

Network Working Group
Request for Comments: 3308
Category: Standards Track

P. Calhoun
Black Storm Networks
W. Luo
Cisco Systems, Inc.
D. McPherson
TCB
K. Peirce
Malibu Networks, Inc.
November 2002

Layer Two Tunneling Protocol (L2TP) Differentiated Services Extension

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 (2002). All Rights Reserved.

Abstract

This document describes mechanisms which enable the Layer Two Tunneling Protocol (L2TP) to negotiate desired Per Hop Behavior (PHB) code for the L2TP control connection, as well as individual sessions within an L2TP tunnel.

L2TP provides a standard method for tunneling PPP packets. The current specification provides no provisions for supporting Differentiated Services (diffserv) over the L2TP control connection or subsequent data sessions. As a result, no standard mechanism currently exists within L2TP to provide L2TP protocol negotiations for service discrimination.

Table of Contents

1.	Specification of Requirements	2
2.	Introduction	2
3.	Control Connection Operation	3
3.1.	Control Connection DS AVP (SCCRQ, SCCRP)	4
4.	Session Operation	4
4.1.	Session DS AVP (ICRQ, ICRP, OCRQ, OCRP)	6
5.	DS AVPs Correlation	6
6.	PHB Encoding	6
7.	DSCP Selection	7
8.	Packet Reordering and Sequence Numbers	7
9.	Crossing Differentiated Services Boundaries	7
10.	IANA Considerations	8
11.	Security Considerations	8
12.	Acknowledgements	8
13.	References	8
14.	Authors' Addresses	9
15.	Full Copyright Statement	10

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 [RFC 2119].

2. Introduction

The L2TP specification currently provides no mechanism for supporting diffserv (DS). This document describes mechanisms that enable L2TP to indicate desired PHB code, as defined in [RFC 3140], to be associated with an L2TP control connection, as well as individual sessions within an L2TP tunnel.

The actual bit interpretation of the DS field is beyond the scope of this document, and is purposefully omitted. This document is concerned only with defining a uniform exchange and subsequent mapping mechanism for the DS AVPs.

3. Control Connection Operation

As defined in [RFC 2661], a control connection operates in-band over a tunnel to control the establishment, release, and maintenance of sessions and of the tunnel itself. As such, this document provides a mechanism to enable discrimination of L2TP control messages from other packets. For this purpose, we introduce the Control Connection DS (CCDS) AVP.

The presence of the CCDS AVP serves as an indication to the peer (LAC or LNS) that the tunnel initiator wishes both the tunnel initiator and terminator to use the per-hop behavior(s) (PHB(s)) indicated by the AVP's PHB code for all packets within the tunnel's control connection. A PHB is a description of the externally observable forwarding behavior of a DS node applied to a particular DS behavior aggregate, as defined in [RFC 2475]. The most simple example of a PHB is one which guarantees a minimal bandwidth allocation of a link to a behavior aggregate.

Upon receipt of a Start-Control-Connection-Request (SCCRQ) containing the CCDS AVP, if the tunnel terminator provides no support for the CCDS AVP it MUST ignore the AVP and send an SCCRCP to the tunnel initiator without the CCDS AVP. The tunnel initiator interprets the absence of the CCDS AVP in the SCCRCP as an indication that the tunnel terminator is incapable of supporting CCDS.

Upon receipt of an SCCRCP that contains no CCDS AVP in response to a SCCRQ that contained a CCDS AVP, if the tunnel initiator wants to continue tunnel establishment it sends an SCCCEN. Otherwise, it sends a StopCCN to the tunnel terminator to end the connection. The StopCCN control message MUST contain the Result Code 8 that indicates CCDS AVP value (47) as the reason for sending the StopCCN.

If the tunnel terminator provides support for CCDS, it SHOULD use the Host Name AVP embedded in SCCRQ to consult its local policy, and to determine whether local policy permits the requested PHB code to be used on this control connection. If it is unwilling or unable to support the requested PHB code after consulting the local policy, the tunnel terminator MUST send an SCCRCP control message containing a CCDS AVP indicating the value it is willing to use. If the CCDS AVP value is the same as the one in the SCCRQ, it signals the acceptance of the requested PHB code. If the value is different it serves as a counter-offer by the tunnel terminator.

If the tunnel initiator receives an SCCRP that contains a CCDS AVP with a value other than that requested in the SCCRQ, the tunnel initiator SHOULD check the PHB code against its own policy. If it is unwilling to use the value, the tunnel initiator MUST send a StopCCN control message containing the Result Code 8 that indicates CCDS AVP value (47) as the reason for sending the StopCCN.

3.1. Control Connection DS AVP (SCCRQ, SCCRP)

The CCDS AVP is encoded as Vendor ID 0, and the Attribute Type is 47.

Each CCDS AVP is encoded as follows:

Vendor ID = 0
Attribute = 47

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
M H 0 0 0 0 0								Length								0															
47																PHB Code															

This AVP MAY be present in the following message types: SCCRQ and SCCRP. This AVP MAY be hidden (the H-bit set to 0 or 1) and is optional (M-bit not set). The length (before hiding) of this AVP MUST be 8 octets. The encoding of the PHB code is described in Section 6.

4. Session Operation

As defined in [RFC 2661], an L2TP session is connection-oriented. The LAC and LNS maintain states for each call that is initiated or answered by an LAC. An L2TP session is created between the LAC and LNS when an end-to-end connection is established between a Remote System and the LNS. Datagrams related to the connection are sent over the tunnel between the LAC and LNS. As such, this document provides a mechanism to enable discrimination for packets within a particular session from those in other sessions. For this purpose, we introduce the Session DS (SDS) AVP.

The presence of the SDS AVP serves as an indication to the peer (LAC or LNS) that the session initiator wishes both the session initiator and terminator to use the per-hop behavior(s) (PHB(s)) indicated by the AVP's PHB code for all packets within the session.

Upon receipt of an Incoming-Call-Request (ICRQ) or Outgoing-Call-Request (OCRQ) containing the SDS AVP if the session terminator provides no support for the requested PHB code, the session terminator MUST ignore the SDS AVP and send an ICRP or OCRP to the session initiator without the SDS AVP. The session initiator interprets the absence of the SDS AVP in the ICRP or OCRP as an indication that the session terminator is incapable of supporting SDS.

Upon receipt of an ICRP or OCRP that contains no SDS AVP in response to an ICRQ or OCRQ that contained an SDS AVP, if the session initiator is willing to omit employing SDS AVP it continues session establishment as defined in [RFC 2661]. Otherwise, it sends a CDN to the session terminator to end the connection. The CDN control message MUST contain the Result Code 12 that indicates SDS AVP value (48) as the reason for sending the CDN.

In order to help the session terminator to distinguish one session from another when consulting the local policy of the PHB code, the session initiator MAY use the identifier or a combination of identifiers embedded in AVPs such as Proxy Authen Name AVP, Calling Number AVP, Called Number AVP, and Sub-Address AVP. When Proxy Authen Name AVP is used as a distinguisher, it SHOULD be present in the ICRQ or OCRQ. The designated DS identifier(s) used for looking up the PHB code SHOULD be configurable.

If the session terminator provides support for SDS, it SHOULD use the the designated DS identification AVP (via out-of-band agreement between the administrators of the LAC and LNS) to consult the local policy and determinate whether the local policy permits the requested PHB code to be used on this session. If it is unwilling or unable to support the requested PHB code the session terminator MUST do one of the following:

- 1) Send a CDN message containing the Result Code 12 that indicates SDS AVP value (48) as the reason for sending the CDN.
- 2) Send an Incoming-Call-Reply (ICRP) or Outgoing-Call-Reply (OCRP) message containing an SDS AVP indicating the PHB code the terminator is willing to use for the session.

If the session terminator supports the PHB code in the SDS AVP session establishment MUST continue as defined in [RFC 2661].

If the session initiator receives an ICRP or OCRP that contains an SDS AVP with a value other than that requested in the ICRQ or OCRQ, and the session initiator is unwilling to use the value, the session initiator MUST send a CDN message containing the Result Code 12 that

indicates SDS AVP value (48) as the reason for sending the CDN. If the session initiator receives an ICRP or OCRP that contains an SDS AVP with a value other than that requested in the ICRQ or OCRQ, and the session initiator is willing to use the value, the session initiator MUST proceed as indicated in [RFC 2661].

4.1. Session DS AVP (ICRQ, ICRP, OCRQ, OCRP)

The SDS AVP is encoded as Vendor ID 0, and the Attribute Value is 48.

Each SDS AVP is encoded as follows:

Vendor ID = 0
Attribute = 48

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
M H 0 0 0 0								Length								0															
48																PHB Code															

This AVP MAY be present in the following message types: ICRQ, ICRP, OCRQ and OCRP. This AVP MAY be hidden (the H-bit set to 0 or 1) and is optional (M-bit is not set 0). The length (before hiding) of this AVP MUST be 8 octets. The encoding of the PHB code is described in Section 6.

5. DS AVPs Correlation

CCDS AVP and SDS AVP are independent of each other. CCDS AVP is used to signal diffserv for the control connection between two L2TP peers, while SDS AVP is used for data connection. The PHB code signaled in one AVP SHOULD NOT have any implication on the PHB code signaled in the other AVP. Implementations MAY choose to implement either or both DS AVPs, and operations MAY choose to enable diffserv on either or both types of connections.

6. PHB Encoding

The PHB code is a left-justified 16-bit field using Per Hop Behavior (PHB) encoding defined in [RFC 3140]. Note that [RFC 3140] and its successor are the ultimate authority defining PHB encoding.

Upon successful establishment of an L2TP tunnel control connection or individual L2TP session employing the appropriate DS AVP defined in this document, both LAC and LNS MUST use their own PHB-to-DSCP mappings of their present DS domains to map the PHB to a DSCP and place it in the DS field of the outer IP header of packets transmitted on the connection.

7. DSCP Selection

The requirements or rules of each service and DSCP mapping are set through administrative policy mechanisms which are outside the scope of this document.

8. Packet Reordering and Sequence Numbers

[RFC 2474] RECOMMENDS that PHB implementations not cause reordering of packets within an individual connection. [RFC 3140] requires that a set of PHBs signaled using a single PHB ID MUST NOT cause additional packet reordering within an individual connection vs. using a single PHB. Since the CCDS and SDS AVPs contain one PHB ID, use of diffserv PHBs in accordance with this specification should not cause additional packet reordering within an L2TP control or data connection.

Sequence numbers are required to be present in all control messages and are used to provide reliable delivery on the control connection, as defined in [RFC 2661]. While packet reordering is inevitably as much a function of the network as it is local traffic conditioning, the probability of it occurring when employing the CCDS AVP is same as when not employing the AVP. Data messages MAY use sequence numbers to reorder packets and detect lost packets.

9. Crossing Differentiated Services Boundaries

With the potential that an L2TP connection traverses an arbitrary number of DS domains, signaling PHBs via L2TP is more appropriate than signaling DSCPs, because it maintains a consistent end-to-end differentiated service for the L2TP connection. As per [RFC 2983], the negotiated PHBs are mapped to locally defined DSCPs of the current DS domain at the tunnel ingress node. At the DS domain boundary nodes, the DSCPs can be rewritten in the DS field of the outer IP header, so that the DSCPs are always with respect to whatever DS domain the packet happens to be in.

As a result, it is perfectly acceptable that the outermost DS field of packets arriving on a given control connection or session are not marked with the same DSCP value that was used by the tunnel ingress node.

10. IANA Considerations

This document defines 2 L2TP Differentiated Services Extension AVPs. The IANA has assigned the value of 47 for the "CCDS AVP" defined in section 5.1 and the value of 48 for SDS AVP defined in section 6.1.

IANA has also assigned L2TP Result Code values of 8 for disconnecting control connection due to mismatching CCDS value (StopCCN), and 12 for disconnecting call due to mismatching SDS value (CDN).

11. Security Considerations

This encoding in itself raises no security issues. However, users of this encoding should consider that modifying a DSCP MAY constitute theft or denial of service, so protocols using this encoding MUST be adequately protected. No new security issues beyond those discussed in [RFC 2474] and [RFC 2475] are introduced here.

12. Acknowledgements

Many thanks to David Black, W. Mark Townsley, Nishit Vasavada, Andy Koscinski and John Shriver for their review and insightful feedback.

13. References

- [RFC 1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC 2474] Nichols, K., Blake, S., Baker, F. and D. Black "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC 2475] Blake, S., Black, D., Carlson, Z., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [RFC 2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G. and B. Palter, "Layer 2 Tunnel Protocol (L2TP)", RFC 2661, August 1999.
- [RFC 2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, October 2000.

[RFC 3140] Black, D., Brim, S., Carpenter, B. and F. Le Faucheur,
"Per Hop Behavior Identification Codes", RFC 3140, June
2001.

14. Authors' Addresses

Pat R. Calhoun
110 Nortech Parkway
San Jose, CA 95134-2307

Phone: +1 408.941.0500
EMail: pcalhoun@bstormnetworks.com

Wei Luo
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134

Phone: +1 408.525.6906
EMail: luo@cisco.com

Danny McPherson
TCB

Phone: +1 303.470.9257
EMail: danny@tcb.net

Ken Peirce
Malibu Networks, Inc.
1107 Investment Blvd, Suite 250
El Dorado Hills, CA 95762

Phone: +1 916.941.8814
EMail: Ken@malibunetworks.com

15. 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

