

Layer 2 Tunneling Protocol Version 3 Technical Overview

Introduction

Layer 2 Tunneling Protocol version 3 (L2TPv3) allows service providers and large enterprises with native IP core networks to offer high-speed Layer 2 tunneling or VPN services to end-user customers, in conjunction with their Layer 3 VPN offerings. L2TPv3 VPN services can be provided without increasing the expenditure for capital equipment by simply upgrading Cisco IOS Software. L2TPv3 is provided as part of the Unified VPN portfolio of leading-edge VPN technologies available over the widest breadth of Cisco routers.

L2TPv3 is emerging as a core tunneling and VPN technology for next-generation networks. L2TPv3 provides the flexibility and scalability of IP with the privacy of Frame Relay and Asynchronous Transfer Mode (ATM). L2TPv3 will allow network services to be delivered over routed IP networks. Service decisions will be made at the VPN and tunnel endpoints and switched without requiring intermediate preprocessing, providing higher efficiency and scalability.

By reducing customer networking complexity and cost, L2TPv3 VPNs allow service providers to serve a more diverse base of small and medium-sized businesses. Rather than setting up and managing individual point-to-point circuits between each office, a business provides only one

connection from its office router to a service-provider edge router. Service providers expand service offerings and generate additional revenue by offering customers VPNs with managed Internet, intranet, and extranet without the complexity that these applications previously required.

L2TPv3 offers the following advantages for service providers:

- Provides a simple tunneling mechanism to implement transparent LAN and IP functionality, offering a simple means for IP VPN services
- Simplifies the interaction between Service Provider networks and Service Provider/Customer networks
- Protect Existing Investment while Building Packet Core Enhanced VPN Support
- Facilitates New Services
- Allows the transport of non-IP protocols, such as Internetwork Packet Exchange (IPX) and SNA, as well as other desktop protocols

L2TPv3 offers the following advantages for enterprise customers:

- Simplifies the interaction between Service Provider and Customer networks.
- Allows customer to selectively utilize service provider or corporate facilities in order to deploy VPNs



- Allows the transport of non-IP protocols, such as Internetwork Packet Exchange (IPX) and SNA, as well as other desktop protocols
- Easy configuration.
- Enhanced VPN Support though the use of IOS features such as Security, QoS and Management VPNs can be tailored to meet customer requirements

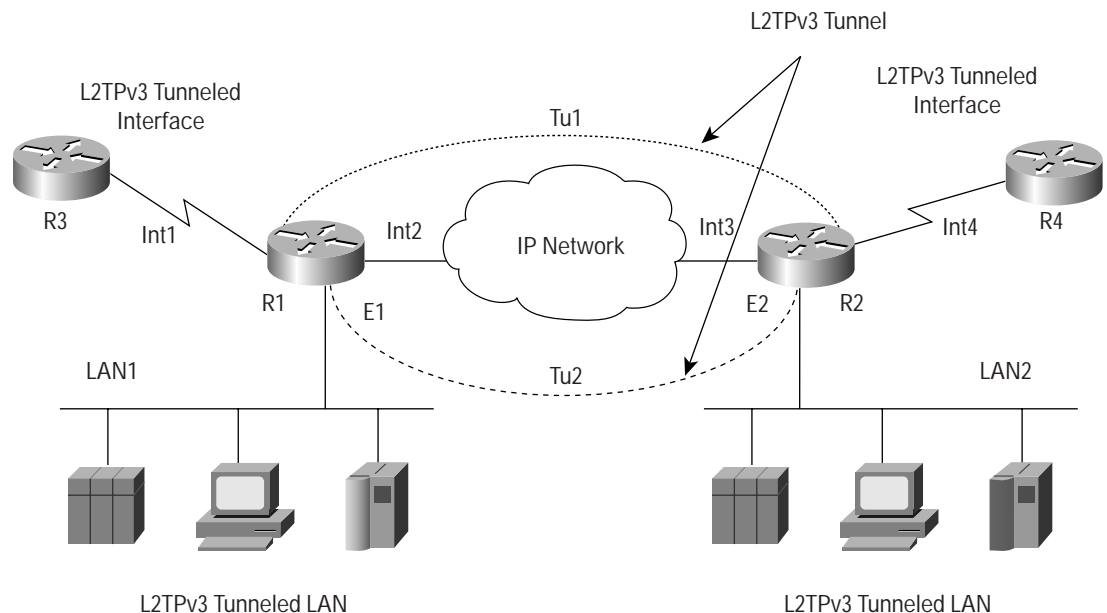
Technical Overview

L2TPv3 technology allows a pair of routers connected via an IP network to provide high-speed transparent Layer 2 connectivity between a pair of interfaces. This functionality can be used to build Layer 2 VPNs or to support traditional (Frame Relay, ATM, leased-line) network migration. L2TPv3 tunnels are available with the IOS basic IP package.

L2TPV3 Operation in Cisco 7000 and 12000 Series Routers

This section discusses interface-based L2TPv3, wherein all traffic between two customer network sites is encapsulated in an IP packet and sent across an IP network. The internal routers of the IP network treat the traffic as any other IP packet and do not need to know anything about the customer networks. This process is known as Layer 2 tunneling (Figure 1).

Figure 1
L2TPv3 Operation



In Figure 1, routers *R1* and *R2* provide L2TPv3 services. These routers communicate with each other using the IP protocol through a path comprising the interface *Int2*, the IP network, and interface *Int3*. In this example, routers *R3* and *R4* communicate through Packet-over-SONET (POS) interfaces using an L2TPv3 tunnel. The L2TPv3 tunnel *Tu1* is configured between interface *Int1* on *R1* and interface *Int4* on *R2*. Any packet arriving on interface *Int1* on *R1* is encapsulated in L2TPv3 and sent via the tunnel (*Tu1*) to *R2*. *R2* decapsulates the packet and transmits it on interface *Int4* to *R4*. When *R4* needs to send a packet to *R3*, the packet follows the same path in reverse.



Note the following regarding L2TPv3 operation:

- All packets received on interface *Int1* will be forwarded to *R4*. *R3* and *R4* cannot see the intervening network.
- In Cisco 12000 Series Internet routers, the other LAN ports on the card that are not being used for L2TPv3 must have a router connected to them. When Media-Access-Control (MAC) filtering assisted by Content-Addressable Memory is turned off to allow L2TPv3 to work, it is turned off on all ports.
- This same method is used for Ethernet interfaces: Any packet received from *LAN1* by *R1* on Ethernet interface *E1* will be encapsulated in L2TPv3 and sent via tunnel *Tu2* to *R2* interface *E2*, where it will be transmitted on *LAN2*.
- This same method is used for Frame Relay subinterfaces: Any packet received from *LAN1* by *R1* on a subinterface will be encapsulated in L2TPv3 and sent via tunnel to *R2* subinterface, where it will be transmitted on *LAN2*.

L2TPv3 Header Description

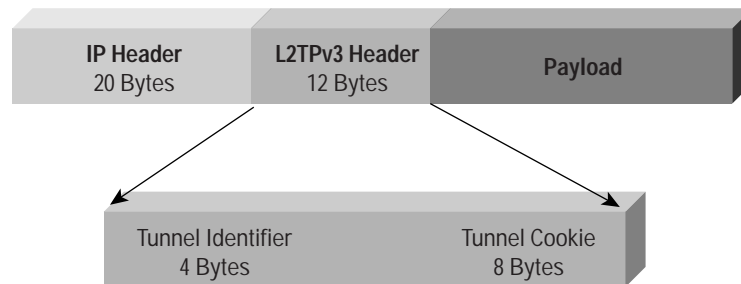
When data enters a L2TPv3 tunnel ingress interface, it is encapsulated with an additional L2TPv3 header (Figure 2).

The L2TPv3 header consists of:

- Payload independent header (12 bytes)
- P delivery header (20 bytes)

The L2TPv3 header takes the format shown in Figure 2.

Figure 2
Packet Encapsulation with L2TPv3



The L2TPv3 header parameters are:

- *Delivery header*—The header needed to carry the L2TPv3 packet across the delivery network. This is an IPv4 header. The delivery header is 20 bytes.
- *L2TPv3 header*—Contains the necessary information needed to uniquely identify the tunnel context at the de-encapsulation point. The payload independent header is 12 bytes.
- *Payload*—To be transported by L2TPv3. It may be a link layer frame or a network layer packet.
- *Tunnel identifier*—Identifies the tunnel context on the decapsulating system. The value of the tunnel ID is selected to optimize the context identification efficiency of the decapsulating system. A decapsulation implementation may therefore elect to support a smaller tunnel-identifier bit field. In this implementation this was achieved by setting an upper value for the L2TPv3 tunnel identifier of 1023. The L2TPv3 tunnel identifier value 0 is reserved for use by the protocol.



- **Tunnel cookie**—An 8-octet signature that is shared between the two endpoints of an L2TPv3 tunnel. This tunnel cookie reduces the chance that contamination of the decapsulated traffic will occur because of an error in configuration. This signature is configured at both the source and destination routers and must match, or the data will be dropped.

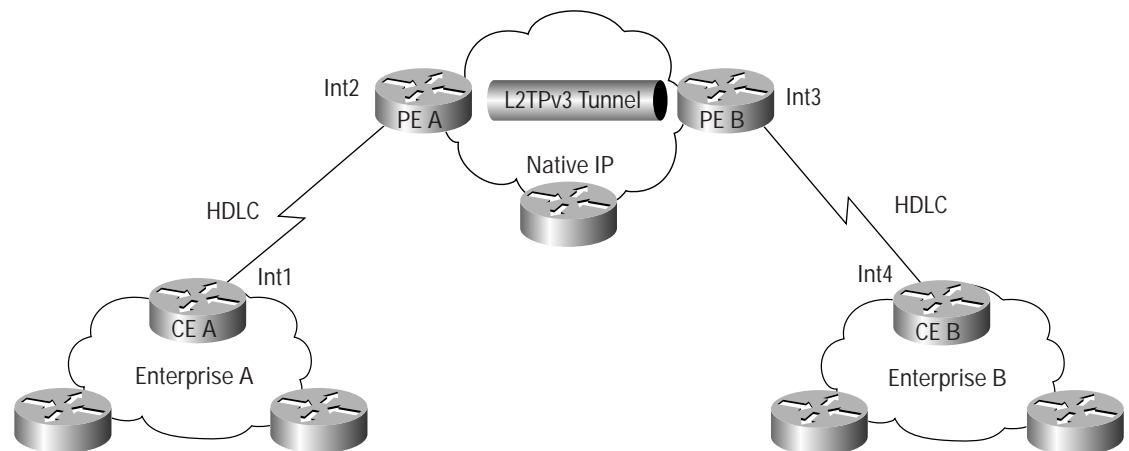
Raw Mode Support

Raw mode is the ability to tunnel information that is arriving over a given physical interface without regard to the type of information being carried. In Raw mode, a physical interface is attached to both ends of an L2TPv3 tunnel. All packets and frames that arrive on this interface are passed through the tunnel. The physical interfaces associated with the tunnel endpoints must be of the same type. Cisco Systems currently supports serial, POS, and Ethernet interfaces in raw mode.

Raw mode can be used effectively to support virtual leased lines. Virtual leased lines are a common requirement from Enterprises wishing to connect remote sites together over a clear channel service.

Figure 3 illustrates the function of L2TPv3 in providing this service.

Figure 3
Virtual Leased Line with L2TPv3 Logical Topology



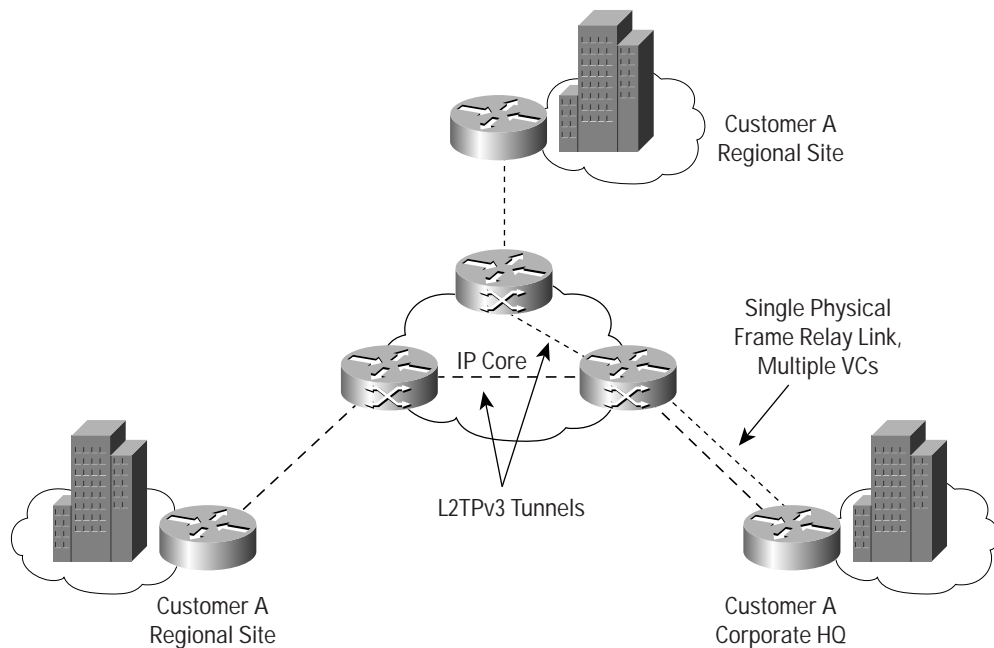
In this case, two DS-3 serial interfaces are connected to the customer's network (Enterprise A). *Int2* and *Int3* form the ingress and egress points of the L2TPv3 tunnel. The service provider maintains IP connectivity between *PE A* and *PE B* using standard routing protocols. This forms the fabric for the Layer 2 VPN to be established. Any packets being sent over the DS-3 from the customer's edge router (*CE A*) will be automatically encapsulated with an L2TPv3 header and forwarded across the IP network to the egress interface on *PE B* and decapsulated. Then the entire original High-Level Data Link Control (HDLC) frame is forwarded out of the serial interface (*Int3*) and on to the customer router *CE B*, thus completing the Layer 2 circuit emulation.



Frame Relay Support

Frame Relay support is designed to support the tunneling of individual virtual circuits (VCs) that are allocated on a single physical interface. In this scenario, the VC associated on a given interface can be tunneled to different destinations. In Figure 4, the ability to tunnel individual VCs is used to deploy a hub-and-spoke type network topology over an IP core.

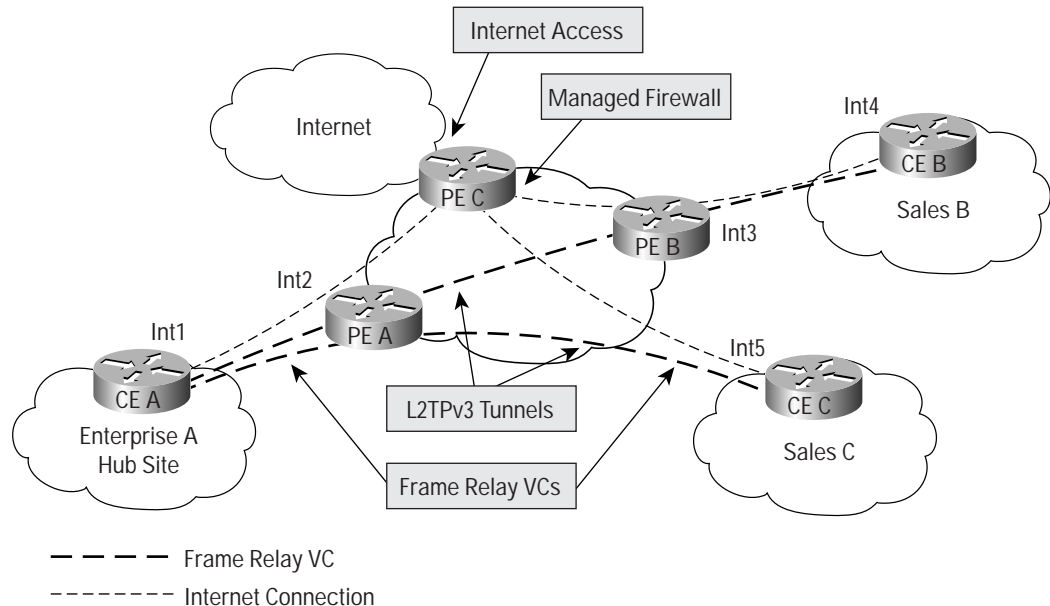
Figure 4
Packet Encapsulation with L2TPv3



The following advanced scenario outlines a service offering that a service provider may wish to support. The customer is an existing enterprise with a Frame Relay hub-and-spoke network. The service provider wishes to offer Internet access coupled with managed firewall and multimedia services. Figure 5 illustrates this service architecture.



Figure 5
Frame Relay Hub-and-Spoke Architecture with Outsourced Internet and Firewall



The enterprise connects to the service provider with a traditional serial interface configuring Frame Relay encapsulations and a subinterface point-to-point configuration. Designating which subinterfaces will be the corporate intranet and which will provide Internet access. This decentralizes Internet access in traditional hub-and-spoke configurations, thus reducing the bandwidth requirements at the hub site. The enterprise is free to run an autonomous routing policy and even add IP Security (IPSec) encryption for enhanced security requirements.

Frame Relay Subinterface Restrictions

- If a Frame Relay subinterface is configured for tunneling, it must be mapped to a unique L2TPv3 tunnel. (Each L2TPv3 tunnel must have one-to-one mapping with a Frame Relay subinterface.)
- The data-link connection identifier (DLCI) at the ingress router must be the same DLCI bound at the egress router.
- L2TPv3 Frame Relay subinterfaces support 10-bit DLCI addresses. Frame Relay extended addressing is not supported.
- Multipoint DLCIs are not supported.


Ethernet Support

L2TPv3 Operation in Cisco 10720 Internet Routers

Support of the L2TPv3 feature in Cisco 10720 Internet routers allows service providers to offer Ethernet services to their customers by extending their Ethernet or virtual LAN (VLAN) from one location to the other using the L2TPv3 tunnel.



Figure 6
L2TPv3 Operation on Cisco 10720 Internet Routers



In Figure 6, the two routers at the end of the L2TPv3 tunnel are connected via a point-to-point POS link. The functionality supported is Layer 2-to-Layer 2 extension over the L2TPv3 tunnel. Either the entire interface or the VLAN subinterface can be mapped to a L2TPv3 tunnel to extend Ethernet over the IP network. This mechanism allows service providers to offer Ethernet services over a wide area.

When L2TPv3 is used to connect customers across the IP backbone, the physical interface connecting to the customer's network becomes the tunnel ingress/egress interface. The interfaces on the Cisco 10720 Internet routers that can be used as an ingress or egress can be either Ethernet or 802.1Q-encapsulated subinterfaces.

The Internet service provider routers communicate normally using the IP routing protocols configured across the IP core network. The customer routers (CE routers) communicate across the configured L2TPv3 tunnels. Any data packet that is received by the Cisco 10720 Internet Router is encapsulated with a L2TPv3 header and sent via L2TPv3 tunnel, either L2TPv3 Tunnel 1 or L2TPv3 Tunnel 2. Data packets arriving from customer sites can travel over the main Ethernet interface. In this scenario, packets are encapsulated with the L2TPv3 header and sent to the other site.

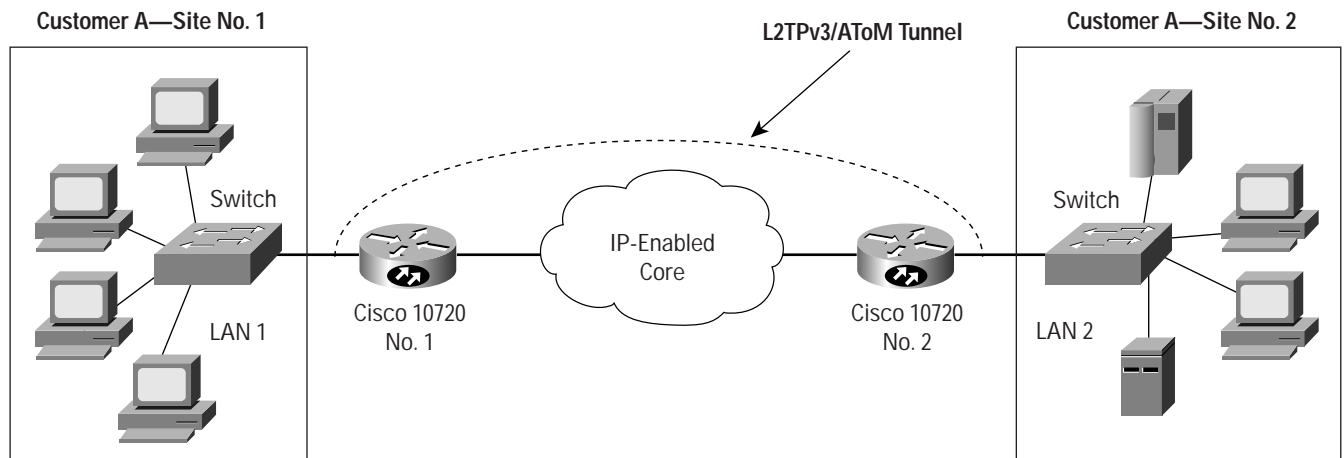
If customer data packets are coming in over an 802.1Q-encapsulated interface, packets will still be encapsulated with an L2TPv3 header and sent to the other Cisco 10720 Internet Router. On the receiving end, the Cisco 10720 Internet Router decapsulates the L2TPv3 header and forwards the encapsulated 802.1Q traffic to the other customer CE router.



Ethernet over L2TPv3

Figure 7 shows the operation of LAN-to-LAN connectivity using an L2TPv3 tunnel. Incoming packets from *LAN1* home in on the Ethernet interface of the Cisco 10720 Internet Router No. 1. Data packets from the customer are encapsulated with an L2TPv3 header and sent across the IP core backbone using IP routing to the egress Ethernet interface of Cisco 10720 Internet Router No. 2.

Figure 7
Ethernet over L2TPv3



The Cisco 10720 Internet Router on the receiving side decapsulates the data packets and forwards the traffic to LAN2. A similar operation occurs if traffic originates from LAN2. LAN1 and LAN2 become connected across the IP backbone network, and the Cisco 10720 Internet Router relays the Layer 2 data packets without the need to get involved with the customer in any routing or knowledge of the customer IP addresses. Both customer sites appear as if they are connected to the same wire. This can be extended to multiple sites if needed. The main difference in this case is that on the customer sites there are no Layer 3 routers, so L2TPv3 extends Layer 2 connectivity across the IP-enabled core for the service provider.

VLAN ID Rewrite at L2TPv3 VLAN Tunnel Egress

The VLAN ID Rewrite feature applies to L2TPv3 tunnels that are attached to 802.1Q VLAN interfaces on a Cisco 10720 Internet Router. The egress side of an L2TPv3 tunnel that is mapped to a VLAN rewrites the VLAN ID in outgoing 802.1Q packets to the ID of the local VLAN.

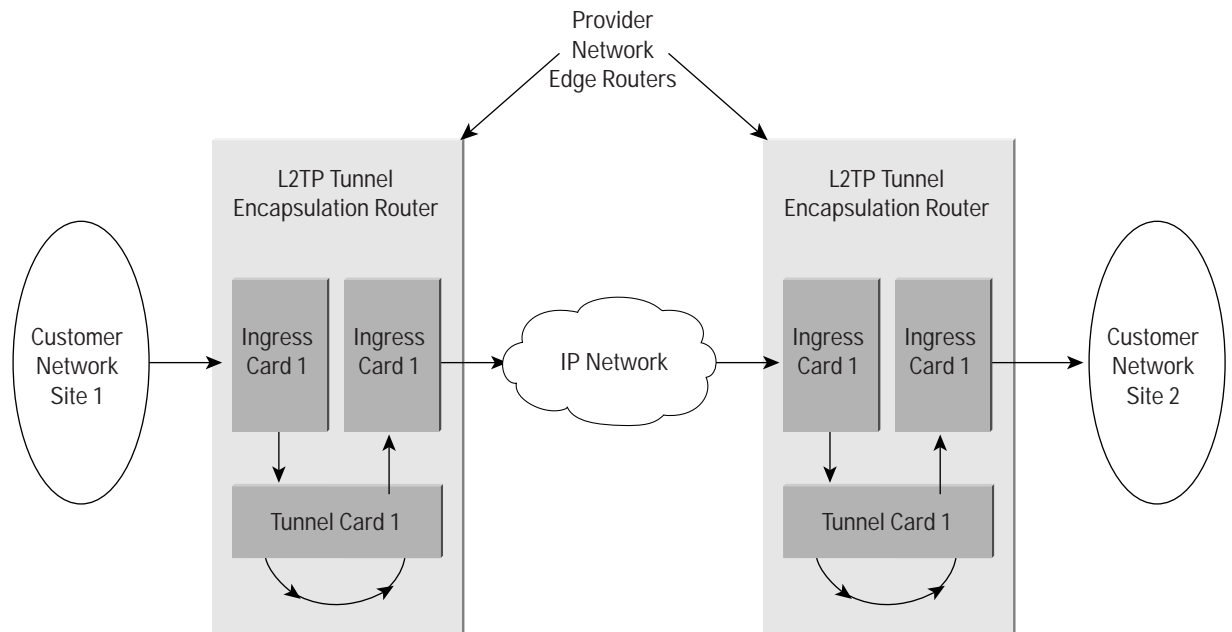
This feature allows you to use VLAN interfaces with different VLAN IDs at both ends of an L2TPv3 tunnel. When you use the VLAN ID Rewrite feature, Cisco recommends using the Spatial Reuse Protocol (SRP) as the backbone interface on the Cisco 10720 Internet Router.

Tunnel Cards in the Cisco 12000 Series Internet Router

The Cisco 12000 Series Internet routers require additional tunnel cards for L2TPv3 operation. Tunnel cards are not used with Cisco 7200 or 7500 Series routers.



Figure 8
L2TPv3 Packet Handling in the Cisco 12000 Series Internet Router



Note: The arrows in Figure 8 represent the flow and direction of a packet in one direction. The actual traffic in the tunnel can flow in either direction.

Actions on the Encapsulation Router

In Figure 8, traffic from the customer network on Site 1 is sent to an ingress interface on the provider network edge router. When the interface is configured for L2TPv3 tunneling, all arriving packets are forwarded to the tunnel card. The tunnel card encapsulates the packet with an encapsulation header containing the IP and L2TPv3 header information. The encapsulated packet is then sent to the appropriate egress card, which sends the packet to the IP network as a normal IP packet.

Actions on the Tunnel Decapsulation Router

When an encapsulated L2TPv3 packet arrives at the tunnel card, the packet is checked for a valid session ID and a matching L2TPv3 key. If any of the two are not correct, the packet is silently dropped. (The user is not notified.) If the session ID and L2TPv3 key are correct, the tunnel card decapsulates the packet (by removing the IP+L2TPv3 header) and sends the packet to the egress card. The egress card then sends the packet to the customer network. It does not add a new Layer 2 header. (The Layer 2 header is carried from the origin of the tunnel.)

Note: If the tunnel card receives non-L2TPv3 packets (other IP/Internet Control Message Protocol [ICMP] packets such as a ping “loopback address”), the packets are sent to the line card CPU and to the route processor.



General Limitations

L2TPv3 has these limitations:

- The maximum number of L2TPv3 tunnels that can be configured is limited to 1022 tunnels.
- The Maximum Transmission Unit (MTU) on the customer's interface must be set so that no fragmentation is required throughout the tunnel path. This is necessary because L2TPv3 tunnels do not support fragmentation. (A solution is being developed.)

The MTU in the IP backbone must be x bytes larger than the MTU that operates on the pseudo-wire. The value for x is:

- 802.1Q = 50
 - Ethernet = 46
 - POS = 36
 - Frame Relay = 34
 - CHDLC = 36
- No inherent signaling or keep-alive mechanisms (work is in progress).

Availability

Supported Platforms and Release

- Cisco 12000 Series Internet routers
- Cisco 10720 Metro Ethernet Router
- Cisco 7200 Series routers
- Cisco 7500 Series routers

Cisco IOS Software Release 12.0(18)ST

- Raw mode—L2TPv3 tunneling at port level; like interfaces on each end of tunnel
- Platforms—Cisco 12000, 7500, and 7200 Series routers

Cisco IOS Software Release 12.0(19)ST

- Frame Relay—L2TPv3 tunneling for Frame Relay point-to-point subinterfaces; each Frame Relay permanent virtual circuit (PVC) maps to a unique tunnel
- Platforms—Cisco 12000, 7500, and 7200 Series routers

Cisco IOS Software Release 12.0(21)ST

- 802.1Q VLAN—L2TPv3 tunneling for 802.1Q point-to-point subinterfaces
- Platforms—Cisco 12000, 7500, and 7200 Series routers

Cisco IOS Software Release 12.0(19)SP

- Raw mode—L2TPv3 tunneling at the port level; like interfaces on each end of the tunnel
- 802.1Q VLAN—L2TPv3 tunneling for 802.1Q point-to-point subinterfaces
- Platforms—Cisco 10720 Internet Router

Note: L2TPv3 tunnels are supported with the basic IP package.



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) 202777/ETMG 03/02