

Layer 2 Tunneling Protocol Version 3 Enables Layer 2 Services for IP Networks

The competitive environment for service providers has changed considerably since the Internet became a global force in the 1990s. Enterprises are no longer signing up for new IP-based services for the novelty or out of fear of being left behind by the competition. The challenge for service providers today is to grow their businesses by expanding their customer base and service revenue in a more cautious spending environment. Most enterprises are taking a more conservative approach to network investments. New IP-based services give enterprises an opportunity to improve their productivity and competitiveness while lowering their existing network expenses. Service providers that offer these services and savings can grow their customer base and service revenue. This white paper focuses on one such opportunity—offering traditional network integration over native IP network cores to lower the costs of maintaining separate traditional networks while adding IP-based services to enterprise customers.

Introduction

Enterprises and governments worldwide use traditional Layer 2 connection services. Services such as ATM, Frame Relay, and leased line provide the point-to-point connectivity upon which private networks are built.

Today's enterprise network managers have many questions and options to consider when implementing and operating the corporate intranet. Do they manage it internally or outsource it to the service provider? Should application-specific service levels be considered? What network security levels are required? What features can ever-shrinking budgets support?

Many enterprise customers use Layer 2 services such as ATM, Ethernet, Frame Relay, and leased line to interconnect the corporate intranet using their service provider. With the build-out of common packet-based infrastructure in the service provider core, these Layer 2 frames that now exist at the edge can be tunneled across the packet-switched network.

Cisco IOS[®] Software offers Layer 2 Tunneling Protocol (L2TP) Version 3, an Internet Engineering Task Force (IETF) standard track protocol that provides this capability. L2TPv3 helps enable service providers to deliver traditional Layer 2 services entirely from their IP infrastructures. It empowers services providers to:

- Drive down the cost of providing traditional Layer 2 services through superior cost efficiencies of multiservice IP infrastructures
- Extend their existing Layer 2 networks without expanding their legacy networks
- Lower the cost of providing multiple services to a customer site through service bundling

Offering a traditional Layer 2 service such as Frame Relay using an IP network infrastructure can lower the cost of providing the same service compared to offering the same service using a dedicated Layer 2 network. IP network infrastructures support multiple service types, and multiservice networks can spread network investments and operating costs across a larger and more diverse customer base. L2TPv3 also allows a service provider to extend the geographic reach of its traditional Layer 2 service to areas where its Layer 2 networks do not currently exist.

Traditional Layer 2 services can now be offered as far as the IP network can reach.


Using L2TPv3, service providers can now enhance their product portfolios to include managed Internet, intranet, and extranet services without adding complexity and expense. Customer equipment investments are protected as customers continue to connect to the service provider through their existing infrastructures.

Cisco Systems® has long been an innovator in internetworking technologies. Advanced hardware platforms running Cisco IOS Software are the primary building blocks for efficient, profitable networks. Cisco is committed to working with customers to develop the equipment, protocols, and support they require, with the technologies of their choice.

Layer 2 Tunneling Protocol Version 3 History

L2TPv3 combines an enhanced L2TPv2 control plane with an optimized 2 field header. L2TPv3 is designed to enable service providers with a native IP-based infrastructure to offer transparent LAN services to their customers (Figure 1). L2TPv3 includes support for multiple Layer 2 encapsulations, including 802.1Q virtual LAN (VLAN), Cisco High-Level Data Link Control (HDLC), Ethernet, Frame Relay, Packet over SONET (POS), and Point-to-Point Protocol (PPP) support. The enhancements allowed service subscribers to essentially connect two similar interfaces back-to-back without any knowledge of the underlying IP network used to deliver their frames.

Figure 1
Example of Transparent LAN Services



The L2TPv3 tunnel provides the transport to allow routers 3 and 4 to appear to be connected back-to-back with POS interfaces (interfaces 1 and 4). The POS interfaces will be completely unaware of the IP transport network being used to form this connection. For a detailed discussion of the protocol capability, see the “Layer 2 Tunneling Protocol Version 3 Overview” section later in this document.

Feature Compatibility

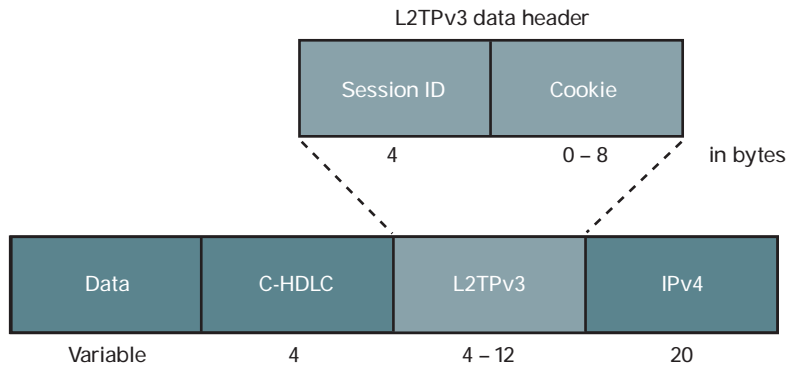
Because L2TPv3 tunnels use IP Version 4 as the delivery protocol, Cisco offers extensive support for popular features such as multicast, NetFlow, and IP-based quality of service (QoS). Using L2TPv3, service providers can use existing compatibility with IP Security (IPSec) for enhanced network security or to allow enterprise customers to manage their own security. Enterprises using multicast and QoS support for differing traffic definitions and priorities are unaffected.

Layer 2 Tunneling Protocol Version 3 Overview

L2TPv3 includes two distinct components or message types. The first is a control connection, a reliable, in-band connection between endpoints responsible for tunnel and session setup, teardown, and maintenance, and is facilitated through “control messages.” The second is a forwarding plane, responsible for the encapsulation of Layer 2 data to be forwarded over the IP network through “data messages.” Either component can be implemented independently.

When the control connection is implemented between a pair of provider edge routers this is referred to as the L2TP Control Connection Endpoint or LCCE. When the control connection is operational, it can negotiate session IDs and other requirements for circuits subject to Layer 2 transport. These are referred to as attachment circuits. After the session ID has been negotiated, it can be prepended to the Layer 2 datagram that is being transported (Figure 2).

Figure 2
L2TP Version 3 Encapsulation in an IP Version 4 Header



Note: The default Layer 2-specific sublayer is not depicted.

The session ID is a 32-bit locally significant field used to identify the call on the destination or egress tunnel endpoint. The session ID will be negotiated by the control connection or statically defined if using the L2TPv3 data plane only.

The cookie is a variable length (with a maximum of eight bytes), word-aligned optional field. The control connection can negotiate this as an additional level of guarantee beyond the regular session ID lookup to make sure that a data message has been directed to the correct session or that any recently reused session ID will not be misdirected.

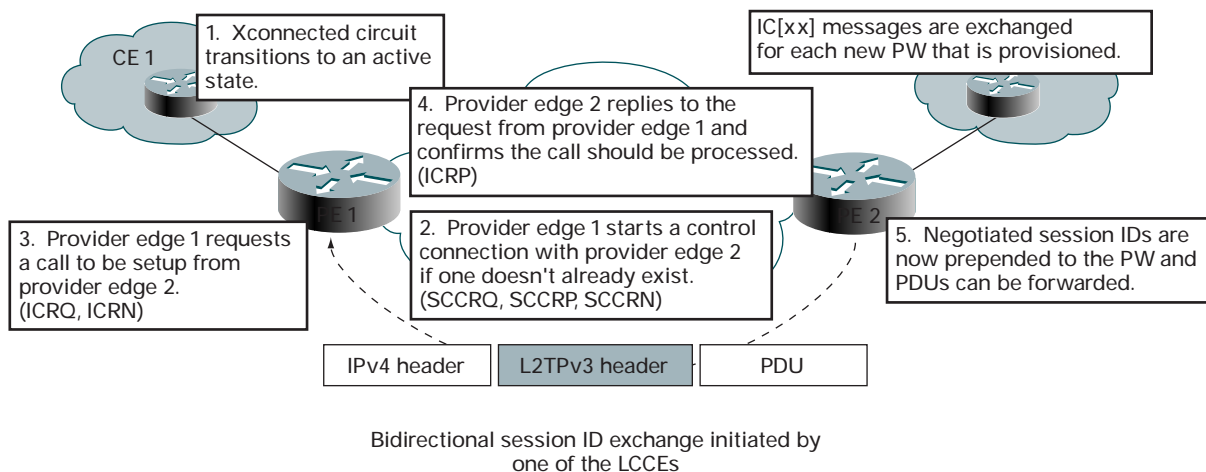
For a detailed description of L2TPv3 control channel operation, see the IETF draft for L2TP at

<http://www.ietf.org/html.charters/l2tpext-charter.html>

How Layer 2 Tunneling Protocol Works

This discussion will focus on the macro processes involved in creating an L2TPv3-based service. Figure 3 depicts the basic protocol operation.

Figure 3
Protocol Operational Overview



1. First, the customer connects to the service provider's edge router via a DS-3 serial interface and sets the encapsulation to High-Level Data Link Control (HDLC). No special configuration is required for the customer's edge router.

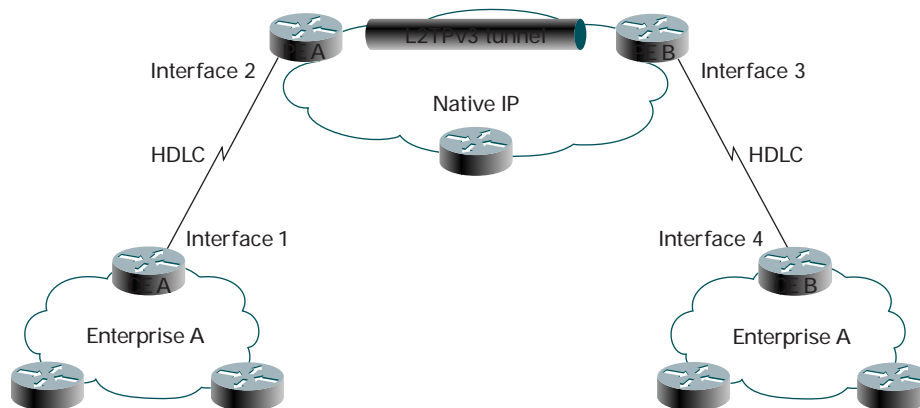
2. In the service provider network on the provider-edge router connecting to the customer site, an L2TPv3 tunnel is configured with the destination IP address of the peering provider-edge router where the customer's egress circuit is attached. This is accomplished using the XConnect command-line interface (CLI). The corresponding XConnect configuration is required on the remote provider-edge router.
3. At this point, L2TPv3 will determine if a control connection exists between the destination provider-edge router. If not, the provider-edge router will send out a Start-Control-Connection-Request message to initiate one. After the tunnel has been established, session negotiation can take place for any attachment circuits requiring Layer 2 transport services between the LCCEs.
4. Next, session ID and cookie values can be negotiated between the LCCEs to give each attachment circuit its unique identification for proper demultiplexing at the remote provider edge. This is a bidirectional process and the session ID is unique to the remote peer only. In the absence of the L2TPv3 control connection, static session IDs may be defined.
5. After session IDs have been successfully negotiated, data received on the ingress interface of the provider-edge router will be prepended with the remote provider-edge router's session ID and forwarded through the outer IP header's destination IP address.
6. Finally, the packet is received at the destination provider-edge router, the L2TPv3 header is demultiplexed based on the session ID and validated against the negotiated cookie value. If the header is valid, it is stripped and the original Layer 2 frame is forwarded through the associated physical port and on to the destination customer-edge router.

Layer 2 Tunneling Protocol Version 3 Applications

Virtual Leased Line

Virtual leased lines are a common requirement of the enterprise that wants to connect remote sites over a clear, dedicated bandwidth. Encapsulations typically employed are Cisco HDLC or PPP. Figure 4 illustrates the function of L2TPv3 in providing this service.

Figure 4
Virtual Leased Line



In this case, two DS-3 serial interfaces are connected to the customer's network (Enterprise A). Interface 2 and Interface 3 form the ingress and egress points of the L2TPv3 tunnel. The service provider maintains IP connectivity between provider-edge routers (PE A and PE B) using standard Interior Gateway Protocols (IGPs) such as Intermediate System-to-Intermediate System (IS-IS) Protocol or Open Shortest Path First (OSPF) Protocol. This forms the fabric for the Layer 2 VPN to be established. Any packets sent over the DS3 from the customer-edge router (CE A) would be automatically encapsulated with a L2TPv3 header and forwarded across the IP network to the egress interface on PE B and decapsulated. Then the entire original HDLC frame is forwarded out of the serial interface (Interface 3) and on to the customer-edge router CE B, thus completing the Layer 2 circuit emulation. Here is a typical configuration using the XConnect CLI:

XConnect CLI Example

PE_A:

```
interface Loopback0
  ip address 172.18.255.1 255.255.255.255
!
pseudowire-class L2TPv3_Default
  encapsulation l2tpv3
  sequencing both
  ip local interface Loopback0

!
interface Serial4/0
  no ip address
  encapsulation hdlc
  dsu bandwidth 44210
  framing c-bit
  cablelength 10
  xconnect 172.18.255.3 600 pw-class L2TPv3_Default
!
...
```

PE_B:


```
interface Loopback0
  ip address 172.18.255.3 255.255.255.255
!
pseudowire-class L2TPv3_Default
  encapsulation l2tpv3
  sequencing both
  ip local interface Loopback0

!
interface Serial4/0
  no ip address
  encapsulation hdlc
  dsu bandwidth 44210
  framing c-bit
  cablelength 10
  xconnect 172.18.255.1 600 pw-class L2TPv3_Default
!
...
```

Benefits of a Virtual Leased Line

- Service providers use common IP-packet infrastructure to offer the virtual leased-line service, thus expanding or consolidating existing services.
- Using port-level policing, a tiered pay-rate plan can be implemented, limiting the input a customer can send into the network and maximizing bandwidth utilization.

Figure 6
Connecting MPLS Networks Across an IP Core



In this scenario, the customer is an MPLS service provider connected to the IP transit provider via CE-PE A and CE-PE B. The IP transit provider treats this connection the same as if an enterprise customer were using the services. The MPLS service provider can use the IP service provider's network to send MPLS-based traffic. Because the IP network is transparent to the MPLS service provider, the IP network provider is able to forward Lightweight Directory Protocol (LDP) traffic and to perform normal label swapping as if the POS connections were directly connected. The entire HDLC frame is forwarded across the IP network.

Benefits of an MPLS Transit Network

- The IP network can be used as a transit network that links together separate MPLS networks or creates virtual Network Access Points (NAPs) without having to establish peering relationships or implement carrier-supporting technologies.
- L2TPv3 allows for a phased migration to an MPLS-enabled core.
- IP service providers don't have to administer or exchange routing information with customer service-provider networks.

Configuration Examples

For more L2TPv3 applications and configuration examples, visit:

<http://www.univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s24/l2tpv3.htm>

References

For more information on L2TP, visit: <http://www.ietf.org/home.html>

CISCO SYSTEMS



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 © 1992–2003 Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo 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. (0403R)