Release Notes for Cisco CPT–CTC and Documentation Release 9.3 and Cisco IOS Release 15.1(01)SA

Revised: January 17, 2012, OL-25133-01

Cisco Carrier Packet Transport Release Notes

The Cisco Carrier Packet Transport (CPT) is released in 9.3. Release notes contain the new features and enhancements for the Cisco Carrier Packet Transport (CPT) platform. For detailed information regarding features, capabilities, hardware, and software introduced with this release, see *Cisco CPT Configuration Guide–CTC and Documentation Release 9.3 and Cisco IOS Release 15.1(01)SA*. For the latest version of the Release Notes for Carrier Packet Transport Release 9.3, visit the following URL:

http://www.cisco.com/en/US/docs/optical/cpt/r9_3/release/notes/cpt93_relnotes.html

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

http://tools.cisco.com/Support/BugToolKit/action.do?hdnAction=searchBugs

Revision History

Date	Notes
September 2011	Added a Critical Bug Fixes section that describes the critical bug fixes made in 9.301 release.
December 2011	Added a Critical Bug Fixes section that describes the critical bug fixes made in 9.302 release.

Critical Bug Fixes

Critical Bug Fixes in Release 9.302

It is recommended to use Release 9.302 that is the latest release of Cisco CPT.

Release 9.302 addresses the following critical fixes over Release 9.301:

- Invalid configuration attempts caused database corruption or automatic reset of the fabric card. In case of database corruption, all the Layer 2 configurations were lost resulting in the node to come up with the default configuration.
- The bandwidth values used by the interfaces were set to high values occasionally by the CPT system. Due to this, the services were not provisioned on these interfaces through CTC because bandwidth availability check failed although the bandwidth was actually available.
- Attaching or removing the same table-map from more than one service instance present on the port-channel was erroneous. Traffic was not marked as expected and attaching the table map to the target was not possible.
- CTC accepted class-map and policy-map names that exceeded 40 characters and resulted in corrupt node database.
- The policy-map queue statistics displayed incorrect rate for channel groups.
- The SNMP ifIndex for the ports is not the same when queried from the TNC card and from the fabric card using the *community_string_configured_from_CTC@fabric_card_slot_number* format.
- When SNMP walk operation was performed on CISCO–EVC–MIB, the active fabric card crashed. Hence, CISCO–EVC–MIB support is added in this release.

- When you configure an EFP and send bi-directional traffic on an interface or channel group, the input/output packet rate and bits per second rate in the **show interface** command displayed erroneous information such as the following:
    - Did not match with the actual traffic rate.
    - Displayed the value as 0 occasionally.
    - Displayed incorrect values.

For information on other bugs fixed in this release, use the Bug ToolKit.

Critical Bug Fixes in Release 9.301

Release 9.301 addresses the following critical fixes over Release 9.3:

- The throughput issue for 4 byte VLAN overhead was identified. The user can now configure a Committed Information Rate (CIR) to be less than the expected CIR for point-to-multipoint traffic, MPLS Label Switch Router (LSR) traffic, and MPLS multisegment pseudowire traffic. The highest percentage error is 6% for 64 byte frames.
- The EMS secure mode was not working. The user can now start CTC even when the node is in EMS secure mode.
- The SFP+ ports on the fabric card share the 13 Gbps bandwidth. The first SFP+ port has more priority than the second SFP+ port. When both the SFP+ ports carry traffic, the first SFP+ port receives 10 Gbps bandwidth whereas the second SFP+ port receives 3 Gbps bandwidth. When the first SFP+ port receives less traffic, the second SFP+ port receives the remaining bandwidth. When only one SFP+ port carries traffic, then the port receives 10 Gbps bandwidth.
- Certain trap object IDs (OIDs) were missing in CERENT-45-MIB.MIB. The following trap OIDs are now present in CERENT-45-MIB.MIB:
    - 7525: PTS-FAIL
    - 7535: SAT-ACT-LINK-FAIL
    - 7540: SAT-COMM-FAIL
- In LSR or multisegment pseudowire traffic flows, the MPLS EXP remarking configuration is lost after the ingress card resets. Now, the table map configuration for MPLS EXP bits in LSR traffic is retained even after the card that carries traffic is reset.
- The remarking of MPLS EXP is not effective on the LSR or SPE node when the ingress and egress MPLS interfaces are present on different cards. The remarking is effective on the LSR or SPE node only when the ingress and egress MPLS interfaces are present on the same card. This issue is not present on the LER node.

For information on other bugs fixed in this release, use the Bug ToolKit.

Software and Hardware Requirements

Before you begin to install CPT, you must check if your system meets the minimum software and hardware requirements. This section describes the software and hardware requirements for CPT.

- Hardware—IBM-compatible PC with a Pentium IV or faster processor, CD-ROM drive, a minimum of 1 GB RAM, 20 GB hard disk with 250 MB of available hard drive space.
- Operating System:
    - Windows 2000 Professional, Windows XP Professional, Windows Vista, or Windows 7, Windows Server 2003 or 2008.

- UNIX workstation with Solaris Version 9 or 10 on an UltraSPARC-III or faster processor, with a minimum of 1 GB RAM and a minimum of 250 MB of available hard drive space.
- Apple Mac OS X.

  (Use the latest patch/Service Pack released by the OS vendor. Check with the vendor for the latest patch/Service Pack.)

- Java Runtime Environment—Java Runtime Environment Version 1.6.
- Browser for PC—Internet Explorer 6.x, 7.x, 8.x. For UNIX Workstation—Mozilla 1.7. For MacOS-X PC—Safari

To install or upgrade CPT, refer to the guides listed in the Related Documentation section.

New Features and Functionality

This section highlights new features and functionality for Release 9.3. For detailed documentation of each of these features, see the *Cisco CPT Configuration Guide–CTC and Documentation Release 9.3 and Cisco IOS Release 15.1(01)SA*.

Hardware

The CPT Release 9.3 supports the following hardware:

Fabric Card

The fabric card is a single–slot card with two 10 Gigabit Ethernet SFP+ ports and two 10 Gigabit Ethernet XFP ports. The SFP+ ports on the fabric card serve as normal ports or InterConnect (IC) ports. When the SFP+ ports are used as IC ports, these ports are used to connect with the SFP+ ports on the CPT 50 panel. The XFP ports on the fabric card support the Optical Transport Network (OTN) protocol. The fabric card, which runs the route processor version of the Cisco IOS software, provides high availability and high switching capacity.

Line Card

The line card has four 10 Gigabit Ethernet SFP+ ports. The SFP+ ports on the line card serve as normal ports or IC ports. The line card runs the line card version of the Cisco IOS software. The line card expands the I/O capacity of CPT 200 and CPT 600 chassis by interconnecting with other line and fabric cards.

CPT 50 Panel

The CPT 50 panel enables the number of ports to be scaled on the CPT system. The CPT 50 panel has four 10 Gigabit Ethernet SFP+ ports and 44 Gigabit Ethernet SFP ports. The four 10 Gigabit Ethernet SFP+ ports can be used to connect with the fabric and line cards.

The CPT 50 panel is not placed in the CPT 200 or CPT 600 shelf. The CPT 50 panel runs the line card version of the Cisco IOS software.

The CPT 50 panel has redundant DC feeds. The CPT 50 panel DC power supply can handle 48 V and 24 V. The 48 V power supply has both ANSI and ETSI versions. The CPT 50 panel also supports AC power supply.

Software

The following CPT software features are added in Release 9.3. You can configure the software features either through CTC or Cisco IOS commands.

Ethernet Virtual Circuit

The Ethernet Virtual Circuit (EVC) represents a logical relationship between Ethernet User–Network interfaces (UNI) in a provider–based Ethernet service. The EVC represents the service offered and is carried through the provider network. Each EVC is configured by its unique name across the provider network.

In the CPT system, the EVC represents a Carrier Ethernet service and is an entity that provides an end–to–end connection between two or more endpoints.

The traffic for the service needs to pass through several switches in the provider network to connect sites across the provider network. The instance of a specific EVC service on the physical interface of each network device through which the EVC passes through is called an Ethernet Flow Point (EFP). An EFP is a logical demarcation point of an EVC on an interface. An EFP can be associated with a bridge domain.

The CPT system supports the following types of EVCs:

- Ethernet Private Line
- Ethernet Virtual Private Line
- Ethernet Private LAN
- Ethernet Virtual Private LAN
- Rooted Multipoint EVC and Split Horizon

Multiprotocol Label Switching

Multiprotocol Label Switching (MPLS) is the technology that scales IP networks for service providers. It provides mechanisms for IP quality-of-service (QoS) and IP traffic engineering. MPLS is an industry standard that uses label switching as a method to forward IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets. The forwarding of MPLS packets is based on preestablished IP routing information.

MPLS enables service providers to offer additional services to their enterprise customers, including VPNs, improved traffic engineering, QoS, Layer 2 tunneling, and multiprotocol support.

There are two ways to set up an MPLS infrastructure: Label Distribution Protocol (LDP) and Multiprotocol Label Switching – Traffic Engineering (MPLS–TE). LDP differs from MPLS–TE in terms of the protocol used to distribute the labels along the path. LDP uses the Label Distribution Protocol whereas MPLS–TE uses the Resource Reservation Protocol – Traffic Engineering (RSVP–TE) protocol to distribute the labels. However, both LDP and RSVP–TE uses Open Shortest Path First (OSPF) for the routing protocol.

The CPT system supports OSPF and OSPF-TE in this release.

Multiprotocol Label Switching – Transport Profile

Multiprotocol Label Switching Transport Profile (MPLS–TP) is a carrier–grade packet transport technology that enables service providers to move from Synchronous Optical Networking (SONET) and Synchronous Digital Hierarchy (SDH) time–division multiplexing (TDM) to packet switching. MPLS–TP enables MPLS to be deployed in a transport network to support packet transport services with a similar degree of predictability to that found in existing transport networks.

The key features of MPLS–TP are as follows:

- Connection–oriented.
- Carries Layer 3 and Layer 2 services.
- Runs over IEEE Ethernet PHYs, OTN, WDM and so on.
- Static and bidirectional label-switched path (LSP) provisioning.

- Operations, Administration, and Maintenance (OAM) functions similar to those available in traditional optical transport networks such as SONET or SDH are provided. These OAM functions belong to the MPLS–TP data plane and are independent from the control plane.
- Fault propagation through Bidirectional Fault Detection (BFD), Link Down Indication (LDI), and Lockout Request (LKR) messages.
- 1:1 revertive path protection.
- IP–less provisioning of tunnels.
- Network provisioning through CTC.
- Traffic switchover time from working LSP to protect LSP and vice versa is up to 50 milliseconds.

Pseudowire

In this release, CPT supports only the forwarding of the Ethernet frames coming from the customer networks under Any Transport over MPLS (AToM). The technique used to transport such a frame is called pseudowire, that is, the emulation of a native service over the MPLS network.

A pseudowire is a tunnel established between two provider edge (PE) routers across the core carrying the Layer 2 payload encapsulated as MPLS data. This helps the carriers migrate from Layer 2 networks, such as Ethernet over MPLS to an MPLS core.

Backup pseudowires can be set. The router can be configured to send the pseudowire status to a peer router, even when the attachment circuit is down. The static or dynamically configured set of two or more pseudowire segments can be defined that behave and function as a single point-to-point pseudowire. The BFD control channel over Virtual Circuit Connection Verification (VCCV) feature provides OAM functions for MPLS pseudowires.

Static and dynamic pseudowires can be created in this release. The static pseudowire can carry traffic over MPLS-TE tunnels, MPLS-TP tunnels, and LDP. The dynamic pseudowire can carry traffic over LDP and MPLS-TE tunnels.

Licensing

A license is a permit for a software feature to be functional or enabled on a device. The "pay as you grow" model enables the hardware and software capacity to be upgraded by using a license key. CTC helps in deploying licenses to the Cisco CPT devices in the network, discovering the devices, and managing and viewing the inventory of licenses and devices. A return merchandise authorization (RMA) process is not required to add new hardware. The license can be purchased, electronically delivered, and used to enable the increased port capacity.

New or upgraded Cisco devices must be registered and must have a product authorization key (PAK) to obtain licenses from Cisco.

High Availability

The CPT system supports Stateful switchover (SSO), Active-Active Data Plane (AADP), Cisco Nonstop Forwarding (NSF), and In-Service Software Upgrade (ISSU).

SSO ensures state synchronization and non-disruptive switchover from an active to a standby fabric card, thereby providing an increase in both system and network availability. In SSO, the standby fabric card is fully initialized and is ready to assume control from the active fabric card when the switchover occurs.

AADP refers to the load sharing between the two fabric cards. The redundant fabric cards run in an active-standby control model. However, both the fabric cards have ports that carry active traffic.

Cisco NSF works with the SSO feature to minimize the amount of time a network is unavailable following a switchover. The main objective of the Cisco NSF feature is to continue forwarding IP packets after the switchover of the active fabric card.

Software upgrade is an important consideration for high availability. CPT supports the ISSU process to perform planned software upgrades within the HA system. ISSU provides the ability to perform a stateful upgrade even when both the fabric cards are in different versions. ISSU is built over the SSO infrastructure.

Quality of Service

Quality of service (QoS) refers to the ability of a network to provide improved service to selected network traffic over various underlying technologies including Ethernet and 802.1 networks, and MPLS networks. In particular, QoS provides improved and more predictable network services by implementing the following services:

- Supporting guaranteed bandwidth.
- Improving loss characteristics.
- Avoiding and managing network congestion.
- Shaping network traffic.
- Setting traffic priorities across the network

Resilient Ethernet Protocol

The Resilient Ethernet Protocol (REP) is a protocol that provides an alternative to the Spanning Tree Protocol (STP) to support Layer 2 resiliency, and fast switchover with Ethernet networks. REP provides a way to control network loops, handle link failures, and improve convergence time.

REP performs the following tasks:

- Controls a group of ports connected in a segment.
- Ensures that the segment does not create any bridging loops.
- Handles single link failure within the segment.
- Improves convergence time.
- Supports VLAN load balancing at the service instance level.

Link Aggregation Group

The Link Aggregation Group (LAG) bundles individual Ethernet links into a single logical link that provides the aggregate bandwidth of up to eight physical links. When an EFP is configured on a LAG, the EFP is protected against link failures. When a link within a LAG fails, the traffic previously carried over the failed link switches to the remaining links within that LAG.

The Link Aggregation Control Protocol (LACP) is a control protocol over LAG to check for any LAG misconfigurations. LACP enables a single Layer 2 link to be formed automatically from two or more Ethernet links. This protocol ensures that both ends of the Ethernet link are functional and agree to be members of the aggregation group. LACP must be enabled at both ends of the link to be operational.

MAC Learning

The CPT system is a distributed system with fabric cards, line cards, and CPT 50 panels. The MAC addresses learned on one line card needs to be learned or distributed on the other line cards. The MAC Learning feature enables the distribution of the MAC addresses learned on one line card to the other line cards.

MAC learning is supported and enabled only for point–to–multipoint bridge domains.

Multicast VLAN Registration

Multicast VLAN Registration (MVR) is designed for applications using wide-scale deployment of multicast traffic across an Ethernet ring-based service-provider network (for example, the broadcast of multiple television channels over a service provider network). MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on the network-wide multicast bridge domain. It allows the single multicast bridge domain to be shared in the network while subscribers remain in separate bridge domains. MVR provides the ability to continuously send multicast streams in the multicast bridge domain and also to isolate the streams from the subscriber bridge domains for bandwidth and security reasons.

IGMP Snooping

As networks increase in size, multicast routing becomes critically important as a means to determine which segments require multicast traffic and which do not. IP multicasting allows IP traffic to be propagated from one source to a number of destinations, or from many sources to many destinations. Rather than sending one packet to each destination, one packet is sent to the multicast group identified by a single IP destination group address. Internet Group Management Protocol (IGMP) snooping restricts flooding of multicast traffic by sending multicast traffic only to the interfaces that are subscribed to a particular multicast group.

Using the Bug ToolKit to Search Bugs

In Cisco Carrier Packet Transport Release 9.3 and later releases, use the Bug ToolKit to view the list of outstanding and resolved bugs in a release.

This section explains how to use the Bug ToolKit to search for a specific bug or to search for all the bugs in a specific release.

Procedure

Step 1 Go to http://tools.cisco.com/Support/BugToolKit/action.do?hdnAction=searchBugs.
You will be prompted to log into Cisco.com. After successful login, the Bug Toolkit page opens.

Step 2 To search for a specific bug, enter the bug ID in the **Search for Bug ID** field and click **Go** in the **Search Bugs** tab.
To search for bugs in a specific release, enter the following search criteria:

- Select Product Category—Select **Optical Networking**.
- Select Products—Select **Cisco Carrier Packet Transport (CPT) System** from the list.
- Software Version—Select **9.30**, **9.301**, **9.302**, or **9.50** to view the list of outstanding and resolved bugs in Cisco CPT software.
- Search for Keyword(s)—Separate search phrases with boolean expressions (AND, NOT, OR) to search within the bug title and details.
- Advanced Options—You can either perform a search using the default search criteria or define custom criteria for an advanced search. To customize the advanced search, select **Use custom settings for severity, status, and others** and provide the following information:
  - Severity—Select the severity level from 1 to 6.
  - Status—Select **Open**, **Fixed**, or **Terminated**.

    Select **Open** to view all the open bugs. To filter the open bugs, uncheck the **Open** check box and select the appropriate sub-options that appear below the Open check box. The sub-options are New, Held, More, Open, Waiting, Assigned, Forwarded, Postponed, Submitted, and Information Required. For example, if you want to view only new bugs in Cisco CPT Release 9.3, only select **New**.

Select **Fixed** to view fixed bugs. To filter fixed bugs, uncheck the **Fixed** check box and select the appropriate sub-options that appear below the Fixed check box. The sub-options are Resolved or Verified.

Select **Terminated** to view terminated bugs. To filter terminated bugs, uncheck the **Terminated** check box and select the appropriate sub-options that appear below the Terminated check box. The sub-options are Closed, Junked, and Unreproducible. Select multiple options as required.

- Advanced—Check the **Show only bugs containing bug details** check box to view only those bugs that contain detailed information, such as symptoms and workarounds.
- Modified Date—Select this option if you want to filter bugs based on the date on which the bugs were last modified.
- Results Displayed Per Page—Select the appropriate option from the list to restrict the number of results that appear per page.

Step 3 Click **Search**. The Bug Toolkit displays the list of bugs based on the specified search criteria.

Export to Spreadsheet

The Bug ToolKit provides the following options to export bugs to a spreadsheet:

- Click **Export All to Spreadsheet** link in the Search Results page under the Search Bugs tab. Specify file name and folder name to save the spreadsheet. All the bugs retrieved by the search will be exported.
- Click **Export All to Spreadsheet** link in the My Notifications tab. Specify file name and folder name to save the spreadsheet. All the saved bugs in all the groups will be exported.

If you are unable to export the spreadsheet, log into the Technical Support Website at http://www.cisco.com/cisco/web/support/index.html for more information or call Cisco TAC (1-800-553-2447).

Related Documentation

Release-Specific Documents

Use the Cisco CPT Release Notes, Release 9.3 in conjunction with the following referenced Release 9.3 publication:

- *Release Notes for Cisco ONS 15454, ONS 15454 M2, and ONS 15454 M6 DWDM, Release 9.3*

Platform-Specific Documents

- *Cisco CPT Configuration Guide-CTC and Documentation Release 9.3 and Cisco IOS Release 15.1(01)SA*
- *Cisco CPT Command Reference Guide-CTC and Documentation Release 9.3 and Cisco IOS Release 15.1(01)SA*
- *Cisco CPT Licensing Configuration Guide*

Additional References

The following link provides additional information on CPT:

- http://www.cisco.com/go/cpt

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as an RSS feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service. Cisco currently supports RSS version 2.0.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: http://www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1110R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands