

Network Working Group
Request for Comments: 4349
Category: Standards Track

C. Pignataro
M. Townsley
Cisco Systems
February 2006

High-Level Data Link Control (HDLC) Frames
over Layer 2 Tunneling Protocol, Version 3 (L2TPv3)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

The Layer 2 Tunneling Protocol, Version 3, (L2TPv3) defines a protocol for tunneling a variety of data link protocols over IP networks. This document describes the specifics of how to tunnel High-Level Data Link Control (HDLC) frames over L2TPv3.

Table of Contents

1. Introduction	2
1.1. Abbreviations	2
1.2. Specification of Requirements	3
2. Control Connection Establishment	3
3. HDLC Link Status Notification and Session Establishment	3
3.1. L2TPv3 Session Establishment	3
3.2. L2TPv3 Session Teardown	5
3.3. L2TPv3 Session Maintenance	5
3.4. Use of Circuit Status AVP for HDLC	6
4. Encapsulation	6
4.1. Data Packet Encapsulation	6
4.2. Data Packet Sequencing	7
4.3. MTU Considerations	7
5. Applicability Statement	8
6. Security Considerations	9
7. IANA Considerations	9
7.1. Pseudowire Type	9
7.2. Result Code AVP Values	9
8. Acknowledgements	9
9. References	10
9.1. Normative References	10
9.2. Informative References	10

1. Introduction

[RFC3931] defines a base protocol for Layer 2 Tunneling over IP networks. This document defines the specifics necessary for tunneling HDLC Frames over L2TPv3. Such emulated circuits are referred to as HDLC Pseudowires (HDLCPWs).

Protocol specifics defined in this document for L2TPv3 HDLCPWs include those necessary for simple point-to-point (e.g., between two L2TPv3 nodes) frame encapsulation, and for simple interface up and interface down notifications.

The reader is expected to be very familiar with the terminology and protocol constructs defined in [RFC3931].

1.1 Abbreviations

HDLC	High-Level Data Link Control
HDLCPW	HDLC Pseudowire
LAC	L2TP Access Concentrator (see [RFC3931])
LCCE	L2TP Control Connection Endpoint (see [RFC3931])
PW	Pseudowire

1.2. Specification of Requirements

In this document, several words are used to signify the requirements of the specification. These words are often capitalized. 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].

2. Control Connection Establishment

In order to tunnel an HDLC link over IP using L2TPv3, an L2TPv3 Control Connection MUST first be established as described in [RFC3931]. The L2TPv3 SCCRP Control Message and corresponding SCCRP Control Message MUST include the HDLC Pseudowire Type of 0x0006 (see Section 7, "IANA Considerations"), in the Pseudowire Capabilities List as defined in 5.4.3 of [RFC3931]. This identifies the control connection as able to establish L2TP sessions to support HDLC Pseudowires (HDLCPWs).

An LCCE MUST be able to uniquely identify itself in the SCCRP and SCCRP messages via a globally unique value. By default, this is advertised via the structured Router ID AVP [RFC3931], though the unstructured Hostname AVP [RFC3931] MAY be used to identify LCCEs as well.

3. HDLC Link Status Notification and Session Establishment

This section specifies how the status of an HDLC interface is reported between two LCCEs, and the associated L2TP session creation and deletion that occurs.

3.1. L2TPv3 Session Establishment

Associating an HDLC serial interface with a PW and its transition to "Ready" or "Up" results in the establishment of an L2TP session via the standard three-way handshake described in Section 3.4.1 of [RFC3931]. For purposes of this discussion, the action of locally associating an interface running HDLC with a PW by local configuration or otherwise is referred to as "provisioning" the HDLC interface. The transition of the interface to "ready" or "up" will be referred to as the interface becoming ACTIVE. The transition of the interface to "not-ready" or "down" will be referred to as the interface becoming INACTIVE.

An LCCE MAY initiate the session immediately upon association with an HDLC interface or wait until the interface becomes ACTIVE before attempting to establish an L2TP session. Waiting until the interface transitions to ACTIVE may be preferred, as it delays allocation of resources until absolutely necessary.

The Pseudowire Type AVP defined in Section 5.4.4 of [RFC3931], Attribute Type 68, MUST be present in the ICRQ messages and MUST include the Pseudowire Type of 0x0006 for HDLCPWs.

The Circuit Status AVP (see Section 3.4) MUST be present in the ICRQ and ICRP messages and MAY be present in the SLI message for HDLCPWs.

Following is an example of the L2TP messages exchanged for an HDLCPW that is initiated after an HDLC interface is provisioned and becomes ACTIVE.

LCCE (LAC) A	LCCE (LAC) B
-----	-----
HDLC Interface Provisioned	
	HDLC Interface Provisioned
HDLC Interface ACTIVE	
	HDLC Interface ACTIVE
	ICRQ (status = 0x03) ---->
	<---- ICRP (status = 0x03)
L2TP session established, OK to send data into tunnel	
	L2TP session established, OK to send data into tunnel
	ICCN ---->

In the example above, an ICRQ is sent after the interface is provisioned and becomes ACTIVE. The Circuit Status AVP indicates that this link is ACTIVE and New (0x03). The Remote End ID AVP [RFC3931] MUST be present in the ICRQ in order to identify the HDLC link (together with the identity of the LCCE itself as defined in Section 2) with which to associate the L2TP session. The Remote End ID AVP defined in [RFC3931] is of opaque form and variable length, though one MUST at a minimum support use of an unstructured four-octet value that is known to both LCCEs (either by direct configuration, or some other means). The exact method of how this value is configured, retrieved, discovered, or otherwise determined at each LCCE is outside the scope of this document.

As with the ICRQ, the ICRP is sent only after the associated HDLC interface transitions to ACTIVE as well. If LCCE B had not been provisioned for the interface identified in the ICRQ, a CDN would have been immediately returned indicating that the associated link was not provisioned or available at this LCCE. LCCE A SHOULD then exhibit a periodic retry mechanism. If so, the period and maximum number of retries MUST be configurable.

An Implementation MAY send an ICRQ or ICRP before an HDLC interface is ACTIVE, as long as the Circuit Status AVP reflects that the link is INACTIVE and an SLI is sent when the HDLC interface becomes ACTIVE (see Section 3.3).

The ICCN is the final stage in the session establishment, confirming the receipt of the ICRP with acceptable parameters to allow bidirectional traffic.

3.2. L2TPv3 Session Teardown

In the event a link is removed (unprovisioned) at either LCCE, the associated L2TP session MUST be torn down via the CDN message defined in Section 3.4.3 of [RFC3931].

General Result Codes regarding L2TP session establishment are defined in [RFC3931]. Additional HDLC result codes are defined as follows:

- 20 - HDLC Link was deleted permanently (no longer provisioned)
- 21 - HDLC Link has been INACTIVE for an extended period of time

3.3. L2TPv3 Session Maintenance

HDLC PWs over L2TP make use of the Set Link Info (SLI) control message defined in [RFC3931] to signal HDLC link status notifications between PEs. The SLI message is a single message that is sent over the L2TP control channel, signaling the interface state change.

The SLI message MUST be sent any time there is a status change of any values identified in the Circuit Status AVP. The only exceptions to this are the initial ICRQ, ICRP, and CDN messages, which establish and teardown the L2TP session itself. The SLI message may be sent from either PE at any time after the first ICRQ is sent (and perhaps before an ICRP is received, requiring the peer to perform a reverse Session ID lookup).

All sessions established by a given control connection utilize the L2TP Hello facility defined in Section 4.4 of [RFC3931] for session keepalive. This gives all sessions basic dead peer and path detection between PEs.

3.4. Use of Circuit Status AVP for HDLC

HDLC reports Circuit Status with the Circuit Status AVP defined in [RFC3931], Attribute Type 71. For reference, this AVP is shown below:

```

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

```

The Value is a 16-bit mask with the two least significant bits defined and the remaining bits reserved for future use. Reserved bits MUST be set to 0 when sending, and ignored upon receipt.

The N (New) bit SHOULD be set to one (1) if the Circuit Status indication is for a new HDLC circuit; to zero (0) otherwise.

The A (Active) bit indicates whether the HDLC interface is ACTIVE (1) or INACTIVE (0).

4. Encapsulation

4.1. Data Packet Encapsulation

HDLCPWs use the default encapsulations defined in [RFC3931] for demultiplexing, sequencing, and flags. The HDLCPW Type over L2TP is intended to operate in an "interface to interface" or "port to port" fashion, passing all HDLC data and control PDUs over the PW. The HDLC PDU is stripped of flags and trailing FCS, bit/byte unstuffing is performed, and the remaining data, including the address, control, and protocol fields, is transported over the PW.

Since all packets are passed in a largely transparent manner over the HDLCPW, any protocol that has HDLC-like framing may utilize the HDLCPW mode, including PPP, Frame-Relay ("port to port" Frame-Relay transport), X.25 (LAPB), etc. In such cases, the negotiations and signaling of the specific protocols transported over the HDLCPW take place between the Remote Systems. A non-exhaustive list of examples and considerations of this transparent nature include:

- o When the HDLCPW transports Point-to-Point Protocol (PPP) traffic, PPP negotiations (Link Control Protocol, optional authentication, and Network Control Protocols) are performed between Remote Systems, and LCCEs do not participate in these negotiations.

- o When the HDLCPW transports Frame-Relay traffic, PVC status management procedures (Local Management Interface) take place between Remote Systems, and LCCEs do not participate in LMI. Additionally, individual Frame-Relay virtual-circuits are not visible to the LCCEs, and the FECN, BECN, and DE bits are transported transparently.
- o When the HDLCPW transports X.25 (LAPB) traffic, LCCEs do not function as either LAPB DCE or DTE devices.

On the other hand, exceptions include cases where direct access to the HDLC interface is required, or modes that operate on the flags, FCS, or bit/byte unstuffing that is performed before sending the HDLC PDU over the PW. An example of this is PPP ACCM negotiation.

4.2. Data Packet Sequencing

Data Packet Sequencing MAY be enabled for HDLCPWs. The sequencing mechanisms described in Section 4.6.1 of [RFC3931] MUST be used for signaling sequencing support. HDLCPWs over L2TP MUST request the presence of the L2TPv3 Default L2-Specific Sublayer defined in Section 4.6 of [RFC3931] when sequencing is enabled, and MAY request its presence at all times.

4.3. MTU Considerations

With L2TPv3 as the tunneling protocol, the packet resulting from the encapsulation is N bytes longer than the HDLC frame without the flags or FCS. The value of N depends on the following fields:

L2TP Session Header:

Flags, Ver, Res	4 octets (L2TPv3 over UDP only)
Session ID	4 octets
Cookie Size	0, 4, or 8 octets
L2-Specific Sublayer	0 or 4 octets (i.e., using sequencing)

Hence the range for N in octets is:

N = 4-16, L2TPv3 data messages are over IP;
N = 16-28, L2TPv3 data messages are over UDP;
(N does not include the IP header.)

The MTU and fragmentation implications resulting from this are discussed in Section 4.1.4 of [RFC3931].

5. Applicability Statement

HDLC Pseudowires support a "port to port" or "interface to interface" deployment model operating in a point-to-point fashion. In addition to the transport of HDLC frames, a natural application of HDLCPWs allows for the transport of any protocol using an HDLC-like framing.

The HDLCPW emulation over a packet-switched network (PSN) has the following characteristics in relationship to the native service:

- o HDLC data and control fields are transported transparently (see Section 4.1). The specific negotiations and signaling of the protocol being transported are performed between Remote Systems transparently, and the LCCE does not participate in them.
- o The trailing FCS (Frame Check Sequence) containing a CRC (Cyclic Redundancy Check) is stripped at the ingress LCCE and not transported over HDLCPWs. It is therefore regenerated at the egress LCCE (see Section 4.1). This means that the FCS may not accurately reflect errors on the end-to-end HDLC link. Errors or corruption introduced in the HDLCPW payload during encapsulation or transit across the packet-switched network may not be detected. This lack of integrity-check transparency may not be of concern if it is known that the inner payloads or upper protocols transported perform their own error and integrity checking. To allow for payload integrity-checking transparency on HDLCPWs using L2TP over IP or L2TP over UDP/IP, the L2TPv3 session can utilize IPSec as specified in Section 4.1.3 of [RFC3931].
- o HDLC link status notification is provided using the Circuit Status AVP in the SLI message (see Section 3.4).
- o The length of the resulting L2TPv3 packet is longer than the encapsulated HDLC frame without flags and FCS (see Section 4.3), with resulting MTU and fragmentation implications discussed in Section 4.1.4 of [RFC3931].
- o The packet-switched network may reorder, duplicate, or silently drop packets. Sequencing may be enabled in the HDLCPW for some or all packets to detect lost, duplicate, or out-of-order packets on a per-session basis (see Section 4.2).
- o The faithfulness of an HDLCPW may be increased by leveraging Quality of Service features of the LCCEs and the underlying PSN.

6. Security Considerations

HDLC over L2TPv3 is subject to the security considerations defined in [RFC3931]. Beyond the considerations when carrying other data link types, there are no additional considerations specific to carrying HDLC.

7. IANA Considerations

7.1. Pseudowire Type

The signaling mechanisms defined in this document rely upon the allocation of an HDLC Pseudowire Type (see Pseudowire Capabilities List as defined in 5.4.3 of [RFC3931] and L2TPv3 Pseudowire Types in 10.6 of [RFC3931]) by the IANA (number space created as part of publication of [RFC3931]). The HDLC Pseudowire Type is defined in Section 2 of this specification:

L2TPv3 Pseudowire Types

0x0006 - HDLC Pseudowire Type

7.2. Result Code AVP Values

This number space is managed by IANA as described in section 2.3 of [BCP0068]. Two new L2TP Result Codes for the CDN message appear in Section 3.2. The following is a summary:

Result Code AVP (Attribute Type 1) Values

20 - HDLC Link was deleted permanently (no longer provisioned)

21 - HDLC Link has been INACTIVE for an extended period of time

8. Acknowledgements

Thanks to Sudhir Rustogi and George Wilkie for valuable input. Maria Alice Dos Santos provided helpful review and comment. Many thanks to Mark Lewis for providing review and clarifying comments during IETF Last Call.

9. References

9.1. Normative References

- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

- [BCP0068] Townsley, W., "Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers Authority (IANA) Considerations Update", BCP 68, RFC 3438, December 2002.

Authors' Addresses

Carlos Pignataro
Cisco Systems
7025 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709

EMail: cpignata@cisco.com

W. Mark Townsley
Cisco Systems
7025 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709

EMail: mark@townsley.net

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 on-line IPR repository at <http://www.ietf.org/ipr>.

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 implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

