

Network Working Group
Request for Comments: 4591
Category: Standards Track

M. Townsley
G. Wilkie
S. Booth
S. Bryant
Cisco Systems
J. Lau
July 2006

Frame Relay 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 Frame Relay over L2TPv3, including frame encapsulation, virtual-circuit creation and deletion, and status change notification.

Table of Contents

1. Introduction	2
1.1. Abbreviations	3
1.2. Specification of Requirements	3
2. Control Connection Establishment	3
3. PVC Status Notification and Session Establishment	3
3.1. L2TPv3 Session Establishment	4
3.2. L2TPv3 Session Teardown	5
3.3. L2TPv3 Session Maintenance	5
3.4. Use of the Circuit Status AVP for Frame Relay	6
3.5. Frame Relay Header Length AVP	7
4. Encapsulation	7
4.1. Data Packet Encapsulation	7
4.2. Data Packet Sequencing	9
4.3. MTU Considerations	9
5. Applicability Statement	10
6. Security Considerations	10
7. IANA Considerations	11
7.1. Pseudowire Type	11
7.2. Result Code AVP Values	11
7.3. Control Message Attribute Value Pairs (AVPs)	11
8. Acknowledgements	11
9. References	12
9.1. Normative References	12
9.2. Informative References	12

1. Introduction

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

Protocol specifics defined in this document for L2TPv3 FRPWs operating in a "virtual circuit-to-virtual circuit" mode include those necessary for frame encapsulation, PVC creation and deletion, and status change notification. Frame Relay traffic may also be transported in a "port-to-port" or "interface-to-interface" fashion using High-Level Data Link Control (HDLC) Pseudowires as defined in [RFC4349]. Support for Switched Virtual Circuits (SVCs) and Switched/Soft Permanent Virtual Circuits (SPVCs) are outside the scope of this document.

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

1.1. Abbreviations

FR Frame Relay
FRPW Frame Relay Pseudowire
LCCE L2TP Control Connection Endpoint (See [RFC3931])
PVC Permanent virtual circuit
PW Pseudowire
VC Virtual circuit

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 a Frame Relay circuit 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 Frame Relay Data Link Connection Identifier (DLCI) PW Type of 0x0001 (see IANA Considerations), in the Pseudowire Capabilities List, as defined in Section 5.4.3 of [RFC3931]. This identifies the control connection as able to establish L2TP sessions to support Frame Relay Pseudowires (FRPWs).

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

3. PVC Status Notification and Session Establishment

This section specifies how the status of a PVC is reported between two LCCEs. This includes what should happen when a PVC is created, deleted or when it changes state between ACTIVE and INACTIVE. When emulating a Frame Relay service, if the procedures for PVC status management defined in [Q933] Annex A are being used between an LCCE and the attached Remote System, an LCCE MUST participate in them (see Section 3.3).

3.1. L2TPv3 Session Establishment

PVC creation (provisioning) results in establishment of an L2TP session via the standard three-way handshake described in Section 3.4.1 of [RFC3931]. An LCCE MAY initiate the session immediately upon PVC creation or wait until the PVC state transitions to ACTIVE before attempting to establish a session for the PVC. Waiting until the PVC transitions to ACTIVE may be preferred, as it delays allocation of L2TP resources until it is absolutely necessary.

The Pseudowire Type AVP defined in Section 5.4.4 of [RFC3931], Attribute Type 68, MUST be present in the Incoming-Call-Request (ICRQ) messages and MUST include the Frame Relay DLCI PW Type of 0x0001 for FRPWs.

The Circuit Status AVP (see Section 3.4) MUST be present in the ICRQ and Incoming-Call-Reply (ICRP) messages and MAY be present in the Set Link Info (SLI) message for FRPWs.

The Frame Relay Header Length AVP (see Section 3.5) MAY be present in the ICRQ and ICRP messages.

The following is an example of the L2TP messages exchanged for an FRPW that is initiated after a new PVC is provisioned and becomes ACTIVE.

LCCE (LAC) A ----- FR PVC Provisioned FR PVC ACTIVE	LCCE (LAC) B ----- FR PVC Provisioned ICRQ (status = 0x03) ----> FR PVC ACTIVE <---- ICRP (status = 0x03)
--	--

L2TP session established,
OK to send data into tunnel

ICCN----->

L2TP session established,
OK to send data into tunnel

In the example above, an ICRQ is sent after the PVC is created and becomes ACTIVE. The Circuit Status AVP indicates that this PVC is ACTIVE and New (0x03). The Remote End ID AVP [RFC3931] MUST be

present in the ICRQ in order to identify the PVC (together with the identity of the LCCE itself, as defined in Section 2) to associate the L2TP session with. 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 FR PVC transitions to ACTIVE as well. If LCCE B had not been provisioned for the PVC identified in the ICRQ, a Call-Disconnect-Notify (CDN) would have been immediately returned indicating that the circuit 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 a PVC is ACTIVE, as long as the Circuit Status AVP reflects that the PVC is INACTIVE and an SLI is sent when the PVC becomes ACTIVE (see Section 3.3).

The Incoming-Call-Connected (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 that a PVC is deleted (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 Frame Relay result codes are defined as follows:

17: FR PVC was deleted permanently (no longer provisioned) 18: FR PVC has been INACTIVE for an extended period of time 19: Mismatched FR Header Length

3.3. L2TPv3 Session Maintenance

FRPW over L2TP makes use of the SLI control message defined in [RFC3931] to signal Frame Relay link status notifications between LCCEs. This includes ACTIVE or INACTIVE notifications of the VC, and any other parameters that may need to be shared between the tunnel endpoints or LCCEs in order to provide proper PW emulation. The SLI message is a single message that is sent over the L2TP control

channel signalling the state change. Since the message is delivered reliably, there is no additional response or action required of the PW subsystem to ensure that the state change notification was received by the tunnel peer.

The SLI message MUST be sent any time there is a circuit status change that may be reported by any values identified in the Circuit Status AVP. The only exceptions to this are the initial ICRQ, ICRP, and CDN messages, which establish and tear down the L2TP session itself when the PVC is created or deleted. The SLI message may be sent from either LCCE at any time after the first ICRQ is sent (and perhaps before an ICRP is received, requiring that the peer to perform a reverse Session ID lookup).

An LCCE participating in the procedures for PVC status management defined in [Q933], Annex A, MUST transmit an SLI message including the Circuit Status AVP (see Section 3.4) when it detects a change in the status for a particular local FR PVC (i.e., when it detects a service-affecting condition or the clearing of such a condition). An LCCE receiving an SLI message indicating a change in the status of a particular FRPW SHOULD generate corresponding updates for the FR PVC towards the Remote System, as defined in [Q933], Annex A.

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 LCCEs.

3.4. Use of the Circuit Status AVP for Frame Relay

Frame Relay circuit status is reported via 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 by the sender and ignored by the receiver.

The N (New) bit indicates whether the Circuit Status indication is for a new FR PVC (1) or an existing FR PVC (0).

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

3.5. Frame Relay Header Length AVP

The "Frame Relay Header Length AVP", Attribute type 85, indicates the number of bytes in the Frame Relay header. The two peer LCCEs MUST agree on the length of the Frame Relay header.

This AVP is exchanged during session negotiation (in ICRQ, ICRP). If the other LCCE supports a different Frame Relay header length, the associated L2TP session MUST be torn down via CDN message with result code 19 (see Section 3.2).

If the Frame Relay Header Length AVP is not signalled, it MUST be assumed that the peer uses a 2-byte Frame Relay header.

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

Frame Relay Header Length (ICRQ, ICRP)

```

      0                               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|   Frame Relay Header Length   |
+---+---+---+---+---+---+---+---+

```

The Frame Relay Header Length Type is a 2-octet unsigned integer with the following values defined in this document:

2: Two-octet Frame Relay Header 4: Four-octet Frame Relay Header

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP MAY be set to 0 but MAY vary (see Section 5.2 of [RFC3931]). The length (before hiding) of this AVP is 8.

4. Encapsulation

4.1. Data Packet Encapsulation

The FR PDU is transported in its entirety, excluding the opening and closing High Level Data Link Control (HDLC) flags and the frame check sequence (FCS). Bit stuffing is undone. The L2TPv3 Session Header is that as defined in [RFC3931]. If sequencing or other features require presence of an L2-Specific Sublayer, the Default format defined in Section 4.6 of [RFC3931] MUST be used.

The FR header is defined in [Q922]; however, the notation used differs from that used in IETF specifications. For reference, the FR header (referred to as Address Field in [Q922]) in IETF notation is

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
| hi dlci   |C|0|lo dlci|F|B|D|1|
+---+---+---+---+---+---+---+---+

```

Two-octet FR Header

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| hi dlci   |C|0| dlci   |F|B|D|0|   dlci   |0| dlci_lo   |0|1|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Four-octet FR Header

C/R (bit 6) FR frame C/R (command/response) bit [Q922].

F - FECN (bit 12): FR FECN (Forward Explicit Congestion Notification) bit [Q922].

B - BECN (bit 13):

FR BECN (Backward Explicit Congestion Notification) bit [Q922].

D - DE (bit 14) FR DE bit indicates the discard eligibility [Q922].

Usage of the C/R, FECN, BECN, and DE bits is as specified in [Q922].

The C/R bit is conveyed transparently. Its value MUST NOT be changed by the LCCE.

The FECN bit MAY be set by the LCCE to notify the receiving end-user that the frames it receives have encountered congestion. The end-user may use this indication for destination-controlled transmit rate adjustment. The bit must never be cleared by the LCCE. If the LCCE does not support FECN, it shall pass the bit unchanged.

The BECN bit MAY be set by the LCCE to notify the receiving end-user that the frames it transmits may encounter congestion. The end-user may use this indication to adjust its transmit rate. The bit must never be cleared by the LCCE. If the LCCE does not support BECN, it shall pass the bit unchanged.

The DE bit MAY be set by a policing function on the LCCE to indicate that this frame SHOULD be discarded in preference to other frames in a congestion situation. The bit must never be cleared by the LCCE. If the LCCE does not support DE, it shall pass the bit unchanged.

The encapsulation of Frame Relay frames with the two-octet FR Header is REQUIRED. The encapsulation of Frame Relay frames with the four-octet FR Header is OPTIONAL. The encapsulation of Frame Relay frames with the three-octet FR Header is outside the scope of this document.

4.2. Data Packet Sequencing

Data Packet Sequencing MAY be enabled for FRPWs. The sequencing mechanisms described in [RFC3931] MUST be used for signalling sequencing support. FRPW over L2TP MUST request the presence of the L2TPv3 Default L2-Specific Sublayer when sequencing is enabled and MAY request its presence at all times.

If the FRPW is known to be carrying data that does not require packet order be strictly maintained (such as IP), then packet sequencing for the FRPW SHOULD NOT be enabled.

4.3. MTU Considerations

With L2TPv3 as the tunneling protocol, the packet resulted from the encapsulation is N bytes longer than Frame Relay frame without the opening and closing HDLC 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., with sequencing)

Thus, 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

The Frame Relay PW emulation described in this document allows a service provider to offer a Frame Relay PVC-based service across an IP packet-switched network (PSN). A Frame Relay port-based service can be offered using [RFC4349].

The FRPW emulation has the following characteristics in relationship to the native service:

- o There is a one-to-one mapping between a Frame Relay PVC and an FRPW, supporting bi-directional transport of variable length frames. The Frame Relay frame is transported in its entirety, including the DLCI and the C/R, FECN, BECN, and DE bits, but excluding the opening and closing flags and the FCS. The egress LCCE re-writes the DLCI and regenerates the FCS.
- o Two- and four-octet address fields are supported. The length is negotiated between LCCEs during session establishment (see Section 3.5).
- o The availability or unavailability of the PVC is signalled between LCCEs using the Circuit Status AVP (see Section 3.4). Loss of connectivity between LCCEs can be detected by the L2TPv3 keepalive mechanism (see Section 4.4 in [RFC3931]). These indications can be used to determine the PVC status to be signalled through [Q933] procedures at the Frame Relay interface.
- o The maximum frame size that can be supported is limited by the PSN MTU, unless fragmentation and reassembly is used (see Section 4.1.4 of [RFC3931]).
- o Sequencing may be enabled on the FRPW to ensure that frames are delivered in order (see Section 4.2).
- o Quality of Service characteristics, such as throughput (CIR), committed burst size (bc), excess burst size (be), and priority, can be provided by leveraging Quality of Service features of the LCCEs and the underlying PSN.

6. Security Considerations

Frame Relay over L2TPv3 is subject to the security considerations defined in [RFC3931]. There are no additional considerations specific to carrying Frame Relay that are not present for carrying other data link types.

7. IANA Considerations

7.1. Pseudowire Type

The following value for the Frame Relay DLCI PW Type (see Pseudowire Capabilities List, as defined in 5.4.3 of [RFC3931], and L2TPv3 Pseudowire Types in 10.6 of [RFC3931]) is allocated by the IANA (number space already created as part of publication of [RFC3931]):

L2TPv3 Pseudowire Types

0x0001: Frame Relay DLCI Pseudowire Type

7.2. Result Code AVP Values

This number space is managed by IANA as described in Section 2.3 of [RFC3438]. Three 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

17: PVC was deleted permanently (no longer provisioned)

18: PVC has been INACTIVE for an extended period of time

19: Mismatched FR Header Length

7.3. Control Message Attribute Value Pairs (AVPs)

This number space is managed by IANA as described in Section 2.2 of [RFC3438]. An additional AVP Attribute, specified in Section 3.5, was allocated for this specification:

Control Message Attribute Value Pairs

85: Frame Relay Header Length

8. Acknowledgements

The first Frame Relay over L2TP document, "Frame Relay Service Type for L2TP", was published in February of 2001, by Nishit Vasavada, Jim Boyle, Chris Garner, Serge Maskalik, and Vijay Gill. This document is substantially different, but the basic concept of carrying Frame Relay over L2TP is the same.

Thanks to Lloyd Wood for a razor-sharp review.

Carlos Pignataro helped with review and editing of the document.

During IETF Last Call, Mark Lewis provided thorough review and comments.

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.
- [RFC4349] Pignataro, C. and M. Townsley, "High-Level Data Link Control (HDLC) Frames over Layer 2 Tunneling Protocol, Version 3 (L2TPv3)", RFC 4349, February 2006.

9.2. Informative References

- [RFC3438] Townsley, W., "Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers Authority (IANA) Considerations Update", BCP 68, RFC 3438, December 2002.
- [Q922] ITU-T Recommendation Q.922, "ISDN Data Link Layer Specification for Frame Mode Bearer Services", ITU, Geneva, 1992.
- [Q933] ITU-T Recommendation Q.933, "Signalling specifications for frame mode switched and permanent virtual connection control and status monitoring", ITU, Geneva, 2003.

Authors' Addresses

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

EMail: mark@townsley.net

George Wilkie
Cisco Systems
96 Commercial Street
Edinburgh, EH6 6LX
United Kingdom

EMail: gwilkie@cisco.com

Skip Booth
Cisco Systems
7025 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709

EMail: ebooth@cisco.com

Stewart Bryant
Cisco Systems
250 Longwater Ave
Green Park
Reading RG2 6GB
United Kingdom

EMail: stbryant@cisco.com

Jed Lau

EMail: jedlau@gmail.com

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).

