorrect interface MTU negotiation is seen on any interface whose default MTU is larger than the Ethernet MTU (for example, FDDI). When the VAX comes up, the router negotiates a block size that is larger than the maximum value that it can process (1524). Consequently, all adjacent routers send larger-sized updates, which the router rejects. This makes all destinations behind the router unreachable. [CSCdi20225]

EXEC and Configuration Parser

The command "lo" has different results with different versions of cs500 software. For instance, in 9.1(3), typing "lo" means logout. In 9.1(6), the command "lo" means lock. In 9.21, "lo" generates the message "ambiguous command". [CSCdi15444]

Translate commands including the tcp port option cannot be removed from the configuration. [CSCdi18159]

Starting in Release 9.21, with the new parser, the **ip split-horizon** command is generated before

the **encapsulation command in the configuration file during NVGEN. Because the encapsulation** command has a different default for turning split horizon on and off for different encapsulations, the **ip split-horizon** or **no ip split-horizon** cisco.local.lts.ltsLineTable.ltsLineEntry.tsLineModem.87 none cisco.local.lts.ltsLineTable.ltsLineEntry.tsLineModem.88 none cisco.local.lts.ltsLineTable.ltsLineEntry.tsLineModem.89 none cisco.local.lts.ltsLineTable.ltsLineEntry.tsLineModem.90 none

Any other modem control is returned properly. [CSCdi17661]

The handling (and thus the Response) to an SNMP Get/Set Request message with any of the following erroneous cases is incorrect: * Extra octets in OID or instance field; * Insufficient values given for instance.

A Get Respoonse with noSuchName error status should be returned rather than responding by ignoring the invalidity. [CSCdi17926]

The **busy-message** config command is not working at all. [CSCdi18586]

Communication Server

When a **username** command is used with an **autocommand** that connects to a remote host using telnet or rlogin, the screen size reported to the remote host will not match the configuration of the line. [CSCdi16661]

Under certain sequences of telnet option negotiation, the telnet local echo fails to work as expected. [CSCdi17185]

Attempts to establish Host Initiated Connections to remote destinations get placed in the terminal port queue if the remote destination is busy. By default, the terminal queue is scanned every 60 seconds. The remote destination for each entry is polled. If the destination has become idle, the connection is established and the entry is removed from the queue.

Teminal Queue Subcommands

The terminal queue subcommands modify the operation of the terminal queueing system. **terminal-queue** subcommands always follow a **terminal-queue** command.

The **entry-retry-interval** subcommand changes the default polling interval to the number of seconds specified on the command line. The maximum interval that can be specified is 255 seconds. The **no terminal-queue entry-retry-interval** command restores the default to 60 seconds. [CSCdi17780]

DECnet

Connect initiate packets sent to a Phase V cluster alias cause the router to display SYS-2-INLIST messages. This does not happen if the packets are sent to a 'real' address. [CSCdi16801]

The DECnet router maintains routes through a router in Initializing state. The correct behaviour is to purge the routes through the 'Intializing' router. The impact is that packets are sent to the router in Initalizing state. [CSCdi17171]

On a router running DECnet IV to V conversion, the phase IV adjacencies are not entered in the phase V table. The impact is the functioning of the phase IV - phase V conversion is affected. [CSCdi17520]

A router running DECnet produces SYS-INLIST errors and could crash. The router crashes when the command **microcode reload** is uses on the 7000 routers. The workaround is to disable decnet routing on the router and then do the required **microcode reload [CSCdi18439]**

Router does not recognize the **decnet advertise** command and prints out a parser error. [CSCdi18468]

In ISO-IGRP adjacent Phase IV-only routers and End-Nodes are NOT recognized as adjacent when the routing-table is calculated, but may instead be reachable via another adjacent router. This

causes the adjacent node to be 1 hop away in the ISO-IGRP routing table even if it is adjacent. [CSCdi18483]

For a router that is connected to a LAN where ANOTHER router is elected designated router, the **show decnet interface** command displays the wrong designated router. The displayed designated router is to the one we received the latest router hello from. This has no practical impact at all, the selection process works fine, even if the router it self becomes the designated router. [CSCdi18497]

When decnet is enabled on token ring and ethernet interfaces, the decnet packets are slow switched on the ethernet interface. The correct behaviour is that unless explicitly stated using the command **no decnet route-cache** , the decnet packets must be fast switched on all interfaces. This behaviour is caused only when decnet is enabled on token ring interfaces along with the ethernet interfaces. With only ethernet interfaces configured for decnet, the packets are fast switched. The workaround is to reload the router. [CSCdi18666]

The symptoms exhibited by this caveat is a large value for the encapsulation failure count maintained by the DECnet process as seen in the output of a **sh decnet traffic** . This happens because the "ALL PHASE IV L2" routers encapsulation type on the token ring was not being recognized by the router. [CSCdi18783]

IBM Connectivity

The SDLLC Local Acknowledgement feature may not enable properly despite being configured correctly due to an internal code error. [CSCdi16675]

When removing SDLC encapsulation from an FSIP interface with the **no encapsulation sdlc-primary** or **no encapsulation sdlc-secondary** configuration commands, the CxBus complex is reset. The fix prevents this from happening by negating the need to reallocate the CxBus buffers for this encapsulation change. [CSCdi16767]

Due to the incorrect frame translation when SR/TLB is enable. Bridging access-list ([input/ouput]-lsap-list) will fail to stop the frame correctly. [CSCdi17037]

During SDLLC start-up, the system does not respond to the XID POLL frame sent by some upstream devices which require a reponse to properly bring up the LLC2 session. [CSCdi17093]

The **netbios access-list** feature does not function properly for permit/deny decisions against NETBIOS names which contain special non-alphanumeric characters. The workaround is to change the NETBIOS names of the end stations to use only alphanumeric characters. [CSCdi17163]

Reverse Ethernet SDLLC configurations do not work properly under certain conditions. [CSCdi17314]

Output from the **show sdllc** command does not display a page at a time, but instead scrolls until complete. [CSCdi17526]

When source-route bridging is enabled, a **show interface tokenring** *unit* display is scrolled off the screen. [CSCdi17746]

Reverse SDLLC fails when the system is polled at an high rate by the primary SDLC device. The problem may be circumvented by lowering the poll rate on the primary device if configurable. The fix prevents this condition from occuring. [CSCdi17763]

A system reload may happen while configure source-route bridging, and a **no vines routing** is entered. [CSCdi17862]

More thoroughly documented System Code to make maintenance much easier. [CSCdi17939]

When command **source-bridge proxy-netbios-only** is entered, it is put in the configuration file as **source-bridge cos-enable** . The latter in turn isn't recognized as a command and you have to enter **no source-bridge proxy-netbios-only** to get **source-bridge cos-enable** out of the config.

[CSCdi17997]

When trying to open more than one FST remote peer version 3 connection, only one peer can open successfully. The remaining FST remote peers stay in a closed state. The work around is to use TCP or direct encapsulation. The problem was due to an internal data structure not being initialized properly and has been fixed. [CSCdi18117]

The router fails to accept **sdlc rts-timeout, sdlc cts-delay** even thought the interface is configured for half duplex. After the command is entered, the following message will be displayed on console: SDLC : rts-timeout only applies to half duplex infaces. SDLC : cts-delay only applies to half duplex interfaces. [CSCdi18135]

The **netbios name-cache query-timeout** configuration command is not accepted as valid by the parser. The code was fixed to allow the command to be entered. [CSCdi18237]

The configuration command does not allow the **no source-bridge active** configuration command to remove the SRB definition from a Token Ring interface. The parser code was changed to properly handle this condition. [CSCdi18280]

When using FST encapsulation for Remote Source Route Bridging the FWD field in the **show source-bridge** output is erroneous. The number of frames forwarded to the local rings is displayed as twice the actual value. This problem [CSCdi18316]

The old default sdlc hold queue depth of 50 is too low in some configurations and does not allow sufficient time to apply back pressure on the LLC connection before packets are dropped. The default sdlc hold queue value has been changed to 200. The value is still user configurable via the **sdlc holdq** interface configuration sub-command. This problem is most likely to be an issue in configurations with high speed RSRB connectivity feeding a slow speed SDLC line. [CSCdi18461]

The local acknowledgement feature of remote source route bridging sends a SABME to the end station with the incorrect direction bit set in the RIF portion of the MAC frame. The fix corrects this behavior, which existed only in interim releases between 9.1(10.1) and 9.1(10.4) inclusive. [CSCdi18617]

Show source-bridge doesn't count autonomously-switched frames. [CSCdi18714]

A false error message: " All STUN peers for TGs must use the same IP address" may be displayed on a console when a STUN TG peer configuration command is entered. [CSCdi18747]

Support Broadcast 0xFF SDLC address on SNRM for TG. [CSCdi18897]

Interfaces and Bridging

When transparently bridging is turned on Multibus Token Ring cards, the monitor bit is not cleared in the token when the packet is flooded to another Token Ring interface. The Active Monitor on the destination ring will see this bit set, assume the packet has already passed around the ring, purge it, and reissue a new free token. The work around for this problem is to add a static bridge table entry for each destination address, e.g.: **bridge 1 addr 0208.6ce2.088e forward t 0**. Note that the address must be in Ethernet canonical format. This ensures that packets destined for this address will not have to be flooded via transparent bridging. This problem may not happen consistently, since the location of the Active Monitor on the destination ring may change over time. [CSCdi12451]

Systems configured for autonomous bridging fail to update their bridging tables properly when a host moves from one subnet to another. The workaround is to issue the **clear bridge** *group* command to remove any le [CSCdi13997]

The interface subcommand **pulse-time** is ignored by the serial interface on the 3104. It is displayed in the output of **show interface**, but not in **write term**. [CSCdi15558]

4000 FDDI interface goes into administratively down status and stays there.

A show interface for the fddi will include a message saying:

Forced FDDI shutdown when CMT rate exceeded 10358 events/sec

And a show controller FDDI will report:

last non zero cmt rate 10358/sec, peak rate 23/sec

This is a bug in how we determine the rate of CMT (connection management) events. In this example the actual rate never exceeded 23 per second.

In early testing of the FDDI interface on the 4000, an excessive CMT rate could use up 100% of the processor and lock it up, so we checked for excessive rates and if the rate exceeded 1000/second would shut down the interface to protect the rest of the router.

This fix prevents an incorrect CMT rate from being reported and causing an interface to be administratively shutdown. [CSCdi17010]

Fast switching cache values for all protocols are incremented when a serial interface cannot send a frame out a serial link. [CSCdi17332]

When setting queue-limits on any interface, the cbus complex will reset itself. This will cause token rings to re-initialize. [CSCdi17646]

In SIF configuration response, the Path Descriptor Parameter returns a wrong value for the Connection Resource Index for the MAC.. [CSCdi17836]

"cmt disconnect" will force bypass switch to bypass mode. But "cmt connect" will fail to set the bypass switch back to connect mode. This fix corrects the mistake. [CSCdi18010]

When configuring systems for autonomous bridging, the configuration parser accepts the interface subcommand **bridge-group** *group* cbus-bridging even if standard bridging has not yet been configured. The end resul [CSCdi18129]

On systems configured for autonomous bridging, bridge cache entry creation and recreation is performed before, rather than after, checking for filtering and forwarding configuration information. The result is the creation of unnecessary or unused [CSCdi18294]

The system software configuration parser incorrectly accepts the interface configuration of bridging over LAPB encapsulation. Interfaces cannot be configured to support both LAPB encapsulation and bridging at the same time. [CSCdi18420]

The system will not allow you to configure early-token-release for any token ring interfaces. [CSCdi18648]

System software fails to suppress the support of autonomous bridging on interfaces setup for LAT compression, circuit groups, or non-MAC address-type access-lists, which all conflict with autonomous bridging. [CSCdi18658]

The command **show interface** when issued on a system with 4T NIM configured for half-duplex operation may cause the system to restart. [CSCdi18752]

IP Routing Protocols

When using route map, the absence of a route-map clause does not act as an "allow everything" like it should. [CSCdi14007]

If a router receives an IGRP update from a host on a different IP network that the IP network(s) of the receiving interface, the routes are accepted anyway. These routes will subsequently be propagated as inaccessible. The workaround is to either set an administrative distance in order to discard those routes or to add a secondary address to the receiving interface so that it belongs to the same IP network of the sending interface. [CSCdi15106]

When OSPF is not configured on an interface and the user says "show ip ospf interface xxx", routers using Motorola processors will just display the bogus data but on a MIPs processor, this is illegal and causes system to restart. [CSCdi16220]

Checking mask in Hello packets received on Virtual-links causes interoperability problem with DEC.

Since Virtual links are like unnumberd links, masking checking is not needed. [CSCdi16693]

OSPF receives a RIP route thru the secondary address on the interface that has OSPF turned on on the primary address, and redistributes the route into OSPF with the forwarding address set to the secondary address. As a result, SPF on this external LSA fails because it cannot reach the forwarding address (or this address is also imported as external route which we cannot use to get to another external). [CSCdi17369]

When redistributing from one IGRP process into another IGRP process, any outbound distribute-lists are applied before auto-summarization. Therefore, given a scenario where a router has two interfaces in different major networks e.g. igrp 1 being run for B1.B2.0.0 and igrp 2 being run for B3.B4.0.0 (both Class B networks) and igrp 1 is being redistributed into igrp 2. If igrp 2 has a distribute-list out which permits only B1.B2.0.0 (with a wildcard of 0.0.0.0), this network will not be advertised by igrp 2 because of this caveat. [CSCdi17440]

OSPF fails to remove a route through an fddi interface when that interface is unconfigured and shutdown. This does not happen with other routing protocols (IGRP or RIP), and does not happen over an ethernet or token ring interface. [CSCdi17712]

Users can ping broadcast address for user level exec. [CSCdi17719]

Using both 'loose source route' and 'record route' options simultaneously in ICMP ECHO requests may confuse cisco routers. This may cause problems with normal IP routing. [CSCdi17879]

Priority queueing does not classify IP fragments as expected. This is because some of the information required to classify the fragment is not in the packet. [CSCdi17905]

When running with synchronization enabled (the default), BGP will install a less preferable external path into the main routing table. If this less preferable external path is installed, there is no way for the internal path to ever get synchronized. Therfore, the behavior should be changed so that lack of synchronization does not affect preferability of a given path, but rather, we should merely never install or advertize an unsynchronized path. [CSCdi18119]

If a router receives a LSA with invalid LSA type, it can cause a crash as the LSA type is used to index into database Hash table to lookup the LSA. [CSCdi18144]

Static routes to loopback interfaces can be created even if IP routing is not running. [CSCdi18165]

When priority queueing, if an access list is set up to classify against TCP/UDP ports, the priority queueing classification code will compare garbage against the port #s when processing a fragmented IP packet.

The fragmented IP packets to not contain a UDP or TCP header. [CSCdi18278]

OSPF: Doesn't show multiple paths when parallel links are unnumbered [CSCdi18313]

Behaviour of IP source routing can cause incorrect behaviour in BSD based telnet programs. [CSCdi18317]

IP static routes may not be updated once a minute as expected. [CSCdi18340]

The router continually reports '%SYS-2-NOBLOCK: event dismiss with blocking disabled' errors preventing the router from processing other information. Reloading the router temporarily resolves the issue. [CSCdi18565]

The **ping** command gives inaccurate timings when sending packets less than 32 bytes. This manifests itself as very large (and sometimes negative) round trip times. This is cosmetic only and has no operational impact. [CSCdi18685]

The default network does not work properly depending on the subnet used. [CSCdi18743]

distribute-list number eigrp process-number out is broken. This command is ignored by EIGRP

process. [CSCdi18779]

ISO CLNS

no ip address command is broken in IS-iS code. The IS-IS reports the ip address afther the command is entered; However, It works OK for adding/changing an IP address. [CSCdi16387]

Static ES enteries not inserted in IS-IS non-psudonode. The impact is the system does not put the end node into its IS-IS database. [CSCdi16553]

On entering the **clns is-neighbour** command the IS-IS process stops redistributing CLNS static routes. The impact is that CLNS routing problems may arise, but to incomplete information sent into the network. [CSCdi16954]

The holding-time value would show up as large as 705032, though the system is configured for larger value in clns holding-time seconds command. [CSCdi17491]

The behaviour of this caveat is that CLNS neighbours over X.25 networks can't see each other. The impact is that no neighbour adjacencies can be formed over X.25 networks. [CSCdi17771]

**** Release Notes ****

In configuring a CLNS host using the ,**clns host** character '+' cannot be used in the Hostname. The impact is that in the configuration file the Hostname appears without a '+' and the host is unreachable from the router when it is ping'ed using its name in the configuration. [CSCdi18192]

When CLNS cluster aliasing was enabled on an interface and there were more than 2 members in the cluster, we would have a memory leak. This fix will take care of the general case when we have more than 2 members in a CLNS cluster. [CSCdi18550]

IS-IS stops sending IS-IS Hellos after 49 days. This leads to loose adjacencies. [CSCdi18757]

Novell IPX, XNS, and Apollo Domain

A number of messages sent when the system failed to allocate memory for a task have been made more general and can be logged using the standard logging facility.

%IPX-3-NOMEMORY: Unable to allocate memory for [event]

Event can be any number of events or procedures which attempt to allocate memory.

-------------Changes --- (was) IPX: Cannot allocate memory for SAP queue! (now) %IPX-3-NOMEMORY: Unable to allocate memory for SAP queue

(was) %%No memory available (now) %IPX-3-NOMEMORY: Unable to allocate memory for routing table (now) %IPX-3-NOMEMORY: Unable to allocate memory for SAP hash table

(was) %%Insufficient memory for SAP update (now) %IPX-3-NOMEMORY: Unable to allocate memory for SAP update

(was) IPX: No memory to allocate SAP update queue entry! (now) %IPX-3-NOMEMORY: Unable to allocate memory for SAP update queue entry

(was) IPX: No memory to allocate RIP update queue entry! (now) %IPX-3-NOMEMORY: Unable to allocate memory for RIP update queue entry

(was) IPXSAP: Insufficient memory to sort SAP queue x (now) %IPX-3-NOMEMORY: Unable to allocate memory for SAP heap sort

(was) %%Insufficient memory (now) %IPX-3-NOMEMORY: Unable to allocate memory for pointer array

(was) IPXRIP: Cannot send update out interface -- out of memory! (now) %IPX-3-NOMEMORY: Unable to allocate memory for RIP update

(was) %% Unable to allocate memory for EIGRP process. (now) %IPX-3-NOMEMORY: Unable to allocate memory for EIGRP process

(was) [no message, just fail to add entry to table] (now) %IPX-3-NOMEMORY: Unable to allocate memory for route entry [CSCdi17837]

IPX access-list permit too large a value on the upper range than the eight bit protocol field. [CSCdi17848]

Novell updates could stop being issued out an interface after three weeks of uptime. [CSCdi18168]

When doing a clear ipx route * immediately start repopulating the table by sending out RIP and SAP general requests to force our neighbors to reply with RIP and SAP updates. [CSCdi18234]

The destination MAC address isn't properly being bit swapped for FDDI media, the result is an incorrect destination MAC address in show XNS cache and tthe outging packet. In the cache the destination MAC address should be bitswapped but is not, and the outgoing FDDI frame has a bit swapped destination MAC address. Turning off xns route-cache on the FDDI interface is a workaround. [CSCdi18273]

The command no ipx helper-address will not take a novell address and will remove all configured ipx helper-addresses on that interface. [CSCdi18290]

When **ipx router eigrp 69 no redist rip** is done, the "no redist" isn't written out on a "write term". [CSCdi18576]

Lotus Notes OS/2 Servers appear to send Service Advertisement Protocol packets with a tc equal to one. Originating hosts are supposed to set the tc field to zero. cisco 9.21 software was ignoring RIP and SAP Updates with tc field greater than zero. A workaround is to configure a static SAP entry on the cisco that is ignoring the Notes OS/2 Advertisment. [CSCdi18737]

If an IPX RIP entry is in holddown in the routing table when RIP to IPX-EIGRP redistribution is turned off, the route will not be removed from the IPX EIGRP topology table as it should be. [CSCdi18786]

In 9.21(2.4) and 10.0(0.13) the bug fix for CSCdi18737 inadvertently caused NetBIOS broadcasts, and helpered broadcasts which have traveled one hop to be dropped instead of forwarded. Directed NetBIOS broadcasts and non-NetBIOS broadcasts helpered using a helper address other than -1.FFFF.FFFF.FFFF are forwarded properly. [CSCdi18806]

The Show Access-List command will not stop at the end of the screen with the "-More-" but continues to scroll off the console screen for Novell/IPX and XNS access-lists. [CSCdi18878]

Protocol Translation

The Protocol Translator does not handle TCP urgent data. This caused a %TN-3-BADSTATE: Illegal state 4 error message to be displayed. [CSCdi17153]

Under some circumstances, the system can reload when user reset the connection. [CSCdi18535]

Incoming PAD calls on to VTYs without passwords fails without issuing the message "Password required but none set." [CSCdi18618]

[CSCdi18908]

VINES

It is possible, to cause VINES to slowly lose memory. It requires a user issue the 'show vines neighbor' or 'show vines route' command, leave the command paused at the 'more' prompt, and then to delete the last neighbor or route displayed. [CSCdi17998]

If there has been a very long interval (greater than one minute) between the time a router proxies a request from a server and the time that the server replies to that proxy, the router may reload. [CSCdi18285]

Flash updates are not being sent when they should be. This does not affect any routes in the routing tables, just the rate at which table changes are propagated. [CSCdi18286]

When a redirect is received, it is entered into the routing table with a wrong metric value. This can cause circular routes in a network. [CSCdi18287]

Add the ability to filter received routing information based upon either source address or content,

and the ability to filter the content of transmitted routing information. [CSCdi18569]

The metric on existing vines table entries does not get updated when the interface metric is changed. [CSCdi18583]

The router always sends a RTP request for information as a broadcast packet, even when it is attempting to solicit information from a single station. This does not have any affect other than using some unnecessary network bandwidth. [CSCdi18766]

Sites with highly dynamic neighbors (i.e, neighbors going up and down a lot) could sould see system crashing. There is a slim chance but it does and can happen. [CSCdi18994]

When running Vines, certain routing activity such as route deletion may cause the router to reload. [CSCdi19079]

Wide-Area Networking

Configuring SMDS on serial lines which are shutdown, and subsequently re-enabling them can in some circumstances cause a reload. A Token Ring interface appears to be required to trigger this problem. [CSCdi15880]

Packets bridged over SMDS are missing two octets of padding, which could cause interoperability problems. [CSCdi17072]

The **show frame-relay pvc** displays a blank in the interface field when displaying statistics for the PVC. However, the title line displays the information properly. [CSCdi17142]

When the timeourt is change during the PPP LCP negotiation with the "PPP restart-timer" command, there is no effect on timeouts used by the control protocols for individual protocols. This can prevent PPP from initializing over certain slow, high-delay links. [CSCdi17416]

The incorrect SMDS SIP3 padding is used at the end of the INFO field. The 1 to 3 bytes inserted at the end of the INFO field are not set to zero. This may cause other router vendors equipment connected to SMDS to reject the frame. [CSCdi17511]

Asynchronous interfaces configured for PPP dialout or SLIP dial-in that are switched back and forth between PPP and another encapsulation type lose memory. [CSCdi17512]

Fastswitching of IP over PPP is not disallowed for DDR. This prevents the idle timer to be reset when packets go through. Eventually the idle timer expires and the connection is torn down, even though traffic is still going through. The workaround is to explicitly disable fast switching on the DDR interface. [CSCdi17683]

Pings fail when using point-to-point subinterfaces. This is a symptom of the router using an incorrect map entry for a subinterface. [CSCdi17686]

The subcommand **clns is-neighbor** *nsap snpa [X25-facilities-info]*, will not be accepted after a reload or **config network** unless the *nudata statement* is placed in quotes. [CSCdi17761]

If an SMDS address is not specified for an interface, the removal of the default mode DXI 3.2 using the command **no smds dxi, will not be saved in the configuration using a write mem. A system reload will cause the system to revert to the default mode. [CSCdi17762]**

The EXEC command**sh x25 vc**does not display the number of sent INTs (Interrupts). [CSCdi17796]

After ISDN DDR connection is already established, sometimes the line gets a DISCONNECT message from the remote end and the line drops. The only way to get the line back to where you can redial the distant end is to issue a **clear int bri 0** command. [CSCdi17908]

CLNS and bridge packets on an interface that is configured for atm-dxi encapsulation fail because the address is not entered properly into the address table. [CSCdi18037]

While the CLNS and bridge map statements can be added using the **atm-dxi map clns/bridge** *vpi vci* command, they cannot be removed using the **no atm-dxi map clns/bridge** *vpi vci* command.

The map statement for other protocols works fine. [CSCdi18040]

A Cisco 2500/3000 series router with an ISDN Basic Rate Interface (BRI) can now support the Australian switchtype. The basic-ts013 switchtype should be used in the configuration file. [CSCdi18128]

Improved the debug output of **debug isdn-Q931** by printing out both the decoded and raw Q931 information. [CSCdi18161]

If the encapsulation on one router in a ppp link between a pair of routers is changed to hdlc, that router's line protocol will go to the down state, while the adjacent router, still running ppp encap, will report line protocol up. [CSCdi18201]

Switching to ppp encapsulation will cause some memory to be allocated. This memory is never freed. The problem is that when you switch to ppp encaps again later the memory is re-allocated. [CSCdi18243]

The commands **clear dialer interface async 0** and **clear dialer interface bri 0** are rejected when they should be accepted and clear the respective dialer statistics.. [CSCdi18256]

If lcp is up and LQM is configured (ppp quality x) the link will go into a CONFREQ/CONFACK loop. The workaround is to shut / no shut the interface. [CSCdi18277]

A cisco router with a BRI interface (2500/3000) configured with a basic-net3 switchtype may encounter problems sending and receiving data on the B-channel.

This may occur if a SETUP_ACK message, instead of a CALL_PROCEEDING message, is received in response to an outgoing SETUP message.

It is also possible for buffers to be lost when running the basic-net3 switchtype. This can occur on the router requesting a disconnect from the network. Eventually the router will run out of available buffers and reload. [CSCdi18423]

Add new feature for incoming ISDN BRI calls. Now able to verify a Called Party Number and Subaddress number in the incoming SETUP message if it is delivered by the switch. The verification will always be done if a number is configured in either an "answer1" or an "answer2" command : interface bri 0 isdn answer1 5552222:1234 or isdn answer2 9991111:9876

The 5552222 and 9991111 are the Called party numbers and the 1234 and the 9876 are the Subaddresses. Note that the ":" is the separator.

If nothing is configured, all calls will be accepted. In case one or both the answer numbers are configured, then the incoming Called party number and the Subaddress will be verified before accepting the call.

It is possible to configure just the Called party number or just the Subaddress. In such a case, only that part will be verified.

The verification will proceed from right to left, for both Called party number and the Subaddress.

It is possible to declare a digit a "don't care" digit by configuring it as a "x" or "X". In such a case, any incoming digit will be allowed. [CSCdi18425]

When switching an X.25 Call, an NUID facility is stripped in all cases. If the Call is switched to a DTE interface, the NUID facility should be forwarded. [CSCdi18526]

Under very obscure cirumstances, X.25 may not free a buffer on failure to route a Call from a remote TCP connection to a DDN interface. [CSCdi18527]

When acting as a frame relay switch, the router does not set the most significant bit of the second DLCI octet in the PVC Information Element. This bit should be set to 1. [CSCdi18553]

A router with an ISDN BRI interface configured for the basic-1tr6 switchtype may have problems connecting on Channel B2. An incoming SETUP message using Channel B2 can be incorrectly answered using Channel B1. This may cause the PPP protocol to keep the BRI channel interface in a

Protocol-Up and Line-Down situation. It will also prevent the B2 channel from receiving any more calls. [CSCdi18562]

The **ppp restart-timer** commands parameter was being used as milliseconds. This should be seconds. Make sure it is scaled accordingly. [CSCdi18599]

The parser does not accept encapsulating PVC configuration commands. [CSCdi18671]

Cisco routers with an ISDN BRI interface using basic-dms100 or basic-ni1 switch types may have problems connecting to the correct B-channel. This may occur if a call is made to another country and the switch responds to an outgoing SETUP message with a SETUP_ACK message.

The router will now use the correct B-channel. It is also possible to append a # character to the end of the called number. This will tell the switch that "no more digits follow, all the digits are in this SETUP message".

The basic-net3 switchtype will no longer include a Diagnostic field in an outgoing Disconnect message. This caused problems for Sweden when running basic-net3 switchtype. [CSCdi18832]

Cisco routers with an ISDN BRI interface used in Norway and New Zealand require a special version of NET3 code. They are just enough different in their interpretation of the spec. For Norway, use basic-nwnet3, and for New Zealand, use basic-nznet3. [CSCdi18903]

IPX over PPP (IPXCP) now negotiates for the IPX network number. If we are requested, we can supply our net number. If our net number is different than the remote net number, we will ask them to become our number if appropriate (their number smaller than ours), or we will issue an warning message and negotiate with no network numbers. The warning is of the form:

mismatched IPX network numbers. Ours = %x, theirs = %x

Where ours and theirs display the IPX network numbers each side desires. [CSCdi18917]

IPX over PPP now negotiates both network number and node number. Our node number is provided to the other side. They cannot change it, nor can we assign then a node number. [CSCdi19077]

9.21(1) Caveats/9.21(2) Modifications

This section describes possibly unexpected behavior by Release 9.21(1). Unless otherwise noted, these caveats apply to all 9.21 releases up to and including 9.21(1). For additional caveats applicable to Release 9.21(1), see the caveats sections for newer 9.21 releases. The caveats for newer releases precede this section.

All the caveats listed in this section are resolved in release 9.21(2).

AppleTalk

AppleTalk packets forwarded across a Cayman Tunnel do not have their DDP hop counts incremented. There is no workaround, but there is no significant system impact. [CSCdi16470]

Configuring iptalk for an interface and, then removing it from configuration file. Can not enable back the iptalk configuration to the router. The workaround is to disable the appletalk process by configuring first "no appletalk" and then configur [CSCdi16711]

Appletalk fast switch feature is broken over FSIP for phase-1 ethernet.As a result, The router drops to process switching. [CSCdi16895]

cisco's AppleTalk implementation is not compliant with the ANVL AppleTalk test suite. None of the nonconformities have any major system impact and are not usually encountered under normal operation. [CSCdi17048]

On the low end platforms, fast switching from Phase 1 encapsulations on Ethernet to Phase 2 encapsualtions on an HDLC encapsulation serial link are written out with the wrong header size. To work around this, issue the command **no apple route-cache** on your HDLC encapsulated serial links.

This problem is fixed in a future maintenance release. [CSCdi17484]

Routes learned from a Cayman Tunnel peer may never disappear from the routing table and can never be cleared. This unexpected behavior occurs only when the peer is a non-cisco (such as a Cayman Systems GatorBox) and when one or more routes learned via the tunnel belong to more than one zone. There is no workaround. There is no performance impact. [CSCdi17611]

NBP FwdReq and LkUp packets outgoing the same interface on which the original request was received may unexpectedly use the source address of the request packet on AppleTalk interfaces configured with a wide cable-range. As a result devices performing MAC address gleaning may enter incorrect information into their AARP caches. A wide cable-range is defined as a cable-range with different starting and ending network numbers. [CSCdi17628]

If there are some other routers in the ring configured for pre-FDDItalk, this would cause buffer leak. The symptoms are, the router may eventually crash due to memory starvation, or produces "SYS-2-INLIST" messages on the console and these messages give a traceback which leads to the AppleTalk RTMP. [CSCdi17724]

ZIP Queries for nonextended networks may never complete. As a result a router may never acquire a network's zonelist. There is no workaround. [CSCdi17751]

Basic System Services

The following snmp router interface command was added: [no] snmp trap link-status

This command allows the administrator to disable the sending of link-status traps. An example would be: int ethernet 0 no snmp trap link-status Link up/down traps will no longer be sent concerning interface ethernet 0. [CSCdi14808]

An SNMP query of routing information returns a value incorrectly of OTHER for routes that have been discovered using OSPF. [CSCdi16024]

The system fails to reply to a DO TIMEMARK when translating from telnet to X.25. This may result in a telnet session hang between th cisco and the machine sending the DO TIMEMARK. [CSCdi16405]

After a system has been up for some time, the small buffer pool will start to record large number of misses (in the **show buffers** command), even though it claims that there are large numbers of buffers in the free list. The most visible effect is that all xremote will slow down considerably, and the client xremote statistics will show that nearly all packets are being transmitted twice. [CSCdi16843]

The snmp variable snmpEnableAuthenTraps accepted illegal enumerations and count them as disable discards. This has been fixed to accept either an Enable or Disable value. All other values are reported as an error. [CSCdi17414]

Enhancements were made to various temporary MIB objects that fix some ifIndex lex-next issues. Also added additional badValue returns. [CSCdi17415]

Certain ifIndex type tables performed get-nexting incorrectly when presented with an integer index of 0xffff

Certain enumeratins failed to return a badValue when given an invalid enumeration. dot1dStpForwardDelay was in the Table as RW. It was corrected to to be RO. [CSCdi17417]

This is a correction to a previous bug fix that corrected the result of an improperly configured field. [CSCdi17454]

Snmp does not respond to an empty varbind list properly. This has been corrected. [CSCdi17468]

Certain MIB variables dealing with modems parse improperly when given an instance of 0xffff. [CSCdi17521]

Communication Server

There should be a way to kick users off of a cisco comm server after a certain timeout, whether or

not they are idle. Such a command is usefull when contention is great for a dialin pool, for example. [CSCdi17530]

DECnet

The router goes into a state where it stops sending out DECnet Level1 routing updates. The number of Level1 adjacencies will show 0 though the router does have adjacent Level1 nodes. In some instances the Level1 adjacency count will go to 65,363. The impact of this is DECnet Level1 routing breaks down. The workaround is to reinitalize the DECnet routing process in the router by removing and re-entering the DECnet routing commands. [CSCdi15732]

On a router with DECnet enabled, when the global command **no decnet routing** is entered, the Token Ring and FDDI interface with no decnet configured flap. The impact is that routing is momentarily affected on these interfaces. [CSCdi17219]

EXEC and Configuration Parser

If you run the **Setup** command on an already-configured Comm Server, any async network interfaces that had been configured will have IP turned off by the resulting configuration. [CSCdi16873]

IBM Connectivity

When Locally Assigned Addresses with the Manufacturers address of 4000.02 is used in conjunction with command **netbios enable-name-cache** there are no packet savings register. If the mac address is changed to the original vendors number 1000.5A then packet savings counters increment nrmally. [CSCdi15794]

When using the SAP prioritization feature for RSRB, the **clear source-bridge** command resets the counters for only the high queue. The low, medium, and normal queue counters will not be reset. This is a cosmetic defect and does not effect the operation of the SAP prioritization feature. [CSCdi16233]

Adding and removing Source-route bridging interface subcommand(**source-bridge** *source-ring# bridge# destination-ring#* cause memory leak in the router. The memory is allocated by adding the command but will not release properly when the command is removing, thus causes the memory leak. [CSCdi16275]

The configuration parser does not properly accept the **sdlc poll-limit-value** or **sdlc poll-pause-time** configuration commands. [CSCdi16710]

RIF's can get in the cache with an invalid/bogus bits set in the Routing Control field. This fix will validate the RIF entry before caching it. [CSCdi16769]

sdlc poll-wait-timeout cannot be set.
When an interface is set as a secondary station, the following error message displays:"SDLC: slow-poll is only applicable on primary interfaces."
When an interface is set as a primary station, the following error message displays:"SDLC: Poll wait value applicable only to the SDLC secondary side." [CSCdi16961]

The system does not respond properly to an I-frame received just after starting a session using STUN Local-Ack. The work around is to disable Local-Ack. [CSCdi16992]

When removing the global **source-bridge ring-group** configuration command, the router may crash due to an invalid memory pointer. [CSCdi17134]

When remote source-bridging is enabled between multiple peers, one or more of the peers maybe stuck in REMOPEN state. This is observable via show source-bridging. The correct behavior is to transit from REMWAIT to OPEN state. [CSCdi17149]

1/When configuring **encapsulation stun** in 9.21 and 10.0, the mtu and ip mtu are mistakenly set back to the default.

2/**tcp-queue-max** option is not configurable in 9.21 and 10.0. A parser chain bug made the option unreachable.

3/When the user configures **sdlc N1**, the sdlc code sets the interface mtu variable instead of the sdlc_n1 variable. [CSCdi17157]

RSRB peer connections using TCP encapsulation drop packets under heavy and sustained load. The **show source-bridge** command will show large amounts of drops on the remote peer connection and the TCP column will show the queue full at 101 packets. [CSCdi17486]

Output from **debug source-bridge** or **debug source-error** may contain lines in which the IP address of the remote peer is not formatted correctly. [CSCdi17581]

A system reload may occur when attempting to remove stun process by entering **no stun peer-name ip-address**. This seems to be occured only after the stun sessions are up and run for a while. The router won't reload if the above command is entered right after a session is established. [CSCdi17657]

Interfaces and Bridging

CTR cards hear their own DECnet hellos, resulting in a "%DNET-3-HEARSELF: Hello type 1 for my address.." error message. This has no operational impact. [CSCdi07368]

Certain combinations of system code and FDDI microcode may cause packet duplication on the FDDI ring. [CSCdi14083]

After net booting there is a chance that the CTR card may go into a reset condition and not completely initialize. Performing a shutdown followed by a no shutdown on the interface will clear the condition. [CSCdi14554]

The transmitter-delay command does not work with the HSSI interface for CbusII and HIP. Setting the command causes the line to go down. Removing the transmitter-delay will allow the line to come back up. It works correctly with HSSI and CbusI. However, the output of "show controller cbus" displays the delay in microseconds. The delay is set in terms of "flags", not microseconds. [CSCdi15266]

The **downstream neighbor** in show interface fddi remains **0000.0000.0000**. The issue is resolved in microcode and an additional workaround was creaded at system image level. [CSCdi15780]

Systems configured for transparent bridging and DECnet routing will not bridge (forward) DEC multicast traffic after DECnet routing has been turned off. The workaround is to delete those DEC multicast addresses from the bridging table using the gl [CSCdi15827]

Corrupt / invalid ethernet frames where incorrectly bridged to the FDDI interface resulting in **FDDI transitions**. The source of the corrupt frame was in this case a Hirschmann hub sending diagnostic packages. [CSCdi15992]

The commands **show interface [***type unit***]** [accounting] and **show interface stats** fail to display autonomous bridging counters. This is because the commands fail to poll the relevant bridge cac [CSCdi16116]

System issues link down traps from the fddi interface when in fact the interface did not go down. **debug fddi-cmt-events** shows link down and link up trap without any corresponding cmt data. [CSCdi16506]

When running 19.2K on FSIP, clock devisor table in the system image was not taking into account the 8.064 Mhz reference crystal. Changing the table in the system image took care of problem. All other serial adapter's use the 8.0 Mhz reference crystal. [CSCdi16664]

FDDI trace counter can increment when there are no beacons on the ring. [CSCdi16744]

Vines may not work properly on CTR interfaces that are also part of a transparent bridge group. [CSCdi16797]

Novell can not be Fast Switched onto token ring interfaces if the packets have a RIF field. [CSCdi16875]

If a **no dialer** command of any type is issued and the interface is not configured as a dial on demand interface, the router may restart. [CSCdi16886]

On low end routers (IGS, 3000, 4000), changing the serial encapsulation from from non-HDLC to HDLC causes all packets to be processed switched, even though fastswitching is possible and configured. The work around is to do a **shutdown** followed by a **no shutdown** on that interface. [CSCdi16887]

The CRC 32 configuration commnad has been ommitted from the 9.21 and 10.0 versions. [CSCdi16964]

A Token ring interface may return an incorrect error message when an error occurs on the ring. [CSCdi17041]

After a beaconing condition the token ring interface of a TRIP card may not reinsert into the ring. [CSCdi17055]

Each time a **slip** or **ppp** exec command is issued, the software will fail to free a small sized chunk of memory. After the cisco has been running for a long time and many such commands have been issued, the software can completely run out of free memory. [CSCdi17395]

On systems configured for bridging from Ethernet to Ethernet, the bridging of ARPA-encapsulated broadcasts does not work. ARPA-encapsulated packets are incorrectly translated into LLC2-like encapsulated packets. [CSCdi17463]

On systems configured for autonomous bridging, the SNMP ifInUcastPkts counter does not get incremented with each packet input on each port. This counter tracks the number of unicast packets delivered to a higher-layer protocol. [CSCdi17595]

The **CRC 32** interface subcommand does not work properly for the HIP card. HIP ucode 162-2 is required for this option to be enabled. [CSCdi17633]

FDDI trace counter can increment when there is no beacon on the ring. [CSCdi17634]

no fddi if-cmt may cause WRAP B failure. [CSCdi17672]

In systems with only one up interface, broadcast or multicast bridged packets received will cause the input hold queue counter for that interface to remain at 38/75 when traffic is stopped. [CSCdi17675]

Code is re-structured so that variables that were initialized with numbers are now initialized with constants that have been defined in header files. This does not affect any functionality of the code. [CSCdi17679]

In modular systems configured with CSC-1R or CSC-2R Token Ring interface cards running STR-MON microcode version 1.2 or older, NetBIOS multicast traffic is not propogated when bridging from Ethernet to Token Ring to Ethernet. [CSCdi17697]

IP Routing Protocols

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

Under rare circumstances, When configuring static routes for a major net and its subnets, IGRP may pick the wrong metric when advertising the routes. [CSCdi15686]

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

The system inside an OSPF stub area receives invalid updates from the area border router (ABR) when the ABR finds that an Autonomous System Boundary Router (ASBR) from other areas is unreachable. This manifests itself by the repeating error messages.

%OSPF-4-ERRRCV: Received invalid packet: Ext update into stub area

The correct behavior is; the ABR will not send any invalid update into the stub area. There is no workaround for this problem.

The system now has the additional capability to control the amount of information distributed into a stub area. The command **area** *area-id* stub no-summary will prevent any summary link state advertisements (LSAs) from being flooded into a stub area, except for the summary LSA for the default route (0.0.0.0). This capability helps to further reduce the size of the database of routers inside the stub area by eliminating inter-area advertisements in addition to external advertisements. [CSCdi16054]

Routing begins before load from envm flash is complete. On long loads from flash some processes have time to open network sessions. [CSCdi16058]

Route-map match commands do not check to see if value has already been entered. Therefore, duplicate values are NV generated. [CSCdi16128]

Some TCP/IP packets do not obey priority queuing configuration. Compressed TCP/IP packets are not recognized as normal packets and therefore get classified as a different packet type.

The workaround is to add to your priority-list:

priority-list 2 protocol compressedtcp medium tcp 21 priority-list 2 protocol compressedtcp medium tcp 23 priority-list 2 protocol compressedtcp medium tcp 513 priority-list 2 protocol compressedtcp normal priority-list 2 protocol compressedtcp medium tcp 767 [CSCdi16291]

BGP distance command does not load in from NVRAM configuration. [CSCdi16361]

Cosmetic

When running EIGRP-IP , Sho ip route x.x.x.x displays incorrect path under the "Routing Descriptor Blocks:". [CSCdi16371]

OSPF does not sufficiently validate received data, which in some cases can cause system failure. There is no workaround to this problem. [CSCdi16521]

OSPF process logs unnecessary message. No further information available. [CSCdi16592]

When using PPP or SLIP, the addresses/routes are redistributed via OSPF to the other CS500's and routers on the network. However, when this PPP/SLIP interface is disconnected, the fact that this address/route is no longer valid is not redistributed. So the other CS500's and routers still think this is a valid path. This causes some problems. This problem has only occurred in software versions 9.21(0.135) to (0.138) Version 9.21(0.134) and below work fine. [CSCdi16604]

When an interface goes down, the system fails to poison the corresponding subnet route in RIP or HELLO routing advertisements sent out other interfaces that are part of the same major network number. The system also fails to poison a network summary route advertised by RIP or HELLO to other networks. The result is that adjacent routers must time out the corresponding route in their tables, instead of being notified of the routing change immediately. [CSCdi16698]

$IGNORE

Just a cosmetic change to add a new error message. [CSCdi16709]

Show ip ospf database does not stop properly. No further information available. [CSCdi16736]

Failed to delete dummy ndbs created by OSPF process due to pdb-route_type field. This has the effect of running SPF on externals every 30 seconds . Work around: Clear the routing table, "clear ip route *" [CSCdi16831]

Ospf_clean_area() skips every other link state db in the hash chain when clean up. Hence, the database for OSPF Link State Advertisements is not cleaned up properly. [CSCdi16888]

The output of the command **show ip bgp summary** displays the incorrect version number before the session is established. Once the session is established, it reports the correct neighbor session version. [CSCdi17042]

AppleTalk discovery port on the tunnel interface will take long time to get appletalk zone information if 'tunnel sequence-datagram' is configured and one side of the tunnel is reloaded. [CSCdi17070]

Issuing the command 'no router ospf' hangs the system. [CSCdi17080]

Duplicate secondary address can be configured on a second interface if the original interface is shutdown at the time of configuration. After the duplicate address is configured it cannot be removed. [CSCdi17109]

IP pings work with default options but if the extended IP ping is used with the record or verbose options it fails. [CSCdi17112]

SNMP access to BGP variables is not available. [CSCdi17131]

[CSCdi17147]

Under certain conditions, floating static routes may be incorrectly redistributed into IGRP and RIP with an inaccessible metric. [CSCdi17285]

Cannot configure IP static route with next hop of 0.0.0.0. [CSCdi17294]

When a communication server is running in host-routing mode, the communication server does not respond with a proxy ARP on behalf of its SLIP clients. [CSCdi17501]

Debugging log information is created for OSPF. [CSCdi17709]

Using the show ip protocol can cause the router reload. [CSCdi17748]

At system boot time, tacacs code dies because it fails to establish a UDP socket with which to talk to the tacacs server. A change has been made so that UDP socket is successfully established. [CSCdi17830]

ISO CLNS

When the iso-igrp routing for a Serial 1 inteface is changed from Level 1 to Level2, in the configuration the iso-igrp tag changes from the original Serial 1 tag to that of the Serial 0 interface with Level 2 routing. This results in incorrect tagson interfaces. [CSCdi15588]

The packet forwarding rate on a 4000 is lower than expected. [CSCdi16639]

If the router hears ES hellos, this may cause the ISO-IGRP process to get called, though the router is not configured for ISO-IGRP. [CSCdi17079]

IS-IS causes crash dumps appears without termination. [CSCdi17169]

IS-IS routes are dropped after a 24.8 days time. The routes will disappear from the routing table resulted loosing connectivity. [CSCdi17253]

IS-IS does not poison out adjacencies upon reload. The router does not send out hello messages with a holdtime time of zero when it is about to reload so that other routers in the network will immediately reconverge elsewhere. [CSCdi17624]

Novell IPX, XNS, and Apollo Domain

Novell IPX echoes (ping) do not work to the router's own IPX address on certain kinds of interfaces. [CSCdi16558]

XNS echoes (ping) do not work to the router's own XNS address on certain kinds of interfaces. [CSCdi16567]

Novell IPX watchdog-spoof requires no novell route-cache to be configured on the serial interface doing the spoofing. When configuring watchdog-spoof when fastswitching is disable via software but not disabled through the configuration editor there can be confusion as show Novell interface says fastswitching is already disabled.

Changes made are:

If no novell route-cache has not been configured and watchdog-spoof is attempted:

% IPX fast switching must be disabled Please configure NO NOVELL ROUTE-CACHE

will be displayed instead of just

% IPX fast switching must be disabled first

The Show novell interface display of

NOVELL Fast switching is enabled

changes to

NOVELL Fast switching is configured (enabled)

or

NOVELL Fast switching is configured (disabled)

or

NOVELL Fast switching is not configured

depending on the circumstances. In 9.21 "Fast" is replace with "Autonomous" if autonomous switching is enabled. [CSCdi16953]

Minor modification required in the novell fast switching code [CSCdi17120]

Configuring **ipx network xxx encapsulation hdlc** writes a redundant IPX network number in the NVRAM [CSCdi17286]

The OS/2 Requestor has a bug where it sends a GetNearestServer request with a tc field of 1, it should be 0. We were categorizing this as a bad novell packet and dropping it. [CSCdi17384]

When wishing to delete multiple IPX networks on the same network interface (an interface with both primary and secondary networks), a problem has been seen when they are deleted with a single **no ipx network** command. The workaround is to delete the networks individually with a series of **ipx network nnn** commands, one for each IPX network on the interface.

This problem has been fixed in a future release. [CSCdi17433]

In 9.21(1.5) a problem was introduced in low end (2000s, 3000s, 4000s) IPX fast-switching where the 802.3 MAC length field was being set incorrectly, certain PC IPX drivers would count these packets as errors, others would accept them. Novell IPX SAP, SNAP, and NOVELL-ETHER encapsulations are affected. This problem exists only in 9.21(1.5) -- any other release does not have it. [CSCdi17645]

IPX static servers attached to static routes are not flashed out immeditely as available when a **no shutdown** command is done on an interface, but instead wait for the normal per-minute update. Similarly, the static servers attached to static routes are not flashed out as poisoned immediately, but were only poisoned in the next SAP periodic update.

This may cause a non-fatal one minute delay in statically defined server availability when static servers are configured, but has no other functional impact. [CSCdi17768]

Protocol Translation

X.25 VC's for PAD calls would sometimes be put into hold-down if **hold-vc-timer** is configured. This could cause random system hangs. This behavior has been corrected.

Work-around and correct configuration is that **hold-vc-timer** should not be configured on X.25 interfaces that might make PAD calls, as on a system running protocol translation software. [CSCdi16965]

TCP/IP Host-Mode Services

When using the printer option for a TCP-LAT translation, one packet erroneously remains in the input queue on the receiving interface for each translation attempt which fails. [CSCdi17681]

VINES

The "vines time xxx" configuration commands are written to non-volatile memory even when vines routing has been disabled. This should not affect the operation of the router in any way. [CSCdi16734]

On an VINES interface that is configured for change updates only, the news of routers timing out is not propagated. [CSCdi16946]

Connectivity problems can occur on multipoint frame relay interfaces when VINES is configured to use "change only" updates. These problems can be eliminated by the use of point-to-point frame relay sub-interfaces, or by disabling vines "change only" updates on multipoint interfaces. [CSCdi17096]

If a ping or trace is attempted to a server that is inaccessible, but has a metric of infinity in the routing tables, it will not be possible to complete a ping or trace to that server for four hours. This happens because the first attempt creates a connection timeout value based upon the infinity metric, approximately 4 hours. [CSCdi17432]

Disabling vines split-horizon does not allow VINES Streettalk broadcasts to be forwarded out an interface that they were received on. This will break "hub-and-spoke" frame relay networks since spoke Streettalk broadcasts will not be forwarded from the hub router to other spoke sites. [CSCdi17488]

Vines does not correctly restart on an interface if it is disabled while the interface is up, and then re-enabled while the interface is shutdown. [CSCdi17708]

Wide-Area Networking

TCP header compression over X.25 is not working properly; the encapsulation VC is established for IP traffic and the datagrams are sent using IP. [CSCdi16322]

There is not way to change the amount of time that the cisco PAD code will wait for a repsonse to its x.29 invitation to clear message before clearing the x.25 call. The default of five seconds is too short for some applications. A **x29 inviteclear-time** *seconds* configuration command should be added to the system in future releases (9.21 and later.) [CSCdi16491]

Frame-relay subinterfaces go into administrative shutdown state for a few minutes upon reload, even though the subinterfaces are configured to be up in NVRAM. By the time it comes up OSPF has tried to allocate a router ID and failed. this causes the router to lose the **router OSPF** statment upon reload and consequently lose its OSPF connectivity. [CSCdi16504]

PAP debug messages are misleading. The message stated "remote passed PAP authentication". It was not clear if we authenticated or the remote authenticated. The new messages state whether or not we sent or received the Auth-Ack. [CSCdi16546]

Under some circumstances, the cisco PAD code will fail to actually send an x.29 invitation-to-clear message, even though the output from **debug pad** indicates that one is being sent. [CSCdi16630]

There is a slight chance of the system crashing if a PVC is being used by one protocol and this one protocol is disabled at the same time the system is checking to see if an inverse ARP request should be sent. [CSCdi16672]

On a frame relay interface configured as a point-topint subinterface, the system is looking at inverse ARP requests and using the contents to build map entries. This is unnecessary on a point-to-point interface and can cause confusing map entries. A point-to-point subinterface has a map entry that covers all protocols. [CSCdi16678]

The system sends out inverse ARP requests on PVCs which are configured as point-to-pint llogical interfaces. Since the inverse ARP reply isn't needed and is ignored, there's no need for the request. The only impact is a very small loss of band width due to the inverse ARP packets. [CSCdi16714]

When the **no frame-relay switching** command is invoked, the individual **frame-relay route** statements can no longer be added or removed. [CSCdi16723]

Frame relay interfaces can be configured as DCE or NNI only when frame relay switching is enabled. However, when frame relay switching is disabled, DCE or NNI interfaces can be changed

back to DTE only by issuing the command **frame-relay intf-type dte** or **no frame-relay intf-type dte**. **no frame-relay intf-type dce** and **no frame-relay intf-type nni** should be accepted as well. [CSCdi16724]

When running AppleTalk over PPP, if AppleTalk is configured after the interface is up, it is not negotiated. Workaround is to shutdown the interface and take it out of shutdown after appletalk is configured. [CSCdi17016]

The Frame-relay MIB variable frDlcmiMaxSupportedVCs now returns the correct value for the number of DLCI's supported. [CSCdi17182]

X.25 Calls received on a serial interface cannot be routed to a CMNS host. [CSCdi17212]

When switching frame relay PVC, statistics show that packets are slow switched when they actually are fastswitched. Statistics should show that switched (not routed) frame relay packets are fastswitched. [CSCdi17268]

The configuration commands to shut down a logical interface are not correctly displayed when issuing **write term**. The **shutdown** command does shut the interface down but the information is not saved. The workaround is to load the configuration files from a tftp server and add the needed shutdown command in this file. [CSCdi17274]

pap authentication if configured on one side only does not work properly. The workaround for this is to configure pap authentication on both ends. [CSCdi17282]

Priority queueing is not supported on X.25 interfaces. [CSCdi17357]

An X.25 interface's X.121 address will be set to the default address when an IP address is defined for it. This should not occur on interfaces that are not operating as a DDN or BFE device. [CSCdi17374]

The **no atm-dxi map** *protocol protocol-addr* does not require VPI and VCI as arguments. They are inherently set to zero and so, the command must bypass the check for non-zero values which is applicable only when maps are configured, not when they are removed. [CSCdi17375]

Debug frame-relay-packet does not display the protocol type for link type of Inverse Arp. [CSCdi17480]

A frame relay DCE or NNI interface can send erroneous STATUS messages to a FR DTE device. This happens when a static map is configured and then deleted on the DCE/NNI interface. The symptom is that the DTE device receives STATUS messages that contain garbage. The work around is to change the encapsulation on the DCE/NNI interface and then reconfigure frame relay (and not configure any static map). [CSCdi17534]

The **frame-relay inverse arp** command is not displayed when used on a subinterface. The PVC for subinterfaces is defined correctly though. For a regular frame-relay interface, the protocols for which the interface was not defined are displayed in the configuration. [CSCdi17563]

A cisco 2500 series router with an ISDN Basic Rate Interface can support VN3 switchtype for France.

Also, initial support for New Zealand and Norway which are very similar to the NET3 switchtype. The switchtype definitions are basic-nznet3 and basic-nwnet3. [CSCdi17579]

When a multicast DLCI is provided by the frame-relay switch via the LMI, the router does not correctly maintain the status of the multicast DLCI. On input packets the check for a valid DLCI is not being done correctly, and packets recieved on the multicast DLCI are being dropped. [CSCdi17583]

When the router is configured as a frame relay switch, some protocols (such as Novell) are process switched when all packets should be fastswitched regardless of the protocol. The symptom is that for these protocols, the packet rate is much lower than for say IP. The workaround is to configured

maps or encapsulation as IETF. [CSCdi17604]

When routing Calls, spurious duplicate VCs may be created; these duplicates show up in state P1 and as a "Half-baked connection". [CSCdi17606]

When an async interface is configured for both SLIP and demand dialing (with the **dialer in-band** command), the link will dial correctly but packets will never be transmitted across the link. **debug ip packet** will show each packet failing encapsulation. [CSCdi17609]

When operating as a DDN or BFE attachment, X.25 may not correctly associate a VC with the translated IP address. More than one VC to a given host may result. [CSCdi17631]

If lcp is up and LQM is configured (ppp quality x) the link will go into a CONFREQ/CONFACK loop. The workaround is to shut / no shut the interface. [CSCdi17690]