

Cisco Any Transport over MPLS

Introduction

Any Transport over MPLS (AToM) is the Cisco solution for transporting Layer 2 traffic over an IP or MPLS backbone. AToM extends the usability of the IP/MPLS backbone by enabling it to offer both Layer 2 and Layer 3 services. The AToM product set accommodates many types of Layer 2 frames, including Ethernet, Frame Relay, ATM, PPP and HDLC, across various Cisco router platforms, such as the Cisco 12000 Series routers and the Cisco 7600, 7500, and 7200 routers.

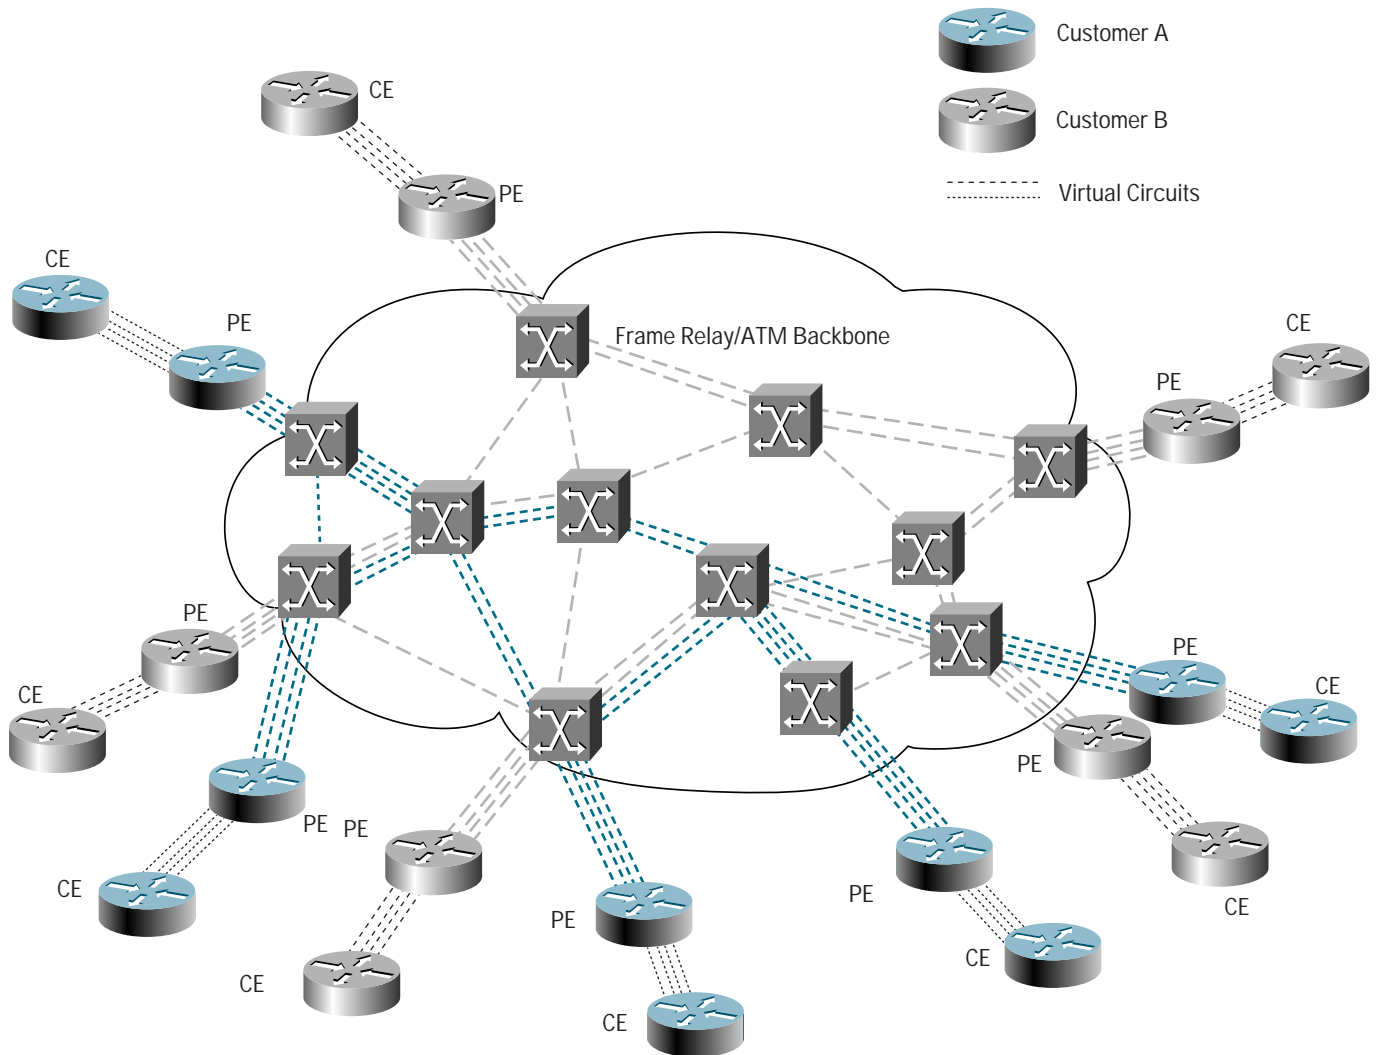
Background

Many service providers currently offer Layer 2 transport services to their customers. These services are offered over a circuit based infrastructure to build Layer 2 VPNs. Figure shows a sample topology where the service provider network provides Layer 2 VPN services to customers.

Initially, VPNs were built using leased lines. Later, service providers offered Layer 2 VPNs based on point-to-point data link layer connectivity, using ATM or Frame Relay virtual circuits. Customers built their own Layer 3 networks to accommodate IP traffic. As a result, separate networks exist for Layer 2 and Layer 3 traffic. But maintaining separate networks for Layer 2 VPNs and Internet traffic is difficult and costly. So service providers want a single IP-based network to provide both Layer 2 and Layer 3 services.



Figure 1
Service Provider Offering VPN Services based on Layer 2 Connectivity



How MPLS Enables Service Providers to Expand Their Networks

One way that service providers can create a better network is to incorporate Layer 2 and Layer 3 services over a common infrastructure, such as an MPLS backbone. MPLS-enabled VPNs (Layer 2 or Layer 3) also offer performance, scalability, and new value-added services using other MPLS enhancements such as traffic engineering, fast reroute, and QoS. With Cisco's implementation of the MSF framework using separate control planes, MPLS can be added to switches and ports without any impact to existing Layer 2 services.

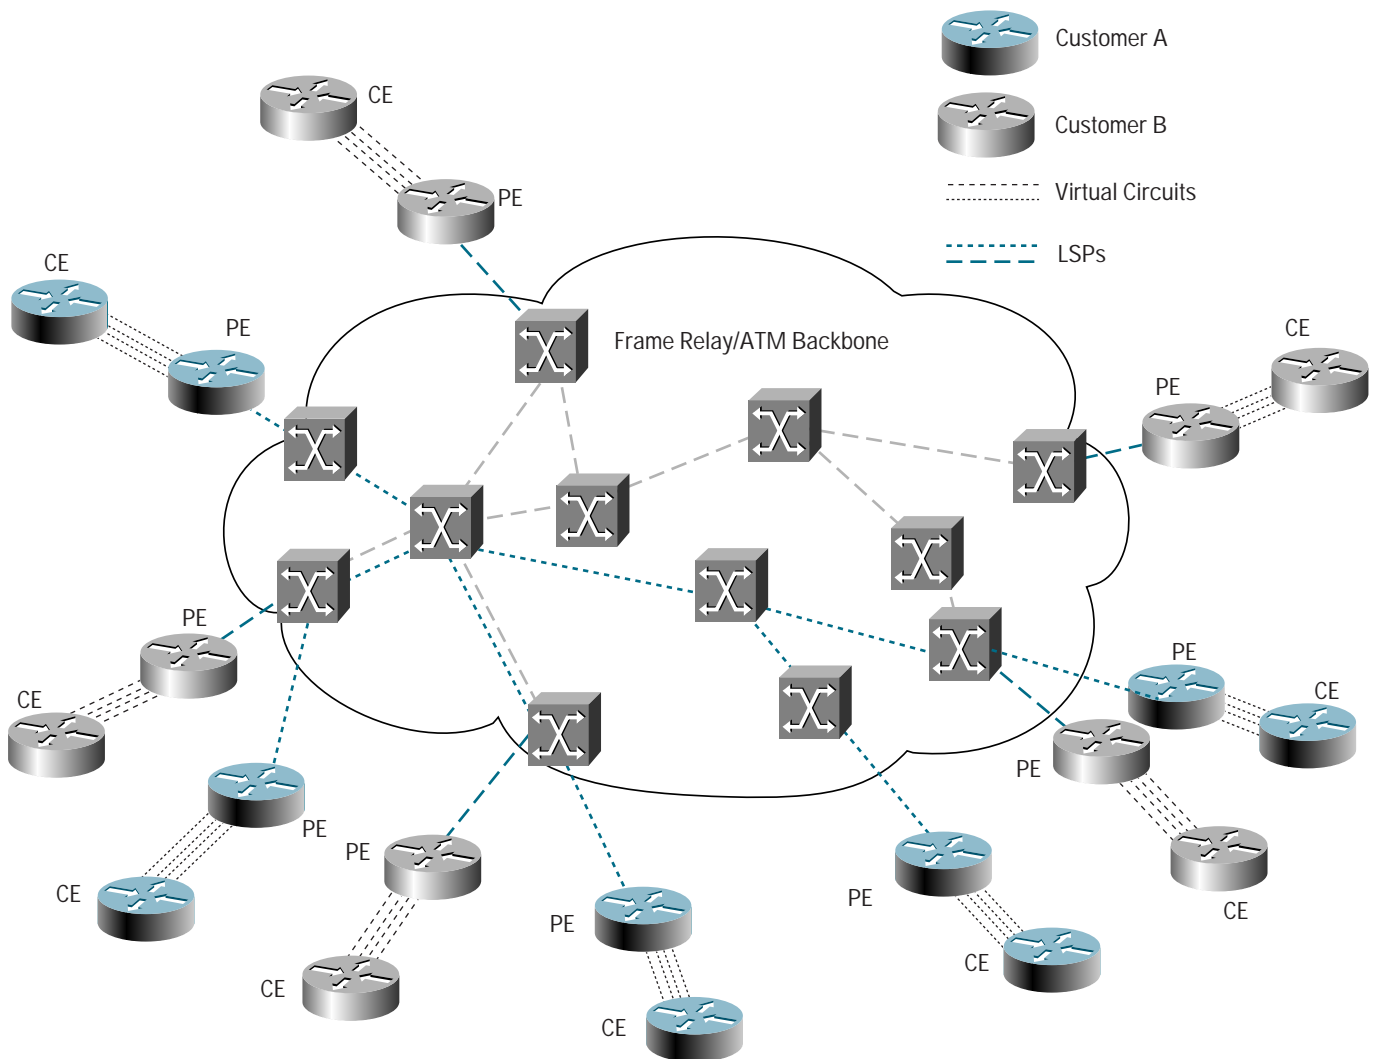


How MPLS AToM Improves Scalability

Many service providers have currently implemented connection-oriented ATM networks in the core. ATM provides QoS, bandwidth, and the ability to perform traffic engineering. However, an ATM network without MPLS does not scale well for VPNs because it relies on virtual circuit (VC) state information in the core. For every VC created as part of an L2 VPN, the core switches keep the complete state for the VC (as shown in Figure 1).

MPLS AToM links any pair of provider edge (PE) routers with a single Label-Switched Path (LSP) instead of a multitude of VCs. That single Label-Switched Path then serves as a carrier of many Emulated Virtual Circuits by means of label stacking, resulting in better scalability.

Figure 2
One Label-Switched Path (LSP) Can Replace Many Virtual Circuits





Simultaneous with this advantage is the fact that MPLS can be implemented on top of existing Layer 2 or Layer 3 networks. There is no need to build a separate network, as MPLS and ATM can run at the same time, even on the same physical port.

How AToM Helps Transition from Layer 2 VPNs to MPLS-based VPNs

As service providers migrate to MPLS-based networks, they need to maintain support for existing services, including the ability to transport Layer 2 traffic. AToM is the architectural framework for accomplishing that, regardless of whether the Layer 2 service is ATM, Frame Relay, Ethernet, PPP or HDLC. AToM encapsulates packets at the PE router, transports them over the backbone to a PE router on the other side of the cloud, unencapsulates them, and sends them to their destination.

The upgrade to AToM is transparent to customers, since this can be enabled on an existing MPLS network without any major disruptions. Unlike Layer 3 VPNs using MPLS, the service provider does not participate in the customer's Layer 3 routing. The service provider provides Layer 2 connectivity only.

Benefits of Using AToM to Build Layer 2 VPNs in MPLS Networks

MPLS AToM has the following benefits.

- A unified network for Layer 2 and Layer 3 VPNs means lower infrastructure and maintenance costs for the service provider. Using an IP or MPLS backbone enables the service provider to offer Layer 2 VPN services along with IP and Layer 3 VPNs in the same network. The service provider can support VPN traffic with other traffic (e.g. Internet) on a single infrastructure.
- Customer sites are independent of the service provider backbone, so problems on a customer network will not affect the this backbone. The service provider backbone is more reliable when it is separate from the customer network. Refer to Figure 2 to see how the service provider network and the customer network work independently:
 - The PE routers assign the initial label associated with a packet as it enters the MPLS core and strip the label off the packet when it exits the core.
 - The P routers perform the forwarding and switching functions (as detailed in Figure 3 and its accompanying description).
 - The CE routers connect via Attachment VCs to the PE routers. CE routers do not need any MPLS enhancements in order to use the MPLS VPN services.
- Because MPLS uses frames or cells, it can work in IP over ATM networks and IP over MPLS networks. MPLS is “agnostic” in its ability to accommodate protocols. Therefore, service providers can run IP over ATM networks and incrementally implement MPLS to their network. Adding MPLS to the network in phases can be a better alternative in many situations than having to make a complete conversion all at once.
- Cisco is committed to supporting the IETF draft, *Architecture for Layer 2 VPNs*. AToM is a critical component in the implementation of that draft. (See the section below, *Standards and Drafts Upon Which AToM Is Based*.) Cisco's IETF-based design has the following benefits:



- Signaling and encapsulation follow the IETF Martini draft, which enables the SP's customers to build their own VPNs without the SP having to take the time and trouble to do so itself. Furthermore, Cisco's use of directed LDP sessions allows peer AToM devices to communicate anywhere on the MPLS cloud, even across SP boundaries. By contrast, Layer 2 VPNs that rely on BGP for signalling (following IETF's Kompella draft) have trouble addressing routing policies across different SP domains. Nor can their customers set up their own VPNs.
- Scalability is enhanced for the same reasons. AToM can add new sites and CE-CE circuits without having to provision every PE on the backbone, and without having to interrupt network service to expand pre-provisioned PEs, as must the systems based on draft-Kompella architecture. With AToM only the PEs on which the service is offered need to be configured.

Core routers store a minimal amount of forwarding information, because in MPLS core routers do not store any VPN information. And each PE router must store only the forwarding information of the VPNs to which it connects, so the service provider core network is not affected by the number of VPNs it services.

A third scalability benefit of AToM is that it allows for unlimited VCs to be created, in line with the IETF draft *Transport of Layer 2 Frames Over MPLS*. Other Layer 2 solutions that do not adhere to this IETF draft are limited in scalability.

- Topology: AToM VPNs are configurable with equal ease on partial-mesh and full-mesh topologies. Draft-Kompella based architectures are not well suited for partial mesh topologies.
- Variety: The IETF architecture Cisco employs enables the SP's customer to specify different characteristics (such as bandwidth) on their different circuits. Such variety is not possible with draft-Kompella based architectures.
- Failure recovery is immediate in AToM. By contrast, other systems must wait until mandatory notification time-outs have occurred before their failed circuits can be re-established.
- Provisioning: Single end-point configuration is a proposed simplification in the *Architecture for Layer 2 VPNs*; it is not available in systems based on other architectures.
- Debugging is easier in AToM, because it does not have to deal with label blocks and active-vs-inactive VCs. When building large Layer 2 VPNs, it is possible to end up with fragmented label blocks in large networks. In such situations debugging gets much more complicated, because you have to use indexes and tables to figure out which label belongs to which VC and then diagnose the problem. That also goes against the philosophy of MPLS, where a label should be an internal value used for dynamic path buildup and switching.
- Upgrading to AToM is transparent to the customer. Because the service provider network is separate from the customer network, the service provider can upgrade to AToM without disruption of service to the customer. The customers assume that they are using a traditional Layer 2 backbone.
- AToM can be combined with QoS and Traffic Engineering to build new revenue generating services such as virtual leased lines that mimic existing Layer 2 services (ATM, Frame Relay) without compromising the scalability and flexibility of the MPLS networks on which they run. The white paper *Virtual Leased Line Services Using Cisco MPLS DiffServ-Aware Traffic Engineering* provides more information on providing leased line services that guarantee bandwidth. You can find that white paper at the following URL:

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/msdvl_wp.htm



Supported Transport Types

AToM enables the following types of Layer 2 frames and cells to be directed across an MPLS backbone:


- Ethernet, Ethernet VLAN
- Frame Relay
- ATM Adaption Layer Type-5 (AAL5)
- ATM cell relay
- Point to Point Protocol (PPP)
- High-Level Data Link Control (HDLC)

The first phase of AToM development in Cisco IOS Software supports like-to-like connectivity. This requires that the same transport type be at each end of the network. In future, AToM will be enhanced to provide inter-working functions that can connect disparate transport types at each end, such as Frame Relay at one end connecting to Ethernet VLAN at the other.

How Any Transport over MPLS Works

Figure shows how a packet travels from Site 1 to Site 2 in a VPN, using the IP/MPLS backbone.

Figure 3
How Packets Are Transported across an IP/MPLS Backbone





The following process shows a packet traveling from a CPE router on the left side of the network (Site 1) across the service provider network, to a CPE router on the right side (Site 2).

1. Packets flow from Site 1 to PE1 on the edge of the service provider network through a traditional Layer 2 virtual circuit, in this case a Frame Relay circuit.
2. In the service provider network, an operator configures a label switched path (LSP) from PE1 to PE2.
3. For AToM, the operator configures:
 - (a) At PE1, a cross-connect between Attachment VC 101 and Emulated VC1 (shown as VC Label 10 above), and the destination PE to be PE2
 - (b) At PE2, a cross-connect between Emulated VC1 and Attachment VC 201, and the source PE to be PE1

Note: No AToM configuration is required on the P routers.

4. At PE1 the following events then take place:
 - (a) An incoming packet on the PE router is stripped of the layer 2 header.
 - (b) A control word and Emulated VC label [10] are pushed onto the packet.
 - (c) An appropriate network facing interface is selected
 - (d) An LSP tunnel label [50] is pushed (for normal MPLS routing through the cloud)

The control word and the emulated virtual circuit label are pertinent only to the PE routers. The routers within the MPLS backbone (the P routers) do not use the control word or the virtual circuit label. Instead, the P routers use the LSP tunnel labels [50 & 90] to move the packet through the MPLS backbone. A P router does not distinguish AToM traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

5. The packet is sent through the service provider network to PE2.
6. The following events takes place on the egress router PE2.
 - (a) The Emulated VC label [10] is stripped.
 - (b) The control word is processed and stripped.
 - (c) The Layer 2 header is reconstructed for Attachment VC DLCI 201.
 - (d) The packet is sent out the appropriate customer-facing interface.

Note: No tunnel label is present in the network-facing side of the router because that label was popped by the penultimate P router.

7. PE2 connects to Site 2 through a traditional Layer 2 virtual circuit, in this case an Frame Relay circuit.

AToM and QoS Support

Background of QoS

Quality of Service (QoS) sorts and classifies packet requests into different traffic classes and allocates the proper resources to direct traffic based on various criteria, including: application type, user or application ID, source or destination IP address, and other variables.



The bits in the packet translate to the priority of the packet. For MPLS packets, the MPLS experimental bits, also known as the EXP bits, allow you to specify the QoS for an MPLS packet. For an IP packet, the IP precedence/DSCP bits allow you to specify the QoS for an IP packet.

If the service provider wants to set QoS for an MPLS packet to a value different from that of the IP precedence bit, the service provider can set the MPLS experimental field instead of overwriting the value in the customer's IP precedence field. The IP header remains available for the customer's use; the QoS for the IP packet is not changed as the packet travels through the MPLS network.

By preserving the IP packet's precedence/DSCP bits from being modified in the service provider network, the customer can differentiate traffic within their network, without buying multiple grades of service from the provider.

QoS with AToM

In the case of AToM, the same QoS classification and marking mechanisms are used. For example, based on type of service of the Attachment VC, the MPLS EXP field can be set to a higher priority that allows better delivery of Layer 2 frames across the MPLS network. Layer 2 QoS, such as the 802.1P field in the IP header, can be easily mapped to MPLS EXP to translate quality of service from Layer 2 to MPLS, thereby providing bandwidth, delay and jitter guarantees. In the case of Frame Relay and ATM, the EXP values can be set by reference to the DE (Discard Eligible) bit marking in the Frame header and to the CLP (Cell Loss priority) bit marking in the ATM cell header.

Future Development of AToM

In its future phases, AToM will provide the following capabilities:

Any-to-any connectivity. This will allow a customer to initiate a VPN at, for example, its frame relay network and have it extend to the desired destination over an ATM, Ethernet or HDLC/PPP network. Likewise, Layer 2 VPNs will be able to traverse the MPLS cloud to connect Ethernet networks with POS and with ATM.

AToM makes Any-to-any possible because PEs can negotiate differences they perceive in the virtual circuit type, and can then choose the most efficient (least overhead) method for inter-working frames and cells across the cloud.

Transparent LAN. This capability allows users in a multi-site enterprise to communicate with any host in their organization as if they were on an extended LAN interface. They do not need to know at which site the host is located. In Ethernet terms, the individual user stays with a single VLAN id, and can broadcast messages to many hosts throughout the system, regardless of their geographic location.

This convenience emerges from AToM's capability to provide the PE with MAC-address learning and with replication, which together make the MPLS cloud act like a distributed switch. As a result, the various sites perceive each other to be within a single broadcast domain.

Full support for IETF draft *Architecture for Layer 2 VPNs*. Utilizing the advantages of this architecture will make it possible for AToM to provide BGP-based autodiscovery, single end-point provisioning models, and zero-configuration Attachment VCs.

Cell packing. Incorporating cell packing into AToM will improve the efficiency of ATM cell relay over MPLS.



Conclusion

AToM is a powerful technology that allows service providers to offer Layer 2 connectivity over MPLS networks, thereby increasing revenue opportunities. It does so by eliminating many of the drawbacks of existing Layer 2 technologies. Combining AToM with QoS and Traffic Engineering allows service providers to build value added services such as virtual leased line, transparent LAN, and interworking on a single, common infrastructure.

Standards and Drafts Upon Which Cisco AToM Is Based

The AToM technology is based on the following IETF draft documents:

- *Transport of Layer 2 Frames Over MPLS*, accessible at:
<http://search.ietf.org/internet-drafts/draft-martini-l2circuit-trans-mpls-xx.txt>
- *Encapsulation Methods for Transport of Layer 2 Frames Over MPLS*, accessible at:
<http://search.ietf.org/internet-drafts/draft-martini-l2circuit-encap-mpls-xx.txt>
- *Architecture for Layer 2 VPNs*, accessible at:
<http://search.ietf.org/internet-drafts/draft-ietf-ppvpn-l2vpn-xx.txt>

Related Documentation

- *Technology Brief: Ethernet over MPLS on the Cisco 7600 Router*
<http://wwwin.cisco.com/ent/isbu/prod/pdfs/EoMPLS-techno-brief.pdf>
- Cisco IOS Software and Multiprotocol Label Switching
http://www.cisco.com/warp/public/cc/pd/iosw/ioft/iowft/prodlit/iosmp_ai.pdf

Glossary

The following terms are common to the AToM product set.

Any Transport over MPLS (AToM)—The name of the Cisco product set that transports many types of frames and cells over an MPLS backbone.

Label switched path (LSP)— A path from one MPLS router to another. Packets travel between MPLS routers through LSPs. LSPs go in only one direction. To provide two-way traffic, configure LSPs in each direction. A label-switched path can be established dynamically, based on normal routing mechanisms, or through configuration.

Label Imposition—The act of putting label(s) onto a packet. Label imposition is done by the PE router. In the case of Layer 2 transport over MPLS, this is the router that receives a packet and encapsulates it with MPLS encapsulation.

Label Disposition—The act of removing label(s) from a packet. This is done by the PE router. In the case of Layer 2 transport over MPLS, this is the router that receives an MPLS encapsulated packet, removes the last label, and transmits the Layer 2 PDU out the appropriate interface.

Label Distribution Protocol (LDP)—A standard protocol between MPLS-enabled routers to negotiate the labels (addresses) used to forward packets.

Virtual circuit (VC)—A logical connection created to ensure reliable communication between two network devices. A virtual circuit denotes the existence of a logical, bidirectional path from one device to another across a network.



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 Europe
11 Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

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.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 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 Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
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 © 1992–2002, Cisco Systems, Inc. All rights reserved. CCIP, the Cisco *Powered* Network mark, the Cisco Systems Verified logo, Cisco Unity, Fast Step, Follow Me Browsing, FormShare, Internet Quotient, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, GigaStack, IOS, IP/TV, LightStream, MGX, MICA, the Networkers logo, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site 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.
(0201R) 201777/ETMG 03/02