

TRILL Working Group  
INTERNET-DRAFT  
Intended Status: Proposed Standard

Sam Aldrin  
Donald Eastlake  
Huawei Technologies  
Tissa Senevirathne  
Ayan Banerjee  
Santiago Alvarez  
CISCO

Expires: September 3, 2012

March 2, 2012

TRILL Data Center Interconnect  
draft-aldrin-trill-data-center-interconnect-00

## Abstract

This document presents a solution suite for TRILL data center sites to be connected over WAN networks. TRILL protocol is primarily designed to work within intra-data centers. Connecting different sites over WAN using overlay tunnel protocols is the primary method employed at present. Though this presents a simple mechanism to extend the LAN sites to be interconnected, it also brings in the problem of scalability for TRILL nicknames exchanged between sites, latency, duplication of traffic etc. This draft proposes a way to extend the TRILL sites without having to reveal the data of the LAN like customer MAC's or provide MAC's over the WAN, but to establish connections between various sites by extending routing protocol to exchange minimal information, thus reducing the information flow to the required sites only. Document also proposes BGP routing protocol extensions as an example to establish paths and information about the essential R Bridges nicknames, over WAN networks like MPLS.

## 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."

INTERNET DRAFT

TRILL datacenter interconnect

March 2, 2012

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>

## Copyright and License Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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>	Solution overview . . . . .	<a href="#">5</a>
<a href="#">2.1</a>	Site inter-connection . . . . .	<a href="#">5</a>
<a href="#">2.2</a>	Requirements overview . . . . .	<a href="#">6</a>
<a href="#">2.2</a>	TRILL campus extension . . . . .	<a href="#">6</a>
<a href="#">2.3</a>	TRILL nickname exhaustion . . . . .	<a href="#">6</a>
<a href="#">3</a>	Solution comparison analysis . . . . .	<a href="#">7</a>
<a href="#">3.1</a>	TRILL campus extension . . . . .	<a href="#">7</a>
<a href="#">3.2</a>	TRILL campus interconnection with E-VPN and PBB-EVPN . . . . .	<a href="#">8</a>
<a href="#">3.3</a>	TRILL campus interconnection over VPN's . . . . .	<a href="#">8</a>
<a href="#">4</a>	Operational Overview . . . . .	<a href="#">9</a>
<a href="#">4.1</a>	Campus and Backbone Areas . . . . .	<a href="#">9</a>
<a href="#">4.2</a>	Unicast forwarding . . . . .	<a href="#">10</a>
<a href="#">4.3</a>	Multicast Forwarding . . . . .	<a href="#">10</a>
<a href="#">4.4</a>	MAC learning . . . . .	<a href="#">12</a>
<a href="#">4.5</a>	TRILL nickname aggregation . . . . .	<a href="#">12</a>

<a href="#">4.6</a>	Route advertisement with BGP . . . . .	<a href="#">12</a>
<a href="#">5.</a>	TRILL campus inter-connectivity . . . . .	<a href="#">13</a>
<a href="#">5.1</a>	Route advertisement . . . . .	<a href="#">13</a>
<a href="#">5.2</a>	Inter-site nickname exchange . . . . .	<a href="#">14</a>
<a href="#">5.3</a>	Border RBridge capability exchange . . . . .	<a href="#">14</a>

<a href="#">5.4</a>	TRILL adjacency resolution . . . . .	<a href="#">14</a>
<a href="#">5.5</a>	Forwarding of data frames . . . . .	<a href="#">15</a>
<a href="#">6.</a>	Proposed additions and extensions . . . . .	<a href="#">15</a>
<a href="#">6.1</a>	BGP extension . . . . .	<a href="#">15</a>
<a href="#">6.2</a>	Border RBridge capability TLV . . . . .	<a href="#">16</a>
<a href="#">6.3</a>	TRILL nickname aggregation sub-TLV . . . . .	<a href="#">16</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="#">17</a>
	Authors' Addresses . . . . .	<a href="#">17</a>

## 1 Introduction

TRILL protocol is primarily designed as an intra-datacenter protocol by leveraging the routing functionality to interconnect bridges. Traditional Ethernet networks provided a single path for forwarding the traffic, which is usually derived using protocols like Spanning Tree. TRILL provided a way to utilize multiple links for forwarding, thus utilizing the resources effectively. Even though TRILL is new protocol, it seamlessly integrates with legacy bridging networks without having to forklift upgrade of all the bridges to support TRILL. By not having to learn the MAC addresses of end stations by intermediate devices, provided a powerful way to interconnect bridges within a datacenter and maximizing the resource usage and providing multipath usage option.

TRILL enabled network creates efficiency by having reduced forwarding table size. By doing TRILL nickname based forwarding created a layer of abstraction and much easier to implement the protocol. This enabled to address the scalability of a L2 domain, where thousands of Rbridges could exist to meet the needs of a datacenter. By leveraging IS-IS protocol, the information exchange and leveraging the path computation technology brought forth a new paradigm into bridging technology. TRILL Base Protocol Specification [[RFC6325](#)] specifies a tree based paradigm to forward broadcast and multicast traffic as well as unknown unicast traffic.

Even though the TRILL is enabled within a datacenter and is not primarily designed to work over WAN, there is a need to interconnect various TRILL data sites. The same datacenter provider could be

having multiple TRILL sites and these TRILL enabled datacenter sites could run independently or could share resources in order to cater to the needs of customers. As such, there exist few proposals based on overlay technologies which interconnect these sites but those solutions require MAC learning at the edge R Bridges and stripping of TRILL nickname on the frame. Another option is to interconnect these TRILL sites using Pseudowires and making a huge TRILL site. This is useful option but the downside of this is when provider would like to maintain independent sites and exchange only the required data to be shared across sites, it becomes complicated to maintain the networks.

This draft solves these core problems of interconnection, site independence, dynamic information exchange and setting up of connections over WAN, leveraging existing WAN technologies like MPLS, etc. Though the primary goal is effective interconnection, there are various efficient schemes, which could enable seamless TRILL network deployments, are also be presented. Solution covers both unicast and multicast data traffic.

## [1.1](#) Terminology

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 [RFC 2119](#) [[RFC2119](#)].

## [2.](#) Solution overview

This section provides the high-level overview of the solution to the problems in various scenarios. More detailed representation of the solution is covered in the later sections of the draft.

TRILL site or TRILL campus uses IS-IS to setup the R Bridge interconnection. A R Bridge knows how to reach another R Bridge within the campus. When two TRILL campuses are interconnected, one could visualize it in two different perspectives. First one is a merger of two TRILL campuses into one. This requires each R Bridge to know about the other and the IS-IS should be able to compute the shortest paths from one to another. The downside of this model is that information exchange explosion and the size of IS-IS db and number of PDU's being exchanged could increase exponentially.

The second perspective is to have these two TRILL campuses being interconnected over a WAN, but their functioning nature is independent to each other. The two campuses exchanges only the required information between border RBridges of the campuses. This will be more optimal and leads to interconnecting multiple campuses without having to redesign the whole network to ensure uniqueness and identity of RBridges.

Solution being proposed is an option of maintaining site independency with a simplified solution and modeled around the existing and proven technologies. Some of the enhancements were already proposed in other drafts like multilevel TRILL [[TISSA-MLEVEL](#)] and the solution leverages by extending those definitions as necessary. The solution addresses the following areas described in the following sections

## [2.1](#) Site inter-connection

Each TRILL campus is considered as an independent site or an L1 IS-IS domain. These TRILL campus sites are interconnected over WAN. Each area will have an appointed border RBridges. These RBridges exchange the information of other border RBridges of different TRILL campus sites to establish connection with each other. In order to exchange the information, the route has to be established. Extension of BGP protocol is done to exchange the border RBridges information and establish a route or vpn/mvpn connection between each of the border

RBridges. More details of the extension are detailed in the below section. These sites are interconnected either over IP or MPLS. As MPLS is a mature WAN technology, this draft references the solution based on MPLS. This does not preclude that other technologies could not be used.

## [2.2](#) Requirements overview

There are various requirements necessary to be met in order to provide a seamless TRILL inter-connectivity across campuses and datacenter. Some of the important requirements are as follows

- o Extend TRILL technology over interconnect

- o Ability to provide the same fast convergence for mobility as it does in intra-TRILL campus
- o Ability to work with various WAN technologies
- o Option for dynamic establishment of connectivity across sites
- o Minimal changes to protocols and definitions
- o Backward compatible to existing networks and their functions.

## [2.2](#) TRILL campus extension

By interconnecting TRILL campus sites over WAN, one could extend the L1 area, but that would cause other issues as detailed in the earlier section. With this solution, all of the TRILL nicknames of one campus are not exchanged with other campuses. Instead, there will be R Bridges which gets appointed, so only those specific nicknames are exchanged with others. Also one could create another hierarchical model like VPN, where participating appointed R Bridges for that VPN could be exchanged with other campuses belonging to that VPN. Appointed or designated R Bridges information will be exchanged with other campus site using the Affinity TLV extension. Detailed description on how to create and the usage are detailed in multilevel draft [[TISSA-MLEVEL](#)]. When each campus site creates this information and exchanged with other campus sites of the network or VPN, this could be used for two purposes, 1. A VPN specific WAN paths could be created between campus sites. 2. The data distribution could be optimized without having the need to be broadcast the data frames from one campus into every campus site.

## [2.3](#) TRILL nickname exhaustion

Though this draft is not meant to provide solution for TRILL nickname exhaustion, it enables provider to deal with the problem effectively and not having to re-design the network, every time a new campus is interconnected. The proposed solution has R Bridges which are not required to be exposed outside of the campus and there are other R Bridges which are also known as border R Bridges. These border R Bridges nicknames are unique globally. Rest of the R Bridges nicknames are significant locally, that includes the appointed


RBridges nicknames for a VPN. When a frame has to be forwarded to an RBridge which resides in another campus, the originating RBridge only knows how to get to the borderRBridge. This border RBridge should have the list of RBridges of other campus sites and thus could select the appropriate MPLS LSP or MVPN or PW and encapsulate the TRILL frame with the label header and forward over that. More details are covered in the unicast and multicast sections of the detailed solution.

### [3.](#) Solution comparison analysis

As eluded to in the earlier sections, there are various methods on interconnecting different TRILL campus sites. Before going into the details of proposed solution a close examination of some of the proposed solutions, provides better perspective of this solution.

#### [3.1](#) TRILL campus extension

In this model TRILL campuses are connected over WAN using technologies like PW. This is the most simplest way of interconnecting the sites. When campuses are interconnected, the TRILL campus will get expanded and each RBridge could reach each other. The main criteria for this will be to maintain unique nickname for RBridges.



As shown in the figure above, two TRILL campuses are interconnected over WAN. Border RBridges establish connection over WAN using PW or other WAN technologies. All the nicknames within each campus sites have to be unique. The WAN in this case is transparent to the TRILL

campuses and the path computation doesn't involve WAN component,




instead it will be like one TRILL campus. When RB1 originates a TRILL frame destined to RB2, it traverses over BRB1 and BRB2 and reaches RB2.

This solution is workable when the campuses are small and do NOT need to change or requires interconnecting more TRILL campuses. The other downside for this model is, when two campuses are interconnected and there is overlap of nicknames, the network has to be re-designed to eliminate the duplicate nicknames and make each RBridge to have a unique nickname.

### [3.2](#) TRILL campus interconnection with E-VPN and PBB-EVPN

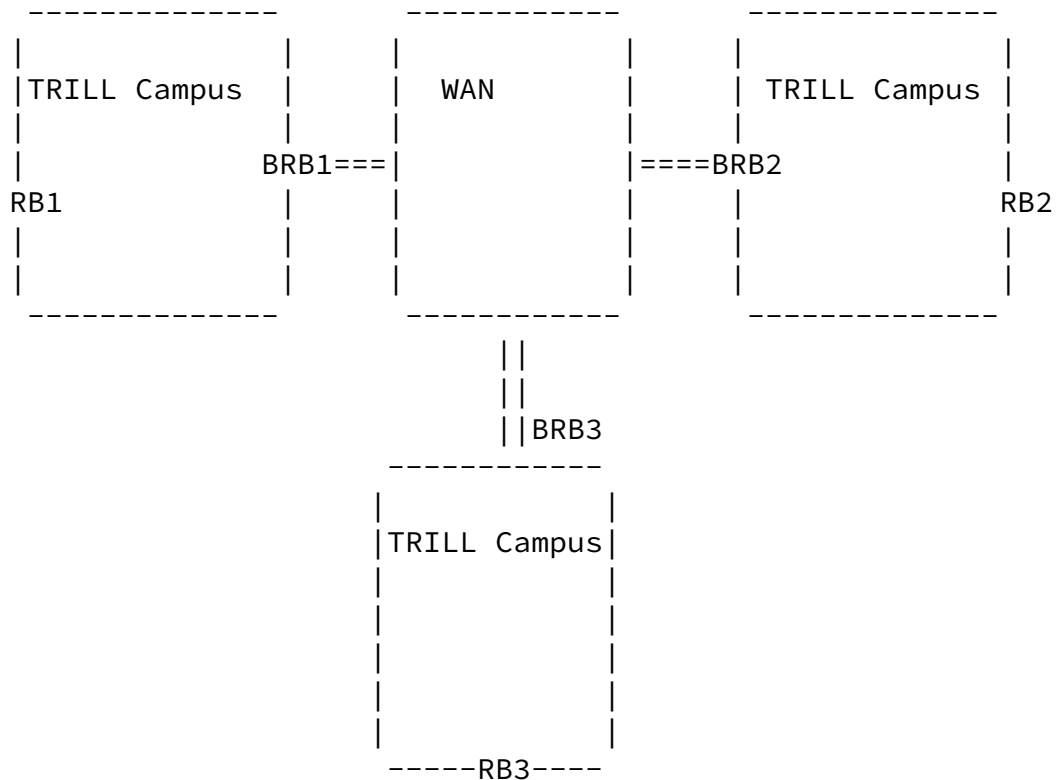
TRILL campuses could be extended over WAN using E-VPN and PBB-EVPN.



The [PBB-EVPN] draft proposes interconnection details on how two TRILL campuses could be interconnected using the E-VPN technology. In this a new BGP route is advertised for reachability of TRILL R Bridges. This technique leverages the PBB technology and also enables to retain TRILL header but is recommended to avoid transmitting TRILL encapsulated frames over the WAN links. The primary downside of this method is the requirement for edge R Bridges to learn MAC Addresses in order to resolve the adjacency.

### [3.3](#) TRILL campus interconnection over VPN's

In this method, TRILL campus sites could be interconnected over VPN's.



These VPN's could be established statically or dynamically. In order to establish dynamically, the border RBridges needs to exchange information of the nicknames and connect different sites. The hierarchical model like H-VPLS could be established as well. [PBB-EVPN] draft provides some details on how to achieve this, but still requires border RBridges to exchange MAC information and resolution for L2 adjacency. One other option is much similar to the first model where campuses exchange TRILL nicknames between campuses over VPN's. Though this model groups different sites according to the way VPN's are configured, avoiding flooding of TRILL nicknames or site independency cannot be achieved.

#### [4. Operational Overview](#)

##### [4.1 Campus and Backbone Areas](#)

Each TRILL campus will be assigned a border RBridge. This is identified using the 'Attached' bit in the IS-IS PDU. The border RBridge has list of the RBridges of each campus site. These list of bridges are exchanged using the TRILL nickname aggregation sub-TLV. Details of the sub-TLV are detailed in the below section.

Every TRILL campus campus need not exchange all the RBridge nicknames

with other campuses. Let us take the scenario of campus A to be interconnected with campus B. In campus A, there are RBridges

RB1...RB10, which are interconnected in L1. These nicknames are not tunneled or exchanged with other L1 campus sites. Similarly campus B has RB11...RB20 and need not be distinct from campus A RBridge nicknames. So, if a new campus is connected to the domain, there is no need for the network to be redesigned or restructured.

The RBridges at each campus advertise the information to other campus to establish a route/MPLS LSP between campus sites. A new BGP extension as defined below is sent out. This reachability TLV is received by other border RBridges and establish the MPLS LSP between them.

The border RBridges will have the complete information of its campus RBridges. Not all of the RBridges nicknames need not be advertised globally. So, the globally exchanged nicknames of RBridges should be unique across campuses. Depending on the policy established, these Border RBridges will exchange the TRILL information with other campus border Bridges, over the routes/MPLS LSP established between different campuses. In IS-IS domain the equivalent of this is the L2 backbone area, which in this case, is established over WAN.

#### [4.2](#) Unicast forwarding

If the destination TRILL nickname is not known, the originating or transit RBridges forwards it to border RBridge. As the border RBridge has all the nicknames of each campus, it forwards the frame to the right campus border RBridge, which in turn could forward within its campus as per base protocol specification [[RFC6325](#)]. In the case where the destination is unknown, the frame is flooded to each campus. Using the MAC learning procedures, the associated RBridge will be learnt for the subsequent frames to be forwarded as unicast, instead of flooding. Flooding into various campuses or TRILL data sites happen only if the the frame is of global ID based on VLAN identification.

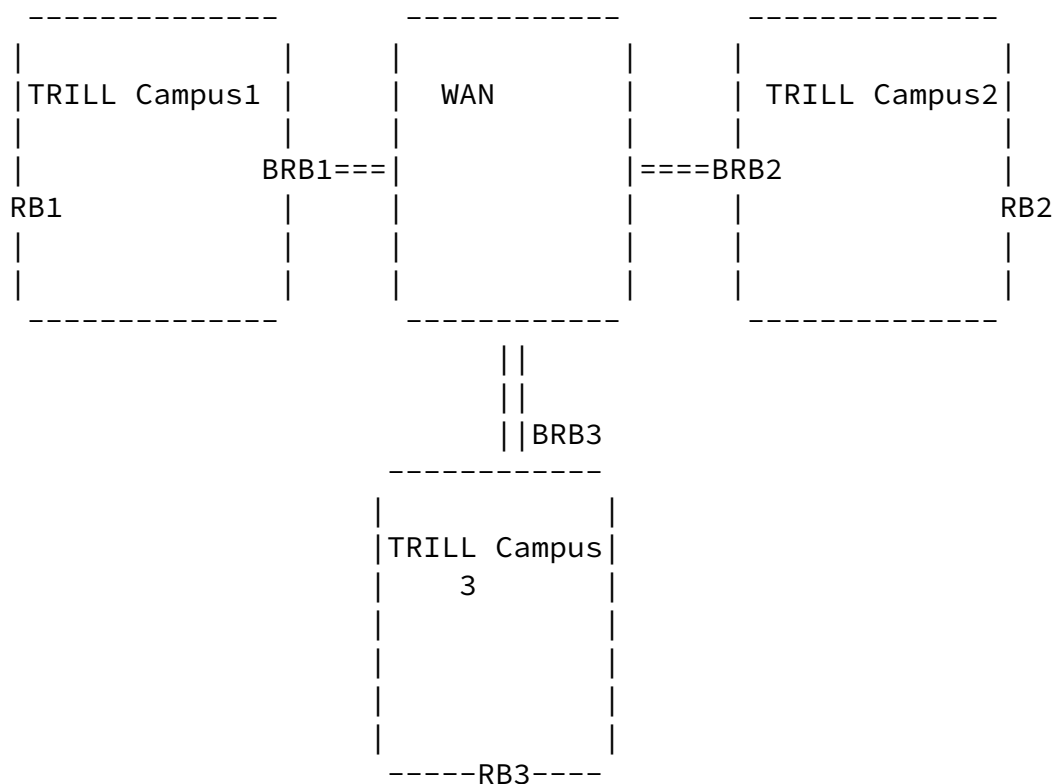
#### [4.3](#) Multicast Forwarding

Whether the traffic scope is local or global across campuses is identified by VLAN or port or fine-grain label. If the traffic is to

be forwarded within campus, it is done using the local tree. But if the forwarding has to be done globally, it needs to use the global tree. When the frame has global context, it will be flooded into other TRILL sites as well.

In MPLS the multicast networks are established within the core networks. In the case of RBridges, which are part of the customer networks and do not participate in the service provider networks and their topologies, the multicast tree could be built using IP

multicast or leverage MVPN services offered by the service provider. The global trees are established between border RBridges with the help of information exchanges between border RBridges. As the IS-IS is limited to individual campus sites, the information for the backbone tree over WAN has to be exchanged between border RBridges either as a IP multicast PIM message or specific TLV indented for other campus border RBridges. More details on this will be added in the later versions of the draft.



In the above figure when the multicast frame has to be sent from campus 1 to campus 2 and 3, the frame arrives at border RBridge BRB1. With the default global tree between border RBridges of different campuses, the forwarding is setup to egress the frame or replicated over MPLS LSP's to all other campus sites. If the frame is destined for non-default global tree, the appropriate MPLS LSP(s) are identified and the frame is forwarded accordingly.

Once the frame is reached on the border router of the campuses, the frame is locally multicast forwarded. The same technique as employed in the multilevel draft [[TISSA-MLEVEL](#)] is used here as well.

If mVPN services are deployed interconnecting campus sites, the multicast tree is built over these services based on the customer VLANs.

#### [4.4](#) MAC learning

When a frame is to be forwarded from customer MAC A to customer MAC B, the frame is set as unknown unicast frame over TRILL networks. If the MAC A and MAC B are connected over WAN, the frame is transmitted over WAN to the other campus. When the frame is reached at the RBridge connecting to MAC B, it will learn about the originator RBridge for MAC A. While responding, the egress RBridge know the originating RBridge, it will unicast the frame to the originator.

#### [4.5](#) TRILL nickname aggregation

Nicknames are allocated or assigned to RBridges in a given campuses using various methods. It could be OSS, CLI or could be a dynamic control protocol which configures the nicknames. As the nicknames are confined to each L1 area, the nickname management, when sites are connected over WAN, it is essential to optimize the name allocation in order to use the name space effectively. Name allocation is not in the scope of this draft. If there is a necessity, the topic could be considered in the later revisions of the draft. For this version we do recommend some of the optimization techniques for nickname allocation defined in the multilevel draft [[TISSA-MLEVEL](#)].

Each border RBridges needs to exchange the nicknames of each campuses with other border RBridges. As the border RBridges are connected over

various types of WAN networks, mandating enhancement to a specific protocol is deemed not the right approach. As the information exchange has to be done, certain characteristics for the data exchange have to be met.

- o The amount of data exchange has to be minimal and optimized.
- o the information exchange has to be quick.
- o Conflicts and duplicate information flow has to be avoided

This draft proposes a generic TLV, which is to be exchanged between border RBridges. If the nickname allocation is done in terms of ranges, the same could be exchanged between border RBridges, seamlessly. As the TLV has to be terminated at the border RBridges, it is better suited to be sent as a unicast message to the neighboring border RBridge. This could be sent as an IP message with TRILL header containing the target border RBridge. More details on how to encapsulate and process the frame should be in the later versions of the draft.

#### [4.6](#) Route advertisement with BGP

BGP route advertisement is to connect border RBridges. As described in earlier sections, each campus site exchanges the TRILL nicknames with other campuses. These nicknames are not leaked in to the campus but will be maintained at the border RBridges. The connectivity between these border RBridges is similar to L2 backbone of IS-IS but in this case it is over WAN and IS-IS is absent.

This draft, unlike some other drafts, do not propose to de-capsulate the TRILL frame to learn the MAC addresses. Instead, it propose to use nicknames themselves to be exchanged between sites.

A new BGP route advertisement is defined to advertise the border RBridge nickname with the other border RBridges in different campuses. This advertisement will let other campuses know its MPLS label, its nickname, IP address etc. Once the route is established, further information of campus nicknames etc could be exchanges to establish the inter-connectivity of the TRILL campuses

## [5. TRILL campus inter-connectivity](#)

The primary reason to interconnect TRILL campuses is to maintain geographically distant, segmented sites and customer specific segregation possible by interconnecting and not having to redo the network and campus redesign for every change and need. With customers being mobile or services offered by service providers could be re-located depending on the time-zones and resource availability, restricting to a specific site is a thing of the past.

These constraints brought forth the need to have different sites interconnected over the WAN, be it MPLS or VPLS or IP and to provide the services on demand to meet the needs of customers and their data center needs. As TRILL has proven to be an effective protocol by bringing routing technologies into bridges or L2 forwarding, the short coming of TRILL interconnect is the immediate need. As eluded to in the earlier sections on different kinds of solutions, meeting all the needs of the TRILL DCI as laid out in the requirements section, is the primary goal of this draft.

### [5.1 Route advertisement](#)

Border RBridges only participate in interconnecting various TRILL campuses. These border RBridges are elected or identified as described in the earlier section i.e. using IS-IS protocol advertisement. These border RBridges, when required to interconnect with other campuses, advertise the route to other site border RBridges using the BGP enhancement. Upon receipt of the route advertisement, the MPLS LSP or IP path or Pseudowire is established between RBridges and they could communicate with each other and

exchange other information.

If L2 connectivity is to be used with protocols like VPLS, a similar method could be employed, where the PWE3 could be established between border RBridges. More details to be added in the later versions of the draft.

### [5.2 Inter-site nickname exchange](#)

There are three types of nicknames which are exchanged between border RBridges.

- o Nickname of border RBridges
- o Nicknames of RBridges for each campus
- o Nicknames of RBridges which are part of a specific customer VLAN or VPN

The nickname aggregation TLV is used as payload to be exchanged between border RBridges. This information is used to establish inter-connectivity between TRILL sites per customer VLAN or default global tree.

### [5.3](#) Border RBridge capability exchange

An additional capability TLV is defined to exchange info on what each of the border RBridge is capable of. This is very essential for forward and backward capability . Capability information not only indicates the capability version but could also force the interconnection to be restricted as per the policy set by the customer. Some of the capability advertisements are as follows.

- o Version.
- o default nick name resolution
- o connect more campuses
- o active-active link support
- o Ability to support multicast forwarding

### [5.4](#) TRILL adjacency resolution

When a frame is to be forwarded from one campus to another, the adjacency resolution has to be done on the border RBridge. When TRILL nicknames are advertised from one border RBridge to another, the

border RBridge keeps the database of all the nicknames. The MPLS LSP's between each RBridges are established by the route advertisements. VPN specific services could be established based on the customer VLAN ID's and also the exchange of nicknames per VPN



between border RBridges.

Once the frame is received on the border RBridge, it will look in the forwarding table to identify the next hop. The adjacency information could indicate an MPLS LSP with a specific label encapsulation. For VPN, there will be additional VPN label in the label stack. The TRILL frame is encapsulated, without removing the TRILL header and is forwarded over the MPLS LSP.

## [5.5](#) Forwarding of data frames

The TRILL frames are forwarded as per the base protocol [[RFC6325](#)] within a campus site. The forwarding of the frames from TRILL campus to campus will be over MPLS LSP's or IP or whichever WAN connection is established between border RBridges. The encapsulation of the Frame with WAN header is based on the adjacency resolution made in the forwarding on the border RBridge.

## [6](#). Proposed additions and extensions

There are certain extensions being proposed in this draft to interconnect TRILL campuses or datacenters. These include new additions to routing and also new TLV's to exchange information between border RBridges. There are few references to the extensions proposed in other drafts which are used in this draft as well.

### [6.1](#) BGP extension

A new BGP route advertisement is done to exchange and establish route/MPLS lsp between border RBridges.

```
+-----+
|RD (8 octets)                               |
+-----+
|Originating RBridge MAC address             |
+-----+
|Originating RBridge IP address(v4/v6)      |
+-----+
| Nickname Length (1 octet)                  |
+-----+
| RBridge Nickname (2 octets)                 |
+-----+
| MPLS Label (n * 3 octets)                  |
+-----+
```

## 6.2 Border RBridge capability TLV

This TLV as defined in earlier section, defines the capability of a border RBridge, to be exchanged with other border RBridges for seamless inter-working across campus sites.

```

    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type = <TBD>                      |      Length = 8      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Version                          |      Flags              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Flags                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Definition of flag bits will be identified and defined later.

## 6.3 TRILL nickname aggregation sub-TLV

The nickname aggregation TLV defined in multilevel draft [TISSA-MLEVEL] is used in advertising the nicknames into other border Routers. Some new additions or changes will be proposed in later versions of the draft.

INTERNET DRAFT

TRILL datacenter interconnect

March 2, 2012

## [7](#) Security Considerations

<Security considerations text>

## [8](#) IANA Considerations

<IANA considerations text>

## [9](#) References

### [9.1](#) Normative References

- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2234] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 2234](#), November 1997.
- [RFC4971] Vasseur, JP., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", [RFC 4971](#), July 2007.
- [RFC6325] Perlman, R., et.al, "Routing Bridges (RBridges): Base Protocol Specification", [RFC 6325](#), July 2011.
- [trillcmt] Senevirathne, T., et.al, "Coordinated Multicast Trees (CMT) for TRILL", Work in Progress, January 2012.

### [9.2](#) Informative References

- [RFC6326] Eastlake, D, et.al, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [RFC 6326](#), July 2011.
- [rfc6326bis] Eastlake, D, et.al, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", Work in Progress, [draft-eastlake-isis-rfc6326bis-04.txt](#), January 2012.

[TISSA-MLEVEL] Senevirathne, "RBridges: Multilevel TRILL", Work in Progress, [draft-tissa-trill-multilevel-00.txt](#), February 201.

## Authors' Addresses

Aldrin, et.al. Expires September 3, 2012 [Page 17]

---

INTERNET DRAFT TRILL datacenter interconnect March 2, 2012

Sam Aldrin  
Huawei Technologies  
2330 Central Express Way  
Santa Clara, CA 95951

Email: [aldrin.ietf@gmail.com](mailto:aldrin.ietf@gmail.com)

Tissa Senevirathne  
CISCO Systems  
375 East Tasman Drive  
San Jose CA 95134

Phone: 408-853-2291  
Email: [tsenevir@cisco.com](mailto:tsenevir@cisco.com)

Ayan Banerjee  
CISCO Systems  
425 East Tasman Drive  
San Jose CA 95134

Phone: 408-527-0539  
Email: [ayabaner@cisco.com](mailto:ayabaner@cisco.com)

Donald Eastlake  
Huawei Technologies  
155 Beaver Street  
Milford, MA 01757 USA

Phone: +1-508-333-2270  
Email: [d3e3e3@gmail.com](mailto:d3e3e3@gmail.com)

Santiago Alvarez  
CISCO systems

Email: [saalvare@cisco.com](mailto:saalvare@cisco.com)

Aldrin, et.al.

Expires September 3, 2012

[Page 18]