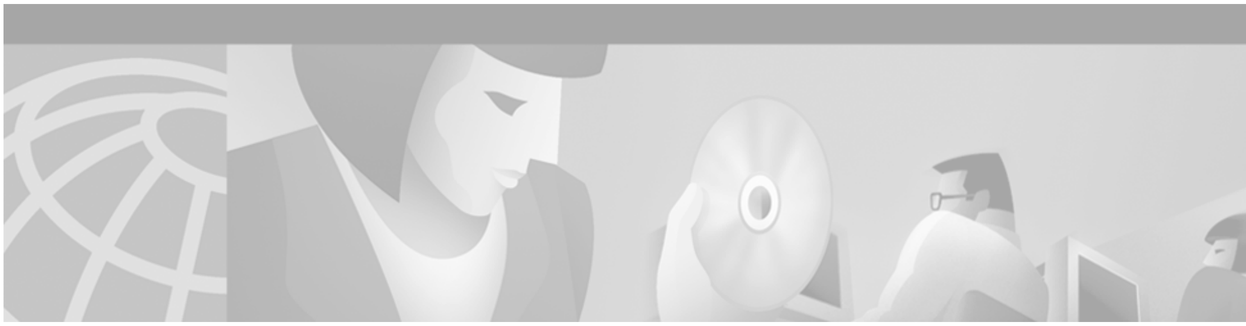

Release Notes for Cisco ONS 15454 SDH Release 3.4.1

November, 2002

Release notes address closed (maintenance) issues, caveats, and new features for the Cisco ONS 15454 SDH multiplexer. For detailed information regarding features, capabilities, hardware, and software introduced with this release, refer to Release 3.4 of the *Cisco ONS 15454 SDH Installation and Operations Guide, and Cisco ONS 15454 SDH Troubleshooting and Reference Guide.* For the most current version of the Release Notes for Cisco ONS 15454 SDH Release 3.4.1, visit the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/ong/15400/454relnt/index.htm

Cisco also provides Bug Toolkit, a web resource for tracking defects. To access Bug Toolkit, visit the following URL:

http://www.cisco.com/cgi-bin/Support/Bugtool/launch_bugtool.pl

Contents

Cisco Systems

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Changes to the Release Notes

This section documents supplemental changes that have been added to the *Release Notes for Cisco ONS 15454 SDH Release 3.4.1* since the production of the Cisco ONS 15454 SDH System Software CD for Release 3.4.1.

The following changes have been added to the release notes for Release 3.4.1.

Changes to Caveats

The following caveat has been added.

Transmission Control Protocol Specification, page 6

Caveats

Review the notes listed below before deploying the ONS 15454 SDH. Caveats with DDTS tracking numbers are known system limitations that are scheduled to be addressed in a subsequent release. Caveats without DDTS tracking numbers are provided to point out procedural or situational considerations when deploying the product.

Hardware

DDTS # CSCdy00622

Very rarely, an Equipment Failure Alarm can occur on an externally timed TCC+ or TCC-I card after a reset. If this occurs, BITS will be displayed as good for one TCC, but bad for the other. If the issue occurs on the standby TCC, a second reset could clear the problem. If the issue occurs on the active TCC, the card must be replaced. This issue is under investigation.

DDTS # CSCdw92634

SDH electrical cards only support a VC4 J1 trace string setting for all VC4s. You cannot set the J1 byte for individual VC4s. When the J1 byte is set for a single VC4, all other VC4s in the same line card will be set to the same value. This issue will be resolved in Release 4.0.

DDTS # CSCdw14501

Interconnection Equipment failure alarms may be generated at 55 degrees C, and 72 volts. When the operating environment is at 55 degrees C and 72 volts, interconnection equipment failure alarms for the following cards can occur:

- STM16SH
- STM64LH
- STM16LH
- XC10G

The alarms could potentially occur on any BTC192 board: OC48AS, GigE, OC192 or OC192LR. This issue will be resolved in Release 4.0.

DDTS # CSCdw65251

Recovery times in excess of 60 ms are possible for an E3 or DS3 circuit when there is a disruption on the fiber span. This occurs when you either remove or soft-reset active span cards. This issue will be resolved in Release 4.0.

DDTS # CSCdw50903

E1 boards with second source components can incur bit errors under extreme environmental conditions. When these boards operate under voltage and temperature stress conditions and a temperature ramp rate of 1 degree per minute, the boards could exhibit dribbling bit errors at high temperatures: BER = 5.5e-6. To avoid this, you must apply the temperature ramp rate at 0.5 degree per minute. This ramp rate complies with the NEBS standard; however, this issue will be revisited in Release 4.0.

Line Cards

DDTS # CSCdy56366 and CSCdy12392

With a Linear Multiplex Section Protection (LMSP) setup with STM-16, STM-64, or STM4-4 cards, when a protection switch occurs, the MS PSC and PSD fields on the STM-16 performance pane do not increment. This issue will be resolved in a future release.

DDTS # CSCdy65482

On the AIC-i card, a volume adjustment on the receive value of a four-wire orderwire circuit will be displayed as the negative of its actual value. To work around this issue, enter the negative of the value you actually want for the receive value. For example, adjust the receive value on CTC to -2 dbm for a gain of 2 dbm. This issue will be resolved in Release 4.0.

SONET and SDH Card Compatibility

Tables 1, 2, and 3 list the cards that are compatible for the ONS 15454 SONET and ONS 15454 SDH platforms. All other cards are platform specific.

Table 1 SDH Data Cards that are SONET Compatible

Product Name	Description
15454E-G1000-4	4 port Gigabit Ethernet Module - need GBICs
15454E-E100T-12	12 port 10/100BT Ethernet Module
15454E-E1000-2	2 port Gigabit Ethernet Module - need GBICs

Table 2 SONET Data Cards that are SDH Compatible

Product Name	Description
15454-G1000-4	4 Port Gigabit Ethernet
15454-E100T-G	10/100BT, 12 Circuit, compatible w/ XC, XCVT and XC-10G
15454-E1000-2-G	Gigabit Ethernet, 2 circuit., GBIC - G

Table 3 Miscellaneous Compatible Cards and Components

Product Name	Description
15454-BLANK	Empty slot Filler Panel
15454-GBIC-LX	1000Base-LX, SM or MM, standardized for 15454/327
15454-GBIC-SX	1000Base-SX, MM, standardized for 15454/327
15454-FIBER-BOOT=	Bag of 15 90 degree fiber retention boots

DDTS # CSCdw44431

No HP-PLM alarms are raised on ONS 15454 SDH optical cards. Because there is no way to set the expected payload label, the optical card cannot anticipate a particular label.

The card accepts all payload labels except for unequipped without raising an HP-PLM alarm. To avoid confusion, note that if a signal is terminated on an electrical card, the card raises a PLM if the label is not correct for the terminated traffic. It is not known at this time when or if this issue will be resolved.

DDTS # CSCdw80652

When one traffic card in a 1:N protection group is reset, and then another card is reset, there will be a loss of traffic on the second card, after the first card completes its reset, lasting until the second card completes its reset. This only occurs when the protect card tries to handle the traffic of a card that is resetting, and that card is carrying traffic because when it reset the protect card was carrying traffic for another card. This loss of traffic occurs because the protect card attempts to set its relays to handle the traffic of the working card, but the relays on the working card are also set to carry the traffic, and since the card is resetting, no software is running to switch its relays. This issue most frequently presents itself when testing a double-failure scenario: resetting two cards in a protection group. Wait until the first card completes its reset sequence before resetting the second card to prevent this problem. Configuring cards in 1:1 instead of 1:N protection should also avoid the problem. This issue will not be resolved.

DDTS # CSCdw57215

In a configuration with STM16 Any Slot cards and an VC4-8c circuit, provisioned between G1000-4 cards with traffic going over the STM16 span, extracting the G1000-4 card at one end of the VC4-8c circuit before deleting the circuit will result in a traffic hit on all existing SDH circuits defined over that same span.

XC10G Boot Process

If you install a new XC10G card to the node and it fails to boot, remove the card and reinsert it. If the card still fails to boot, return it using the RMA procedure. This issue will be resolved in future hardware.

Jitter Performance with XC10G

During testing with the XC10G, jitter generation above 0.10 UI p-p related to temperature gradient testing has been observed. This effect is not expected to be seen under standard operating conditions. Changes are being investigated to improve jitter performance in a subsequent version of the XC10G cross connect card. DDTS numbers related to this issue include CSCdv50357, CSCdv63567, CSCdv68418, CSCdv68441, CSCdv68389, CSCdv59621, and CSCdv73402.

E Series and G Series Cards

DDTS # CSCdy32536

No PDI-P alarm is raised against an STM-64 card upon a G1000 circuit failure. Note, however, that the PDI-P alarm is raised against the terminating G1000 card. The PDI-P condition on the terminating card may be used for root cause analysis. This issue is under investigation.

E1000-2/E100T

Do not use the repair circuit option with provisioned stitched Ethernet circuits. It is not known at this time when or if this issue will be resolved.

Single-card EtherSwitch

Each E100/E1000 card can be configured as a single-card EtherSwitch configuration to allow VC4-4c of bandwidth to be dropped at each card. The following scenarios for provisioning are available:

VC4-4c

VC4-2c, VC4-2c

VC4-2c, VC4, VC4

VC4, VC4, VC4, VC4

When configuring scenario 3, the VC4-2c must be provisioned before either of the VC4 circuits.

Multicard EtherSwitch

When deleting and recreating Ethernet circuits that have different sizes, you must delete all VC4 circuits provisioned to the EtherSwitch before you create the new circuit scenario. (See the preceding “Single-card EtherSwitch” section on page 6 for details on the proper order of circuit creation.) Enable front ports so that the VLANs for the ports are carried by the largest circuit first. A safe approach is to enable the front port before you create any circuits and then retain the front port VLAN assignment afterwards. If you break the rules when creating a circuit, or if you have to delete circuits and recreate them again, delete all circuits and start over with the largest first.

Maintenance and Administration

> **Caution** VxWorks is intended for qualified Cisco personnel only. Customer use of VxWorks is not recommended, nor is it supported by Cisco's Technical Assistance Center. Inappropriate use of VxWorks commands can have a negative and service affecting impact on your network. Please consult the troubleshooting guide for your release and platform for appropriate troubleshooting procedures. To exit without logging in, enter a Control-D (hold down the Control and D keys at the same time) at the Username prompt. To exit after logging in, type "logout" at the VxWorks shell prompt.

Transmission Control Protocol Specification

A vulnerability in the Transmission Control Protocol (TCP) specification (RFC793) has been discovered by an external researcher. The successful exploitation enables an adversary to reset any established TCP connection in a much shorter time than was previously discussed publicly. Depending on the application, the connection might be automatically reestablished. In other cases, a user must repeat the action (for example, open a new Telnet or SSH session). Depending on the attacked protocol, a successful attack might have consequences beyond terminated connection that also must be considered. This attack vector is only applicable to those sessions that terminate on a device (such as a router, switch, or computer) and not to those sessions that only pass through the device (for example, transit traffic that is being routed by a router). Also, this attack vector does not directly compromise data integrity or confidentiality.

All Cisco products that contain TCP stack are susceptible to this vulnerability.

This advisory is available at http://www.cisco.com/warp/public/707/cisco-sa-20040420-tcp-ios.shtml, and describes the vulnerability as it applies to Cisco products that run Cisco IOS® software.

A companion advisory that describes the vulnerability for products that do not run Cisco IOS software is available at http://www.cisco.com/warp/public/707/cisco-sa-20040420-tcp-nonios.shtml.

This issue is resolved in Releases 2.3.5, 4.1.4 and 4.6.2.

DDTS # CSCdz36843

When a power supply is connected on one port only, the Power Failure alarm is reported in reverse. For example, if power is connected at the A side, the alarm is "PWR-A failure." This issue is resolved in Release 4.0.

DDTS # CSCdz35812

When using the Defaults Editor (in the Node view, Provisioning tab of CTC), an exception may be thrown after which the following values cannot be set using the Defaults Editor:

- STM64.pmthresholds.rs.nearend.15min.BBE
- STM64.pmthresholds.rs.nearend.15min.EB
- STM64.pmthresholds.rs.nearend.15min.ES
- STM64.pmthresholds.rs.nearend.15min.SES
- STM16.pmthresholds.rs.nearend.15min.BBE
- STM16.pmthresholds.rs.nearend.15min.EB

- STM16.pmthresholds.rs.nearend.15min.ES
- STM16.pmthresholds.rs.nearend.15min.SES
- STM4.pmthresholds.rs.nearend.15min.BBE
- STM4.pmthresholds.rs.nearend.15min.EB
- STM4.pmthresholds.rs.nearend.15min.ES
- STM4.pmthresholds.rs.nearend.15min.SES
- STM1.pmthresholds.rs.nearend.15min.BBE
- STM1.pmthresholds.rs.nearend.15min.EB
- STM1.pmthresholds.rs.nearend.15min.ES
- STM1.pmthresholds.rs.nearend.15min.SES

These thresholds may still be set from the card view of CTC. This issue will be resolved in a future release.

DDTS # CSCdy10030

CVs are not positively adjusted after exiting a UAS state. When a transition has been made from counting UAS, at least 10 seconds of non-SES must be counted to exit UAS. When this event occurs, Telcordia GR-253 specifies that CVs that occurred during this time be counted, but they are not. There are no plans to resolve this issue at this time.

DDTS # CSCdy65172

No high order path (HP) TCAs are reported for E3 ports. HP PMs are correctly reported for all E3 ports. However, Threshold Crossing Alarms are only generated for Port 1. To work around this issue, examine the PM to see if there are HP errors on an E3 card. This issue will be resolved in a future release.

DDTS # CSCdy63135

If a mixed protection situation arises within the network and CTC freezes while you are routing a circuit automatically, end the CTC session, then restart CTC and route manually. Note that this is only likely to occur if you route automatically from source to destination when there is no bandwidth available from the chosen source to the chosen destination, a loop exists within the network with mixed protection, and the destination exists outside of the loop. This issue will be resolved in Release 4.0.

DDTS # CSCdy21992

Rarely, an STM4-4 card configured to measure PMs could incorrectly mark the previous 1 day PM counts as yellow, or invalid. This issue will be resolved in a future release.

DDTS # CSCdx35561

CTC is unable to communicate with an ONS 15454 SDH that is connected via an Ethernet craft port. CTC does, however, communicate over an SDCC link with an ONS 15454 SDH that is Ethernet connected, yielding a slow connection. This situation occurs when multiple nodes are on a single Ethernet segment and the nodes have different values for any of the following features:

- Enable OSPF on the LAN

- Enable Firewall
- Craft Access Only

When any of these features are enabled, the proxy ARP service on the node is also disabled. The ONS 15454 SDH proxy ARP service assumes that all nodes are participating in the service.

This situation can also occur immediately after the aforementioned features are enabled. Other hosts on the Ethernet segment (for example, the subnet router) may retain incorrect ARP settings for the ONS 15454 SDHs.

To avoid this issue, all nodes on the same Ethernet segment must have the same values for Enable OSPF on the LAN, Enable Firewall, and Craft Access Only. If any of these values have changed recently, it may be necessary to allow connected hosts (such as the subnet router) to expire their ARP entries.

You can avoid waiting for the ARP entries to expire on their own by removing the SDCC links from the affected ONS 15454 SDH nodes. This will disconnect them for the purposes of the proxy ARP service and the nodes should become directly accessible over the Ethernet. Network settings on the nodes can then be provisioned as desired, after which the SDCC can be restored.

This issue is under investigation.

DDTS # CSCdy38603

VT Cross-connects downstream from a E1 can automatically transition from the OOS-AINS state to the IS state even though the E1 signal is not clean (for example, when there is an LOS present). This can occur when you have created a VT circuit across multiple nodes with E1s at each end, and you have not yet applied a signal to the E1 ports, and then you place the E1 ports in OOS-AINS, OOS-MT, or IS. When you then place the circuit in OOS-AINS, the circuit state changes to IS (within one minute). This issue will be resolved in a future release.

DDTS # CSCdy11012

When the topology host is connected to multiple OSPF areas, but CTC is launched on a node that is connected to fewer areas, the topology host appears in CTC, and all nodes appear in the network view, but some nodes remain disconnected. This can occur when the CTC host does not have routing information to connect to the disconnected nodes. (This can happen, for example, if automatic host detection was used to connect the CTC workstation to the initial node.)

CTC will be able to contact the topology host to learn about all the nodes in all the OSPF areas, but will be unable to contact any nodes that are not in the OSPF areas used by the launch node. Therefore, some nodes will remain disconnected in the CTC network view.

To work around this issue, if no firewall enabled, then the network configuration of the CTC host can be changed to allow CTC to see all nodes in the network. The launch node must be on its own subnet to prevent network partitioning, and craft access must not be enabled. The CTC host must be provisioned with an address on the same subnet as the initial node (but this address must not conflict with any other node in the network), and with the default gateway of the initial node. CTC will now be able to contact all nodes in the network.

If a firewall is enabled on any node in the network, then CTC will be unable to contact nodes outside of the initial OSPF areas. This issue is under investigation.

DDTS # CSCdy12392

In a Linear Multiplex Section Protection (LMSP) setup with STM-16, STM64 or STM4-4 cards, PSC and PSD counts do not increment after a protection switch. This issue will be resolved in a future release.

DDTS # CSCdy47232

If an F-UDC is setup on a 1+1 protection path or an F-UDC circuit is changed to 1+1, the path does not function. Note that MS-UDC functions correctly, but F-UDC will not work with 1+1. This issue will be resolved in a future release.

DDTS # CSCdy57891

An LOP-P alarm can be inadvertently cleared by an LOS that is raised and cleared. On STM-N cards, when an LOP condition and an LOS condition are both present on the input, an LOS will be raised as per Telcordia GR253 alarm hierarchy. However, upon clearing the LOS with the LOP still present, the LOP alarm is not raised. An AIS-P condition will be visible. This issue will be resolved in a future release.

DDTS # CSCdy53835

When you delete an overhead AIC-I circuit that crosses an SDCC enabled link, the OSPF area IDs of both SDCCs on the link are lost. If this issue occurs, in the SDCC provisioning tab, click Edit and uncheck the Enable OSPF checkbox on both sides of the link. Return to the SDCC provisioning tab, click Edit once more, and check the Enable OSPF checkbox on both sides of the link again. This will restore the default OSPF area ID for both SDCCs. This issue will be resolved in Release 4.0.

DDTS # CSCdy46980

If you attempt to reseat both working and protect trunk cards at the same time while there are orderwire circuits running through a trunk card that is being reseated, the AIC-I card will sound an alarm until the booting process completes. To avoid this, ensure that any trunk card being reseated is protected. This issue will be resolved in a future release.

DDTS # CSCdw38283

If a node has one good BITS reference and is running in a normal state, and you configure a second BITS reference, then reconfigure the second reference within 30 seconds of applying the first configuration, the node will enter FAST START SYNC mode. To avoid this problem, wait a minute before configuring the second reference a second time. This issue is a hardware limitation, and there are no current plans to resolve it.

DDTS # CSCdw23208

The following table summarizes B1, B2, and B3 error count reporting for SDH optical cards. Note that not all reporting is done according to ITU specifications. In particular, ITU specifies error counts for B1 and B3 as the number of blocks with errors (refer to ITU-T G.826 for paths and ITU-T G.829 for RS and MS).

Table 0-4 Error Count Reporting

	B1	B2	B3
ITU Specification	block	bit	block
STM1	block	bit	block
STM4	bit	bit	bit
STM16 trunk	bit	bit	bit
STM16 AS	block	bit	bit
STM64	block	bit	bit

DDTS # CSCdw82689

After creating 509 VLANs and provisioning many Ethernet circuits, Ethernet circuit provisioning can become very slow, or possibly fail. Ethernet traffic may also incur an outage of a few minutes. To avoid this problem, delete any VLANs that are created but not used, and do not recreate them. There is no resolution planned for this issue.

DDTS # CSCdw47506

A CTC communications failure on the network during circuit creation can cause a "Circuit Provisioning Error" exception. An attempt to continue with the errored circuit creation results in other exceptions that occur repeatedly on each attempt to continue. This issue has been seen infrequently, and only on large networks. To correct the problem, abandon the attempted circuit creation and start over. This issue will be resolved in a future release.

DDTS # CSCdv10824: Netscape Plugins Directory

If you use CTC, JRE, and the Netscape browser with a Microsoft Windows platform, you must ensure that any new installation of Netscape uses the same Netscape directory as the previous installation did, if such an installation existed. If you install Netscape using a different path for the plugins directory, you will need to reinstall JRE so that it can detect the new directory.

"Are you sure" Prompts

Whenever a proposed change occurs, the "Are you sure" dialog box appears to warn the user that the action can change existing provisioning states or can cause traffic disruptions.

MS-SPRing Functionality

DDTS # CSCdy65890

If you have PCA circuits over two-fiber or four-fiber MS-SPRing protect channels, an incorrect auto-inservice transition occurs after traffic preemption. You may place the circuit back into the OOS-AINS state after the BLSR has returned to the unswitched mode, using the Circuit Editing pane of the CTC. This issue will be resolved in a future release.

DDTS # CSCdy63060

In a 4 node, two-fiber MS-SPRing configuration, the E100 unstitched circuit state can become stuck at OOS-AINS-PARTIAL, even if there are no alarms and conditions raised.

This issue has been seen under the following conditions:

Step 1 Set up a 4 node, two-fiber MS-SPRing.

Step 2 Provision an E100 point to point circuit starting with the OOS-AINS state and the longer

Step 3 path as the working path. The working path should have at least one pass-through node.

Step 4 Ensure that Ethernet ports and OC-N ports are all in service, no alarms or conditions are raised, and traffic is running clear.

If the state does not change automatically, use the Circuit Edit Window to explicitly set the circuit state to IS. This issue will be resolved in Release 4.0

DDTS # CSCdy62992

The E3 circuit state does not transition correctly when the destination drop is in the OOS state.

For example:

Step 1 Create a 3 node, two-fiber, OC-48 MS-SPRing.

Step 2 Provision E3 cards as drop nodes. The destination drop card ports should be provisioned as OOS.

Step 3 Create a port group circuit selecting the OOS_AINS circuit state. The circuit states will transition to IS for all three VC3 circuits and the VCT.

This issue will be resolved in a future release.

DDTS # CSCdy56668

Ethernet circuits may appear in the CTC circuit table with an INCOMPLETE status after an MS-SPRing span is upgraded. The circuits, when this occurs, are not truly incomplete. They are unaffected and continue to carry traffic. To see the circuit status correctly, restart CTC. This issue is under investigation.

DDTS # CSCdy55349

Rarely, some DS3i cards may fail to carry traffic after a power cycle. This can be seen when power cycling an entire chassis, inserting an unprotected DS3i card, or protection switching to a DS3i card that has not carried traffic since it was powered up. Executing a cross connect protect switch will restore traffic. This issue will be resolved in a future release.

DDTS # CSCdx76262

In a two fiber MS-SPRing, an E3 traffic loss can exceed 60 ms upon an optical switch or XC10G side switch. This can occur whenever an optical switch or XC10G side switch occurs, even on a passthrough node; however, it does not occur consistently. This issue will be resolved in a future release.

DDTS # CSCdy59242

Under some circumstances, if a fail-to-switch alarm is raised upon introducing SF-R with the existing Lockout Span command, the alarm becomes stuck after the SF-R and Lockout Span are cleared.

The following example illustrates how this can occur.

In a two fiber MS-SPRing, say the east side of Node 1 is connected to the west side of Node 2.

Step 1 Perform a Lockout Span on the east side on Node 1.

Step 2 Remove the transmit fiber on the east span of Node 1. (Node 2 detects signal failure on its west side.) Traffic is lost, as expected, due to the Lockout Span on the ring. A Fail-to-Switch alarm is raised.

Step 3 Re-insert the transmit fiber. Traffic comes back, but the fail-to-switch alarm is still reported.

Step 4 Clear the Lockout Span. The Fail-to-Switch alarm becomes stuck.

The issue is that Node 2 ignores a long-path Lockout Span on its east side and initiates a ring switch with a local SF-R request, then fails.

To avoid this issue, make sure the ring is in the idle state and issue an Exercise Ring command on the span that reports Fail-to-Switch alarm to clear that alarm. This issue will be resolved in Release 4.0.

DDTS # CSCdv89939 and CSCdy46597

After a MS-SPRing span or ring switch, traffic is switched to a different set of nodes and a protection STS is used. At this point, any ongoing J1 monitoring does not follow the switch. As a result, there is no J1 monitoring on the protection path. If there is a mismatch of the J1 string on the protection path, the TIM_P alarm will not be raised. Also, you can retrieve the actual captured J1 string on the working path, but if MS-SPRing switched from working to protect, you cannot retrieve the J1 string on the protect path. MS-SPRing support for J1 trace is a feature request that will be addressed in a future release.

DDTS # CSCdy41904

After a ring switch (where PCA traffic is preempted), putting a PCA circuit in the out of service (OOS) state will not stop traffic flow for that circuit once the ring switch is cleared. To avoid this issue, delete the circuit or place the circuit in the IS, then the OOS state. This issue will be resolved in a future release.

DDTS # CSCdy10805

If you upgrade one of the rings in a two by two MS-SPRing configuration, an EXTRA-TRAF-PREEMPT alarm may be raised and subsequently fail to clear on one of the rings. If the ring that has the stuck alarm already has some PCA circuits on it, you can issue and then clear a Force Ring. This should clear the stuck alarm. If no PCA circuits exist on the ring, then create one temporarily, and follow the above procedure to clear the alarm. After the alarm clears, you can remove the Force Ring, then delete the PCA circuit. This issue will be resolved in Release 4.0.

DDTS # CSCdy30125

In a two by two MS-SPRing configuration, with PCA circuits passing through the common node, if one of the rings is a two fiber MS-SPRing and you upgrade it, a PCA connection will be promoted to become protected on the upgraded ring side. In this scenario, you can end up with a circuit that is PCA on one ring and protected on the other ring.

This can occur with any colliding STSs; in other words, any situation where the STS from the working side is going to overlay the STS from the protection side when a ring or span switch occurs. On a span switch in a four fiber MS-SPRing this would be STS #1 on the working and STS #1 on the protect on the same side (i.e. east or west). For a ring switch on a four fiber MS-SPRing it would be STS #1 on the working and STS #1 on the protect on the opposite side of the ring. In a two fiber BSLR there is only a ring switch, so the colliding STSs would be STS #1 on one side of the ring and STS #7 on the opposite side (for an STM-4 ring, for example). Symptoms of a failure will be protected traffic that is dropped or that has a stuck AIS-P.

When you perform a two fiber MS-SPRing upgrade in a two by two configuration, ensure that no PCA circuits cross through the common node before you start the upgrade. Note that the PCA circuits that are added and dropped on the same ring are safe, as they will be promoted to become fully protected. All PCA circuits that cross the common node to go to another ring must be deleted before the upgrade, then recreated once the upgrade is successfully finished. This issue will be resolved in Release 4.0.

DDTS # CSCdy35901

In a four-node, STM-64, four fiber MS-SPRing, traffic remains lost after a lockout is cleared on an adjacent node when the local node has an SF-R raised. To correct this problem if it occurs, issue a force ring on the side of the SF-R affected node that the SF-R is raised on. This issue will be resolved in Release 4.0.

DDTS # CSCdy37939 and CSCdy01642

In STM-4 MS-SPRing configurations, a WKSWPR alarm that occurs can take several seconds before it appears. The workaround is to simply wait for the alarm, which should appear after a brief delay. This issue will be resolved in a future release.

DDTS # CSCdy48872

Issuing a lockout span in one direction while a ring switch (SF-R) is active in the other direction may result in a failure to restore PCA circuits on the ring.

To see this issue, on a node participating in a two fiber MS-SPRing with PCA circuits terminating at the node over the two fiber MS-SPRing, cause an SF-R by failing the receive fiber in one direction (say, west). Then issue a lockout span in the other direction (in our example, east). Since the lockout span has higher priority than the SF-R, the ring switch should clear and PCA traffic should be restored on spans without a fiber fault. The ring switch does clear, but PCA traffic does not restore. To correct this issue, clear the fiber fault. All traffic restores properly. This issue will be resolved in Release 4.0.

DDTS # CSCdy48463

Protected traffic loss can occur when a PCA monitor circuit is created on one ring to monitor a protected circuit on another ring and the two rings switch. This can also occur on a common node when a circuit is PCA on one ring side and protected on the other side.

You can see this issue in one of two ways:

1. With a two-ring configuration, create a PCA monitor circuit on one ring to monitor a protected circuit on another ring, then set the monitor point to a trunk access point anywhere on that circuit. A ring switch on both rings can trigger the traffic loss.

2. In a two-MS-SPRing configuration, with PCA circuits passing through the common node, if one of the rings is a two fiber MS-SPRing and you upgrade it, then the previous PCA connection is promoted to become protected. The result is that now a protected circuit is connected to a PCA circuit, just as in the case of the monitor circuit, and the issue occurs.

This can occur with any colliding STSs; in other words, any situation where the STS from the working side is going to overlay the STS from the protection side when a ring or span switch occurs.

To avoid this issue, when you create PCA monitor circuits, create them to monitor the drop side of the connection and never monitor the trunk side. Also when you perform two fiber MS-SPRing ring upgrades, make sure that no PCA circuits cross through the common node before you start the upgrade. Any PCA circuits that are added and dropped on the same ring are safe, as they will be promoted to become fully protected. All PCA circuits that cross the common node to go to another ring must be deleted before the upgrade, then recreated once the upgrade is successfully completed. This issue will be resolved in Release 4.0.

DDTS # CSCdw53481

Two MS-Rs are not allowed to coexist. If you execute a manual ring switch command on one side of an MS-SPRing node and apply another manual ring switch command on other side of the node, the second manual ring switch command is rejected. This works as designed. The implementation complies with Telcordia GR-1230, R6-102.

DDTS # CSCdx45851

On a four fiber MS-SPRing, VC4-16c traffic fails switch after restoring the database to all nodes. This only happens when you are restoring the database for all nodes at the same time. The proper procedure for restoring the database for multiple nodes is to restore one node at a time. This procedure is documented in the user documentation.

DDTS # CSCdx19598

A rare hardware failure on an STM16AS card transmitter can trigger SEF on the receiving STM16AS card in a four fiber MS-SPRing (or BLSR) configuration. The BER calculations are suspended when SEF is detected, so SD or SF is never raised. Likewise SEF is not considered a signal failure condition like LOS or LOF, so a protection switch will not occur. If this occurs, use the CTC GUI to force a protection switch on the MS-SPRing (or BLSR). This issue will be resolved in a future release.

DDTS # CSCdv53427

In a two ring, two fiber MS-SPRing (or BLSR) configuration (or a two ring MS-SPRing or BLSR configuration with one two fiber and one four fiber ring) it is possible to provision a circuit that begins on one ring, crosses to a second ring, and returns to the original ring. Such a circuit can have protection vulnerabilities if one of the common nodes is isolated, or if a ring is segmented in such a way that two non-contiguous segments of the circuit on the same ring are each broken. This issue will be resolved in a future release.

Database Restore on an MS-SPRing (or BLSR)

When restoring the database on an MS-SPRing (or BLSR), follow these steps:

Step 1 To isolate the failed node, issue a force switch toward the failure node from the adjacent east and west nodes.

Step 2 If more than one node has failed, restore the database one node at a time.

Step 3 After the TCCi has reset and booted up, release the force switch from each node.

SNCP Functionality

DDTS # CSCdw66071

After a switch to protect is cleared for a revertive SNCP circuit, the WTR alarm is not raised, although the wait period is observed and the circuit reverts back to working. This issue will be resolved in Release 4.0.

Active Cross Connect or TCCi Card Removal

As in MS-SPRing (or BLSR) and 1+1, you must perform a lockout on SNCP (or UPSR) before removing an active cross connect or TCCi card. The following rules apply to SNCP (or UPSR).

Active cross connect cards should not generally be removed. If the active cross connect or TCCi card must be removed, you can first perform an XC10G side switch and then remove the card once it is in standby, or you can perform a lockout on all circuits that originate from the node whose active cross connect or active TCCi will be removed (performing a lockout on all spans will also accomplish the same goal). No lockout is necessary for switches initiated through CTC.

Resolved Caveats for Release 3.4.1

The following items are resolved in Release 3.4.1

Hardware

DDTS # CSCdw74751

An STM16 line card might exhibit signal fail after the unit is power-cycled twice or more. This issue can arise when multiple power cycles have occurred. The STM16 card displays a signal fail LED after these power cycles. CTC reports LOS on the affected card. When you replace that card with another STM16, traffic returns to normal. When you insert the affected STM16 card in another node within the circuit, it now also reports signal fail (and CTC reports LOS). This issue has only been seen once, and may be specific to the card it was seen on. This issue is resolved in Release 3.4.

Line Cards

DDTS # CSCdw60129

When a terminal loopback is placed on an E1 port, but the port is not in service, the circuit will not reliably carry traffic. This only occurs when a port is put into terminal loopback but the port is not put into service. Placing the port in service either before or after creating the terminal loopback will cause the circuit to pass traffic correctly. This issue is resolved by the enhanced state model in Release 3.4, which allows a terminal loopback when the port is in OOS-MT or OOS-AINS, but not in OOS.

DDTS # CSCdw80537

When the active TCCI and XC10G are removed simultaneously, the DS3i card might drop traffic. The issue does not arise if the second card is removed after the first card has completed switching. Also, this issue has only been reported when DS3i traffic entered the node through an SNCP optical ring. To recover traffic, reset or reseat the DS3i card. This issue is resolved in Release 3.4.

DDTS # CSCdx26746

When you insert an electrical card into an unprovisioned slot of a node with XC10G cross connect cards, in some cases, the card will fail to boot up and auto-provision itself. To recover in this situation, provision the card through the user interface. This issue is resolved in Release 3.4.

DDTS # CSCdx00506

After monitoring traffic for several days, B3 errors may be seen on E3, STM-1, STM-4 and Ethernet traffic with the XC10G cross-connect card. This issue was resolved in a subsequent build of the Release 3.3 software. It is also resolved in Release 3.4 and forward.

DDTS # CSCdw45637

Rarely, if an E1 card has been running for at least 12 hours and is in a 1:N protection group, the switch time to protection can take approximately 500 ms. This issue is resolved in Release 3.4.

E Series and G Series Cards

DDTS # CSCdx68385

Ethernet traffic on E series cards is subject to possible loss while performing a span upgrade from STM-16 (not including STM-16AS) to STM-64. To recover from such a loss, perform an XC10G switch. This issue was due to a bad power supply in the tester's setup. The bug report was later discarded as invalid.

DDTS # CSCdx53004

An STM-1 circuit is sometimes not allowed to be provisioned on a G1000-4 card if there are certain other circuits already existing on the same card. This can happen under one of two scenarios:

If a G1000 card already has some circuits which have been provisioned via a Release 3.2 image with some large

circuits (such as STM-8c, STM-4c or STM-3c) then, if new STM-1 circuits are attempted with a Release 3.3 image, these circuits may be disallowed.

Also, occasionally even if all circuits were provisioned by a Release 3.3 image but a few large circuits (like those above) were provisioned first then STM-1 circuits may be prevented from being provisioned.

In some cases similar symptoms may appear if the problem is due to a known initial hardware limitation (refer to the G1000-4 section of the user reference guide for details). The way to distinguish the two cases is that with the known hardware limitation the total sum of the circuit sizes of existing circuits and the new circuit has to be STM-12c or greater. If the total is less than STM-12c then you have this problem.

If, using the above test, you can determine with certainty that you have this problem, you can recover from it by deleting all the existing circuits on the affected card and then re-provisioning all of them, as well as the new circuit, in the order of smallest circuit size first. However, deletion of all existing circuits may not be necessary if you can delete existing circuits until the total provisioned bandwidth is STM-8c or less and then start re-provisioning circuits in order of smallest through largest. This issue is resolved in Release 3.4.

Throughput/Latency Testing

When testing the G1000-4 for latency/throughput at, near, or above the maximum allowable line rate per the guiding specifications, IEEE 802.3 and 802.3z. Customers testing for Throughput or Latency may see throughput calculations that can vary from 100% to 99.98% throughput, depending on the accuracy of the test set clock and the variability of the clock on the G1000-4. As described in the text below, the G1000-4 is fully compliant with the specification for line rate gigabit Ethernet. However, during testing in the lab environment, technicians need to be cognizant of the throughput settings and accuracy of the clock on the test set to ensure that the variances in throughput seen on the G1000-4 are not mistakenly perceived as being out of specification. Further, it needs to be understood that such testing is not reflective of traffic conditions that would be experienced in real world networks.

IEEE 802.3 allows for a variation in the clock rate of +/- 100 parts per million (ppm), allowing a range of speeds to be considered conforming to the specification.

The legal range of for Gigabit Ethernet is an follows:

- Minimum Speed—1,487,946 Frames Per Second
- Nominal Speed—1,488,095 Frames Per Second

- Maximum Speed—1,488,244 Frames Per Second

Conforming devices may not vary the preamble size, start frame delimiter size, or reduce the inter packet gap. The G1000-4 is fully compliant with these parameters.

During lab testing with a throughput testing device (Spirent Smartbits, Ixia test devices, etc.), because of a speed variance between the ingress packets from the external device and the egress speed from the G1000-4, throughput can vary from 100 percent to 99.98%, depending on the difference in clock speeds between the devices. Due to the allowable variation of clock tolerance, Some G1000 cards transmit below the nominal clock speed for Gigabit Ethernet, but well within the IEEE specification. In fact, although the specification allows for +/- 100ppm of tolerance, the oscillator on the G1000-4 has been found to vary only between +/- 40ppm on average (G1000-4 clock never runs below the minimum speed of 1,487,946 frames per second outlined in the IEEE specification). We guarantee the +/- 100ppm per the specification.

Short duration traffic bursts that are above the nominal rate are buffered, thus traffic isn't dropped for bursty traffic above the nominal rate. However, sustained traffic that is above wirespeed will be buffered and at some point the buffers will overflow can result in a nominal amount of dropped packets. The G1000 card will never drop a single frame with test equipment that is running at -100 ppm of line rate.

This issue can only be witnessed in a lab environment, as it would require all of the following conditions to occur simultaneously in a real network in order to cause frame loss.

Sustained traffic that is above the minimum clock speed possible. For example, if the clock on the G1000 was running -100 ppm or 1,487,946 frames per second, the sustained traffic would have to last 53.69 seconds in order to cause frame loss. This is because there is a 149 frame per second mismatch and we can buffer 8,000 64 byte frames.

Traffic patterns that are fixed frame sizes with a constant minimum Inter frame Gap. This is not real world traffic and can only be produced by high end test equipment.

This issue is resolved in Release 3.4.

DDTS # CSCdx05444

If a circuit already exists on a G1000-4 card, the provisioning of any subsequent circuit can cause bit errors up to 1 ms on the existing circuit(s). The only affected circuit size found in testing is VC4-8c. This issue is resolved in Release 3.4.

DDTS # CSCdw43919

When running traffic at 100% line rate with a repetitive data pattern, frame corruption and loss may occur for approximately one to three percent of the frames. This issue can occur when all of the following conditions are present:

- 100% line rate traffic.
- Attached device clock rate is greater than the G1000-4 clock rate (even if within the 100 ppm tolerance range of IEEE 802.3).
- Repetitive data patterns--this issue is not seen with random data patterns, even if the other conditions are met.

There are two ways to avoid this issue:

1. Use varying or random data frames for line rate performance measurements.

2. If fixed repetitive data patterns must be used for testing, use an “all zeros” data pattern in the frame, including all-zero source and destination MAC addresses. This pattern is known to not exhibit the problem. You can also experiment with varying the data patterns one bit at a time in order to determine other fixed patterns that will not exhibit the problem.

In summary, this problem will only be exhibited in test scenarios, for which the above workarounds can be used, and the probability of occurrence is extremely remote.

This issue is resolved in Release 3.4.

Maintenance and Administration

Performance Monitoring Using Cisco Transport Manager

In Release 3.4, Cisco Transport Manager users that performed PM retrievals might have encountered any or all of the following issues:

- G1000 statistics appearing unpredictably in the wrong fields
- Missing PM data
- Correct PMs falsely marked as invalid
- Incorrect PMs not marked as invalid

These issues were most likely to occur with SONET path data. SDH path data was unaffected. All of these issues have been resolved in Release 3.4.1.

DDTS # CSCdx53993

A less than 1 ms traffic disruption can occur when deleting a DS3i 1:1 or 1:N protection group. This issue is resolved in Release 3.4.

DDTS # CSCdw82921

The ITU specification requires that if working and protect optical cards both have an SD or SF, the traffic should pass on the working card. The SONET specification does not require this. The ONS 15454 SDH platform conforms to the SONET specification, but not the ITU specification. Specifically, when an SD or SF condition occurs on a working optical card, the ONS 15454 SDH switches traffic to the protect card. When an SD or SF condition then occurs on the protect card, traffic will stay on the protect card. The ITU specification requires that traffic should revert to the working card, even though that card has an SD or SF as well. There is no degradation of service, so no workaround is really needed; however, a force command can be issued to force the traffic onto the working card if desired. This issue is resolved in Release 3.4.

DDTS # CSCdw78048

The boot menu's host IP address is added to the node's routing table. The host IP is used when a node downloads a software image from a network in a development environment. This occurs whenever the manufacturer's default host IP is not null in the bootline (boot menu). In general, this is not a problem since the default host IP is a reserved IP address. However, an installation engineer may want to change the default host and gateway IP to a customer's host and gateway IP, or to null, by changing the bootline through the console. In Release 3.4, the host IP is not added to the routing table if a node is not booting from a network.

DDTS # CSCdw95301

When there are large numbers of VC4 circuits (greater than 100) and when there is a lot of circuit activity (for example, when there are a lot of updates), display of the Circuits pane can be slow and the user interface may be unresponsive for several minutes. This issue is resolved in Release 3.4.

DDTS # CSCdw52185

Symptom: User initiated switches (Manual or Force) are not cleared if overridden by higher priority switch requests. If a request is overridden by a higher priority request, it must be cleared manually in order to prevent possible switch oscillation. This issue is resolved in Release 3.4.

DDTS # CSCdw58748

In the “Audit Trail” under the Maintenance tab, there is a status field. Sometimes this field contains an “X” character; however, the X character is not documented in the user documentation. To understand the meanings for each result in the Pass/Fail column, please consult the following:

P = function completed successfully

F = function completed with failure

X = function did not complete

This issue is resolved in the Release 3.4 documentation.

DDTS # CSCdw45674

When there is a facility loopback in place, the E3 L3M chip will pass line traffic towards the backplane, and the facility loopback will then result in a terminal loopback. You cannot generate AIS towards the backplane in this scenario. The L3M chip does not support facility and terminal loopback at the same time. The workaround is to set loopback type equal to “none” before changing the loopback from facility to terminal and vice versa. This issue is resolved in Release 3.4.

DDTS # CSCdx03404

If you perform a manual switch from the working to the protect card, upon removing the protect card, traffic will switch back to working within 50 ms. When the protect card is replaced and the card has rebooted, the manual switch will trigger a traffic switch back to the protect card which can cause a service disruption on the order of several hundred ms. To avoid this service disruption, remove the manual switch command before replacing the protect card. The manual switch command is cleared upon detection of a higher priority failure (as per Telcordia GR-253) in Release 3.4.

DDTS # CSCdw71844

When a Force or Manual switch request is made while a higher priority request is present (in other words, SD/SF or Lockout), the user-initiated switch request will not be denied. This issue is resolved in Release 3.4.

DDTS # CSCdw64191

When testing throughput and latency of VC4-8c circuits on the G1000-4 card, Gigabit Ethernet utilization must be no more than 99.98%. If utilization exceeds this rate, an increase in latency will result. This is an unlikely scenario in a production network, considering dynamic frame sizes, patterns, utilization rates, and interframe gaps. This issue is resolved in Release 3.4.

DDTS # CSCdw03281

Under certain conditions, the CTC GUI freezes. To recover from this condition, you must restart CTC. This behavior has only been seen when all of the following conditions are met:

- Two sets of 6 nodes, each node connected to 4 of the other nodes in its set
- Circuits total at least 850
- Several operations occur over a short period
- JRE 1.2.2 is running on the workstation running CTC

This issue is resolved in Release 3.4.

DDTS # CSCct03396 Ring Map Change Dialog Box

When you add a node to an MS-SPRing, CTC displays a Ring Map Change dialog box asking you to accept the change. If you browse away from the node view before this dialog box has appeared, the dialog box may fail to appear, or may come up behind another window. This issue is resolved in Release 3.4.

DDTS # CSCdt94185

CTC can fail to drop user initiated switch requests (Manual or Force) when a higher priority request is initiated. This issue can arise when a switch request is made by the user and then another, higher priority request is made. CTC should preempt the user request with the higher priority request. If CTC fails to clear the request, manually clear the request. This issue is resolved in Release 3.4.

MS-SPRing Functionality

DDTS # CSCdx29951

Normally, a node is isolated by unidirectional failures on each of its working and protect spans; however, if this isolation fails to complete for some reason, passthrough traffic can be lost. To preserve traffic, issue a force ring from the two adjacent nodes, towards the problem node. Clear the force ring on each adjacent node once the isolated node recovers from failures. This issue is resolved in Release 3.4.

DDTS # CSCdx39743

High switch times can occur on STM1 traffic passing through MS-SPRings when the active cross connect card is pulled. Issue a manual switch of the cross connect cards, or reset the active cross connect card before removing it. Locking the spans prior to pulling the cross connect card also reduces switch time in this scenario. This issue is resolved in Release 3.4.

DDTS # CSCdx42388

Splitting an MS-SPRing by creating a signal failure ring (SF-R) on both the east and west spans causes the cross connect card and the TCC to reboot simultaneously. To avoid this, issue a force ring (FS-R) on the side that detects signal failure for each node affected. This issue is resolved in Release 3.4.

DDTS # CSCdx42436

Permanent failure of protect STM64 after both XCs self-boot can occur on a four fiber MS-SPRing with four nodes. If you introduce SF on both the working and protect spans on opposite links (to segment the ring into two smaller rings), the XCs and the TCC in Slot 7 reboot. This issue is resolved in Release 3.4.

DDTS # CSCdx42694

Promoting signal degrade span (SD-S) to signal degrade ring (SD-R) on the same node while a higher priority span request exists in the ring can cause traffic loss. For example, if you:

1. Introduce signal degrade span (SD-S) on the west working span.
2. Introduce signal degrade span (SD-S) on the east working span.
3. Introduce signal degrade protection (SD-P) on the east working span.

Then, to recover traffic to the working span on the east, you must issue a lockout protection span on the east; or, to recover traffic to the protect span on the east, you must issue a force span on the east.This issue is resolved in Release 3.4.

DDTS # CSCdx42710

A cross connect card reboot and traffic loss can occur on a four fiber MS-SPRing when signal degrade span (SD-S) is promoted to signal failure span (SF-S) while signal degrade ring (SD-R) exists in the other span. The following example illustrates the conditions that can cause this issue to occur:

1. Using variable attenuators, introduce SD-R on the east span (working & protect); by attenuating the working span first and the protect span later.
2. Introduce SD-S on the west working span using a variable attenuator.
3. Escalate SD-S to SF-S by increasing the attenuation; increase a bit further and you get LOF on the west working span.
4. In this situation, WTR is raised and cleared continuously on the west working span.
5. After a few minutes, the active XC in one node attempts to reboot and all traffic is lost and then regained, but the WTR problem persists.

To prevent signal oscillation on the west working span, issue a force span command on that span. Release the force when the fiber is replaced. This issue is resolved in Release 3.4.

DDTS # CSCdx42820

Following issuing a force ring on a four fiber MS-SPRing on one side and creating a signal failure ring by pulling both the working and protect receive fibers on the other side, restoral of the working fiber causes the cross connect card to reboot. The same symptom can occur if you issue two force ring commands on the east and west of a node and then fail the protect fiber on either side of that node. To avoid this situation, in the first case, release the force ring command before restoring the working fiber on the other side. In the second case, issuing two force ring commands on the side node to enable isolation of a node is not recommended. Rather, issue the force ring command from the nodes adjacent to the node to be isolated. This issue is resolved in Release 3.4.

DDTS # CSCdv26389

Symptom: The CLEAR command does not clear a WTR (Wait To Restore) condition on MS-SPRing (or BLSR). If the MS-SPRing (or BLSR) gets into a WTR state, there is no preemptive message that can be sent to the system to remove the WTR. This issue is resolved in Release 3.4.

DDTS # CSCdx20789

Clearing a force switch span (FS-S) on a four fiber MS-SPRing during a signal failure span (SF-S) can result in a condition where the FS-S still exists on the same span, and a traffic hit is incurred. To avoid this issue, do not clear a force switch span while a signal failure working is detected. This issue is resolved in Release 3.4.

DDTS # CSCdx02278

A four fiber node with signal failure or signal degrade on the local protect might raise a K-byte PT event when a PT span K-byte is received. This issue is resolved in Release 3.4.

DDTS # CSCdw81494, CSCdw81592, CSCdw82811, CSCdv09279, CSCdv83805

A Manual or Force switch is not released when SD or SF occurs. This occurs under the following conditions:

In a non-revertive linear 1+1 bidirectional link between two nodes, A and B,

Step 1 Issue a manual or force switch to protect span on Node A.

Step 2 Generate a SD or SF on the protection receiver side of Node A by pulling the receive fiber.

Step 3 The line does switch from the protection to the working while the SD or SF lasts.

Step 4 The manual or force switch is never released. The requested switch to protect hangs in the CTC GUI.

Step 5 After the SD or SF is released, the line switches back to protection.

This issue is resolved in Release 3.4.

DDTS # CSCdw62602

Node isolation caused by four unidirectional failures on the four spans of a node can result in traffic loss. To regain traffic, issue a “Force ring” on both sides of the isolated node and proceed with repairing the failures. This issue is resolved in Release 3.4.

SNCP Functionality

DDTS # CSCdw75434

If you perform a Force switch after performing a Manual switch on an SNCP link, you may receive an error message that “all circuits are failed to switch.” To correct this situation, clear the Manual switch on the protected link and then perform the Force switch on other link. This issue is resolved in Release 3.4.

SNMP

DDTS # CSCdx27495

When traps 3580 (peerCardNotResponding) and 3590 (alarmsAndEventsSuppressedForThisObject) are received from a Cisco ONS 15454 SDH node, the NMS does not display the associated cerent454ObjectName varbind (the rest of the varbinds are displayed correctly). This problem occurs only if the NMS uses the trap definitions from the “CERENT-*-MIB.mib” MIB file to build a display template to display all the Cisco 15454 SDH proprietary traps. To work around this issue, manually add additional varbind display information to the trap display template in the NMS for these traps. This issue is resolved in Release 3.4.

Documentation

SNCP Traffic Patterns

The *Cisco ONS 15454 SDH Installation and Operations Guide*, Release 3.3, Chapter 5, Section 5.2, “Creating SNCP Circuits” incorrectly states that:

> Working traffic flows in one direction and protection traffic flows in the opposite direction. If a problem occurs in the working traffic path, the receiving node switches to the path coming from the opposite direction.

With Release 3.3, all primary (working) SNCP circuits traverse the ring in the same direction on the fiber. This means that both circuits do not traverse the same fiber immediately after provisioning. One significant benefit of this feature is that a single fiber cut does not cause both directions of a two-way SNCP circuit to switch. Another benefit is that the ONS 15454 SDH SNCP circuit behaves more in line with current industry expectations.

This issue is resolved in Release 3.4.

Ethernet Card Names

In the user documentation, E1000-2 (product name 15454E-E1000-2) is called E1000-2-G; and E100T-12 (product name 15454E-E100T-12 is called E100-12-G. This issue is resolved in Release 3.4.

New Features and Functionality

This section highlights new features and functionality for Release 3.4.x. For detailed documentation of each of these features, consult the user documentation.

New Hardware

AIC-I

The Alarm Interface Controller – International (AIC-I) card expands the system management capabilities of the AIC card for the ONS 15454 SDH platform.

The AIC-I (Alarm Interface Controller – International) module is an optional card that expands systems management capabilities for customer-defined alarm I/O, user data, and orderwire functionality. This card resides in one of the common slots (Slot 9), but is not required for node operation unless the expanded AIC-I features are desired.

The AIC-I supports 8 light indicators as follows.

Table 5 AIC-I Light indicators

Light Indicator	States
FAIL light indicator	OFF when module is operating properly, RED when the module is sensed as failed, or when the card is coming up.
ACT light indicator	OFF when the module is not operational, GREEN when the module is active.
PWR light indicator (A or B feed)	GREEN when the power is within the normal range, AMBER when the power is below normal but still functional, RED when the power feed is out of range (either too low or too high).
INPUT/OUTPUT light indicators	INPUT light is OFF when there is no input alarm raised, YELLOW when an alarm exists. OUTPUT light is OFF when no external controls have been triggered, YELLOW when one has.
RING light indicators (one for LOW and one for EOW)	OFF when no call is present on the orderwire, GREEN Flashing when a call is received.

AIC-I Features

Table 6 lists the available ANSI support for AIC-I card features.

Table 6 AIC-I Feature Support

AIC-I Feature	Support in SDH
Input/Output alarm contact support	4 I/O
Input only contact support	16-Input
F1 64K User Data Channel	Yes

Table 6 AIC-I Feature Support

AIC-I Feature	Support in SDH
D4-D12 576K Data Communication Channel	Yes
A-Law support	Yes
Mu-Law support	Yes
Selective station calling	Yes
Interoperability with existing AIC	NA

Input and Input/Output Alarm Contacts Support

The AIC-I card provides additional input/output alarm contact closures for customer use: up to 16 user-defined input, and 4 user-defined input/output contacts. The 4 input/output contacts are provisionable either as all inputs or all outputs. They default to input contacts.The alarms are user-definable via CTC. An LED for inputs and another for outputs are included on the front panel to indicate the overall status of the alarm contacts. Input alarms are typically used for external sensors like open doors, temperature sensors, flood sensors and other environmental conditions. Output contacts are typically used to drive visual or audible devices like bells, lights, pagers, or even to control generators, heaters, fans, etc.

All of the contacts (input and input/output) can be programmed separately via the AIC-I card view Provisioning tab.

F1 User Data Channel (F-UDC)

The user data channel allows the local user to physically connect a dedicated data channel of 64 Kbps (F1 byte in SOH) between two nodes in a 15454 SDH network. Each AIC-I supports two F-UDCs, and each UDC can be routed to a unique and separate optical interface on the ONS 15454 SDH system. The F1 UDC is a 64 Kbps point-to-point channel and is routed between optical interfaces via the TCC/TCC+ module in intermediate locations. Provisioning of a UDC function is accomplished via the Overhead Circuits tab within the Provisioning tab in the Network view.

Data Communication Channel (MS-UDC)

The DCC utilizes the DCC-M and allows you to physically connect a dedicated data channel of 576 Kbps (D4-D12 bytes in MSOH) between two nodes in an ONS 15454 SDH network. Each AIC-I supports two MS-UDCs, and each UDC can be routed to an individual and separate optical interface on the ONS 15454 SDH. The MS-UDC is a 576 Kbps channel and is routed between optical interfaces via the TCC-I modules in intermediate locations. Provisioning of an MS-UDC function is accomplished via the Overhead Circuits tab within the Provisioning tab in the Network view.

Orderwire Functionality

Orderwire provides a craftsperson the ability to plug a standard phone set into an ONS 15454 SDH and communicate with one or more other craftspeople working at other ONS 15454 SDHs. The orderwire is a PCM-encoded voice channel that rides on bytes E1 or E2 in the section and line overhead. The orderwire interface on the AIC-I supports both 4-wire and 2-wire connection via an RJ-11 jack.

The AIC-I allows simultaneous use of both local (Multiplex Section overhead signal) and express (Regenerator Section overhead channel) orderwire channels on a SDH ring or particular optics facility. If the AIC-I card is not equipped at regenerator stations, the TCC will pass-through the E1/E2 byte.

Provisioning of an orderwire function is accomplished via the Overhead Circuits tab within the Provisioning tab in the Network view. Both of the EOW and LOW channels can also be adjusted via the AIC-I card view.

Note The OC3/STM-1 card does not support the express orderwire channel.

STM4-4 SH 1310nm Card

The STM4-4 SH card provides the same functionality as the legacy STM4 SH 1310 card, but with four times the port density. This card increases optical sensitivity compared to the legacy STM4 SH 1310 card. The STM4-4 uses angled SC connectors. SDCC terminations and DCC tunnels can be provisioned on all four ports.

The STM4-4 card requires the XC10G card. STM4-4 cards are only supported in the multi-speed slots (1-4, 14-17) of the ETSI shelf as indicated by the “star” code on the lower extractor. If the card is installed in a high-speed slot, its graphical representation will not appear in CTC and a Mismatch of Equipment and Attributes (MEA) alarm will be raised.

STM4-4 Card and Span Upgrades

A legacy STM4 card can be upgraded to an STM4-4 without first removing DCC, timing, ring, protection, and/or circuit provisioning by right-clicking on the card image from the CTC node view, selecting the Change Card option from the STM4 Card popup menu, and selecting STM4_4 from the Change To pull down menu. The STM4 port will be mapped to Port 1 on the STM4-4.

An STM4 to STM4-4 span upgrade can be performed without disrupting traffic by right-clicking on the STM4 span from the CTC network view and selecting STM4_4 from the Upgrade To pull-down menu. The STM4 port will be mapped to Port 1 on the STM4-4.

Card and Span Protection

CTC will not allow provisioning of a working and protect path or channel on the same STM4-4 card. This has implications for MS-SPRING, SNCP and 1+1 protection schemes as follows.

- Two fiber MS-SPRING is supported on the STM4-4 but the east and west spans cannot be provisioned on the same STM4-4 card. As is the case with legacy STM4 cards, four fiber MS-SPRING is not supported.
- Working and protect SNCP paths cannot use the same STM4-4 card; neither automatic nor manual circuit creation will allow circuits to be created this way.
- A 1+1 protection group must use the same port number on both the working and protect cards. For example, Port 1 on the protect card must protect Port 1 on the working card.

New Software Features and Functionality

Network Time Protocol

The Network Time Protocol (NTP) feature enhances the SNTP (Simple Network Time Protocol) functionality of the ONS 15454 SDH for Release 3.4.x. Now NTP servers are supported. Previously, the ONS 15454 SDH supported only an SNTP server.

Note The ONS 15454 SDH does not act as a time server. Rather, it acts as a client, obtaining time from the provisioned server.

Spanning Tree Control

Release 3.4 adds the ability to Turn Spanning Tree off for Ethernet circuits. You can disable or enable spanning tree on a circuit-by-circuit basis on unstitched Ethernet cards in a point-to-point configuration. This feature allows you to mix spanning tree protected circuits with unprotected circuits on the same card, to reuse VLANS, and to set up two single-card E-series Ethernet cards on the same node to form an intranode circuit.

AIS Off Mode for J1 Path Trace

In the previous release, when a J1 trace mismatch is detected, the ONS 15454 SDH inserts AIS path in the timeslot. The AIS off mode feature will allow customers to choose between sending or not sending AIS-P downstream from the TIM-P defect. When AIS-P is inserted downstream the likelihood of a traffic outage increases. The AIS off mode feature does not require any special configurations other than J1 compatible modules. See the *Cisco ONS 15454 SDH Reference Guide* for details on J1 compatibility.

Microsoft Windows XP

Release 3.4.x supports the Microsoft Windows XP operating system.

MS-SPRing Wizard

Release 3.4 introduces the MS-SPRing wizard, which allows you to create, edit, and delete a MS-SPRing from the network view of CTC. The MS-SPRing wizard reduces common errors in creating rings from distant nodes. The wizard also facilitates creating and deleting rings over a much shorter period than it took in previous releases to individually turn up MS-SPRing attributes on a node-by-node basis. For specific functions and limitations of the MS-SPRing wizard, consult the user documentation for Release 3.4.

Queued & Preemptive Switching

In releases prior to 3.4, the node accepted and stored switch commands regardless of higher priority requests. Once the higher priority request was cleared, the lower priority command was applied. Also, although the software preempted lower priority user commands for higher priority requests, it reapplied the lower priority command once the higher priority request had been cleared or completed.

Release 3.4 complies with Telcordia GR-253 Issue 2, allowing a higher priority, local or remote request to preempt (override) an external, lower priority request. The preempted request is not retained in memory or in a queue for completion (in other words, when the higher priority request is cleared, the preempted switch request will not be reinitiated). Thus, when you attempt to apply a switch command under these circumstances, the request will be denied.

In Release 3.4, you will be notified immediately if a condition occurs in which a command is overridden. The software will deny a switch request immediately if a higher priority request already exists.

Note This behavior applies to all protection types: 1:1, 1:N, 1+1, MS-SPRing and SNCP. (Note that in Release 3.3, 1:1 and 1:N are fully compliant).

For details on possible switch commands and their associated priority levels, as well as other actions can affect the Automatic Selector Criteria switch state, consult the user documentation for Release 3.4.

Multiple OSPF Areas

In Release 3.3, only one OSPF area was supported within a data communication channel (DCC-R) network. With Release 3.4, you can configure multiple areas on different DCC-R links for the same node. This type of configuration limits the amount and size of flooded Link State Advertisement (LSA) updates to individual areas that occur each time there is a topology change or scheduled update. This gives you the ability to better control the amount of traffic over each DCC-R link.

RIP Support

In Release 3.3, only static routes and OSPF routing protocols were supported for a TCC-I LAN. Many deployed networks today use RIP to exchange IP routing information. Release 3.4 provides RIP as a routing protocol option, giving the network designer increased flexibility and more choices for network design. In a small network, RIP has the advantage of very little overhead in terms of bandwidth used and configuration and management time. RIP is also easy to configure and implement.

RIP is a distance vector routing protocol, in which the router only exchanges routing information between connected neighbors. RIP Version 1 advertises routes by sending updates to the broadcast address 255.255.255.255. All devices on the LAN receive and process broadcasts. RIP Version 1 is a classful routing protocol. Classful routing always summarizes routes by the major network numbers and always considers the network class. This is always done at network boundaries. Subnets are not advertised to other major networks. Non-contiguous subnets are not visible to each other.

RIP Version 2 advertises routes by sending updates to IP multicast address 224.0.0.9. To reduce unnecessary load on those hosts that are not listening to RIP-2 messages, the IP multicast address is used for periodic broadcasts. RIP Version 2 is a classless routing protocol. Classless routing differs from classful routing in that the prefix length is transmitted. The prefix length is evaluated at each point it is encountered throughout the network. Thus, the prefix length can be changed to advertise routes differently at different locations within a network. Classless routing enables more efficient use of IP address space and reduces routing traffic.

The RIP-2 feature can be enabled on a LAN management interface through CTC to advertise to a router on the network. You can choose between OSPF and RIP, with “None” as the default.

Up to 25 occurrences each of the address-family identifier (AFI), address, and metric fields are permitted in a single IP RIP packet. That is, up to 25 routing table entries can be listed in a single RIP packet. If the AFI specifies an authenticated message, only 24 routing table entries can be specified.

RIP uses hop count to rate the value of each different route. The hop count is the number of routers that can be traversed in a route. A directly connected network has a metric of zero; an unreachable network has a metric of 16.

To avoid a potential routing loop when distributing routes between RIP and OSPF, a node only advertises routes it knows through RIP. Any RIP updates the node receives from the routers on the LAN are discarded. For any network behind the directly connected router, static routes must be provisioned on the node.

Static Route Enhancement

In releases prior to 3.4, nodes discovered each other on the network using OSPF through IP, over PPP, over SDCC. In Release 3.4, static route enhancement allows ONS 15454 SDHs ONS 15454s and ONS 15327s to communicate with 3rd party equipment using IP over PPP over SDCC. You can disable OSPF on the CTC SDCC Termination screen. You can then create route entries in the Static Route tab to access IP-enabled 3rd party equipment.

Note Static Route Enhancement will not allow visibility to 3rd party equipment using IP over PPP over SDCC when the proxy/firewall feature is enabled.

UCP

Unified Control Plane (UCP) is a modular software feature set that provides functions that increase networking control capabilities in the areas of routing, signaling, and provisioning. UCP also provides resource and service discovery. In past releases, provisioning was a manual process requiring users of management systems to set up multiple segments in an end-to-end circuit configuration. UCP was developed to automate end-to-end optical network provisioning, with the addition of mesh restoration capability and drive intelligence within nodes.

UCP O-UNI (Optical User Network Interface) is a well-defined set of protocols used for signaling and routing between a service provider network and equipment in a client network based on an emerging standard. The UNI 1.0 specification describes the set of services, interfaces, and signaling capabilities.

In the ONS 15454 SDH O-UNI 1.0 implementation, services are used to invoke:

- Connection creation
- Connection deletion
- Connection status inquiry
- Autonomous connection status change notification update

Note Connection modification is not supported in this release.

Clients request services of the optical network using O-UNI signaling. UNI signaling agents (that is, client side agents), referred to as O-UNI-C, make requests of network side agents (O-UNI-N) using O-UNI signaling messages. These messages are transported over an IP control channel “in band” (IB) or “out of band” (OB) using RSVP-TE.

O-UNI-C provides signaling termination functionality. O-UNI-N provides signaling pass-through and inter-working functionality, circuit routing, and reachability information to O-UNI-C. O-UNI-C is implemented in Release 3.4. O-UNI-N will be implemented in a future release.

Protection Channel Access

In compliance with Telcordia GR-1230-CORE, section 3.4, protection channel access (PCA) is supported with the ONS 15454 SDH Release 3.4. Normal MS-SPRing utilizes only 50% of ring bandwidth, while the other 50% remains idle until a ring or span switch occurs. PCA circuits can run through the idle bandwidth. These circuits will be preempted when a switch occurs, making room for protected circuits. PCA circuits are provisionable on the protection channels of 2, or 4 fiber MS-SPRing configurations. PCA circuits are restored after the wait-to-restore timer times out following a switch. This feature allows service providers to utilize their networks more efficiently.

Circuit routing preferences for Release 3.4 are enhanced to support PCA circuit creation. Circuit routing preferences can be viewed as being divided into primary and secondary preferences. The primary routing preference (PRP) determines if the entire path is protected (fully protected path) or unprotected. The PRP is set when the circuit is created via CTC and cannot be changed. The Secondary Routing Preference (SRP) depends on the PRP and is also set when the circuit is created via CTC.

Table 7 Routing Preferences

Primary Routing Preference (PRP)	Secondary Routing Preference (SRP)
Protected (fully protected path)	Node Diverse Required, Node Diverse Desired, Link Diverse
Unprotected	PCA enabled/disabled

For details on provisioning PCA, consult the user documentation for Release 3.4.

Enhanced State Model

The Release 3.4 enhanced state model adds increased control of the service state for ports and circuits. This state model provides increased options for entities (ports, equipment, or circuits) out of service, awaiting automatic activation, or out of service and under maintenance. The new state model provides the ability to provision an entity as service-ready while awaiting the arrival of an additional required item (traffic or physical card) before going into service.

In addition to the established states, IS (In Service) and OOS (Out of Service), the enhanced state model adds the Out of Service-Auto In Service (OOS-AINS) and Out of Service-Maintenance (OOS-MT) states.

Note Loopbacks are only allowed when the entity is in the OOS_MT, or OOS_AINS state.

Maintenance Mode

The OOS-MT mode is the same as the IS mode except that alarms are not reported autonomously, yet they can still be retrieved. Maintenance is allowed while an entity is in this state. This OOS-MT state applies to the port level and the circuit level.

Auto In-Service Provisioning

The enhanced state allows any entity to be in an Auto In-Service (OOS-AINS) state. This state allows you the ability to provision an entity (port, equipment, or circuit) to be ready to be placed in service, but to await the arrival of the required item (traffic or physical card) before actually going into service. This allows pre-provisioning of circuits and cards, which then automatically activate upon the detection of the appropriate signal or hardware (for example, when a card is inserted). The OOS-AINS state saves carriers from the need to filter alarms due to the pre-provisioning of circuits before the signal is received from their customers. In the case of cards, the feature permits accurate reflection of the expected status of the card while the card itself has yet to be inserted. When an entity is in the OOS_AINS state, alarms associated with the entity are reported in the conditions pane, rather than the alarms pane.

For details on the uses and behaviors of this state, consult the user documentation for Release 3.4.

Node Defaults

In Release 3.4.x, you can override the system default values for the node and card level that exist on the ONS 15454, ONS 15454 SDH, or ONS 15327. This function is provided at the node provisioning pane level and will change the value which the node will use for the parameter setting. Many default provisioning values are now configurable. For example, you can decide whether ports on a certain type of card should default to OOS, OOS-AINS, OOS-MT, or IS when the card is pre-provisioned or inserted.

In Release 3.4.x CTC there is a Defaults Editor tab accessible from the Node View > Provisioning tabs. Default values can be changed, exported, imported, and applied. Default values can also be reset to revert the defaults from the most previous "Apply" to the node. The export file is an ASCII text file, similar to the ".ctcrc" file. CTC can save and load the default overrides to or from a file.

Application of new, card level and lower defaults does not affect items already provisioned or pre-provisioned. These defaults only apply to entities created after them.

Application of new node level defaults is an alternate way of provisioning those values. This method is made available because there is no way to apply the new values when the node is created later, since applying the values to the node requires that the node already exist. The exceptions to the node level defaults are the node.protection and node.circuits defaults, which are used only when 1+1 or MS-SPRing protection is provisioned, or when a SNCP circuit is provisioned. Previously provisioned 1+1 or MS-SPRing will not be affect by these changes to defaults, nor will any previously provisioned SNCP circuits.

Filtering of Circuit Table

Release 3.4 adds options in the Circuit window to filter to a specific port on a card. These are in addition to the options to filter by network, node, or card level. These options will restrict the circuits listed to only those items allowed by the filter and associated with the current view.

Overhead Circuits Provisioning

Release 3.4 introduces A-Z provisioning of overhead circuits. Consult the user documentation for further details on this enhancement.

SNMP Enhancements

The SNMP Agent has been modified in Release 3.4 to accommodate the new enhanced state model changes for the ONS 15454 SDH, ONS 15454, and ONS 15327. The SNMP MIBS have been modified to accommodate the various state changes.

SNMP Agent modifications for the enhanced state model only affect one MIB variable, ifAdminStatus, which is part of the ifEntry table.

The new enhanced states and the corresponding return values for the ifAdminStatus states are outlined in Table 8.

Table 8 IfAdminStatus

Enhanced State Model	IfAdminStatus return value
IS	up(1)
OOS	down(2)

Table 8 IfAdminStatus

Enhanced State Model	IfAdminStatus return value
OOS_MT	testing(3)
OOS_AINS	down(2)

Note These states are also displayed in CTC when provisioning a port in or out of service.

CTC Security Enhancements

Prevention of Identical User ID and Password

As of Release 3.4, CTC prevents the creation of a userid and password that are identical. The userid and password are identical if they contain the same characters in the same numbers and sequence, irrespective of case. For example, "betsy" and "BeTSy" are considered to be the same, while "betsy", "ysteb" and "betssy" are all different.

As of Release 3.4, CTC prevents the creation of a password containing as a subset of characters the associated userid. The password contains the associated userid if it contains anywhere within it the same characters in the same numbers and sequence that make up the entire userid, irrespective of case. For example, password "SBeTSyXC" and userid "betsy" are disallowed, while password "betsy" and userid "SBeTSyXC", or password "bet3sy" and userid "BeTSy" are allowed.

Enforce Password Complexity

CTC will not allow creation of new passwords that do not comply with Telcordia GR-815, which states that passwords must be at least six characters long, contain at least one alphabetic, one numeric and one special character (+, # or %).

Note CTC warns on entry only when a pre-existing non-GR815 compliant password is used, permitting the user to continue with the older password.

Password Toggling Prevention

As of Release 3.4 CTC prevents users from changing a password to the current password value. For example, if the existing password is "*@nite", the new password cannot also be "*@nite".

Login Warning

In Release 3.4.x the CTC login screen warning has been changed. In Release 3.4.x the CTC Login window is provisionable up to 250 characters. Only a system administrator/superuser can edit the login warning window. This is done via the node view, in the Provisioning > Security > Legal Disclaimer > HTML tabs. The system administrator or superuser can edit the log-in warning message as required preview it by clicking on the preview sub-tab. All users logging into the node will see the new, modified warning message thereafter.

Related Documentation

Release-Specific Documents

- *Release Notes for Cisco ONS 15454 SDH Release 3.4*
- *Release Notes for Cisco ONS 15454 Release 3.4.1*
- *Release Notes for Cisco ONS 15327 Release 3.4.1*

Platform-Specific Documents

- *Cisco ONS 15454 SDH Installation and Operations Guide, Release 3.4*
- *Cisco ONS 15454 SDH Troubleshooting and Maintenance Guide, Release 3.4*
- *Cisco ONS 15454 SDH Product Overview, Release 3.4*

Obtaining Documentation

The following sections provide sources for obtaining documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following sites:

- http://www.cisco.com
- http://www-china.cisco.com
- http://www-europe.cisco.com

Documentation CD-ROM

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or as an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco Product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, in North America, by calling 800 553-NETS(6387).

Documentation Feedback

If you are reading Cisco product documentation on the World Wide Web, you can submit technical comments electronically. Click **Feedback** in the toolbar and select **Documentation**. After you complete the form, click **Submit** to send it to Cisco.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Attn Document Resource Connection
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools. For Cisco.com registered users, additional troubleshooting tools are available from the TAC website.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information and resources at anytime, from anywhere in the world. This highly integrated Internet application is a powerful, easy-to-use tool for doing business with Cisco.

Cisco.com provides a broad range of features and services to help customers and partners streamline business processes and improve productivity. Through Cisco.com, you can find information about Cisco and our networking solutions, services, and programs. In addition, you can resolve technical issues with online technical support, download and test software packages, and order Cisco learning materials and merchandise. Valuable online skill assessment, training, and certification programs are also available.

Customers and partners can self-register on Cisco.com to obtain additional personalized information and services. Registered users can order products, check on the status of an order, access technical support, and view benefits specific to their relationships with Cisco.

To access Cisco.com, go to the following website:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC website is available to all customers who need technical assistance with a Cisco product or technology that is under warranty or covered by a maintenance contract.

Contacting TAC by Using the Cisco TAC Website

If you have a priority level 3 (P3) or priority level 4 (P4) problem, contact TAC by going to the TAC website:

http://www.cisco.com/tac

P3 and P4 level problems are defined as follows:

- P3—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- P4—You need information or assistance on Cisco product capabilities, product installation, or basic product configuration.

In each of the above cases, use the Cisco TAC website to quickly find answers to your questions.

To register for Cisco.com, go to the following website:

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

If you cannot resolve your technical issue by using the TAC online resources, Cisco.com registered users can open a case online by using the TAC Case Open tool at the following website:

http://www.cisco.com/tac/caseopen

Contacting TAC by Telephone

If you have a priority level 1 (P1) or priority level 2 (P2) problem, contact TAC by telephone and immediately open a case. To obtain a directory of toll-free numbers for your country, go to the following website:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

P1 and P2 level problems are defined as follows:

- P1—Your production network is down, causing a critical impact to business operations if service is not restored quickly. No workaround is available.
- P2—Your production network is severely degraded, affecting significant aspects of your business operations. No workaround is available.

This document is to be used in conjunction with the documents listed in the "Related Documentation" section.