

TRILL Working Group
INTERNET-DRAFT
Intended status: Informational

Lucy Yong
Donald Eastlake
Sam Aldrin
Huawei R&D USA
Jon Hudson
Brocade
March 10, 2012

Expires: September 9, 2012

Transparent Interconnection of Lots of Links (TRILL) over
MPLS Pseudo Wires
<[draft-yong-trill-trill-o-mpls-01.txt](#)>

Abstract

This informational document describes ways to interconnect TRILL R Bridges by using MPLS Pseudo Wire (PW) services with existing TRILL and MPLS standards so as to form a unified TRILL campus.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Distribution of this document is unlimited. Comments should be sent to the TRILL working group mailing list <rbridge@postel.org>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

INTERNET-DRAFT

TRILL over MPLS PW

Table of Contents

1.	Introduction.....	3
1.1	Conventions used in this document.....	3
2.	Use Cases.....	4
2.1	Point-To-Point Interconnection.....	4
2.1.1	Direct Point-to-Point.....	5
2.1.2	Provider Point-to-Point Service.....	5
2.2	Multi-Access Link Interconnection.....	6
3.	RBridge Behavior for MPLS Pseudo Wire.....	9
4.	IANA Considerations.....	10
5.	Security Considerations.....	10
6.	Acknowledgements.....	10
7.	Normative References.....	11
8.	Informative References.....	12

INTERNET-DRAFT

TRILL over MPLS PW

1. Introduction

The IETF TRILL (Transparent Interconnection of Lots of Links) standard [[RFC6325](#)] [[RFC6326](#)] provides optimal pair-wise data frame forwarding without configuration in multi-hop networks with arbitrary topology and link technology, and supports multipathing of both unicast and multicast traffic. TRILL enables a new method to construct a campus or data center network. Devices that implement TRILL are called RBridges (Routing Bridges) or TRILL Switches.

This document describes the use of MPLS Pseudo Wire or VPLS links by TRILL.

1.1 Conventions used in this document

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](#)].

Acronyms used in this document include the following:

AC - Attachment Circuit

IS-IS - Intermediate System to Intermediate System

MPLS - Multi-Protocol Label Switching

PE - Provider Edge

PPP – Point-to-Point Protocol

PW – Pseudo Wire

QoS – Quality of Service

RB – RBridge

RBridge – Routing Bridge

TRILL – Transparent Interconnection of Lots of Links

TRILL Switch – An alternative term for an RBridge

VPLS – Virtual Private LAN Service

VSI – Virtual Service Instance

[2.](#) Use Cases

TRILL campuses at different locations may interconnect by networks that are implemented with different technologies to form a unified RBridge campus. This section describes use cases assuming that IP/MPLS technology is available. From the MPLS network view, a pair of RBridges can be directly connected with a pseudo wire or an RBridge can act as a Customer Edge device that connects to a Provider Edge device via an attachment circuit. RBridge ports [[RFC6325](#)], by default, support both point-to-point links and multi-access links.

[Section 2.1](#) describes point-to-point links, i.e. TRILL over an Ethernet or PPP point-to-point link that is over an MPLS network. [Section 2.2](#) describes TRILL over a bridged LAN or equivalent that is implemented by MPLS/VPLS.

[2.1](#) Point-To-Point Interconnection

Either an Ethernet or PPP link over an MPLS network interconnects two RBridge ports. This can be either a direct pseudo wire between the RBridges or by attachment circuits from each RBridge port to Provider

Edge devices that provide a transparent tunnel between the provider edge attachment points.

MPLS already supports many pseudo wire transport encapsulations [[RFC4446](#)]. Two types of TRILL links between RBridges have been standardized are Ethernet [[RFC6325](#)] and PPP [[RFC6361](#)]. Pseudo wire encapsulations for these two interfaces are specified in [[RFC4448](#)] and [[RFC4618](#)], respectively.

The method described in 2.1.1 below is typically suitable when the TRILL and MPLS facilities have common management while the method described in 2.1.2 is typically suitable when the TRILL and MPLS facilities are separately managed. In the case of different management, the core MPLS operator can sell a VPWS service to an RBridge operator.

In both cases, the MPLS label switched routers involved need no awareness of TRILL.

A pseudo wire may cross multiple MPLS domains [[RFC5659](#)]. In these cases, RBridges may be considered to connect to T-PEs and it works in the same way as a single domain. The MPLS network can provide transport resiliency for a pseudo wire. The dual homing (two attachment circuits) can be used for attachment circuit protection. In this case, two TRILL links are established; RBridges can perform load balancing over two links.

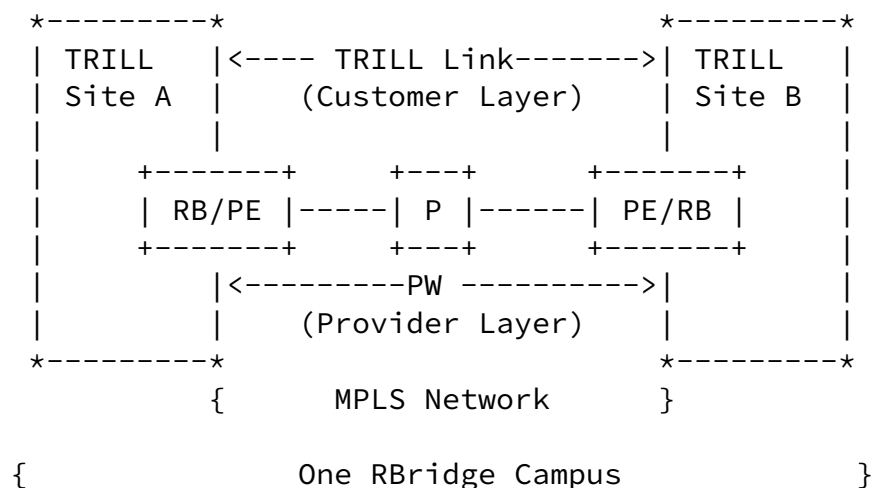
[2.1.1](#) Direct Point-to-Point

Two RBridge ports can be connected directly by an MPLS pseudo wire. This implies that the RBridges, which are TRILL routers, are also acting as label switched routers. The pseudo wire can be either Ethernet over MPLS or PPP over MPLS but PPP over MPLS is recommended because it saves 16 bytes per frame. The pseudo wire between two RBridge ports can be auto-configured [[RFC4447](#)] or manually configured; the two RBridges then appear directly interconnected with a transparent link.

(Technically speaking, it is possible to create a specially designated TRILL encapsulated pseudo wire for point-to-point TRILL over MPLS. However, the authors think that this is not worth the effort in this case because of available technologies, particularly

the highly-efficient PPP link technology.)


From a customer/provider point of view, this can also be thought of as shown in the following diagram:




The interworking between the RBridge network and the MPLS network is within the combined TRILL/MPLS device. This has a similar architecture to MPLS/VPLS [[RFC4762](#)].

2.1.2 Provider Point-to-Point Service

Two RBridge ports may also be connected by attachment circuits (ACs) to Provider Edge (PE) devices that are part of a typically separately managed provider network. The provider network then provides a transparent path between these attachment circuits, connecting the RBridge ports. The following diagram illustrates this arrangement:





If the attachment circuits are Ethernet, the two PEs would be configured to form a pseudo wire with Ethernet encapsulation [[RFC4448](#)]. This pseudo wire can be auto-configured [[RFC4447](#)] or manually configured; the two RBridges then appear directly interconnected with an Ethernet link. The provider edge devices SHOULD use the Raw mode and non-service-delimiting, which provides a transparent Ethernet transport.

If the attachment circuits are PPP link type, the two PEs must be configured to form a PW with PPP encapsulation [[RFC4618](#)]. After the pseudo wire is established between two PEs, the two RBridges then appear directly connected with a PPP link. The PPP link configuration will be more efficient than the Ethernet point-to-point configuration; it saves about 16 bytes per frame by replacing the TRILL Outer.MacDA, Outer.MacSA, Outer.VLAN, and outer Ethertype with a PPP code point [[RFC6361](#)].

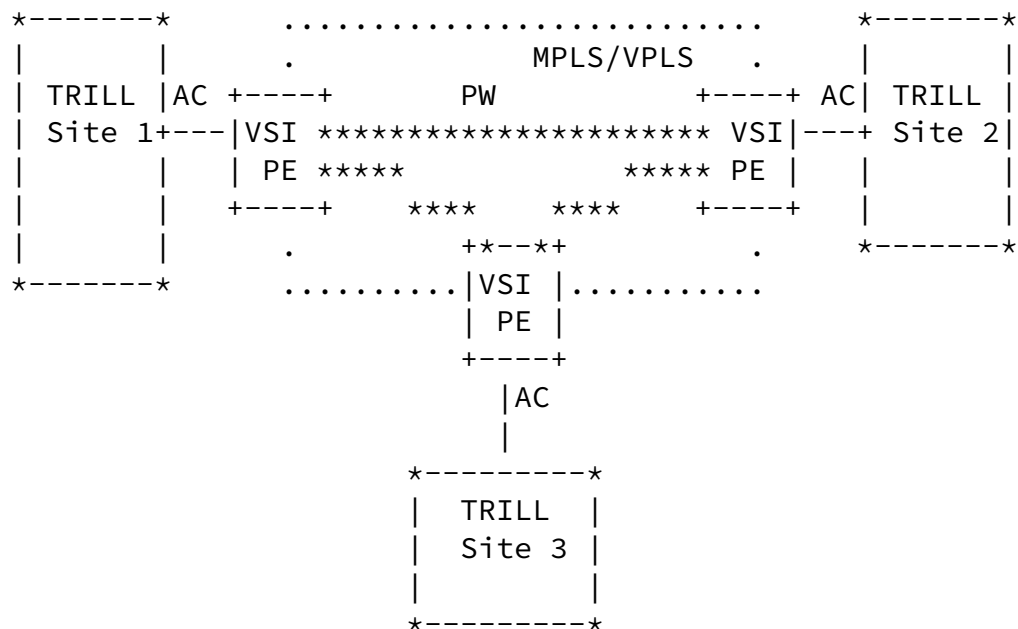
In both case, the pseudo wire provides transparent transport between attachment circuits. The TRILL link and the connectivity over that link is automatically discovered and established through TRILL IS-IS control messages.

[2.2](#) Multi-Access Link Interconnection

Multiple RBridges may interconnect via an [[802.10](#)] Bridged LAN that acts as a hub. The bridged LAN simply forwards on the outer Ethernet header of the TRILL over Ethernet frames. This configuration creates what appears to each connected RBridge as a multi-access link. In other words, each RBridge connecting to a bridged LAN normally has connectivity to every other RBridge connecting to the same bridged LAN.

MPLS/VPLS can provide the same capability when multiple parts of an RBridge campus are interconnected over an IP/MPLS network and make each RBridge attaching to the VPLS appear as having a multi-access TRILL link. The diagram below shows one RBridge campus is split between three different sites and the use of MPLS/VPLS for RBridge

interconnection. Ethernet attachment circuits and pseudo wires are assumed.



One VPLS instance is configured on three Provider Edge (PE) devices and the pseudo wires are configured for the VPLS instance. Each RBridge Site connects to the VSI on a PE via an attachment circuit. The VSI on a PE forwards TRILL frames based on the outer Ethernet header of the frames [\[RFC6325\]](#). Either BGP [\[RFC4761\]](#) or LDP [\[RFC4762\]](#) protocol can be used to automatically construct the VPLS instance on the PEs.

The choice of three VSIs and three singly connected sites was for illustrative purposes. There could be two or more than three. Furthermore, each TRILL site could have multiple connections to the VPLS network and/or direct connection via other technologies or VPLS networks to one or more of the other sites. TRILL sorts all this out and router properly.

A PE may connect to several different RBridge campuses that belong to different customers. Separated VPLS instances are configured for individual customers and customer traffic is isolated by VPLS instance. The PE treats an RBridge as a generic Ethernet customer devices and has no awareness of TRILL. The outer Ethernet MAC of TRILL frames may be either a next-hop RBridge MAC address (for unicast frames) or one of TRILL defined multicast addresses (ALL-IS-IS-RBridges and All-RBridges) [\[RFC6325\]](#). The VSI at each PE learns the source MAC addresses on each VSI interface and forward the frame based on the destination MAC. For the multicast frames, the VSI replicates the frames to all pseudo wires it associates. If a VPLS is configured with some optimization capability [\[VPLS-BCAST\]](#), the multicast frames can be delivered over a point-to-multipoint pseudo wire while unicast frames are carried over a point-to-point pseudo

INTERNET-DRAFT

TRILL over MPLS PW

wire.

The scenario above can also be extended to multiple RBridge interconnections when a device serves both the RBridge and PE functions, similarly to the case in [Section 2.1.2](#) and 2.1.1 above.

Note: If the customer devices associated with one VPLS instances happen to include some RBridges and some end stations or IEEE 802.1Q bridges providing paths to end stations, TRILL will, by default, be able to handle this by providing both through service and end station service. However, the end station addresses will be visible to the VPLS instance. If, in such a case, all the RBridge ports connected to the VPLS are configured as trunk ports (see [Section 4.9.2 of \[RFC6325\]](#)), then they will not provide any end station service.

INTERNET-DRAFT

TRILL over MPLS PW

3. RBridge Behavior for MPLS Pseudo Wire

This section describes RBridge behaviors for TRILL Ethernet or TRILL PPP links over MPLS pseudo wire (PW) as described in Sections [2.1](#).

1. For two RBridge ports connecting via a PPP pseudo wire, the ports MUST be configured as IS-IS point-to-point because there are no subnetwork point of attachment (SNPA/MAC) addresses of the end points at the PPP protocol layer. Thus TRILL will use IS-IS P2P Hellos that, as described in "Point-to-Point IS to IS Hello PDU" (section 9.7 of [\[IS-IS\]](#)), do not use Neighbor TLVs or require SNPAs. However, as described [section 4.2.4.1 of \[RFC6325\]](#), three-way IS-IS handshake using extended circuit IDs is required.
2. Any MPLS forwarder within an MPLS label switched path does not change the TRILL Header Hop Count. RBridges are not aware of the packet forwarders in with the MPLS network.
3. If it is desired for MPLS label switched routers to perform QoS in the same way as RBridges do, an Ethernet path MUST be used and RBridges MUST be configured to send an Outer.VLAN tag on the RBridge port leading to the pseudo wire. The PE can then copy the priority value from the Outer.VLAN tag to the COS field of the pseudo wire label prior to the forwarding [\[RFC5462\]](#).
4. TRILL MTU-probe and TRILL MTU-ack messages ([section 4.3.2 of \[RFC6325\]](#)) are not needed on a pseudo wire link. Implementations MUST NOT send MTU-probe and SHOULD NOT reply to these messages. The MTU pseudo wire interface parameter SHOULD be used instead. PE MUST configure the MTU size as the originating RBridges Size specified in [Section 4.3.1 of \[RFC6325\]](#).

[4.](#) IANA Considerations

No IANA action is required by this document. RFC Editor: Please remove this section before publication.

[5.](#) Security Considerations

The IS-IS authentication mechanism [[RFC5304](#)] [[RFC5310](#)], at the TRILL IS-IS layer, can be used to prevent fabrication of link-state control messages over TRILL links including those discussed in this document.

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

Use cases in which the path between RBridges transits a provider network under separate administration may represent a substantial increase in the threat of observation, deletion, modification, or insertion of data or control information. Under such circumstances consideration should be give to the use of security at the TRILL link level, such as [[802.1AE](#)] if the path between the RBridge ports is Ethernet or security as suggested in [[RFC6361](#)] if that path is PPP.

[6.](#) Acknowledgements

The authors sincerely acknowledge the contributions of Ben Mack-Crane and Sue Hares.

7. Normative References

- [RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", [BCP 116](#), [RFC 4446](#), April 2006.
- [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.
- [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.
- [RFC4618] Martini, L., "Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks",

[BCP 116](#), [RFC 4618](#), September 2006.

[RFC4761] Kompella, K. and Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC4761](#), January 2007.

[RFC4762] Lasserre, M. and Kompella, V, "Virtual Private LAN Service(VPLS) Using Label Distribution Protocol (LDP) Signaling",[RFC4762](#), January 2007

[RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels," [BCP 14](#) and [RFC 2119](#), March 1997

[RFC5304] Li, T. and Atkinson, R, "IS-IS Cryptographic Authentication," [RFC 5304](#), October 2008

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), February 2009

[RFC5462] Andersson, L. and Asati, R., "Multiprotocol Label Switching (MPLS) Label Stack entry: "Exp" Field Rename to "Traffic Class" Field", [RFC5462](#), February 2009

[RFC5659] Bocci, M and Bryant, S, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", [RFC 5659](#), October 2009.

[RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", [RFC6325](#), July 2011.

[RFC6326] Eastlake 3rd, D., Banerjee, A., Dutt, D., Perlman, R.,

D. Eastlake, et al

[Page 11]

INTERNET-DRAFT

TRILL over MPLS PW

and Ghanwani, A. "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [RFC6326](#), July 2011.

[RFC6361] Carlson, J., and D. Eastlake, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", [RFC6361](#), August 2011.

[8](#). Informative References

- [802.1AE] "IEEE Standard for Local and metropolitan area networks / Media Access Control (MAC) Security", 802.1AE-2006, 18 August 2006.
- [802.1Q] IEEE 802.1, "IEEE Standard for Local and metropolitan area networks - Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, May 2011.
- [IS-IS] International Organization for Standardization, "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)", ISO/IEC10589:2002, Second Edition, Nov 2002
- [VPLS-BCAST] Delord, S, and Key, R., "Extension to LDP-VPLS for Ethernet Broadcast and Multicast", [draft-ietf-l2vpn-ldp-vpls-broadcast-exten-02](#), work in progress, 2011.

Lucy Yong
Huawei R&D USA
5340 Legacy Drive
Plano, TX 75025

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

Donald E. Eastlake, 3rd
Huawei R&D USA
155 Beaver Street
Milford, MA 01757 USA

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

Sam Aldrin
Huawei R&D USA
2330 Central Expressway
Santa Clara, CA 95050

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

Jon Hudson
Brocade
130 Holger Way
San Jose, CA 95134

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

Copyright and IPR Provisions

Copyright (c) 2012 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. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of [RFC 5378](#). No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under [RFC 5378](#), shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

