

INTERNET-DRAFT  
Intended Status: Informational

Mohammed Umair  
Kingston Smiler S  
Shaji Ravindranathan  
IPInfusion  
Lucy Yong  
Donald Eastlake 3rd  
Huawei Technologies  
February 17, 2016

Expires: August 20, 2016

**Date Center Interconnect using TRILL**  
**<[draft-muks-trill-dci-02.txt](#)>**

**Abstract**

This document describes a TRILL based Data Center Interconnect (DCI) solution. TRILL is predominantly used inside a data center for providing intra-data center switching optimally by utilizing multiple links. This draft described a way to use TRILL to extend a network across different data center via MPLS service provider network using Virtual TRILL Service/Switch Domain (VTSD) as described in draft [[VTSD](#)].

**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

<a href="#">1</a>	Introduction . . . . .	<a href="#">4</a>
<a href="#">1.1</a>	Terminology . . . . .	<a href="#">5</a>
<a href="#">2</a>	Date Center Topology . . . . .	<a href="#">7</a>
<a href="#">3</a>	Requirements of DCI Protocol . . . . .	<a href="#">7</a>
<a href="#">3.1</a>	Multi-homing with all-active forwarding . . . . .	<a href="#">8</a>
<a href="#">3.2</a>	Effectively scaling the bandwidth by adding more links . . . . .	<a href="#">8</a>
<a href="#">3.3</a>	Auto-discovery of services . . . . .	<a href="#">9</a>
<a href="#">3.4</a>	Delivering Layer 2 and Layer 3 services over the same interface . . . . .	<a href="#">9</a>
<a href="#">3.5</a>	BUM traffic optimization . . . . .	<a href="#">9</a>
<a href="#">3.6</a>	Control plane learning of MAC . . . . .	<a href="#">10</a>
<a href="#">3.7</a>	Virtualisation and isolation of different instances . . . . .	<a href="#">10</a>
<a href="#">3.8</a>	MAC mass-withdrawal . . . . .	<a href="#">10</a>
<a href="#">3.9</a>	Significantly larger Name-Space in the Overlay (16M segments) . . . . .	<a href="#">10</a>
<a href="#">3.10</a>	Extensive OAM Capabilities . . . . .	<a href="#">11</a>
<a href="#">3.11</a>	Supporting Ring topology in the Core Network . . . . .	<a href="#">11</a>
<a href="#">4</a>	TRILL DCI Operations . . . . .	<a href="#">11</a>
<a href="#">4.1</a>	TRILL data center . . . . .	<a href="#">12</a>
<a href="#">4.2</a>	Layer2 data center . . . . .	<a href="#">12</a>
<a href="#">4.2.1</a>	TRILL Access load-balancing . . . . .	<a href="#">13</a>
<a href="#">4.2.1.1</a>	Appointed Forwarder Mechanism . . . . .	<a href="#">13</a>
<a href="#">4.2.1.2</a>	TRILL Active-Active Access . . . . .	<a href="#">13</a>
<a href="#">4.2.2</a>	TRILL Pseudowire load-balancing . . . . .	<a href="#">13</a>
<a href="#">5</a>	MPLS encapsulation and Loop free provider PSN/MPLS . . . . .	<a href="#">14</a>
<a href="#">6</a>	Frame Processing . . . . .	<a href="#">15</a>
<a href="#">6.1</a>	Frame processing between data center T2 switch and TIR. . . . .	<a href="#">15</a>
<a href="#">6.2</a>	Frame processing between TIR's . . . . .	<a href="#">15</a>
<a href="#">7</a>	Security Considerations . . . . .	<a href="#">17</a>
<a href="#">8</a>	IANA Considerations . . . . .	<a href="#">17</a>
<a href="#">9</a>	References . . . . .	<a href="#">17</a>
<a href="#">9.1</a>	Normative References . . . . .	<a href="#">17</a>



<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">19</a>
	Authors' Addresses . . . . .	<a href="#">20</a>

## **1 Introduction**

The IETF Transparent Interconnection of Lots of Links (TRILL) protocol [[RFC6325](#)] [[RFC7177](#)] [[RFC7180bis](#)] 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. Devices implementing TRILL are called Routing Bridges(RBridges)or TRILL Switches.

TRILL is predominantly used inside data centers for providing intra-data center switching optimally by utilizing multiple links. Though TRILL usually runs within a data center, there is a need to interconnect various data center sites to provide connectivity across data centers. This draft describes a solution to this problem by using VTSD (Virtual TRILL Switch Domain) as described in the draft [[VTSD](#)].

The draft [[VTSD](#)] introduces a new terminology called VTSD (Virtual TRILL Switch Domain) and VPTS (Virtual Private TRILL Service).

The (Virtual Private TRILL Service) VPTS is a L2 TRILL service, that emulates TRILL service across a Wide Area Network (WAN) over MPLS PWE3. VPTS is similar to what VPLS does for bridge domain.

The VTSD [Virtual Trill Switch Domain] is similar to VSI (layer 2 bridge) in VPLS model, but this acts as a TRILL RBridge. 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 adjacency. The VTSD must be capable of performing all the operations that a standard TRILL Switch can do.

Pseudo Wire Emulation Edge-to-Edge (PWE3) is a mechanism that emulates the essential attributes of a service such as Ethernet over a Packet Switched Network (PSN). The required functions of PWS include encapsulating service-specific PDUs arriving at an ingress port, and carrying them across a path or tunnel, managing their timing and order, and any other operations required to emulate the behavior and characteristics of the service as faithfully as possible.

The PEs may be connected by an MPLS Label Switched Path (LSP) infrastructure, which provides the benefits of MPLS technology. The PEs may also be connected by an IP network, in which case IP/GRE (Generic Routing Encapsulation) tunneling or other IP tunneling can



be used to provide PSN functionality. The detailed procedures in this document are specified only for MPLS LSPs as PSN. However, these procedures are designed to be extensible to use IP tunnels as PSN, which is outside the scope of this document.

## **1.1 Terminology**

Acronyms used in this document include the following:

AC	- Attachment Circuit [ <a href="#">RFC4664</a> ]
Access Port	- A TRILL switch port configured with the "end station service enable" bit on, as described in <a href="#">Section 4.9.1 of [RFC6325]</a> . All AC's, VTSD ports connected to CE's, should configured as TRILL Access port.
AF	- Appointed Forwarder [ <a href="#">RFC6325</a> ] and [ <a href="#">RFC6439bis</a> ].
BUM	- Broadcast, Unknown destination Unicast and Multicast
Data Label	- VLAN or FGL
DCI	- Data Center Interconnect
ECMP	- Equal Cost Multi Pathing
FGL	- Fine-Grained Labeling [ <a href="#">RFC7172</a> ]
IS-IS	- Intermediate System to Intermediate System [ <a href="#">IS-IS</a> ]
L2	- Layer 2
LAN	- Local Area Network
Link	- The means by which adjacent TRILL switches or VTSD is connected. May be a bridged LAN.
MLAG	- Multi-Chassis Link Aggregation





MPLS	- Multi-Protocol Label Switching
PE	- Provider Edge Device
PSN	- Packet Switched Network
PW	- Pseudowire [ <a href="#">RFC4664</a> ]
RBridge	- An alternative name for TRILL Switch
TIR	- TRILL Intermediate Router (Devices where Pseudowire starts and Terminates)
TRILL	- Transparent Interconnection of Lots of Links OR Tunneled Routing in the Link Layer [ <a href="#">RFC6325</a> ]
TRILL Site	- A part of a TRILL campus that contains at least one RBridge.
TRILL switch	- A device implementing the TRILL protocol. An alternative name for an RBridge.
Trunk port	- A TRILL switch port configured with the "end station service disable" bit on, as described in <a href="#">Section 4.9.1 of [RFC6325]</a> . All pseudowires should be configured as TRILL Trunk port.
VLAN	- Virtual Local Area Network
VPLS	- Virtual Private LAN Service
VPTS	- Virtual Private TRILL Service
VSI	- Virtual Service Instance [ <a href="#">RFC4664</a> ]
VTSI	- Virtual TRILL Service Instance
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.



VTSD-AP

- A VTSD TRILL Access port can be a AC or a logical port connected with CE's. it can be a combination of physical port and Data Label.  
OR just Physical port connected to CE's

## 2. Data Center Topology

The reference topology used in this document is a 3 tier traditional topology. Although other topologies may be utilized within the data center, most of such L2 based data centers may be modeled as a 3 tier traditional topology. The reference topology is illustrated in Figure 1. To keep terminologies simple and uniform, in this document these layers will be referred to as the Tier-1, Tier-2 and Tier-3 "tiers", and the switches in these layers will be termed as T1SW, T2SW etc. For simplicity reasons, the entire data center topology will not be mentioned in the further sections. Only the relevant nodes will be shown identified by this nomenclature.




Figure 1: Typical data center network topology

## 3. Requirements of DCI Protocol



This section describes in detail the key requirements of a data center interconnect (DCI) solution.

### **3.1. Multi-homing with all-active forwarding**

One of the key requirements of a DCI layer in data center is to efficiently use all the links in the network by spreading load across them. There are two types of links in the DCI layer. The links that provides connectivity to the data center and the links that provides connectivity to other data center via backbone network. The DCI layer should use both of these link types optimally in an all-active forwarding manner. Typically the links towards the data center will have multiple connectivity towards the peer for providing redundancy in the network. TRILL supports multiple active parallel links between the TRILL R-Bridges / traditional L2 bridges. For providing active load-balance for traffic from layer2 data center TRILL uses the Appointed Forwarder mechanism.

The Appointed forwarder mechanism defined in [[RFC6439bis](#)] provides a way to actively forwarding end station traffic by a RBridge, so that native traffic in each VLAN is handled by at most one RBridge. By default, the DRB (Designated RBridge) on a link is in-charge of native traffic for all VLANs on the link. The DRB may, if it wishes, act as Appointed Forwarder for any VLAN and it may appoint other R Bridges that have ports on the link as Appointed Forwarder for one or more VLANs with any one of the mechanism described in [[RFC6439bis](#)].

A RBridge on a multi-access link forms adjacency [[RFC7177](#)] with other RBridge if the VLAN's configured/enabled between them are common. For example there are four R Bridges attached to multi-access link, say RB1, RB2, RB3 and RB4. RB1 and RB2 are configured with single VLAN "VLAN 2", whereas RB3 and RB4 are configured with "VLAN 3". Assume that there are no Native VLAN's present on any of the R Bridges connected to multi-access link. Since TRILL Hellos are sent with VLAN Tag enabled on the interface, RB3 and RB4 drops the hellos of RB1 and RB2 (since they are not configured for VLAN 2). Similarly RB1 and RB2 drops the Hellos of RB3 and RB4. This results in RB1 and RB2 not forming adjacency with RB3 and RB4. RB1 and RB2 after electing DRB and forming adjacency between them, will decide about VLAN 2 AF. Similarly RB3 and RB4 decide about the VLAN 3 AF.

### **3.2. Effectively scaling the bandwidth by adding more links**

As more and more services are deployed over the cloud, there is a clear requirement for easily scaling the bandwidth in the network without disturbing the existing running services. One of the ways to scale the bandwidth is to add a link (either physical or logical)



across the devices which require higher bandwidth rate. The same requirement is applicable in the DCI layer for interfaces towards the backbone network and towards the data center.

TRILL protocol, by design itself inherently takes care of this by optimally utilizing multiple links. As PWE3 interface, which provides connectivity to different data center is also part of TRILL network, it is possible to dynamically scale the bandwidth in the backbone network / DCI interface by adding more PWE3 to the VTSD instance.

### **3.3. Auto-discovery of services**

This is one of the key requirement of data center so DCI services will be provisioned with minimal configuration effort. Currently the TRILL protocol doesn't have any mechanism to discover peer VTSD / TIR nodes. A new draft should be proposed to address this in TRILL.

### **3.4. Delivering Layer 2 and Layer 3 services over the same interface**

It is desirable to provide a mechanism to advertise both layer 2 and layer 3 forwarding information (Route (IP prefix) in case of L3 and MAC in case of L2) to the peer nodes across the data centers. The control plane way of distributing the forwarding information provides multiple benefits in terms of scale, performance and security. [ARP/ND-Optimization] and [[IA-APPsub-TLV](#)] provides mechanism to distribute both MAC and IP information over TRILL network.

### **3.5. BUM traffic optimization**

A key design consideration in a DCI network is handling BUM (Broadcast, Unknown Unicast and Multicast) traffic. If the BUM packets are handled as in the traditional layer 2 network (by forwarding to all the ports which are part of the same broadcast domain), this will create excessive overhead in the network. TRILL takes care of this issue using the distribution tree and pruning mechanism.

Any unknown unicast, multicast or broadcast frames inside VTSD should be processed or forwarded through any one of the distribution tree's path. If any multi-destination frame is received from the wrong pseudowire at a VTSD, the TRILL protocol running in VTSD should perform a RPF check as specified in [[RFC7180bis](#)] and drops the packet.

Pruning (VLAN (or FGL) and multicast pruning) mechanism of Distribution Tree as specified in [[RFC6325](#)] and [[RFC7180bis](#)] can also be used for forwarding of multi-destination data frames on the branches that are not pruned.





Also the ARP/ND-proxy and control plane MAC address learning mechanism mentioned in Sections [3.4](#) and [3.6](#) will help the VTSD/RBridges in the network to learn the unicast MAC address from ARP/ND packets. This reduces the unknown unicast flow.

### **[3.6.](#) Control plane learning of MAC**

Learning MAC addresses in data plane brings the scaling limitations of the devices to the DCI layer. Hence the protocol that provides DCI should provide control plane learning to overcome the data plane learning limitation. Mac address learning through ESADI (End Station Address Distribution Information Protocol) is described in [[RFC7357](#)] and requires no changes to the protocol. However the following optional extensions can be provided for controlling the MAC learning mechanism.

a) Way to disable remote MAC Addresses learning through data plane and b) Control over the number of MAC Addresses to be advertised to the remote VTSD's.

The mechanism to provide these optional extensions is out side the scope of this document.

### **[3.7.](#) Virtualisation and isolation of different instances**

VTSD provides a way to isolate the TRILL link state databases and the forwarding information across different TRILL sites. As VPTS is similar to VPLS, the mac address and the nickname learnt on a particular VPTS is isolated from other VPTS instance in the system.

### **[3.8.](#) MAC mass-withdrawal**

It is desirable in the data center to have a mechanism to flush a set of MAC addresses from the network, in the event of node failures in the network. This Mass MAC-Address withdrawal may also be applicable when there is any movement in the End-stations. Mass MAC-Address withdrawal is specified in draft [[Address-Flush](#)] and require no changes to the protocol.

### **[3.9.](#) Significantly larger Name-Space in the Overlay (16M segments)**

Layer2 DCI technologies can be used to provide overlay connectivity between Top of Rack switches over an IP underlay. When a DCI protocol is used as an overlay, there is a clear requirement to have a larger namespace to provide services to different customers. TRILL FGL [[RFC7172](#)] provides  $2^{24}$  data labels to isolate different TRILL services.



### **3.10.    Extensive OAM Capabilities**

It is desirable to provide extensive OAM support in the data center network for fault indication, fault localization, performance information and diagnostic functions. TRILL already provides extensive support for OAM capabilities as specified in [[RFC7174](#)] and [[RFC7175](#)]. These mechanisms can be used for fault indication, localization and performance information.

### **3.11.    Supporting Ring topology in the Core Network**

In most cases, the backbone network that provides connectivity to the data center is deployed as a ring topology to provide fault tolerance. It is highly desirable to provide a similar kind of service with the DCI protocol. Most of the existing DCI technologies like VPLS doesn't provide this support due to split horizon rule running inside the backbone network.

In case of VTSD, as TRILL takes care of forming loop-free topology, there is no need to run split horizon in the backbone network. This paves the way for traffic to move from one PW to another PW and eases the formation of service over ring topology in the core, without having any mesh or hub-spoke connections.

## **4.    TRILL DCI Operations**

The below diagram represents a high level overview of TRILL DCI. In the below diagram there are two data centers as DataCenter1 and DataCenter2. DataCenter1 has two core routers (which are also part of the backbone network) as TIR1 and TIR2. Similarly DataCenter2 has only one core router as TIR3. These TIR devices are connected via the backbone network using PSN Tunnels. Pseudowires are configured across these devices.

Operations of VTSD is described in draft [[VTSD](#)]. There are multiple attachment circuit interfaces which are configured from T2SW to TIR1 and TIR2. The T2SW can be part of Layer2 network (Layer2 data center) or TRILL network (TRILL data center).



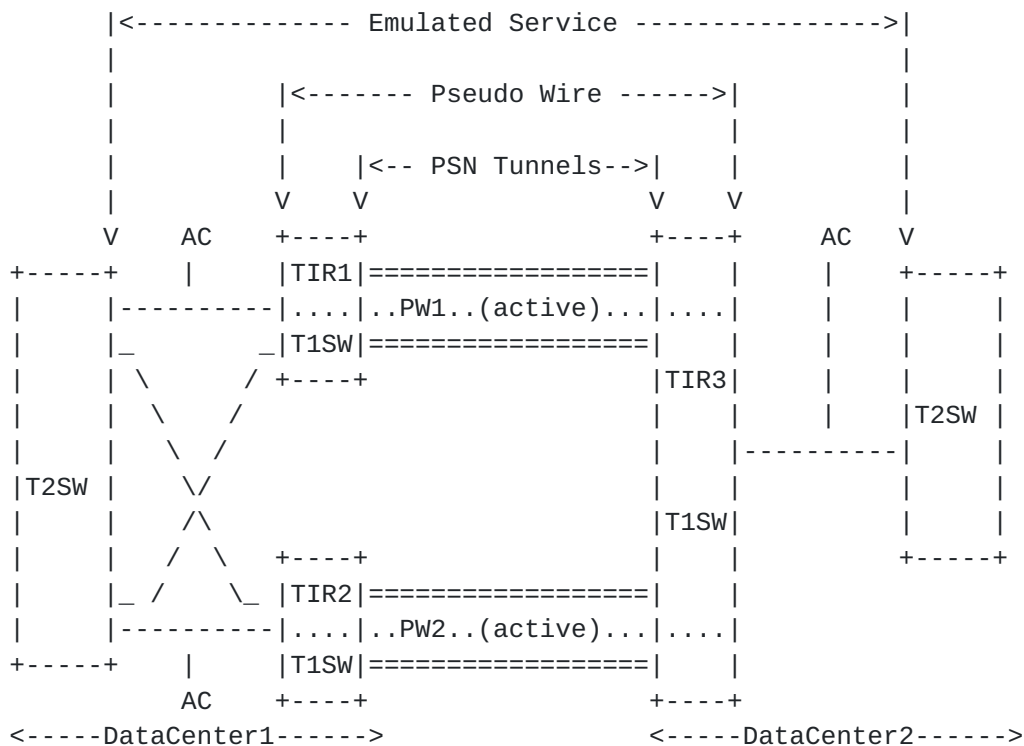


Figure 1: Data Center DCI

#### 4.1. TRILL data center

In this case, the VTSD or TIR will form TRILL adjacencies with other VTSDs present in its peer VPTS neighbor, and also with the RBridges present in the TRILL sites (here it is T2SW). As the entire network runs the TRILL protocol (from data center1 to data center2), TRILL will take care of efficiently using the multiple attachment circuit interfaces and PWE3 interfaces. Load balancing of frames between Tier-2 switch and VTSD will be taken care by the TRILL protocol running inside the RBridges (Tier-2 Switch) and VTSD (Tier-1 Switch) as described in draft [[VTSD](#)].

#### 4.2. Layer2 data center

In case of layer2 data center, as Tier-2 switches are Layer-2 Ethernet switches, an Attachment Circuit should work as a normal TRILL Access port. As the data center is not running the TRILL protocol, the mechanism to provide active load balancing for Attachment Circuits differs from the TRILL data center. The subsequent sections describe in detail the operation of TRILL access load-balancing in a layer2 data center.



#### **4.2.1. TRILL Access load-balancing**

This section describes the mechanism to provide active load balancing across Tier1 and Tier2 switch. There are two ways to provide load balancing.

a) Using the Appointed Forwarder mechanism [RFC6339bis] (load-balancing based on VLAN), and b) Using TRILL Active-Active Access mechanism [[RFC7379](#)] (similar to MC-LAG solution).

##### **4.2.1.1. Appointed Forwarder Mechanism**

The Appointed forwarder mechanism defined in [[RFC6439bis](#)] provides a way for actively forwarding the traffic by a RBridge, with the intent that native traffic in each VLAN be handled by at most one RBridge. By default, the DRB (Designated RBridge) on a link is in-charge of native traffic for all VLANs on the link. The DRB may, if it wishes, act as Appointed Forwarder for any VLAN and it may appoint other RBridges that have ports on the link as Appointed Forwarder for one or more VLANs. The DRB may appoint other RBridges on the link with any one of the mechanism described in [[RFC6439bis](#)]. Based on the type of attachment circuit (port-based or VLAN based), the DRB chooses the appointed forwarder RBridges/VTSDs, which can distribute the traffic based on the VLANs.

##### **4.2.1.1.1. Port-based AC operations.**

In this case, the VTSDs in TIR1 and TIR2 will form TRILL adjacency via AC ports. If the attachment circuit port can carry N number of end-station service VLANs, then TIR1 and TIR2's VTSDs can equally distribute them using AF Mechanism of TRILL.

##### **4.2.1.1.2. VLAN-based AC operations.**

Likewise in Port-based AC, in this case also the VTSDs in TIR1 and TIR2 will form TRILL adjacency via AC ports. Since only one VLAN end-station service is enabled per VTSD, only one TIR's VTSD can become AF for that VLAN. Hence native traffic can be processed by any one of the AC.

##### **4.2.1.2. TRILL Active-Active Access**

TRILL Active-Active Access is specified in [[RFC7379](#)], [[Pseudo-Nickname](#)], [[Centralized-replication](#)], and [[AA-Multi-Attach](#)]. Mechanisms specified in these drafts can be utilized effectively to provide TRILL Active-Active Access.

#### **4.2.2. TRILL Pseudowire load-balancing**





TRILL supports multiple parallel adjacencies between neighbor R Bridges. [Appendix C of \[RFC6325\]](#) and [section 3.5 of \[RFC7177\]](#) describes this in detail. Multipathing across such parallel connections can be done for unicast TRILL Data traffic on a per-flow basis, but is restricted for multi-destination traffic. VTSD should also support this functionality.

TRILL DCI Pseudowires which belong to the same VTSD instance in a TIR and connect to same remote TIR are referred to as parallel pseudowires. These parallel pseudowires corresponds to a single link inside VTSD.

Here all pseudowires should be capable of carrying traffic.

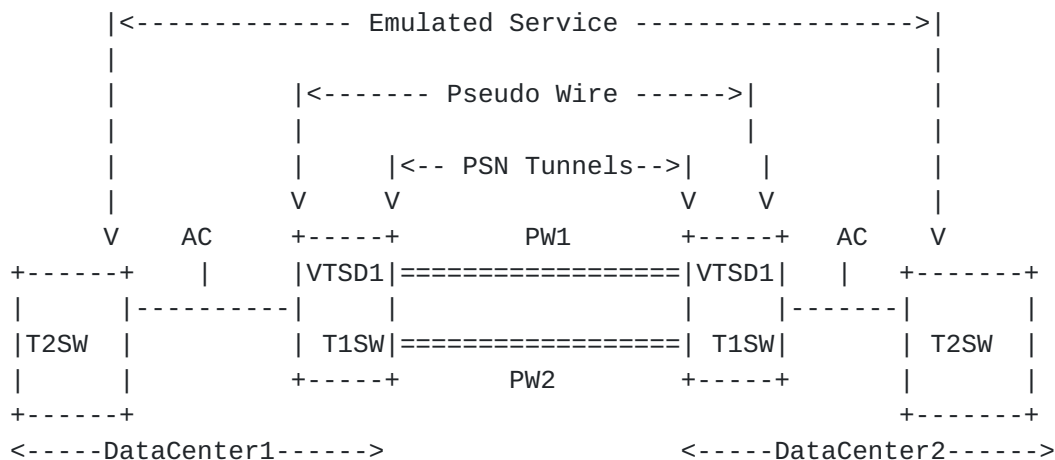


Figure 2: Parallel pseudowires with TRILL DCI

In above Figure 2, PW1 and PW2 are parallel pseudowires, as these pseudowires belongs to same VTSD and provides a connectivity across same TIRs.

This mechanism provides a way for actively increasing and optimally utilizing the bandwidth in the backbone network without affecting the existing traffic.

## 5. MPLS encapsulation and Loop free provider PSN/MPLS

TRILL use of MPLS encapsulation over pseudowire is specified in



[[RFC7173](#)], and requires no changes in the frame format.

TRILL DCI doesn't require a Split Horizon mechanism in the backbone PSN network, as TRILL takes care of Loop free topology using Distribution Trees. Any multi-destination frame will traverse a distribution tree path. All distribution trees are calculated based on TRILL base protocol standard [[RFC6325](#)] as updated by [[RFC7180bis](#)].

## **6. Frame Processing**

This section specifies frame processing from data center T2 switch and TIR's

### **6.1. Frame processing between data center T2 switch and TIR.**

In a multi-homed topology, where in a data center switch (T2SW) is connected to two TIRs, the AF mechanism described in [section 4.2.1.1](#) will be used to decide which TIR/VTSD will carry the traffic for a particular VLAN. This is applicable to the case wherein the data center switch is connected to a PE/TIR device via multiple layer 2 interfaces to increase the bandwidth.

As a frame gets ingressed into a TIR (or any one of the TIR, when the T2SW switches are connected to multiple TIR's) after passing the AF check, the TIR encapsulates the frame with TRILL and MPLS headers and forwards the frame on a pseudowire. If parallel pseudowires are present, the TRILL protocol running in VTSD will select any one of the pseudowires and forward the TRILL Data packet over it. Multi-destination packets will be forwarded on a distribution tree's path [[RFC7180bis](#)]

Even if any of the paths or links fails between T2SW switch and TIR's or between TIR's, frames can be always be forwarded to any of available UP links or paths through other links/pseudowires. This is one of the key advantage provided by TRILL DCI mechanism.

If multiple equal paths are available, TRILL will distribute traffic among all the paths.

Also VTSD doesn't depend on the routing or signaling protocol that is running between TIRs, provided there is a PSN tunnel available with proper encapsulation mechanism.

Any multi-destination frames, when ingressed to TIR's, will traverse one of the distribution trees, with strong RPF Checks. The Hop count field in TRILL Header will avoid loops or duplication of traffic.

### **6.2. Frame processing between TIR's**



When a frame arrives from T2SW switch to a VTSD inside TIR, the TRILL protocol will forward the frames to the proper pseudowire. When multiple paths/pseudowires are available between the TIR's then, the shortest path calculated by TRILL protocol will be used. If multiple paths are of equal cost, then TRILL protocol will do ECMP load spreading. If any multi-destination frame gets received by the VTSD through a pseudowire, TRILL will do an RPF check.

When a frame arrives from peer TIR/VTSD through a pseudowire, the MPLS header will be de-capsulated and further action will be taken depending on the egress nickname field of TRILL header. If egress nickname is the nickname of this VTSD, MAC address table and AF lookup will be performed and the frame will be forwarded by decapsulating the TRILL header. If egress nickname belongs to some other VTSD, frame will be forwarded on a pseudowire connected to that VTSD by encapsulating with an MPLS header.



## **7. Security Considerations**

This document does not change the TRILL protocol and thus has minimal security effects.

See [[RFC6325](#)] for general TRILL Security Considerations.

## **8. IANA Considerations**

This document requires no IANA actions.

RFC Editor: Please delete this section before publication

## **9. References**

### **9.1. 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>>.
  
- [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>>.
  
- [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>>.
  
- [RFC7174] Salam, S., Senevirathne, T., Aldrin, S., and D. Eastlake 3rd, "Transparent Interconnection of Lots of Links (TRILL) Operations, Administration, and Maintenance (OAM) Framework", [RFC 7174](#), DOI 10.17487/RFC7174, May 2014, <<http://www.rfc-editor.org/info/rfc7174>>.





[RFC7175] Manral, V., Eastlake 3rd, D., Ward, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Bidirectional Forwarding Detection (BFD) Support", [RFC 7175](http://www.rfc-editor.org/info/rfc7175), DOI 10.17487/RFC7175, May 2014, <<http://www.rfc-editor.org/info/rfc7175>>.

[RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", [RFC 7177](http://www.rfc-editor.org/info/rfc7177), DOI 10.17487/RFC7177, May 2014, <<http://www.rfc-editor.org/info/rfc7177>>.

[RFC7180bis] Eastlake, D., et al, "TRILL: Clarifications, Corrections, and Updates", [draft-ietf-trill-rfc7180bis](#), work in progress.

[VTSD] Umair, M., Smiler, K., Eastlake, D., Yong, L., "TRILL Transparent Transport over MPLS" [draft-muks-trill-transport-over-mpls](#), work in progress.

[RFC6439bis] Eastlake, D., et al., "TRILL: Appointed Forwarders", [draft-ietf-trill-rfc6439bis](#), work in progress.

[IA-APPsub-TLV] Eastlake, D., et al., "TRILL: Interface Addresses APPsub-TLV", [draft-ietf-trill-ia-appsubtlv](#), work in progress.

[ARP/ND-Optimization] Li, Y., Eastlake, D., et al., "TRILL: ARP/ND Optimization", [draft-ietf-trill-arp-optimization](#), work in progress.

[Address-Flush] Hao, W., Eastlake, D., "TRILL: Address Flush Protocol", [draft-hao-trill-address-flush](#), work in progress.

[Pseudo-Nickname] Zhai, H., et al., "TRILL: Pseudo-Nickname for Active-Active Access", [draft-ietf-trill-pseudonode-nickname](#), work in progress.

[Centralized-replication] Hao, W., Li, Y., et al., "Centralized



Replication for BUM traffic in active-active edge connection", [draft-ietf-trill-centralized-replication](#), work in progress.

[AA-Multi-Attach] Zhang, M., et al., "TRILL Active-Active Edge Using Multiple MAC Attachments", [draft-ietf-trill-aa-multi-attach](#), work in progress.

## **9.2. 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>>.
- [RFC7357] Zhai, H., Hu, F., Perlman, R., Eastlake 3rd, D., and O. Stokes, "Transparent Interconnection of Lots of Links (TRILL): End Station Address Distribution Information (ESADI) Protocol", [RFC 7357](#), DOI 10.17487/RFC7357, September 2014, <<http://www.rfc-editor.org/info/rfc7357>>.
- [RFC7379] Li, Y., Hao, W., Perlman, R., Hudson, J., and H. Zhai, "Problem Statement and Goals for Active-Active Connection at the Transparent Interconnection of Lots of Links (TRILL) Edge", [RFC 7379](#), DOI 10.17487/RFC7379, October 2014, <<http://www.rfc-editor.org/info/rfc7379>>.



Authors' Addresses

Mohammed Umair  
IPInfusion  
RMZ Centennial  
Mahadevapura Post  
Bangalore - 560048 India  
EMail: mohammed.umair2@gmail.com

Kingston Smiler S  
IPInfusion  
RMZ Centennial  
Mahadevapura Post  
Bangalore - 560048 India  
EMail: kingstonsmiler@gmail.com

Shaji Ravindranathan  
IPInfusion  
3965 Freedom Circle, Suite 200  
Santa Clara, CA 95054 USA  
EMail: srnathan2014@gmail.com

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

