

Internet Engineering Task Force (IETF)
Request for Comments: 8385
Category: Informational
ISSN: 2070-1721

M. Umair
Cisco
S. Kingston Smiler
PALC Networks
D. Eastlake 3rd
Huawei
L. Yong
Independent
June 2018

Transparent Interconnection of Lots of Links (TRILL)
Transparent Transport over MPLS

Abstract

This document specifies methods to interconnect multiple TRILL (Transparent Interconnection of Lots of Links) sites with an intervening MPLS network using existing TRILL and VPLS (Virtual Private LAN Service) standards. This document addresses two problems: 1) providing connection between more than two TRILL sites that are separated by an MPLS provider network and 2) providing a single logical virtualized TRILL network for different tenants that are separated by an MPLS provider network.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8385>.

Copyright Notice

Copyright (c) 2018 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 (<https://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	3
1.1. Terminology	3
2. TRILL-over-MPLS Model	5
3. VPLS Model	5
3.1. Entities in the VPLS Model	6
3.2. TRILL Adjacency for VPLS Model	7
3.3. MPLS Encapsulation for VPLS Model	7
3.4. Loop-Free Provider PSN/MPLS	7
3.5. Frame Processing	7
4. VPTS Model	7
4.1. Entities in the VPTS Model	9
4.1.1. TRILL Intermediate Router (TIR)	10
4.1.2. Virtual TRILL Switch/Service Domain (VTSD)	10
4.2. TRILL Adjacency for VPTS Model	10
4.3. MPLS Encapsulation for VPTS Model	10
4.4. Loop-Free Provider PSN/MPLS	11
4.5. Frame Processing	11
4.5.1. Multi-destination Frame Processing	11
4.5.2. Unicast Frame Processing	11
5. VPTS Model versus VPLS Model	11
6. Packet Processing between Pseudowires	12
7. Efficiency Considerations	12
8. Security Considerations	12
9. IANA Considerations	13
10. References	13
10.1. Normative References	13
10.2. Informative References	14
Acknowledgements	15
Authors' Addresses	16

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 document, in conjunction with [RFC7173] on TRILL transport using pseudowires, addresses two problems:

- 1) providing connection between more than two TRILL sites that belong to a single TRILL network and are separated by an MPLS provider network using [RFC7173]. (Herein, this is 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, this is for providing connection between TRILL sites belonging to a tenant/tenants over a MPLS provider network. (Herein, this is 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 document also addresses the problem of multi-tenancy by isolating one tenant's traffic from the other.

[RFC7173] mentions how to interconnect a pair of 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].

1.1. Terminology

Acronyms and terms used in this document include the following:

AC - Attachment Circuit [RFC4664]
Data Label - VLAN Label or Fine-Grained Label

database	- IS-IS link state database
ECMP	- Equal-Cost Multipath
FGL	- Fine-Grained Labeling [RFC7172]
IS-IS	- Intermediate System to Intermediate System [IS-IS]
LAN	- Local Area Network
MPLS	- Multiprotocol Label Switching
PBB	- Provider Backbone Bridging
PE	- Provider Edge device
PSN	- Packet Switched Network
PW	- Pseudowire [RFC4664]
TIR	- TRILL Intermediate Router (Device that has both IP/MPLS and TRILL functionality)
TRILL	- Transparent Interconnection of Lots of Links OR Tunneled 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 in two different ways:

- a) the VPLS Model for TRILL b) the VPTS Model / TIR Model for TRILL

Both these models can be used to solve problem statements 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 the VPLS model. The PE routers in the below topology model should support all the functional components mentioned in [RFC4664].

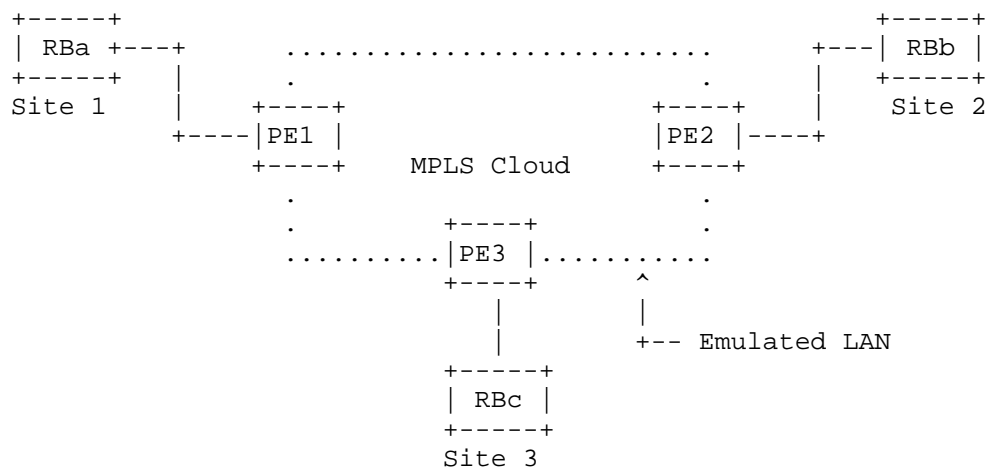


Figure 1: Topological Model of TRILL over MPLS
Connecting 3 TRILL Sites

Figure 2 below shows the topological model of TRILL over MPLS to connect multiple TRILL sites belonging to a tenant. ("Tenant" here is a TRILL campus, not a specific 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.

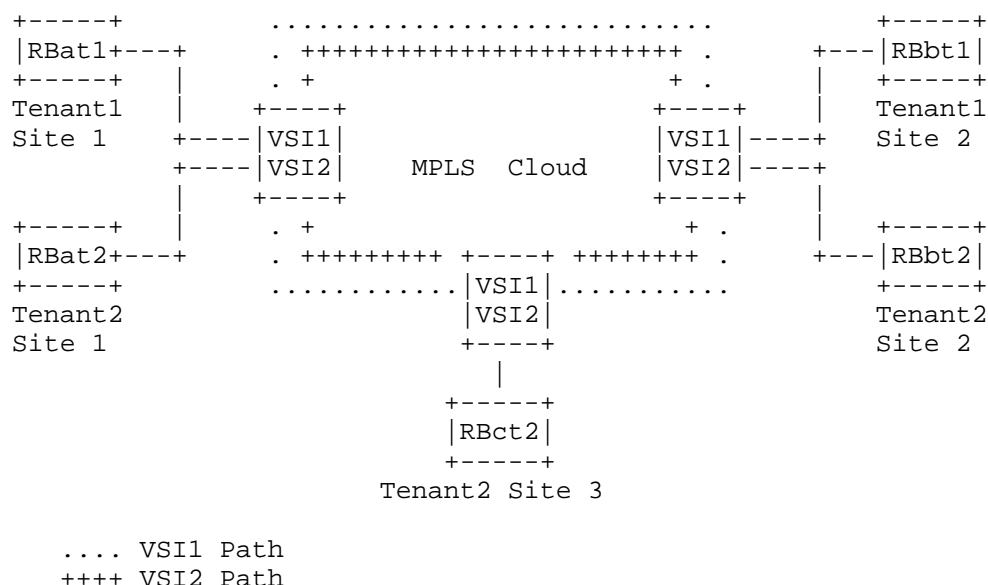


Figure 2: Topological Model for VPLS Model
Connecting 2 Tenants with 3 Sites Each

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 are connected via a single bridged Ethernet. These PE devices are the same as the 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 by a Layer 2 bridge domain (multi-access link) as specified in [RFC6325]. This model doesn't require 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 Customer Edge (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 campus are defined in [RFC6325] as updated by [RFC7780].

3.2. TRILL Adjacency for VPLS Model

As specified in Section 3, the MPLS cloud looks like an emulated LAN (also called multi-access link or broadcast link). This results in RBridges at different sites looking like they are connected by a multi-access link. With such interconnection, the TRILL adjacencies 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 as specified in Section 3.3.

3.3. MPLS Encapsulation for VPLS Model

Use of VPLS [RFC4762] [RFC4761] 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 RBridge (i.e., CE) and is not aware that the CE is an RBridge device. As an example, an MPLS-encapsulated TRILL packet within the MPLS network can use the format illustrated in Appendix A of [RFC7173] for the non-PBB case. For the PBB case, additional header fields illustrated in [RFC7041] can be added by the entry PE and removed by the exit PE.

3.4. Loop-Free Provider PSN/MPLS

No explicit handling is required to avoid a loop-free topology. The "split horizon" technique specified in [RFC4664] will take care of avoiding loops in the provider PSN network.

3.5. Frame Processing

The PE devices 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 Layer 2 TRILL service that emulates TRILL service across a Wide Area Network (WAN). VPTS is similar to what VPLS does for a bridged core but provides a TRILL core. 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 TRILL Intermediate Routers (TIRs), and the VSIs are replaced with Virtual TRILL Switch Domains (VTSDs). The TIR devices must be capable of supporting both

MPLS and TRILL as specified 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 an RBridge.

Below are some of the reasons for interconnecting TRILL sites without isolating the TRILL control plane of one TRILL site from other sites.

- 1) Nickname uniqueness: One of the basic requirements of TRILL is that RBridge nicknames are unique within the campus [RFC6325]. If we segregate the control plane of one TRILL site from other TRILL sites and provide interconnection between these sites, it may result in nickname collision.
- 2) Distribution trees and their pruning: When a TRILL Data packet traverses a Distribution Tree, it will stay on it even in other TRILL sites. If no end-station service is enabled for a particular Data Label in a TRILL site, the distribution tree may be pruned and TRILL data packets of that particular Data Label might never get to another TRILL site where the packets had no receivers. The TRILL Reverse Path Forwarding (RPF) check will always be performed on the packets that are received by TIRs 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 3 TRILL Sites

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 the VPTS model. VTSD1 and VTSD2 are two Virtual TRILL Switch Domains (Virtual R Bridges) 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.



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

This model introduces two new entities, TIR and VTSD, which are described below.

4.1.1. TRILL Intermediate Router (TIR)

The TIRs 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 that provides transparent bridging functionality in the PE device is replaced with VTSD in a TIR.

4.1.2. Virtual TRILL Switch/Service Domain (VTSD)

The VTSD is similar to the VSI (Layer 2 bridge) in the VPLS model, but the VTSD acts as a TRILL RBridge. The VTSD is a superset of the 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 forming 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 kept by the RBridge on a per-tenant basis. The VTSD must also take care of segregating one tenant's traffic from another's. Each VTSD will have its own nickname for each tenant. If a TIR supports 10 TRILL tenants, it needs to be assigned with 10 TRILL nicknames, one for the nickname space of each of its tenants, and run 10 copies of TRILL protocols, one for each tenant. It is possible that it would have the same nickname for two or more tenants, but, since the TRILL data and control traffic are separated for the tenants, there is no confusion.

4.2. TRILL Adjacency for VPTS Model

The VTSD must be capable of forming a TRILL adjacency with the corresponding VTSDs present in its peer VPTS neighbor and also with the neighboring RBridges of 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 or Ethernet pseudowires for MPLS encapsulation as specified in [RFC7173] and requires no changes in the packet format in that RFC. In accordance with [RFC7173], the PPP encapsulation is the default.

4.4. Loop-Free Provider PSN/MPLS

This model isn't required to employ the "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 the TRILL base protocol standard [RFC6325] as updated by [RFC7780].

4.5. Frame Processing

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

4.5.1. Multi-destination Frame Processing

Any multi-destination (unknown unicast, multicast, or broadcast, as indicated by the multi-destination bit in the TRILL header) packets inside a VTSD will be processed or forwarded through the distribution tree for which they were encapsulated on TRILL ingress. If any multi-destination packet is received from the wrong pseudowire at a VTSD, the TRILL protocol running in the VTSD will perform an RPF check as specified in [RFC7780] and drop the packet.

The 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 are forwarded in the 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 the destination, then VTSD should be capable of doing ECMP for those equal-cost paths.

5. VPTS Model versus VPLS Model

The 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. In contrast, the VPTS model uses distribution trees for loop-free topology. As this is an emulated TRILL service, for interoperability purposes, the VPTS model is the default.

6. Packet Processing between Pseudowires

Whenever a packet gets received over a pseudowire, a VTSD will decapsulate the MPLS headers then check the TRILL header. If the egress nickname in the TRILL header is for a TRILL site located beyond another pseudowire, then the VTSD will encapsulate the packet 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 that are destined to TIR3 are received from the pseudowire between TIR1 and TIR3, the VTSD inside TIR3 will decapsulate the MPLS headers, then check the TRILL header's egress nickname field. If the egress nickname indicates it is destined for the RBridge in Site 3, then the packet will be sent to RBC; if the egress nickname is located at Site 2, VTSD will add MPLS headers for the pseudowire between TIR3 and TIR2 and forward the packet on that pseudowire.

7. Efficiency Considerations

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 located in every TRILL site. This will prevent data packets from being received at TRILL sites where end-station service is not enabled for that data packet.

8. Security Considerations

This document specifies methods using existing standards and facilities in ways that do not create new security problems.

For general VPLS security considerations, including discussion of isolating customers from each other, see [RFC4761] and [RFC4762].

For security considerations for transport of TRILL by pseudowires, see [RFC7173]. In particular, since pseudowires are supported by MPLS or IP, which are in turn supported by a link layer, that document recommends using IP security, such as IPsec [RFC4301] or DTLS [RFC6347], or the lower link-layer security, such as MACSEC [802.1AE] for Ethernet links.

Transmission outside the customer environment through the provider environment, as described in this document, increases risk of compromise or injection of false data through failure of tenant isolation or by the provider. In the VPLS model (Section 3), the use of link encryption and authentication between the CEs of a tenant that is being connected through provider facilities should be a good

defense. In the VPTS model (Section 4), it is assumed that the CEs will peer with virtual TRILL switches of the provider network, and thus link security between TRILL switch ports is inadequate as it will terminate at the edge PE. Thus, encryption and authentication from end station to end station and authentication are more appropriate for the VPTS model.

For added security against the compromise of data, end-to-end encryption and authentication should be considered; that is, encryption and authentication from source end station to destination end station. This would typically be provided by IPsec [RFC4301] or DTLS [RFC6347] or other protocols convenient to protect the information of concern.

For general TRILL security considerations, see [RFC6325].

9. IANA Considerations

This document has no IANA actions.

10. References

10.1. Normative References

- [IS-IS] ISO, "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.
- [RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [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, <<https://www.rfc-editor.org/info/rfc4762>>.
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBriges): Base Protocol Specification", RFC 6325, DOI 10.17487/RFC6325, July 2011, <<https://www.rfc-editor.org/info/rfc6325>>.

- [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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc7780>>.

10.2. Informative References

- [802.1AE] IEEE, "IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Security", IEEE Std 802.1AE, DOI 10.1109/IEEESTD.2006.245590.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, DOI 10.17487/RFC4026, March 2005, <<https://www.rfc-editor.org/info/rfc4026>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4664] Andersson, L., Ed., and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <<https://www.rfc-editor.org/info/rfc4664>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC7041] Balus, F., Ed., Sajassi, A., Ed., and N. Bitar, Ed., "Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging", RFC 7041, DOI 10.17487/RFC7041, November 2013, <<https://www.rfc-editor.org/info/rfc7041>>.

[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, <<https://www.rfc-editor.org/info/rfc7172>>.

Acknowledgements

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

Authors' Addresses

Mohammed Umair
Cisco Systems
SEZ, Cessna Business Park
Sarjapur - Marathahalli Outer Ring road
Bengaluru - 560103
India

Email: mohammed.umair2@gmail.com

S. Kingston Smiler
PALC NETWORKS PVT LTD
Envision Technology Center
#119, 1st Floor, Road No.3
EPIP Area Phase 1, Whitefield
Near Vydehi Hospital
Bengaluru - 560066, Karnataka
India

Email: kingstonsmiler@gmail.com

Donald Eastlake 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757
United States of America

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

Lucy Yong
Independent

Phone: +1-469-227-5837
Email: lucyyong@gmail.com

