

INTERNET-DRAFT
Intended Status: Informational

Mohammed Umair
Kingston Smiler Selvaraj
IPInfusion
Donald Eastlake 3rd
Lucy Yong
Huawei Technologies
October 12, 2016

Expires: April 15, 2017

TRILL Transparent Transport over MPLS
<[draft-ietf-trill-transport-over-mpls-02](#)>

Abstract

This document specifies how to interconnect multiple Transparent Interconnection of Lots of links (TRILL) sites with an intervening MPLS network using existing TRILL and VPLS standards. This draft addresses two problems as follows:

- 1) Providing connection between more than two TRILL sites that are separated by MPLS provider network.
- 2) Providing a single logical virtualized TRILL network for different tenants that are separated by an MPLS provider network.

Status of this Memo

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

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/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2016 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	4
1.1.	Terminology	4
2.	TRILL Over MPLS Model	5
3.	VPLS Model	6
3.1.	Entities in the VPLS Model	7
3.3.	TRILL Adjacency for VPLS model	8
3.4.	MPLS encapsulation for VPLS model	8
3.5.	Loop Free provider PSN/MPLS.	8
3.6.	Frame processing.	8
4.	VPTS Model	8
4.1.	Entities in the VPTS Model	10
4.1.1.	TRILL Intermediate Routers (TIR)	11
4.1.2.	Virtual TRILL Switch/Service Domain (VTSD)	11
4.2.	TRILL Adjacency for VPTS model	11
4.3.	MPLS encapsulation for VPTS model	12
4.4.	Loop Free provider PSN/MPLS.	12
4.5.	Frame processing.	12
4.5.1.	Multi-Destination Frame processing	12
4.5.2.	Unicast Frame processing	12
5.	Extensions to TRILL Over Pseudowires [RFC7173]	12
6.	VPTS Model Versus VPLS Model	13
7.	Packet processing between pseudowires	13
8.	Security Considerations	14
9.	IANA Considerations	14
10.	References	14
10.1.	Normative References	14

10.2 . Informative References	15
Acknowledgements	15
Authors' Addresses	15

1. Introduction

The IETF Transparent Interconnection of Lots of Links (TRILL) protocol [[RFC6325](#)] [[RFC7177](#)] [[RFC7780](#)] provides transparent forwarding in multi-hop networks with arbitrary topology and link technologies using a header with a hop count and link-state routing. TRILL provides optimal pair-wise forwarding without configuration, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. Intermediate Systems (ISs) implementing TRILL are called Routing Bridges(RBridges) or TRILL Switches

This draft, in conjunction with [[RFC7173](#)] on TRILL Transport using Pseudowires, addresses two problems:

- 1) Providing connection between more than two TRILL sites belongs to a single TRILL network that are separated by an MPLS provider network using [[RFC7173](#)]. (Herein also called problem statement 1.)
- 2) Providing a single logical virtualized TRILL network for different tenants that are separated by an MPLS provider network. In short providing connection between TRILL sites belonging to a tenant/tenants over a MPLS provider network. (Herein also called problem statement 2.)

A tenant is the administrative entity on whose behalf their associated services are managed. Here tenant refers to a TRILL campus that is segregated from other tenants for security reasons.

A key multi-tenancy requirement is traffic isolation so that one tenant's traffic is not visible to any other tenant. This draft also addresses the problem of multi-tenancy by isolating one tenant's traffic from the other.

1.1. Terminology

Acronyms used in this document include the following:

AC	- Attachment Circuit [RFC4664]
Data Label	- VLAN or FGL
ECMP	- Equal Cost Multi Path
FGL	- Fine-Grained Labeling [RFC7172]
IS-IS	- Intermediate System to Intermediate

System [[IS-IS](#)]

LDP	- Label Distribution Protocol
LAN	- Local Area Network
MPLS	- Multi-Protocol Label Switching
PE	- Provider Edge Device
PSN	- Packet Switched Network
PW	- Pseudowire [RFC4664]
TIR	- TRILL Intermediate Router (Devices that has both IP/MPLS and TRILL functionality)
TRILL	- Transparent Interconnection of Lots of Links OR Tunnelled Routing in the Link Layer
TRILL Site	- A part of a TRILL campus that contains at least one RBridge.
VLAN	- Virtual Local Area Network
VPLS	- Virtual Private LAN Service
VPTS	- Virtual Private TRILL Service
VSI	- Virtual Service Instance [RFC4664]
VTSD	- Virtual TRILL Switch Domain OR Virtual TRILL Service Domain A Virtual RBridge that segregates one tenant's TRILL database as well as traffic from the other.
WAN	- Wide Area Network

[2.](#) TRILL Over MPLS Model

TRILL Over MPLS can be achieved by two different ways.

a) VPLS Model for TRILL

b) VPTS Model/TIR Model

Both these models can be used to solve the problem statement 1 and 2. Herein the VPLS Model for TRILL is also called Model 1 and the VPTS Model/TIR Model is also called Model 2.

3. VPLS Model

Figure 1 shows the topological model of TRILL over MPLS using VPLS model. The PE routers in the below topology model should support all the functional Components mentioned in [[RFC4664](#)].

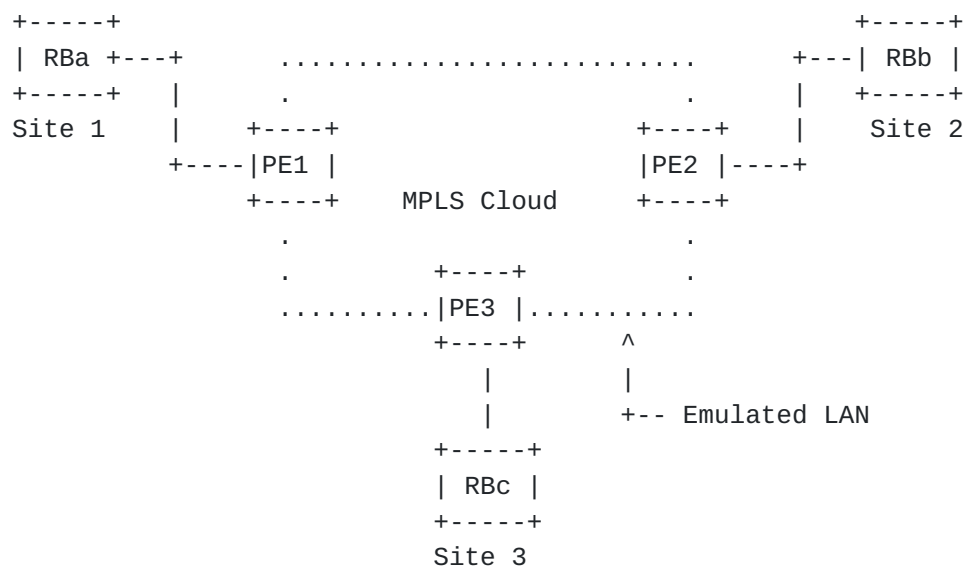


Figure 1: Topological Model of TRILL over MPLS
connecting three TRILL Sites

Figure 2 below shows the topological model of TRILL over MPLS to connect multiple TRILL sites belonging to a tenant (tenant here is a campus, not a Data label). VSI1 and VSI2 are two Virtual Service Instances that segregate Tenant1's traffic from other tenant traffic. VSI1 will maintain its own database for Tenant1, similarly VSI2 will maintain its own database for Tenant2.




Figure 2: Topological Model for VPLS Model
connecting 2 Tenants with 3 sites each

In this model, TRILL sites are connected to VPLS-capable PE devices that provide a logical interconnect, such that TRILL R Bridges belonging to a specific tenant connected via an single bridged Ethernet. These devices are the same as PE devices specified in [RFC4026]. The Attachment Circuit ports of PE Routers are layer 2 switch ports that are connected to the R Bridges at a TRILL site. Here each VPLS instance looks like an emulated LAN. This model is similar to connecting different R Bridges (TRILL sites) by a layer 2 bridge domain (multi access link) as specified in [RFC6325]. This model doesn't requires any changes in PE routers to carry TRILL packets, as TRILL packets will be transferred transparently.

3.1. Entities in the VPLS Model

The PE (VPLS-PE) and CE devices are defined in [RFC4026].

The Generic L2VPN Transport Functional Components like Attachment Circuits, Pseudowires, VSI etc. are defined in [RFC4664].

The RB (R Bridge) and TRILL Sites are defined in [RFC6325] as updated by [RFC7780].

3.3. TRILL Adjacency for VPLS model

As specified in [section 3](#) of this document, the MPLS cloud looks like an emulated LAN (also called multi-access link or broadcast link). This results in R Bridges at different sites looking like they are connected by a multi-access link. With such interconnection, the TRILL adjacency over the link are automatically discovered and established through TRILL IS-IS control messages [[RFC7177](#)]. These IS-IS control messages are transparently forwarded by the VPLS domain, after doing MPLS encapsulation specified in the [section 3.4](#).

3.4. MPLS encapsulation for VPLS model

Use of VPLS [[RFC4762](#)] to interconnect TRILL sites requires no changes to a VPLS implementation, in particular the use of Ethernet pseudowires between VPLS PEs. A VPLS PE receives normal Ethernet frames from an R Bridge (i.e., CE) and is not aware that the CE is an R Bridge device. As a result, an MPLS-encapsulated TRILL packet within the MPLS network will use the format illustrated in [Appendix A of \[RFC7173\]](#).

3.5. Loop Free provider PSN/MPLS.

No explicit handling is required to avoid loop free topology. Split Horizon technique specified in [[RFC4664](#)] will take care of avoiding loops in the provider PSN network.

3.6. Frame processing.

The PE device transparently process the TRILL control and data frames. Procedures to forward the frames are defined in [[RFC4664](#)]

4. VPTS Model

The (Virtual Private TRILL Service) VPTS is a L2 TRILL service, that emulates TRILL service across a Wide Area Network (WAN). VPTS is similar to what VPLS does for bridge domain. VPLS provides "Virtual Private LAN Service" for different customers. VPTS provides "Virtual Private TRILL Service" for different TRILL tenants.

Figure 3 shows the topological model of TRILL over MPLS using VPTS. In this model the PE routers are replaced with TIR (TRILL Intermediate Router) and VSI is replaced with VTSD (Virtual TRILL Switch Domain). The TIR devices must be capable of supporting both MPLS and TRILL and is defined in [section 4.1.1](#). The TIR devices are interconnected via PWs and

appear as a unified emulated TRILL campus with each VTSD inside a TIR equivalent to a RBridge.

Some of the reasons for interconnecting TRILL Sites without isolating TRILL Control plane of one TRILL site with other is as described below.

- 1) Nickname Uniqueness: One of the basic requirements of TRILL is that, RBridge Nicknames should be unique with in the campus. If we segregate control plane of one TRILL site from other TRILL site and provide interconnection between these sites, it may result in Nickname collision.
- 2) Distribution Tree and its pruning: When a TRILL Data packet traverse a Distribution Tree, it will adhere to it even in other TRILL site. If no end-station service is enabled for a particular Data Label in a TRILL site, Distribution Tree may be pruned and TRILL data packets of that particular Data Label might never get to other TRILL sites. RPF check will always be performed on the packets that are received TIR's through pseudowires.
- 3) Hop Count values: When a TRILL data packet is received over a pseudowire by a TIR, the TIR does the processing of Hop Count defined in [\[RFC6325\]](#) and will not perform any resetting of Hop Count.

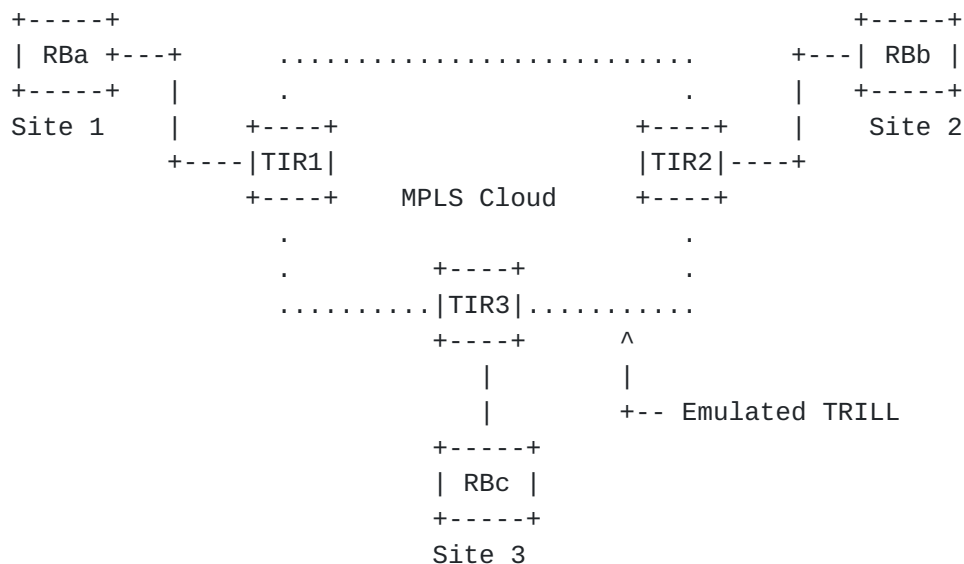


Figure 3: Topological Model of VPTS/TIR connecting three TRILL Sites

In the above figure (Figure 3) Site1, Site2 and Site3 (running TRILL protocol) are connected to TIR Devices. These TIR devices along with the MPLS cloud looks like an unified emulated TRILL campus. Only the PE devices in the MPLS network should be replaced with TIRs so the intermediate P routers are agnostic to the TRILL protocol.

Figure 4 below extends the topological model of TRILL over MPLS to connect multiple TRILL sites belonging to a tenant (tenant here is a campus, not a Data label) using VPTS model. VTSD1 and VTSD2 are two Virtual TRILL Switch Domains (Virtual RBridges) that segregate Tenant1's traffic from Tenant2's traffic. VTSD1 will maintain its own TRILL database for Tenant1, similarly VTSD2 will maintain its own TRILL database for Tenant2.

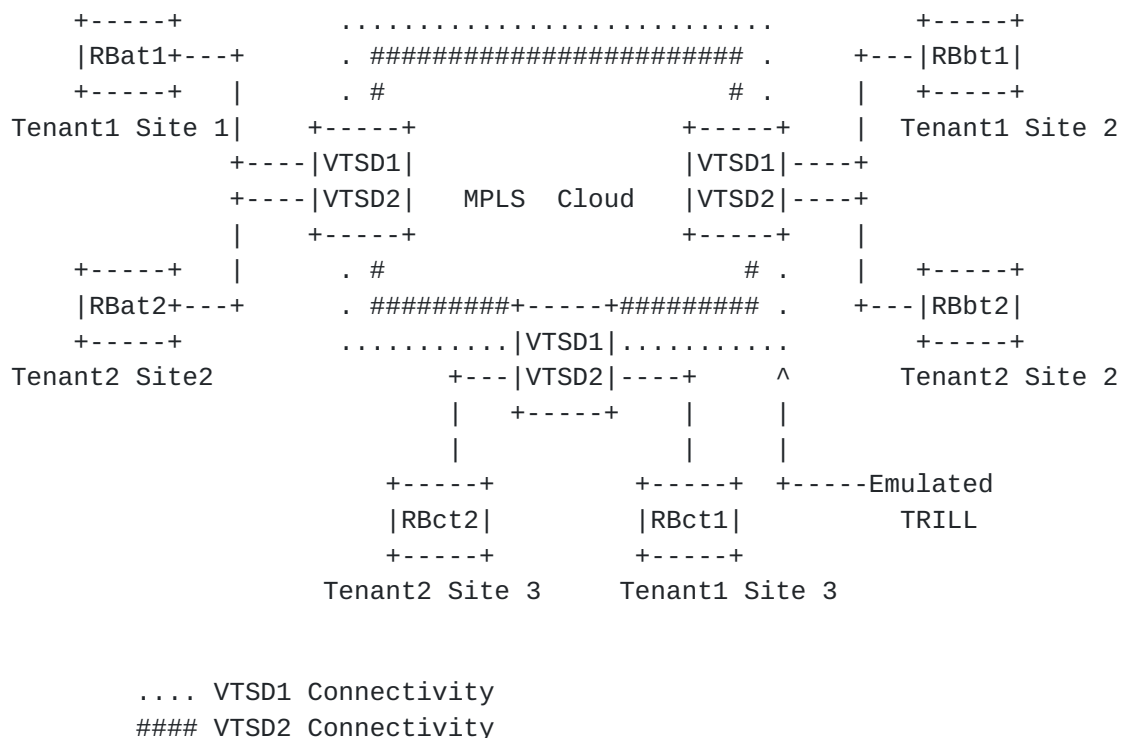


Figure 4: Topological Model of VPTS/TIR
connecting 2 tenants with three TRILL Sites

4.1. Entities in the VPTS Model

The CE devices are defined in [[RFC4026](#)].

The Generic L2VPN Transport Functional Components like Attachment Circuits, Pseudowires etc. are defined in [[RFC4664](#)].

The RB (RBridge) and TRILL Campus are defined in [[RFC6325](#)] as updated by [[RFC7780](#)].

This model introduces two new entities called TIR and VTSD.

4.1.1. TRILL Intermediate Routers (TIR)

The TIRs [TRILL Intermediate Routers] must be capable of running both VPLS and TRILL protocols. TIR devices are a superset of the VPLS-PE devices defined in [[RFC4026](#)] with the additional functionality of TRILL. The VSI instance that provides transparent bridging functionality in the PE device is replaced with VTSD in TIR.

4.1.2. Virtual TRILL Switch/Service Domain (VTSD)

The VTSD [Virtual Trill Switch Domain] is similar to VSI (layer 2 bridge) in VPLS model, but this acts as a TRILL RBridge. The VTSD is a superset of VSI and must support all the functionality provided by the VSI as defined in [[RFC4026](#)]. Along with VSI functionality, the VTSD must be capable of supporting TRILL protocols and form TRILL adjacencies. The VTSD must be capable of performing all the operations that a standard TRILL Switch can do.

One VTSD instance per tenant must be maintained, when multiple tenants are connected to a TIR. The VTSD must maintain all the information maintained by the RBridge on a per tenant basis. The VTSD must also take care of segregating one tenant traffic from other. Each VTSD should have its own nickname, If a TIR supports 10 TRILL tenants, it needs to be assigned with ten TRILL nicknames and run ten copies of TRILL protocols, one for each.

4.2. TRILL Adjacency for VPTS model

The VTSD must be capable of forming TRILL adjacency with other VTSDs present in its peer VPTS neighbor, and also the neighbor R Bridges present in the TRILL sites. The procedure to form TRILL Adjacency is specified in [[RFC7173](#)] and [[RFC7177](#)].

4.3. MPLS encapsulation for VPTS model

The VPTS model uses PPP pseudowires for MPLS encapsulation as specified in [[RFC7173](#)], and requires no changes in the packet format in that RFC.

4.4. Loop Free provider PSN/MPLS.

This model isn't required to employ Split Horizon mechanism in the provider PSN network, as TRILL takes care of Loop free topology using Distribution Trees. Any multi-destination packet will traverse a distribution tree path. All distribution trees are calculated based on TRILL base protocol standard [[RFC6325](#)] as updated by [[RFC7780](#)].

4.5. Frame processing.

This section specifies multi-destination and unicast frame processing in VPTS/TIR model.

4.5.1. Multi-Destination Frame processing

Any multi-destination (unknown unicast, multicast or broadcast) packets inside VTSD should be processed or forwarded through any one of the distribution trees. If any multi-destination packet is received from the wrong pseudowire at a VTSD, the TRILL protocol running in VTSD will perform a RPF check as specified in [[RFC7780](#)] and drops the packet.

Pruning mechanism in Distribution Trees, as specified in [[RFC6325](#)] and [[RFC7780](#)] can also be used to avoid forwarding of multi-destination data packets on the branches where there are no potential destinations.

4.5.2. Unicast Frame processing

Unicast packets must be forwarded in same way they get forwarded in a standard TRILL Campus as specified in [[RFC6325](#)]. If multiple equal cost paths are available over pseudowires to reach destination, then VTSD should be capable of doing ECMP for them.

5. Extensions to TRILL Over Pseudowires [[RFC7173](#)]

The [[RFC7173](#)] mentions how to interconnect a pair of Transparent Interconnection of Lots of Links (TRILL) switch ports using pseudowires. This document explains, how to connect multiple TRILL sites (not limited to only two sites) using the mechanisms and encapsulations defined in [[RFC7173](#)].

6. VPTS Model Versus VPLS Model

VPLS Model uses a simpler loop breaking rule: the "split horizon" rule, where a PE must not forward traffic from one PW to another in the same VPLS mesh.

An issue with the above rule is that if a pseudowire between PEs fails, frames will not get forwarded between the PEs where pseudowire went down.

VPTS solves this problem, as the VPTS Model uses distribution Trees for loop free topology. Hence the frames reach all TIRs even when any one of the pseudowires fails in the mesh.

If equal cost paths are available to reach a site over pseudowires, VPTS Model can use ECMP for processing of frames over pseudowires.

7. Packet processing between pseudowires

Whenever a packet gets received over a pseudowire, VTSD will decapsulate the MPLS headers followed by checking the TRILL header. If the egress nickname in the TRILL header is for a TRILL site located beyond another pseudowire, then VTSD will encapsulate with new MPLS headers and send it across the proper pseudowire.

For example in figure 3, consider that the pseudowire between TIR1 and TIR2 fails, Then TIR1 will communicate with TIR2 via TIR3, whenever packets which are destined to TIR3 gets received from pseudowire between TIR1 and TIR3, VTSD inside TIR3 will decapsulate the MPLS headers, then check the TRILL header's egress nickname field. If egress nickname is destined for the RBridge in site3 then the packet will be sent to RBc, if egress nickname is located at site2, VTSD will add MPLS headers for the pseudowire between TIR3 and TIR2 and forward the packet on that pseudowire.

8. Security Considerations

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

For transport of TRILL by Pseudowires security consideration, see [[RFC7173](#)].

For general VPLS security considerations, see [[RFC4762](#)]

Since the VPTS Model uses Distribution trees for processing of multi-destination data packets, it is always advisable to have at least one Distribution tree root to be located in every TRILL site. This will avoid data packets getting received at TRILL sites where end-station service is not enabled for that data packet.

9. IANA Considerations

This document requires no IANA actions. RFC Editor: Please delete this section before publication

10. References

10.1. Normative References

- [IS-IS] "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, 2002".

- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", [RFC 6325](#), DOI 10.17487/RFC6325, July 2011, <<http://www.rfc-editor.org/info/rfc6325>>.

- [RFC4762] Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), DOI 10.17487/RFC4762, January 2007, <<http://www.rfc-editor.org/info/rfc4762>>.

- [RFC7173] Yong, L., Eastlake 3rd, D., Aldrin, S., and J. Hudson, "Transparent Interconnection of Lots of Links (TRILL) Transport Using Pseudowires", [RFC 7173](#), DOI 10.17487/RFC7173, May 2014, <<http://www.rfc->

[editor.org/info/rfc7173](http://www.rfc-editor.org/info/rfc7173)>.

- [RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", [RFC 7177](#), DOI 10.17487/RFC7177, May 2014, <<http://www.rfc-editor.org/info/rfc7177>>.
- [RFC7780] Eastlake 3rd, D., Zhang, M., Perlman, R., Banerjee, A., Ghanwani, A., and S. Gupta, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", [RFC 7780](#), DOI 10.17487/RFC7780, February 2016, <<http://www.rfc-editor.org/info/rfc7780>>.

10.2. Informative References

- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", [RFC 4026](#), DOI 10.17487/RFC4026, March 2005, <<http://www.rfc-editor.org/info/rfc4026>>.
- [RFC4664] Andersson, L., Ed., and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", [RFC 4664](#), DOI 10.17487/RFC4664, September 2006, <<http://www.rfc-editor.org/info/rfc4664>>.
- [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](#), DOI 10.17487/RFC7172, May 2014, <<http://www.rfc-editor.org/info/rfc7172>>.

Acknowledgements

The contributions of Andrew G. Malis is gratefully acknowledged in improving the quality of this document.

Authors' Addresses

Mohammed Umair
IPInfusion
RMZ Centennial
Mahadevapura Post
Bangalore - 560048 India

E-Mail: mohammed.umair2@gmail.com

Kingston Smiler Selvaraj
IPInfusion
RMZ Centennial
Mahadevapura Post
Bangalore - 560048 India

E-Mail: kingstonsmiler@gmail.com

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

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

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

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

