

Network Working Group
Request for Comments: 5641
Updates: 3931, 4349, 4454, 4591, 4719
Category: Standards Track

N. McGill
C. Pignataro
Cisco Systems
August 2009

Layer 2 Tunneling Protocol Version 3 (L2TPv3)
Extended Circuit Status Values

Abstract

This document defines additional Layer 2 Tunneling Protocol Version 3 (L2TPv3) bit values to be used within the "Circuit Status" Attribute Value Pair (AVP) to communicate finer-grained error states for Attachment Circuits (ACs) and pseudowires (PWs). It also generalizes the Active bit and deprecates the use of the New bit in the Circuit Status AVP, updating RFC 3931, RFC 4349, RFC 4454, RFC 4591, and RFC 4719.

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) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	2
1.1.	Specification of Requirements	2
1.2.	Abbreviations	3
2.	L2TPv3 Extended Circuit Status Values	3
3.	Circuit Status Usage and Clarifications	7
4.	Updates to Existing RFCs	8
5.	Security Considerations	9
6.	IANA Considerations	9
7.	Acknowledgements	9
8.	References	9
8.1.	Normative References	9
8.2.	Informative References	10

1. Introduction

Currently, the L2TPv3 Circuit Status AVP [RFC3931] is able to convey the UP/DOWN status of an access circuit. However, a finer granularity is often useful to determine the direction of the fault, as has been added for MPLS-based pseudowires and is used in the pseudowire control protocol using the Label Distribution Protocol (LDP); see Section 3.5 of [RFC4446] and Section 5.4.2 of [RFC4447].

Additionally, it is useful (in session-level redundancy scenarios) to be able to indicate if a pseudowire is in a standby state, where it is fully established by signaling and allows Operations, Administration, and Maintenance, but is not switching data. Again, such functionality is available for MPLS-based pseudowires using LDP, see [PREF-FWD].

This document provides extended circuit status bit values for L2TPv3 and adds them in a manner such that it is backwards compatible with the current Circuit Status AVP. These new bits are applicable to all pseudowire types that use the Circuit Status AVP.

1.1. 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].

1.2. Abbreviations


The following abbreviations are used in this document and in the documents that it updates. L2TPv3 Control Message Types are listed in Section 6 of [RFC3931].

- AC Attachment Circuit
- AVP Attribute Value Pair
- LCCE L2TP Control Connection Endpoint
- NNI Network-Network Interface
- PE Provider Edge
- PSN Packet Switched Network
- PW Pseudowire

2. L2TPv3 Extended Circuit Status Values

The Circuit Status AVP (ICRQ, ICRP, ICCN, OCRQ, OCRP, OCCN, SLI), Attribute Type 71, 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, currently defined in [RFC3931], has the following format:



Bit	Bit-Value	Name
(A)	15 0x0001	Active
(N)	14 0x0002	New

As currently defined in [RFC3931] and replicated in [RFC4349], [RFC4454], [RFC4591], and [RFC4719], the two bits have the following meanings:

- o The A (Active) bit indicates whether the circuit is up/active/ready (1) or down/inactive/not-ready (0).
- o The N (New) bit indicates whether the circuit status indication is for a new circuit (1) or an existing circuit (0).


This document updates the semantics of the A and N bits as follows (see also Section 4):

The A (Active) bit indicates whether the local pseudowire endpoint (including the local Attachment Circuit (AC) and local Packet Switched Network (PSN)-facing pseudowire termination) has no faults present and is up/active/ready (1) or has faults present and is down/inactive/not-ready (0).

The N (New) bit indicates if the notification is for a new circuit (1) or an existing circuit (0), and is provided to emulate Network-Network Interface (NNI) signaling between Provider Edge (PE) routers, e.g., Frame Relay NNI. It MAY be used to convey that a circuit has been re-provisioned or newly provisioned at the PE, which can already be inferred from the L2TP control message type. It is therefore uncertain as to what use the receiving PE can make of this bit, although it MAY include logging. This document deprecates this bit as it is of little or no use, hence this bit SHOULD be ignored on receipt and is OPTIONAL to set on sending. For reference, see Section 3.4 of [RFC4591], which does not specify any additional usage beyond the setting of the N bit in the ICRQ, ICRP (and OCRQ, OCRP) and the clearing of it in all other control messages.

This document also extends this bitmap of values to allow for finer granularity of local pseudowire (i.e., Attachment Circuit or PSN-facing endpoint) status reporting.


The Attribute Value field for the Circuit Status AVP, including the new values, has the following format:



Bit	Bit-Value	Name
(A)	15 0x0001	Active: Pseudowire has no faults
(N)	14 0x0002	New [use deprecated]
(R)	13 0x0004	Local Attachment Circuit (ingress) Receive Fault
(T)	12 0x0008	Local Attachment Circuit (egress) Transmit Fault
(I)	11 0x0010	Local PSN-facing PW (ingress) Receive Fault
(E)	10 0x0020	Local PSN-facing PW (egress) Transmit Fault
(S)	9 0x0040	Pseudowire is in Standby mode

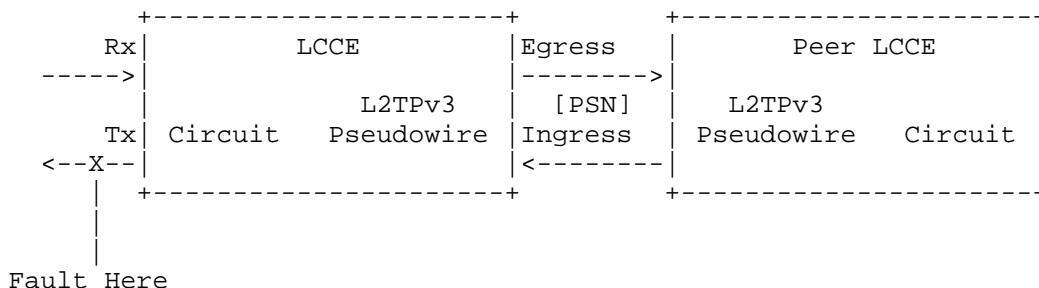
The new bit values have the following meanings:

(R), Local Attachment Circuit (ingress) Receive Fault



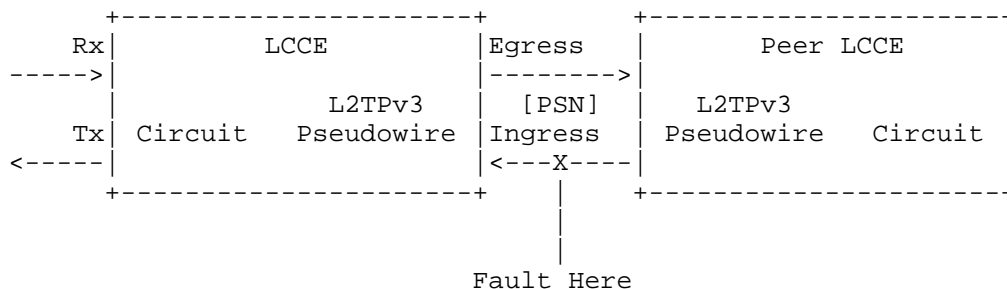
An alarm or fault has occurred at the local Attachment Circuit such that it is unable to receive traffic. It can still transmit traffic.

(T), Local Attachment Circuit (egress) Transmit Fault



A fault has occurred at the local Attachment Circuit such that it is unable to transmit traffic. It can still receive traffic.

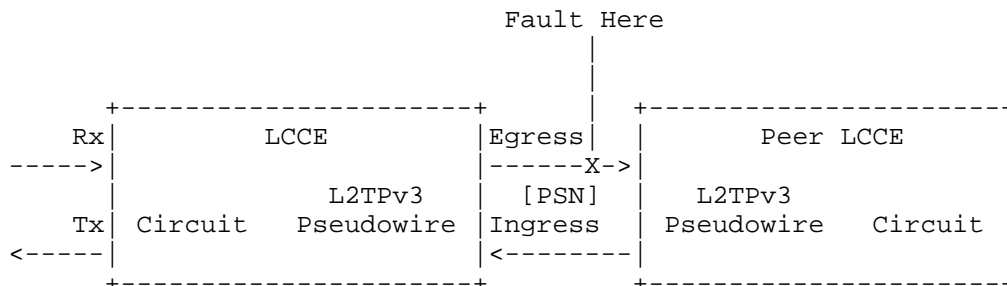
(I), Local PSN-facing PW (ingress) Receive Fault



A fault has occurred in the receive direction between the local endpoint and the remote L2TP endpoint.

Note that a fault at the session level would not necessarily trigger an L2TP control connection timeout. The means of detecting this fault are outside the scope of this document; as an example, detection may be via PW Type-specific means, Bidirectional Forwarding Detection (BFD), or other methods.

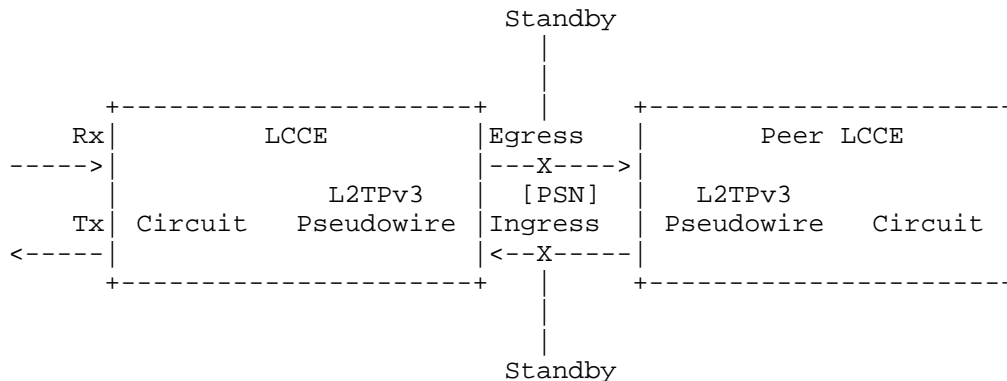
(E), Local PSN-facing PW (egress) Transmit Fault



A fault has occurred in the transmit direction between the local endpoint and the remote L2TP endpoint.

Note that a fault at the session level would not necessarily trigger an L2TP control connection timeout. The means of detecting this fault are outside the scope of this document; as an example, detection may be via PW Type-specific means, BFD, or other methods.

(S), Pseudowire is in Standby mode



The pseudowire has been placed into a Standby mode, which means that although it was signaled (during setup of the PW) and is operational, it is NOT switching user traffic. Any received user traffic SHOULD be dropped. User traffic MUST NOT be transmitted.

A standby pseudowire also allows for means to check its data plane liveness in order to ensure its ability to switch data packets end-to-end. This is achieved, for example, as detailed in [RFC5085] or [VCCV-BFD]. However, data is not forwarded from an Attachment Circuit (AC) into the L2TPv3 session, or from the L2TPv3 session out to the AC.

3. Circuit Status Usage and Clarifications

In implementations prior to this specification, bits 0-13 MUST be set to zero (see Section 5.4.5 of [RFC3931]). This allows for legacy implementations to interwork properly with new implementations.

The following are clarifications regarding the usage of the Circuit Status AVP bits as defined in this specification:

- o The (R), (T), (I), and (E) bits are collectively referred to as "fault status bits".
- o [RFC3931] defined the (A) bit as pertaining to local access circuit state only. This document redefines it as meaning that "no faults are present on the local pseudowire endpoint."
- o If multiple faults occur, all the fault status bits corresponding to each fault MUST be set (i.e., they MUST be bitwise ORed together).
- o The (A) bit MUST NOT be set until all fault status bits are cleared. This behavior allows an endpoint to be backwards compatible with a remote endpoint that does not understand these new status bits.
- o If any of the fault status bits are set, then the (A) bit MUST be cleared. That is, the fault status bits (R, T, I, E) are a more granular definition of (A), such that ORing the bits provides an inverted (A).
- o If (A) is clear and the fault status bits (R, T, I, E) are clear, it means that there is no extended circuit status. That is, the circuit is down/inactive/not-ready (from the (A) bit), without a more granular (extended) indication.
- o The (S) bit can be set in conjunction with any other bit, including (A). A pseudowire endpoint in Standby (S bit set) can be up/active/ready (A bit set) or experiencing a fault (A bit cleared and one or more of the fault status bits (R, T, I, E) set).
- o Leaving Standby mode is indicated by the clearing of the (S) bit.

- o The usage of the (N) bit has been deprecated.

4. Updates to Existing RFCs

This document updates existing RFCs that define (either generically or in the context of a specific set of PW Types) the Active and New bits of the Circuit Status AVP. The Active and New bits of the Circuit Status AVP are specified in Section 5.4.5 of [RFC3931]. Those definitions are adapted to specific Attachment Circuits and replicated in Section 3.4 of [RFC4349] (High-Level Data Link Control Frames over L2TPv3), Section 8 of [RFC4454] (Asynchronous Transfer Mode over L2TPv3), Section 3.4 of [RFC4591] (Frame Relay over L2TPv3), and Section 2.3.3 of [RFC4719] (Ethernet Frames over L2TPv3). This document updates the definitions in all five of these references to say:

The A (Active) bit indicates whether the local pseudowire endpoint (including the local Attachment Circuit and local PSN-facing pseudowire termination) has no faults present and is up/active/ready (1) or has faults present and is down/inactive/not-ready (0).

The N (New) bit usage is deprecated; it SHOULD be ignored on receipt and is OPTIONAL to set on sending.

This document also updates Section 2.2 (bullet c) of [RFC4719], removing the following two sentences:

For ICRQ and ICRP, the Circuit Status AVP MUST indicate that the circuit status is for a new circuit (refer to N bit in Section 2.3.3).

For ICCN and SLI (refer to Section 2.3.2), the Circuit Status AVP MUST indicate that the circuit status is for an existing circuit (refer to N bit in Section 2.3.3) and reflect the current status of the link (refer to A bit in Section 2.3.3).

And finally, this document updates Section 3.1 of [RFC4349], Section 3.1 of [RFC4454], Section 3.1 of [RFC4591], and Section 2.2 of [RFC4719] with the following paragraph addition:

The usage of the N bit in the Circuit Status AVP is deprecated. Therefore, for ICRQ and ICRP, the Circuit Status AVP need not indicate on sending (nor check on receipt) that the circuit status is for a new circuit, and for ICCN and SLI, the Circuit Status AVP need not indicate on sending (nor check on receipt) that the circuit status is for an existing circuit.

5. Security Considerations

Security considerations for the Circuit Status AVP are covered in the base L2TPv3 specification (see Section 8 of [RFC3931]). No additional security considerations exist with extending this attribute.

6. IANA Considerations

The Circuit Status Bits number space [IANA-l2tp] is managed by IANA as per Section 10.7 of [RFC3931]. Five new bits (bits 9 through 13) and one updated bit (bit 14) have been assigned as follows:

Circuit Status Bits - per [RFC3931]

Bit 9 - S (Standby) bit
Bit 10 - E (Local PSN-facing PW (egress) Tx Fault) bit
Bit 11 - I (Local PSN-facing PW (ingress) Rx Fault) bit
Bit 12 - T (Local AC (egress) Tx Fault) bit
Bit 13 - R (Local AC (ingress) Rx Fault) bit
Bit 14 - N (New) bit [use deprecated]

7. Acknowledgements

The authors wish to thank Muhammad Yousuf, Mark Townsley, George Wilkie, Prashant Jhingran, Pawel Sowinski, and Ignacio Goyret for useful comments received.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [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.
- [RFC4454] Singh, S., Townsley, M., and C. Pignataro, "Asynchronous Transfer Mode (ATM) over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4454, May 2006.

- [RFC4591] Townsley, M., Wilkie, G., Booth, S., Bryant, S., and J. Lau, "Frame Relay over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4591, August 2006.
- [RFC4719] Aggarwal, R., Townsley, M., and M. Dos Santos, "Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4719, November 2006.

8.2. Informative References

- [IANA-l2tp] Internet Assigned Numbers Authority, "Layer Two Tunneling Protocol 'L2TP'", <<http://www.iana.org>>.
- [PREF-FWD] Muley, P., Bocci, M., and L. Martini, "Preferential Forwarding Status bit definition", Work in Progress, September 2008.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.
- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC5085] Nadeau, T. and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.
- [VCCV-BFD] Nadeau, T. and C. Pignataro, "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", Work in Progress, July 2009.

Authors' Addresses

Neil McGill
Cisco Systems
7025-4 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709
USA

EEmail: nmcgill@cisco.com

Carlos Pignataro
Cisco Systems
7200-12 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709
USA

EEmail: cpignata@cisco.com