

Internet Engineering Task Force (IETF)
Request for Comments: 7173
Category: Standards Track
ISSN: 2070-1721

L. Yong
D. Eastlake 3rd
S. Aldrin
Huawei
J. Hudson
Brocade
May 2014

Transparent Interconnection of Lots of Links (TRILL) Transport Using Pseudowires

Abstract

This document specifies how to interconnect a pair of Transparent Interconnection of Lots of Links (TRILL) switch ports using pseudowires under existing TRILL and Pseudowire Emulation End-to-End (PWE3) standards.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7173>.

Copyright Notice

Copyright (c) 2014 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction.....	2
1.1. Conventions Used in This Document.....	2
2. PWE3 Interconnection of TRILL Switches.....	3
2.1. PWE3 Type-Independent Details.....	3
2.2. PPP PWE3 Transport of TRILL.....	4
3. Security Considerations.....	6
Appendix A. Use of Other Pseudowire Types	7
Acknowledgements	8
Normative References	9
Informative References	10

1. Introduction

The Transparent Interconnection of Lots of Links (TRILL) protocol [RFC6325] provides optimal pair-wise data frame routing without configuration in multi-hop networks with arbitrary topology. TRILL supports multipathing of both unicast and multicast traffic. Devices that implement TRILL are called TRILL switches or Routing Bridges (RBridges).

Links between TRILL switches can be based on arbitrary link protocols, for example, PPP [RFC6361], as well as Ethernet [RFC6325]. A set of connected TRILL switches together form a TRILL campus that is bounded by end stations and Layer 3 routers.

This document specifies how to interconnect a pair of TRILL switch ports using a pseudowire under existing TRILL and PWE3 (Pseudowire Emulation End-to-End) standards.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Acronyms used in this document include the following:

IS-IS - Intermediate System to Intermediate System [IS-IS]

MPLS - Multi-Protocol Label Switching

PPP - Point-to-Point Protocol [RFC1661]

PW - Pseudowire [RFC3985]

PWE3 - PW Emulation End-to-End

RBridge - Routing Bridge, an alternative name for a TRILL switch

TRILL - Transparent Interconnection of Lots of Links [RFC6325]

TRILL Switch - A device implementing the TRILL protocol

2. PWE3 Interconnection of TRILL Switches

When a pseudowire is used to interconnect a pair of TRILL switch ports, a PPP [RFC4618] pseudowire is used as described below. The pseudowire between such ports can be signaled [RFC4447] or manually configured. In this context, the TRILL switch ports at the ends of the pseudowire are acting as native service processing (NSP) elements [RFC3985] and, assuming that the pseudowires are over MPLS or IP [RFC4023] networks, as label switched or IP routers at the TRILL switch ports.

Pseudowires provide transparent transport, and the two TRILL switch ports appear directly interconnected with a transparent link. With such an interconnection, the TRILL adjacency over the link is automatically discovered and established through TRILL IS-IS control messages [RFC7177].

A pseudowire is carried over a packet switched network tunnel [RFC3985], for example, an MPLS or MPLS-TP label switched path tunnel in MPLS networks. Either a signaling protocol or manual configuration can be used to configure a label switched path tunnel between two TRILL switch ports. This application needs no additions to the existing pseudowire standards.

2.1. PWE3 Type-Independent Details

The sending pseudowire TRILL switch port SHOULD map the inner priority of the TRILL Data packets being sent to the Traffic Class field of the pseudowire label [RFC5462] so as to minimize the probability that higher priority TRILL Data packets will be discarded due to excessive TRILL Data packets of lower priority.

TRILL IS-IS PDUs critical to establishing and maintaining adjacency (Hello and MTU PDUs) SHOULD be sent with the MPLS Traffic Class that calls for handling with the maximum priority. Other TRILL IS-IS PDUs SHOULD be sent with the MPLS Traffic Class denoting the highest priority that is less than the maximum priority. TRILL Data packets SHOULD be sent with appropriate MPLS Traffic Classes, typically mapped from the TRILL Data packet priority, such that TRILL Data packet Traffic Classes denote priorities less than the priorities

used for TRILL IS-IS PDUs. This minimizes the probability of other traffic interfering with these important control PDUs and causing false loss of adjacency or other control problems.

If a pseudowire supports fragmentation and reassembly (a feature that has received little or no deployment), then there is no reason to do TRILL MTU testing on it, and the pseudowire will not be a constraint on the TRILL campus-wide MTU size (Sz) (see Section 4.3.1 of [RFC6325]). If the pseudowire does not support fragmentation (the more common case), then the available TRILL IS-IS packet payload size over the pseudowire (taking into account MPLS encapsulation with a control word) or some lower value, MUST be used in helping to determine MTU size (Sz) (see Section 5 of [RFC7180]).

An intervening MPLS label switched router or similar packet switched network device has no awareness of TRILL. Such devices will not change the TRILL Header hop count.

2.2. PPP PWE3 Transport of TRILL

For a PPP pseudowire (PW type = 0x0007), the two TRILL switch ports being connected are configured to form a pseudowire with PPP encapsulation [RFC4618]. After the pseudowire is established and TRILL use is negotiated within PPP, the two TRILL switch ports appear directly connected with a PPP link [RFC1661] [RFC6361].

If pseudowire interconnection of two TRILL switch ports is signaled [RFC4447], the initiating TRILL switch port MUST attempt the connection setup with pseudowire type PPP (0x0007).

Behavior for TRILL with a PPP pseudowire continues to follow that of TRILL over PPP as specified in Section 3 of [RFC6361].

The following figures show what a TRILL Data packet and TRILL IS-IS packet look like over such a pseudowire in the MPLS case, assuming no TRILL Header extensions:

Server MPLS Tunnel Label(s)	n*4 octets (4 octets per label)
PW Label	4 octets
Control Word	4 octets
PPP Header 0x005d	2 octets
TRILL Header	6 octets
Destination MAC Address	6 octets
Source MAC Address	6 octets
Data Label	4 or 8 octets
Payload Body	variable

Figure 1: TRILL Data Packet in Pseudowire

"Data Label" is the VLAN Label or Fine-Grained Label [RFC7172] of the payload.

Server MPLS Tunnel Label(s)	n*4 octets (4 octets per label)
PW Label	4 octets
Control Word	4 octets
PPP Header 0x405d	2 octets
Common IS-IS Header	8 octets
IS-IS PDU Type Specific Header	variable
IS-IS TLVs	variable

Figure 2: TRILL IS-IS Packet in Pseudowire

The PPP Header fields (0x005d and 0x405d, respectively) for TRILL Data and IS-IS packets shown above are specified in [RFC6361].

3. Security Considerations

TRILL-level security mechanisms, such as the ability to use authentication with TRILL IS-IS PDUs [RFC6325], are not affected by link technology, such as the use of pseudowire links as specified in this document.

Link security may be useful in improving TRILL campus security. TRILL is transported over pseudowires as TRILL over PPP over pseudowires, pseudowires are over MPLS or IP, and MPLS and IP are over some lower-level link technology. Thus, link security below the TRILL level for a pseudowire link could be provided by PPP security, pseudowire security, MPLS or IP security, or security of the link technology supporting MPLS or IP.

PPP TRILL security considerations are discussed in [RFC6361]. For security considerations introduced by carrying PPP TRILL links over pseudowires, see [RFC3985], which discusses the risks introduced by sending protocols that previously assumed a point-to-point link on a pseudowire built on a packet switched network (PSN). However, the PPP layer in TRILL transport by pseudowire is somewhat vestigial and intended primarily as a convenient way to use existing PPP code points to identify TRILL Data packets and TRILL IS-IS packets. Furthermore, existing PPP security standards are arguably questionable in terms of current security criteria. For these reasons, it is NOT RECOMMENDED to use PPP security in the transport of TRILL by pseudowires as specified in this document.

It is RECOMMENDED that link security be provided at the layers supporting pseudowires transporting TRILL, that is, at the MPLS or IP layer or the link layer transporting MPLS or IP.

For applications involving sensitive data, end-to-end security should always be considered, in addition to link security, to provide security in depth. In this context, such end-to-end security should be between the end stations involved so as to protect the entire path to, through, and from the TRILL campus.

For general TRILL protocol security considerations, see [RFC6325].

Appendix A. Use of Other Pseudowire Types

This informational appendix briefly discusses the use of pseudowire types other than PPP for the transport of TRILL.

The use of Ethernet pseudowires [RFC4448] was examined by the authors and would be possible without change to such pseudowires; however, this would require an additional 12 or 16 bytes per packet within the payload being transmitted over the pseudowire for a TRILL Data packet (Figure 3) and a TRILL IS-IS packet (Figure 4) over such an Ethernet pseudowire in the MPLS case, assuming no TRILL Header extensions (compare with Figures 1 and 2):

Server MPLS Tunnel Label(s)	n*4 octets (4 octets per label)
PW Label	4 octets
Optional Control Word	4 octets
TRILL Hop Dest. MAC Address	6 octets
TRILL Hop Source MAC Address	6 octets
Optional VLAN and/or other tags	variable
TRILL Ethertype (0x22f3)	2 octets
TRILL Header	6 octets
Destination MAC Address	6 octets
Source MAC Address	6 octets
Data Label	4 or 8 octets
Payload Body	variable

Figure 3: TRILL Data Packet in Ethernet Pseudowire

"Data Label" is the VLAN Label or Fine-Grained Label [RFC7172] of the payload.

Server MPLS Tunnel Label(s)	n*4 octets (4 octets per label)
PW Label	4 octets
Optional Control Word	4 octets
TRILL Hop Dest. MAC Address	6 octets
TRILL Hop Source MAC Address	6 octets
Optional VLAN and/or other tags	variable
Layer 2 IS-IS Ethertype 0x22f4	2 octets
Common IS-IS Header	8 octets
IS-IS PDU Type Specific Header	variable
IS-IS TLVs	variable

Figure 4: TRILL IS-IS Packet in Ethernet Pseudowire

It would also be possible to specify a new pseudowire type for TRILL traffic, but the authors feel that any efficiency gain over PPP pseudowires would be too small to be worth the complexity of adding such a specification. Furthermore, using PPP pseudowire encoding means that any traffic dissector that understands TRILL PPP encoding [RFC6361] and PPP pseudowires [RFC4618] will automatically be able to recursively decode TRILL transported by pseudowire.

Acknowledgements

Thanks for the valuable comments from the following, who are listed in alphabetic order:

Stewart Bryant, Stephen Farrell, Brian Haberman, Christer Holmberg, Joel Jaeggli, Barry Leiba, Erik Nordmark, Yaron Sheffer, and Yaakov (J) Stein.

Normative References

- [RFC1661] Simpson, W., Ed., "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.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC4618] Martini, L., Rosen, E., Heron, G., and A. Malis, "Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks", RFC 4618, September 2006.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, February 2009.
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.
- [RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011.
- [RFC7172] Eastlake 3rd, D., Zhang, M., Agarwal, P., Perlman, R., and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL): Fine-Grained Labeling", RFC 7172, May 2014.
- [RFC7180] Eastlake 3rd, D., Zhang, M., Ghanwani, A., Manral, V., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", RFC 7180, May 2014.

Informative References

- [IS-IS] ISO/IEC 10589:2002, Second Edition, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", 2002.
- [RFC3985] Bryant, S., Ed., and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.
- [RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, April 2006.
- [RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, May 2014.

Authors' Addresses

Lucy Yong
Huawei Technologies
5340 Legacy Drive
Plano, TX 75024
USA

Phone: +1-469-227-5837
EMail: lucy.yong@huawei.com

Donald E. Eastlake 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757
USA

Phone: +1-508-333-2270
EMail: d3e3e3@gmail.com

Sam Aldrin
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Phone: +1-408-330-4517
EMail: sam.aldrin@huawei.com

Jon Hudson
Brocade
130 Holger Way
San Jose, CA 95134
USA

Phone: +1-408-333-4062
EMail: jon.hudson@gmail.com