

Network Working Group
Request for Comments: 4618
Category: Standards Track

L. Martini
E. Rosen
Cisco Systems, Inc.
G. Heron
A. Malis
Tellabs
September 2006

Encapsulation Methods for Transport of
PPP/High-Level Data Link Control (HDLC) over MPLS Networks

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

A pseudowire (PW) can be used to carry Point to Point Protocol (PPP) or High-Level Data Link Control (HDLC) Protocol Data Units over a Multiprotocol Label Switching (MPLS) network without terminating the PPP/HDLC protocol. This enables service providers to offer "emulated" HDLC, or PPP link services over existing MPLS networks. This document specifies the encapsulation of PPP/HDLC Packet Data Units (PDUs) within a pseudowire.

Table of Contents

1. Introduction	2
2. Specification of Requirements	2
3. Applicability Statement	5
4. General Encapsulation Method	6
4.1. The Control Word	6
4.2. MTU Requirements	8
5. Protocol-Specific Details	9
5.1. HDLC	9
5.2. Frame Relay Port Mode	9
5.3. PPP	10
6. Using an MPLS Label as the Demultiplexer Field	11
6.1. MPLS Shim EXP Bit Values	11
6.2. MPLS Shim S Bit Value	11
7. Congestion Control	12
8. IANA Considerations	12
9. Security Considerations	12
10. Normative References	13
11. Informative References	13

1. Introduction

A PPP/HDLC pseudowire (PW) allows PPP/HDLC Protocol Data Units (PDUs) to be carried over an MPLS network. In addressing the issues associated with carrying a PPP/HDLC PDU over an MPLS network, this document assumes that a PW has been set up by some means outside the scope of this document. This may be via manual configuration, or using a signaling protocol such as that defined in [RFC4447].

The following figure describes the reference models that are derived from [RFC3985] to support the HDLC/PPP PW emulated services. The reader is also assumed to be familiar with the content of the [RFC3985] document.

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].

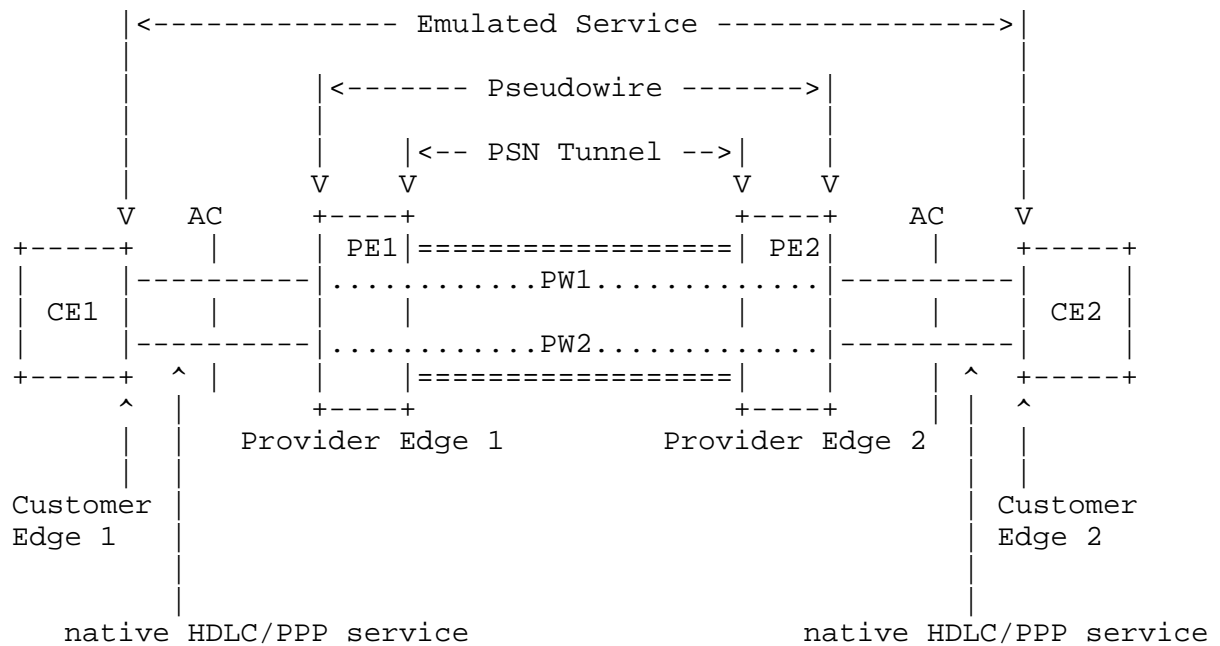


Figure 1. PWE3 HDLC/PPP interface reference configuration

This document specifies the emulated PW encapsulation for PPP and HDLC; however, quality of service related issues are not discussed in this document. For the purpose of the discussion in this document, PE1 will be defined as the ingress router and PE2 as the egress router. A layer 2 PDU will be received at PE1, encapsulated at PE1, transported across the network, decapsulated at PE2, and transmitted out on an attachment circuit at PE2.

The following reference model describes the termination point of each end of the PW within the PE:

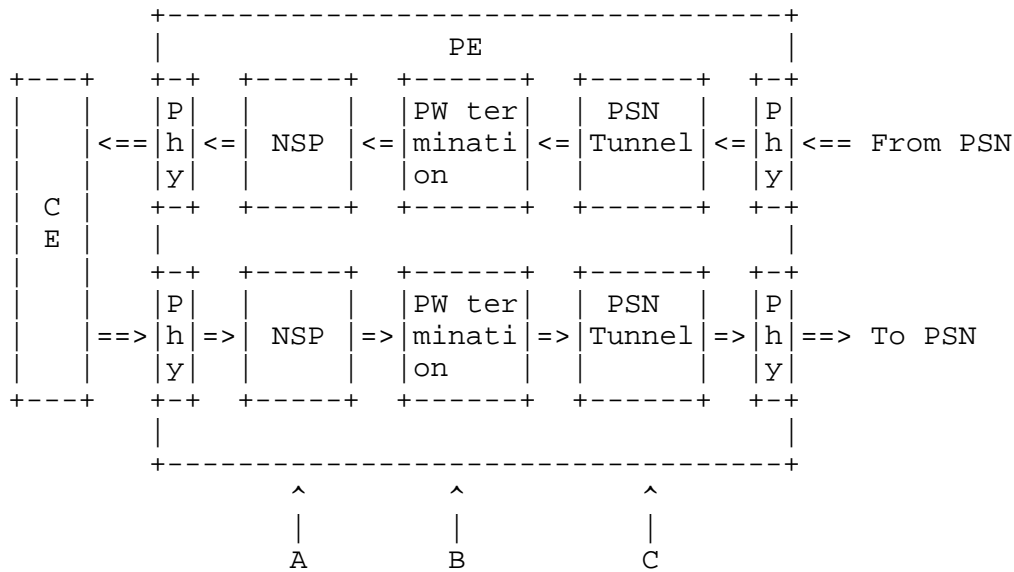


Figure 2. PW reference diagram

The PW terminates at a logical port within the PE, defined at point B in the above diagram. This port provides an HDLC Native Service Processing function that will deliver each PPP/HDLC packet that is received at point A, unaltered, to the point A in the corresponding PE at the other end of the PW.

The Native Service Processing (NSP) function includes packet processing that is required for the PPP/HDLC packets that are forwarded to the PW termination point. Such functions may include bit stuffing, PW-PW bridging, L2 encapsulation, shaping, and policing. These functions are specific to the native packet technology and may not be required for the PW emulation service.

The points to the left of B, including the physical layer between the CE and PE, and any adaptation (NSP) functions between it and the PW terminations, are outside of the scope of PWE3 and are not defined here.

"PW Termination", between A and B, represents the operations for setting up and maintaining the PW, and for encapsulating and decapsulating the PPP/HDLC packets as necessary to transmit them across the MPLS network.

3. Applicability Statement

PPP/HDLC transport over PW service is not intended to emulate the traditional PPP or HDLC service perfectly, but it can be used for some applications that require PPP or HDLC transport service.

The applicability statements in [RFC4619] also apply to the Frame Relay port mode PW described in this document.

The following are notable differences between traditional PPP/HDLC service, and the protocol described in this document:

- Packet ordering can be preserved using the OPTIONAL sequence field in the control word; however, implementations are not required to support this feature.
- The Quality of Service model for traditional PPP/HDLC links can be emulated, however this is outside the scope of this document.
- A Frame Relay Port mode PW, or HDLC PW, does not process any frame relay status messages or alarms as described in [Q922] [Q933].
- The HDLC Flags are processed locally in the PE connected to the attachment circuit.

The HDLC mode is suitable for port-to-port transport of Frame Relay User Network Interface (UNI) or Network Node Interface (NNI) traffic. Since all packets are passed in a largely transparent manner over the HDLC PW, any protocol that has HDLC-like framing may use the HDLC PW mode, including PPP, Frame-Relay, and X.25. Exceptions include cases where direct access to the HDLC interface is required, or modes that operate on the flags, Frame Check Sequence (FCS), or bit/byte unstuffing that is performed before sending the HDLC PDU over the PW. An example of this is PPP Asynchronous-Control-Character-Map (ACCM) negotiation.

For PPP, since media-specific framing is not carried, the following options will not operate correctly if the PPP peers attempt to negotiate them:

- Frame Check Sequence (FCS) Alternatives
- Address-and-Control-Field-Compression (ACFC)
- Asynchronous-Control-Character-Map (ACCM)

Note, also, that PW LSP Interface MTU negotiation, as specified in [RFC4447], is not affected by PPP Maximum Receive Unit (MRU)

advertisement. Thus, if a PPP peer sends a PDU with a length in excess of that negotiated for the PW tunnel, that PDU will be discarded by the ingress router.

4. General Encapsulation Method

This section describes the general encapsulation format for PPP and HDLC packets over MPLS pseudowires.




Figure 3. General format for PPP/HDLC encapsulation over PSNs

The PSN Transport Header depends on the particular tunneling technology in use. This header is used to transport the encapsulated PPP/HDLC information through the packet-switched core.

The Pseudowire Header identifies a particular PPP/HDLC service on a tunnel. In case the of MPLS, the Pseudowire Header is the MPLS label at the bottom of the MPLS label stack.

The Control Word is inserted before the PPP/HDLC service payload. It may contain a length and sequence number.

4.1. The Control Word

There are four requirements that may need to be satisfied when transporting layer 2 protocols over an MPLS PSN:

- i. Sequentiality may need to be preserved.
- ii. Small packets may need to be padded in order to be transmitted on a medium where the minimum transport unit is larger than the actual packet size.
- iii. Control bits carried in the header of the layer 2 packet may need to be transported.

- iv. Creating an in-band associated channel for operation and maintenance communications.

The Control Word defined in this section is based on the Generic PW MPLS Control Word, as defined in [RFC4385]. It provides the ability to sequence individual packets on the PW and avoidance of equal-cost multiple-path load-balancing (ECMP) [RFC2992] and enables Operations and Management (OAM) mechanisms, including [VCCV].

[RFC4385] states, "If a PW is sensitive to packet mis-ordering and is being carried over an MPLS PSN that uses the contents of the MPLS payload to select the ECMP path, it MUST employ a mechanism which prevents packet mis-ordering." This is necessary because ECMP implementations may examine the first nibble after the MPLS label stack to determine whether the content of the labeled packet is IP. Thus, if the PPP protocol number of a PPP packet carried over the PW without a control word present begins with 0x4 or 0x6, it could be mistaken for an IPv4 or IPv6 packet. This could, depending on the configuration and topology of the MPLS network, lead to a situation where all packets for a given PW do not follow the same path. This may increase out-of-order packets on a given PW or cause OAM packets to follow a different path from that of actual traffic.

The features that the control word provides may not be needed for a given PPP/HDLC PW. For example, ECMP may not be present or active on a given MPLS network, and strict packet sequencing may not be required. If this is the case, the control word provides little value and is therefore optional. Early PPP/HDLC PW implementations have been deployed that do not include a control word or the ability to process one if present. To aid in backwards compatibility, future implementations **MUST** be able to send and receive packets without the control word.

In all cases, the egress PE MUST be aware of whether the ingress PE will send a control word over a specific PW. This may be achieved by configuration of the PEs, or by signaling, as defined in [RFC4447].

The control word is defined as follows:




Figure 4. MPLS PWE3 control word

In the above diagram, the first 4 bits are set to 0 in indicate a CW [RFC4385].

The next 4 bits provide space for carrying protocol-specific flags. These are not used for HDLC/PPP, and they MUST be set to 0 for transmitting and MUST be ignored upon receipt.

The next 2 bits are defined in [RFC4623].

The next 6 bits provide a length field, which is used as follows: If the packet's length (defined as the length of the layer 2 payload plus the length of the control word) is less than 64 bytes, the length field MUST be set to the packet's length. Otherwise, the length field MUST be set to zero. The value of the length field, if not zero, is used to remove any padding that may have been added by the MPLS network. If the control word is used and padding was added to the packet in transit on the MPLS network, then when the packet reaches the egress PE the padding MUST be removed before forwarding the packet.

The next 16 bits provide a sequence number that can be used to guarantee ordered packet delivery. The processing of the sequence number field is OPTIONAL.[RFC4385]

The sequence number space is a 16-bit, unsigned circular space. The sequence number value 0 is used to indicate an unsequenced packet.[RFC4385]

The procedures described in Section 4 of [RFC4385] MUST be followed to process the sequence number field.

4.2. MTU Requirements

The network MUST be configured with an MTU that is sufficient to transport the largest encapsulation packets. When MPLS is used as the tunneling protocol, for example, this is likely to be 12 or more bytes greater than the largest packet size. The methodology described in [RFC4623] MAY be used to fragment encapsulated packets that exceed the PSN MTU. However, if [RFC4623] is not used, then if the ingress router determines that an encapsulated layer 2 PDU exceeds the MTU of the PSN tunnel through which it must be sent, the PDU MUST be dropped.

If a packet is received on the attachment circuit that exceeds the interface MTU subTLV value [RFC4447], it MUST be dropped. It is also RECOMMENDED that PPP devices be configured to not negotiate PPP MRUs larger than that of the AC MTU.

5. Protocol-Specific Details

5.1. HDLC

HDLC mode provides port-to-port transport of HDLC-encapsulated traffic. The HDLC PDU is transported in its entirety, including the HDLC address and control fields, but excluding HDLC flags and the FCS. Bit/Byte stuffing is undone. If the OPTIONAL control word is used, then the flag bits in the control word are not used and MUST be set to 0 for transmitting and MUST be ignored upon receipt.

When the PE detects a status change in the attachment circuit status, such as an attachment circuit physical link failure, or if the AC is administratively disabled, the PE MUST send the appropriate PW status notification message that corresponds to the HDLC AC status. In a similar manner, the local PW status MUST also be reflected in a respective PW status notification message, as described in [RFC4447].


The PW of type 0x0006 "HDLC" will be used to transport HDLC packets. The IANA allocation registry of "Pseudowire Type" is defined in the IANA allocation document for PWs [RFC4446] along with initial allocated values.

5.2. Frame Relay Port Mode

Figure 5 illustrates the concept of frame relay port mode or many-to-one mapping, which is an OPTIONAL capability.

Figure 5a shows two frame relay devices physically connected with a frame relay UNI or NNI. Between their two ports, P1 and P2, n frame relay Virtual Circuits (VCs) are configured.

Figure 5b shows the replacement of the physical frame relay interface with a pair of PEs and a PW between them. The interface between a Frame Relay (FR) device and a PE is either an FR UNI or an NNI. All FR VCs carried over the interface are mapped into one HDLC PW. The standard frame relay Link Management Interface (LMI) procedures happen directly between the CEs. Thus with port mode, we have many-to-one mapping between FR VCs and a PW.



[Pn]: A port

Figure 5a. FR interface between two FR devices




Figure 5b. Pseudowires replacing the FR interface

FR VCs are not visible individually to a PE; there is no configuration of individual FR VC in a PE. A PE processes the set of FR VCs assigned to a port as an aggregate.

FR port mode provides transport between two PEs of a complete FR frame using the same encapsulation as described above for HDLC mode.

Although frame relay port mode shares the same encapsulation as HDLC mode, a different PW type is allocated in [RFC4446]: 0x000F Frame-Relay Port mode.

All other aspects of this PW type are identical to the HDLC PW encapsulation described above.

5.3. PPP

PPP mode provides point-to-point transport of PPP-encapsulated traffic, as specified in [RFC1661]. The PPP PDU is transported in its entirety, including the protocol field (whether compressed using Protocol Field Compression or not), but excluding any media-specific framing information, such as HDLC address and control fields or FCS.

If the OPTIONAL control word is used, then the flag bits in the control word are not used and MUST be set to 0 for transmitting and MUST be ignored upon receipt.

When the PE detects a status change in the attachment circuit (AC) status, such as an attachment circuit physical link failure, or if the AC is administratively disabled, the PE MUST send the appropriate PW status notification message that corresponds to the PPP AC status. Note that PPP negotiation status is transparent to the PW and MUST NOT be communicated to the remote MPLS PE. In a similar manner, the local PW status MUST also be reflected in a respective PW status notification message, as described in [RFC4447].

A PW of type 0x0007 "PPP" will be used to transport PPP packets.

The IANA allocation registry of "Pseudowire Type" is defined in the IANA allocation document for PWs [RFC4446] along with initial allocated values.

6. Using an MPLS Label as the Demultiplexer Field

To use an MPLS label as the demultiplexer field, a 32-bit label stack entry [RFC3032] is simply prepended to the emulated PW encapsulation and thus appears as the bottom label of an MPLS label stack. This label may be called the "PW label". The particular emulated PW identified by a particular label value must be agreed by the ingress and egress LSRs, either by signaling (e.g., via the methods of [RFC4447]) or by configuration. Other fields of the label stack entry are set as described below.

6.1. MPLS Shim EXP Bit Values

If it is desired to carry Quality of Service information, the Quality of Service information SHOULD be represented in the EXP field of the PW label. If more than one MPLS label is imposed by the ingress LSR, the EXP field of any labels higher in the stack MUST also carry the same value.

6.2. MPLS Shim S Bit Value

The ingress LSR, PE1, MUST set the S bit of the PW label to a value of 1 to denote that the PW label is at the bottom of the stack.

7. Congestion Control

As explained in [RFC3985], the PSN carrying the PW may be subject to congestion, the characteristics of which are dependent upon PSN type, network architecture, configuration, and loading. During congestion, the PSN may exhibit packet loss that will impact the service carried by the PPP/HDLC PW. In addition, since PPP/HDLC PWs carry an unspecified type of services across the PSN, they cannot behave in a TCP-friendly manner prescribed by [RFC2914]. In the presence of services that reduce transmission rate, PPP/HDLC PWs will thus consume more than their fair share and SHOULD be halted.

Whenever possible, PPP/HDLC PWs should be run over traffic-engineered PSNs providing bandwidth allocation and admission control mechanisms. IntServ-enabled domains providing the Guaranteed Service (GS) or DiffServ-enabled domains using EF (expedited forwarding) are examples of traffic-engineered PSNs. Such PSNs will minimize loss and delay while providing some degree of isolation of the PPP/HDLC PW's effects from neighboring streams.

The PEs SHOULD monitor for congestion (by using explicit congestion notification, [VCCV], or by measuring packet loss) in order to ensure that the service using the PPP/HDLC PW may be maintained. When significant congestion is detected, the PPP/HDLC PW SHOULD be administratively disabled. If the PW has been set up using the protocol defined in [RFC4447], then procedures specified in [RFC4447] for status notification can be used to disable packet transmission on the ingress PE from the egress PE. The PW may be restarted by manual intervention, or by automatic means after an appropriate waiting time.

8. IANA Considerations

This document has no new IANA Actions. All necessary IANA actions have already been included in [RFC4446].

9. Security Considerations

The PPP and HDLC pseudowire type is subject to all the general security considerations discussed in [RFC3985][RFC4447]. This document specifies only encapsulations, and not the protocols that may be used to carry the encapsulated packets across the MPLS network. Each such protocol may have its own set of security issues, but those issues are not affected by the encapsulations specified herein.

10. Normative References

- [RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [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.
- [RFC4619] Martini, L., Ed., Kawa, C., Ed., and A. Malis, Ed., "Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks", RFC 4619, September 2006.
- [RFC4623] Malis, A. and M. Townsley, "Pseudowire Emulation Edge-to-Edge (PWE3) Fragmentation and Reassembly", RFC 4623, August 2006.

11. Informative References

- [Q922] ITU-T Recommendation Q.922 Specification for Frame Mode Basic call control, ITU Geneva 1995.
- [Q933] ITU-T Recommendation Q.933 Specification for Frame Mode Basic call control, ITU Geneva 2003.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, September 2000.
- [RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path Algorithm", RFC 2992, November 2000.
- [RFC3985] Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.

[VCCV] Nadeau, T., et al., "Pseudo Wire Virtual Circuit Connection Verification (VCCV)", Work in Progress, October 2005.

Contributing Author Information

Yeongil Seo
463-1 KT Technology Lab
Jeonmin-dong Yusung-gu
Daejeon, Korea

EMail: syil@kt.co.kr

Toby Smith
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205

EMail: tob@laurelnetworks.com

Authors' Addresses

Luca Martini
Cisco Systems, Inc.
9155 East Nichols Avenue, Suite 400
Englewood, CO, 80112

EMail: lmartini@cisco.com

Giles Heron
Tellabs
Abbey Place
24-28 Easton Street
High Wycombe
Bucks
HP11 1NT
UK

EMail: giles.heron@tellabs.com

Eric C. Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719

EMail: erosen@cisco.com

Andrew G. Malis
Tellabs
1415 West Diehl Road
Naperville, IL 60563

EMail: Andy.Malis@tellabs.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).

