Layer 2 Tunnel Protocol Version 3 Software Configuration Guide for Cisco 1000 Series Connected Grid Routers (Cisco IOS)

January 2014

OL-31238-01

This document describes how to configure Layer 2 Tunnel Protocol Version 3 (L2TPv3) on the Cisco 1000 Series Connected Grid Router (hereafter referred to as the CGR 1000).

This document includes the following sections:

Information About L2TPv3

The L2TPv3 feature expands Cisco's support of Layer 2 VPNs. Layer 2 Tunneling Protocol Version 3 (L2TPv3) is an IETF l2tpext working group draft that provides several enhancements to L2TP to tunnel any Layer 2 payload over L2TP. Specifically, L2TPv3 defines the L2TP protocol for tunneling Layer 2 payloads over an IP core network by using Layer 2 VPNs.

Note The L2TPv3 feature on the CGR 1000 supports only Ethernet, Ethernet subinterface, and 802.1q (VLAN). Other Layer 2 technologies are not supported, so they are not described in this document.

This section includes the following topics:

Cisco Systems, Inc.
www.cisco.com

Benefits

Benefits of this feature include the following:

- Simplifies deployment of VPNs
- Does not require Multiprotocol Label Switching (MPLS)
- Supports Layer 2 tunneling over IP for Ethernet and 802.1q (VLAN)
- Provides cookies for authentication
- Provides session state updates and multiple sessions
- Supports interworking (Ethernet-VLAN, Ethernet-QinQ, and VLAN-QinQ)

L2TPv3 Operation

L2TPv3 provides a method for delivering L2TP services over an IPv4 (non-UDP) backbone network. It encompasses the signaling protocol as well as the packet encapsulation specification. L2TPv3 provides the following features:

- Xconnect for Layer 2 tunneling through a pseudowire over an IP network
- Layer 2 VPNs for provider edge (PE)-to-PE router service using xconnect that supports Ethernet and 802.1q (VLAN) Layer 2 circuits, including dynamic forwarded sessions

The figure below shows how you can use the L2TPv3 feature for setting up VPNs using Layer 2 tunneling over an IP network. All traffic between two customer network sites is encapsulated in IP packets carrying L2TP data messages and sent across an IP network. The backbone routers of the IP network treat the traffic as any other IP traffic and need not know anything about the customer networks.

Figure 1 L2TPv3 Operation Example

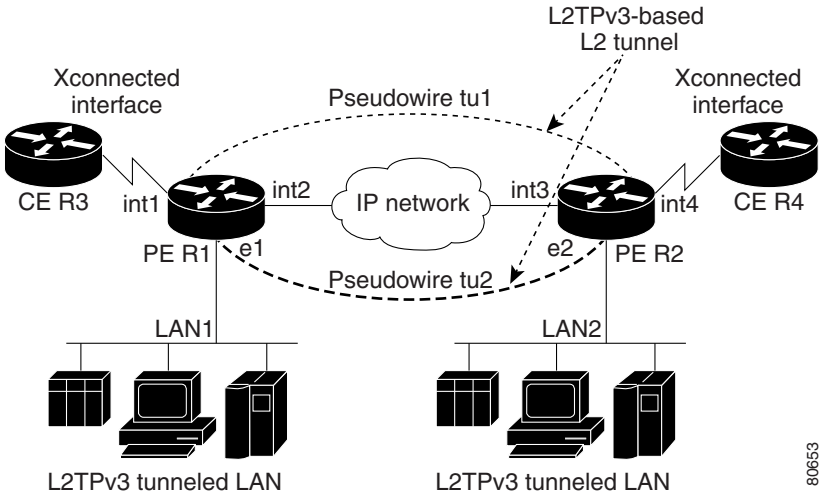

In the figure above, the PE routers R1 and R2 provide L2TPv3 services. The R1 and R2 routers communicate with each other using a pseudowire over the IP backbone network through a path comprising the interfaces int1 and int2, the IP network, and interfaces int3 and int4.

In this example, the customer edge (CE) routers R3 and R4 communicate through a pair of xconnect Ethernet or VLAN interfaces using an L2TPv3 session. The L2TPv3 session tu1 is a pseudowire configured between interface int1 on R1 and interface int4 on R2. Any packet arriving on interface int1 on R1 is encapsulated and sent through the pseudowire control channel (tu1) to R2. R2 decapsulates the packet and sends 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 features regarding L2TPv3 operation:

- All packets received on interface int1 are forwarded to R4. R3 and R4 cannot detect the intervening network.
- For Ethernet interfaces, any packet received from LAN1 by R1 on Ethernet interface e1 are encapsulated directly in IP and sent through the pseudowire session tu2 to R2 interface e2, where it is sent on LAN2.
- A VLAN on an Ethernet interface can be mapped to an L2TPv3 session.

L2TPv3 Header Description

The L2TPv3 header has the format shown in the figure below.

Figure 2 L2TPv3 Header Format

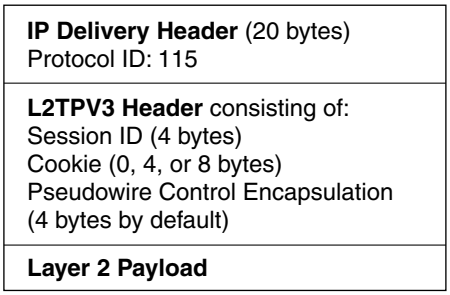

Each L2TPv3 packet contains an L2TPv3 header that includes a unique session ID representing one session and a variable cookie length. The L2TPv3 session ID and the Tunnel Cookie field length are assigned through the CLI. See the section "Configuring L2TPv3" section on page 9 for more information on the CLI commands for L2TPv3.

Session ID

The L2TPv3 session ID identifies the session context on the decapsulating system. For dynamic sessions, the value of the session ID is selected to optimize the context identification efficiency of the decapsulating system. A decapsulation implementation may therefore elect to support a smaller session ID bit field. In this L2TPv3 implementation, an upper value for the L2TPv3 session ID was set at 023. The L2TPv3 session ID value 0 is reserved for use by the protocol.

Note The local session ID must be unique on the decapsulating system and is restricted to the least significant ten bits.

Session Cookie

The L2TPv3 header contains a control channel cookie field that has a variable length of 0, 4, or 8 bytes according to the cookie length supported by a given platform for packet decapsulation. The control channel cookie length is dynamically determined for dynamic sessions.

The variable cookie length does not present a problem when the same platform is at both ends of an L2TPv3 control channel. However, when different platforms interoperate across an L2TPv3 control channel, both platforms need to encapsulate packets with a 4-byte cookie length.

Pseudowire Control Encapsulation

The L2TPv3 pseudowire control encapsulation consists of 32 bits (4 bytes) and contains information used to sequence L2TP packets. For the purposes of sequencing, only the first bit and bits 8 to 31 are relevant. Bit 1 indicates whether the Sequence Number field, bits 8 to 31, contains a valid sequence number and is to be updated.

L2TPv3 Features

L2TPv3 on the CGR 1000 provides xconnect support for Ethernet and 802.1q (VLAN).

For details about L2TPv3 features, see the following sections in the chapter Layer 2 Tunneling Protocol Version 3 in the Wide-Area Networking Configuration Guide: Layer 2 Services, Cisco IOS Release 15M&T:

- Control Channel Parameters
- L2TPv3 Control Channel Authentication Parameters
- Dynamic L2TPv3 Sessions
- Sequencing
- Local Switching
- L2TPv3 Type of Service Marking
- Keepalive
- L2TPv3 Control Message Hashing
- L2TPv3 Control Message Rate Limiting
- L2TPv3 Digest Secret Graceful Switchover
- L2TPv3 Pseudowire
- Manual Clearing of L2TPv3 Tunnels
- L2TPv3 Tunnel Management
- L2TPv3 Protocol Demultiplexing
- Color Aware Policer on Ethernet over L2TPv3
- Site of Origin for Border Gateway Protocol VPNs
- L2TPv3 Custom Ethertype for Dot1q and QinQ Encapsulations

Note L2TPv3 Layer 2 Fragmentation is not supported on the CGR 1000. Therefore, fragmentation-related commands such as **ip dfbit set** and **ip pmtu** are not supported.

Supported L2TPv3 Payloads

The L2TPv3 feature on the CGR 1000 supports the following payloads:

- Ethernet
- VLAN
- IPv6

Note Each L2TPv3 tunneled packet includes the entire Layer 2 frame of the payloads described in this section. If sequencing is required (see the Sequencing section), a Layer 2-specific sublayer (see the "Pseudowire Control Encapsulation" section on page 4) is included in the L2TPv3 header to provide the Sequence Number field.

Ethernet

An Ethernet frame arriving at a PE device is simply encapsulated in its entirety with an L2TP data header. At the other end, a received L2TP data packet is stripped of its L2TP data header. The payload, an Ethernet frame, is then forwarded to the appropriate attachment circuit.

Because the L2TPv3 tunneling protocol serves essentially as a bridge, it need not examine any part of an Ethernet frame. Any Ethernet frame received on an interface is tunneled, and any L2TP-tunneled Ethernet frame is forwarded out of the interface.

Note Because of the way in which L2TPv3 handles Ethernet frames, an Ethernet interface must be configured to promiscuous mode to capture all traffic received on the Ethernet segment attached to the device. All frames are tunneled through the L2TP pseudowire.

VLAN

L2TPv3 supports VLAN memberships in the following ways:

- Port-based, in which undated Ethernet frames are received
- VLAN-based, in which tagged Ethernet frames are received

In L2TPv3, xconnect Ethernet supports port-based VLAN membership and the reception of tagged Ethernet frames. A tagged Ethernet frame contains a tag header (defined in 802.1Q), which is 4 bytes long and consists of a 2-byte tag protocol identifier (TPID) field and a 2-byte tag control information (TCI) field. The TPID indicates that a TCI follows. The TCI is further broken down into the following three fields:

- User priority field
- Canonical format indicator (CFI)
- A 12-bit VLAN ID (VID)

For L2TPv3, an Ethernet subinterface configured to support VLAN switching can be bound to an xconnect service so that all Ethernet traffic, tagged with a VID specified on the subinterface, is tunneled to another PE. The VLAN Ethernet frames are forwarded in their entirety. The receiving PE can rewrite the VID of the tunneled traffic to another value before forwarding the traffic onto an attachment circuit.

To successfully rewrite VLANs, it may be necessary to disable the Spanning Tree Protocol (STP). This can be done on a per-VLAN basis by using the **no spanning-tree vlan** command.

Note Because of the way in which L2TPv3 handles VLAN packets, the Ethernet interface must be configured in promiscuous mode to capture all traffic received on the Ethernet segment attached to the device. All frames are tunneled through the L2TP pseudowire.

IPv6 Protocol Demultiplexing

The Protocol Demultiplexing for L2TPv3 feature introduces the ability to provide native IPv6 support by setting up a specialized IPv6 network and offloading IPv6 traffic from the IPv4 network. IPv6 traffic is tunneled transparently to the IPv6 network using L2TPv3 pseudowires without affecting the configuration of the CE devices. IPv4 traffic is routed as usual within the IPv4 network, maintaining the existing performance and reliability of the IPv4 network.

Figure 3 shows a network deployment that offloads IPv6 traffic from the IPv4 network to a specialized IPv6 network. The PE devices demultiplex the IPv6 traffic from the IPv4 traffic. IPv6 traffic is routed to the IPv6 network over an L2TPv3 pseudowire, while IPv4 traffic is routed normally. The IPv4 PE devices must be configured to demultiplex the incoming IPv6 traffic from the IPv4 traffic. The PE devices facing the IPv6 network do not require the IPv6 configuration.

Figure 3 Protocol Demultiplexing of IPv6 Traffic from IPv4 Traffic

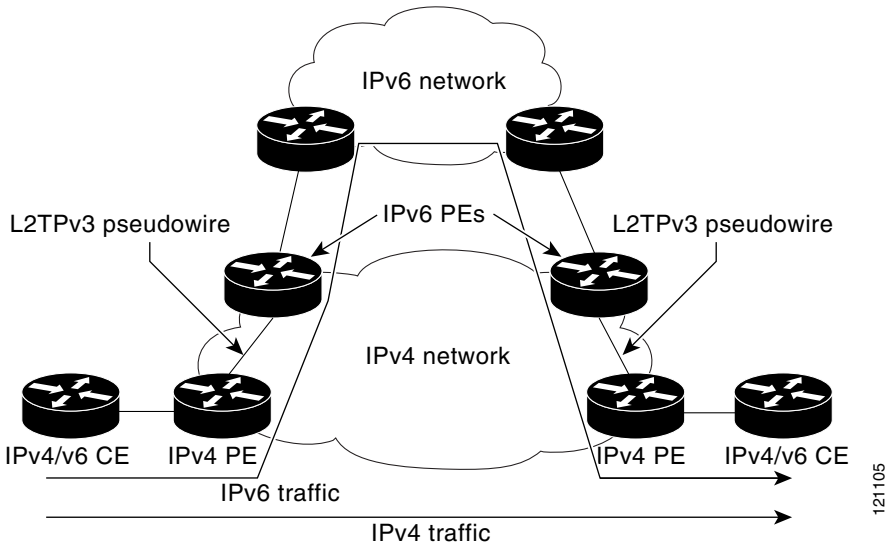

If no IP address is configured, the protocol demultiplexing configuration is rejected. If an IP address is configured, the **xconnect** command configuration is rejected unless protocol demultiplexing is enabled in xconnect configuration mode before exiting that mode. If an IP address is configured with an **xconnect** command configuration and protocol demultiplexing is enabled, the IP address cannot be removed. To change or remove the configured IP address, the **xconnect** command configuration must first be disabled.

Table 1 shows the valid combinations of configurations.

Table 1 Valid Configuration Scenarios

Scenario	IP Address	xconnect Configuration	Protocol Demultiplexing Configuration
Routing	Yes	No	–
L2VPN	No	Yes	No
IPv6 Protocol Demultiplexing	Yes	Yes	Yes

Prerequisites

- Before you configure an xconnect attachment circuit for a provider edge (PE) device (see Configuring the Xconnect Attachment Circuit), you must enable the Cisco Express Forwarding (formerly known as CEF) feature. To enable Cisco Express Forwarding on an interface, use the **ip cef** or **ip cef distributed** command.
- You must configure a loopback interface on the router for originating and terminating the L2TPv3 traffic. The loopback interface must have an IP address that is reachable from the remote PE device at the other end of an L2TPv3 control channel.

Guidelines and Limitations

General L2TPv3 Restrictions

- Cisco Express Forwarding must be enabled for the L2TPv3 feature to function. The xconnect configuration mode is blocked until Cisco Express Forwarding is enabled. On distributed platforms, such as the Cisco 7500 series, if Cisco Express Forwarding is disabled while a session is established, the session is torn down. The session remains down until Cisco Express Forwarding is reenabled. To enable Cisco Express Forwarding, use the **ip cef** or **ip cef distributed** command.
- The IP local interface must be a loopback interface. Configuring any other interface with the **ip local interface** command results in a nonoperational setting.
- The number of sessions on PPP, High-Level Data Link Control (HDLC), Ethernet, or 802.1q VLAN ports is limited by the number of interface descriptor blocks (IDBs) that the router can support. For PPP, HDLC, Ethernet, and 802.1q VLAN circuit types, an IDB is required for each circuit.
- The interface keepalive feature is automatically disabled on the interface to which xconnect is applied.
- The CGR 1000 does not support fragmentation of IP packets through L2TPV3 tunnels. Therefore, fragmentation-related commands such as **ip dfbit set** and **ip pmtu** are not supported.

VLAN-Specific Restrictions

- A PE device is responsible only for static VLAN membership entries that are configured manually on the device. Dynamic VLAN membership entries, entry aging, and membership discovery are not supported.
- Implicit tagging for VLAN memberships operating on other layers, such as membership by MAC address, protocol type at Layer 2, or membership by IP subnet at Layer 3, is not supported.
- Point-to-multipoint and multipoint-to-point configurations are not supported. There is a 1:1 relationship between an attachment circuit and an L2TPv3 session.

Default Settings

Feature	Default Setting
L2TP Control Channel Parameters	
Receive window size	Upper limit
Initial retries	2
Retries	15
Timeout	Default maximum interval: 8 Default minimum interval: 1
Timeout setup	300
Sequencing	Disabled
Authentication for the L2TP Control Channel	
Input format of shared secret	0 (plain-text secret)
L2TPv3 Control Message Hashing	

Feature	Default Setting
Input format of shared secret	0 (plain-text secret)
Hash function	md5
Validation of the message digest	Enabled
AV pair hiding	Disabled
L2TP Control Channel Maintenance Parameters	
Hello interval	60
L2TPv3 Pseudowire	
Protocol	l2tpv3
ToS byte value	0
TTL byte value	255

Configuring L2TPv3

To configure L2TPv3 on the CGR 1000, complete the following procedures in the Layer 2 Tunneling Protocol Version 3 chapter in the Wide-Area Networking Configuration Guide: Layer 2 Services, Cisco IOS Release 15M&T:

- Configuring L2TP Control Channel Parameters
  - Configuring L2TP Control Channel Timing Parameters
  - Configuring L2TPv3 Control Channel Authentication Parameters
  - Configuring L2TP Control Channel Maintenance Parameters
- Configuring the L2TPv3 Pseudowire
- Configuring the Xconnect Attachment Circuit
- Manually Configuring L2TPv3 Session Parameters
- Configuring Protocol Demultiplexing for L2TPv3
- Configuring an L2TPv3 Custom Ethertype for Dot1q and QinQ Encapsulations
- Manually Clearing L2TPv3 Tunnels

Verifying Configuration

Command	Purpose
show l2tun session brief	Displays information about current L2TPv3 sessions on a router.
show l2tun session all	Displays detailed information about current L2TPv3 sessions on a router.

Command	Purpose
show l2tun tunnel	Displays information about the L2TP control channels that are set up to other L2TP-enabled devices for all L2TP sessions on the router.
show l2tun tunnel all	Displays detailed information about the L2TP control channels that are set up to other L2TP-enabled devices for all L2TP sessions on the router.

Configuration Example

See the following sections in the Layer 2 Tunneling Protocol Version 3 chapter in the Wide-Area Networking Configuration Guide: Layer 2 Services, Cisco IOS Release 15M&T for configuration examples:

- Example: Configuring a Negotiated L2TPv3 Session for an Xconnect VLAN Subinterface
- Example: Verifying an L2TPv3 Session
- Example: Verifying an L2TP Control Channel
- Example: Configuring L2TPv3 Control Channel Authentication
- Example: Configuring L2TPv3 Digest Secret Graceful Switchover
- Example: Verifying L2TPv3 Digest Secret Graceful Switchover
- Configuring Protocol Demultiplexing for L2TPv3 Examples
- Example: Manually Clearing an L2TPv3 Tunnel
- Example: Configuring an L2TPv3 Custom Ethertype for Dot1q and QinQ Encapsulations

Related Documents

- Layer 2 Tunneling Protocol Version 3 chapter in the Wide-Area Networking Configuration Guide: Layer 2 Services, Cisco IOS Release 15M&T

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

This document is to be used in conjunction with the documents listed in the "Related Documents" section.