

Release Notes for Cisco MGX Route Processor Module (RPM-XF) IOS Release 12.2(15)T5 for MGX 8850 and MGX 8950 Release 4.0.10 (PXM45)

Contents

- About These Release Notes..... 2
 - Special Notes 2
- New Features 2
 - MGX 8950 Support for RPM-XF 3
 - EiBGP Load Balancing 3
 - IP Accounting Counters 4
 - Multiple Actions under Police Command 4
 - QoS Sub Optimal Link Utilization 5
- RPM-XF Redundancy Support..... 5
- Features Not Supported in This Release..... 6
 - Network Management Features 6
- SNMP MIB..... 6
- RPM-XF Limitations and Restrictions 6
- Notes and Cautions 8
 - RPM-XF auto_config File Management..... 9
 - Card Management 10
 - RPM-XF Bootflash Precautions 10
- Resolved Caveats—Release 12.2(15)T5..... 11
- Open Caveats—Release 12.2(15)T5 12
- Compatibility Notes 20
 - RPM-XF Boot File and Firmware File Names and Sizes 20



RPM-XF Compatibility Matrix	20
MGX RPM-XF Hardware.....	20
Cisco IOS Release Compatibility Information	21
Using XModem to Download Flash to RPM-XF Cards	21
Historical Information for 12.2.x IOS Baseline.....	23
Resolved Caveats in Release 12.2.15T	23
Resolved Caveats Prior to Release 12.2.15T	26
Related Documentation.....	40
Obtaining Documentation	40
Cisco.com	40
Documentation Feedback.....	41
Obtaining Technical Assistance	41
Cisco TAC Website	41
Opening a TAC Case	42
TAC Case Priority Definitions.....	42
Obtaining Additional Publications and Information.....	42

About These Release Notes

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM, a member of the Cisco Connection Family, is updated monthly. Therefore, it might be more current than printed documentation. To order additional copies of the Documentation CD-ROM, contact your local sales representative or call customer service. The CD-ROM package is available as a single package or as an annual subscription.

Note that for Release 4.0.10 , the user documentation (command reference, overview, and installation and configuration guides) use the MGX Release 4 and Cisco IOS documents in addition to this release note.

Product documentation for MGX 8850 is available at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/wanbu/8850px45/re14/index.htm>

<http://www.cisco.com/univercd/cc/td/doc/product/wanbu/mgx8250/12/rpmpr/reNotes/index.htm>

Special Notes

Please refer to the this Web page for the latest updates on software and firmware issues.

http://www.cisco.com/cgi-bin/Software/Iosplanner/Planner-tool/prints.pl?get_crypto=&data_from=&hardware_name=&software_name=&release_name=12.2.15T&majorRel=12.2&state=:RL&type=Early%20Deployment&file=12.2.15T.c.html

New Features

The MGX RPM-XF is a next-generation, high performance model of the RPM for the MGX 8850 and MGX 8950 platforms, using PXM45 processor modules. It is a router module based on an RM7000A MIPS processing engine that fits into slots 1-6 and slots 9-16 in the MGX 8850 and slots 1-6 and slots 11-16 in the MGX 8950.

The RPM-XF hardware provides forwarding technology for packet switching capabilities in excess of 2-million pps. The forwarding engine is packet based and is interfaced to the midplane of the system through a combination of switch interface technologies. For more information on the RPM-XF, refer to the *Cisco MGX Route Processor Module (RPM-XF) Installation and Configuration Guide, Release 4*.

MGX 8950 Support for RPM-XF

In Cisco MGX Release 3.0.10 and later, the MGX 8950 supports the MGX RPM-XF card set. The MGX RPM-XF card set can occupy any of the available service module slots in the MGX 8950, which are slots 1-6 and slots 11-16.

EiBGP Load Balancing

To balance load by BGP on multiple paths to destination, traffic is directed on multiple available paths between autonomous systems (AS) by gateway routers.

The following CLI are used to implement this feature.

Command	Description
maximum-path <nums>	Configure maximum number of EiBGP parallel routes. For example: <pre>bgpbox-zenith-CE1(config)#router bgp 4 bgpbox-zenith-CE1(config-rout)#maximum-paths 3 bgpbox-zenith-CE1(config-rout)#end</pre>
show ip bgp	This command has been enhanced to show the multipaths. Each multipath is marked as 'multipath'. The bestpath is marked as 'multipath' and 'bestpath'. The output also has what flavour of multipath is enabled. For example: <pre>bgpbox-zenith-CE1#sh ip bgp 141.22.0.0 BGP routing table entry for 141.22.0.0/16, version 18 Paths: (2 available, best #1) Multipath: eBGP Advertised to non peer-group peers: 7.0.76.9 100 5 7.0.76.2 from 7.0.76.2 (100.0.0.2) Origin IGP, localpref 100, valid, external, multipath, best 100 5 7.0.76.9 from 7.0.76.9 (100.0.0.9) Origin IGP, localpref 100, valid, external, multipath</pre>

Limitations:

- If there are multiple alternate paths for the peering point, only one of the paths is utilized for a given prefix.
- Only per-flow load balancing is supported. Per-packet load balancing is not supported.
- Supported only in MPLS/VPN networks.
- Load balancing proportional to the link's bandwidth [dmz-linkbw] is not supported. The load balancing will be performed on the available links with equal costs.
- The maximum number of paths that can be used for load sharing is 6. This is the current IOS limitation too.
- Load balance will not work if RDs are same with RR. RDs have to be different if RRs are used.
- CEs in different VPNs using same RDs does not work.

IP Accounting Counters

MGX RPM-XF stores the packet/byte counters based on precedence/dscp values on a per interface level at input ONLY.

The following CLI have been added or enhanced for this release to implement this feature:

Command	Description
ip accounting ?	pop20-slot6 (config-if) # ip accounting ? <pre> precedence Count packets by IP precedence on this interface dscp Count packets by dscp on this interface </pre>
ip accounting precedence ?	pop20-slot6 (config-if) # ip accounting precedence ? <pre> input received packets and bytes </pre>
ip accounting dscp ?	pop20-slot6 (config-if) # ip accounting dscp ? <pre> input received packets and bytes </pre>
show int [interface] precedence	pop20-slot5# show int [interface] precedence
show int [interface] dscp	pop20-slot5# show int [interface] dscp
clear counters	pop20-slot5# clear counters

Limitations:

- Counters are maintained ONLY at input per interface.
- There is no count of dropped/transmitted based onDSCP/PREC packets per interface.

Multiple Actions under Police Command

The MGX RPM-XF **police** command is now similar to the IOS/RPM command. Therefore, multiple **exceed** and **conform** actions can be applied on the **police** command.

The **police** CLI command has been enhanced as follows:

Prior to this release, the **police** command had no menus and all parameters were listed on one line, as shown in the following example.

```
domino80p01-z001#sh policy test1z
```

```
.....
police 128000 8000 8000 conform-action transmit exceed-action drop
```

This allows only one value for *conform-action* and *exceed-action*.

The newly enhanced **police** command functions as shown in the following example.

```
ipftrt90r14-01(config-pmap-c)#police 128000 8000 8000
ipftrt9(config-pmap-c-police)#conform-action transmit
ipftrt9(config-pmap-c-police)#exceed-action set-dscp 28
ipftrt9(config-pmap-c-police)#exceed-action set-mpls 2
```

Note that multiple **conform-action** and **exceed-action** parameters can be configured.

QoS Sub Optimal Link Utilization

RPM-XF uses VTMS as a scheduling algorithm. VTMS schedules queues based on the current link utilization in real time. The previous version of the VTMS algorithm was efficient and mapped well in an ASIC or network processor. However, it did not fully utilize the link.

In this release, a fix is made (via a new CLI command) to allow the user to specify the over subscription factor on a queue. The factor is in the range on 1-31 and can be denoted as 2n. An over subscription factor of n = 2 on any queue, essentially means to subscribe that queue by a factor of 4 (2n where n (in this example) is 2; so 2 raised to power 2=4).

The new syntax for bandwidth and priority commands are:

```
[no] bandwidth {<kbps> | percent <percentage> | remaining percent <percentage>}
[maximize-utilization [<max-shift>]]
```

```
[no] priority {<kbps> | percent <percentage>} [maximize-utilization
```

RPM-XF Redundancy Support

RPM-XF 1:N redundancy is used to switch configuration and traffic from one RPM-XF module to another RPM-XF module. Route processing continues with minimal traffic loss even if an RPM-XF fails and there is no operator or direct access to swap the failed card or fix the problem. Currently we support RPM-XF warm redundancy, which ensures Layer 2 state restoration. Layer 3 state is restored via convergence.



Note

When resetting a shelf with RPM-XFs configured for 1:N redundancy, it is recommended that you bring up the primary slots in active state.

The main benefits of 1:N Redundancy are:

- An RPM-XF card with hardware problems can be fixed while the redundant standby card takes over its functionality.
- Software upgrades are easier and can be done with less downtime.
- LAN interface redundancy supported with MAC addresses of primary RPM-XF copied to standby RPM-XF.
- 1:N Redundancy support for Gigabit Ethernet interface backcards during front card switchover.
- Y cable redundancy support for POS backcards during front card switchover. (With Y cable, 1:N redundancy is restricted to N = 1).

The following are the general guidelines for redundancy on the RPM--XF:

- **Adred** is not allowed between RPM-PR and RPM-XF.
- To configure redundancy, the Primary RPM-XF should be in Active state and Secondary RPM-XF card must be in Active/Standby state.
- Removal of the Active RPM-XF back card does not cause switchover to the standby RPM-XF.
- User has to make sure that **E:RPM/auto_config_slot#** is created before adding redundancy. This may require a login to primary card through the command line and manually adding **boot config e:auto_config_slot#** followed by a **write mem**.
- Executing **switchcc** back-to-back with **switchredcd** can cause problems. We recommend giving a gap of at least 5 seconds between **switchredcd** and a **switchcc**.
- IOS software on a standby card should be the same or higher version than the Active RPM-XF card.
- Booting the card from an image on tftp server is not recommended when the card is in redundancy group. The card should be booted from image in bootflash or PXM disk only.
- Configuring the standby RPM-XF is not recommended.

Features Not Supported in This Release

The following features are not supported in this release.

- LSC Redundancy
- MPLS TE tunnels on ATM Interfaces
- VC Merge
- RPM-PR to RPM-XF upgrade
- OIR of backcards without interfaces in shutdown mode
- Per packet load balancing
- Modem connectivity on Auxiliary port
- ROMMON's Xmodem functionality does not support the Speed option.

Network Management Features

Network management features are detailed in the Release Notes for *Cisco WAN Manager 12.0.00* at:

<http://cisco.com/univercd/cc/td/doc/product/wanbu/svplus/12/rnotes/index.htm>

SNMP MIB

SNMP MGX Release 4.0.10 MIB are provided with the delivery of this release. The MIB is in standard ASN.1 format and is located in the same directory within the release bundle on CCO. These files may be compiled with most standards-based MIB compilers. The tar file for MIB contains the file that contains the MIB release notes. This contains only MGX MIBS.

Cisco IOS MIBS are not part of this bundle. They are part of 12.2(15)T5 CCO release.

RPM-XF Limitations and Restrictions

The RPM-XF limitations and restrictions that apply to this release are as follows:

- **E: RPM/auto_config_slot#** must be created before adding redundancy. This may require a login through the CLI and manually adding the **boot config** command followed by a **write mem**.
- PVPs can not operate at a rate greater than 599039 kbps.
- High speed VC (SCR greater than or equal to 599,039 kbps) would not get full-configured rate for single flow (unique source and destination IP address). This happens because for high speed VCs, the PXF creates two queues and these queues can't be shared for same stream. Sharing two queues for same stream would cause out of sequence packets.
- PXF queue selection algorithm may cause traffic drop for multiple stream going to same destination via multiple paths. When the PXF gets a packet, it selects the output queue based on source and destination IP address. These addresses hash into one of the queues for the selected destination. So if there are multiple paths for the same destination, there is a possibility that multiple streams would hash to one queue, causing some queues to overflow, while others might be under-utilized.
- PXF buffer depletion may occur if packets of the same size (especially packets greater than 640 bytes) are sent to a congested interface.
- Currently VBR-nrt and VBR-rt are treated with same priority system wide.
- RPM-XF PVP only supports UBR.
- PVP in RPM-XF is not OAM managed.
- If out-of-sync SPVC or SPVP exist on RPM-XF, shrinking of PNNI partition would not be permitted.
- A single RPM-XF can only function as either an Edge LSR or as an LSC, but not as both.
- Because RPM-XF only supports UBR, VBR-rt and VBR-nrt, on the PXM, **dsppnportsrc** for RPM-XF port will show 0 available resource for CBR, ABR and signaling service types. Also, **cnfnpportcac** for CBR and ABR will be rejected.
- If RPM-XF is configured as an eLSR, RPM-XF does not support incoming VC-merge LVCs. There is a problem logged against LSC module that it cannot support both VC-merge/non-VC-merge supporting VSI slaves at the same time. So for now, if RPM-XF eLSR is part of a cell based MPLS network (with RPM-PRs or AXSMs in the same node), disable the VC-merge feature on LSC. (Note that VC-merge is enabled on LSC by default).
- RPM-XF eLSR only supports at most two MPLS sub-interfaces. Attempting to configure over the limit will result in an error message.
- Although RPM-XF VSI slave supports connections statistics Get command, only packets and bytes counts are available. Therefore, **show xtag cross-connect traffic int xtagatm** connection statistic display on LSC are actually packet counts from RPM-XF eLSR.
- OIR of MGX-1GE and MGX-1OC12POS-IR back cards are supported only with interfaces in *shutdown* state.
- MGX-1GE back card does not have the capability to provide line loopback.
- Flow Control Option is not configurable with MGX-1GE back card.
- MGX-1GE back card does not support SFP security.
- Line loopback and internal loopback cannot be set at the same time for the MGX-1OC12POS-IR back card with AMCC Mux.
- **pos ais-shut** command is not supported on MGX-1OC12POS-IR back card.

- Traffic rate per flow is at half the interface speed for POS/GigE interfaces in this new release.
- Autonegotiation is not recommended to be toggled on GigE interfaces when traffic is going through. This may result in permanent disruption of traffic.
- For UBR 2 queues always exist which would result in half the Flow rate for each flow because of hashing algorithm hashing into single queue .

The performance limits supported in this release are the following:

- 2K ATM SPVC Connection endpoints
- 2K IDBs
- 4K LVCs
- 100 VPCs
- 256 Policymap
- 100 OSPF neighbors
- 6 IOS-based cards in MGX shelf
- 500 VRFs: 500
- 500 BGP CE Peers
- 100 RIP CE sessions
- 500 Static CEs
- 100,000 VPN Routes per PE
- 250K non-VPN Routes per RPM-XF
- 50 Xtag interfaces per RPM-XF
- 300 OAM enabled connections

For more RPM-XF performance details, contact your sales representative.

Notes and Cautions

The following notes and cautions should be reviewed before using this release.

- Attempting to initiate RPM-XF switchover when **write mem** is in progress on the active RPM-XF card may lead to the card coming up with a partial configuration. When an **addred** is executed, an automatic **write mem** is triggered on the primary RPM-XF. If the primary card fails when the **write mem** is in progress, the card may come up with a partial configuration. The duration of **write mem** depends on the configuration size and can take up to 4 minutes to complete.
- There is a new stable “Boot-Hold” state displayed on the PXM45 when **dspecds** is executed. This state indicates that the RPM-XF is running only boot image. This state is reached when config register is set to 0x1 or when the bootldr cannot find the run-time image, but found the boot image. Enter **cc** to access the RPM-XF from the PXM45.
- Valid boot image need not be the first file in the boot flash. The RPM-XF will load from any valid boot image from the **bootflash:**. The run-time image can be the first file in the boot flash and RPM-XF will come up with that image.
- Trying to change PCR value of VP tunnel or changing MTU of switch interface with more than 4K VCs may cause CPU hog.

- If there is a large number of VCs (PVCs or LVCs or both) on RPM-XF card, executing disruptive operations on the main switch interface (int switch1) may cause flapping of protocols that run on these VCs. Examples of disruptive operations are **clear int switch1** and modification of PVP parameters. These operations cause deactivation and re-activation of all VCs under the main switch interface. Depending on the number of VCs, the time required to complete such operations may exceed certain protocol timeout limit. Examples of protocols that may be affected are OSPF and TDP/LDP.
- RPM-XF VSI slave tends to output informational warning/trace back messages caused by misconfigurations and CAC failures (onto console/IOS log file). These messages are mostly for information/debugging purpose. When these messages are observed, confirm that connection status is still intact and traffic is still passing successfully.
- Due to PXF scr granularity, the configured scr on IOS *pvc* CLI may not be the same as the actual scr programmed in the PXF. PXF bandwidth chunk size is 18 kbps; all PXF VC scr will be programmed as multiples of 18 kbps. For instance, if the PVCs were configured with 50 kbps as pcr, 54 kbps would be programmed in PXF. **show atm pvc** display will show 50 kbps, and VSI Slave will account 50 kbps during CAC. However, 54 kbps is actually being used. So as a result, when bandwidth usage is reaching the maximum value, both VSI Slave and PNNI will continue to allow connection provisioning, because VSI Slave and PNNI available bandwidth shows more than PXF actually has left.
- **Saveallcnf** (issued on the PXM45/B card) captures configuration data saved by the RPM-XF card (as well as AXSM and PXM45 cards), and saves it on the active PXM45/B card's hard disk. Configure the RPM-XF to store its configuration on the PXM45/B hard disk (E:/RPM) by entering **boot config e:auto_config_slot#** in the running configuration of the RPM-XF. To ensure that the saved file contains the latest RPM-XF configuration, execute the **write mem** command on each RPM-XF card prior to the entering **saveallcnf** command. This also ensures that the RPM-XF files on the active PXM45 hard disk will contain the latest configuration to be saved.
- For ELSR to LSC connectivity, the default control VC used is 32. If PNNI partition exists with VCI 32 as part of its partition range, when an MPLS partition is added, there are two options to handle the situation:
 - Add the MPLS controller and define its partition with available range. On eLSR, define control VC from any VCI value within the range defined in partition. The same VC should be defined on LSC on xTag interface.
 - Reconfigure PNNI partition to spare the control VC usage both on RPM-XF and AXSM, AXSM/B or AXSM-E APS Management Information.
- Whenever the RPM-XF configuration is changed, enter the **write mem** command on the RPM-XF to save the configuration. If this is not done, the changed configuration will be lost on an RPM-XF card reboot or RPM-XF switchover, in the case of redundancy.

RPM-XF auto_config File Management

The RPM-XF *auto_config_slot#* file stores the configuration for the RPM-XF card. The *slot#* portion of the name should be set to the logical slot number that corresponds to the RPM-XF card. This file can be stored in bootflash or in the E:/RPM directory on the PXM45 hard disk. The configuration is also stored in NVRAM using the name startup-config.

When the RPM-XF card is inserted or rebooted, it searches for the configuration file in the following sequence:

1. If there is an auto_config file corresponding to its logical slot on the PXM45 hard disk, the RPM-XF card uses the configuration stored on the hard disk.

2. If boot variable points to configuration stored in the PXM45 hard disk or Bootflash and if the file is not found, the card comes up as Active-F with the default configuration.
3. If there is no auto_config file on the hard disk, then the NVRAM version is used.

**Note**

In case of RPM-XF redundancy, the configuration should always be stored in `auto_config_slot#` file in the E:RPM directory of the PXM45 hard disk. Failure to find the auto_config file will lead to aborting of a user-initiated switchover (**switchredcd**) and a fatal error will be flagged.

Card Management

The following card management notes and cautions should be reviewed before using this release.

- There is a new stable state displayed on the PXM **dspcds** command—**Boot-Hold**, which signifies that the RPM-XF is running the boot image only. On the RPM-XF, the prompt will display as `boot>`
- The run-time IOS image cannot be used as a bootloader to load a different IOS image.
- Change of console speed on the terminal server may cause the card to end up in the ROMMON state. To avoid this, set the config register to `0x2102`.

Another workaround is to enter **cont** on the ROMMON within 2 minutes of going into ROMMON state. This will bring the card to its original stable state.

**Note**

It is recommended to always use 9600 baud as the console speed.

- The IOS version of the runtime as well as the boot image will be displayed in the **dspcd**, **dsprevs**, and **dsprevs -s** output. The version will be displayed under the heading of IOS version. Revision Control is not available for RPM-XF (like RPM-PR).

**Note**

The commands **loadrev** and **setrev** do not apply for RPM-XF.

RPM-XF Bootflash Precautions

The RPM-XF bootflash is used to store boot image, configuration and *run-time* files. Erasing the boot image from the Flash will cause the card to not boot.

The RPM-XF boot image, which comes loaded on the Flash, will work for all RPM-XF IOS images. Therefore, there is no reason to delete or move the factory installed boot image.

In order to avoid any unnecessary failures that would require card servicing, do the following:

- Never erase the boot file from the RPM Flash
- Never change the position of the boot file on the RPM Flash
- Use care when “squeezing” the Flash to clean it up.

As long as the boot file remains intact in the first position on the flash, the RPM-XF will boot successfully.

If the bootflash is corrupted, use the **tftpdnld** procedure described in the *Cisco MGX Route Processor Module (RPM-XF) Installation and Configuration Guide* or xmodem procedure described in [“Using XModem to Download Flash to RPM-XF Cards”](#) later in this document to download a new boot image.

Resolved Caveats—Release 12.2(15)T5

Table 1 lists the resolved caveats in Cisco IOS Release 12.2(15)T5.

Table 1 Resolved Caveats in Cisco IOS Release 12.2(15)T

CSCeb06375	<p>Symptoms:</p> <p>An access list may fail to work as configured.</p> <p>Conditions:</p> <p>This symptom may be observed when virtual circuits (VCs) are torn down and recreated. The symptom occurs after the switch subinterface is shut down and then brought up.</p> <p>Workaround:</p> <p>Remove the access list and add it again.</p>
CSCea49948	<p>Symptoms:</p> <p>Multiple crashinfo files may be saved on a Cisco Route Processor Module XF (RPM-XF) because of a reused IPC buffer (a second IPC send). If multiple crashinfo files are generated, the available storage space in the bootflash might get all consumed. If the card is part of a redundant pair, the card may fail during an attempt to switch back to the primary card from the secondary card in the redundant pair.</p> <p>Conditions:</p> <p>The symptom is observed when a switchover occurs from an RPM-XF to the redundant RPM-XF. If there is not enough bootflash free storage space available to load the configuration file, RPM-XF card might not go to Active state. After multiple switchover attempts, the Processor Switch Module (PXM) marks the first RPM-XF as failed.</p> <p>Workaround:</p> <p>Delete some crashinfo files to clear space on the bootflash before attempting to switch back to the active RPM-XF card.</p>
CSCea93735	<p>Symptoms:</p> <p>Control traffic may not be dequeued from a Parallel Express Forwarding (PXF) processor towards a Route Processor (RP).</p> <p>Conditions:</p> <p>This symptom is observed on a Cisco MGX 8800 series Route Processor Module XF (RPM-XF) in a Multiprotocol Label Switching (MPLS) environment. In a cell-based MPLS network, the symptom occurs when an MP LS packet with a Time To Live (TTL) setting below 2 reaches a provider edge (PE) router. In a frame-based MPLS network, the symptom occurs when an MPLS packet with explicit null labels and with a TTL setting below 2 reaches a provider (P) or PE router.</p> <p>Workaround:</p> <p>Configure the "no mpls ip propagate-ttl" global configuration command on all the routers in the MPLS network to prevent MPLS packets with a TTL setting below 2 from being generated.</p>

Open Caveats—Release 12.2(15)T5

Table 2 lists the open caveats in Cisco IOS Release 12.2(15)T5.

Table 2 Open Caveats in Cisco IOS Release 12.2(15)T5

Caveat Number	Description
CSCec16481	<p>Symptom:</p> <p>A Cisco device running Internetwork Operating System (IOS) and enabled for the Open Shortest Path First (OSPF) Protocol is vulnerable to a Denial of Service (DoS) attack from a malformed OSPF packet. The OSPF protocol is not enabled by default.</p> <p>Conditions:</p> <p>The vulnerability is only present in IOS release trains based on 12.0S, 12.2, and 12.3. Releases based on 12.0, 12.1 mainlines and all IOS images prior to 12.0 are not affected. Refer to the Security Advisory for a complete list of affected release trains.</p> <p>Workaround:</p> <p>Further details and the workarounds to mitigate the effects are explained in the Security Advisory which is available at the following URL: http://www.cisco.com/warp/public/707/cisco-sa-20040818-ospf.shtml.</p>
CSCea27838	<p>Symptoms:</p> <p>The following tracebacks were observed during router boot-up:</p> <pre>%SYS-3-CPUHOG: Task ran for 3100 msec (87338/197), process = TC-ATM Proc, PC = 40B2EAE8. -Traceback= 40B2EAF0 40B30BA0</pre> <p>Conditions:</p> <p>This traceback was observed on a Cisco MGX 8850 Route Processor Module (RPM-XF) during boot-up time. If this traceback is observed, RPM-XF card may hang for a while with user not able to "cc" to this card because of CPU hogging all cycles. While the card is in hang state for a while, LDP/OSPF could go down with CPUHOG error/traceback logged.</p> <p>Workaround:</p> <p>There is no workaround.</p>
CSCeb05118	<p>Symptom:</p> <p>RPM-XF router, while configured as eLSR, might reload when deleting MPLS sub-interfaces</p> <p>Conditions:</p> <p>On routes removal, stray LVCs are not cleaned up/removed. Under such condition deletion of MPLS interface may lead to RPM-XF reset.</p> <p>Workaround:</p> <p>There is no workaround</p>

Table 2 Open Caveats in Cisco IOS Release 12.2(15)T5 (continued)

CSCea60343	<p>Symptom: Connection goes into mismatch after adding from xf to pr/b</p> <p>Conditions: addcon between RPM-XF to RPM-PR/B 12.2(15)t1 3.0(20.100)</p> <p>Workaround: under investigation</p>
CSCdv59661	<p>Symptom: An SPVP connection goes into the mismatch state after modifying the PCR of the corresponding "atm pvp" tunnel.</p> <p>Conditions: The SPVP was either an incomplete connection, the master endpoint of a complete dax connection or any (master/slave) endpoint of a non-dax connection. The PCR of the corresponding "atm pvp" was modified to a new value.</p> <p>Workaround: Delete and readd the "sw conn vpc".</p>
CSCdw45040	<p>Symptom: RPM-XF comes up with partial configuration</p> <p>Condition: This could happen if writemem is issued on the RPM-XF that is transitioning from standby to Active state.</p> <p>Workaround : As a precautionary measure do not execute writemem while the card is not in Active state. To restore the original configuration, use the back-up config (if available).</p>
CSCdw68935	<p>Symptoms: Once in a while VSI Error appears on RPM-XF: VSICORE-4-VSICGENERR: VSICORE: VsiError: VsiErr:Connection Reassert Error (VcoEntry), 0x5011,2327,810C1802,10D1802,D,0,0,38,46,41,131,1</p> <p>Conditions: Occurs once in a while when connection exists between RPM-XF and AXSM/AXSM-E and AXSM/AXSM-E is reloaded.</p> <p>Workaround: There is no workaround for this problem. When the VSI Error is displayed, connection is not impacted. It is confirmed that traffic still flow through connection correctly.</p>

Table 2 Open Caveats in Cisco IOS Release 12.2(15)T5 (continued)

CSCdx06018	<p>Symptom:</p> <p>Output traffic drop on VBR VCs is at or below the SCR when multiple streams are going to same destination via multiple VBR VCs.</p> <p>Conditions:</p> <p>This happens when there are equal cost paths via VBR VCs. The PXF queue selection algorithm may cause traffic to drop for multiple streams going to the same destination via multiple paths.</p> <p>When PXF gets a packet, it selects the out going interface based on the source and the destination IP addresses. These addresses are hashed into only one of the queues for the selected destination. If multiple paths also exist for the same destination, multiple streams may be hashed into one queue, causing some queues to overflow and others to be under-utilized.</p> <p>Workaround: None</p>
CSCdx07534	<p>Symptom:</p> <p>VSI reassert error happened on ELSR when "switchredcd" or "resetcd" on ELSR is executed.</p> <p>Example of errors:</p> <pre data-bbox="618 926 1520 1010">-Process= "RPMXF VSIS", ipl= 0, pid= 103 %VSICORE-4-VSICGENERR: VSICORE: VsiError: VsiErr:Connection Reassert Error (VcoEntry)</pre> <p>The number of errors depends on the number of LVCs being established.</p> <p>Conditions:</p> <p>VSI reassert error/warning messages show up on console terminal when switchredcd or resetcd was performed on the ELSR.</p> <p>Workaround:</p> <p>None. The error message is harmless.</p>
CSCdx92871	<p>Symptom:</p> <p>iBGP load balancing does not work when two CEs are configured to be in different VPNs.</p> <p>Conditions:</p> <p>iBGP multi-path was configured on the PE with two paths to one of the CEs say CE2. The traffic was passed from CE1 (connected directly to the PE) to CE2 where the two CEs were configured to be in different VPNs.</p> <p>Workaround: None</p>
CSCdy05346	<p>Symptom:</p> <p>Shelf name, sw_type_str are not present in the Switch Get Configuration. Response from RPM-XF VSI slave.</p> <p>Conditions:</p> <p>LSC shows some of the fields in the display commands 'show contr xtag' and 'show contr vsi ses {sesn #}' as empty for an RPM-XF VSI slave.</p> <p>Workaround: None</p>

Table 2 Open Caveats in Cisco IOS Release 12.2(15)T5 (continued)

CSCdy17457	<p>Symptoms: "sh int" command output shows wrong packet number after shut/no shut interface</p> <p>Conditions: If a user executes a shut/no-shut on the interface, when traffic is being pumped, one might see wrong output from "sh int" command.</p> <p>Workaround: Clearing the counters on that interface should fix the problem.</p>
CSCdy26882	<p>Symptom: Interface counters on the Gigabit Ethernet interface show much higher than expected values for 30 second output packets and bits/s.</p> <p>Conditions: The interface is shutdown and the backcard is removed and inserted again. Then the interface is enabled. Bidirectional data flow being established in quick succession with interface being enabled.</p> <p>Workaround: None</p>
CSCdy38285	<p>Symptom: RPMXF has different maximum allowed packet size for input and output traffic. The output maximum allowed size is smaller.</p> <p>Conditions: It has been observed on the 1-port POS card. Not clear if it happens on other interfaces.</p> <p>Workaround: Increase the MTU size of the interface to allow packets of larger size to go through.</p>
CSCdy42274	<p>Symptom: Performance may drop for max data rates if the PXF is reloaded.</p> <p>Conditions: When the traffic is pumped at a high rate to get the max throughput possible on the interface, after PXF reload performance may deteriorate.</p> <p>Workaround: Perform a shut/no shut on the interface.</p>
CSCdy48760	<p>Symptom: Execution of `show controller gig1/0' command output shows that remote end auto-negotiate mode is set to `auto' rather than to `force'.</p> <p>Conditions: When the auto-negotiation is enabled on the local GigE interface and disabled on the remote GigE interface with both interfaces UP, `show controller gig1/0' command output shows remote end set to "auto" rather than to "force".</p> <p>Workaround: None</p>

Table 2 Open Caveats in Cisco IOS Release 12.2(15)T5 (continued)

CSCdy55392	<p>Symptom: PXF reload while traffic is passing through the Gigabit Ethernet interface.</p> <p>Conditions: When traffic is passing through the Gigabit Ethernet interface and the interface MAC Address on Gigabit Ethernet interface is modified through CLI, PXF may reload.</p> <p>Workaround: 'shut' the GigE interface, configure the MAC address and 'no shut' the interface.</p>
CSCdy56575	<p>Symptom: Cannot remove "mpls atm vpi 0-10" from XTagATM interface on LSC.</p> <p>Condition: "no mpls atm vpi 0-10" command can not remove the vpi range configuration from XTagATM interface.</p> <p>Workaround: None</p>
CSCdy73751	<p>Symptom: The outgoing traffic on GigabitEthernet interface stopped flowing.</p> <p>Conditions: With auto-negotiation enabled on the GigE interface and outward traffic flowing; if the hardware sends certain error interrupts, traffic may stop.</p> <p>Workaround: Performing a shut/noshut on GigE interface should restart the traffic.</p>
CSCdy81782	<p>Symptom: "no shut" on the PPP interface before VA goes down causes PXF to drop certain packets. This results in being unable to have successful ping through routes involving such a PPPoA subinterface from a remote node. Arriving ICMP packets are dropped in PXF.</p> <p>Conditions: Under a PPPoA subinterface, if "shut" and "no shut" commands are issued in succession such that "no shut" is completed before the associated virtual- access goes down. When node is in this condition, pings originating from the local node and non-ICMP traffic appears to work correctly. Only pings from a remote node would fail.</p> <p>Workaround: After a "shut", wait until the associated virtual-access changed to go down before issuing "no shut". If symptom occurs, go to the subinterface, perform a "shut", wait for the associated virtual-access to go down and then perform a "no shut".</p>

Table 2 Open Caveats in Cisco IOS Release 12.2(15)T5 (continued)

CSCdz36589	<p>Symptom: All traffic is dropped on VBR connection if PCR is configured to 800 Mbps.</p> <p>Conditions: Provision VBR connection with PCR equal to 800 Mbps. Pass traffic below SCR rate. All traffic will be dropped.</p> <p>Workaround: Configure PCR at 790 Mbps.</p>
CSCdz63683	<p>Symptom: When issuing the "sh policy-map int <sub-interface>" command, matched number of packets does not equal to output packets + dropped packets.</p> <p>Conditions: (1) attach a policy map with classes to a sub-interface; (2) start traffic; and ((3) issue "sh policy-map int <sub-interface name>" command.</p> <p>Workaround: None</p>
CSCdz86609	<p>Symptom: Packet drop was observed at switch interface1 when traffic flowing through.</p> <p>Condition: This occurred when pumping through traffic of about 150Mbps from one ELSR to another ELSR then to a CE and the IP traffic rate is below the available PNNI partition bandwidth.</p> <p>Workaround: None</p>
CSCea15938	<p>Symptom: GTS shapes traffic rate too aggressively for POS/GigE interfaces.</p> <p>Conditions: Pass traffic with dscp bit set to af13, GTS will drop packets when traffic rate reaches 50-60% of the shaping rate.</p> <p>Workaround: GTS works properly when sending traffic to multiple VPI/VCIs. If sending a traffic to a single VPI/VCI, double the shaping rate in order to receive the correct amount of traffic.</p>
CSCea74420	<p>Symptom: Renaming a policy map does not reflect on the running configuration of the VC the policy map is attached to.</p> <p>Conditions: A policy map is renamed using the "rename" command.</p> <p>Workaround: Delete the old policy map and add a new policy map.</p>

Compatibility Notes

RPM-XF Boot File and Firmware File Names and Sizes

The following table displays the RPM-XF boot and firmware file names and sizes for this release.

Table 3 RPM Boot and Firmware File Names and Sizes

	File Name	File Size (in bytes)
Boot File	rpmxf-boot-mz.122-15.T5	3049044
Firmware File	rpmxf-p12-mz.122-15.T5	8213336

RPM-XF Compatibility Matrix

Table 4 RPM Boot and Firmware File Names and Sizes

	File Name	File Size (in bytes)
Boot File	rpmxf-p12-mz.122-15.T	8205732
Firmware File	rpmxf-boot-mz.122-15.T	3047076

MGX SW version	3.0.10	4.0.10
IOS Version	12.2.15T	12.2(15)T5
CWM	11.0.10	12.0.00.1

MGX RPM-XF Hardware

[Table 5](#) shows the front card and back card compatibility for the RPM-XF hardware supported in this release. The table lists the card model/ name, part numbers, the minimum version and the minimum revisions of each card supported. Note that there may be more than one 800 level part numbers for the same front cards. The minimum version is identified by the last 2 digits of the 800 level numbers.

Table 5 Hardware Compatibility Matrix

Front Cards	Part Number/ Min. Version	Rev.	Back Cards	Part Number/ Min. Version	Rev.
MGX-RPM-XF-512	800-09307-03	A0	MGX-XF-UI	800-09492-01	A0
			MGX-1GE	800-18420-03	A0
			MGX-1OC12POS-IR	800-08359-05	A0

Table 6 SFP Compatibility Matrix for MGX-1GE

SFPs	Part Number/ Min. Version	Rev.
MGX-GE-SX	30-1301-01	A0
MGX-GE-LHLX	30-1299-01	A0
MGX-GE-ZX	10-1439-01	A0

Cisco IOS Release Compatibility Information

All IOS firmware can be downloaded from CCO from the following location:

<http://www.cisco.com/kobayashi/sw-center/sw-ios.shtml>

Using XModem to Download Flash to RPM-XF Cards

Use the xmodem feature to download the flash to an RPM-XF card. During this process, the card should be connected to a target machine through HyperTerminal with settings of 9600, n, 8, and 1.

- Step 1** Put the node in monitor mode by entering the **priv** command to gain access to the privileged commands as follows:

```
rommon 1> priv
You now have access to the full set of monitor commands. Warning:
some commands will allow you to destroy your configuration and/or
system images and could render the machine unbootable.
```

- Step 2** The xmodem command becomes available and the general syntax of this command and availability of this can be checked by giving xmodem command without any parameters on the CLI, as follows:

```
rommon 2 > xmodem
usage: xmodem [-cy]
-c CRC-16
-y ymodem-batch protocol
rommon 3 >
```

The command line options for xmodem are as follows:

Option	Definition
-c	xmodem performs the download using CRC-16 error checking to validate packets. Default is 8-bit CRC.
-y	xmodem uses Ymodem-batch protocol for downloading, which uses CRC-16 error checking.



Note If you do not find the xmodem commands, then the xmodem feature is not available on this rommon version. In that case, you must return the card to Cisco.

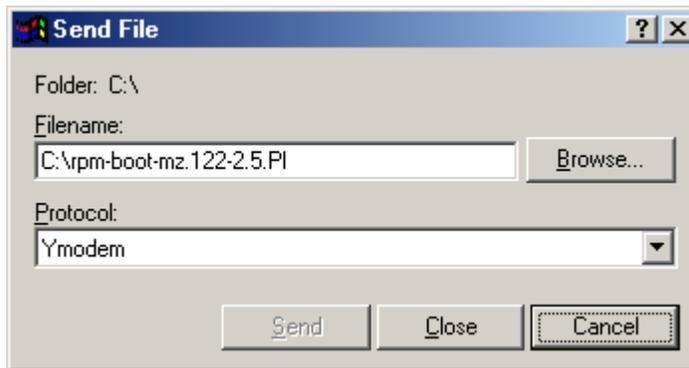


Note The rommon “xmodem/ymodem” transfer only works on the console port. You can only download files to the router. You cannot use “xmodem/ymodem” to get files from the router.

For example:

```
rommon 4> xmodem -cy
Do not start sending the image yet...
Invoke this application for disaster recovery. Do you wish to
continue? y/n [n]: y
```

Step 3 Use the Transfer-->Send File option in HyperTerminal to start the image transfer.



In the Filename box, browse and choose the image file to be downloaded. Also since we used the “y” option while invoking the xmodem, set the transfer protocol to ymodem or use Xmodem protocol by not specifying the –y option on the command line.

The transfer screen comes up and transfer starts. (The transfer may not start immediately; wait for some time and it should start.)

The system will reset itself from here and will boot with new software image.

Historical Information for 12.2.x IOS Baseline

Resolved Caveats in Release 12.2.15T

Table 7 lists the resolved caveats in Cisco IOS Release 12.2.15T.

Table 7 Resolved Caveats in Cisco IOS Release 12.2.15T

Caveat Number	Description
CSCea49948	<p>Symptoms:</p> <p>Multiple Crashinfo due to IPC Messages. Multiple crash info files are generated, filling the bootflash of RPM-XF card.</p> <p>Conditions: It may happen due to polling that involves IPC communication between Processor Switch Module 45 (PXM45) controller and RPM-XF card. RPM-XF switchover may lead to this issue.</p> <p>Workaround: None</p>
CSCdz82543	<p>Symptoms:</p> <p>Cannot cc to RPM-XF due to Messages on the Console. You may not be able to log into a Cisco Route Processor Module XF (RPM-XF), although when you display the status of the module from a Processor Switch Module 45 (PXM45) controller, no irregularities are shown. If you manage to establish a console connection into the RPM-XF, continuous traceback messages may be displayed on screen.</p> <p>Conditions:</p> <p>These symptoms are observed when RPM-XFs are installed in an MGX45 shelf, Cisco WAN Manager is configured to poll the RPM-XFs for statistics and scripts run on the node that log into every module sequentially.</p> <p>Workaround: None</p>
CSCdx08155	<p>Symptom:</p> <p>On LSC, querying of lvc statistics for an xtagatm interface would not abort command upon user entering a ctrl-c. If user use "show xtagatm cross-connect traffic" to query on lvc statistics, normally, user can quit the command in the middle by giving the ctrl-c sequence. However, the CLI would not return the prompt until the VSI Master logic complete requesting statistics for all lvc.</p> <p>Conditions:</p> <p>If the number of lvc on LSC is small, user may not be able to notice the impact. However, if the number of lvc reach a few thousands (i.e. 4000 lvc), user may experience no response from the cli session up to 2-3 minutes. But the rest of the system will still function normally.</p> <p>Workaround:</p> <p>If user needs to query on lvc statistics, try to let the command run to completion. Avoid aborting the command in the middle of execution with ctrl-c.</p>

Table 7 Resolved Caveats in Cisco IOS Release 12.2.15T (continued)

Caveat Number	Description
CSCdy26703	<p>Symptoms:</p> <p>A ping from a customer edge (CE) to a provider edge (PE) may fail, and Parallel Express Forwarding (PXF) may stall.</p> <p>Conditions:</p> <p>These symptoms are observed on an MGX Route Processor Module (RPM) that is installed in a Cisco MGX 8000 series switch that is functioning as a PE router in a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) network and that has VPN routing/forwarding (VRF) enabled on the virtual template. The PE router is connected to the CE router via a PPP over ATM (PPPoATM) link. PXF may stall even though the PPPoATM session is established.</p> <p>Workaround: None</p>
CSCdy27120	<p>Symptom:</p> <p>Traffic doesn't flow thru a VLAN on GigE interface</p> <p>Conditions:</p> <p>If a shutdown is executed on an interface before defining the encapsulation type, traffic might not flow thru a VLAN on a GigE interface.</p> <p>Workaround:</p> <p>On the VLAN interface, disable and re-enable the encapsulation.</p>
CSCdy65600	<p>Symptoms:</p> <p>The output and input flow control parameters of a Gigabit Ethernet interface are displayed as ? "aused".?</p> <p>Conditions:</p> <p>This symptom is observed on the Gigabit Ethernet interface a Cisco router that has auto-negotiation enabled.</p> <p>Workaround:</p> <p>Disable auto-negotiation on the Gigabit Ethernet interface.</p>
CSCdz70762	<p>Symptom:</p> <p>Multi-vc traffic coming on a particular precedence goes out to a queue with wrong precedence.</p> <p>Condition:</p> <p>The mac rewrite string is found to be incorrectly configured.</p> <p>Workaround: None</p>

Table 7 Resolved Caveats in Cisco IOS Release 12.2.15T (continued)

Caveat Number	Description
CSCdz82543	<p>Symptoms:</p> <p>You may not be able to log into a Cisco Route Processor Module XF (RPM-XF), although when you display the status of the module from a Processor Switch Module 45 (PXM45) controller, no irregularities are shown. If you manage to establish a console connection into the RPM-XF, continuous traceback messages may be displayed on screen.</p> <p>Conditions:</p> <p>These symptoms are observed when RPM-XFs are installed in an MGX45 shelf and Cisco WAN Manager is configured to poll the RPM-XFs for statistics, causing scripts to run on the node and log into every module sequentially.</p> <p>Workaround: None</p>
CSCea05477	<p>Symptom:</p> <p>After policy map is created, RPM-XF reset.</p> <p>Conditions:</p> <p>Created a policy map with police option with CIR rate configured to 125000000 and bc configured to 20000 for confirm action and exceed action was drop. Soon after creating policy map it gave</p> <pre>Conform burst size increased to 62500</pre> <p>and card resetted.</p> <p>Workaround: None</p>
CSCin32860	<p>Symptom:</p> <p>Access list info of snmp-server community lost after RPM-XF reset</p> <p>Conditions:</p> <p>After RPM-XF is reset , the access list info of the snmp-server community is lost.</p> <p>Workaround:</p> <p>After RPM-XF is reset, re-configure the snmp-server community string as required on RPM-XF.</p>

Resolved Caveats Prior to Release 12.2.15T

Table 8 lists the resolved caveats prior to Cisco IOS Release 12.2.15T.

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T

Bug ID	Description
CSCdw20568	<p>Symptom: Cisco Class-Based QoS mib - CISCO-CLASS-BASED-QOS MIB - is not supported on RPM-PR and RPM-XF.</p> <p>Condition: SNMP walk on CISCO-CLASS-BASED-QOS mib on RPM-PR and RPM-XF cards return no values for mib objects.</p> <p>Workaround: None.</p>
CSCdw55382	<p>Symptom: The output of the command sh swi conn vcc/vpc doesn't show the value of the maximum cost field correctly.</p> <p>Conditions: Maximum cost was configured explicitly on the master endpoint of the connection to a value of 4294967295 i.e. 0xFFFFFFFF.</p> <p>Workaround: Use dspon command on the PXM to look at the configured value.</p>
CSCdw57105	<p>Symptom: Show sub-interface counter shows incorrect value.</p> <p>Condition: When there is some drop on the sub-interface, value becomes negative, causing the numbers to be incorrect.</p> <p>Workaround: Issue clear counter command.</p>
CSCdw68738	<p>Symptom: Cobalt From RP Own Errors counter increments in show hard pxf dma count output. This does not affect data/traffic.</p> <p>Conditions: Spontaneous increments of these counters have been noticed always with no specific side effect.</p> <p>Workaround: None</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdw69661	<p>Symptoms:</p> <p>Invalid Epid Error message seen.</p> <pre>00:00:10: %P2IPC-4-COMEPDELETED: ssi_ipc_epid_idx_validate() Non-existing CommEp 60010F8 has invalid tag 4096; Expected tag is 0 -Process= "P2IPC Receive Process", ipl= 0, pid= 17</pre> <p>Condition:</p> <p>After RPM-XF switchover, standby card prints invalid Epid Error messages as VSI Slave Epid on Standby is not in use.</p> <p>Work around:</p> <p>None</p> <p>Further problem description:</p> <p>The error messages are harmless traceback as it happens on the Standby card.</p>
CSCdw88019	<p>Symptom:</p> <p>Loopbacks provided on GigE backcard should be renamed to 'internal' and 'external' from 'mac' and 'driver', respectively</p> <p>Conditions:</p> <p>Previously, RPM-XF software used "mac" and "driver" as loopback command parameters to mean internal and external loopbacks respectively. These parameters have been replaced by "internal" and "external" as loopback command parameters.</p> <p>Workaround:</p> <p>Use show controllers gigabitEthernet 1/0 command and user documentation to correlate the loopback type with the loopback configuration.</p>
CSCdw88767	<p>Symptom:</p> <p>Humvee counters show improper value and counters cannot be cleared.</p> <p>Conditions:</p> <p>While traffic is passing through, counters wrap around and become negative.</p> <p>Workaround:</p> <p>None.</p>
CSCdw95563	<p>Symptom:</p> <p>After increasing the PCR value of PVP, traffic is getting dropped at new rate.</p> <p>Condition:</p> <p>When PCR value is changed, it does not take affect so VP is still shaped at old rate.</p> <p>Workaround:</p> <p>Delete the VP and readd it.</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdx00982	<p>Symptom:</p> <p>SNMP get returns a different value for pcr/scr from what was configured.</p> <p>Conditions:</p> <p>The connection was added through CWM/SNMP with a value such that on conversion from cps to kbps and then back to cps does not give back the original cps value.</p> <p>Workaround:</p> <p>While adding connections through CWM/SNMP, do not use such values (in cps) for scr/pcr for which conversion from cps to kbps and then back to cps does not give back the original value.</p>
CSCdx12730	<p>Symptom:</p> <p>All the PVCs on the switch1 interface go to INACTIVE state.</p> <p>Conditions:</p> <p>Any destructive command can cause this condition. Some of them are:</p> <ol style="list-style-type: none"> 1. clear interface switch1 command was issued. 2. Modification of the atm pvp tunnel. 3. Changing MTU value of Switch1 interface or sub-interface. <p>Workaround:</p> <p>Issue clear interface switch1 command.</p>
CSCdx16897	<p>Symptom:</p> <p>Performance issue observed in cleaning up and creating LVCS.</p> <p>Conditions:</p> <p>With RPM-XF as headend eLSR, it takes about 1 minute to cleanup and recreate 500 LVCS.</p> <p>Workaround:</p> <p>None</p>
CSCdx44836	<p>Symptom:</p> <p>Modifying an existing PVP may cause the following VSI error to be displayed on console or logged:</p> <pre>04:57:14: %VSI_VRM-4-GENERR_NUM: VSIRmGetXConnectInfo, line 6658: Vsis RM error <Failed to search Vco database for lcn =>, info=1</pre> <p>Conditions:</p> <p>This error only happens if the existing PVP has a configured switch-conn-vpc but is not routed yet (only slave end is added).</p> <p>This error is informational only to indicate that there is no remote cross-connect information to be returned because the vpc is not routed yet. There is no service impact.</p> <p>Workaround:</p> <p>Complete the VPC connection by adding the remote end.</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdx46583	<p>Symptom: Need to verify IOS images on PXM Hard Drive and in RPM-XF Flash.</p> <p>Conditions: In order to do the health check of the IOS images on PXM Hard disk & Bootflash, a new CLI command is being introduced: debug rpm check_image now x:<image-name>.</p> <p>Workaround: None.</p>
CSCdx49122	<p>Symptom: dspcd <slot#> for RPM-XF slot doesn't show the full CLEI code / Serial number. One character at the end is missing.</p> <p>Condition: When dspcd command is executed on PXM.</p> <p>Workaround: cc to RPM-XF card and do sh rpm cdmgmt scmExtPollInfo to find the correct CLEI code and Serial number.</p>
CSCdx52025	<p>Symptom: Could not correlate output packets dropped on sub-interface with switch1 interface packet drop counters.</p> <p>Condition: Sub interface counters don't have any drop counters, so if any packet is dropped on a particular sub-interface, it is shown on switch1 interface but not on that sub-interface.</p> <p>Workaround: Use sh atm pvc <vpi/vci> to see all the drops.</p>
CSCdx55586	<p>Symptom: Setting ccCopyEntryRowStatus to ACTIVE returns "general error" status even if the row is correctly configured.</p> <p>Conditions: Attempting to use the CISCO-CONFIG-COPY-MIB to copy a configuration file.</p> <p>Workaround: None.</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdx58504	<p>Symptom:</p> <p>RPM-XF show switch conn vcc vpc displays NSAP in following format: 47.0091.8100.0000.0001.6443.6c58.0000.0109.1802.00</p> <p>which is not consistent with PXM dsicons display.</p> <p>Conditions:</p> <p>RPM-XF "show switch conn vcc vpc" displays NSAP in following format: 47.0091.8100.0000.0001.6443.6c58.0000.0109.1802.00</p> <p>which is not consistent with PXM "dsicons" command output: 47.009181000000000164436c58.000001091802.00</p> <p>Workaround:</p> <p>None.</p>
CSCdx62385	<p>Symptom:</p> <p>flapping of BGP causes RPM-XF reload.</p> <p>Condition:</p> <p>The reload happened after the routes exceeded 102K vrf route limit and clear ip bgp * is executed.</p> <p>Workaround:</p> <p>Do not exceed the given limit of Maximum number of vrf routes which is 102k.</p>
CSCdx64337	<p>Symptom:</p> <p>After changing the console baud rate the console may behave unpredictably.</p> <p>Conditions:</p> <p>The console responsiveness is unpredictable. You may need to hit enter many times to get a prompt back.</p> <p>Workaround:</p> <p>Avoid changing the baud rate of the console port. To recover from the erratic console behavior, you will need to perform a system reload.</p>
CSCdx64361	<p>Symptom:</p> <p>ROMMON console can lock up after pasting a large buffer.</p> <p>Conditions:</p> <p>The ROMMON console driver can run out of descriptors with a sufficiently large paste buffer (~1k). After this happens, the console appears hung.</p> <p>Workaround:</p> <p>Paste data to the ROMMON console in increments less than 1k. Once the ROMMON console driver runs out of descriptors, the only way to recover is to perform a system reset or power cycle.</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdx69702	<p>Symptom:</p> <p>The output counters displayed under show policy-map int <swl.x> are not incremented.</p> <p>Condition:</p> <p>When an output service policy is applied to a subinterface and the policy map includes class maps with no actions.</p> <p>Workaround:</p> <p>Either do not configure class maps with no actions or rearrange the class maps within a policy map so the class maps with no actions are last.</p>
CSCdx71190	<p>Symptoms:</p> <p>A software-forced reload may occur on a router, and the Open Shortest Path First (OSPF) process may fail.</p> <p>Conditions:</p> <p>This symptom is observed on a Cisco 7200 series router that is running Cisco IOS Release 12.2(10.7)T1 and that is configured for Border Gateway Protocol (BGP) tag switching with OSPF in an autonomous system. The router may reload when the main link bandwidth is changed to switch to the shortest path. This behavior may cause the router OSPF process to fail.</p> <p>Workaround:</p> <p>There is no workaround.</p>
CSCdx76951	<p>Symptom:</p> <p>There was humvee error on RPM-XF card.</p> <p>Condition:</p> <p>When ever switchcc is done or card comesup, humvee generates some errors. These errors are harmless, unless they are incrementing coutinously.</p> <p>Workaround:</p> <p>None.</p>
CSCdx80500	<p>Symptom:</p> <p>A CLI command is needed to show the history of the messages that RPM-XF received from the Shelf Manager on PXM. Also, a second command is needed to clear this log.</p> <p>Condition:</p> <p>NONE</p> <p>Workaround:</p> <p>None</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdx87265	<p>Symptom: Deletion trap is not sent out for notOnRpm connections.</p> <p>Condition: RPM-XF reset or switch over before provisioned connections get saved into configuration.</p> <p>Workaround: None.</p>
CSCdx91454	<p>Symptom: The status LEDs for the management backcard are not illuminated correctly.</p> <p>Conditions: When the boot image is used, the status LED on the back of the management back card is not illuminated as it should be. With the system image status LED on the back of the management back card is illuminated but the LM1OK LED on the front card is off instead of green.</p> <p>Workarounds: Use the 'show interface FastEthernet2/0' and 'show interface FastEthernet2/1' commands to view the status of the management backcard.</p>
CSCdx93773	<p>Symptoms: Packet drop on egress sub-interface below configured rate.</p> <p>Conditions: Sending traffic at SCR rate on a vbr pvc with high SCR value configured.</p> <p>Workaround: None</p>
CSCdy02182	<p>Symptom: When the GigabitEthernet device driver detects an error with the link to the front card, it does not automatically try to correct the situation properly.</p> <p>Conditions: When the GigabitEthernet device driver detects an error with the link to the front card you may see messages similar to: <pre>BARIUM ERROR EVENT: send reason 7 slot 1</pre> <p>The device driver should then try to reset the link to the front card and continue but does not.</p> <p>Workaround: When you receive an error message similar to the one described in conditions, you should shutdown and restart the interface to clear the condition.</p> </p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdy03275	<p>Symptom: Traffic is not passing in frame-based MPLS network when RPM-XF is configured as P router</p> <p>Conditions: RPM-XF configured as P router.</p> <p>Workaround: None</p>
CSCdy05871	<p>Symptom: Tail drops on PXF queue while sending traffic at OC12 rate.</p> <p>Conditions: Configure interface speed as more than 500Mbps (X) and send more than X/2 single flow traffic</p> <p>Workaround: Use multiple flows (different source-destination)</p>
CSCdy09544	<p>Symptom: LLQ (Low Latency Queue) starves low priority traffic.</p> <p>Condition: Traffic exceed configured bandwidth for the LLQ, but still below the configured SCR (Sustain Cell Rate). There are drops on other low priority classes even though they are in their configured limit.</p> <p>Workaround: None</p>
CSCdy11581	<p>Symptom: Received traps for "Fast Ethernet Interface Down(60662)" and "Fast Ethernet Interface up(60661)" have incorrect <i>ifName</i> contents.</p> <p>Conditions: Generate "Fast Ethernet Interface up(60661)" by doing admin up on interface and notice the contents of "<i>ifName</i>" in the trap PDU on CWM or any other SNMP trap client. Do the same for "Fast Ethernet Interface Down(60662)" by doing admin down on the interface. The contents for "<i>ifName</i>" is "FastEthernet/12/0". The correct value should be FastEthernet/<log-slot>/<bay>/<port>.</p> <p>Workaround: None.</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdy15295	<p>Symptom: cbQosQueueingStats and cbQosREDClassStats MIB entries are not populated.</p> <p>Condition: SNMP mib walk on RPM-XF for cbQosQueueingStats and cbQosREDClassStats mib objects return no entries.</p> <p>Workaround: None.</p>
CSCdy23757	<p>Symptom: Data stops flowing from VLAN after removal and insertion of GigabitEthernet backcard is done.</p> <p>Conditions: After removal and insertion of GigabitEthernet backcard, data traffic from the configured VLAN stops.</p> <p>Workaround: Perform `no encaps dot1q` on the VLAN interface and then re-add all the configuration defined on that VLAN.</p>
CSCdy26495	<p>Symptom: class-map output queue packet counter doesn't show the correct number of packets.</p> <p>Conditions: With policy-map attached, send the traffic with the rate not to see drops</p> <p>Workaround: Use the class-map total number of packet counters.</p>
CSCdy26755	<p>Symptom: Execution of PXM command "dspcd" for the RPM-XF card does not show 800 Level Rev number for the frontcard and the backcard</p> <p>Conditions: Execution of PXM command "dspcd" on a particular slot, even when backcard and front card is present, the 800 level Rev information is blank.</p> <p>Workaround: `cc` to RPM and use the command `sh rpm eeprom`</p>
CSCdy27852	<p>Symptoms: Excessive delay for LLQ packets.</p> <p>Conditions: Send LLQ packets during congestion of the non LLQ queues. The issue only happens for vbr-nrt connections with rate <= 128 kbps, if excessive traffic is pumped on queues other than LLQ and default queue. The issue would happen for all rates if excessive traffic is pumped on the default queue.</p> <p>Workaround: none</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdy28132	<p>Symptoms: Traffic forwarding stops. You may see the traffic forwarded to the wrong VC.</p> <p>Conditions: Shut followed by no shut on the MPLS interface with multi-vc could lead to this issue. This is an intermittent issue. This could also happen if the LVCs go down and come back up for any other reason. The issue is seen only with VRF prefixes.</p> <p>Workaround: shut/no shut the interface. If problem persists, switch to a redundant card if one is available.</p>
CSCdy30260	<p>Symptom: Protocol flap is observed and data labeled transfer stops temporarily on RPM-XF when RPM-XF card in adjacent slot is removed. Traffic resumes after around 100 seconds.</p> <p>Conditions The issue happens when RPM-XF card in the adjacent slot is pulled out, the RPM-XF cards being in one of the following pair of slots: (1,2), (3,4), (5,6), (9,10), (11,12), (13,14), (15,16). (MGX 8850) (1,2), (3,4), (5,6), (11,12), (13,14), (15,16). (MGX 8950)</p> <p>This happens with or without 1:N redundancy.</p> <p>Workaround Do not use RPM-XF cards in adjacent slots with combination as above.</p>
CSCdy31406	<p>Symptom: RPM-XF frame-based P router's PXF reloaded after "shut" PE subinterface.</p> <p>Conditions: RPM-XF configured as P router. Upon "shut" on PE's subinterface, PXF module on P router got reloaded.</p> <p>Workaround: None.</p>
CSCdy37576	<p>Symptom Cannot add a dax connection between the RPM-XF (10) and the AXSM (1)</p> <p>Condition From the PXM perspective, the available CPS on the axsm is different than that of the rpm-xf. however when i am picking the lower value to add the pvc the pvc is still failing to add. the rpm-xf is showing reason as no resources.</p> <p>Work-around Unknown</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdy38362	<p>Symptom: Line Alarm seen on Gigabit ethernet interface on MGX-1GE even when administratively down.</p> <p>Conditions: If interface was in alarm prior to being shutdown, it continues to persist even after the shutdown.</p> <p>Workaround: Doing a 1:N redundancy switch over (if redundancy is configured) to the standby RPM-XF clears this alarm on the card and at the node level. This is to be done after the Gigabit Ethernet interface is shutdown, and a "write mem" has been done on the active. The alarm can be cleared by regular means such as ensuring proper cabling, SFP type match with other end, and administratively enabling other end prior to shutting down this end. If the other end is also not running the same software release that includes this bug then the alarm will move to the other end on trying such workaround. If redundancy is not configured, resetting the front card (disruptive) would be needed to clear the alarm.</p>
CSCdy39423	<p>Symptom: Traffic stops going out on GigE interface when enabling autonegotiation parameter.</p> <p>Conditions: With traffic flowing, if <i>autonegotiation</i> parameter is enabled on GigE interface, traffic may stop.</p> <p>Workaround: shut/no shut on the GigE interface</p>
CSCdy39806	<p>Symptom: "No switch partition configured" trace back error logged.</p> <p>Condition: Setup eLSR with vp tunnel mode on RPM-XF</p> <p>Workaround: None</p>
CSCdy39861	<p>Symptom: Spurious memory trace back error logged when disable VRF forwarding under switch sub-interface</p> <p>Conditions: Disable VRF forwarding under switch sub-interface on RPM-XF LER</p> <p>Workaround: None</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdy40930	<p>Symptoms: LLQ packets dropped on SAR because of lack of buffers</p> <p>Condition: Congest the pvc with the LLQ traffic</p> <p>Workaround: none</p> <p>Further problem description: The scheduler calculates the time it took to send a burst of traffic on a VC. This calculation results in an integer value which is a rounded down of the actual value. This rounding causes error in the scheduling time, which causes shaping inaccuracy. The fix is to try to keep the rounding error to a minimal.</p>
CSCdy41773	<p>Symptom: In case of GIGE backcard initialization failure, further configuration on it may cause the RPM-XF card to reboot.</p> <p>Conditions: If the GIGE backcard initialization fails, loading start-up config or manual configuration that changes the uninitialized GIGE backcard configuration, might cause RPM-XF to get reset.</p> <p>Workaround: None</p>
CSCdy45515	<p>Symptom: Connection endpoint on RPM-XF Does Not Generate RDI Upon Receiving AIS</p> <p>Conditions: RPM-XF endpoint of a xpvc does not generate RDI upon receiving AIS. The RPM-XF atm pvc erroneously reports that RDI is generated. However, a display of the connection count at the corresponding AXSM on the node indicates that the RDI is, in fact, NOT generated.</p> <p>Workaround: unknown</p>
CSCdy51893	<p>Symptom: Class queues do not get programmed correctly.CBWFQ may not work correctly.</p> <p>Condition: Modify the policy map values associated to a particular subif.</p> <p>Workaround: shut/no shut on sub-interface</p>

Table 8 Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)

CSCdy53728	<p>Symptom: LLQ when defined with class queues does not achieve full SCR and also improper traffic distribution between the queues.</p> <p>Conditions: With less traffic on class queues and more traffic on LLQ and with rates such that SCR is congested, its observed that the bandwidth is not being shared accurately among the queues as per the ratios assigned to them. LLQ sometimes get less traffic or more traffic and its not deterministic.</p> <p>Workaround: Enable MaxUtilization and OverSubscription on LLQ and Class Queues.</p>
CSCdy55202	<p>Symptom: sh pol int shows zero bandwidth for all the classes.</p> <p>Conditions: When ever class default is configured with some bandwidth and policy is removed and reattached, all other class bandwidth becomes zero. This happens for mpls interfaces.</p> <p>Workaround: None.</p>
CSCdy56345	<p>Symptom: After removal and insertion of POS backcard, `Assertion Failure' tracebacks were observed.</p> <p>Conditions: When a policy map is configured on POS i/f and removal and insertion of the POS backcard is done, "Assertion Failure" tracebacks will be observed on the console window.</p> <p>Workaround: Remove the policy map and reapply it on the POS i/f.</p> <p>Further Problem Description: The issue is applicable to GigE backcard also.</p>

Table 8 *Resolved Caveats Prior to Cisco IOS Release 12.2.15T (continued)*

CSCdy71426	<p>Symptom: All the traffic on PXF stops.</p> <p>Conditions: If a service policy was attached to an interface, with no bandwidth leftover for default class. Any traffic pumped on this interface default class would cause PXF to freeze and hence stopping traffic.</p> <p>Workaround: Reconfigure the policymap to have some bandwidth for class-default and then reload the PXF microcode using 'micro reload pxf'.</p>
CSCdy75485	<p>Symptoms: All Layer 2 management packets are dropped, which causes all interfaces that depend upon keepalives to transition to the down state.</p> <p>Conditions: This symptom is observed on a Cisco 10000 series router in a configuration with a large numbers of interfaces.</p> <p>Workaround: Unknown</p>

Related Documentation

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM, a member of the Cisco Connection Family, is updated monthly. Therefore, it might be more current than printed documentation. To order additional copies of the Documentation CD-ROM, contact your local sales representative or call customer service. The CD-ROM package is available as a single package or as an annual subscription.

Note that for Release 3, the user documentation (command reference, overview, and installation and configuration guides) use the MGX Release 4 and Cisco IOS documents in addition to this release note.

Product documentation for MGX 8850 is available at the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/wanbu/8850px45/re14/index.htm>

<http://www.cisco.com/univercd/cc/td/doc/product/wanbu/mgx8250/12/rpmpr/reNotes/index.htm>

Obtaining Documentation

Cisco provides several ways to obtain documentation, technical assistance, and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

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

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order annual or quarterly subscriptions through the online Subscription Store:

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

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

<http://www.cisco.com/en/US/partner/ordering/index.shtml>

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

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

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

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, the Cisco Technical Assistance Center (TAC) provides 24-hour, award-winning technical support services, online and over the phone. Cisco.com features the Cisco TAC website as an online starting point for technical assistance.

Cisco TAC Website

The Cisco TAC website (<http://www.cisco.com/tac>) provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The Cisco TAC website is available 24 hours a day, 365 days a year.

Accessing all the tools on the Cisco TAC website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a login ID or password, register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Opening a TAC Case

The online TAC Case Open Tool (<http://www.cisco.com/tac/caseopen>) is the fastest way to open P3 and P4 cases. (Your network is minimally impaired or you require product information). After you describe your situation, the TAC Case Open Tool automatically recommends resources for an immediate solution. If your issue is not resolved using these recommendations, your case will be assigned to a Cisco TAC engineer.

For P1 or P2 cases (your production network is down or severely degraded) or if you do not have Internet access, contact Cisco TAC by telephone. Cisco TAC engineers are assigned immediately to P1 and P2 cases to help keep your business operations running smoothly.

To open a case by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete listing of Cisco TAC contacts, go to this URL:

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

TAC Case Priority Definitions

To ensure that all cases are reported in a standard format, Cisco has established case priority definitions.

Priority 1 (P1)—Your network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Priority 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Priority 3 (P3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Priority 4 (P4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The Cisco Product Catalog describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the Cisco Product Catalog at this URL:

http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- Cisco Press publishes a wide range of networking publications. Cisco suggests these titles for new and experienced users: Internetworking Terms and Acronyms Dictionary, Internetworking Technology Handbook, Internetworking Troubleshooting Guide, and the Internetworking Design Guide. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

<http://www.ciscopress.com>

- Packet magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access Packet magazine at this URL:
<http://www.cisco.com/go/packet>
- iQ Magazine is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:
<http://www.cisco.com/go/iqmagazine>
- Internet Protocol Journal is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:
http://www.cisco.com/en/US/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html
- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:
<http://www.cisco.com/en/US/learning/index.html>

This document is to be used in conjunction with the Cisco WAN Switching publications.

CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0709R)

Copyright © 2003 Cisco Systems, Inc.
All rights reserved.

