

Network Working Group  
Internet-Draft  
Category: Standards Track  
<[draft-ietf-l2tpext-l2tp-base-11.txt](#)>

J. Lau  
M. Townsley  
cisco Systems  
I. Goyret  
Lucent Technologies  
Editors  
October 2003

## Layer Two Tunneling Protocol (Version 3)

### Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

### Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

### Abstract

This document describes "version 3" of the Layer Two Tunneling Protocol (L2TPv3). L2TPv3 defines the base control protocol and encapsulation for tunneling multiple layer 2 connections between two IP connected nodes. Additional documents detail the specifics for each link-type being emulated.

INTERNET DRAFT

L2TPv3

October 2003

## Contents

Status of this Memo.....	<a href="#">1</a>
<a href="#">1.</a> Introduction.....	<a href="#">4</a>
<a href="#">1.1</a> Changes from <a href="#">RFC 2661</a> .....	<a href="#">5</a>
<a href="#">1.2</a> Specification of Requirements.....	<a href="#">5</a>
<a href="#">1.3</a> Terminology.....	<a href="#">5</a>
<a href="#">2.</a> Topology.....	<a href="#">9</a>
<a href="#">3.</a> Protocol Overview.....	<a href="#">10</a>
<a href="#">3.1</a> Control Message Types.....	<a href="#">11</a>
<a href="#">3.2</a> L2TP Header Formats.....	<a href="#">12</a>
<a href="#">3.2.1</a> L2TP Control Message Header.....	<a href="#">12</a>
<a href="#">3.2.2</a> L2TP Data Message.....	<a href="#">13</a>
<a href="#">3.3</a> Control Connection Management.....	<a href="#">14</a>
<a href="#">3.3.1</a> Control Connection Establishment.....	<a href="#">15</a>
<a href="#">3.3.2</a> Control Connection Teardown.....	<a href="#">15</a>
<a href="#">3.4</a> Session Management.....	<a href="#">16</a>
<a href="#">3.4.1</a> Session Establishment for an Incoming Call.....	<a href="#">16</a>
<a href="#">3.4.2</a> Session Establishment for an Outgoing Call.....	<a href="#">16</a>
<a href="#">3.4.3</a> Session Teardown.....	<a href="#">16</a>
<a href="#">4.</a> Protocol Operation.....	<a href="#">17</a>
<a href="#">4.1</a> L2TP Over Specific Packet-Switched Networks (PSN)....	<a href="#">17</a>
<a href="#">4.1.1</a> L2TPv3 over IP.....	<a href="#">18</a>
<a href="#">4.1.2</a> L2TP over UDP.....	<a href="#">19</a>
<a href="#">4.1.3</a> IP Fragmentation Issues.....	<a href="#">21</a>
<a href="#">4.2</a> Reliable Delivery of Control Messages.....	<a href="#">21</a>
4.3 Control Connection and Control Message Authentication	23
<a href="#">4.4</a> Keepalive (Hello).....	<a href="#">24</a>
<a href="#">4.5</a> Forwarding Session Data Frames.....	<a href="#">25</a>
<a href="#">4.6</a> Default L2-Specific Sublayer.....	<a href="#">25</a>
<a href="#">4.6.1</a> Sequencing Data Packets.....	<a href="#">26</a>
<a href="#">4.7</a> L2TPv2/v3 Interoperability and Migration.....	<a href="#">27</a>
<a href="#">4.7.1</a> L2TPv3 over IP.....	<a href="#">27</a>
<a href="#">4.7.2</a> L2TPv3 over UDP.....	<a href="#">27</a>
<a href="#">4.7.3</a> Automatic L2TPv2 Fallback.....	<a href="#">28</a>
<a href="#">5.</a> Control Message Attribute Value Pairs.....	<a href="#">28</a>
<a href="#">5.1</a> AVP Format.....	<a href="#">28</a>
<a href="#">5.2</a> Mandatory AVPs and Setting the M Bit.....	<a href="#">30</a>

<a href="#">5.3</a>	Hiding of AVP Attribute Values.....	<a href="#">31</a>
<a href="#">5.4</a>	AVP Summary.....	<a href="#">33</a>
<a href="#">5.4.1</a>	General Control Message AVPs.....	<a href="#">33</a>
<a href="#">5.4.2</a>	Result and Error Codes.....	<a href="#">37</a>
<a href="#">5.4.3</a>	Control Connection Management AVPs.....	<a href="#">40</a>

<a href="#">5.4.4</a>	Session Management AVPs.....	<a href="#">45</a>
<a href="#">5.4.5</a>	Circuit Status AVPs.....	<a href="#">53</a>
<a href="#">6.</a>	Control Connection Protocol Specification.....	<a href="#">56</a>
<a href="#">6.1</a>	Start-Control-Connection-Request (SCCRQ).....	<a href="#">56</a>
<a href="#">6.2</a>	Start-Control-Connection-Reply (SCCRP).....	<a href="#">56</a>
<a href="#">6.3</a>	Start-Control-Connection-Connected (SCCCN).....	<a href="#">57</a>
<a href="#">6.4</a>	Stop-Control-Connection-Notification (StopCCN).....	<a href="#">57</a>
<a href="#">6.5</a>	Hello (HELLO).....	<a href="#">58</a>
<a href="#">6.6</a>	Incoming-Call-Request (ICRQ).....	<a href="#">58</a>
<a href="#">6.7</a>	Incoming-Call-Reply (ICRP).....	<a href="#">59</a>
<a href="#">6.8</a>	Incoming-Call-Connected (ICCN).....	<a href="#">59</a>
<a href="#">6.9</a>	Outgoing-Call-Request (OCRQ).....	<a href="#">60</a>
<a href="#">6.10</a>	Outgoing-Call-Reply (OCRP).....	<a href="#">61</a>
<a href="#">6.11</a>	Outgoing-Call-Connected (OCCN).....	<a href="#">61</a>
<a href="#">6.12</a>	Call-Disconnect-Notify (CDN).....	<a href="#">62</a>
<a href="#">6.13</a>	WAN-Error-Notify (WEN).....	<a href="#">62</a>
<a href="#">6.14</a>	Set-Link-Info (SLI).....	<a href="#">63</a>
<a href="#">6.15</a>	Explicit-Acknowledgement (ACK).....	<a href="#">63</a>
<a href="#">7.</a>	Control Connection State Machines.....	<a href="#">64</a>
<a href="#">7.1</a>	Malformed Control Messages.....	<a href="#">64</a>
<a href="#">7.2</a>	Control Connection States.....	<a href="#">65</a>
<a href="#">7.3</a>	Incoming Calls.....	<a href="#">67</a>
<a href="#">7.3.1</a>	ICRQ Sender States.....	<a href="#">68</a>
<a href="#">7.3.2</a>	ICRQ Recipient States.....	<a href="#">69</a>
<a href="#">7.4</a>	Outgoing Calls.....	<a href="#">70</a>
<a href="#">7.4.1</a>	OCRQ Sender States.....	<a href="#">71</a>
<a href="#">7.4.2</a>	OCRQ Recipient (LAC) States.....	<a href="#">72</a>
<a href="#">7.5</a>	Termination of a Control Connection.....	<a href="#">73</a>
<a href="#">8.</a>	Security Considerations.....	<a href="#">74</a>
<a href="#">8.1</a>	Control Connection Endpoint and Message Security.....	<a href="#">74</a>
<a href="#">8.2</a>	Data Channel Security.....	<a href="#">74</a>
<a href="#">8.3</a>	End-to-End Security.....	<a href="#">75</a>
<a href="#">8.4</a>	L2TP and IPsec.....	<a href="#">75</a>

<a href="#">8.5</a>	Impact of L2TPv3 Features on <a href="#">RFC 3193</a> .....	<a href="#">75</a>
<a href="#">9.</a>	Internationalization Considerations.....	<a href="#">76</a>
<a href="#">10.</a>	IANA Considerations.....	<a href="#">76</a>
<a href="#">10.1</a>	Control Message Attribute Value Pairs (AVPs).....	<a href="#">76</a>
<a href="#">10.2</a>	Message Type AVP Values.....	<a href="#">77</a>
<a href="#">10.3</a>	Result Code AVP Values.....	<a href="#">77</a>
<a href="#">10.3.2</a>	Error Code Field Values.....	<a href="#">77</a>
<a href="#">10.4</a>	AVP Header Bits.....	<a href="#">77</a>
<a href="#">10.5</a>	L2TP Control Message Header Bits.....	<a href="#">77</a>
<a href="#">10.6</a>	Pseudowire Types.....	<a href="#">78</a>

INTERNET DRAFT

L2TPv3

October 2003

<a href="#">10.7</a>	Application Code.....	<a href="#">78</a>
<a href="#">10.8</a>	Circuit Status Bits.....	<a href="#">78</a>
<a href="#">10.9</a>	Default L2-Specific Sublayer bits.....	<a href="#">78</a>
<a href="#">10.10</a>	L2-Specific Sublayer Type.....	<a href="#">78</a>
<a href="#">10.11</a>	Data Sequencing Level.....	<a href="#">79</a>
<a href="#">13.</a>	Acknowledgments.....	<a href="#">79</a>
<a href="#">11.</a>	References.....	<a href="#">80</a>
<a href="#">11.1</a>	Normative References.....	<a href="#">80</a>
<a href="#">11.2</a>	Informative References.....	<a href="#">81</a>
<a href="#">12.</a>	Editors' Addresses.....	<a href="#">82</a>
<a href="#">Appendix A:</a>	Control Slow Start and Congestion Avoidance.....	<a href="#">82</a>
<a href="#">Appendix B:</a>	Control Message Examples.....	<a href="#">83</a>
<a href="#">Appendix C:</a>	Processing Sequence Numbers.....	<a href="#">85</a>
<a href="#">Appendix D:</a>	Intellectual Property Notice.....	<a href="#">87</a>
<a href="#">Appendix E:</a>	Full Copyright Statement.....	<a href="#">87</a>

## [1.](#) Introduction

The Layer Two Tunneling Protocol (L2TP) provides a dynamic mechanism for tunneling Layer 2 (L2) "circuits" across a packet-oriented data

network (e.g., over IP). L2TP, as originally defined in [RFC 2661](#), is a standard method for tunneling Point to Point Protocol (PPP) [\[RFC1661\]](#) sessions. L2TP has since been adopted for tunneling a number of other L2 protocols. In order to provide greater modularity, this document describes the base L2TP protocol, independent of the L2 payload that is being tunneled.

The base L2TP protocol defined in this document consists of (1) the control protocol for dynamic creation, maintenance, and teardown of L2TP sessions, and (2) the L2TP data encapsulation to multiplex and demultiplex L2 data streams between two L2TP nodes across an IP network. Additional documents are expected to be published for each layer 2 data link emulation type (a.k.a. pseudowire-type) supported by L2TP (i.e., PPP, Ethernet, Frame Relay, etc.). These documents will contain any individual details that are outside the scope of this base specification.

## [1.1](#) Changes from [RFC 2661](#)

Many of the protocol constructs described in this document are carried over from [RFC 2661](#). Changes include clarifications based on years of interoperability and deployment experience as well as modifications to either improve protocol operation or provide a clearer separation from PPP. The intent of these modifications is to achieve a healthy balance between code reuse, interoperability experience, and a directed evolution of L2TP as it is applied to new tasks.

When the designation between L2TPv2 and L2TPv3 is necessary, L2TP as defined in [RFC 2661](#) will be referred to as "L2TPv2", corresponding to the value in the Version field of an L2TP header. (Layer 2 Forwarding, L2F, [\[RFC2341\]](#) was defined as "version 1".) At times, L2TP as defined in this document will be referred to as "L2TPv3". Otherwise, the acronym "L2TP" will refer to L2TPv3 or L2TP in general.

Notable differences between L2TPv2 and L2TPv3 include:

Separation of all PPP-related AVPs, references, etc., including a

portion of the L2TP data header that was specific to the needs of PPP. The PPP-specific constructs are described in a companion document.

Transition from a 16-bit Session ID and Tunnel ID to a 32-bit Session ID and Control Connection ID, respectively.

Extension of the Tunnel Authentication mechanism to cover the entire control message rather than just a portion of certain messages.

Details of these changes and a recommendation for transitioning to L2TPv3 may be found in [Section 4.7](#).

## [1.2](#) Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

## [1.3](#) Terminology

Attribute Value Pair (AVP)

The variable-length concatenation of a unique Attribute

(represented by an integer), a length field, and a Value containing the actual value identified by the attribute. Zero or more AVPs make up the body of control messages, which are used in the establishment, maintenance, and teardown of control connections. This basic construct is sometimes referred to as a Type-Length-Value (TLV) in some specifications. (See also: Control Connection, Control Message.)

Call (Circuit Up)

The action of transitioning a circuit on an L2TP Access Concentrator (LAC) to an "up" or "active" state. A call may be dynamically established through signaling properties (e.g., an incoming or outgoing call through the Public Switched Telephone Network (PSTN)) or statically configured (e.g., provisioning a

Virtual Circuit on an interface). A call is defined by its properties (e.g., type of call, called number, etc.) and its data traffic. (See also: Circuit, Session, Incoming Call, Outgoing Call, Outgoing Call Request.)

## Circuit

A general term identifying any one of a wide range of L2 connections. A circuit may be virtual in nature (e.g., an ATM PVC, an ethernet VLAN, or an L2TP session), or it may have direct correlation to a physical layer (e.g., an RS-232 serial line). Circuits may be statically configured with a relatively long-lived uptime, or dynamically established with signaling to govern the establishment, maintenance, and teardown of the circuit. For the purposes of this document, a statically configured circuit is considered to be essentially the same as a very simple, long-lived, dynamic circuit. (See also: Call, Remote System.)

## Client

(See Remote System.)

## Control Connection

An L2TP control connection is a reliable control channel that is used to establish, maintain, and release individual L2TP sessions as well as the control connection itself. (See also: Control Message, Data Channel.)

## Control Message

An L2TP message used by the control connection. (See also: Control Connection.)

## Data Message

Message used by the data channel. (a.k.a. Data Packet, See also: Data Channel.)

## Data Channel

The channel for L2TP-encapsulated data traffic that passes between

two LCCEs over a Packet Switched Network (i.e. IP). (See also: Control Connection, Data Message.)

## Incoming Call

The action of receiving a call (circuit up event) on an LAC. The call may have been placed by a remote system (e.g., a phone call over a PSTN), or it may have been triggered by a local event (e.g., interesting traffic routed to a virtual interface). An incoming call that needs to be tunneled (as determined by the LAC) results in the generation of an L2TP ICRQ message. (See also: Call, Outgoing Call, Outgoing Call Request.)

## L2TP Access Concentrator (LAC)

If an L2TP Control Connection Endpoint (LCCE) is being used to cross-connect an L2TP session directly to a data link, we refer to it as an L2TP Access Concentrator (LAC). An LCCE may act as both an L2TP Network Server (LNS) for some sessions and an LAC for others, so these terms must only be used within the context of a given set of sessions unless the LCCE is in fact single purpose for a given topology. (See also: LCCE, LNS.)

## L2TP Control Connection Endpoint (LCCE)

An L2TP node which exists at either end of an L2TP control connection. May also be referred to as an LAC or LNS, depending on whether tunneled frames are processed at the data link (LAC) or network layer (LNS). (See also: LAC, LNS.)

## L2TP Network Server (LNS)

If a given L2TP session is terminated at the L2TP node and the encapsulated network layer (L3) packet processed on a virtual interface, we refer to this L2TP node as an L2TP Network Server (LNS). A given LCCE may act as both an LNS for some sessions and an LAC for others, so these terms must only be used within the context of a given set of sessions unless the LCCE is in fact single purpose for a given topology. (See also: LCCE, LAC.)

## Outgoing Call



The action of placing a call by an LAC, typically in response to policy directed by the peer in an Outgoing Call Request message. (See also: Call, Incoming Call, Outgoing Call Request.)

#### Outgoing Call Request

A request sent to an LAC to place an outgoing call. The request contains specific information not known a priori by the LAC (i.e., a number to dial). (See also: Call, Incoming Call, Outgoing Call.)

#### Packet-Switched Network (PSN)

A network that uses packet-switching technology for data delivery. For L2TPv3, this layer is principally IP. Other examples include MPLS, Frame Relay, and ATM.

#### Peer

When used in context with L2TP, Peer refers to the far end of an L2TP control connection (i.e., the remote LCCE). An LAC's peer may be either an LNS or another LAC. Similarly, an LNS's peer may be either an LAC or another LNS. (See also: LAC, LCCE, LNS.)

#### Pseudowire (PW)

An emulated circuit as it traverses a PSN. There is one Pseudowire per L2TP Session. (See also: Packet-Switched Network, Session.)

#### Pseudowire Type


The payload type being carried within an L2TP session. Examples include PPP, Ethernet, and Frame Relay. (See also: Session.)

#### Remote System

An end-system or router connected by a circuit to an LAC.

#### Session

An L2TP session is the entity which is created between two LCCEs in order to exchange parameters for and maintain an emulated L2 connection. Multiple sessions may be associated with a single Control Connection.



```

|<- emulated service ->|
|<----- L2 service ----->|

```

INTERNET DRAFT

L2TPv3

October 2003

(c) LNS-LNS Reference Model: This model has two LNSs as the LCCEs. A user-level, traffic-generated, or signaled event typically drives session establishment from one side of the tunnel. For example, a tunnel generated from a PC by a user, or automatically by customer premises equipment.

```

+-----+                               +-----+
[home network]...| LNS |.....[ IP ].....| LNS |...[home network]
+-----+                               +-----+
                    |<- emulated service ->|
                    |<----- L2 service ----->|

```

Note: In L2TPv2, user-driven tunneling of this type is often referred to as "voluntary tunneling" [[RFC2809](#)]. Further, an LNS acting as part of a software package on a host is sometimes referred to as an "LAC Client" [[RFC2661](#)].

### 3. Protocol Overview

L2TP utilizes two types of messages, control messages and data messages (sometimes referred to as "control packets" and "data packets", respectively). Control messages are used in the establishment, maintenance, and clearing of control connections and sessions. These messages utilize a reliable control channel within L2TP to guarantee delivery (see [Section 4.2](#) for details). Data messages are used to encapsulate the L2 traffic being carried over the L2TP session. Unlike control messages, data messages are not retransmitted when packet loss occurs.

The L2TPv3 control message format defined in this document borrows largely from L2TPv2. These control messages are used in conjunction with the associated protocol state machines that govern the dynamic setup, maintenance, and teardown for L2TP sessions. The data message format for tunneling data packets may be utilized with or without the L2TP control channel, either via manual configuration or other signaling methods to pre-configure or distribute L2TP session information. Utilization of the L2TP data message format with other signaling methods is outside the scope of this document.

INTERNET DRAFT

L2TPv3

October 2003

Figure 3.0: L2TPv3 Structure




Figure 3.0 depicts the relationship of control messages and data messages over the L2TP control and data channels, respectively. Data messages are passed over an unreliable data channel, encapsulated by an L2TP header, and sent over a Packet-Switched Network (PSN) such as IP, UDP, Frame Relay, ATM, MPLS, etc. Control messages are sent over a reliable L2TP control channel, which operates over the same PSN.

The necessary setup for tunneling a session with L2TP consists of two steps: (1) Establishing the control connection, and (2) establishing a session as triggered by an incoming call or outgoing call. An L2TP session **MUST** be established before L2TP can begin to forward session frames. Multiple sessions may be bound to a single control connection, and multiple control connections may exist between the same two LCCEs.

### [3.1](#) Control Message Types

The Message Type AVP (see [Section 5.4.1](#)) defines the specific type of control message being sent.

This document defines the following control message types (see Sections [6.1](#) through [6.13](#) for details on the construction and use of each message):

#### Control Connection Management

- 0 (reserved)
- 1 (SCCRQ) Start-Control-Connection-Request
- 2 (SCCRP) Start-Control-Connection-Reply
- 3 (SCCCN) Start-Control-Connection-Connected
- 4 (StopCCN) Stop-Control-Connection-Notification
- 5 (reserved)
- 6 (HELLO) Hello
- TBA-M1 (ACK) Explicit Acknowledgement

#### Call Management

- 7 (OCRQ) Outgoing-Call-Request
- 8 (OCRP) Outgoing-Call-Reply
- 9 (OCCN) Outgoing-Call-Connected
- 10 (ICRQ) Incoming-Call-Request
- 11 (ICRP) Incoming-Call-Reply
- 12 (ICCN) Incoming-Call-Connected
- 13 (reserved)
- 14 (CDN) Call-Disconnect-Notify

#### Error Reporting

- 15 (WEN) WAN-Error-Notify

#### Link Status Change Reporting

- 16 (SLI) Set-Link-Info

### [3.2](#) L2TP Header Formats

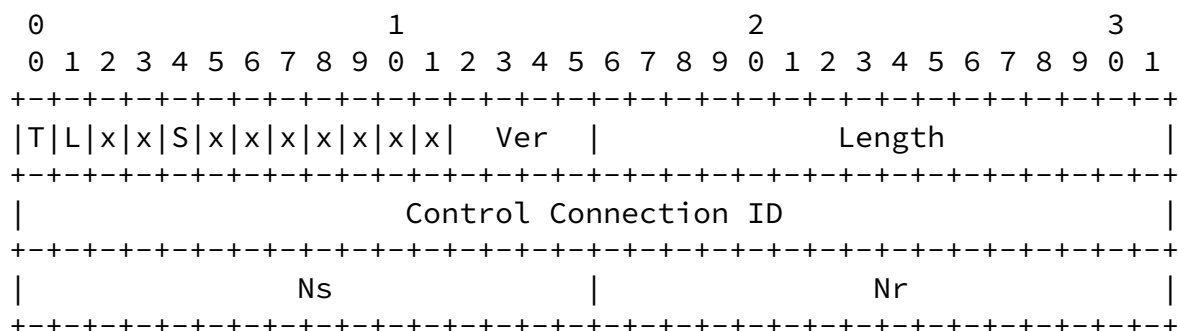
This section defines header formats for L2TP control messages and L2TP data messages. All values are placed into their respective fields and sent in network order (high-order octets first).

### 3.2.1 L2TP Control Message Header

The L2TP control message header provides information for the reliable transport of messages that govern the establishment, maintenance, and teardown of L2TP sessions. By default, control messages are sent over the underlying media in-band with L2TP data messages.

The L2TP control message header is formatted as follows:

Figure 3.2.1: L2TP Control Message Header



The T bit MUST be set to 1, indicating that this is a control message.

The L and S bits MUST be set to 1, indicating that the Length field and sequence numbers are present.

The x bits are reserved for future extensions. All reserved bits MUST be set to 0 on outgoing messages and ignored on incoming messages.

The Ver field indicates the version of the L2TP control message header described in this document. On sending, this field MUST be set to 3 for all messages (unless operating in an environment which includes L2TPv2 [[RFC2661](#)] and/or L2F [[RFC2341](#)] as well, see [Section 4.1](#) for details).

The Length field indicates the total length of the message in octets, always calculated from the start of the control message header itself (beginning with the T bit).

The Control Connection ID field contains the identifier for the

control connection. L2TP control connections are named by identifiers that have local significance only. That is, the same control connection will be given unique Control Connection IDs by each LCCE from within each endpoint's own Control Connection ID number space. As such, the Control Connection ID in each message is that of the intended recipient, not the sender. Non-zero Control Connection IDs are selected and exchanged as Assigned Control Connection ID AVPs during the creation of a control connection.

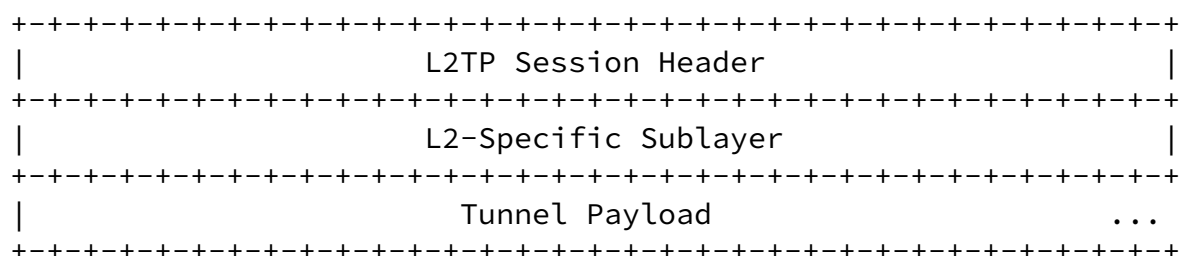
Ns indicates the sequence number for this control message, beginning at zero and incrementing by one (modulo  $2^{16}$ ) for each message sent. See [Section 4.2](#) for more information on using this field.

Nr indicates the sequence number expected in the next control message to be received. Thus, Nr is set to the Ns of the last in-order message received plus one (modulo  $2^{16}$ ). See [Section 4.2](#) for more information on using this field.

### 3.2.2 L2TP Data Message

In general, an L2TP data message consists of a (1) Session Header, (2) an optional L2-Specific Sublayer, and (3) the Tunnel Payload, as depicted below.

Figure 3.2.2: L2TP Data Message Header



The L2TP Session Header is specific to the encapsulating PSN over which the L2TP traffic is delivered. The Session Header MUST provide (1) a method of distinguishing traffic among multiple L2TP data sessions and (2) a method of distinguishing data messages from control messages.

Each type of encapsulating PSN MUST define its own session header, clearly identifying the format of the header and parameters necessary to setup the session. [Section 4.1](#) defines two session headers, one for transport over UDP and one for transport over IP.

The L2 Specific Sublayer is an intermediary layer between the L2TP session header and the start of the tunneled frame. It contains control fields that are used to facilitate the tunneling of each frame (e.g. sequence numbers or flags). The default L2-Specific Sublayer for L2TPv3 is defined in [Section 4.6](#).

The Data Message Header is followed by the Tunnel Payload, including any necessary L2 framing as defined in the payload-specific companion documents.

### [3.3](#) Control Connection Management

The L2TP Control Connection handles dynamic establishment, teardown, and maintenance of the L2TP sessions and of the control connection itself. The reliable delivery of control messages is described in [Section 4.2](#).

This section describes typical control connection establishment and teardown exchanges. It is important to note that, in the diagrams that follow, the reliable control message delivery mechanism exists independently of the L2TP state machine. For instance, Explicit Acknowledgement (ACK) messages may be sent after any of the control messages indicated in the exchanges below if an acknowledgment is not piggybacked on a later control message.

LCCEs are identified during control connection establishment either

by the Host Name AVP, the Router ID AVP, or a combination of the two (see [Section 5.4.3](#)). The identity of a peer LCCE is central to selecting proper configuration parameters (i.e. Hello interval, window size, etc.) for a control connection, as well as for



determination of how to setup associated sessions within the control connection, password lookup for control connection authentication, control connection level tie-breaking, etc.

### 3.3.1 Control Connection Establishment

Establishment of the control connection involves an exchange of AVPs that identifies the peer and its capabilities.

A three-message exchange is used to establish the control connection. The following is a typical message exchange:

LCCE A	LCCE B
-----	-----
SCCRQ ->	
	<- SCCRQ
SCCCN ->	

### 3.3.2 Control Connection Teardown

Control connection teardown may be initiated by either LCCE and is accomplished by sending a single StopCCN control message. As part of the reliable control message delivery mechanism, the recipient of a StopCCN MUST send an ACK message to acknowledge receipt of the message and maintain enough control connection state to properly accept StopCCN retransmissions over at least a full retransmission cycle (in case the ACK message is lost). The recommended time for a full retransmission cycle is at least 31 seconds (see [Section 4.2](#)). The following is an example of a typical control message exchange:

LCCE A	LCCE B
-----	-----
StopCCN ->	
(Clean up)	
	(Wait)
	(Clean up)

An implementation may shut down an entire control connection and all sessions associated with the control connection by sending the StopCCN. Thus, it is not necessary to clear each session individually when tearing down the whole control connection.

### [3.4](#) Session Management

After successful control connection establishment, individual sessions may be created. Each session corresponds to a single data stream between the two LCCEs. This section describes the typical call establishment and teardown exchanges.

#### [3.4.1](#) Session Establishment for an Incoming Call

A three-message exchange is used to establish the session. The following is a typical sequence of events:

LCCE A	LCCE B
-----	-----
(Call Detected)	
ICRQ ->	
	<- ICRP
(Call Accepted)	
ICCN ->	

#### [3.4.2](#) Session Establishment for an Outgoing Call

A three-message exchange is used to set up the session. The following is a typical sequence of events:

LCCE A	LCCE B
-----	-----
	<- OCRQ
OCRP ->	
(Perform Call Operation)	
OCCN ->	
(Call Operation Completed Successfully)	

#### [3.4.3](#) Session Teardown

Session teardown may be initiated by either the LAC or LNS and is

accomplished by sending a CDN control message. After the last

INTERNET DRAFT

L2TPv3

October 2003

session is cleared, the control connection MAY be torn down as well (and typically is). The following is an example of a typical control message exchange:

```
LCCE A      LCCE B
-----
CDN ->
(Clean up)

              (Clean up)
```

## [4.](#) Protocol Operation

### [4.1](#) L2TP Over Specific Packet-Switched Networks (PSN)

If necessary, L2TP may operate over a variety of Packet Switched Networks. There are two modes described for operation over IPv4, L2TP over IP ([Section 4.1.1](#)) and L2TP over UDP ([Section 4.1.2](#)). L2TPv3 implementations MUST support L2TP over IP and SHOULD support L2TP over UDP for better NAT and FW traversal, and easier migration from L2TPv2.

L2TP over other PSNs may be defined, but the specifics are outside the scope of this document. Examples of L2TPv2 over other PSNs include [[RFC3070](#)] and [[RFC3355](#)].

The following field definitions are defined for use in all L2TP Session Header encapsulations.

#### Session ID

A 32-bit field containing a non-zero identifier for a session. L2TP sessions are named by identifiers that have local significance only. That is, the same logical session will be given different Session IDs by each end of the control connection for the life of the session. When the L2TP control connection is used for session establishment, Session IDs are selected and exchanged as Local Session ID AVPs during the creation of a session.

## Cookie

The optional Cookie field contains a variable length (maximum 64 bits), value used to check the association of a received data message with the session identified by the Session ID. The Cookie MUST be set to the configured or signaled random value for this session utilizing all bits in the field. The Cookie provides an

additional level of guarantee that a data message has been directed to the proper session by the Session ID. A well-chosen Cookie may prevent inadvertent misdirection of stray packets with recently reused Session IDs, Session IDs subject to packet corruption, etc. The Cookie may also provide protection against some specific malicious packet insertion attacks, as described in [section 8.2](#).

When the L2TP control connection is used for session establishment, random Cookie values are selected and exchanged as Assigned Cookie AVPs during session creation.


### [4.1.1](#) L2TPv3 over IP

L2TPv3 over IP utilizes the IANA assigned IP protocol ID 115.

#### [4.1.1.1](#) L2TPv3 Session Header Over IP

Unlike L2TP over UDP, the L2TPv3 session header over IP is free of any restrictions imposed by coexistence with L2TPv2 and L2F. As such, the header format has been redesigned to optimize packet processing. The following session header format is utilized when operating L2TPv3 over IP:

Figure 4.1.1.1: L2TPv3 Session Header Over IP



```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Session ID and Cookie fields are as defined in [Section 4.1](#). The Session ID of zero is reserved for use by L2TP control messages (see [Section 4.1.1.2](#)).

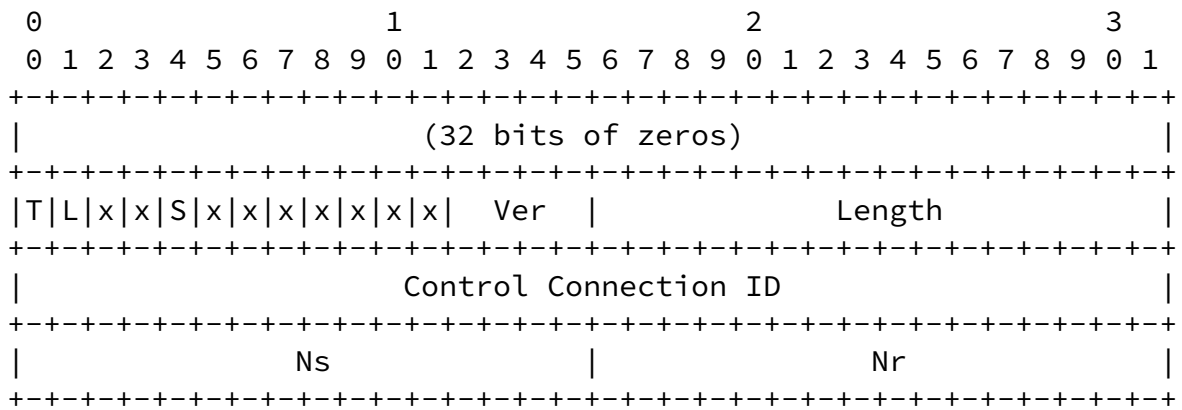
[4.1.1.2](#) L2TP Control and Data Traffic over IP

Unlike L2TP over UDP which uses the common T bit to distinguish between L2TP control and data packets, L2TP over IP uses the reserved Session ID of zero (0) when sending control messages. It is presumed that checking for the zero Session ID is more efficient -- both in header size for data packets and in processing speed for

distinguishing between control and data messages -- than checking for the presence of a given single bit.

The entire control message header over IP, including the zero session ID, appears as follows:

Figure 4.1.1.2: L2TPv3 Control Message Header Over IP



Named fields are as defined in [Section 3.2.1](#). Note that the Length field is still calculated from the beginning of the control message header, beginning with the T bit. It does NOT include the "(32 bits of zeros)" depicted above.

When operating directly over IP, L2TP packets lose the ability to

take advantage of the UDP checksum as a simple packet integrity check. This is of particular concern for L2TP control messages. Control Message Authentication ([Section 4.3](#)), even with an empty password field, provides for a sufficient packet integrity check and SHOULD always be enabled.

#### [4.1.2](#) L2TP over UDP

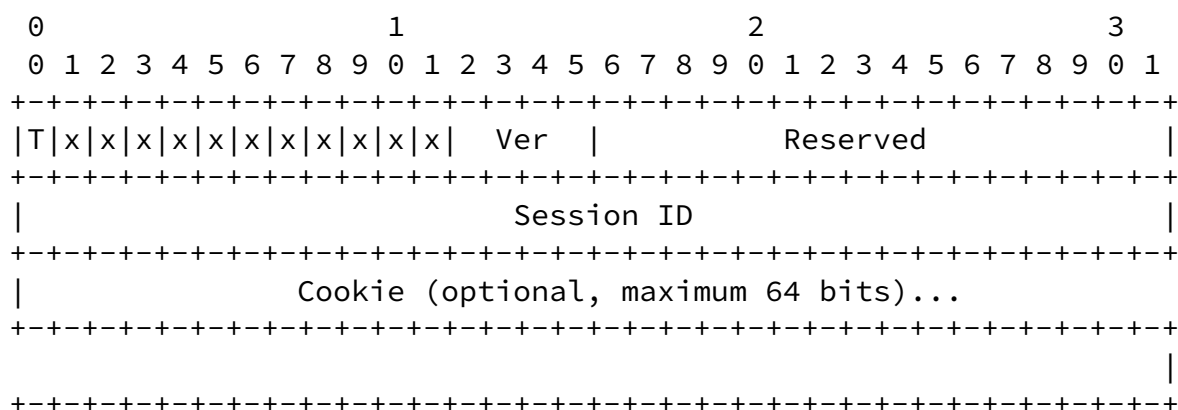
L2TPv3 over UDP must consider other L2 tunneling protocols that may be operating in the same environment, including L2TPv2 [[RFC2661](#)] and L2F [[RFC2341](#)].

While there are efficiencies gained by running L2TP directly over IP, there are possible side effects as well. For instance, L2TP over IP is not as NAT-friendly as L2TP over UDP.

##### [4.1.2.1](#) L2TP Session Header Over UDP

The following session header format is utilized when operating L2TPv3 over UDP:

Figure 4.1.2.1: L2TPv3 Session Header over UDP



The T bit MUST be set to 0, indicating that this is a data message.

The x bits and Reserved field are reserved for future extensions. All reserved values MUST be set to 0 on outgoing messages and ignored on incoming messages.

The Ver field MUST be set to 3, indicating an L2TPv3 message.

Note that the initial bits 1, 4, 6 and 7 have meaning in L2TPv2 [[RFC2661](#)], and are deprecated and marked as reserved in L2TPv3. Thus, for UDP mode on a system that supports both versions of L2TP, it is important that the Ver field be inspected first to determine the Version of the header before acting upon any of these bits.

The Session ID and Cookie fields are as defined in [Section 4.1](#).

#### [4.1.2.2](#) UDP Port Selection

The method for UDP Port Selection defined in this section is identical to than defined for L2TPv2 [[RFC2661](#)].

When negotiating a control connection over UDP, control messages MUST be sent as UDP datagrams using the registered UDP port 1701 [[RFC1700](#)]. The initiator of an L2TP control connection picks an available source UDP port (which may or may not be 1701), and sends to the desired destination address at port 1701. The recipient picks a free port on its own system (which may or may not be 1701) and sends its reply to the initiator's UDP port and address, setting its own source port to the free port it found.

Any subsequent traffic associated with this control connection (either control traffic or data traffic from a session established through this control connection) must use these same UDP ports.

It has been suggested that having the recipient choose an arbitrary source port (as opposed to using the destination port in the packet initiating the control connection, i.e., 1701) may make it more difficult for L2TP to traverse some NAT devices. Implementations should consider the potential implication of this capability before choosing an arbitrary source port. A NAT device that can pass TFTP traffic with variant UDP ports should be able to pass L2TP UDP traffic since both protocols employ similar policies with regard to UDP port selection.

#### [4.1.2.3](#) UDP Checksum

UDP checksums **MUST** be enabled for control messages and **MAY** be enabled for data messages. It should be noted that enabling checksums on data packets may significantly increase the data packet processing burden.

#### [4.1.3](#) IP Fragmentation Issues

IP fragmentation may occur as the L2TP packet travels over the IP substrate. L2TP makes no special efforts defined in this document to optimize this.

#### [4.2](#) Reliable Delivery of Control Messages

L2TP provides a lower level reliable delivery service for all control messages. The Nr and Ns fields of the control message header (see [Section 3.2.1](#)) belong to this delivery mechanism. The upper level functions of L2TP are not concerned with retransmission or ordering of control messages. The reliable control messaging mechanism is a sliding window mechanism that provides control message retransmission and congestion control. Each peer maintains separate sequence number state for each control connection.

The message sequence number, Ns, begins at 0. Each subsequent message is sent with the next increment of the sequence number. The sequence number is thus a free-running counter represented modulo 65536. The sequence number in the header of a received message is considered less than or equal to the last received number if its value lies in the range of the last received number and the preceding 32767 values, inclusive. For example, if the last received sequence number was 15, then messages with sequence numbers 0 through 15, as well as 32784 through 65535, would be considered less than or equal. Such a message would be considered a duplicate of a message already received and ignored from processing. However, in order to ensure that all messages are acknowledged properly (particularly in the case of a lost ACK message), receipt of duplicate messages **MUST** be acknowledged by the reliable delivery mechanism. This acknowledgment

may either piggybacked on a message in queue or sent explicitly via an ACK message.

All control messages take up one slot in the control message sequence number space, except the ACK message. Thus, Ns is not incremented



after an ACK message is sent.

The last received message number, Nr, is used to acknowledge messages received by an L2TP peer. It contains the sequence number of the message the peer expects to receive next (e.g. the last Ns of a non-ACK message received plus 1, modulo 65536). While the Nr in a received ACK message is used to flush messages from the local retransmit queue (see below), the Nr of the next message sent is not updated by the Ns of the ACK message. Nr SHOULD be sanity-checked before flushing the retransmit queue. For instance, if the Nr received in a control message is greater than the last Ns sent plus 1 modulo 65536, the control message is clearly invalid.

The reliable delivery mechanism at a receiving peer is responsible for making sure that control messages are delivered in order and without duplication to the upper level. Messages arriving out of order may be queued for in-order delivery when the missing messages are received. Alternatively, they may be discarded, thus requiring a retransmission by the peer. When dropping out of order control packets, Nr MAY be updated before the packet is discarded.

Each control connection maintains a queue of control messages to be transmitted to its peer. The message at the front of the queue is sent with a given Ns value and is held until a control message arrives from the peer in which the Nr field indicates receipt of this message. After a period of time (a recommended default is 1 second but SHOULD be configurable) passes without acknowledgment, the message is retransmitted. The retransmitted message contains the same Ns value, but the Nr value MUST be updated with the sequence number of the next expected message.

Each subsequent retransmission of a message MUST employ an exponential backoff interval. Thus, if the first retransmission occurred after 1 second, the next retransmission should occur after 2 seconds has elapsed, then 4 seconds, etc. An implementation MAY place a cap upon the maximum interval between retransmissions. This cap SHOULD be no less than 8 seconds per retransmission. If no peer response is detected after several retransmissions (a recommended default is 10, but MUST be configurable), the control connection and all associated sessions MUST be cleared. As it is the first message to establish a control connection, the SCCRQ MAY employ a different retransmission maximum than other control messages in order to help facilitate failover to alternate LCCEs in a timely fashion.

When a control connection is being shut down for reasons other than loss of connectivity, the state and reliable delivery mechanisms MUST be maintained and operated for the full retransmission interval after the final message StopCCN message has been sent (e.g. 1 + 2 + 4 + 8 + 8... seconds), or until the StopCCN message itself has been acknowledged.

A sliding window mechanism is used for control message transmission and retransmission. Consider two peers, A and B. Suppose A specifies a Receive Window Size AVP with a value of N in the SCCRQ or SCCRP message. B is now allowed to have a maximum of N outstanding (e.g. unacknowledged) control messages. Once N messages have been sent, B must wait for an acknowledgment from A that advances the window before sending new control messages. An implementation may advertise a non-zero receive window as small or as large as it wishes, depending on its own ability to process incoming messages before sending an acknowledgment. Each peer MUST limit the number of unacknowledged messages it will send before receiving an acknowledgment by this Receive Window Size. The actual internal unacknowledged message send-queue depth may be further limited by local resource allocation or by dynamic slow-start and congestion-avoidance mechanisms.

When retransmitting control messages, a slow start and congestion avoidance window adjustment procedure SHOULD be utilized. A recommended procedure is described in [Appendix A](#). A peer MAY drop messages, but MUST NOT actively delay acknowledgment of messages as a technique for flow control of control messages. [Appendix B](#) contains examples of control message transmission, acknowledgment, and retransmission.

### [4.3](#) Control Connection and Control Message Authentication

L2TP incorporates an optional authentication and integrity check for all control messages. This mechanism consists of a computed one-way hash over the header and body of the L2TP control message, a pre-configured shared secret, and a local and remote nonce (random value) exchanged via the Nonce AVP. This per-message authentication and integrity check is designed to perform a mutual authentication between L2TP nodes, integrity checking of all control messages, and guard against control message spoofing and replay attacks that would otherwise be trivial to mount.

A shared secret (password) MUST exist between communicating L2TP nodes to enable the authentication mode. See [Section 5.4.3](#) for details on calculation of the Message Digest and construction of the Nonce and Message Digest AVPs.

INTERNET DRAFT

L2TPv3

October 2003

L2TPv3 Control Connection and Control Message Authentication is similar to L2TPv2 [[RFC2661](#)] Tunnel Authentication in its use of a shared secret for peer authentication, use of a one-way hash calculation, and exchange of a random value. The principal difference is that, instead of computing the hash over selected contents of a received control message (e.g. the Challenge AVP and Message Type) as in L2TPv2, the entire message is used in the hash in L2TPv3. In addition, instead of including the hash digest in just the SCCRP and SCCCN messages, it is now included in all L2TP messages.

The Control Message Authentication mechanism is optional, and may be disabled if both peers agree. For example, if IPsec is already being used for security and integrity checking between the LCCes, the function of the L2TP mechanism becomes redundant and may be disabled.

Presence of the Control Message Authentication Nonce AVP in an SCCRQ or SCCRP message serves as indication to a peer that Control Message Authentication is enabled. If an SCCRQ or SCCRP contains a Control Message Authentication Nonce AVP, the receiver of the message MUST respond with a Message Digest AVP in all subsequent messages sent. Control Connection and Control Message Authentication is always bidirectional, either both sides participate in authentication or neither do.

If the Control Message Authentication is disabled, the Message Digest AVP still MAY be sent containing an integrity check of the message. The integrity check is calculated as in [Section 5.4.3](#), with an empty and zero length shared secret, local nonce, and remote nonce. If a erroneous Message Digest AVP is received, it should be assumed that a bit error has occurred and the control message dropped.

Implementations MAY rate-limit control messages, particularly SCCRQ messages, upon receipt for performance reasons or to protect against denial of service attacks.

#### [4.4](#) Keepalive (Hello)

A keepalive mechanism is employed by L2TP to detect loss of connectivity between a pair of LCCes. This is accomplished by injecting Hello control messages (see [Section 6.5](#)) after a period of time has elapsed since the last data message or control message was received on an L2TP session or control connection, respectively. As

with any other control message, if the Hello message is not reliably delivered, the sending LCCE declares that the control connection is down and resets its state for the control connection. This behavior ensures that a connectivity failure between the LCCEs is detected independently by each end of a control connection.

Since the control channel is operated in-band with data traffic over the PSN, this single mechanism can be used to infer basic data connectivity between a pair of LCCEs for all sessions associated with the control connection.

Periodic keepalive for the control connection **MUST** be implemented by sending a Hello if a period of time (a recommended default is 60 seconds, but **MUST** be configurable) has passed without receiving any message (data or control) from the peer. An LCCE sending Hello messages across multiple control connections between the same LCCE endpoints **MUST** employ a jittered timer mechanism to prevent grouping of Hello messages.

#### [4.5](#) Forwarding Session Data Frames

Once session establishment is complete, circuit frames are received at an LCCE, encapsulated in L2TP (with appropriate attention to framing as described in documents for the particular pseudowire type), and forwarded over the appropriate session. For every outgoing data message, the sender places the identifier specified in the Local Session ID AVP (received from peer during session establishment) in the Session ID field of the L2TP data header. In this manner, session frames are multiplexed and demultiplexed between a given pair of LCCEs. Multiple control connections may exist between a given pair of LCCEs, and multiple sessions may be associated with a given control connection.

The peer LCCE receiving the L2TP data packet identifies the session with which the packet is associated by the Session ID in the data packet's header. The LCCE then checks the Cookie field in the data packet against the Cookie value received in the Assigned Cookie AVP during session establishment. It is important for implementers to note that the Cookie field check occurs after looking up the session context by the Session ID, and as such consists merely of a value match of the Cookie field and that stored in the retrieved context.



#### [4.6.1](#) Sequencing Data Packets

The Sequence Number field may be used to detect lost, duplicate, or out-of-order packets within a given session.

When L2 frames are carried over an L2TP-over-IP or L2TP-over-UDP/IP data channel, this part of the link has the characteristic of being able to reorder, duplicate, or silently drop packets. Reordering may break some non-IP protocols or L2 control traffic being carried by the link. Silent dropping or duplication of packets may break protocols that assume per-packet indications of error, such as TCP header compression. While a common mechanism for packet sequence detection is provided, the sequence dependency characteristics of individual protocols are outside the scope of this document.

If any protocol being transported by over L2TP data channels cannot tolerate misordering of data packets, packet duplication, or silent packet loss, sequencing may be enabled on some or all packets by using the S bit and Sequence Number field defined in the default L2-

Specific Sublayer(see [Section 4.6](#)). For a given L2TP session, each LCCE is responsible for communicating to its peer the level of sequencing support that it requires of data packets that it receives. Mechanisms to advertise this information during session negotiation are provided (see Data Sequencing AVP in [Section 5.4.4](#)).

When determining whether a packet is in or out of sequence, an implementation SHOULD utilize a method that is resilient to temporary dropouts in connectivity coupled with high per-session packet rates. The recommended method is outlined in [Appendix C](#).

#### [4.7](#) L2TPv2/v3 Interoperability and Migration

L2TPv2 and L2TPv3 environments should be able to coexist while a migration to L2TPv3 is made. Migration issues are discussed for each media type in this section. Most issues apply only to implementations that require both L2TPv2 and L2TPv3 operation. However, even L2TPv3-only implementations must at least be mindful of these issues in order to interoperate with implementations that support both versions.

#### [4.7.1](#) L2TPv3 over IP

L2TPv3 implementations running strictly over IP with no desire to interoperate with L2TPv2 implementations may safely disregard most migration issues from L2TPv2. All control messages and data messages are sent as described in this document, without normative reference to [RFC2661](#).

If one wishes to tunnel PPP over L2TPv3, and fallback to L2TPv2 only if it is not available, then L2TPv3 over UDP with the automatic fallback as described in [section 4.7.3](#) MUST be used. There is no deterministic method for automatic fallback from L2TPv3 over IP to either L2TPv2 or L2TPv3 over UDP. One could infer whether L2TPv3 over IP is supported by sending an SCCRQ and waiting for a response, but this could be problematic during periods of packet loss between L2TP nodes.

#### [4.7.2](#) L2TPv3 over UDP

The format of the L2TPv3 over UDP header is defined in [Section 4.1.2.1](#).

When operating over UDP, L2TPv3 uses the same port (1701) as L2TPv2 and shares the first two octets of header format with L2TPv2. The Ver field is used to distinguish L2TPv2 packets from L2TPv3 packets. If an implementation is capable of operating in L2TPv2 or L2TPv3 modes, it is possible to automatically detect whether a peer can support

L2TPv2 or L2TPv3 and operate accordingly. The details of this fallback capability is defined in the following section.

#### [4.7.3](#) Automatic L2TPv2 Fallback

When running over UDP, an implementation may detect whether a peer is L2TPv3-capable by sending a special SCCRQ that is properly formatted for both L2TPv2 and L2TPv3. This is accomplished by sending an SCCRQ with its Ver field set to 2 (for L2TPv2), and ensuring that any L2TPv3-specific AVPs (i.e. AVPs present within this document and not defined within [RFC 2661](#)) within the message are sent with each M bit set to 0, and all L2TPv2 AVPs present as they would be for L2TPv2. This is done so that L2TPv3 AVPs will be ignored by an L2TPv2-only implementation. Note that, in both L2TPv2 and L2TPv3, the value





Attribute Type	Attribute Value ...
(until Length is reached)	

The first six bits comprise a bit mask that describes the general attributes of the AVP. Two bits are defined in this document; the remaining bits are reserved for future extensions. Reserved bits MUST be set to 0. An AVP received with a reserved bit set to 1 MUST be treated as an unrecognized AVP.

**Mandatory (M) bit:** Controls the behavior required of an implementation that receives an unrecognized AVP. The M bit of a given AVP MUST only be inspected and acted upon if the AVP is unrecognized (see [Section 5.2](#)).

**Hidden (H) bit:** Identifies the hiding of data in the Attribute Value field of an AVP. This capability can be used to avoid the passing of sensitive data, such as user passwords, as cleartext in an AVP. [Section 5.3](#) describes the procedure for performing AVP hiding.

**Length:** Contains the number of octets (including the Overall Length and bit mask fields) contained in this AVP. The Length may be calculated as 6 + the length of the Attribute Value field in octets. The field itself is 10 bits, permitting a maximum of 1023 octets of data in a single AVP. The minimum Length of an AVP is 6. If the Length is 6, then the Attribute Value field is absent.

**Vendor ID:** The IANA assigned "SMI Network Management Private Enterprise Codes" [[RFC1700](#)] value. The value 0, corresponding to IETF adopted attribute values, is used for all AVPs defined within this document. Any vendor wishing to implement its own L2TP extensions can use its own Vendor ID along with private Attribute values, guaranteeing that they will not collide with any other vendor's extensions or future IETF extensions. Note that there are 16 bits allocated for the Vendor ID, thus limiting this feature to the first 65,535 enterprises.

**Attribute Type:** A 2-octet value with a unique interpretation across all AVPs defined under a given Vendor ID.

Attribute Value: This is the actual value as indicated by the Vendor ID and Attribute Type. It follows immediately after the Attribute Type field and runs for the remaining octets indicated in the Length (i.e., Length minus 6 octets of header). This field is absent if the Length is 6.

## [5.2](#) Mandatory AVPs and Setting the M Bit

If the M bit is set on an AVP that is unrecognized by its recipient, the session or control connection associated with the control message containing the AVP MUST be shutdown. If the control message containing the unrecognized AVP is associated with a session (e.g. an ICRQ, ICRP, ICCN, SLI, etc.) then the session MUST be issued a CDN with a Result Code of 2 and Error Code of 8 as defined in [section 5.4.2](#). and shutdown. If the control message containing the unrecognized AVP is associated with establishment or maintenance of a Control Connection (e.g. SCCRQ, SCCRP, SCCCEN, Hello) then the associated Control Connection MUST be issued a StopCCN with Result Code of 2 and Error Code of 8 as defined in [section 5.4.2](#). and shutdown. If the M bit is not set on an unrecognized AVP, the AVP MUST be ignored when received, processing the control message as if the AVP was not present.

Receipt of an unrecognized AVP that has the M bit set is catastrophic to the session or control connection with which it is associated. Thus, the M bit should only be set for AVPs that are deemed crucial to proper operation of the session or control connection by the sender. AVPs that are considered crucial by the sender may vary by application and configured options. In no case shall a receiver of an AVP "validate" if the M bit is set on a recognized AVP. If the AVP is recognized (as all AVPs defined in this document MUST be for a compliant L2TPv3 specification), then by definition the M bit is of no consequence.

The sender of an AVP is free to set its M bit to 1 or 0 based on whether the configured application strictly requires the value contained in the AVP to be recognized or not. For example, "Automatic L2TPv2 Fallback" ([Section 4.7.3](#)), requires the setting of the M bit on all new L2TPv3 AVPs to zero if fallback to L2TPv2 is supported and desired, and 1 if not.

The M bit is useful as extra assurance for support of critical AVP extensions. However, more explicit methods may be available to determine support for a given feature rather than using the M bit alone. For example, if a new AVP is defined in a message for which there is always a message reply (i.e. an ICRQ, ICRP, SCCRQ or SCCRP message) rather than simply sending an AVP in the message with the M bit set, availability of the extension may be identified by sending

INTERNET DRAFT

L2TPv3

October 2003

an AVP in the request message and expecting a corresponding AVP in a reply message. This more explicit method, when possible, is preferred.

The M bit also plays a role in determining whether or not a malformed or out-of-range value within an AVP should be ignored or result in termination of a session or control channel. See [Section 7.1](#) for more details on this.

### [5.3](#) Hiding of AVP Attribute Values

The H bit in the header of each AVP provides a mechanism to indicate to the receiving peer whether the contents of the AVP are hidden or present in cleartext. This feature can be used to hide sensitive control message data such as user passwords, IDs, or other vital information.


The H bit MUST only be set if (1) a shared secret exists between the LCCEs and (2) Control Connection and Control Message Authentication is enabled (see [Section 4.3](#)). If the H bit is set in any AVP(s) in a given control message, a Random Vector AVP must also be present in the message and MUST precede the first AVP having an H bit of 1.

The shared secret between LCCEs is used to derive a unique shared key for hiding and unhiding calculations. The derived shared key is obtained via a one-way HMAC-MD5 hash [[RFC1321](#)] on the shared secret concatenated with a single octet containing the value 1.

$$\text{shared\_key} = \text{HMAC\_MD5}(\text{shared\_secret} \parallel 1)$$

Hiding an AVP value is done in several steps. The first step is to take the length and value fields of the original (cleartext) AVP and encode them into a Hidden AVP Subformat as follows:

Figure 5.3: Hidden AVP Subformat





octets, the Subformat is transformed as if the Attribute Value field had been padded to 16 octets before the XOR. Only the actual octets present in the Subformat are modified, and the length of the AVP is not altered.

If the Subformat is longer than 16 octets, a second one-way MD5 hash is calculated over a stream of octets consisting of the shared key followed by the result of the first XOR. That hash is XORed with the second 16-octet (or less) segment of the Subformat and placed in the corresponding octets of the Value field of the Hidden AVP.

If necessary, this operation is repeated, with the shared key used

along with each XOR result to generate the next hash to XOR the next segment of the value with.

The hiding method was adapted from [[RFC2865](#)], which was taken from the "Mixing in the Plaintext" section in the book "Network Security" by Kaufman, Perlman and Speciner [[KPS](#)]. A detailed explanation of the method follows:

Call the shared key S, the Random Vector RV, and the Attribute Value AV. Break the value field into 16-octet chunks p1, p2, etc., with the last one padded at the end with random data to a 16-octet boundary. Call the ciphertext blocks c(1), c(2), etc. We will also define intermediate values b1, b2, etc.

$$\begin{array}{ll}
 b_1 = \text{MD5}(\text{AV} + \text{S} + \text{RV}) & c(1) = p_1 \text{ xor } b_1 \\
 b_2 = \text{MD5}(\text{S} + c(1)) & c(2) = p_2 \text{ xor } b_2 \\
 \cdot & \cdot \\
 \cdot & \cdot \\
 \cdot & \cdot \\
 b_i = \text{MD5}(\text{S} + c(i-1)) & c(i) = p_i \text{ xor } b_i
 \end{array}$$

The String will contain c(1)+c(2)+...+c(i), where + denotes concatenation.

On receipt, the random vector is taken from the last Random Vector AVP encountered in the message prior to the AVP to be unhidden. The above process is then reversed to yield the original value.

#### [5.4](#) AVP Summary

The following sections contain a list of all L2TP AVPs defined in this document.

Following the name of the AVP is a list indicating the message types that utilize each AVP. After each AVP title follows a short description of the purpose of the AVP, a detail (including a graphic) of the format for the Attribute Value, and any additional information needed for proper use of the AVP.

#### [5.4.1](#) General Control Message AVPs

##### Message Type (All Messages)

The Message Type AVP, Attribute Type 0, identifies the control message herein and defines the context in which the exact meaning of the following AVPs will be determined.

The Attribute Value field for this AVP has the following format:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|           Message Type           |
+---+---+---+---+---+---+---+---+

```

The Message Type is a 2-octet unsigned integer.

The Message Type AVP MUST be the first AVP in a message, immediately following the control message header (defined in [Section 3.2.1](#)). See [Section 3.1](#) for the list of defined control message types and their identifiers.

The Mandatory (M) bit within the Message Type AVP has special meaning. Rather than an indication as to whether the AVP itself should be ignored if not recognized, it is an indication as to whether the control message itself should be ignored. If the M bit is set within the Message Type AVP and the Message Type is unknown to the implementation, the control connection MUST be cleared. If the M bit is not set, then the implementation may ignore an unknown message type. The M bit MUST be set to 1 for all message types defined in this document. This AVP MAY NOT be hidden (the H bit MUST be 0).

A vendor-specific control message may be defined by setting the Vendor ID of the Message Type AVP to a value other than the IETF Vendor ID of 0 (see [Section 5.1](#)). The Message Type AVP MUST still be the first AVP in the control message.

The Message Digest AVP, Attribute Type AVP-TBA-1, is used as an integrity check and authentication of the L2TP Control Message header and body.

[illegible]

- 0 HMAC-MD5 [[RFC1321](#)]
- 1 HMAC-SHA-1 [[RFC2104](#)]

The Message Digest is of variable length and contains the result of the control message authenticity and integrity calculation. For Digest Type 0 (HMAC-MD5) the length of the digest MUST be 16 bytes. For Digest Type 1 (SHA-1) the digest MUST be 20 bytes.

If Control Connection and Control Message Authentication is enabled, the Message Digest AVP MUST present in all messages and

MUST be placed immediately after the Message Type AVP. This forces the Message Digest to be present within each message at a well-known and fixed offset.

The shared secret between LCCEs is used to derive a unique shared key for Control Connection and Control Message Authentication calculations. The derived shared key is obtained via a one-way HMAC-MD5 hash [[RFC1321](#)] on the shared secret concatenated with a single octet containing the value 2.

$$\text{shared\_key} = \text{HMAC\_MD5}(\text{shared\_secret} \parallel 2)$$

Calculation of the digest is as follows for all messages other than the SCCRQ:

$$\text{Digest} = \text{Hash}(\text{local\_nonce} \parallel \text{remote\_nonce} \parallel \text{shared\_key} \parallel \text{control\_message})$$

Hash: Hashing algorithm identified by the Digest Type

local\_nonce: Nonce chosen locally and advertised to the remote LCCE.

remote\_nonce: Nonce received from the remote LCCE

(The local\_nonce and remote\_nonce are advertised via the Control Message Authentication Nonce AVP, also defined in this section.)

shared\_key: Derived shared key for this Control Connection

control\_message: The entire contents of the L2TP Control Message, including the Control Message header and all AVPs.

Note that the Control Message header in this case begins after the 0 Session ID (see [Section 4.1.1.2](#)) when running over IP, and after the UDP header when running over UDP (see [Section 4.1.2.1](#)).

When calculating the Message Digest, the Message Digest AVP MUST be present within the control message with the Digest Type set to its proper value, but the Message Digest itself set to zeros.



When receiving a control message, the contents of the Message Digest AVP MUST be compared against the expected digest value based on local calculation. This is done by performing the same digest calculation above, with the local\_nonce and remote\_nonce reversed. This message authenticity and integrity checking MUST be performed before utilizing any information contained within the control message. If the calculation fails, the message MUST be dropped.

The SCCRQ has special treatment as it is the initial message commencing a new Control Connection. As such, there is only one nonce available. Since the nonce is present within the message itself as part of the Control Message Authentication Nonce AVP, there is no need to use it in the calculation explicitly. Calculation of the SCCRQ Digest is performed as follows:

Digest = Hash (shared\_key | control\_message)

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length is 23 for Digest Type 1 (HMAC-MD5), and 27 for Digest Type 2 (SHA-1). Control Message Authentication Nonce (SCCRQ, SCCRP)

The Control Message Authentication Nonce AVP, Attribute Type AVP-TBA-15, MUST contain a cryptographically random value [[RFC1750](#)]. This value is used for Control Message Authentication.

The Attribute Value field for this AVP has the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Nonce ... (arbitrary number of octets)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Nonce is of arbitrary length, though at least 16 octets is recommended. The Nonce contains the random value for use in

the Control Message Authentication hash calculation (see

Message Digest AVP definition in this section).

If Control Connection and Message Authentication is enabled, this AVP MUST be present in the SCCRQ and SCCRP messages.

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 6 plus the length of the Nonce.

#### Random Vector (All Messages)

The Random Vector AVP, Attribute Type 36, MUST contain a cryptographically random value [[RFC1750](#)]. This value is used for AVP Hiding.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Random Octet String ... (arbitrary number of octets)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The Random Octet String is of arbitrary length, though at least 16 octets is recommended. The string contains the random vector for use in computing the MD5 hash to retrieve or hide the Attribute Value of a hidden AVP (see [Section 5.3](#)).

More than one Random Vector AVP may appear in a message, in which case a hidden AVP uses the Random Vector AVP most closely preceding it. As such, at least one Random Vector AVP MUST precede the first AVP with the H bit set.

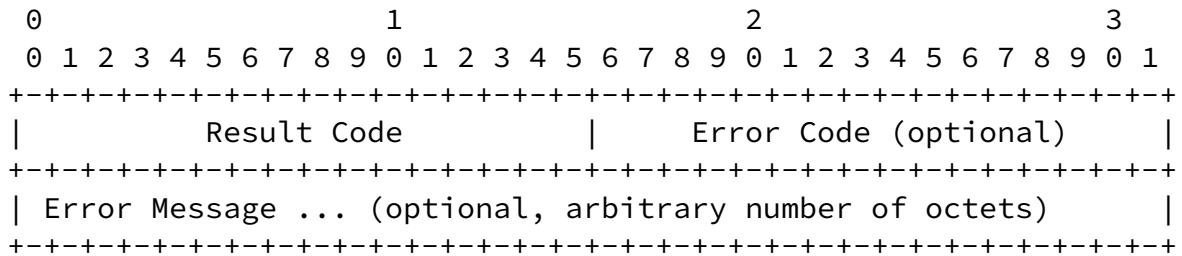
This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 6 plus the length of the Random Octet String.

#### [5.4.2](#) Result and Error Codes

##### Result Code (StopCCN, CDN)

The Result Code AVP, Attribute Type 1, indicates the reason for terminating the control channel or session.

The Attribute Value field for this AVP has the following format:



The Result Code is a 2-octet unsigned integer. The optional Error Code is a 2-octet unsigned integer. An optional Error Message can follow the Error Code field. Presence of the Error Code and Message is indicated by the AVP Length field. The Error Message contains an arbitrary string providing further (human-readable) text associated with the condition. Human-readable text in all error messages MUST be provided in the UTF-8 charset using the Default Language [[RFC2277](#)].

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length is 8 if there is no Error Code or Message, 10 if there is an Error Code and no Error Message, or 10 plus the length of the Error Message if there is an Error Code and Message.

Defined Result Code values for the StopCCN message are as follows:

- 0 - Reserved.
- 1 - General request to clear control connection.
- 2 - General error, Error Code indicates the problem.
- 3 - Control channel already exists.
- 4 - Requester is not authorized to establish a control channel.
- 5 - The protocol version of the requester is not supported, Error Code indicates highest version supported.
- 6 - Requester is being shut down.
- 7 - Finite state machine error or timeout

General Result Code values for the CDN message are as follows:

- 0 - Reserved.
- 1 - Session disconnected due to loss of carrier or circuit disconnect.
- 2 - Session disconnected for the reason indicated in Error Code.
- 3 - Session disconnected for administrative reasons.
- 4 - Session establishment failed due to lack of appropriate facilities being available (temporary condition).
- 5 - Session establishment failed due to lack of appropriate facilities being available (permanent condition).
- 6 - 11 Reserved (PPP-specific codes defined outside this document).

RC-TBA-1 - Session not established due to losing tie breaker.  
RC-TBA-2 - Session not established due to unsupported PW type.

INTERNET DRAFT

L2TPv3

October 2003

RC-TBA-3 - Session not established, sequencing required without valid L2-Specific Sublayer.  
RC-TBA-4 - Finite state machine error or timeout.

Additional service-specific Result Codes are defined outside this document.

The Error Codes defined below pertain to types of errors that are not specific to any particular L2TP request, but rather to protocol or message format errors. If an L2TP reply indicates in its Result Code that a general error occurred, the General Error value should be examined to determine what the error was. The currently defined General Error codes and their meanings are as follows:

- 0 - No general error.
- 1 - No control connection exists yet for this pair of LCCEs.
- 2 - Length is wrong.
- 3 - One of the field values was out of range.
- 4 - Insufficient resources to handle this operation now.
- 5 - Invalid Session ID.
- 6 - A generic vendor-specific error occurred.
- 7 - Try another. If initiator is aware of other possible responder destinations, it should try one of them. This can be used to guide an LAC or LNS based on policy.
- 8 - The session or control connection was shut down due to receipt of an unknown AVP with the M bit set (see [Section 5.2](#)). The Error Message SHOULD contain the attribute of the offending AVP in (human-readable) text form.
- 9 - Try another directed. If an LAC or LNS is aware of other possible destinations, it should inform the initiator of the control connection or session. The Error Message MUST contain a comma-separated list of addresses from which the initiator may choose. If the L2TP data channel runs over IPv4, then this would be a comma-separated list of IP addresses in the canonical dotted-decimal format (e.g. "10.0.0.1, 10.0.0.2, 10.0.0.3") in the UTF-8 charset using the Default Language [[RFC2277](#)]. If there are no servers for the LAC or LNS to suggest, then Error Code 7 should be used. The delimiter between addresses MUST be precisely a

single comma and a single space.

When a General Error Code of 6 is used, additional information about the error SHOULD be included in the Error Message field. A vendor-specific AVP MAY be sent to more precisely detail a vendor-specific problem.

### [5.4.3](#) Control Connection Management AVPs

#### Control Connection Tie Breaker (SCCRQ)

The Control Connection Tie Breaker AVP, Attribute Type 5, indicates that the sender desires a single control connection to exist between a given pair of LCCEs.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Control Connection Tie Breaker Value ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
                                                    ... (64 bits)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The Control Connection Tie Breaker Value is an 8-octet random value that is used to choose a single control connection when two LCCEs request a control connection concurrently. The recipient of a SCCRQ must check to see if a SCCRQ has been sent to the peer; if so, a tie has been detected. In this case, the LCCE must compare its Control Connection Tie Breaker value with the one received in the SCCRQ. The lower value "wins", and the "loser" MUST discard its control connection, sending a StopCCN if the SCCRQ that it had sent was acknowledged by the receiving peer. In the case in which a tie breaker is present on both sides and the value is equal, both sides MUST discard their control connections and restart control connection negotiation with a new, random tie breaker value.

If a tie breaker is received and an outstanding SCCRQ has no tie

breaker value, the initiator that included the Control Connection Tie Breaker AVP "wins". If neither side issues a tie breaker, then two separate control connections are opened.

Applications which employ a distinct and well-known initiator have no need for tie-breaking, and this AVP MAY be omitted and the tie-breaking functionality disabled. Applications which require tie-breaking also require that an LCCE be uniquely identifiable upon receipt of an SCCRQ. For L2TP over IP, this MUST be accomplished via the Router ID AVP.

Note that in [[RFC2661](#)], this AVP was referred to as the "Tie-Breaker AVP". Here, the AVP serves the same purpose and has the same attribute value and composition. The name was changed simply to distinguish between the Session and Control Connection Tie Breaker AVP.

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 14.

## Host Name (SCCRQ, SCCRP)

The Host Name AVP, Attribute Type 7, indicates the name of the issuing LAC or LNS, encoded in the US-ASCII charset.

The Attribute Value field for this AVP has the following format:

[illegible]

The Host Name is of arbitrary length, but MUST be at least 1 octet.

This name should be as broadly unique as possible; for hosts participating in DNS [[RFC1034](#)], a host name with fully qualified domain would be appropriate. The Host Name AVP and/or Router ID AVP MUST be used to identify an LCCE as described in [Section 3.3](#).

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for

this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 6 plus the length of the Host Name.

#### Router ID (SCCRQ, SCCRP)

The Router ID AVP, Attribute Type AVP-TBA-2, is an identifier used to identify an LCCE for control connection setup, tie breaking, and/or tunnel authentication.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Router Identifier                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Router Identifier is a 4-octet unsigned integer. Its value is unique for a given LCCE, per [Section 8.1 of \[RFC2072\]](#). The Host Name AVP and/or Router ID AVP MUST be used to identify an LCCE as described in [Section 3.3](#).

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for

this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 10.

#### Vendor Name (SCCRQ, SCCRP)

The Vendor Name AVP, Attribute Type 8, contains a vendor-specific (possibly human-readable) string describing the type of LAC or LNS being used.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Vendor Name ... (arbitrary number of octets)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Vendor Name is the indicated number of octets representing the





the peer earlier (i.e. in the SCCRQ) MUST be sent as the Assigned Control Connection ID AVP in the StopCCN. This policy allows the peer to try to identify the appropriate control connection via a reverse lookup.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

#### Receive Window Size (SCCRQ, SCCRP)

The Receive Window Size AVP, Attribute Type 10, specifies the receive window size being offered to the remote peer.

The Attribute Value field for this AVP has the following format:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|           Window Size           |
+---+---+---+---+---+---+---+---+

```

The Window Size is a 2-octet unsigned integer.

If absent, the peer must assume a Window Size of 4 for its transmit window.

The remote peer may send the specified number of control messages before it must wait for an acknowledgment. See [Section 4.2](#) for more information on reliable control message delivery.


This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 8.

#### Pseudowire Capabilities List (SCCRQ, SCCRP)

The Pseudowire Capabilities List (PW Capabilities List) AVP, Attribute Type AVP-TBA-4, indicates the L2 payload types the sender

can support. The specific payload type of a given session is identified by the Pseudowire Type AVP.

The Attribute Value field for this AVP has the following format:



Defined PW types that may appear in this list are managed by IANA and will appear in associated pseudowire-specific documents for each PW type.


If a sender includes a given PW type in the PW Capabilities List AVP, the sender assumes full responsibility for supporting that particular payload, such as any payload-specific AVPs, L2-Specific Sublayer, or control messages that may be defined in the appropriate companion document.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 8 octets with one PW type specified, plus 2 octets for each additional PW type.

## Preferred Language (SCCRQ, SCCRP)

The Preferred Language AVP, Attribute Type AVP-TBD-14, provides a method for an LCCE to indicate to the peer the language in which human-readable messages it sends SHOULD be composed. This AVP contains a single language tag or language range [[RFC3066](#)].

The Attribute Value field for this AVP has the following format:



The Preferred Language is the indicated number of octets representing

the language tag or language range, encoded in the US-ASCII charset.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 6 plus the length of the Preferred Language.

#### [5.4.4](#) Session Management AVPs

Local Session ID (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN, CDN, WEN, SLI)

The Local Session ID AVP (analogous to the Assigned Session ID in L2TPv2), Attribute Type AVP-TBA-5, contains the identifier being assigned to this session by the sender.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Local Session ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Local Session ID is a 4-octet non-zero unsigned integer.

The Local Session ID AVP establishes the two identifiers used to multiplex and demultiplex sessions between two LCCEs. Each LCCE chooses any free value it desires, and sends it to the remote LCCE using this AVP. The remote LCCE MUST then send all data packets associated with this session using this value. Additionally, for all session-oriented control messages sent after this AVP is received (e.g. ICRP, ICCN, CDN, SLI, etc.), the remote LCCE MUST echo this value in the Remote Session ID AVP.

Note that a Session ID value is unidirectional. Because each LCCE chooses its Session ID independent of its peer LCCE, the value does not have to match in each direction for a given session.

See [Section 4.1](#) for additional information about the Session ID.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be 1 set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

Remote Session ID (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN, CDN, WEN, SLI)

The Remote Session ID AVP, Attribute Type AVP-TBA-6, contains the

identifier that was assigned to this session by the peer.

The Attribute Value field for this AVP has the following format:

```

      0             1             2             3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Remote Session ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The Remote Session ID is a 4-octet non-zero unsigned integer.

The Remote Session ID AVP MUST be present in all session-level control messages. The AVP's value echoes the session identifier advertised by the peer via the Local Session ID AVP. It is the same value that will be used in all transmitted data messages by this side of the session. In most cases, this identifier is sufficient for the peer to look up session-level context for this control message.

When a session-level control message must be sent to the peer before the Local Session ID AVP has been received, the value of the Remote Session ID AVP MUST be set to zero. Additionally, the Local Session ID AVP (sent in a previous control message for this session) MUST be included in the control message. The peer must then use the Local Session ID AVP to perform a reverse lookup to find its session context. Session-level control messages defined in this document that might be subject to a reverse lookup by a receiving peer include the CDN, WEN, and SLI.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5](#), but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

#### Assigned Cookie (ICRQ, ICRP, OCRQ, OCRP)

The Assigned Cookie AVP, Attribute Type AVP-TBA-7, contains the Cookie value being assigned to this session by the sender.

The Attribute Value field for this AVP has the following format:

```

      0             1             2             3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
```

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Assigned Cookie (32 or 64 bits) ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Assigned Cookie is a 4-octet or 8-octet random value.

The Assigned Cookie AVP contains the value used to check the association of a received data message with the session identified by

the Session ID. All data messages sent to a peer MUST use the Assigned Cookie sent by the peer in this AVP. The value's length (0, 32, or 64 bits) is obtained by the Length of the AVP.

A missing Assigned Cookie AVP or Assigned Cookie Value of zero length indicates that the Cookie field should not be present in any data packets sent to the LCCE sending this AVP.

See [Section 4.1](#) for additional information about the Assigned Cookie.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP may be 6, 10, or 14 octets.

## Serial Number (ICRQ, OCRQ)

The Serial Number AVP, Attribute Type 15, contains an identifier assigned by the LAC or LNS to this session.

The Attribute Value field for this AVP has the following format:

```

0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Serial Number
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Serial Number is a 32-bit value.

The Serial Number is intended to be an easy reference for administrators on both ends of a control connection to use when investigating session failure problems. Serial Numbers should be set to progressively increasing values, which are likely to be unique for

a significant period of time across all interconnected LNSs and LACs.

Note that in [RFC 2661](#), this value was referred to as the "Call Serial Number AVP". It serves the same purpose and has the same attribute value and composition.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

## Remote End ID (ICRQ, OCRQ)

The Remote End ID AVP, Attribute Type AVP-TBA-8, contains an identifier used to bind L2TP sessions to a given circuit, interface, or bridging instance. It also may be used to detect session-level

ties.

The Attribute Value field for this AVP has the following format:

[illegible]

The Remote End Identifier field is a variable-length field whose value is unique for a given LCCE peer, as described in [Section 3.3](#).

A session-level tie is detected if an LCCE receives an ICRQ or OCRQ with an End ID AVP whose value matches that which was just sent in an outgoing ICRQ or OCRQ to the same peer. If the two values match, an LCCE recognizes that a tie exists (e.g. both LCCEs are attempting to establish sessions for the same circuit). The tie is broken by the Session Tie Breaker AVP.

By default, the LAC-LAC cross-connect application (see [section 2.0\(b\)](#)) of L2TP over an IP network MUST utilize the Router ID AVP and Remote End ID AVP to associate a circuit to an L2TP session. Other AVPs MAY be used for LCCE or circuit identification as specified in companion documents.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 6 plus the length of the Remote End Identifier value.

#### Application Code (ICRQ, OCRQ)

The Application Code AVP, Attribute Type AVP-TBA-9, is a 2 octet value for enumerating application types for a given L2TP session.

The Attribute Value field for this AVP has the following format:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+---+---+---+---+---+---+---+---+---+
|           Application Code           |
+---+---+---+---+---+---+---+---+

```

The Application Code is a 2 octet value used to identify a specific application for an L2TP session, perhaps causing certain values within AVPs defined in this document to be interpreted or acted upon in a different manner dictated by the Application Code. For example,

a given Application Code could instruct an LCCE to perform a specific directory lookup on the Hostname and/or Router ID AVP information associated with this session (perhaps even encoding the destination address of the given directory server).

An Application Code of 0, or absence of this AVP in any control message, indicates that all AVPs should be interpreted as defined in this document.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 8.

#### Session Tie Breaker (ICRQ, OCRQ)

The Session Tie Breaker AVP, Attribute Type TBD, is used to break ties when two peers concurrently attempt to establish a session for the same circuit.

The Attribute Value field for this AVP has the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Session Tie Breaker Value ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
                                                    ... (64 bits)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Session Tie Breaker Value is an 8-octet random value that is used to choose a session when two LCCEs concurrently request a session for the same circuit. A tie is detected by examining the peer's identity (described in [Section 3.3](#)) plus the per-session shared value communicated via the End ID AVP. In the case of a tie, the recipient of an ICRQ or OCRQ must compare the received tie breaker value with the one that it sent earlier. The LCCE with the lower value "wins", and the "loser" MUST send a CDN with result code set to RC-TBA-1 (as defined in [Section 5.4.2](#)) to tear down the session it instigated. In the case in which a tie is detected, tie breakers are sent by both sides, and the tie breaker values are equal, both sides MUST discard their sessions and restart session negotiation with new random tie breaker values.

If a tie is detected but only one side sends a Session Tie Breaker AVP, the session initiator that included the Session Tie Breaker AVP "wins". If neither side issues a tie breaker, then both sides MUST tear down the session.

This AVP MUST NOT be hidden (the H bit MUST be 0). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length of this AVP is 14.

#### Pseudowire Type (ICRQ, OCRQ)

The Pseudowire Type (PW Type) AVP, Attribute Type AVP-TBA-10, indicates the L2 payload type of the packets that will be tunneled using this L2TP session.

The Attribute Value field for this AVP has the following format:



```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
|           PW Type           |
+---+---+---+---+---+---+---+---+

```

A peer MUST NOT request an incoming or outgoing call with a PW Type AVP specifying a value not advertised in the PW Capabilities List AVP it received during control connection establishment. Attempts to do so MUST result in the call being rejected via a CDN with the Result Code set to RC-TBA-2 (see [Section 5.4.2](#)).

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 8.

## L2-Specific Sublayer (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN)

The L2-Specific Sublayer AVP, Attribute Type AVP-TBA-11, indicates the the presence and format of the L2-Specific Sublayer the sender of this AVP requires on all incoming data packets for this L2TP session.

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
|   L2-Specific Sublayer Type   |
+---+---+---+---+---+---+---+---+

```

The L2-Specific Sublayer Type is a 2-octet unsigned integer with the following values defined in this document:

- 0 - There is no L2-Specific Sublayer present.
- 1 - The default L2-Specific Sublayer (defined in [Section 4.6](#)) is used.

If this AVP is received and has a value other than zero, the

receiving LCCE MUST include the identified L2-Specific Sublayer in its outgoing data messages. If the AVP is not received, it is assumed that there is no sublayer present.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this

AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 8.

#### Data Sequencing (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN)

The Data Sequencing AVP, Attribute Type AVP-TBA-12, indicates that the sender requires some or all of the data packets that it receives to be sequenced.

The Attribute Value field for this AVP has the following format:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
|           Data Sequencing Level           |
+---+---+---+---+---+---+---+---+

```

The Data Sequencing Level is a 2-octet unsigned integer indicating the degree of incoming data traffic that the sender of this AVP wishes to be marked with sequence numbers.

The following values are valid data sequencing levels:

- 0 - No incoming data packets require sequencing.
- 1 - Only non-IP data packets require sequencing.
- 2 - All incoming data packets require sequencing.

If a data sequencing level of 0 is specified, there is no need to send packets with sequence numbers. If sequence numbers are sent, they will be ignored upon receipt. If no Data Sequencing AVP is received, a data sequencing level of 0 is assumed.

If a data sequencing level of 1 is specified, only non-IP traffic carried within the tunneled L2 frame should have sequence numbers applied. Non-IP traffic here refers to any packets that cannot be classified as an IP packet within their respective L2 framing (i.e., a PPP control packet or NETBIOS frame encapsulated by Frame Relay before being tunneled). All traffic that can be classified as IP MUST be sent with no sequencing (e.g. the S bit in the L2-Specific Sublayer is set to zero). If a packet is unable to be classified at all (e.g. due to it being compressed or encrypted at layer 2) or if an implementation is unable to perform such classification within L2 frames, all packets MUST be provided with sequence numbers

(essentially falling back to a data sequencing level of 2).

If a data sequencing level of 2 is specified, all traffic MUST be sequenced.

Data sequencing may only be requested when there is an L2-Specific Sublayer present that can provide sequence numbers. If sequencing is requested without requesting a L2-Specific Sublayer AVP, the session MUST be disconnected with a Result Code of RC-TBA-3.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 1, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 6.

#### Tx Connect Speed (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN)

The Tx Connect Speed BPS AVP, Attribute Type 24, contains the speed of the facility chosen for the connection attempt.

The Attribute Value field for this AVP has the following format:

0								1								2								3							
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
Tx Connect Speed BPS																															

The Tx Connect Speed BPS is a 4-octet value indicating the speed in bits per second. A value of zero indicates that the speed is indeterminable or that there is no physical point-to-point link.

When the optional Rx Connect Speed AVP is present, the value in this AVP represents the transmit connect speed from the perspective of the LAC (e.g. data flowing from the LAC to the remote system). When the optional Rx Connect Speed AVP is NOT present, the connection speed between the remote system and LAC is assumed to be symmetric and is represented by the single value in this AVP.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

#### Rx Connect Speed (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN)

The Rx Connect Speed AVP, Attribute Type 38, represents the speed of the connection from the perspective of the LAC (e.g. data flowing from the remote system to the LAC).

INTERNET DRAFT

L2TPv3

October 2003

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Rx Connect Speed BPS                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Rx Connect Speed BPS is a 4-octet value indicating the speed in bits per second. A value of zero indicates that the speed is indeterminable or that there is no physical point-to-point link.

Presence of this AVP implies that the connection speed may be asymmetric with respect to the transmit connect speed given in the Tx Connect Speed AVP.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

#### Physical Channel ID (ICRQ, ICRP, OCRP)

The Physical Channel ID AVP, Attribute Type 25, contains the vendor-specific physical channel number used for a call.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Physical Channel ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Physical Channel ID is a 4-octet value intended to be used for logging purposes only.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 10.

#### [5.4.5](#) Circuit Status AVPs

The Circuit Status AVP, Attribute Type AVP-TBA-13, indicates the initial status of or a status change in the circuit to which the session is bound.

The Attribute Value field for this AVP has the following format:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|           Reserved           |N|A|
+---+---+---+---+---+---+---+---+---+

```

The A (Active) bit indicates whether the circuit is up/active/ready (1) or down/inactive/not-ready (0).

The N (New) bit indicates whether the circuit status indication is for a new circuit (1) or an existing circuit (0). Links which have a similar mechanism available (e.g. Frame Relay) MUST map the setting of this bit to the associated signaling for that link. Otherwise, the New bit SHOULD still be set the first time the L2TP session is established after provisioning.

The remaining bits are reserved for future use. Reserved bits MUST be set to 0 when sending and ignored upon receipt.

The Circuit Status AVP is used to advertise whether a circuit or interface bound to an L2TP session is up and ready to send and/or receive traffic. Different circuit types have different names for status types. For example, HDLC primary and secondary stations refer to a circuit as being "Receive Ready" or "Receive Not Ready", while Frame Relay refers to a circuit as "Active" or "Inactive". This AVP adopts the latter terminology, though the concept remains the same regardless of the PW type for the L2TP session.

In the simplest case, the circuit referred by this AVP is a single physical interface, port, or circuit depending on application and how the session was setup. The status indication in this AVP may then be used to provide simple ILMI interworking for a variety of circuit types. For virtual or multipoint interfaces, the Circuit Status AVP



```

|                                     Buffer Overruns                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Timeout Errors                                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Alignment Errors                                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The following fields are defined:

Reserved: 2 octets of Reserved data is present (providing longword alignment within the AVP of the following values). Reserved data MUST be zero on sending and ignored upon receipt.

Hardware Overruns: Number of receive buffer overruns since call was established.

Buffer Overruns: Number of buffer overruns detected since call was established.

Timeout Errors: Number of timeouts since call was established.

Alignment Errors: Number of alignment errors since call was established.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see [Section 5.2](#)). The Length (before hiding) of this AVP is 32.

## [6.](#) Control Connection Protocol Specification

The following control messages are used to establish, maintain, and tear down L2TP control connections. All data packets are sent in network order (high-order octets first). Any "reserved" or "empty" fields MUST be sent as 0 values to allow for protocol extensibility.

The exchanges in which these messages are involved are outlined in [Section 3.3](#).

### [6.1](#) Start-Control-Connection-Request (SCCRQ)

Start-Control-Connection-Request (SCCRQ) is a control message used to initiate a control connection between two LCCEs. It is sent by either the LAC or the LNS to begin the control connection establishment process.

The following AVPs MUST be present in the SCCRQ:

- Message Type
- Host Name
- Router ID
- Assigned Control Connection ID
- Pseudowire Capabilities List

The following AVPs MAY be present in the SCCRQ:

- Random Vector
- Nonce
- Message Digest
- Control Connection Tie Breaker
- Vendor Name
- Receive Window Size
- Preferred Language

## [6.2](#) Start-Control-Connection-Reply (SCCRP)

Start-Control-Connection-Reply (SCCRP) is the control message sent in reply to a received SCCRQ message. The SCCRP is used to indicate that the SCCRQ was accepted and establishment of the control connection should continue.

The following AVPs MUST be present in the SCCRP:

- Message Type
- Host Name
- Router ID
- Assigned Control Connection ID
- Pseudowire Capabilities List

The following AVPs MAY be present in the SCCRP:

- Random Vector
- Nonce
- Message Digest
- Vendor Name
- Receive Window Size
- Preferred Language



### [6.3](#) Start-Control-Connection-Connected (SCCCN)

Start-Control-Connection-Connected (SCCCN) is the control message sent in reply to an SCCRP. The SCCCN completes the control connection establishment process.

The following AVP MUST be present in the SCCCN:

Message Type

The following AVP MAY be present in the SCCCN:

Random Vector  
Message Digest

### [6.4](#) Stop-Control-Connection-Notification (StopCCN)

Stop-Control-Connection-Notification (StopCCN) is the control message sent by either LCCE to inform its peer that the control connection is being shut down and that the control connection should be closed. In addition, all active sessions are implicitly cleared (without sending any explicit session control messages). The reason for issuing this request is indicated in the Result Code AVP. There is no explicit reply to the message, only the implicit ACK that is received by the reliable control message delivery layer.

The following AVPs MUST be present in the StopCCN:

Message Type  
Result Code

The following AVPs MAY be present in the StopCCN:

Random Vector  
Message Digest  
Assigned Control Connection ID

Note that the Assigned Control Connection ID MUST be present if the StopCCN is sent after an SCCRQ or SCCRP message has been sent.

## [6.5](#) Hello (HELLO)

The Hello (HELLO) message is an L2TP control message sent by either peer of a control connection. This control message is used as a "keepalive" for the control connection. See [Section 4.2](#) for a description of the keepalive mechanism.

HELLO messages are global to the control connection. The Session ID in a HELLO message MUST be 0.

The following AVP MUST be present in the HELLO:

Message Type

The following AVP MAY be present in the HELLO:

Random Vector  
Message Digest

## [6.6](#) Incoming-Call-Request (ICRQ)

Incoming-Call-Request (ICRQ) is the control message sent by an LCCE to a peer when an incoming call is detected (although the ICRQ may also be sent as a result of a local event). It is the first in a three-message exchange used for establishing a session via an L2TP control connection.

The ICRQ is used to indicate that a session is to be established between an LCCE and a peer. The sender of an ICRQ provides the peer with parameter information for the session. However, the sender makes no demands about how the session is terminated at the peer (i.e. whether the L2 traffic is processed locally, forwarded, etc.).

The following AVPs MUST be present in the ICRQ:

Message Type  
Local Session ID  
Remote Session ID  
Serial Number  
Pseudowire Type  
Circuit Status

The following AVPs MAY be present in the ICRQ:

- Random Vector
- Message Digest
- Assigned Cookie
- Remote End ID
- Application ID
- Session Tie Breaker
- L2-Specific Sublayer
- Data Sequencing
- Tx Connect Speed
- Rx Connect Speed
- Physical Channel ID

#### [6.7](#) Incoming-Call-Reply (ICRP)

Incoming-Call-Reply (ICRP) is the control message sent by an LCCE in response to a received ICRQ. It is the second in the three-message exchange used for establishing sessions within an L2TP control connection.

The ICRP is used to indicate that the ICRQ was successful and that the peer should establish (i.e. answer) the incoming call if it has not already done so. It also allows the sender to indicate specific parameters about the L2TP session.

The following AVPs MUST be present in the ICRP:

- Message Type
- Local Session ID
- Remote Session ID
- Circuit Status

The following AVPs MAY be present in the ICRP:

- Random Vector
- Message Digest
- Assigned Cookie
- L2-Specific Sublayer
- Data Sequencing
- Tx Connect Speed
- Rx Connect Speed
- Physical Channel ID

#### [6.8](#) Incoming-Call-Connected (ICCN)

Incoming-Call-Connected (ICCN) is the control message sent by the LCCE that originally sent an ICRQ upon receiving an ICRP from its

INTERNET DRAFT

L2TPv3

October 2003

peer. It is the final message in the three-message exchange used for establishing L2TP sessions.

The ICCN is used to indicate that the ICRP was accepted, that the call has been established, and that the L2TP session should move to the established state. It also allows the sender to indicate specific parameters about the established call (parameters that may not have been available at the time the ICRQ is issued).

The following AVPs MUST be present in the ICCN:

- Message Type
- Local Session ID
- Remote Session ID

The following AVPs MAY be present in the ICCN:

- Random Vector
- Message Digest
- L2-Specific Sublayer
- Data Sequencing
- Tx Connect Speed
- Rx Connect Speed
- Circuit Status

## [6.9](#) Outgoing-Call-Request (OCRQ)

Outgoing-Call-Request (OCRQ) is the control message sent by an LCCE to an LAC to indicate that an outbound call at the LAC is to be established based on specific destination information sent in this message. It is the first in a three-message exchange used for establishing a session and placing a call on behalf of the initiating LCCE.

Note that a call may be any L2 connection requiring well-known destination information to be sent from an LCCE to an LAC. This call could be a dialup connection to the PSTN, an SVC connection, the IP address of another LCCE, or any other destination dictated by the sender of this message.

The following AVPs MUST be present in the OCRQ:

- Message Type

Local Session ID  
Remote Session ID  
Serial Number  
Pseudowire Type  
Circuit Status

INTERNET DRAFT

L2TPv3

October 2003

The following AVPs MAY be present in the OCRQ:

Random Vector  
Message Digest  
Assigned Cookie  
Remote End ID  
Application ID  
Tx Connect Speed  
Rx Connect Speed  
Session Tie Breaker  
L2-Specific Sublayer  
Data Sequencing

#### [6.10](#) Outgoing-Call-Reply (OCRP)

Outgoing-Call-Reply (OCRP) is the control message sent by an LAC to an LCCE in response to a received OCRQ. It is the second in a three-message exchange used for establishing a session within an L2TP control connection.

OCRP is used to indicate that the LAC has been able to attempt the outbound call. The message returns any relevant parameters regarding the call attempt. Data MUST NOT be forwarded until the OCCN is received indicating that the call has been placed.

The following AVPs MUST be present in the OCRP:

Message Type  
Local Session ID  
Remote Session ID  
Circuit Status

The following AVPs MAY be present in the OCRP:

Random Vector  
Message Digest

- Assigned Cookie
- L2-Specific Sublayer
- Tx Connect Speed
- Rx Connect Speed
- Data Sequencing
- Physical Channel ID

#### [6.11](#) Outgoing-Call-Connected (OCCN)

Outgoing-Call-Connected (OCCN) is the control message sent by an LAC to another LCCE after the OCRP and after the outgoing call has been completed. It is the final message in a three-message exchange used

for establishing a session.

OCCN is used to indicate that the result of a requested outgoing call was successful. It also provides information to the LCCE who requested the call about the particular parameters obtained after the call was established.

The following AVPs MUST be present in the OCCN:

- Message Type
- Local Session ID
- Remote Session ID

The following AVPs MAY be present in the OCCN:

- Random Vector
- Message Digest
- L2-Specific Sublayer
- Tx Connect Speed
- Rx Connect Speed
- Data Sequencing
- Circuit Status

#### [6.12](#) Call-Disconnect-Notify (CDN)

The Call-Disconnect-Notify (CDN) is a control message sent by an LCCE to request disconnection of a specific session. Its purpose is to inform the peer of the disconnection and the reason for the disconnection. The peer MUST clean up any resources, and does not

send back any indication of success or failure for such cleanup.

The following AVPs MUST be present in the CDN:

- Message Type
- Result Code
- Local Session ID
- Remote Session ID

The following AVP MAY be present in the CDN:

- Random Vector
- Message Digest

### [6.13](#) WAN-Error-Notify (WEN)

The WAN-Error-Notify (WEN) is a control message sent from an LAC to an LNS to indicate WAN error conditions. The counters in this message are cumulative. This message should only be sent when an error

occurs, and not more than once every 60 seconds. The counters are reset when a new call is established.

The following AVPs MUST be present in the WEN:

- Message Type
- Local Session ID
- Remote Session ID
- Circuit Errors

The following AVP MAY be present in the WEN:

- Random Vector
- Message Digest

### [6.14](#) Set-Link-Info (SLI)

The Set-Link-Info control message is sent by an LCCE to convey link or circuit status change information regarding the circuit associated with this L2TP session. For example, if PPP renegotiates LCP at an LNS or between an LAC and a Remote System, or if a forwarded Frame Relay VC transitions to Active or Inactive at an LAC, an SLI message

SHOULD be sent to indicate this event. Precise details of when the SLI is sent, what PW type-specific AVPs must be present, and how those AVPs should be interpreted by the receiving peer are outside the scope of this document. These details should be described in the associated pseudowire-specific documents that require use of this message.

The following AVPs MUST be present in the SLI:

- Message Type
- Local Session ID
- Remote Session ID

The following AVPs MAY be present in the SLI:

- Random Vector
- Message Digest
- Circuit Status

#### [6.15](#) Explicit-Acknowledgement (ACK)

The Explicit Acknowledgement (ACK) message is used only to acknowledge receipt of a message or messages on the Control Connection (e.g. for purposes of updating Ns and Nr values). Receipt of this message does not trigger an event for the L2TP protocol state machine.

A message received without any AVPs (including the Message Type AVP), is referred to as a Zero Length Body (ZLB) message, and serves the same function as the Explicit Acknowledgement. ZLB messages are only permitted when the Control Message Authentication defined in [Section 4.3](#) is not enabled.

The following AVPs MAY be present in the ACK message:

- Message Type
- Message Digest

mi.sp

## [7](#). Control Connection State Machines

The state tables defined in this section govern the exchange of control messages defined in [Section 6](#). Tables are defined for



incoming call placement and outgoing call placement, as well as for initiation of the control connection itself. The state tables do not encode message timeout and retransmission behavior, as this is handled in the underlying reliable control message delivery mechanism (see [Section 4.2](#)).

Timers MAY be employed to enforce a maximum period of time allowed in a transitional state (i.e. between idle and established). Generally, this is a protection mechanism for cleaning up state when an error has occurred, and should be treated as such. The precise period of time to wait is application dependent and MUST be configurable, with the notable default of no timeout at all. If a session is shutdown by a state transition timeout, a CDN MUST be sent with a Result Code set to RC-TBA-4 (see [Section 5.4.2](#)). If a control connection is shutdown by a state transition timeout, a StopCCN MUST be sent with a Result Code set to 7 (see [Section 5.4.2](#)).

## [7.1](#) Malformed Control Messages

Receipt of an invalid or unrecoverable malformed control message SHOULD be logged appropriately and the control connection cleared to ensure recovery to a known state. The control connection may then be restarted by the initiator. An invalid control message is defined as (1) a message that contains a Message Type marked as mandatory (see [Section 5.4.1](#)) but that is unknown to the implementation, or (2) a control message that is received in the wrong state.

Examples of malformed control messages include (1) a message that has an invalid value in its header, (2) a message that contains an AVP that is formatted incorrectly or whose value is out of range, and (3) a message that is missing a required AVP. A control message with a malformed header MUST be discarded.

If a malformed AVP is received with the M bit set, the session or control connection MUST be terminated with a proper Result or Error Code sent (see [Section 5.2](#)). A malformed yet non-mandatory (M bit is not set) AVP within a control message should be handled like an unrecognized non-mandatory AVP. That is, the AVP MUST be ignored (with the exception of logging a local error message), and the message MUST be accepted.

It is impossible to list all potential malformations of a given message. However, an example of a recoverable, malformed AVP might be if the Rx Connect Speed AVP, attribute 38, is received with a length of 8 rather than 10, and the BPS given in 2 octets rather than 4. In the Rx Connect Speed is sent with the M bit set to 0, this malformation should not be considered catastrophic. As such, the control message should be accepted as if the AVP had not been received (with the exception of a local error message being logged). This example is by no means a license to create malformed AVPs, but simply a guideline for how liberal one should be in acceptance of messages containing errors.

In the following tables, there are several cases where a protocol message is sent and then a "clean up" occurs. Note that, regardless of the initiator of the control connection shutdown, the reliable delivery mechanism must be allowed to run (see [Section 4.2](#)) before destroying the control connection. This permits the control connection management messages to be reliably delivered to the peer.

[Appendix B.1](#) contains an example of lock-step control connection establishment.

## 7.2 Control Connection States

The L2TP control connection protocol is not distinguishable between the two LCCEs but is distinguishable between the originator and receiver. The originating peer is the one that first initiates establishment of the control connection. (In a tie breaker situation, this is the winner of the tie.) Since either the LAC or the LNS can be the originator, a collision can occur. See the Control Connection Tie Breaker AVP in [Section 5.4.3](#) for a description of this and its resolution.

State -----	Event -----	Action -----	New State -----
idle	Local open request	Send SCCRQ	wait-ctl-reply
idle	Receive SCCRQ, acceptable	Send SCCRP	wait-ctl-conn

idle	Receive SCCRQ,	Send StopCCN,	idle
------	----------------	---------------	------

	not acceptable	clean up	
idle	Receive SCCRP	Send StopCCN, clean up	idle
idle	Receive SCCCN	Send StopCCN, clean up	idle
wait-ctl-reply	Receive SCCRP, acceptable	Send SCCCN, send control-conn open event to waiting sessions	established
wait-ctl-reply	Receive SCCRP, not acceptable	Send StopCCN, clean up	idle
wait-ctl-reply	Receive SCCRQ, lose tie breaker, SCCRQ acceptable	Send SCCRP, Clean up losing connection	wait-ctl-conn
wait-ctl-reply	Receive SCCRQ, lose tie breaker, SCCRQ unacceptable	Send StopCCN, Clean up losing connection	idle
wait-ctl-reply	Receive SCCRQ, win tie breaker	Send StopCCN for losing connection	wait-ctl-reply
wait-ctl-reply	Receive SCCCN	Send StopCCN, clean up	idle
wait-ctl-conn	Receive SCCCN, acceptable	Send control-conn open event to waiting sessions	established
wait-ctl-conn	Receive SCCCN, not acceptable	Send StopCCN, clean up	idle
wait-ctl-conn	Receive SCCRP, SCCRQ	Send StopCCN, clean up	idle
established	Local open request (new call)	Send control-conn open event to waiting sessions	established
established	Administrative control-conn close event	Send StopCCN, clean up	idle

established	Receive SCCRQ, SCCRP, SCCCN	Send StopCCN, clean up	idle
idle, wait-ctl-reply, wait-ctl-conn, established	Receive StopCCN	Clean up	idle

The states associated with an LCCE for control connection establishment are as follows:

#### idle

Both initiator and recipient start from this state. An initiator transmits an SCCRQ, while a recipient remains in the idle state until receiving an SCCRQ.

#### wait-ctl-reply

The originator checks to see if another connection has been requested from the same peer, and if so, handles the collision situation described in [Section 5.4.3](#).

#### wait-ctl-conn

Awaiting an SCCCN. Upon receipt, the challenge response contained in the message is checked. The control connection is established if authentication succeeds; otherwise, it is torn down.

#### established

An established connection may be terminated by either a local condition or the receipt of a StopCCN. In the event of a local termination, the originator MUST send a StopCCN and clean up the control connection. If the originator receives a StopCCN, it MUST also clean up the control connection.

### [7.3](#) Incoming Calls

An ICRQ is generated by an LCCE, typically in response to an incoming call or a local event. Once the LCCE sends the ICRQ, it waits for a response from the peer. However, it may choose to postpone establishment of the call (e.g. answering the call, bringing up the circuit) until the peer has indicated with an ICRP that it will accept the call. The peer may choose not to accept the call if, for instance, there are insufficient resources to handle an additional session.

If the peer chooses to accept the call, it responds with an ICRP. When the local LCCE receives the ICRP, it attempts to establish the

call. A final call connected message, the ICCN, is sent from the local LCCE to the peer to indicate that the call states for both

INTERNET DRAFT

L2TPv3

October 2003

LCCEs should enter the established state. If the call is terminated before the peer can accept it, a CDN is sent by the local LCCE to indicate this condition.

When a call transitions to a "disconnected" or "down" state, the call is cleared normally, and the local LCCE sends a CDN. Similarly, if the peer wishes to clear a call, it sends a CDN and cleans up its session.

#### [7.3.1](#) ICRQ Sender States

State -----	Event -----	Action -----	New State -----
idle	Call signal or ready to receive incoming conn	Initiate local control-conn open	wait-control-conn
idle	Receive ICCN, ICRP, CDN	Clean up	idle
wait-control- conn	Bearer line drop or local close request	Clean up	idle
wait-control- conn	control-conn-open	Send ICRQ	wait-reply
wait-reply	Receive ICRP, acceptable	Send ICCN	established
wait-reply	Receive ICRP, Not acceptable	Send CDN, clean up	idle
wait-reply	Receive ICRQ, lose tie breaker	Process as ICRQ Recipient ( <a href="#">Section 7.3.2</a> )	idle
wait-reply	Receive ICRQ,	Send CDN	wait-reply

	win tie breaker	for losing session	
wait-reply	Receive CDN, ICCN	Clean up	idle
wait-reply	Local close request	Send CDN, clean up	idle

established	Receive CDN	Clean up	idle
established	Receive ICRQ, ICRP, ICCN	Send CDN, clean up	idle
established	Local close request	Send CDN, clean up	idle

The states associated with the ICRQ sender are as follows:

#### idle

The LCCE detects an incoming call on one of its interfaces (e.g. an analog PSTN line rings, or an ATM PVC is provisioned), or a local event occurs. The LCCE initiates its control connection establishment state machine and moves to a state waiting for confirmation of the existence of a control connection.

#### wait-control-conn

In this state, the session is waiting for either the control connection to be opened or for verification that the control connection is already open. Once an indication that the control connection has been opened is received, session control messages may be exchanged. The first of these messages is the ICRQ.

#### wait-reply

The ICRQ sender receives either (1) a CDN indicating the peer is not willing to accept the call (general error or do not accept) and moves back into the idle state, or (2) an ICRP indicating the call is accepted. In the latter case, the LCCE sends an ICCN and enters the established state.

#### established

Data is exchanged over the session. The call may be cleared by any of the following:

- + An event on the connected interface: The LCCE sends a CDN.
- + Receipt of a CDN: The LCCE cleans up, disconnecting the call.
- + A local reason: The LCCE sends a CDN.

### [7.3.2](#) ICRQ Recipient States

State -----	Event -----	Action -----	New State -----
idle	Receive ICRQ, acceptable	Send ICRP	wait-connect
idle	Receive ICRQ, not acceptable	Send CDN, clean up	idle

idle	Receive ICRP	Send CDN clean up	idle
idle	Receive ICCN	Clean up	idle
wait-connect	Receive ICCN acceptable	Prepare for data	established
wait-connect	Receive ICCN not acceptable	Send CDN, clean up	idle
wait-connect	Receive ICRQ, ICRP	Send CDN, clean up	idle
idle, wait-connect, established	Receive CDN	Clean up	idle
wait-connect established	Local close request	Send CDN, clean up	idle
established	Receive ICRQ, ICRP, ICCN	Send CDN, clean up	idle

The states associated with the ICRQ recipient are as follows:

idle

An ICRQ is received. If the request is not acceptable, a CDN is sent back to the peer LCCE, and the local LCCE remains in the idle state. If the ICRQ is acceptable, an ICRP is sent. The session moves to the wait-connect state.

wait-connect

The local LCCE is waiting for an ICCN from the peer. Upon receipt of the ICCN, the local LCCE moves to established state.

established

The session is terminated either by sending a CDN or by receiving a CDN from the peer. Clean up follows on both sides regardless of the initiator.

## [7.4](#) Outgoing Calls

Outgoing calls instruct an LAC to place a call. There are three messages for outgoing calls: OCRQ, OCRP, and OCCN. An LCCE first sends an OCRQ to an LAC to request an outgoing call. The LAC MUST respond to the OCRQ with an OCRP once it determines that the proper facilities exist to place the call and that the call is

administratively authorized. Once the outbound call is connected, the LAC sends an OCCN to the peer indicating the final result of the call attempt.

### [7.4.1](#) OCRQ Sender States

State -----	Event -----	Action -----	New State -----
idle	Local open request	Initiate local control-conn-open	wait-control-conn
idle	Receive OCCN, OCRP	Clean up	idle
wait-control-conn	control-conn-open	Send OCRQ	wait-reply
wait-reply	Receive OCRP,	none	wait-connect



	acceptable		
wait-reply	Receive OCRP, not acceptable	Send CDN, clean up	idle
wait-reply	Receive OCCN	Send CDN, clean up	idle
wait-reply	Receive OCRQ, lose tie breaker	Process as OCRQ Recipient ( <a href="#">Section 7.4.2</a> )	idle
wait-reply	Receive OCRQ, win tie breaker	Send CDN for losing session	wait-reply
wait-connect	Receive OCCN	none	established
wait-connect	Receive OCRQ, OCRP	Send CDN, clean up	idle
idle, wait-reply, wait-connect, established	Receive CDN	Clean up	idle
established	Receive OCRQ, OCRP, OCCN	Send CDN, clean up	idle

wait-reply, wait-connect, established	Local close request	Send CDN, clean up	idle
wait-control- conn	Local close request	Clean up	idle

The states associated with the OCRQ sender are as follows:

idle, wait-control-conn

When an outgoing call request is initiated, a control connection

is created as described above, if not already present. Once the control connection is established, an OCRQ is sent to the LAC, and the session moves into the wait-reply state.

wait-reply

If a CDN is received, the session is cleaned up and returns to idle state. If an OCRP is received, the call is in progress, and the session moves to the wait-connect state.

wait-connect

If a CDN is received, the session is cleaned up and returns to idle state. If an OCCN is received, the call has succeeded, and the session may now exchange data.

established

If a CDN is received, the session is cleaned up and returns to idle state. Alternatively, if the LCCE chooses to terminate the session, it sends a CDN to the LAC, cleans up the session, and moves the session to idle state.

7.4.2 OCRQ Recipient (LAC) States

State	Event	Action	New State
-----	-----	-----	-----
idle	Receive OCRQ, acceptable	Send OCRP, Place call	wait-cs-answer
idle	Receive OCRQ, not acceptable	Send CDN, clean up	idle
idle	Receive OCRP	Send CDN, clean up	idle
idle	Receive OCCN, CDN	Clean up	idle

INTERNET DRAFT		L2TPv3	October 2003
wait-cs-answer	Call placement successful	Send OCCN	established
wait-cs-answer	Call placement	Send CDN,	idle

	failed	clean up	
wait-cs-answer	Receive OCRQ, OCRP, OCCN	Send CDN, clean up	idle
established	Receive OCRQ, OCRP, OCCN	Send CDN, clean up	idle
wait-cs-answer, established	Receive CDN	Clean up	idle
wait-cs-answer, established	Local close request	Send CDN, clean up	idle

The states associated with the LAC for outgoing calls are as follows:

#### idle

If the OCRQ is received in error, respond with a CDN. Otherwise, place the call, send an OCRP, and move to the wait-cs-answer state.

#### wait-cs-answer

If the call is not completed or a timer expires while waiting for the call to complete, send a CDN with the appropriate error condition set, and go to idle state. If a circuit-switched connection is established, send an OCCN indicating success, and go to established state.

#### established

If the LAC receives a CDN from the peer, the call MUST be released via appropriate mechanisms, and the session cleaned up. If the call is disconnected because the circuit transitions to a "disconnected" or "down" state, the LAC MUST send a CDN to the peer and return to idle state.

## 7.5 Termination of a Control Connection

The termination of a control connection consists of either peer issuing a StopCCN. The sender of this message SHOULD wait a full control message retransmission cycle (e.g. 1 + 2 + 4 + 8 ... seconds) for the acknowledgment of this message before releasing the control information associated with the control connection. The recipient of this message should send an acknowledgment of the message to the

peer, then release the associated control information.

When to release a control connection is an implementation issue and is not specified in this document. A particular implementation may use whatever policy is appropriate for determining when to release a control connection. Some implementations may leave a control connection open for a period of time or perhaps indefinitely after the last session for that control connection is cleared. Others may choose to disconnect the control connection immediately after the last call on the control connection disconnects.

## [8.](#) Security Considerations

This section addresses some of the security issues that L2TP encounters in its operation.

### [8.1](#) Control Connection Endpoint and Message Security

The LCCEs may configure a shared secret (password) in order to perform a mutual authentication of one another, and construct an authentication and integrity check of all arriving Control Messages. This mechanism is built-in to L2TPv3, and is described in [section 4.3](#) and in the definition of the Message Digest and Nonce AVPs in [section 5.4.3](#).

This mechanism provides strong mutual peer authentication, and authentication and integrity checking for individual Control Messages.

### [8.2](#) Data Channel Security

As described in [section 4.1](#), the Assigned Cookie sent with each data packet MUST be selected in an unpredictable manner (with the added restriction that two same Cookie values not be selected within a short period of time for a given Session ID).

A 64-bit Cookie provides effective protection against a blind packet insertion attack on a given PE. This is useful as a security feature only within networks where sniffing and correlating packets between L2TP nodes is considered impossible, though inserting IP packets destined to an LCCE may be considered possible (and perhaps trivial by an individual armed with the proper hacking tools). In such cases, the Cookie provides an effective barrier against packet insertion into a VPN by enforcing that a given Session ID match the random 64 bit Cookie. A 32 bit Cookie is vulnerable to brute force guessing at high packet rates, and as such should not be considered an effective barrier to insertion attacks (it still provides an additional

integrity check for the Session ID, as described in [section 4.1](#)).

INTERNET DRAFT

L2TPv3

October 2003

The L2TPv3 Cookie MUST NOT be regarded as a substitute for packet-level security such as that of IPsec when operating over an open or untrusted network where packets may be sniffed and values correlated to spoofed packets. L2TPv3 does not attempt to provide data packet encryption of any kind (without the aid of IPsec).

### [8.3](#) End-to-End Security

Protecting the L2TP packet stream with IPsec does, in turn, also protect the data within the tunneled session packets while transported from one LCCE to the other. Such protection MUST NOT be considered a substitution for end-to-end security between communicating hosts or applications.

### [8.4](#) L2TP and IPsec

When running over IP, IPsec [[RFC2401](#)] provides packet-level security via ESP [[RFC3193](#)]. All L2TP control and data packets for a particular control connection appear as homogeneous UDP or IP data packets to the IPsec system.

In addition to IP transport security, IPsec defines a mode of operation that allows tunneling of IP packets. The packet-level encryption and authentication provided by IPsec tunnel mode and that provided by L2TP secured with IPsec provide an equivalent level of security for these requirements.

IPsec also defines access control features that are required of a compliant IPsec implementation. These features allow filtering of packets based upon network and transport layer characteristics such as IP address, ports, etc. In the L2TP tunneling model, analogous filtering is logically performed at the network layer above L2TP. These network layer access control features may be handled at an LCCE via vendor-specific authorization features based upon the authenticated user, or at the network layer itself by using IPsec transport mode end-to-end between the communicating hosts. The requirements for access control mechanisms are not a part of the L2TP specification and as such are outside the scope of this document.

### [8.5](#) Impact of L2TPv3 Features on [RFC 3193](#)

[RFC3193] defines the recommended method for securing L2TPv2. L2TPv3 possesses identical characteristics to IPsec as L2TPv2 when running on UDP/IP. When operating over IP directly, the principles defined in [\[RFC3193\]](#) still apply, though references to UDP port selection (in particular [Section 4](#) "IPsec Filtering details when protecting L2TP") become far simpler as there are two less variable parameters (source and destination UDP ports) to be concerned with when applying

filters. Specific details for operating L2TPv3 with IPsec will be specified in an update to [\[RFC3193\]](#).

## [9.](#) Internationalization Considerations

The Host Name and Vendor Name AVPs are not internationalized. The Vendor Name AVP, although intended to be human-readable, would seem to fit in the category of "globally visible names" [\[RFC2277\]](#) and so is represented in US-ASCII.

The Preferred Language AVP is not mandatory. If an LCCE does not signify a language preference by the inclusion of this AVP in the SCCRQ or SCCRQ, the Preferred Language AVP is unrecognized, or the requested language is not supported by the peer LCCE, the default language [\[RFC2277\]](#) MUST be used for all internationalized strings sent by the peer.

## [10.](#) IANA Considerations

This document defines a number of "magic" numbers to be maintained by the IANA. This section explains the criteria to be used by the IANA to assign additional numbers in each of these lists. The following subsections describe the assignment policy for the namespaces defined elsewhere in this document.

### [10.1](#) Control Message Attribute Value Pairs (AVPs)

This number space is managed by IANA as per [\[RFC3438\]](#).

New AVPs requiring assignment in this document are defined in the "AVP Summary," [Section 5.4](#), with the encoding "AVP-TBA-x," where "x" is 1, 2, 3...

A summary of the new AVPs follows:

AVP-TBA-1	Message Digest
AVP-TBA-2	Router ID,
AVP-TBA-3	Assigned Control Connection ID
AVP-TBA-4	Pseudowire Capabilities List
AVP-TBA-5	Local Session ID
AVP-TBA-6	Remote Session ID
AVP-TBA-7	Assigned Cookie
AVP-TBA-8	Remote End ID
AVP-TBA-9	Application Code
AVP-TBA-10	Pseudowire Type
AVP-TBA-11	L2-Specific Sublayer
AVP-TBA-12	Data Sequencing
AVP-TBA-13	Circuit Status

AVP-TBA-14	Preferred Language
AVP-TBA-15	Control Message Authentication Nonce

## [10.2](#) Message Type AVP Values

This number space is managed by IANA as per [[RFC3438](#)]. There is one new message type, defined in [section 3.1](#), necessary to be allocated for this specification:

TBA-M1 (ACK)	Explicit Acknowledgement
--------------	--------------------------

## [10.3](#) Result Code AVP Values

This number space is managed by IANA as per [[RFC3438](#)].

New Result Code values for the CDN message are defined in [section 5.4](#). Following is a summary:

RC-TBA-1	- Session not established due to losing tie breaker.
RC-TBA-2	- Session not established due to unsupported PW type.
RC-TBA-3	- Session not established, sequencing required without valid L2-Specific Sublayer.

There are a few cases in [Section 5](#) where these values are referred to directly within the document text with the RC-TBA-x format. The assigned values should be inserted within the text for these cases.

### [10.3.2](#) Error Code Field Values

This number space is managed by IANA as per [\[RFC3438\]](#).

### [10.4](#) AVP Header Bits

There are four remaining reserved bits in the AVP header. Additional bits should only be assigned via a Standards Action [\[RFC2434\]](#).

### [10.5](#) L2TP Control Message Header Bits

There are nine remaining reserved bits in the control message header. Additional bits should only be assigned via a Standards Action [\[RFC2434\]](#).

Care should be taken before using reserved bits 6 and 7 in the L2TPv3 control message header since these bits have meaning for L2TPv2 data messages. Using these two bits in L2TPv3 MAY trigger an unforeseen interoperability problem with L2TPv3 implementations based on L2TPv2. Therefore, it is recommended that these two bits be utilized last, after the other reserved bits have been assigned roles.

### [10.6](#) Pseudowire Types

The Pseudowire Type (PW Type, [Section 5.4](#)) is a two-octet value used in the Pseudowire Type AVP and Pseudowire Capabilities List AVP defined in [Section 5.4.3](#). 0 to 32767 are assignable by Expert Review [\[RFC2434\]](#), 32768 to 65535 by a First Come First Served policy [\[RFC2434\]](#). There are no specific pseudowire types assigned within this document. Each pseudowire-specific document MUST allocate its own PW types from IANA as necessary.

### [10.7](#) Application Code

The Application Code ([Section 5.4](#)) is a two-octet value used in the Application Code AVP. Value 0 is assigned to the base application defined in this document. Additional Application Codes may be assigned by IETF Consensus [\[RFC2434\]](#).

### [10.8](#) Circuit Status Bits



The Circuit Status ([Section 5.4](#)) field is a 16 bit mask, with the two high order bits assigned.

Bit 15 - A (Active) bit

Bit 16 - N (New) bit

Additional bits may be assigned by IETF Consensus [[RFC2434](#)].

#### [10.9](#) Default L2-Specific Sublayer bits

The Default L2 Specific Sublayer defined in [Section 4.6](#) contains 8 bits in the low-order portion of the header, one of which has been assigned and 7 remain.

Bit 1 - S (Sequence) bit

Additional values may be assigned by IETF Consensus [[RFC2434](#)].

#### [10.10](#) L2-Specific Sublayer Type

The L2-Specific Sublayer Type is a 2 octet unsigned integer of which two values have been assigned.

0 - No L2-Specific Sublayer

1 - Default L2-Specific Sublayer present

Additional values may be assigned by Expert Review [[RFC2434](#)].

#### [10.11](#) Data Sequencing Level

The Data Sequencing Level is a 2 octet unsigned integer of which three values have been assigned.

0 - No incoming data packets require sequencing.

1 - Only non-IP data packets require sequencing.

2 - All incoming data packets require sequencing.

Additional values may be assigned by Expert Review [[RFC2434](#)].

### [13.](#) Acknowledgments

Many of the protocol constructs were originally defined in, and the text of this document began with, [RFC 2661](#), "L2TPv2". [RFC 2661](#) authors are W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn and B. Palter.

The basic concept for L2TP and many of its protocol constructs were adopted from L2F [[RFC2341](#)] and PPTP [[RFC2637](#)]. Authors of these drafts are A. Valencia, M. Littlewood, T. Kolar, K. Hamzeh, G. Pall, W. Verthein, J. Taarud, W. Little, and G. Zorn.

Danny Mcpherson and Suhail Nanji published the first "L2TP Service Type" draft which defined the use of L2TP for tunneling of various L2 payload types (initially, Ethernet and Frame Relay).

The team for splitting [RFC 2661](#) into this base document and the companion PPP document consisted of Ignacio Goyret, Jed Lau, Bill Palter, Mark Townsley, and Madhvi Verma. Skip Booth also provided very helpful review and comment.

Some constructs of L2TPv3 were based in part on UTI (Universal Transport Interface), which was originally conceived by Peter Lothberg and Tony Bates.

Stewart Bryant and Simon Barber provided valuable input for the L2TPv3 over IP header.

Juha Heinanen provided helpful review, and input for the Application ID AVP.

Jan Vilhuber, Scott Fluhrer, David McGrew, Scott Wainner, and Skip Booth contributed to the Control Message and Authentication Mechanism as well as general discussions of security.

James Carlson, Thomas Narten and Maria Dos Santos provided very helpful review of the final text.

A number of people provided valuable input and effort for [RFC2661](#), on which this document was based:

John Bray, Greg Burns, Rich Garrett, Don Grosser, Matt Holdrege, Terry Johnson, Dory Leifer, and Rich Shea provided valuable input and

review at the 43rd IETF in Orlando, FL, which led to improvement of the overall readability and clarity of [RFC 2661](#).

Thomas Narten provided a great deal of critical review and formatting. He wrote the first version of the IANA Considerations section.

Dory Leifer made valuable refinements to the protocol definition of L2TP and contributed to the editing of early drafts leading to [RFC 2661](#).

Steve Cobb and Evan Caves redesigned the state machine tables.

Barney Wolff provided a great deal of design input on the endpoint authentication mechanism.

## [11](#). References

### [11.1](#) Normative References

- [RFC1958] Carpenter, B., "Architectural Principles of the Internet", [RFC 1958](#), June 1996.
- [RFC1994] Simpson, W., "PPP Challenge Handshake Authentication Protocol (CHAP)", [RFC 1994](#), August 1996.
- [RFC2072] Berkowitz, H., "Router Renumbering Guide", [RFC 2072](#), January 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and Languages", [BCP 18](#), [RFC 2277](#), January 1998.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [RFC2661] Townsley, W., et al., "Layer Two Tunneling Layer Two Tunneling Protocol (L2TP)", [RFC 2661](#), August 1999.

- [RFC2865] Rigney, C., Rubens, A., Simpson, W., and Willens, S., "Remote Authentication Dial In User Service (RADIUS)", [RFC 2865](#), June 2000.
- [RFC3066] Alvestrand, H., "Tags for the Identification of Languages", [RFC 3066](#), January 2001.
- [RFC3193] Patel, B., Aboba, B., Dixon, W., Zorn, G., and Booth, S., "Securing L2TP using IPsec", [RFC 3193](#), November 2001.

## 11.2 Informative References

- [KPS] Kaufman, C., Perlman, R., and Speciner, M., "Network Security: Private Communications in a Public World", Prentice Hall, March 1995, ISBN 0-13-061466-1.
- [RFC1034] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1321] R. Rivest, "The MD5 Message-Digest Algorithm", [RFC 1321](#), 04/16/1992
- [RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, [RFC 1661](#), July 1994.
- [RFC1700] Reynolds, J., and Postel, J., "Assigned Numbers", STD 2, [RFC 1700](#), October 1994. See also: <http://www.iana.org/numbers.html>.
- [RFC1750] D. Eastlake III, S. Crocker, J. Schiller, "Randomness Recommendations for Security", [RFC 1750](#), December 1994
- [RFC2104] H. Krawczyk, M. Bellare, R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997
- [RFC2138] Rigney, C., Rubens, A., Simpson, W., and Willens, S., "Remote Authentication Dial In User Service (RADIUS)", [RFC 2138](#), April 1997.
- [RFC2341] Valencia, A., Littlewood, M., and Kolar, T., "Cisco Layer Two Forwarding (Protocol) L2F", [RFC 2341](#), May 1998.
- [RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.

INTERNET DRAFT

L2TPv3

October 2003

- [RFC2581] Allman, M., Paxson, V., Stevens, W., "TCP Congestion Control", [RFC 2581](#), April 1999
- [RFC2637] Hamzeh, K., Pall, G., Verthein, W., Taarud, J., Little, W., and Zorn, G., "Point-to-Point Tunneling Protocol (PPTP)", [RFC 2637](#), July 1999.
- [RFC2809] Aboba, B., and Zorn, G., "Implementation of L2TP Compulsory Tunneling via RADIUS", [RFC 2809](#), April 2000.
- [RFC3070] Rawat, V., Tio, R., Nanji, S., and Verma, R., "Layer Two Tunneling Protocol (L2TP) over Frame Relay", [RFC 3070](#), February 2001.
- [RFC3355] Singh, A., Turner, R., Tio, R., Nanji, S., "Layer Two Tunneling Protocol (L2TP) Over ATM Adaptation Layer 5 (AAL5)", [RFC 3355](#), August 2002
- [STEVENS] Stevens, W. Richard, "TCP/IP Illustrated, Volume I: The Protocols", Addison-Wesley Publishing Company, Inc., March 1996, ISBN 0-201-63346-9.

## [12.](#) Editors' Addresses

Jed Lau  
cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
[jedlau@cisco.com](mailto:jedlau@cisco.com)

W. Mark Townsley  
cisco Systems  
[mark@townsley.net](mailto:mark@townsley.net)

Ignacio Goyret  
Lucent Technologies  
[igoyret@lucent.com](mailto:igoyret@lucent.com)

## Appendix A: Control Slow Start and Congestion Avoidance

Although each side has indicated the maximum size of its receive

window, it is recommended that a slow start and congestion avoidance method be used to transmit control packets. The methods described here are based upon the TCP congestion avoidance algorithm as described in [section 21.6](#) of TCP/IP Illustrated, Volume I, by W. Richard Stevens [[STEVENS](#)] (this algorithm is also described in [[RFC2581](#)]).

---

INTERNET DRAFT

L2TPv3

October 2003

Slow start and congestion avoidance make use of several variables. The congestion window (CWND) defines the number of packets a sender may send before waiting for an acknowledgment. The size of CWND expands and contracts as described below. Note, however, that CWND is never allowed to exceed the size of the advertised window obtained from the Receive Window AVP. (In the text below, it is assumed any increase will be limited by the Receive Window Size.) The variable SSTHRESH determines when the sender switches from slow start to congestion avoidance. Slow start is used while CWND is less than SSHTRESH.

A sender starts out in the slow start phase. CWND is initialized to one packet, and SSHTRESH is initialized to the advertised window (obtained from the Receive Window AVP). The sender then transmits one packet and waits for its acknowledgment (either explicit or piggybacked). When the acknowledgment is received, the congestion window is incremented from one to two. During slow start, CWND is increased by one packet each time an ACK (explicit ACK message or piggybacked) is received. Increasing CWND by one on each ACK has the effect of doubling CWND with each round trip, resulting in an exponential increase. When the value of CWND reaches SSHTRESH, the slow start phase ends and the congestion avoidance phase begins.

During congestion avoidance, CWND expands more slowly. Specifically, it increases by  $1/\text{CWND}$  for every new ACK received. That is, CWND is increased by one packet after CWND new ACKs have been received. Window expansion during the congestion avoidance phase is effectively linear, with CWND increasing by one packet each round trip.

When congestion occurs (indicated by the triggering of a retransmission) one-half of the CWND is saved in SSTHRESH, and CWND is set to one. The sender then reenters the slow start phase.

## B.1: Lock-Step Control Connection Establishment

In this example, an LCCE establishes a control connection, with the exchange involving each side alternating in sending messages. This example shows the final acknowledgment explicitly sent within an ACK message. An alternative would be to piggyback the acknowledgment within a message sent as a reply to the ICRQ or OCRQ that will likely follow from the side that initiated the control connection.

INTERNET DRAFT

L2TPv3

October 2003

LCCE A		LCCE B
-----		-----
SCCRQ	->	
Nr: 0, Ns: 0		
		<- SCCRQ
		Nr: 1, Ns: 0
SCCCN	->	
Nr: 1, Ns: 1		
		<- ACK
		Nr: 2, Ns: 1

## B.2: Lost Packet with Retransmission

An existing control connection has a new session requested by LCCE A. The ICRP is lost and must be retransmitted by LCCE B. Note that loss of the ICRP has two effects: It not only keeps the upper level state machine from progressing, but also keeps LCCE A from seeing a timely lower level acknowledgment of its ICRQ.

LCCE A		LCCE B
-----		-----
ICRQ	->	
Nr: 1, Ns: 2		
	(packet lost)	<- ICRP
		Nr: 3, Ns: 1
(pause; LCCE A's timer started first, so fires first)		

ICRQ           ->  
Nr: 1, Ns: 2

(Realizing that it has already seen this packet,  
LCCE B discards the packet and sends an ACK message)

<-           ACK  
Nr: 3, Ns: 2

(LCCE B's retransmit timer fires)

<-           ICRP  
Nr: 3, Ns: 1

ICCN           ->  
Nr: 2, Ns: 3

<-           ACK  
Nr: 4, Ns: 2

## Appendix C: Processing Sequence Numbers

The Default L2-Specific Sublayer, defined in [Section 4.6](#), provides a 24-bit field for sequencing of data packets within an L2TP session. L2TP data packets are never retransmitted, so this sequence is used only to detect packet order, duplicate packets, or lost packets.

The 24-bit Sequence Number field of the Default L2-Specific Sublayer contains a packet sequence number for the associated session. Each sequenced data packet that is sent must contain the sequence number, incremented by one, of the previous sequenced packet sent on a given L2TP session. Upon receipt, any packet with a sequence number equal to or greater than the current expected packet (the last received in-order packet plus one) should be considered "new" and accepted. All other packets are considered "old" or "duplicate" and discarded. Note that the 24-bit sequence number space includes zero as a valid sequence number (as such, it may be implemented with a masked 32-bit counter if desired). All new sessions MUST begin sending sequence numbers at zero.

Larger or smaller sequence number fields are possible with L2TP if an



alternative format to the Default L2-Specific Sublayer defined in this document is used. While 24 bits may be adequate in a number of circumstances, a larger sequence number space will be less susceptible to sequence number wrapping problems for very high session data rates across long dropout periods. The sequence number processing recommendations below should hold for any size sequence number field.

When detecting whether a packet sequence number is "greater" or "less" than a given sequence number value, wrapping of the sequence number must be considered. This is typically accomplished by keeping a window of sequence numbers beyond the current expected sequence number for determination of whether a packet is "new" or not. The window may be sized based on the link speed and sequence number space and SHOULD be configurable with a default equal to one half the size of the available number space (e.g.  $2^{(n-1)}$ , where  $n$  is the number of bits available in the sequence number).

Upon receipt, packets which exactly match the expected sequence number are processed immediately and the next expected sequence number incremented. Packets that fall within the window for new packets may either be processed immediately and the next expected sequence number updated to one plus that received in the new packet, or held for a very short period of time in hopes of receiving the missing packet(s). This 'very short period' should be configurable, with a default corresponding to a time lapse which is at least an order of magnitude less than the retransmission timeout periods of

higher layer protocols such as TCP.

For typical transient packet mis-orderings, dropping out-of-order packets alone should suffice and generally requires far less resources than actively reordering packets within L2TP. An exception is a case where a pair of packet fragments are persistently retransmitted and sent out-of-order. For example, if an IP packet has been fragmented into a very small packet followed by a very large packet before being tunneled by L2TP, it is possible (though admittedly wrong) that the two resulting L2TP packets may be consistently mis-ordered by the PSN in transit between L2TP nodes. If sequence numbers were being enforced at the receiving node without any buffering of out-of-order packets, then the fragmented IP packet may never reach its destination. It may be worth noting here that

this condition is true for any tunneling mechanism of IP packets which include sequence number checking on receipt (i.e. GRE [[RFC2890](#)])).

Utilization of a Data Sequencing Level (see [Section 5.4.3](#)) of 1 (only non-IP data packets require sequencing) allows IP data packets being tunneled by L2TP to not utilize sequence numbers, while utilizing sequence numbers and enforcing packet order for any remaining non-IP data packets. Depending on the requirements of the link-layer being tunneled, and the network data traversing the data-link, this is sufficient in many cases to enforce packet order on frames which require it (such as end-to-end data-link control messages), while not on IP packets which are known to be resilient to packet reordering.

If a large number of packets (e.g. more than one new packet window) are dropped due to an extended outage, or loss of sequence number state on one side of the connection (perhaps as part of a forwarding plane reset or failover to a standby node), it is possible that a large number of packets will be sent in-order, but be wrongly detected by the peer as out-of-order. This can be generally characterized for a window size,  $w$ , sequence number space,  $s$ , and number of packets lost in transit between L2TP endpoints,  $p$ , as follows:

If  $s > p > w$ , then an additional  $(s - p)$  packets that were otherwise received in-order, will be incorrectly classified as out-of-order and dropped. Thus, for a sequence number space,  $s = 128$ , window size,  $w = 64$ , and number of lost packets,  $p = 70$ ;  $128 - 70 = 58$  additional packets would be dropped after the outage until the sequence number wrapped back to the current expected next sequence number.

To mitigate this additional packet loss, one MUST inspect the sequence numbers of packets dropped due to being classified as "old" and reset the expected sequence number accordingly. This may be

accomplished by counting the number of "old" packets dropped that were in sequence among themselves and upon reaching a threshold, resetting the next expected sequence number to that seen in the arriving data packets. Packet timestamps may also be used as an indicator to reset the expected sequence number by detecting a period of time over which "old" packets have been received in-sequence. The ideal thresholds will vary depending on link speed, sequence number

space, and link tolerance to out-of-order packets, and MUST be configurable.

#### Appendix D: Intellectual Property Notice

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

#### Appendix E: Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing

the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

