

Metro Ethernet Technologies on the Cisco Catalyst 6500 Series Switch

Introduction

In more than 25 years, *Ethernet* technology has experienced several changes, evolving from a shared-media half-duplex 10 Mbps LAN technology to a full-duplex switched 10/100/1000/10000 Mbps LAN technology, as the demand for higher networking speeds has grown. Ethernet switches, too, have evolved with advanced capabilities; notably, advanced *quality of service (QoS)* and *security*, which previously were available only for other more expensive media.

By offering cost effective bandwidth, Ethernet has made it possible for digital applications to flourish by taking advantage of a high-performance infrastructure and delivering unprecedented services. Because of its intrinsic qualities, Ethernet has been the desktop technology of choice since it emerged as the most cost effective delivery option for high-speed connectivity.

Recently, a similar trend has emerged in the *metropolitan-area networks (MANs)*, where dark fiber allows providers to offer cost-effective high-speed Ethernet connections to business and residential customers.

MAN architectures using Ethernet technology are based on two types of network design: a classic IP-centric Layer 3 design or a pure Layer 2 design. The pure Layer 2 design is popular because of its simplicity and its “transparency” to user traffic and protocols. This is referred to as *Transparent LAN Services (TLS)*.

In the evolution of Ethernet switching and routing in metropolitan environments, the contribution of Cisco Systems® is dependable and is marked by milestones,

such as the introduction of several important metro technologies. Cisco® has summarized these milestones in *Metro Solutions for Service Providers* guides (for more details, see [1]).

These guides provide methodologies for building an advanced metropolitan network that takes advantage of the latest Cisco networking products and technologies.

As a cornerstone of Cisco metro solutions, the *Cisco Catalyst® 6500 Series Multilayer Switch* spearheads the evolution toward a highly scalable, multifaceted network architecture that can adapt easily to evolving needs by growing in flexibility and performance and by offering customers unique investment protection.

The flexibility of the Cisco Catalyst 6500 Series makes it suitable to use in any tier of a multitiered network architecture. It supports the powerful features that are oriented to the core, distribution, and access layers of a network. When deployed end-to-end in a network, the Cisco Catalyst 6500 Series provides maximum benefits and a high return on investment (ROI).



Primary Benefits of Cisco Catalyst 6500 Series Switch

The most flexible network device available for any metro solution is Cisco Catalyst 6500 Series Switch. Several effective reasons account for this flexibility (see Table 1).

Table 1 The Cisco Catalyst 6500 Series Switch Advantages

Advanced Layer 2 features	Productivity—Enables new services and provides better network resiliency.
Advanced QoS capabilities	Availability of mission critical applications—Ensure the night applications and users receive the correct service levels throughout the network infrastructure.
Layer 3 through 4 features	Scalability and protection—Allow the network to support IP services and secure information assets by preventing unauthorized access at the source.
High availability	Productivity—High availability maximizes productivity and minimizes support costs.
Advanced multicast support	Productivity—Enable new innovation in voice and video applications while maximizing network efficiency.
Support for 10 GbE, coarse and dense wavelength division multiplexing (CWDM and DWDM) interfaces	Scalability—Maximize network bandwidth and facilitate network design.

All the benefits in Table 1 result from the sophisticated hardware and software architecture that has been the foundation of the Cisco Catalyst 6500 platform since its introduction to the market.

Some of the Cisco Catalyst 6500 Series most important features for metropolitan networks are described in the following sections.

802.1Q Tunneling and TLS

*802.1Q tunneling*¹ (QinQ), also known as *tag stacking*, allows the deployment of secure TLS by building on the standard capabilities of the IEEE 802.1Q protocol (please see reference [2]) that is included on all Cisco switches. In particular, 802.1Q tunneling or tag stacking enables service providers to offer “virtual private LANs” that appear as a logical wire or pipe to their customers. Although some customers use overlapping VLAN ranges, traffic remains isolated from one customer to another customer. Point-to-point and point-to-multipoint topologies are possible and easy to deploy. With the introduction of Layer 2 Protocol Tunneling (L2PT), resilient network designs can be implemented. See the following sections for more details on L2PT.

The main advantage of 802.1Q tunneling is that it enables service providers to segregate traffic from different size (enterprise, medium, or small) customers in their infrastructure, while significantly reducing the number of VLANs required to support individual customer connections. Multiple customer VLANs can be transparently carried inside a *single* provider VLAN configured on a Cisco Catalyst 6500 Series without losing their unique VLAN IDs. In addition, the number of VLANs required to support 802.1Q tunnels in the service provider network can be reduced

1. Supported on all Cisco Catalyst 6500 supervisors and linecards starting with Cisco Catalyst OS Release 6.1 and with Cisco IOS® Software Release 12.1(11b)EX.

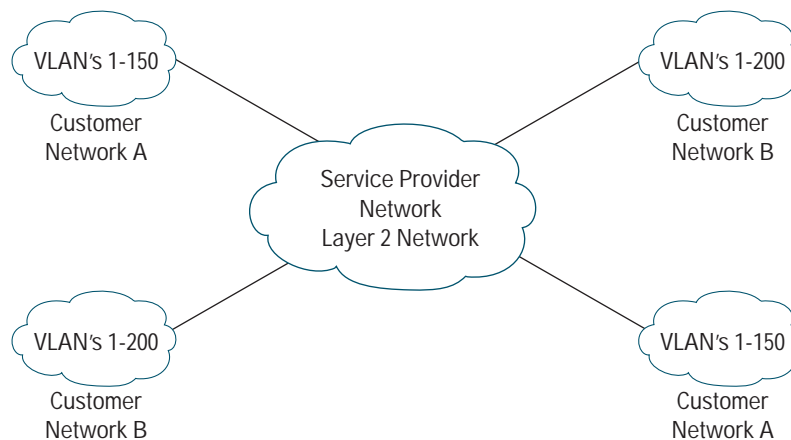


significantly, while the aggregate number of available VLANs can jump from 4096 up to a theoretical maximum of more than 16 million ($= 4096^2$). By using these *Layer 2 tunnels*, it is possible to deliver enterprise-scale connectivity deployed on a shared infrastructure with the same security, prioritization, reliability, and manageability of a private network.

Network Design Without Cisco IEEE 802.1Q Tunneling

Service providers deploying TLS service networks have two main issues: *VLAN overlap and VLAN transparency*. Enterprise customers with VLANs transiting the service provider cloud may use ranges of VLANs that overlap with the service provider (shown in Figure 1), a situation that may cause inadvertent mixing of different customers' traffic through the Layer 2-based network infrastructure. Assigning a unique range of VLANs to each customer in the service provider network solves this problem but rapidly consumes the 4096 VLAN space supported by the 802.1Q standard.

Figure 1
Different Customers Using Overlapping VLAN Ranges




802.1Q tunneling technology addresses this limitation by using a double 802.1Q tagging scheme, allowing service providers to assign a separate VLAN to each customer without losing the original customer VLAN IDs inside the tunnel. For each service provider VLAN, up to 4096 enterprise VLANs can be encapsulated, providing better scalability of the TLS network.

Network Design with 802.1Q Tunneling

In TLS networks, the ideal scenario to support multiple enterprise customers in the service provider environment is to have customers utilize any range of VLAN numbers while the service provider forwards the traffic independently of those VLAN IDs (*transparently*, that is). See Figure 2.



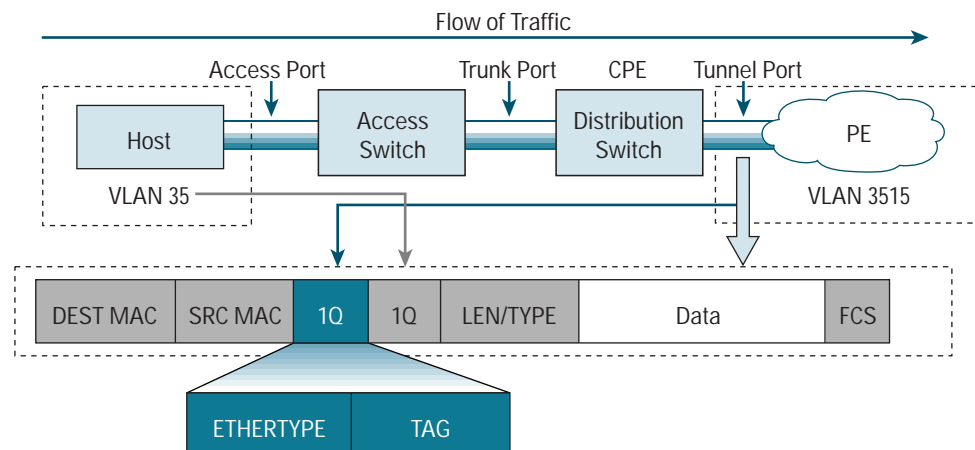
Figure 2
Different Customer VLANs Carried Within Service Providers L2 Tunnels



By assigning a unique VLAN number to each customer, the identity of the VLAN IDs received from customer sites can be preserved. Therefore, this mechanism builds Layer 2 tunnels inside the service provider core where traffic from different business customers is segregated and where it is marked with an appropriate *tunnel ID*.

802.1Q tunneling is also called *tag stacking* because it is a *double encapsulation technique*. It expands the VLAN space by adding an additional 802.1Q tag (the tunnel ID) to all previously-tagged packets when they enter the service provider infrastructure, as illustrated in Figure 3.

Figure 3
Different Customer VLANs Carried Within Service Providers Layer 2 Tunnels



The interfaces on the *customer premises equipment* (CPE) are configured as 802.1Q trunks (note: Inter-Switch Link [ISL] encapsulation technology is not supported), while the interfaces on *provider edge* switches are configured as special *802.1Q tunnel ports*. Because one side of this point-to-point link is a trunk and the other is not, it is called an *asymmetric link*. The special tunnel port on the provider edge side is configured with a unique VLAN that the



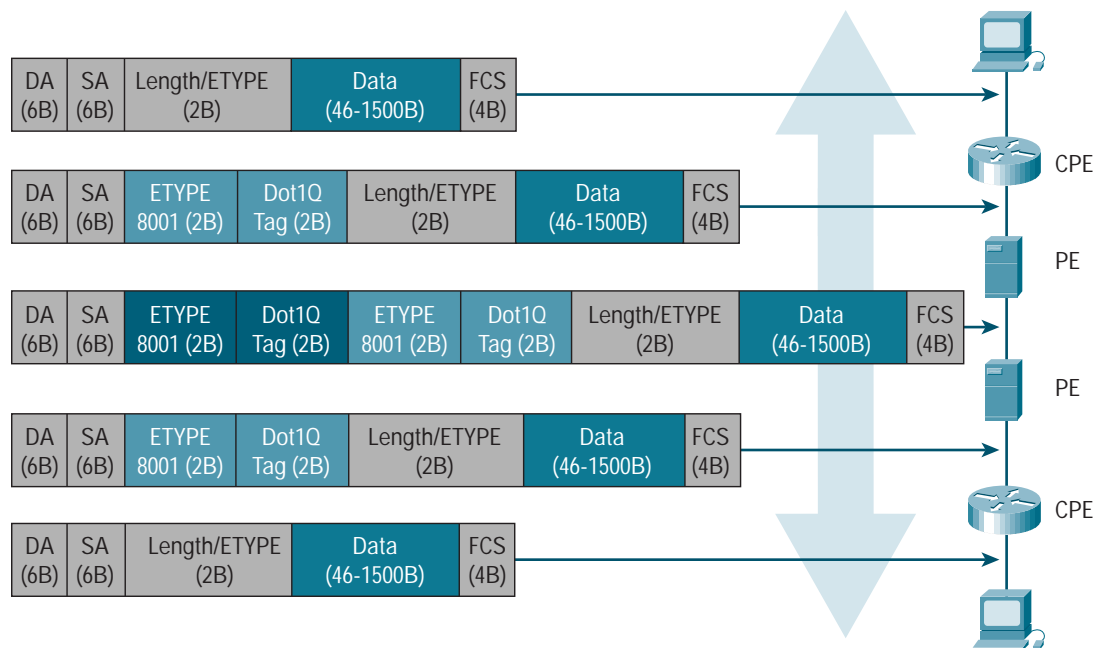
service provider assigns to the directly-attached customer. When this port receives tunnel traffic from the neighboring device, it does not follow the usual procedure of stripping the 802.1Q tags from the frame header for internal forwarding processes. Instead, it leaves the incoming 802.1Q tag intact and puts all the received 802.1Q traffic into the VLAN assigned to the tunnel port. As a result, traffic is double-tagged² as it enters the first port of the service provider infrastructure.

The CPE may be any switch capable of generating 802.1Q-encapsulated frames. Only the provider edge switches inside the service provider network need to be “tag-stacking capable.” Other service provider core switches (called *P switches*) do not need to be aware of the double-tagging scheme, because they switch traffic based on the outermost 802.1Q tag; therefore, no special software is required on P switches.

802.1Q Tunneling Technology Overview

To understand the functions and operation of the 802.1Q tunneling technology, examine a packet flow to see how the feature works when a packet traverses the service provider network. Figure 4 illustrates the process. The steps are described later in the section.

Figure 4
Frame Passing Through an 802.1Q Tunnel



The following steps demonstrate how 802.1Q Tunneling works:

2. Traffic is *logically* double-tagged. In fact, a user cannot see how many tags are used inside the switch when the packets are received.



Step A

An Ethernet frame on a specific VLAN arrives on the provider edge switch from the CPE. The CPE switch is any switch capable of supporting 802.1Q trunks. ISL trunks are not supported because the egress frame from the CPE needs to be 802.1Q-encapsulated, as depicted below:]

DA (6B)	SA (6B)	Etype (8100) (2B)	802.1Q Tag (2B)	Length/Etype (2B)	Data (46 ¹ -1500 Bytes)	FCS (4B)
---------	---------	-------------------	-----------------	-------------------	------------------------------------	----------

1. The minimum payload size of an 802.1Q packet is actually 42 bytes and can be used by certain 802.1Q-capable hosts. Devices sending untagged traffic must use a minimum payload size of 46.

(DA = destination address, SA = source address, Etype = Ethertype, FCS = frame control sequence)

Step B

The ingress port of the provider edge does not remove any 802.1Q headers from the incoming frame; instead, it adds a second tag. In the example below, a second VLAN is assigned to a customer on the provider edge switch and, therefore, its tag (*outer tag*) is appended to the original tag (*inner tag*). ISL or 802.1Q links can carry tunneled traffic between provider edge (PE) and provider core (P) switches inside the service provider infrastructure. Below is the diagram of the frames leaving the provider edge Cisco Catalyst 6500 Series switch. Additions to the original frame are shown in blue and in boldface. If the ISL trunk encapsulation is used from provider edge to provider core devices, the result is:

ISL (26B)	DA (6B)	SA (6B)	Etype 8100 ¹ (2B)	CPE 802.1Q Tag (2B)	Length/Etype (2B)	Data (46-1500 Bytes)	FCS (4B)
------------------	---------	---------	------------------------------	---------------------	-------------------	----------------------	----------

1. 0x8100 is a reserved Ethertype indicating that an 802.1Q tag follows.

If, instead of ISL, we use the 802.1Q encapsulation, the result is:³

DA (6B)	SA (6B)	Etype (8100) (2B)	PE 802.1Q Tag (2B) OUTER TAG	Etype (8100) (2B)	CPE 802.1Q Tag (2B) INNER TAG	Len/Etype (2B)	Data (46-1500 bytes)	FCS (4B)
---------	---------	--------------------------	-------------------------------------	-------------------	-------------------------------	----------------	----------------------	----------

Step C

The egress provider edge switch removes the outer 802.1Q tag before delivering traffic to the destination CPE. In this example, the outer tag is removed but the original inner tag is retained:

DA (6B)	SA (6B)	Etype (8100) (2B)	802.1Q Tag (2B)	Length/Etype (2B)	Data (46-1500 bytes)	FCS (4B)
---------	---------	-------------------	-----------------	-------------------	----------------------	----------

3. When using the 802.1Q encapsulation between a provider edge and a provider core switch or between two provider edge switches, two 802.1Q tags are present in the frames leaving the ingress provider edge switch. The provider edge-assigned tag is referred to as the *outer tag*, and the CPE-assigned tag is referred to as the *inner tag*. Provider core switches forward double-tagged frames based exclusively on the outer tag.



Finally, the CPE switch removes the original tag and forwards the frame to the end host, with the following result:

DA (6B)	SA (6B)	Length/Etype (2B)	Data (46–1500 bytes)	FCS (4B)
---------	---------	-------------------	----------------------	----------

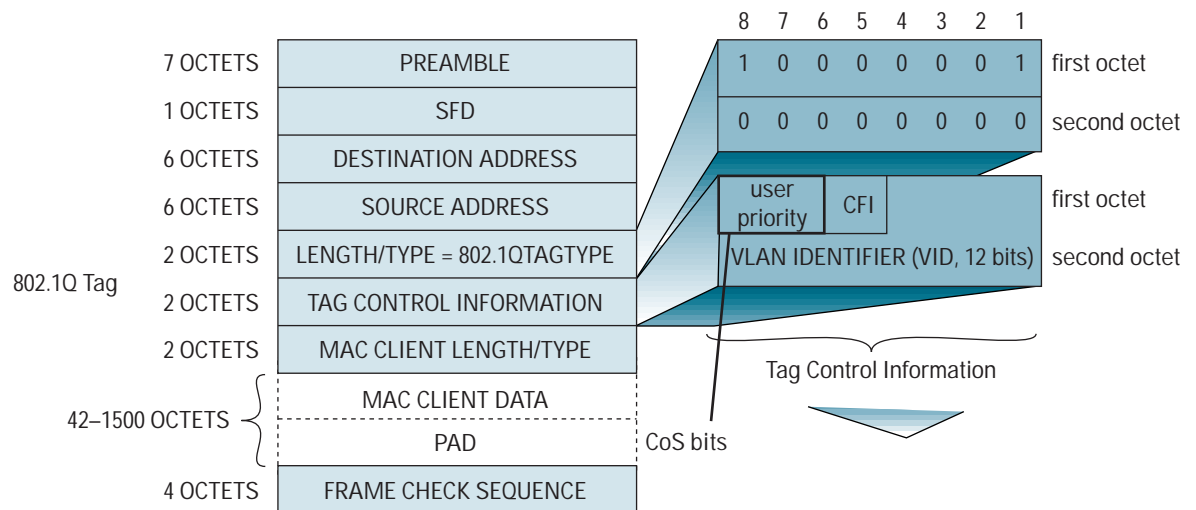
QoS and CoS Mutation with 802.1Q Tunneling

When transiting through the service provider core, double-tagged frames always retain the inner 802.1Q tag originally imposed by the CPE device. As a result, packets appear on the switching fabric of Cisco Catalyst 6500 Series acting as provider edge (PE) or provider core (P) switches with two 802.1Q tags. This creates a *byte-realignment* within the frame, causing the IP header information to be unidentifiable by provider edge or provider core switches. This imposes the following conditions:

- Layer 3 and higher information within the double-tagged Layer 2 frame cannot be identified
- Tunnel traffic (e.g. between PE and P, or between P and P switches) cannot be routed or Layer 3 policed
- PE and P switches can filter tunnel traffic using only Layer 2 parameters (VLANs or Source/Destination MAC addresses). A MAC-based access list will match all types of traffic: MAC-based, IP, IPX, etc.

The inner tag's *user priority* bits (also known as the three bits of Class of Service [CoS]) are preserved end-to-end, from CPE to CPE. Figure 5 shows where CoS bits are located in the 802.1Q tag.

Figure 5
Position of the CoS Bits



It is always possible for the service provider to assign CoS values to the outer tag, either statically or dynamically. To mark the outer tag CoS bits statically, the network administrator assigns a default CoS value to the CPE-facing 802.1Q tunnel port. This results in one selectable CoS value per CPE-facing port.

When using a Cisco Catalyst 6500 Series Switch equipped with a Supervisor Engine 2 (or higher), it is also possible for the network administrator to dynamically assign the outer tag CoS value, based on the incoming traffic rate conformance to a bandwidth contract. To assign the outer tag dynamically, a dual-rate policer specifying a normal and an excess rate is defined. Different actions can be specified, depending on the traffic behavior. A typical scenario

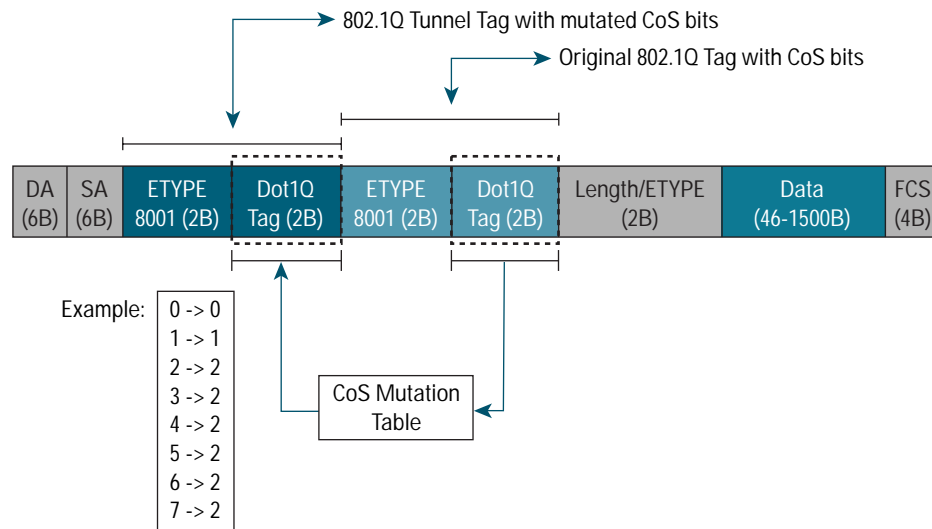


consists in selling a customer three classes of service: bronze, silver, and gold. First, a static CoS value (two, for example) is assigned to the CPE-facing port. When traffic conforms to the normal rate, the policer takes no action, the port sets the CoS to two in the outer tag, and traffic flows through. When traffic exceeds the normal rate but is below the excess rate, the outer tag CoS is set (“marked down”) to one, and traffic is transmitted. When traffic exceeds the excess rate, the CoS is set to zero.

When using a Cisco Catalyst 6500 Series equipped with Supervisor Engine 720 and the new WS-X6700 series of line cards or the new WS-X6802 10 Gigabit Ethernet modules, an additional function—*CoS Mutation*—allows the service provider’s provider edge switch to *inspect* the CPE-assigned inner tag CoS value and use it to assign the outer tag CoS value. In this scenario, the ingress *802.1Q tunnel* port would have to be configured as a *trusted* port (meaning that the received CoS bits will be inspected) and an ingress *CoS mutation map* would have to be used to specify which outer tag CoS value to use in correspondence with each inner tag CoS value.

The CoS Mutation functionality is described in Figure 6.

Figure 6
CoS Mutation on 802.1Q Tunnel Ports



The default CoS mutation map contains the identity mappings: 1 to 1, 2 to 2, 3 to 3, etc. In this configuration, the feature is sometimes informally called “copy of inner to outer CoS bits.”

L2PT

Prior to Catalyst OS Release 7.1 and Cisco IOS® Software Release 12.1(11b)EX, the 802.1Q tunneling feature did *not* tunnel customer-originated spanning tree bridge protocol data units (BPDUs), Cisco Discovery Protocol or VLAN Trunking Protocol (VTP) frames through the service provider cloud. Customer BPDUs were tunneled only from CPE-facing to CPE-facing port on the same provider edge switch, as shown in Figure 7.

Without protocol tunneling, Cisco Discovery Protocol and VTP PDUs are terminated at the provider edge switch; in other words, referring to Figure 7, a `show cdp neighbor` command performed on any CPE device would show the provider edge as a neighbor instead of another CPE. Catalyst OS Release 7.1 and Cisco IOS Software




Release 12.1(11b)EX introduced L2PT, which allows customer PDUs (such as BPDUs, Cisco Discovery Protocol, and VTP frames) to be transparently carried from CPE to CPE across the service provider network creating a *logical connection* between the two ends.

The L2PT feature provides a scalable approach to PDU tunneling by software encapsulating the PDUs on the ingress provider edge switches and then multicasting them out, using the hardware forwarding mechanism of the switch. All P switches inside the service provider network treat these encapsulated frames as regular data packets and forward them out appropriately. The egress provider edge switch listens for these specially encapsulated frames and decapsulates them before forwarding them out of the tunnel.

The L2PT encapsulation mechanism rewrites the destination MAC address of CPE-originated PDUs. In particular, the ingress provider edge switch rewrites the destination MAC address of the PDUs received on a tunnel port with a Cisco reserved multicast MAC address (01-00-0c-cd-cd-d0). The PDU is then flooded out on the tunnel port VLAN (for example, if the 802.1Q tunnel port is in VLAN 5, the frame is flooded in VLAN 5). Figure 7 illustrates the L2PT concept with the three types of special frames.

Figure 7
With L2PT, BPDUs, Cisco Discovery Protocol, and VTP Messages Are Passed Across the Layer 2 Service Provider Network

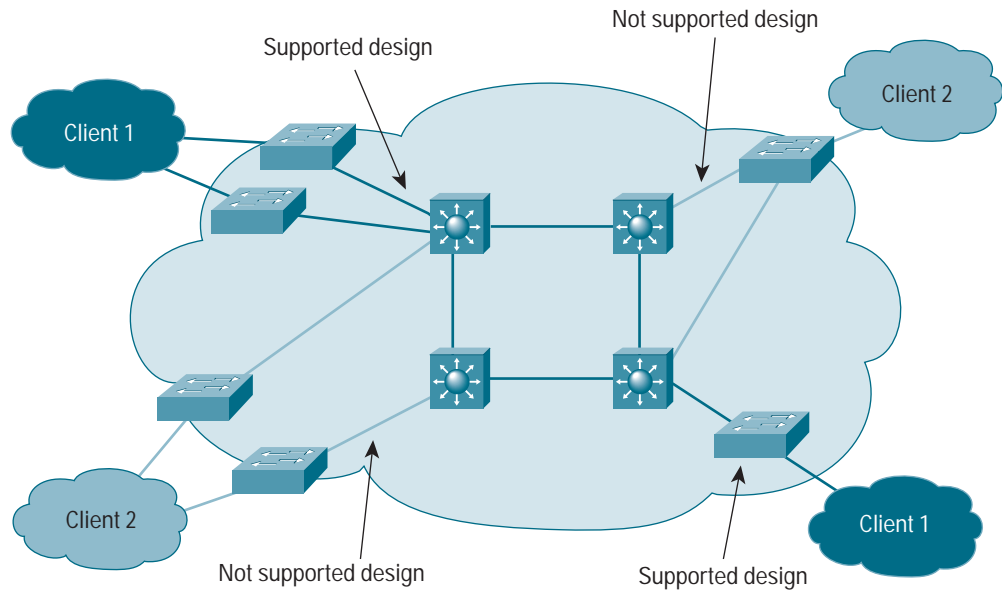


Spanning Tree and L2PT

If the L2PT feature is not enabled, providing Layer 2 redundancy to customers is difficult because, by default, 802.1Q tunneling does not transport customer BPDUs between two sites across the service provider backbone. Therefore, loops are not allowed between client sites that are interconnected by 802.1Q tunneling. The only redundant design that is allowed is when the CPE switches of a client site are connected to the *same* provider edge in the service provider core as depicted in the figure below.

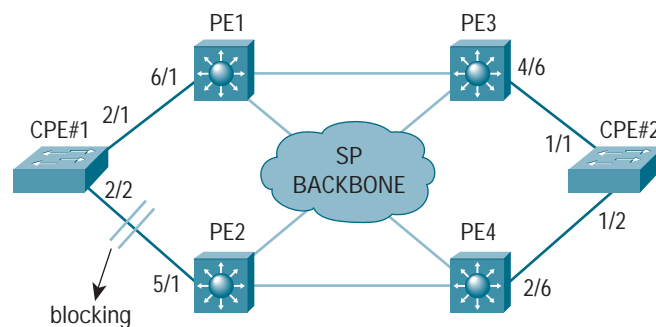


Figure 8
Redundant Layer 2 Design Without L2PT



The issue of a redundant design primarily relates to the spanning tree interaction between the enterprise and the service provider. Within the service provider cloud itself, the normal rules of spanning tree, either 802.1D or 802.1S/W, apply. The introduction of L2PT offers more flexibility to the service provider in terms of resilient network design. The service provider can offer dual homing to CPE switches as illustrated in Figure 9.

Figure 9
CPE-to-CPE Redundancy With L2PT



In Figure 9, the service provider has enabled L2PT for BPDUs on all four provider edge switches on ports 6/1, 5/1, 4/6 and 2/6. These ports are also 802.1Q tunnel ports and are all part of the same VLAN: this creates a L2 VPN between CPE#1 and CPE#2. BPDUs sent from CPE#2 (assumed to be the root for the customer VLANs) are now received by CPE#1, which in turn can block port 2/2 to break the logical loop. Without L2PT, the loop would have been undetected and traffic would have circled endlessly around the topology, paralyzing the network.



Rootguard at the Ingress Provider Edge Switch


When L2PT is not enabled,⁴ the spanning tree of the native⁵ VLAN of the CPE switch merges with the spanning-tree of the 802.1Q tunnel port VLAN of the provider edge switch, even if the native CPE switch VLAN does not match the service provider-assigned VLAN on the provider edge. BPDUs for other VLANs are simply dropped when arriving on the PE switch. This is undesirable for a service provider, because any CPE device potentially becomes the root of the service provider backbone VLAN (the VLAN of the 802.1Q tunnel port). Cisco developed *Rootguard* to prevent root bridge takeover.

Basic spanning tree calculations involve two fields: Bridge ID (BID) and Path Cost. A BID is a single, 8-byte field that is composed of two subfields shown below:

Bridge Priority (2 bytes)	MAC address (6 bytes)
---------------------------	-----------------------

The default bridge priority is 32,768. Cisco Catalyst 6500 Series switches choose a single root bridge per VLAN, based on the lowest BID. If the CPE switch in the tag stacking scenario has a lower MAC address and bridge priority than the provider edge switch, it becomes the spanning tree root for the service provider-assigned VLAN, resulting in a modified service provider-backbone topology.

Figure 10
Selection of an Incorrect Root Switch Outside of the Service Provider Network



Rootguard can be enabled on a per-port basis on the provider edge switch to prevent customer switches from becoming the root. If the client side attempts to take root ownership, the 802.1Q tunnel port goes into a root-inconsistent state. When the port stops receiving superior BPDUs from the CPE switch, it leaves the root-inconsistent state after a period of *max_age* seconds.

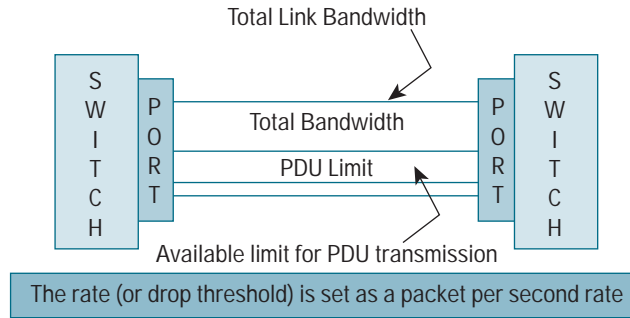
L2PT Thresholds and Rate Limiters
To ensure that the PDUs do not flood the network, L2PT supports input per-port *drop thresholds* that can be configured to set a predefined level (in packets per second [pps]). PDUs to be tunneled that are above this level are dropped. This limits the total processing bandwidths that PDUs can consume.

4. With L2PT, BPDUs from the CPE switch native VLAN can no longer merge with the provider edge switch's spanning tree. In that case, Rootguard is unnecessary.

5. The native VLAN of an 802.1Q port is the VLAN which all untagged frames are assigned to.



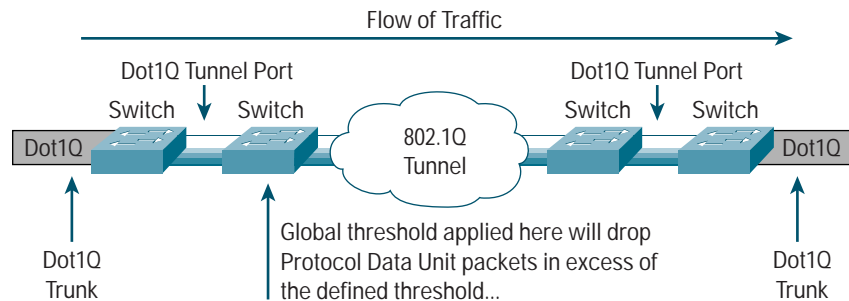
Figure 11
Drop Thresholds for Tunneled PDUs



An input *shutdown threshold* can also be configured on a port and, if the volume of PDUs exceeds the shutdown threshold, the port is put into an *errdisable* state.

Beginning with Cisco IOS Software Release 12.2(17a)SX, L2PT has been enhanced with a global rate limiting ability. Based on a defined *global threshold*, any VTP, Cisco Discovery Protocol, or STP packets in excess of that threshold limit are dropped to better shield the management CPU from dangerous surges of traffic coming from a *multitude of* CPE devices. This capability, coupled with Supervisor Engine 720's "*mls rate-limit layer-2*" hardware functionality, provides a good protection of the control plane from possible overloads due to an excess of PDU traffic.

Figure 12
Global Drop Thresholds for Tunneled PDUs

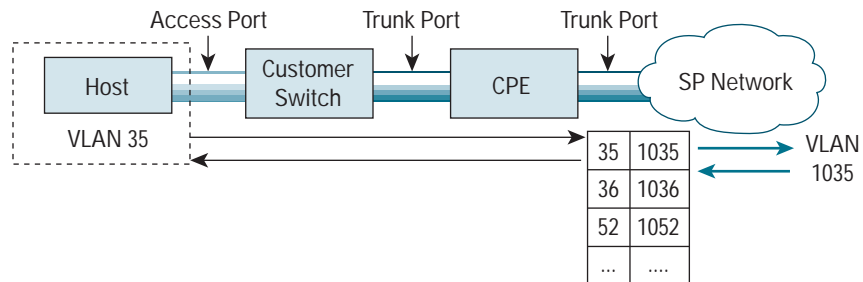


VLAN Translation

Occasionally, service providers must provide an 802.1Q VLAN to their customers without being able to tunnel it through an 802.1Q tunnel. That creates a constraint on the VLAN number to use, because that VLAN is allocated globally in the service provider network and can no longer be used by another customer. To overcome this limitation, beginning with Cisco IOS Software Release 12.2(17b)SXA, the Cisco Catalyst 6500 Series provides *VLAN translation*, which allows the service provider to rewrite the packet VLAN when the traffic ingresses or egresses the network, as shown in Figure 13.



Figure 13
VLAN Translation at the Ingress Point of the Service Provider Network




VLAN translation is a hardware function available on most of the port application-specific integrated circuits (ASICs) of the Cisco Catalyst 6500 Series line cards. Its support varies, depending on the type of line card in use. For more details, refer to Cisco IOS Software Release 12.2(17b)SXA's documentation [3] and release notes.

DSCP Transparency

Due to its L3 heritage, the default behavior of the Cisco Catalyst 6500 Series architecture links the traffic CoS value to a corresponding differentiated services code point (DSCP) value, which is treated as a superset of the original CoS. Most metro Ethernet solutions require that the customer DSCP value be maintained *unaltered end-to-end*, and only the Layer 2 CoS value be used to mark or classify the traffic within the service provider network.

First available with Cisco IOS Software Release 12.1(19)E1, the *mls queuing-only* feature was introduced to maintain CoS and DSCP unchanged while disabling the policy feature card (PFC) QoS functions (marking and policing). This behavior is undesirable and, starting with Cisco Catalyst OS Release 8.2 and Cisco IOS Software Release 12.2(17b)SXA on Supervisor Engine 720, a hardware-based feature called *DSCP Transparency* was implemented to break the linkage between CoS and DSCP, even when PFC QoS functions are enabled. Figure 14 shows the behavior of a network where the DSCP value is preserved end-to-end by using DSCP Transparency.

Figure 14
DSCP Transparency




Ethernet over MPLS

Ethernet over *Multiprotocol Label Switching (EoMPLS)* is used to transport Ethernet frames across an MPLS network along a point-to-point path.



When the edge switch receives an Ethernet packet, it is encapsulated into an MPLS packet to be transported across the network over a *virtual circuit*. Two labels are imposed on the original Ethernet packet; one represents the tunnel the packet is to use, and the other represents the MPLS path to forward the packet along. An overview of an EoMPLS application is shown in Figure 15.

Figure 15
EoMPLS Example



EoMPLS has four main types of transport:

- 802.1Q VLAN to 802.1Q VLAN, often referred to as *VLAN mode*
- Ethernet port to Ethernet port, often referred to as *port mode*
- 802.1Q VLAN to Ethernet port
- *Switched Virtual Interface (SVI)*⁶ mode, in which both Layer 2 switching and EoMPLS tunneling are performed on the traffic

EoMPLS on Supervisor Engine 720 is supported in two different configurations:

- With Supervisor Engine 720's PFC3bXL and DFC3bXL performing the EoMPLS feature. In this case any module can be the uplink to the MPLS core.⁷
- With a FlexWAN or an optical service module (OSM) parallel express forwarding (PXF)-enabled interface as the uplink to the MPLS core performing the EoMPLS feature.

EoMPLS on Supervisor Engine 2 is supported, instead, with only a FlexWAN or an OSM PXF-enabled interface as the uplink to the MPLS core.

In an advanced metro deployment scenario, it is possible to combine the point-to-multipoint capabilities of 802.1Q tunneling with the scalability of EoMPLS (derived from its MPLS foundation) to gain the maximum benefit from the deployment of these two powerful technologies. The logical evolution of this type of design is described in the following section.

6. Switched Virtual Interface. A routed interface associated to a L2 switched VLAN.

7. However, in this case SVI mode is not supported.

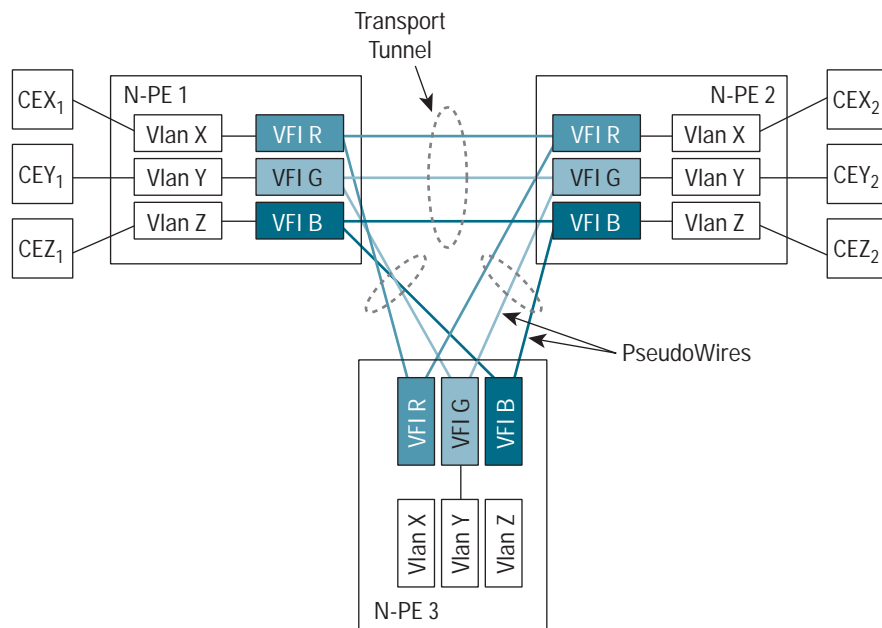


Virtual Private LAN Service

Virtual Private LAN Service (VPLS) provides a multipoint solution for Layer 2 connectivity across an MPLS network. Although EoMPLS provides a point-to-point *logical wire* or “*circuit*” that can be terminated on a single provider edge device, VPLS provides either a *full mesh* or a *partial mesh* of connection circuits (or *pseudo-wires*) between all the participating provider edges, to allow all customer equipment to connect to each other as if they were attached to the same Ethernet segment.

Figure 16 illustrates the VPLS architecture.

Figure 16
VPLS Architecture



In Figure 16, the customer equipment that is attached to the provider edges by using VLANs coupled to the so-called *virtual forwarding instances* (VFIs). VFIs replace the virtual circuits used with regular EoMPLS and are the new abstractions for the configuration of the multipoint tunnels in VPLS.

With VPLS, broadcast, multicast, and unknown unicast traffic is handled by the provider edge switches by flooding the traffic to all pseudo-wires. To prevent loops, the provider edge switches use the *split horizon algorithm*; therefore, they do not send traffic back into the MPLS cloud.

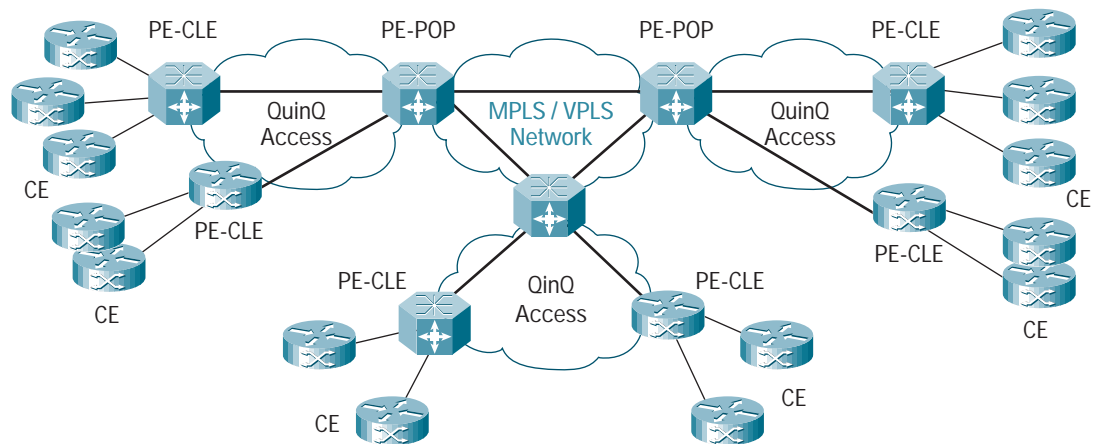


Beginning with Cisco IOS Software Release 12.2(17b)SXA, Supervisor Engine 720⁸ supports VPLS when using the following OSM PXF-enabled interfaces as uplinks to the MPLS core:

OSM modules supporting VPLS
OSM-2-4GE-WAN+
OSM-4OC3-POS-MM+, -SI+
OSM-8OC3-POS-MM+, -SI+
OSM-16OC3-POS-MM+, -SI+
OSM-2OC12-POS-MM+, -SI+, -SL+
OSM-4OC12-POS-MM+, -SI+, -SL+
OSM-1OC48-POS-SS+, -SI+, -SL+
OSM-2OC48-POS/DPT-SS, -SI, -SL (packet over SONET [POS] mode only)

As with EoMPLS, in the most sophisticated metro deployments it is possible to combine the Layer 2 multipoint scalability and the simplicity of 802.1Q tunneling to the robustness of a multipoint MPLS core by tying 802.1Q tunneling and VPLS together in a hierarchical fashion, as shown in Figure 17.

Figure 17
Integration of 802.1Q Tunneling with VPLS



Hardware Flexibility and Scalability

The Cisco Catalyst 6500 Series supports a wide variety of module options for business and residential service deployments. Cisco Catalyst 6500 Series switches ship in 3-slot, 6-slot, 9-slot, and 13-slot configurations and support a range of Ethernet interfaces from 10/100 Mbps to Gigabit to 10 Gigabit Ethernet with copper and fiber interfaces.

8. And in a subsequent release, also Supervisor Engine 2.



Recent additions to the already-ample portfolio of modules include the advanced 6700 and 6800 series line cards, which require Supervisor Engine 720:

- Cisco Catalyst 6500 Series 2-port 10 Gigabit Ethernet Module
- Cisco Catalyst 6500 Series 4-port 10 Gigabit Ethernet Module
- Cisco Catalyst 6500 Series 24-port mixed media Gigabit Ethernet Module
- Cisco Catalyst 6500 Series 48-port 10/100/100 Gigabit Ethernet Module


These new modules support a variety of media interfaces, ranging from fixed RJ-45 connectors for Category 5 twisted-pair cables to the flexible media modules known as Small Form-Factor Pluggables (SFPs) that support 1000BASE-SX/LH/ZX, as well as coarse wavelength-division multiplexing (CDWM) transmission.

On the 10 Gigabit Ethernet side, new XenPak pluggable modules support the latest transceivers defined by IEEE, including the 10GBASE-LR (1310-nm long-reach technology that supports distances up to 10 kilometers (km) and the 10GBASE-ER (1550-nm extended reach technology that supports distances up to 40 km).

In some metro applications, a 10GBASE-ER XenPak can be used for long-haul deployments in conjunction with a Cisco Erbium Doped Fiber Amplifier (EDFA) to extend the reach of a 10 Gigabit Ethernet Single Mode Fiber (SMF) link, even in excess of the standard maximum distance.

Figure 18 summarizes Cisco Catalyst 6500 Series wide range of hardware technologies for any Metro Ethernet deployment.

Figure 18
Media Options Offered by the Newest Cisco Catalyst 6500 Modules Module Series





Conclusion

As the premier Cisco intelligent multilayer modular switch, the Cisco Catalyst 6500 Series delivers secure, converged network services from the metro access to the WAN/MAN edge for service provider metro Ethernet deployments.

Supervisor Engine 720 is part of Cisco Catalyst 6500 Series third-generation (3G) suite of modules to continue lowering total cost of ownership (TCO) while increasing the total system bandwidth and packet-forwarding capability. A Cisco Catalyst 6500 Series switch using a Supervisor Engine 720 offers up to 410 Gigabit Ethernet ports, with a maximum throughput of 400 Mpps for IPv4 and 230 Mpps for IPv6.

Cisco Catalyst 6500 Series optimizes infrastructure utilization and maximizes ROI by using forward-thinking architecture that supports an unparalleled range of services, including data and voice integration and WAN/MAN convergence with these additional benefits:

- Maximum network uptime for higher user productivity and business resiliency
- Comprehensive network security using proven multigigabit Cisco technology
- Investment protection and long product life cycle supporting multiple generations of interfaces and packet forwarding engines
- Operational consistency allowing customers to standardize on a single platform that addresses all network deployment requirements
- Leading services integration supporting the application-aware convergence of data, voice, and video onto a single highly-manageable platform.

In a metro Ethernet Layer 2 end-to-end deployment, Cisco Catalyst 6500 Series can be deployed in the metro access, aggregation, edge, and core layers of the network because of the advanced metro technologies that it supports.

Similarly, in a metro environment with a Layer 2 access connected to a Layer 3 edge and core network or in a Layer 3 end-to-end deployment, the Cisco Catalyst 6500 Series can show its full flexibility by fitting in any layer of the network.

References

1. Metro Solutions for Service Providers,
http://www.cisco.com/en/US/netsol/ns341/ns396/ns223/ns227/networking_solutions_sub_solution.html
2. IEEE P802.1Q/D11 (July 30, 1998), Draft Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks,
<http://standards.ieee.org/reading/ieee/std/lanman>
3. Configuring VLAN Translation,
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080160a7c.html#1044990
4. IEEE Std 802.3ae-2002, IEEE Amendment to Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications,
<http://standards.ieee.org/reading/ieee/std/lanman>



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.
Capital Tower
168 Robinson Road
#22-01 to #29-01
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 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 © 1992–2004 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, the Cisco Systems logo, Catalyst, and Cisco IOS 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.
(0304R) ETMG 203144—JR 01.04