



Q & A

Cisco IOS MPLS Embedded Management

Q. What is Cisco IOS® MPLS Embedded Management?

A. Cisco IOS MPLS Embedded Management is a set of standards and value-added services that facilitate the deployment, operation, administration, and management of Multiprotocol Label Switching (MPLS)-based networks in line with the fault, configuration, accounting, performance, and security (FCAPS) model.

Cisco IOS MPLS Embedded Management is enabled by combining Cisco® MPLS Ping and Traceroute, Cisco Virtual Circuit Connectivity Verification (VCCV), Cisco MPLS Traffic Engineering AutoTunnel and AutoMesh, and Cisco IP SLAs.

The combination of these features provides the necessary tools to proactively manage an MPLS network combined with the ability to troubleshoot specific problems (reactive management) for network operators (refer to Table 1).

Table 1. Cisco IOS MPLS Embedded Management Components and FCAPS

Description	Specification
Fault management	<ul style="list-style-type: none">• Cisco MPLS Ping and MPLS Multipath Traceroute• Cisco VCCV• MIBs• Cisco IP SLAs• Cisco IP SLAs LSP Health Monitor
Configuration	<ul style="list-style-type: none">• Cisco MPLS Traffic Engineering AutoTunnel• Cisco AutoTunnel mesh groups• Cisco IP SLAs• Cisco IP SLAs LSP Health Monitor
Accounting	<ul style="list-style-type: none">• NetFlow• MIBs
Performance	<ul style="list-style-type: none">• Cisco IP SLAs• Cisco IP SLAs LSP Health Monitor• NetFlow• MIBs
Security	<ul style="list-style-type: none">• Resource Reservation Protocol (RSVP) message authentication• Label Distribution Protocol (LDP) message authentication• Message Digest Algorithm 5 (MD5) authentication for routing protocols:<ul style="list-style-type: none">– Border Gateway Protocol (BGP)– Open Shortest Path First (OSPF)• NetFlow Export Information

Q. Why is Cisco IOS MPLS Embedded Management important to the service provider?

A. As carriers and service providers worldwide converge services and disparate networks onto an MPLS-based infrastructure, MPLS operations, administration, and maintenance (OA&M) becomes pivotal for enabling them to provide service-level agreement (SLA) guarantees, service assurance, quality-of-service (QoS) assurance, and overall internetworking service management. Network operators need the ability to reliably conduct SLA testing, detect MPLS control- and user-plane defects, and check MPLS forwarding path integrity in real time. A service provider that is planning to offer managed services on an MPLS-based infrastructure must carefully consider MPLS OA&M to support premium SLAs.

The functions offered by these Cisco Systems® technologies are unique in the industry, and they allow service providers to easily deploy, operate, and monitor MPLS-enhanced services.

MPLS Ping and Traceroute

Q. Which MPLS ping and traceroute draft does Cisco support? Does Cisco support draft-ietf-mpls-lsp-ping-10?

A. The current Cisco implementation (available in Cisco IOS Software Release 12.0(32)SY) is based on draft 09 of MPLS ping and traceroute; earlier implementation is based on draft 03. For interoperability purposes, on the latest Cisco IOS Software release the user can choose, if necessary, the version to be used.

Q. Which Cisco IOS Software release provides support for MPLS ping and traceroute?

A. MPLS ping and traceroute are supported as of Cisco IOS Software Release 12.0(27)S on the Cisco 7200 Series, the Cisco 7500 Series, and the Cisco 12000 Series routers. Support for other platforms, such as the Cisco 7600 Series, will be available with Cisco IOS Software Release 12.2(RLS6)S.

Q. How do MPLS ping and traceroute work?

A. MPLS ping operation is based on the echo request and echo reply (similar to Internet Control Message Protocol [ICMP] ping).

A MPLS echo reply is sent in reply to a MPLS echo request. The MPLS Ping/Traceroute echo request and replies payload are all User Datagram Protocol (UDP) packets, which are forwarded by MPLS within the MPLS network.

The mechanism allows the use of a different return path, which can be specified by the node that sends the echo request packet. The echo reply can be forwarded as an IP packet or MPLS packet to the label switch router (LSR) that originates the MPLS echo request.

An MPLS echo request is a UDP packet that is sent to a target router using the appropriate label stack that is associated with the label-switched path (LSP) to be tested. The destination address of the MPLS echo request UDP packet is different from the address used to select the label stack. It is not used for forwarding; instead, the IP destination address is defined as a 127/8 address and used to:

- Force the packet to be consumed by the router where a LSP breakage occurs
- Force processing of the packet at the terminal point of the LSP if the LSP is intact
- Influence load balancing during forwarding when the transit routers use the destination address in the IP header for load balancing

LSP ping and traceroute provide diagnostics and troubleshooting capability for MPLS LSPs. These tools provide basic building blocks for the MPLS OA&M capabilities, enabling verification of the MPLS data-plane consistency.

LSP ping is a data-plane verification tool that verifies the LSP connectivity and the integrity of the MPLS network. Ping mode can test the integrity of connectivity using the verification on the Forward Equivalence Class (FEC) entity between the ping origin and the egress node for this FEC. This test is carried out by sending an MPLS echo request along the same data path as other packets belonging to this FEC. When the ping packet reaches the end of the path, it is sent to the control plane of the egress LSR, which then verifies that it is indeed an egress for the FEC. The MPLS echo request contains information about the FEC whose MPLS path is being verified.

Cisco Systems, Inc.

All contents are Copyright © 1992–2005 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

Currently, the following FECs are supported (in the MPLS ping and traceroute draft): LDP IPv4 prefix, LDP IPv6 prefix, RSVP IPv4 Session Query, RSVP IPv6 Session Query, VPN IPv4 prefix, and VPN IPv6 prefix.

LSP traceroute is a data-plane verification tool that traces LSPs in the MPLS networks. In the traceroute LSP verification, the packet is sent to the control plane of each transit LSR, which performs various checks, including one that determines if it is a transit LSR for this path. Furthermore, each transit LSR also returns extra information related to the FEC being tested (that is, the label bound to the FEC). This information helps in checking the control plane against the data plane (for example, in checking to see if the local forwarding information matches what the routing protocols determined as the path). Traceroute operation is performed by manipulating the time to live (TTL).

Q. What reply modes does Cisco support?

A. The MPLS ping and traceroute draft allows provision for the following reply modes:

- Do not reply
- Reply via IPv4 UDP packet
- Reply via IPv4 UDP packet with router alert
- Reply via control plane (only RSVP)

The *do not reply* mode is useful for the keepalive type of application running at the remote end, which triggers a state change if it does not receive a LSP ping packet within a predefined time.

Reply via UDP packet implies that an IPv4 UDP packet should be sent in reply to an MPLS echo request.

Reply via IPv4 UDP packet with router alert forces the packet to traverse back to the destination to be processed by the router processor at each intermediate hop.

Reply via control plane is specific to RSVP and requires extension to RSVP signaling.

Cisco supports the following reply modes:

- *Do not reply* is supported but is not used.
- *Reply via IPv4 UDP packet* is supported.
- *Reply via IPv4 UDP packet with router alert* is supported.
- *Reply via control plane* is supported (RSVP traffic engineering).

Q. What return codes are generated by Cisco routers? Is it possible to localize the errors?

A. Table 2 lists the possible return codes, indicating the possible errors (for example, FEC mismatch, etc.).

Table 2.

Value	Meaning
0	The error code is contained in the error code type length value (TLV).
1	A malformed echo request has been received.
2	One or more of the TLVs was not understood.
3	The replying router is an egress node for the FEC.
4	The replying router has no mapping for the FEC.
5	The replying router is not one of the downstream routers.
6	The replying router is one of the downstream routers, and its mapping for this FEC on the received interface is the given label.
7	The replying router is one of the downstream routers, but its mapping for this FEC is not the given label.

Error code TLVs are not defined yet, so a value of 0 is expected but not implied because it is conceivable that a transit router or destination router might be running a later implementation of the draft.

The originating router just reports that it received an error code TLV that is not understood.

Q. What target FEC stacks does Cisco support?

A. As of Cisco IOS Software Release 12.0(27)S, the following FECs are supported:

- LDP IPv4 prefix (echo request, echo reply, and traceroute)
- RSVP IPv4 Session Query (echo request, echo reply, and traceroute)
- VCCV Circuit ID (used by Any Transport over MPLS [AToM] VCCV) (echo request and echo reply)

Support for VPN IPv4 prefix (echo request, echo reply, and traceroute) will be available soon.

Support for other FEC stacks (LDP IPv6, etc.) is under study.

Q. Is there any limitation regarding label stack depth?

A. MPLS ping and traceroute do not introduce any extra MPLS label – they just use the MPLS forwarding infrastructure to convey the packet to the ingress LSR. (Refer to the question “How do MPLS ping and traceroute work?”)

Q. Is downstream mapping supported? Are there limitations?

A. Downstream mapping TLV is used during MPLS traceroute to keep track of the label used going from the headend to the tail end. It also allows MPLS traceroute to detect alternative paths.

Q. Does MPLS Ping – Traceroute allow the verification of all possible paths (equal-cost multipath [ECMP]) between the source and the destination?

A. The MPLS Ping – Traceroute draft suggests a nonoptimized way of discovering ECMPs within an MPLS network from a source to a destination.

MPLS LSP ping and traceroute as delivered in Cisco IOS Software Release 12.0(27)S provides the ability to ping and trace an LSP but does not guarantee the ability to trace all possible paths from the source to the destination. It relies on the ability to define various 127 addresses and assumes a large enough address space configured such that, by hashing those addresses, all load-balanced paths are tested.

If the configured address space is large enough, it always provides a result. But depending on the hash, it may require a lot of time before all possible paths are found.

If there were a way to trace the entire ECMP tree from the source to the destination, it would be simpler to point each path along the tree to check for liveliness.

Cisco implements a more optimized version of that algorithm, called MPLS Multipath Traceroute (ECMP Tree Trace). Cisco MPLS Multipath Traceroute discovers possible ECMPs within an MPLS network between a source and a destination and, therefore, allows testing of all those paths. This feature is available in Cisco IOS Software releases 12.0(32)SY and 12.2SBD (for the following platforms: Cisco 10000, 7500, 7304, 7301, and 7200).


Cisco MPLS VCCV

Q. What is Cisco MPLS VCCV?

A. Cisco MPLS VCCV enhances the monitoring and troubleshooting of Layer 2 services across an MPLS network. Figure 1 shows the pseudowire infrastructure.

Figure 1

Pseudowire Infrastructure



Q. Is Cisco MPLS VCCV available?

A. Cisco MPLS VCCV is available as of Cisco IOS Software Release 12.0(27)S.

Q. How does Cisco MPLS VCCV work?

A. Cisco MPLS VCCV creates a control channel between the two termination-point pseudowire provider edges (PEs) to uniquely identify the connectivity verification packets from the regular Layer 2 payloads.

Ideally such a control channel would be completely in-band.

When a control word is present on a virtual circuit (as defined in draft-ietf-pwe3-cw-xx.txt), it is possible to indicate the control channel by setting a bit in the control header. However, to ensure smooth interoperability between the different devices participating in the pseudowire service, an MPLS router alert label is used to indicate the control channel.

Q. Is it mandatory to support both in-band and router alert control channels?

A. To accommodate the huge number of already deployed devices, both options are available. When Cisco MPLS VCCV is used (between two provider edges), the first step is to agree on capability (in-band or router alert). Afterward, any connectivity verification packets use the agreed-on mode.

Note: *Control word* is the preferred mode because it does not require processing of the packet up to the route processor whenever connectivity verification is used. The use of the *router alert* option implies that each packet is processed by the route processor.

Q. Does Cisco support the in-band or the router alert mode?

A. Cisco supports both modes, but depending on hardware, both modes might not be available.

On the Cisco 12000 Series:

Control word option can be used when the imposition card is Engine3.

Router alert option is used on all other line cards and when the card used for egress is not an engine 3 card.

On the Cisco 7200 and 7500 series routers, *control word* is the preferred mode.

Cisco MPLS Traffic Engineering AutoTunnel Primary and Backup Mesh Groups

Q. What are Cisco MPLS Traffic Engineering AutoTunnel Primary and Backup?

A. Cisco MPLS Traffic Engineering AutoTunnel provides the ability to set up traffic engineering tunnels automatically. This feature provides support for both primary and backup tunnels. Backup tunnels are deployed to protect the primary ones. Cisco MPLS Traffic Engineering AutoTunnel automates the configuration tasks in the deployment of Cisco MPLS Traffic Engineering Fast Reroute (FRR).

Q. Are Cisco MPLS Traffic Engineering AutoTunnel primary and backup mesh groups available?

A. Yes, as of Cisco IOS Software Release 12.0.(27)S.

Cisco MPLS Traffic Engineering AutoTunnel Primary

Cisco MPLS Traffic Engineering AutoTunnel Primary is a one-hop primary tunnel that, when used in conjunction with Cisco MPLS Traffic Engineering FRR protection, protects any traffic steered through the primary “one-hop tunnel.” (Basically, any traffic, including IP traffic going through the physical link, is protected by Cisco MPLS Traffic Engineering FRR.)

Cisco MPLS Traffic Engineering AutoTunnel Backup

Cisco MPLS Traffic Engineering AutoTunnel Backup provides the capability to automatically build MPLS traffic engineering backup tunnels for the primary traffic engineering tunnel. These backup tunnels are set up mainly using *next hop* or *next next hop* protection, whenever available. A manually configured backup tunnel is preferred and thus provides “tweaking capabilities” for autotunnel features.

Q. How do I enable Cisco MPLS Traffic Engineering AutoTunnel Primary?

A. You can enable Cisco MPLS Traffic Engineering on any interface on which you want MPLS traffic engineering with the command: `mpls traffic-eng tunnel`.

To globally enable Cisco MPLS Traffic Engineering AutoTunnel, use the following command: `mpls traffic-eng auto-tunnel primary onehop`.

Q. How do I enable a backup tunnel for the primary autotunnel?

A. The one-hop traffic engineering tunnel can use either any manually configured backup tunnel(s) or backup tunnel(s) automatically created with the command `mpls traffic-eng auto-tunnel backup`.


A manually configured backup tunnel is preferred because it provides tweaking capabilities for autotunnel features.

Q. What is the result of the command `mpls traffic-eng auto-tunnel backup`?

A. This command automatically sets up a backup tunnel for the MPLS traffic engineering one-hop autotunnel. These backup tunnels are set up mainly using *next hop* or *next next hop* protection, whenever available, as shown in Figure 2.

Figure 2

Cisco MPLS AutoTunnel Primary and Backup



Q. What is the Cisco MPLS Traffic Engineering AutoTunnel Mesh Group?

A. The Cisco MPLS Traffic Engineering AutoTunnel Mesh Group increases the amount of bandwidth available over the same MPLS infrastructure and automates the configuration tasks in deployment of full-mesh MPLS traffic engineering tunnels. A full mesh of similar (sharing the same attributes) MPLS traffic engineering tunnels is automatically built between the router members of a specific mesh group.

Such a tool is needed typically when either transitioning an MPLS network to a fully meshed MPLS traffic engineering group (requires heavy configuration), or adding a new router in a fully meshed MPLS traffic engineering core where traffic engineering tunnels to (respectively from) all existing routers from (respectively to) the new one are needed.

Q. How do I enable a Cisco MPLS Traffic Engineering AutoTunnel Mesh Group?

A. All the routers within the same mesh group (for example, voice mesh group or data mesh group) can share the same MPLS traffic engineering characteristics. MPLS traffic engineering characteristics are defined with the use of *Auto-Template*. The set of router members of a given mesh group is defined by using a standard IP access list, which lists the IP addresses of those routers. As of Cisco release 12.0(29)S, the MPLS Traffic Engineering – AutoTunnel Mesh Groups feature uses an address-based discovery to identify mesh group members. This release introduces an Open Shortest Path First (OSPF)-based discovery for identifying mesh group members.

For OSPF to advertise or flood mesh group information, you need to configure a mesh group in OSPF and add that mesh group to an autotemplate interface. When the configuration is complete, OSPF advertises the mesh group IDs to all LSRs. MPLS TE LSPs automatically connect the edge LSRs in each mesh group.

The following are the steps to enable the Cisco MPLS Traffic Engineering AutoTunnel Mesh Group:

- Enable traffic engineering on all routers (members of the mesh group).

- Enable Cisco AutoTunnel mesh groups (global level):

```
router(config)# mpls traffic-eng auto-tunnel mesh.
```
- Configure ACCESS-LIST (standard IP access list), which defines the set of possible tunnel destinations

```
Router(config)# ip access-list standard 1
Router(config-std-nacl)# permit 10.0.0.0 0.255.255.255
```

or enable IGP flooding for AutoTunnel mesh group via the command “mpls traffic-eng mesh-group”.

```
Router(config)# router ospf 100
Router(config-router)# mpls traffic-eng mesh-group 10 loopback 0 area 100
```
- Configure Auto-Template:
 - Configuration when mesh group membership is defined through ACCESS-LIST:

```
router(config)#interface Auto-Template 1
router(config-if)#ip unnumbered Loopback0
router(config-if)#tunnel mode mpls traffic-eng
router(config-if)#tunnel mpls traffic-eng autoroute announce
router(config-if)#tunnel mpls traffic-eng priority 1 1
router(config-if)#tunnel mpls traffic-eng auto-bandwidth
router(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
router(config-if)#tunnel destination access-list 1
```
 - Configuration when mesh group membership is flooded via IGP:

```
router(config)#interface Auto-Template 1
router(config-if)#ip unnumbered Loopback0
router(config-if)#tunnel mode mpls traffic-eng
router(config-if)#tunnel mpls traffic-eng autoroute announce
router(config-if)#tunnel mpls traffic-eng priority 1 1
router(config-if)#tunnel mpls traffic-eng auto-bandwidth
router(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
router(config-if)#tunnel destination mesh-group 10
```
 - Traffic engineering LSP is automatically set up using the locally configured templates.

Cisco IOS IP SLAs

Q. What are Cisco IOS IP SLAs?

A. Embedded within Cisco IOS Software, Cisco IOS IP SLAs provide a scalable, cost-effective solution for performance measurement. Cisco IOS IP SLAs active monitoring can continuously measure the network performance, providing an ongoing baseline to indicate how the network is performing. Cisco IOS IP SLAs collect network performance information in real time: response time, one-way latency, one-way jitter, one-way packet loss, voice-quality measurement, as well as other network statistics. Cisco IOS IP SLAs provide unidirectional and bidirectional measurements and support measurements per class of service. Also available are proactive notification and threshold violation monitoring for jitter, packet loss, latency, and connectivity. All IP SLAs performance statistics are available in the Simple Network Management Protocol (SNMP) MIBs.

Q. Are Cisco IOS IP SLAs measurements Virtual Route Forwarding (VRF)-aware?

A. Yes, Cisco IP SLAs measurements can be sent to a destination that is found in a VRF routing table. This capability is useful for provider edge or point of presence (POP)-to-customer edge measurements and provider edge-to-remote customer edge measurements within a Layer 3 MPLS VPN. VRF-aware Cisco IP SLAs are also used on routers that use multi-VRF customer edge capability from Cisco.

Q. Do Cisco IOS IP SLAs support LSP ping?

A. A new feature from Cisco called IP SLAs LSP Health Monitor supports LSP and traceroute ping as a mechanism for MPLS network connectivity and performance measurements.

Q. What is Cisco IP SLAs LSP Health Monitor?

A. Cisco IP SLAs are widely used when network performance measurement and SLA monitoring data such as jitter statistics, packet loss, and round-trip time (RTT) are required within an IP-based network. In addition, Cisco IP SLAs provide unique tools that allow monitoring of MPLS Layer 3-based VPNs. Cisco IP SLAs LSP Health Monitor enhances the already feature-rich IP SLAs capabilities by simplifying the deployment and configuration of the IP SLAs measurements whenever performance measurement and SLA monitoring for an MPLS Layer 3 VPN infrastructure are required. The extended features include automatic generation of measurement to measure connectivity and latency between MPLS provider edge routers, proactive monitoring of the MPLS network, and automatic optimization of measurement scheduling, giving a better scan coverage time.

Q. How does Cisco IP SLAs LSP Health Monitor help monitor the MPLS core?

A. The feature uses BGP endpoint discovery in a Layer 3 MPLS VPN network to understand the provider edge endpoints in the network. When discovered, measurements are set up from the source provider edge to other destination provider edge devices. The health monitor is scheduled over a time period to measure from the source to all destinations with connectivity and latency measurements. This solution is scalable because each provider edge-to-provider edge measurement is scheduled sequentially at a rate specified by the user. If fault is found by the LSP ping or traceroute infrastructure for any individual path between provider edge devices, then the measurement frequency is increased at a higher rate to isolate the connectivity problem. If a problem is identified, an SNMP trap is sent to a fault system, allowing the operator to pinpoint and diagnose the problem. The unique capability of this feature will allow many service providers to reduce troubleshooting time and isolate problems in the MPLS core network.

Q. Does the Cisco IP SLAs LSP Health Monitor work for multiple equal-cost paths?

A. The feature supports multiple equal-cost paths and uses this LSP discovery mechanism to poll all possible customer paths between the Layer 3 VPN provider edges. Contact your Cisco account team about release schedules for platform and Cisco IOS Software trains involved.

MPLS-Aware NetFlow

Q. What is NetFlow?

A. NetFlow capitalizes on the flow nature of traffic in the network to provide detailed IP information with minimal impact on router and switch performance. NetFlow monitors each IP packet to characterize IP flows in the router and switch and exports the flows in UDP format to a NetFlow collector. NetFlow is unlike other performance features in that it pushes or exports data to collectors or network management stations, in contrast to traditional polling. The NetFlow collector can correlate, aggregate, and report on the data received from the network. NetFlow data can be used for a variety of purposes, including security monitoring, traffic analysis and planning, usage-based billing, data warehousing and mining for marketing purposes, etc. NetFlow gives Cisco customers a scalable method to understand what traffic is being sent in the network, where it is sent when, and how much traffic there is. This information has proven invaluable for network professionals in recent years. Most Cisco platforms support NetFlow in hardware today, allowing extremely scalable operation. NetFlow is also used extensively through the command-line interface (CLI) for troubleshooting a network.

Q. What is NetFlow Version 9?

A. NetFlow Version 9 is a new, flexible and extensible export format for NetFlow flow data. Historically NetFlow has used a fixed-field export format, but this format is not extensible. NetFlow Version 9 uses a template to describe the data in the export packet and allows flexibility in the content of flow export records. NetFlow Version 9 extends NetFlow to support MPLS label information, egress NetFlow, Layer 2 and security exports, IPv6 information, BGP next hop, and multicast information. Version 9 has recently been proposed as an IETF standard.

Q. What is MPLS-aware NetFlow?

A. MPLS-aware NetFlow uses NetFlow Version 9 to export information needed for MPLS capacity planning and IP accounting. It can be used to determine and account for traffic to a particular destination in the MPLS cloud. It allows export of information about the complete IP flow, along with information pertaining to up to three labels (label destination prefix information, including MPLS EXP). The user can account for MPLS traffic that contains IP or non-IP packets and can include the MPLS header as part of the accounting information. MPLS-aware NetFlow is available on the Cisco 12000 Series in Cisco IOS Software Release 12.0(24)S and other platforms in Cisco IOS Software Release 12.0(26)S.

Q. What is MPLS Label Information Export?

A. NetFlow Version 9 format can be used along with MPLS Label Information Export to understand routing of flow in the MPLS network. MPLS Label Information Export effectively exports the Label Forwarding Information Base (LFIB), which includes labels in the device and FEC or next hop for the labels. This capability can be used to monitor MPLS next-hop information and can be used along with the accounting or packets and byte information available in MPLS-Aware NetFlow to understand a provider edge-to-provider edge traffic matrix for the MPLS network.

Q. What is MPLS Egress NetFlow?

A. MPLS Egress NetFlow allows users to account for egress traffic exiting an MPLS network on a provider edge router (that is, from the provider edge toward a customer edge in a Layer 3 MPLS VPN). MPLS Egress NetFlow can be combined with IP NetFlow used for Layer 3 MPLS VPN accounting.

For More Information

For more information, refer to the following:

Cisco IOS MPLS: <http://www.cisco.com/go/mpls>

Cisco IOS Software Release 12.0S documentation: <http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1829/index.html>

Advanced performance measurement with Cisco IOS IP SLAs: [IOS documentation for IP SLAs LSP Health Monitor](http://www.cisco.com/go/ipsla) or <http://www.cisco.com/go/ipsla>

VCCV draft, "Pseudowire Virtual Circuit Connectivity Verification (VCCV)," IETF Webpage: <http://www.ietf.org/internet-drafts/draft-ietf-pwe3-vcv-07.txt>

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Website at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR
Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia
Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2005 Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks or 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. (0502R)
KW/LW9998 12/05