

Internet Engineering Task Force (IETF)  
Request for Comments: 7783  
Updates: 6325  
Category: Standards Track  
ISSN: 2070-1721

T. Senevirathne  
Consultant  
J. Pathangi  
Dell  
J. Hudson  
Brocade  
February 2016

Coordinated Multicast Trees (CMT)  
for Transparent Interconnection of Lots of Links (TRILL)

Abstract

TRILL (Transparent Interconnection of Lots of Links) facilitates loop-free connectivity to non-TRILL networks via a choice of an Appointed Forwarder for a set of VLANs. Appointed Forwarders provide VLAN-based load sharing with an active-standby model. High-performance applications require an active-active load-sharing model. The active-active load-sharing model can be accomplished by representing any given non-TRILL network with a single virtual RBridge (also referred to as a virtual Routing Bridge or virtual TRILL switch). Virtual representation of the non-TRILL network with a single RBridge poses serious challenges in multi-destination RPF (Reverse Path Forwarding) check calculations. This document specifies required enhancements to build Coordinated Multicast Trees (CMT) within the TRILL campus to solve related RPF issues. CMT, which only requires a software upgrade, provides flexibility to RBridges in selecting a desired path of association to a given TRILL multi-destination distribution tree. This document updates RFC 6325.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7783>.

## Copyright 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 .....	3
1.1. Scope and Applicability .....	4
2. Conventions Used in This Document .....	5
2.1. Acronyms and Phrases .....	5
3. The Affinity Sub-TLV .....	6
4. Multicast Tree Construction and Use of Affinity Sub-TLV .....	6
4.1. Update to RFC 6325 .....	7
4.2. Announcing Virtual RBridge Nickname .....	8
4.3. Affinity Sub-TLV Capability .....	8
5. Theory of Operation .....	8
5.1. Distribution Tree Assignment .....	8
5.2. Affinity Sub-TLV Advertisement .....	9
5.3. Affinity Sub-TLV Conflict Resolution .....	9
5.4. Ingress Multi-Destination Forwarding .....	10
5.4.1. Forwarding when $n < k$ .....	10
5.5. Egress Multi-Destination Forwarding .....	11
5.5.1. Traffic Arriving on an Assigned Tree to RBk-RBv .....	11
5.5.2. Traffic Arriving on Other Trees .....	11
5.6. Failure Scenarios .....	11
5.6.1. Edge RBridge RBk Failure .....	11
5.7. Backward Compatibility .....	12
6. Security Considerations .....	13
7. IANA Considerations .....	13
8. References .....	14
8.1. Normative References .....	14
8.2. Informative References .....	15
Acknowledgments .....	16
Contributors .....	16
Authors' Addresses .....	16

## 1. Introduction

TRILL (Transparent Interconnection of Lots of Links), as presented in [RFC6325] and other related documents, provides methods of utilizing all available paths for active forwarding, with minimum configuration. TRILL utilizes IS-IS (Intermediate System to Intermediate System) [IS-IS] as its control plane and uses a TRILL Header that includes a Hop Count.

[RFC6325], [RFC7177], and [RFC6439] provide methods for interoperability between TRILL and Ethernet end stations and bridged networks. [RFC6439] provides an active-standby solution, where only one of the RBridges (aka Routing Bridges or TRILL switches) on a link with end stations is in the active forwarding state for end-station traffic for any given VLAN. That RBridge is referred to as the Appointed Forwarder (AF). All frames ingress into a TRILL network via the AF are encapsulated with the TRILL Header with a nickname held by the ingress AF RBridge. Due to failures, reconfigurations, and other network dynamics, the AF for any set of VLANs may change. RBridges maintain forwarding tables that contain bindings for destination Media Access Control (MAC) addresses and Data Labels (VLAN or Fine-Grained Labels (FGLs)) to egress RBridges. In the event of an AF change, forwarding tables of remote RBridges may continue to forward traffic to the previous AF. That traffic may get discarded at the egress, causing traffic disruption.

High-performance applications require resiliency during failover. The active-active forwarding model minimizes impact during failures and maximizes the available network bandwidth. A typical deployment scenario, depicted in Figure 1, may have end stations and/or bridges attached to the RBridges. These devices typically are multi-homed to several RBridges and treat all of the uplinks independently using a Local Active-Active Link Protocol (LAALP) [RFC7379], such as a single Multi-Chassis Link Aggregation (MC-LAG) bundle or Distributed Resilient Network Interconnect [802.1AX]. The AF designation presented in [RFC6439] requires each of the edge RBridges to exchange TRILL Hello packets. By design, an LAALP does not forward packets received on one of the member ports of the MC-LAG to other member ports of the same MC-LAG. As a result, the AF designation methods presented in [RFC6439] cannot be applied to the deployment scenario depicted in Figure 1.

An active-active load-sharing model can be implemented by representing the edge of the network connected to a specific edge group of RBridges by a single virtual RBridge. Each virtual RBridge MUST have a nickname unique within its TRILL campus. In addition to an active-active forwarding model, there may be other applications that may require similar representations.

Sections 4.5.1 and 4.5.2 of [RFC6325], as updated by [RFC7780], specify distribution tree calculation and RPF (Reverse Path Forwarding) check calculation algorithms for multi-destination forwarding. These algorithms strictly depend on link cost and parent RBridge priority. As a result, based on the network topology, it may be possible that a given edge RBridge, if it is forwarding on behalf of the virtual RBridge, may not have a candidate multicast tree on which it (the edge RBridge) can forward traffic, because there is no tree for which the virtual RBridge is a leaf node from the edge RBridge.

In this document, we present a method that allows RBridges to specify the path of association for real or virtual child nodes to distribution trees. Remote RBridges calculate their forwarding tables and derive the RPF for distribution trees based on the distribution tree association advertisements. In the absence of distribution tree association advertisements, remote RBridges derive the SPF (Shortest Path First) based on the algorithm specified in Section 4.5.1 of [RFC6325], as updated by [RFC7780]. This document updates [RFC6325] by changing, when Coordinated Multicast Trees (CMT) sub-TLVs are present, [RFC6325] mandatory provisions as to how distribution trees are constructed.

In addition to the above-mentioned active-active forwarding model, other applications may utilize the distribution tree association framework presented in this document to associate to distribution trees through a preferred path.

This proposal requires (1) the presence of multiple multi-destination trees within the TRILL campus and (2) that all the RBridges in the network be updated to support the new Affinity sub-TLV (Section 3). It is expected that both of these requirements will be met, as they are control-plane changes and will be common deployment scenarios. In case either of the above two conditions is not met, RBridges MUST support a fallback option for interoperability. Since the fallback is expected to be a temporary phenomenon until all RBridges are upgraded, this proposal gives guidelines for such fallbacks and does not mandate or specify any specific set of fallback options.

### 1.1. Scope and Applicability

This document specifies an Affinity sub-TLV to solve RPF issues at the active-active edge. Specific methods in this document for making use of the Affinity sub-TLV are applicable where a virtual RBridge is used to represent multiple RBridges connected to an edge Customer Equipment (CE) device through an LAALP, such as MC-LAG or some similar arrangement where the RBridges cannot see each other's Hellos.

This document does not provide other required operational elements to implement an active-active edge solution, such as MC-LAG methods. Solution-specific operational elements are outside the scope of this document and will be covered in other documents. (See, for example, [RFC7781].)

Examples provided in this document are for illustration purposes only.

## 2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lowercase uses of these words are not to be interpreted as carrying [RFC2119] significance.

### 2.1. Acronyms and Phrases

The following acronyms and phrases are used in this document:

AF: Appointed Forwarder [RFC6439].

CE: Customer Equipment device, that is, a device that performs forwarding based on 802.1Q bridging. This also can be an end station or a server.

Data Label: VLAN or FGL.

FGL: Fine-Grained Label [RFC7172].

LAALP: Local Active-Active Link Protocol [RFC7379].

MC-LAG: Multi-Chassis Link Aggregation. A proprietary extension to [802.1AX] that facilitates connecting a group of links from an originating device (A) to a group of discrete devices (B). Device (A) treats all of the links in a given MC-LAG bundle as a single logical interface and treats all devices in Group (B) as a single logical device for all forwarding purposes. Device (A) does not forward packets received on the multi-chassis link bundle out of the same multi-chassis link bundle. Figure 1 depicts a specific use case example.

RPF: Reverse Path Forwarding. See Section 4.5.2 of [RFC6325].

Virtual RBridge: A purely conceptual RBridge that represents an active-active edge group and is in turn represented by a nickname. For example, see [RFC7781].

3. The Affinity Sub-TLV

Association of an RBridge to a multi-destination distribution tree through a specific path is accomplished by using a new IS-IS sub-TLV, the Affinity sub-TLV.

The Affinity sub-TLV appears in Router Capability TLVs or MT Capability TLVs that are within Link State PDUs (LSPs), as described in [RFC7176]. [RFC7176] specifies the code point and data structure for the Affinity sub-TLV.

4. Multicast Tree Construction and Use of Affinity Sub-TLV

Figures 1 and 2 below show the reference topology and a logical topology using CMT to provide active-active service.

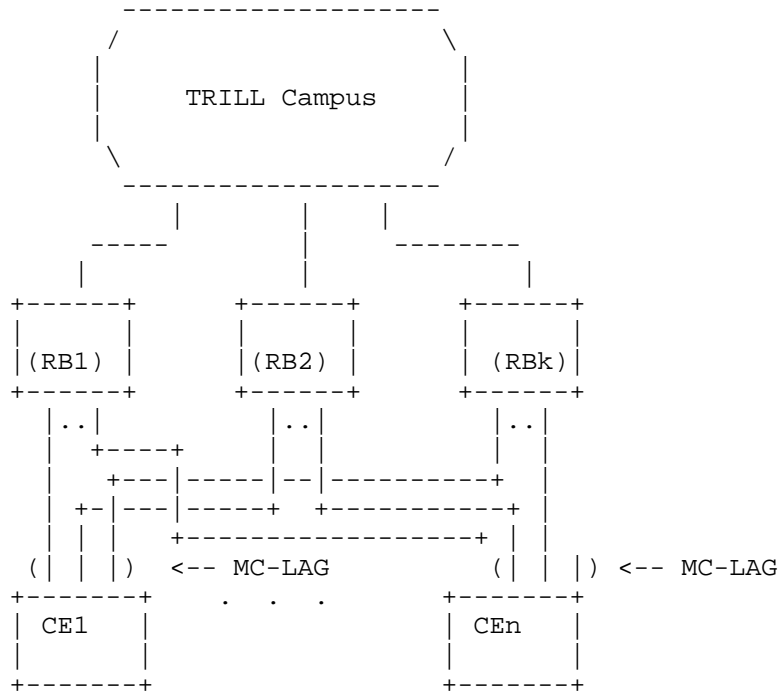


Figure 1: Reference Topology

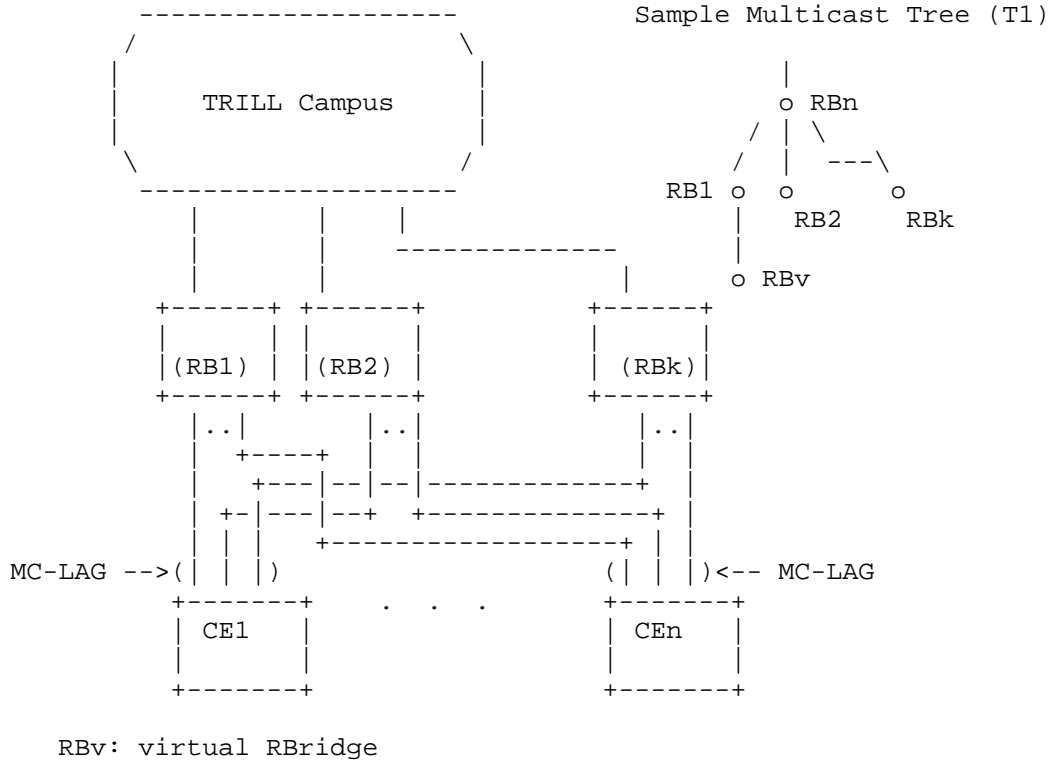


Figure 2: Example Logical Topology

4.1. Update to RFC 6325

This document updates Section 4.5.1 of [RFC6325] and changes the calculation of distribution trees, as specified below:

Each RBridge that desires to be the parent RBridge for a child RBridge (RBy) in a multi-destination distribution tree (Tree x) announces the desired association using an Affinity sub-TLV. The child is specified by its nickname. If an RBridge (RB1) advertises an Affinity sub-TLV designating one of its own nicknames (N1) as its "child" in some distribution tree, the effect is that nickname N1 is ignored when constructing other distribution trees. Thus, the RPF check will enforce the rule that only RB1 can use nickname N1 to do ingress/egress on Tree x. (This has no effect on least-cost path calculations for unicast traffic.)

When such an Affinity sub-TLV is present, the association specified by the Affinity sub-TLV MUST be used when constructing the multi-destination distribution tree, except in the case of a

conflicting Affinity sub-TLV; such cases are resolved as specified in Section 5.3. In the absence of such an Affinity sub-TLV, or if there are any RBridges in the campus that do not support the Affinity sub-TLV, distribution trees are calculated as specified in Section 4.5.1 of [RFC6325], as updated by [RFC7780]. Section 4.3 below specifies how to identify RBridges that support the Affinity sub-TLV.

#### 4.2. Announcing Virtual RBridge Nickname

Each edge RBridge (RB1 to RBk) advertises its LSP virtual RBridge nickname (RBv) by using the Nickname sub-TLV (6) [RFC7176], along with their regular nickname or nicknames.

It will be possible for any RBridge to determine that RBv is a virtual RBridge, because each RBridge (RB1 to RBk) that appears to be advertising that it is holding RBv is also advertising an Affinity sub-TLV asking that RBv be its child in one or more trees.

Virtual RBridges are ignored when determining the distribution tree roots for the campus.

All RBridges outside the edge group assume that multi-destination packets with their TRILL Header Ingress Nickname field set to RBv might use any of the distribution trees that any member of the edge group advertises that it might use.

#### 4.3. Affinity Sub-TLV Capability

RBridges that announce the TRILL Version sub-TLV [RFC7176] and set the Affinity capability bit (Section 7) support the Affinity sub-TLV, calculation of multi-destination distribution trees, and RPF checks, as specified herein.

### 5. Theory of Operation

#### 5.1. Distribution Tree Assignment

Let's assume that there are  $n$  distribution trees and  $k$  edge RBridges in the edge group of interest.

If  $n \geq k$

Let's assume that edge RBridges are sorted in numerically ascending order by IS-IS System ID such that  $RB1 < RB2 < RBk$ . Each RBridge in the numerically sorted list is assigned a monotonically increasing number  $j$  such that  $RB1 = 0$ ,  $RB2 = 1$ ,



$R_{Bi} = j$ , and  $R_{Bi} + 1 = j + 1$ . (See Section 4.5 of [RFC6325], as updated by Section 3.4 of [RFC7780], for how tree numbers are determined.)

Assign each tree to  $R_{Bi}$  such that tree number  $((\text{tree\_number}) \% k) + 1$  is assigned to edge group  $R_{Bridge\ i}$  for tree\_number from 1 to  $n$ , where  $n$  is the number of trees,  $k$  is the number of edge group  $R_{Bridges}$  considered for tree allocation, and "%" is the integer division remainder operation.

If  $n < k$

Distribution trees are assigned to edge group  $R_{Bridges\ R_{B1}}$  to  $R_{Bn}$ , using the same algorithm as the  $n \geq k$  case.  $R_{Bridges\ R_{Bn} + 1}$  to  $R_{Bk}$  do not participate in the active-active forwarding process on behalf of  $R_{Bv}$ .

### 5.2. Affinity Sub-TLV Advertisement

Each  $R_{Bridge}$  in the  $R_{B1}$  through  $R_{Bk}$  domain advertises an Affinity sub-TLV for  $R_{Bv}$  to be its child.

As an example, let's assume that  $R_{B1}$  has chosen Trees  $t_1$  and  $t_{k+1}$  on behalf of  $R_{Bv}$ .

$R_{B1}$  advertises the Affinity sub-TLV;  
{ $R_{Bv}$ , Num of Trees = 2,  $t_1$ ,  $t_{k+1}$ }.

Other  $R_{Bridges}$  in the  $R_{B1}$  through  $R_{Bk}$  edge group follow the same procedure.

### 5.3. Affinity Sub-TLV Conflict Resolution

In TRILL, multi-destination distribution trees are built outward from the root by each  $R_{Bridge}$  so that they all derive the same set of distribution trees [RFC6325].

If  $R_{Bridge\ R_{B1}}$  advertises an Affinity sub-TLV with an AFFINITY RECORD that asks for  $R_{Bridge\ R_{Broot}}$  to be its child in a tree rooted at  $R_{Broot}$ , that AFFINITY RECORD is in conflict with TRILL distribution tree root determination and MUST be ignored.

If  $R_{Bridge\ R_{B1}}$  advertises an Affinity sub-TLV with an AFFINITY RECORD that asks for nickname  $R_{Bn}$  to be its child in any tree and  $R_{B1}$  is not adjacent to  $R_{Bn}$  nor does nickname  $R_{Bn}$  identify  $R_{B1}$  itself, that AFFINITY RECORD is in conflict with the campus topology and MUST be ignored.

If different RBridges advertise Affinity sub-TLVs that try to associate the same virtual RBridge as their child in the same tree or trees, those Affinity sub-TLVs are in conflict with each other for those trees. The nicknames of the conflicting RBridges are compared to identify which RBridge holds the nickname that is the highest priority to be a tree root, with the System ID as the tiebreaker.

The RBridge with the highest priority to be a tree root will retain the Affinity association. Other RBridges with lower priority to be a tree root MUST stop advertising their conflicting Affinity sub-TLVs, recalculate the multicast tree affinity allocation, and, if appropriate, advertise new non-conflicting Affinity sub-TLVs.

Similarly, remote RBridges MUST honor the Affinity sub-TLV from the RBridge with the highest priority to be a tree root (using System ID as the tiebreaker in the event of conflicting priorities) and ignore the conflicting Affinity sub-TLV entries advertised by the RBridges with lower priorities to be tree roots.

#### 5.4. Ingress Multi-Destination Forwarding

If there is at least one tree on which RBv has affinity via RBk, then RBk performs the following operations for multi-destination frames received from a CE node:

1. Flood to locally attached CE nodes subjected to VLAN and multicast pruning.
2. Ingress (by encapsulating with a TRILL Header) and set the Ingress Nickname field of the TRILL Header to RBv (the nickname of the virtual RBridge).
3. Forward to one of the distribution trees, Tree x, in which RBv is associated with RBk.

##### 5.4.1. Forwarding when $n < k$

If there is no tree on which RBv can claim affinity via RBk (probably because the number of trees (n) built is less than the number of RBridges (k) announcing the Affinity sub-TLV), then RBk MUST fall back to one of the following:

1. This RBridge (RBk) should stop forwarding frames from the CE nodes and should mark its port towards those CE nodes as disabled. This will prevent the CE nodes from forwarding data to this RBridge. Thus, the CE nodes will only use those RBridges that have been assigned a tree.

-OR-

2. This RBridge tunnels multi-destination frames received from attached native devices to an RBridge RBy that has an assigned tree. The tunnel destination should forward it to the TRILL network and also to its local access links. (The mechanism of tunneling and handshaking between the tunnel source and destination are out of scope for this specification and may be addressed in other documents, such as [ChannelTunnel].)

The above fallback options may be specific to the active-active forwarding scenario. However, as stated above, the Affinity sub-TLV may be used in other applications. In such an event, the application SHOULD specify applicable fallback options.

## 5.5. Egress Multi-Destination Forwarding

### 5.5.1. Traffic Arriving on an Assigned Tree to RBk-RBv

Multi-destination frames arriving at RBk on a Tree x, where RBk has announced the affinity of RBv via x, MUST be forwarded to CE members of RBv that are in the frame's VLAN. Forwarding to other end nodes and RBridges that are not part of the network represented by the virtual RBridge nickname (RBv) MUST follow the forwarding rules specified in [RFC6325].

### 5.5.2. Traffic Arriving on Other Trees

Multi-destination frames arriving at RBk on a Tree y, where RBk has not announced the affinity of RBv via y, MUST NOT be forwarded to CE members of RBv. Forwarding to other end nodes and RBridges that are not part of the network represented by the virtual RBridge nickname (RBv) MUST follow the forwarding rules specified in [RFC6325].

## 5.6. Failure Scenarios

The failure recovery algorithm below is presented only as a guideline. An active-active edge group MAY use other failure recovery algorithms; it is recommended that only one algorithm be used in an edge group at a time. Details of such algorithms are outside the scope of this document.

### 5.6.1. Edge RBridge RBk Failure

Each of the member RBridges of a given virtual RBridge edge group is aware of its member RBridges through configuration, LSP advertisements, or some other method.

Member RBridges detect nodal failure of a member RBridge through IS-IS LSP advertisements or lack thereof.

Upon detecting a member failure, each of the member RBridges of the RBv edge group start recovery timer  $T_{rec}$  for failed RBridge  $R_{Bi}$ . If the previously failed RBridge  $R_{Bi}$  has not recovered after the expiry of timer  $T_{rec}$ , member RBridges perform the distribution tree assignment algorithm specified in Section 5.1. Each of the member RBridges re-advertises the Affinity sub-TLV with the new tree assignment. This action causes the campus to update the tree calculation with the new assignment.

$R_{Bi}$ , upon startup, advertises its presence through IS-IS LSPs and starts a timer  $T_i$ . Other member RBridges of the edge group, detecting the presence of  $R_{Bi}$ , start a timer  $T_j$ .

Upon expiry of timer  $T_j$ , other member RBridges recalculate the multi-destination tree assignment and advertise the related trees using the Affinity sub-TLV. Upon expiry of timer  $T_i$ ,  $R_{Bi}$  recalculates the multi-destination tree assignment and advertises the related trees using the Affinity sub-TLV.

If the new RBridge in the edge group calculates trees and starts to use one or more of them before the existing RBridges in the edge group recalculate, there could be duplication of packets (for example, more than one edge group RBridge could decapsulate and forward a multi-destination frame on links into the active-active group) or loss of packets (for example, due to the Reverse Path Forwarding check, in the rest of the campus, if two edge group RBridges are trying to forward on the same tree, those from one will be discarded). Alternatively, if the new RBridge in the edge group calculates trees and starts to use one or more of them after the existing RBridges recalculate, there could be loss of data due to frames arriving at the new RBridge being black-holed. Timers  $T_i$  and  $T_j$  should be initialized to values designed to minimize these problems, keeping in mind that, in general, duplication of packets is a more serious problem than dropped packets. It is RECOMMENDED that  $T_j$  be less than  $T_i$ , and a reasonable default is 1/2 of  $T_i$ .

### 5.7. Backward Compatibility

Implementations MUST support a backward compatibility modes to interoperate with "pre-Affinity sub-TLV" RBridges in the network. Such backward compatibility operations MAY include, but are not limited to, tunneling and/or active-standby modes of operation. It should be noted that tunneling would require silicon changes. However, CMT in itself is a software upgrade.

Example:

Step 1. Stop using the virtual RBridge nickname for traffic ingressing from CE nodes.

Step 2. Stop performing active-active forwarding. Fall back to active standby forwarding, based on locally defined policies. The definition of such policies is outside the scope of this document and may be addressed in other documents.

## 6. Security Considerations

In general, the RBridges in a campus are trusted routers, and the authenticity of their link-state information (LSPs) and link-local PDUs (e.g., Hellos) can be enforced using regular IS-IS security mechanisms [IS-IS] [RFC5310]. This includes authenticating the contents of the PDUs used to transport Affinity sub-TLVs.

The particular security considerations involved with different applications of the Affinity sub-TLV will be covered in the document(s) specifying those applications.

For general TRILL security considerations, see [RFC6325].

## 7. IANA Considerations

This document serves as the reference for the registration of "Affinity sub-TLV support" (bit 0) in the "TRILL-VER Sub-TLV Capability Flags" registry.

This document mentions the registration of "AFFINITY" (value 17) in the "Sub-TLVs for TLV 144" registry, but it is intentionally not listed as a reference for that registration; the reference remains [RFC7176].

## 8. References

### 8.1. Normative References

- [IS-IS] International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, November 2002.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<http://www.rfc-editor.org/info/rfc5310>>.
- [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>>.
- [RFC6439] Perlman, R., Eastlake, D., Li, Y., Banerjee, A., and F. Hu, "Routing Bridges (RBridges): Appointed Forwarders", RFC 6439, DOI 10.17487/RFC6439, November 2011, <<http://www.rfc-editor.org/info/rfc6439>>.
- [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>>.
- [RFC7176] Eastlake 3rd, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 7176, DOI 10.17487/RFC7176, May 2014, <<http://www.rfc-editor.org/info/rfc7176>>.

- [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>>.
- [RFC7781] Zhai, H., Senevirathne, T., Perlman, R., Zhang, M., and Y. Li, "Transparent Interconnection of Lots of Links (TRILL): Pseudo-Nickname for Active-Active Access", RFC 7781, DOI 10.17487/RFC7781, February 2016, <<http://www.rfc-editor.org/info/rfc7781>>.

## 8.2. Informative References

- [802.1AX] IEEE, "IEEE Standard for Local and metropolitan area networks - Link Aggregation", IEEE Std 802.1AX-2014, DOI 10.1109/IEEESTD.2014.7055197, December 2014.
- [ChannelTunnel] Eastlake 3rd, D., Umair, M., and Y. Li, "TRILL: RBridge Channel Tunnel Protocol", Work in Progress, draft-ietf-trill-channel-tunnel-07, August 2015.
- [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>>.

## Acknowledgments

The authors wish to extend their appreciations towards individuals who volunteered to review and comment on the work presented in this document and who provided constructive and critical feedback. Specific acknowledgments are due for Anoop Ghanwani, Ronak Desai, Gayle Noble, and Varun Shah. Very special thanks to Donald Eastlake for his careful review and constructive comments.

## Contributors

The work in this document is a result of many passionate discussions and contributions from the following individuals, listed in alphabetical order by their first names:

Ayan Banerjee, Dinesh Dutt, Donald Eastlake, Hongjun Zhai, Mingui Zhang, Radia Perlman, Sam Aldrin, and Shivakumar Sundaram.

## Authors' Addresses

Tissa Senevirathne  
Consultant

Email: [tsenevir@gmail.com](mailto:tsenevir@gmail.com)

Janardhanan Pathangi  
Dell/Force10 Networks  
Olympia Technology Park  
Guindy Chennai 600 032  
India

Phone: +91-44-42208400  
Email: [Pathangi\\_Janardhanan@Dell.com](mailto:Pathangi_Janardhanan@Dell.com)

Jon Hudson  
Brocade  
130 Holger Way  
San Jose, CA 95134  
United States

Phone: +1-408-333-4062  
Email: [jon.hudson@gmail.com](mailto:jon.hudson@gmail.com)