

Internet Engineering Task Force (IETF)
Request for Comments: 7781
Category: Standards Track
ISSN: 2070-1721

H. Zhai
JIT
T. Senevirathne
Consultant
R. Perlman
EMC
M. Zhang
Y. Li
Huawei Technologies
February 2016

Transparent Interconnection of Lots of Links (TRILL):
Pseudo-Nickname for Active-Active Access

Abstract

The IETF TRILL (Transparent Interconnection of Lots of Links) protocol provides support for flow-level multipathing for both unicast and multi-destination traffic in networks with arbitrary topology. Active-active access at the TRILL edge is the extension of these characteristics to end stations that are multiply connected to a TRILL campus as discussed in RFC 7379. In this document, the edge RBridge (Routing Bridge, or TRILL switch) group providing active-active access to such an end station is represented as a virtual RBridge. Based on the concept of the virtual RBridge, along with its pseudo-nickname, this document specifies a method for TRILL active-active access by such end stations.

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/rfc7781>.

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 | 4 |
| 1.1. Terminology and Acronyms | 6 |
| 2. Overview | 7 |
| 3. Virtual RBridge and Its Pseudo-Nickname | 9 |
| 4. Auto-Discovery of Member RBridges | 10 |
| 4.1. Discovering Member RBridge for an RBv | 11 |
| 4.2. Selection of Pseudo-Nickname for an RBv | 13 |
| 5. Distribution Trees and Designated Forwarder | 14 |
| 5.1. Different Trees for Different Member RBridges | 15 |
| 5.2. Designated Forwarder for Member RBridges | 16 |
| 5.3. Ingress Nickname Filtering | 18 |
| 6. TRILL Traffic Processing | 19 |
| 6.1. Ingressing Native Frames | 19 |
| 6.2. Egressing TRILL Data Packets | 20 |
| 6.2.1. Unicast TRILL Data Packets | 20 |
| 6.2.2. Multi-Destination TRILL Data Packets | 21 |
| 7. MAC Information Synchronization in Edge Group | 22 |
| 8. Member Link Failure in an RBv | 23 |
| 8.1. Link Protection for Unicast Frame Egressing | 24 |
| 9. TLV Extensions for Edge RBridge Group | 24 |
| 9.1. PN-LAALP-Membership APPsub-TLV | 24 |
| 9.2. PN-RBv APPsub-TLV | 26 |
| 9.3. PN-MAC-RI-LAALP Boundary APPsub-TLVs | 27 |
| 9.4. LAALP IDs | 29 |
| 10. OAM Packets | 29 |
| 11. Configuration Consistency | 29 |
| 12. Security Considerations | 30 |
| 13. IANA Considerations | 31 |
| 14. References | 31 |
| 14.1. Normative References | 31 |
| 14.2. Informative References | 33 |
| Acknowledgments | 34 |
| Contributors | 34 |
| Authors' Addresses | 35 |

1. Introduction

The IETF TRILL (Transparent Interconnection of Lots of Links) protocol [RFC6325] provides optimal pair-wise data frame forwarding without configuration, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. TRILL accomplishes this by using IS-IS [IS-IS] [RFC7176] link-state routing and encapsulating traffic using a header that includes a Hop Count. Devices that implement TRILL are called RBridges (Routing Bridges) or TRILL switches.

In the base TRILL protocol, an end node can be attached to the TRILL campus via a point-to-point link or a shared link such as a bridged LAN (Local Area Network). Although there might be more than one edge RBridge on a shared link, to avoid potential forwarding loops, one and only one of the edge RBridges is permitted to provide forwarding service for end-station traffic in each VLAN (Virtual LAN). That RBridge is referred to as the Appointed Forwarder (AF) for that VLAN on the link [RFC6325] [RFC6439]. However, in some practical deployments, to increase the access bandwidth and reliability, an end station might be multiply connected to several edge RBridges, and all of the uplinks are handled via a Local Active-Active Link Protocol (LAALP [RFC7379]) such as Multi-Chassis Link Aggregation (MC-LAG) or Distributed Resilient Network Interconnect (DRNI) [802.1AX]. In this case, it is required that traffic can be ingressed into, and egressed from, the TRILL campus by any of the RBridges for each given VLAN. These RBridges constitute an Active-Active Edge (AAE) RBridge group.

With an LAALP, traffic with the same VLAN and source Media Access Control (MAC) address but belonging to different flows will frequently be sent to different member RBridges of the AAE group and then ingressed into the TRILL campus. When an egress RBridge receives such TRILL Data packets ingressed by different RBridges, it learns different correspondences between a {Data Label and MAC address} and nickname continuously when decapsulating the packets if it has data-plane address learning enabled. This issue is known as "MAC address flip-flopping"; it makes most TRILL switches behave badly and causes the returning traffic to reach the destination via different paths, resulting in persistent reordering of the frames. In addition to this issue, other issues, such as duplicate egressing and loopback of multi-destination frames, may also disturb an end station multiply connected to the member RBridges of an AAE group [RFC7379].

This document addresses the AAE issues of TRILL by specifying how members of an edge RBridge group can be represented by a virtual RBridge (RBv) and assigned a pseudo-nickname. A member RBridge of such a group uses a pseudo-nickname instead of its own nickname as

the ingress RBridge nickname when ingressing frames received on attached LAALP links. Other methods are possible: for example, the specification in this document and the specification in [RFC7782] could be simultaneously deployed for different AAE groups in the same campus. If the method defined in [RFC7782] is used, edge TRILL switches need to support the capability indicated by the Capability Flags APPsub-TLV as specified in Section 4.2 of [RFC7782]. If the method defined in this document is adopted, all TRILL switches need to support the Affinity sub-TLV defined in [RFC7176] and [RFC7783]. For a TRILL campus that deploys both of these AAE methods, TRILL switches are required to support both methods. However, it is desirable to only adopt one method in a TRILL campus so that the operating expense, complexity of troubleshooting, etc., can be reduced.

The main body of this document is organized as follows:

- o Section 2 provides an overview of the TRILL active-active access issues and the reason that a virtual RBridge (RBv) is used to resolve the issues.
- o Section 3 describes the concept of a virtual RBridge (RBv) and its pseudo-nickname.
- o Section 4 describes how edge RBridges can support an RBv automatically and get a pseudo-nickname for the RBv.
- o Section 5 discusses how to protect multi-destination traffic against disruption due to Reverse Forwarding Path (RPF) check failure, duplication, forwarding loops, etc.
- o Section 6 covers the special processing of native frames and TRILL Data packets at member RBridges of an RBv (also referred to as an Active-Active Edge (AAE) RBridge group).
- o Section 7 describes the MAC information synchronization among the member RBridges of an RBv.
- o Section 8 discusses protection against downlink failure at a member RBridge.
- o Section 9 lists the necessary TRILL code points and data structures for a pseudo-nickname AAE RBridge group.

1.1. Terminology and Acronyms

This document uses the acronyms and terms defined in [RFC6325] and [RFC7379], as well as the following additional acronyms:

AAE: Active-active Edge RBridge group. A group of edge RBridges to which at least one Customer Equipment (CE) node is multiply attached with an LAALP. AAE is also referred to as "edge group" or "virtual RBridge" in this document.

Campus: A TRILL network consisting of TRILL switches, links, and possibly bridges bounded by end stations and IP routers. For TRILL, there is no "academic" implication in the name "campus".

CE: Customer Equipment (end station or bridge). The device can be either physical or virtual equipment.

Data Label: VLAN or Fine-Grained Label (FGL).

DF: Designated Forwarder.

DRNI: Distributed Resilient Network Interconnect. A link aggregation specified in [802.1AX] that can provide an LAALP between (a) one, two, or three CEs and (b) two or three RBridges.

E-L1FS: Extended Level 1 Flooding Scope [RFC7356].

ESADI: End-Station Address Distribution Information.

FGL: Fine-Grained Labeling or Fine-Grained Labeled or Fine-Grained Label [RFC7172].

LAALP: Local Active-Active Link Protocol [RFC7379], e.g., MC-LAG or DRNI.

MC-LAG: Multi-Chassis Link Aggregation. Proprietary extensions of Link Aggregation [802.1AX] that can provide an LAALP between one CE and two or more RBridges.

OE-flag: A flag used by a member RBridge of a given LAALP to tell other edge RBridges of this LAALP whether this LAALP is willing to share an RBv with other LAALPs that multiply attach to the same set of edge RBridges as the given LAALP does. When this flag for an LAALP is 1, it means that the LAALP needs to be served by an RBv by itself and is not willing to share, that is, it should Occupy an RBv Exclusively (OE).

RBv: Virtual RBridge. An alias for "active-active edge RBridge group" in this document.

vDRB: The Designated RBridge in an RBv. It is responsible for deciding the pseudo-nickname for the RBv.

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. Overview

To minimize impact during failures and maximize available access bandwidth, Customer Equipment (referred to as "CE" in this document) may be multiply connected to the TRILL campus via multiple edge RBridges.

Figure 1 shows such a typical deployment scenario, where CE1 attaches to RB1, RB2, ... RBk and treats all of the uplinks as an LAALP bundle. RB1, RB2, ... RBk then constitute an AAE RBridge group for CE1 in this LAALP. Even if a member RBridge or an uplink fails, CE1 will still get frame forwarding service from the TRILL campus if there are still member RBridges and uplinks available in the AAE group. Furthermore, CE1 can make flow-based load balancing across the available member links of the LAALP bundle in the AAE group when it communicates with other CEs across the TRILL campus [RFC7379].

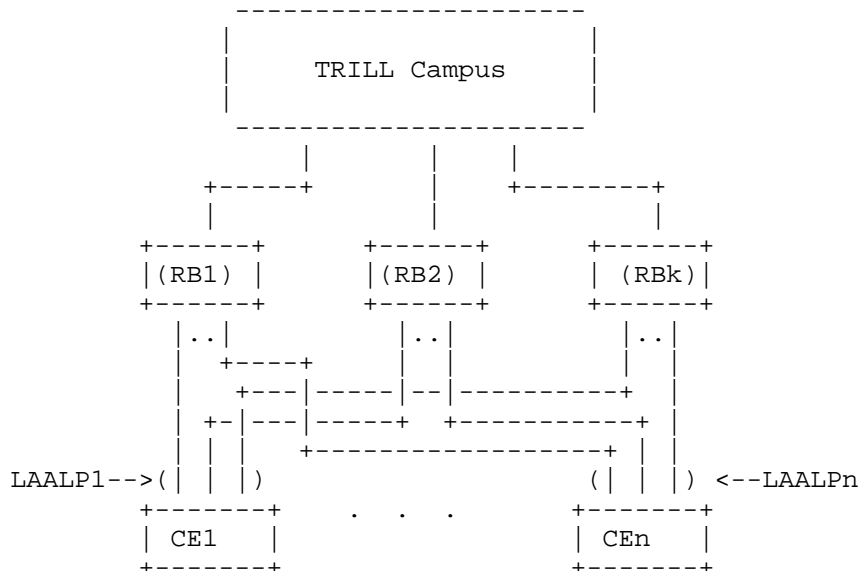


Figure 1: Active-Active Connection to TRILL Edge RBridges

By design, an LAALP (say LAALP1) does not forward packets received on one member port to other member ports. As a result, the TRILL Hello messages sent by one member RBridge (say RB1) via a port to CE1 will not be forwarded to other member RBridges by CE1. That is to say, member RBridges will not see each other's Hellos via the LAALP. So, every member RBridge of LAALP1 thinks of itself as Appointed Forwarder for all VLANs enabled on an LAALP1 link and can ingress/egress frames simultaneously in these VLANs [RFC6439].

The simultaneous flow-based ingressing/egressing can cause some problems. For example, simultaneous egressing of multi-destination traffic by multiple member RBridges will result in frame duplication at CE1 (see Section 3.1 of [RFC7379]); simultaneous ingressing of frames originated by CE1 for different flows in the same VLAN with the same source MAC address will result in MAC address flip-flopping at remote egress RBridges that have data-plane address learning enabled (see Section 3.3 of [RFC7379]). The flip-flopping would in turn cause packet reordering in reverse traffic.

Edge RBridges learn correspondences between a {Data Label and MAC address} and nickname by default when decapsulating TRILL Data packets (see Section 4.8.1 of [RFC6325], as updated by [RFC7172]). Assuming that the default data-plane learning is enabled at edge RBridges, MAC address flip-flopping can be solved by using a virtual RBridge together with its pseudo-nickname. This document specifies a way to do so.

3. Virtual RBridge and Its Pseudo-Nickname

A virtual RBridge (RBv) represents a group of edge RBridges to which at least one CE is multiply attached using an LAALP. More precisely, it represents a group of ports on the edge RBridges providing end-station service and the service provided to the CE(s) on these ports, through which the CE(s) is multiply attached to the TRILL campus using LAALP(s). Such end-station service ports are called RBv ports; in contrast, other access ports at edge RBridges are called regular access ports in this document. RBv ports are always LAALP connecting ports, but not vice versa (see Section 4.1). For an edge RBridge, if one or more of its end-station service ports are ports of an RBv, that RBridge is a member RBridge of that RBv.

For the convenience of description, a virtual RBridge is also referred to as an Active-Active Edge (AAE) group in this document. In the TRILL campus, an RBv is identified by its pseudo-nickname, which is different from any RBridge's regular nickname(s). An RBv has one and only one pseudo-nickname. Each member RBridge (say RB1, RB2 ..., RBk) of an RBv (say RBvn) advertises RBvn's pseudo-nickname using a Nickname sub-TLV in its TRILL IS-IS LSP (Link State PDU) [RFC7176] and SHOULD do so with maximum priority of use (0xFF), along with their regular nickname(s). (Maximum priority is recommended to avoid the disruption to an AAE group that would occur if the nickname were taken away by a higher-priority RBridge.) Then, from these LSPs, other RBridges outside the AAE group know that RBvn is reachable through RB1 to RBk.

A member RBridge (say RBi) loses its membership in RBvn when its last port in RBvn becomes unavailable due to failure, reconfiguration, etc. RBi then removes RBvn's pseudo-nickname from its LSP and distributes the updated LSP as usual. From those updated LSPs, other RBridges know that there is no path to RBvn through RBi now.

When member RBridges receive native frames on their RBv ports and decide to ingress the frames into the TRILL campus, they use that RBv's pseudo-nickname instead of their own regular nicknames as the ingress nickname to encapsulate them into TRILL Data packets. So, when these packets arrive at an egress RBridge, even if they are originated by the same end station in the same VLAN but ingressed by

different member RBridges, no address flip-flopping is observed on the egress RBridge when decapsulating these packets. (When a member RBridge of an AAE group ingresses a frame from a non-RBv port, it still uses its own regular nickname as the ingress nickname.)

Since an RBv is not a physical node and no TRILL frames are forwarded between its ports via an LAALP, pseudonode LSP(s) MUST NOT be created for an RBv. An RBv cannot act as a root when constructing distribution trees for multi-destination traffic, and its pseudo-nickname is ignored when determining the distribution tree root for the TRILL campus [RFC7783]. So, the tree root priority of the RBv's nickname MUST be set to 0, and this nickname MUST NOT be listed in the "s" nicknames (see Section 4.5 of [RFC6325]) by the RBridge holding the highest-priority tree root nickname.

NOTE: In order to reduce the consumption of nicknames, especially in a large TRILL campus with lots of RBridges and/or active-active accesses, when multiple CEs attach to exactly the same set of edge RBridges via LAALPs, those edge RBridges should be considered a single RBv with a single pseudo-nickname.

4. Auto-Discovery of Member RBridges

Edge RBridges connected to a CE via an LAALP can automatically discover each other with minimal configuration through the exchange of LAALP connection information.

From the perspective of edge RBridges, a CE that connects to edge RBridges via an LAALP can be identified by the ID of the LAALP that is unique across the TRILL campus (for example, the MC-LAG or DRNI System ID [802.1AX]), which is referred to as an LAALP ID in this document. On each such edge RBridge, the access port to such a CE is associated with an LAALP ID for the CE. An LAALP is considered valid on an edge RBridge only if the RBridge still has an operational downlink to that LAALP. For such an edge RBridge, it advertises a list of LAALP IDs for its valid local LAALPs to other edge RBridges via its E-L1FS FS-LSP(s) [RFC7356] [RFC7780]. Based on the LAALP IDs advertised by other RBridges, each RBridge can know which edge RBridges could constitute an AAE group (see Section 4.1 for more details). One RBridge is then elected from the group to allocate an available nickname (the pseudo-nickname) for the group (see Section 4.2 for more details).

4.1. Discovering Member RBridge for an RBv

Take Figure 2 as an example, where CE1 and CE2 multiply attach to RB1, RB2, and RB3 via LAALP1 and LAALP2, respectively; CE3 and CE4 attach to RB3, and RB4 via LAALP3 and LAALP4, respectively. Assume that LAALP3 is configured to occupy a virtual RBridge by itself.

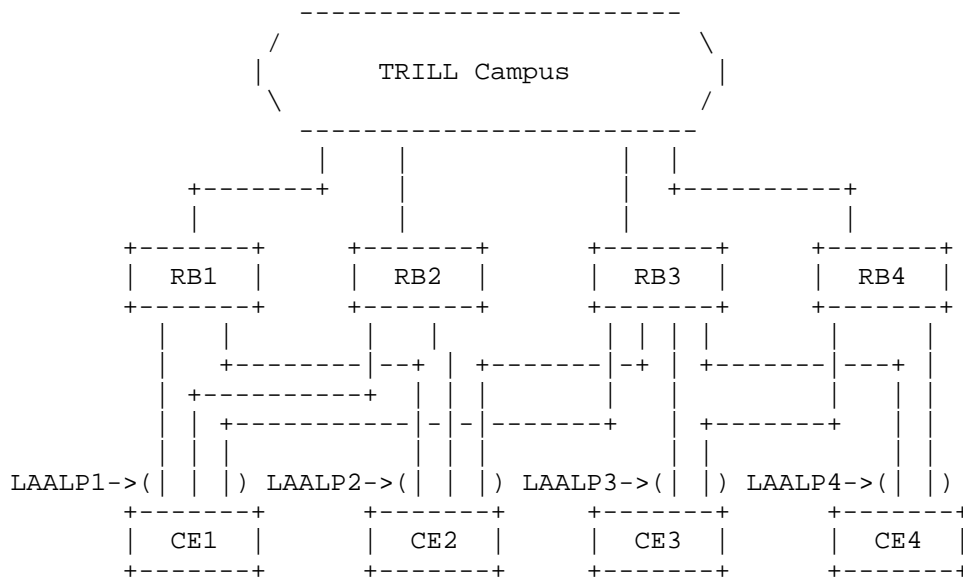


Figure 2: Different LAALPs to TRILL Campus

RB1 and RB2 advertise {LAALP1, LAALP2} in the PN-LAALP-Membership APPsub-TLV (see Section 9.1 for more details) via their TRILL E-L1FS FS-LSPs, respectively; RB3 announces {LAALP1, LAALP2, LAALP3, LAALP4}, and RB4 announces {LAALP3, LAALP4}, respectively.

An edge RBridge is called an "LAALP related RBridge" if it has at least one LAALP configured on an access port. On receipt of the PN-LAALP-Membership APPsub-TLVs, RBn ignores them if it is not an LAALP related RBridge; otherwise, RBn SHOULD use the LAALP information contained in the sub-TLVs, along with its own PN-LAALP-Membership APPsub-TLVs, to decide which RBv(s) it should join and which edge RBridges constitute each such RBv. Based on the information received, each of the four RBridges knows the following:

| LAALP ID | OE-flag | Set of edge RBridges |
|----------|---------|----------------------|
| ----- | ----- | ----- |
| LAALP1 | 0 | {RB1, RB2, RB3} |
| LAALP2 | 0 | {RB1, RB2, RB3} |
| LAALP3 | 1 | {RB3, RB4} |
| LAALP4 | 0 | {RB3, RB4} |

where the OE-flag indicates whether a given LAALP is willing to share an RBv with other LAALPs that multiply attach to the same set of edge RBridges as the given LAALP does.

For an LAALP (for example, LAALP3), if its OE-flag is one, it means that LAALP3 does not want to share, so it MUST Occupy an RBv Exclusively (OE). Support of OE is optional. RBridges that do not support OE ignore the OE-flag and act as if it were zero (see Section 11 ("Configuration Consistency")).

Otherwise, the LAALP (for example, LAALP1) will share an RBv with other LAALPs if possible. By default, this flag is set to zero. For an LAALP, this flag is considered 1 if any edge RBridge advertises it as (value) 1 (see Section 9.1).

In the above table, there might be some LAALPs that attach to a single RBridge due to misconfiguration or link failure, etc. Those LAALPs are considered to be invalid entries. Each of the LAALP related edge RBridges then performs the following algorithm to decide which valid LAALPs can be served by an RBv.

Step 1: Take all the valid LAALPs that have their OE-flags set to 1 out of the table and create an RBv for each such LAALP.

Step 2: Sort the valid LAALPs left in the table in descending order based on the number of RBridges in their associated set of multihomed RBridges. If several LAALPs have the same number of RBridges, these LAALPs are then ordered in ascending order in the proper places of the table, based on their LAALP IDs considered to be unsigned integers. (For example, in the above

table, both LAALP1 and LAALP2 have three member RBridges, assuming that the LAALP1 ID is smaller than the LAALP2 ID, so LAALP1 is followed by LAALP2 in the ordered table.)

Step 3: Take the first valid LAALP (say LAALP_i) with the maximum set of RBridges, say S_i, out of the table and create a new RBv (say RBv_i) for it.

Step 4: Walk through the remaining valid LAALPs in the table one by one, pick up all the valid LAALPs whose sets of multi-homed RBridges contain exactly the same RBridges as that of LAALP_i, and take them out of the table. Then, appoint RBv_i as the servicing RBv for those LAALPs.

Step 5: Repeat Steps 3 and 4 for any LAALPs left, until all the valid entries in the table are associated with an RBv.

After performing the above steps, all the four RBridges know that LAALP3 is served by an RBv, say RBv1, which has RB3 and RB4 as member RBridges; LAALP1 and LAALP2 are served by another RBv, say RBv2, which has RB1, RB2, and RB3 as member RBridges; and LAALP4 is served by RBv3, which has RB3 and RB4 as member RBridges, shown as follows:

| RBv | Serving LAALPs | Member RBridges |
|-------|------------------|-----------------|
| ----- | ----- | ----- |
| RBv1 | {LAALP3} | {RB3, RB4} |
| RBv2 | {LAALP1, LAALP2} | {RB1, RB2, RB3} |
| RBv3 | {LAALP4} | {RB3, RB4} |

In each RBv, one of the member RBridges is elected as the vDRB (referred to in this document as the Designated RBridge of the RBv). Then, this RBridge picks up an available nickname as the pseudo-nickname for the RBv and announces it to all other member RBridges of the RBv via its TRILL E-L1FS FS-LSPs (refer to Section 9.2 for the relative extended sub-TLVs).

4.2. Selection of Pseudo-Nickname for an RBv

As described in Section 3, in the TRILL campus, an RBv is identified by its pseudo-nickname. In an AAE group, one member RBridge is elected for the duty of selecting a pseudo-nickname for this RBv; this RBridge will be the vDRB. The winner in the group is the RBridge with the largest IS-IS System ID considered to be an unsigned integer. Then, based on its TRILL IS-IS link-state database and the potential pseudo-nickname(s) reported in the PN-LAALP-Membership APPsub-TLVs by other member RBridges of this RBv (see Section 9.1 for more details), the vDRB selects an available nickname as the pseudo-nickname for this RBv and advertises it to the other RBridges

via its E-L1FS FS-LSP(s) (see Section 9.2 and [RFC7780]). Except as provided below, the selection of a nickname to use as the pseudo-nickname follows the usual TRILL rules given in [RFC6325], as updated by [RFC7780].

To reduce the traffic disruption caused by the changing of nicknames, if possible, the vDRB SHOULD attempt to reuse the pseudo-nickname recently used by the group when selecting nickname for the RBv. To help the vDRB to do so, each LAALP related RBridge advertises a reusing pseudo-nickname for each of its LAALPs in its PN-LAALP-Membership APPsub-TLV if it has used such a pseudo-nickname for that LAALP recently. Although it is up to the implementation of the vDRB as to how to treat the reusing pseudo-nicknames, the following are RECOMMENDED:

- o If there are multiple available reusing pseudo-nicknames that are reported by all the member RBridges of some LAALPs in this RBv, the available one that is reported by the largest number of such LAALPs is chosen as the pseudo-nickname for this RBv. If a tie exists, the reusing pseudo-nickname with the smallest value considered to be an unsigned integer is chosen.
- o If only one reusing pseudo-nickname is reported, it SHOULD be chosen if available.

If there is no available reusing pseudo-nickname reported, the vDRB selects a nickname by its usual method.

The selected pseudo-nickname is then announced by the vDRB to other member RBridges of this RBv in the PN-RBv APPsub-TLV (see Section 9.2).

5. Distribution Trees and Designated Forwarder

In an AAE group, as each of the member RBridges thinks it is the Appointed Forwarder for VLAN x, without changes made for active-active connection support, they would all ingress frames into, and egress frames from, the TRILL campus for all VLANs. For multi-destination frames, more than one member RBridge ingressing them may cause some of the resulting TRILL Data packets to be discarded due to failure of the Reverse Path Forwarding (RPF) check on other RBridges; for multi-destination traffic, more than one RBridge egressing it may cause local CE(s) to receive duplicate frames. Furthermore, in an AAE group, a multi-destination frame sent by a CE (say CEi) may be ingressed into the TRILL campus by one member RBridge, and another member RBridge will then receive it from the TRILL campus and egress it to CEi; this will result in loopback of the frame for CEi. These problems are all described in [RFC7379].

In the following subsections, the first two issues are discussed in Sections 5.1 and 5.2, respectively; the third issue is discussed in Section 5.3.

5.1. Different Trees for Different Member RBridges

In TRILL, RBridges normally use distribution trees to forward multi-destination frames. (Under some circumstances, they can be unicast, as specified in [RFC7172].) An RPF check, along with other types of checks, is used to avoid temporary multicast loops during topology changes (Section 4.5.2 of [RFC6325]). The RPF check mechanism only accepts a multi-destination frame ingressed by an RBridge (say R_{Bi}) and forwarded on a distribution tree if it arrives at another RBridge (say R_{Bn}) on the expected port. If the frame arrives on any other port, the frame **MUST** be dropped.

To avoid address flip-flopping on remote RBridges, member RBridges use the RBv's pseudo-nickname instead of their regular nicknames as the ingress nickname to ingress native frames, including multi-destination frames. From the view of other RBridges, these frames appear as if they were ingressed by the RBv. When multi-destination frames of different flows are ingressed by different member RBridges of an RBv and forwarded along the same distribution tree, they may arrive at R_{Bn} on different ports. Some of them will violate the RPF check principle at R_{Bn} and be dropped, which will result in lost traffic.

In an RBv, if a different member RBridge uses different distribution trees to ingress multi-destination frames, the RPF check violation issue can be fixed. The Coordinated Multicast Trees (CMT) document [RFC7783] proposes such an approach and makes use of the Affinity sub-TLV defined in [RFC7176] to tell other RBridges which trees a member RBridge (say R_{Bi}) may choose when ingressing multi-destination frames; all RBridges in the TRILL campus can then calculate RPF check information for R_{Bi} on those trees, taking the tree affinity information into account [RFC7783].

This document uses the approach proposed in [RFC7783] to fix the RPF check violation issue. Please refer to [RFC7783] for more details regarding this approach.

5.2. Designated Forwarder for Member RBridges

Take Figure 3 as an example, where CE1 and CE2 are served by an RBv that has RB1 and RB2 as member RBridges. In VLAN x, the three CEs can communicate with each other.

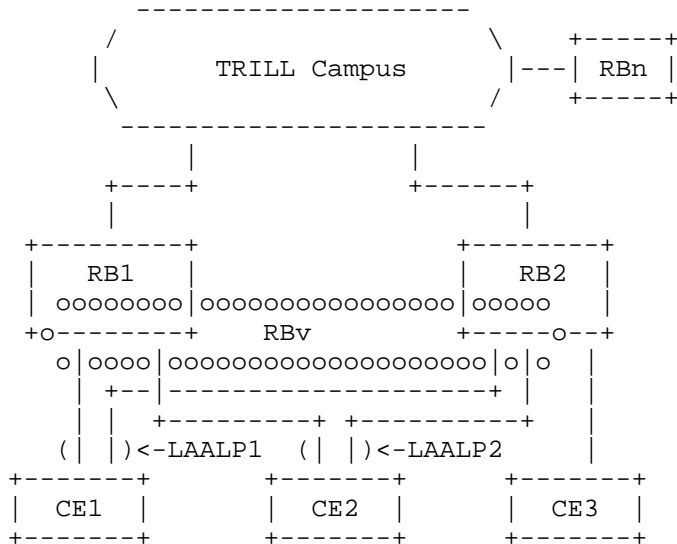


Figure 3: A Topology with Multihomed and Single-Homed CEs

When a remote RBridge (say RBn) sends a multi-destination TRILL Data packet in VLAN x (or the FGL that VLAN x maps to, if the packet is FGL), both RB1 and RB2 will receive it. As each of them thinks it is the Appointed Forwarder for VLAN x, without changes made for active-active connection support, they would both forward the frame to CE1/CE2. As a result, CE1/CE2 would receive duplicate copies of the frame through this RBv.

In another case, assume that CE3 is single-homed to RB2. When it transmits a native multi-destination frame onto link CE3-RB2 in VLAN x, the frame can be locally replicated to the ports to CE1/CE2, and also encapsulated into TRILL Data packet and ingressed into the TRILL campus. When the packet arrives at RB1 across the TRILL campus, it will be egressed to CE1/CE2 by RB1. CE1/CE2 then receives duplicate copies from RB1 and RB2.

In this document, the Designated Forwarder (DF) for a VLAN is introduced to avoid duplicate copies. The basic idea of the DF is to elect one RBridge per VLAN from an RBv to egress multi-destination TRILL Data traffic and replicate locally received multi-destination native frames to the CEs served by the RBv.

Note that the DF has an effect only on the egressing/replicating of multi-destination traffic. It has no effect on the ingressing, forwarding, or egressing of unicast frames. Furthermore, the DF check is performed only for RBv ports, not on regular access ports.

Each RBridge in an RBv elects a DF using the same algorithm; this guarantees that, per VLAN, the same RBridge is elected as the DF by all members of the RBv.

If we assume that there are m LAALPs and k member RBridges in an RBv, then (1) each LAALP is referred to as "LAALPi", where $0 \leq i < m$, and (2) each RBridge is referred to as "RBj", where $0 \leq j < k$. The DF election algorithm per VLAN is as follows:

Step 1: For LAALPi, sort all the RBridges in numerically ascending order based on $\text{SHA-256}(\text{System ID}_j \mid \text{LAALP ID}_i)$ considered to be an unsigned integer, where SHA-256 is the hash function specified in [RFC6234], "System ID_j" is the 6-byte IS-IS System ID of RBj, "|" means concatenation, and "LAALP ID_i" is the LAALP ID for LAALPi. The System ID and LAALP ID are considered to be byte strings. In the case of a tie, the tied RBridges are sorted in numerically ascending order by their System IDs considered to be unsigned integers.

Step 2: Each RBridge in the numerically sorted list is assigned a monotonically increasing number j , such that increasing number j corresponds to its position in the sorted list, i.e., the first RBridge (the one with the smallest $\text{SHA-256}(\text{System ID} \mid \text{LAALP ID})$) is assigned zero and the last is assigned $k-1$.

Step 3: For each VLAN ID n , choose the RBridge whose number equals $(n \bmod k)$ as the DF.

Step 4: Repeat Steps 1-3 for the remaining LAALPs until there is a DF per VLAN per LAALP in the RBv.

For any multi-destination native frames of VLAN x that are received, if RBi is an LAALP attached RBridge, there are three cases where RBi replicates the multi-destination frame, as follows:

- 1) Local replication of the frame to regular (non-AAE) access ports as per [RFC6325] (and [RFC7172] for FGL).
- 2) RBv ports associated with the same pseudo-nickname as that of the incoming port, no matter whether RBi is the DF for the frame's VLAN on the outgoing ports, except that the frame MUST NOT be replicated back to the incoming port. RBi cannot simply depend on the DF to forward the multi-destination frame back into the AAEs associated with the pseudo-nickname, as that would cause the source CE to get the frame back, which is a violation of basic Ethernet properties. The DF will not forward such a frame back into the AAE due to ingress nickname filtering as described in Section 5.3.
- 3) RBv ports on which RBi is the DF for the frame's VLAN while they are associated with different pseudo-nickname(s) than that of the incoming port.

For any multi-destination TRILL Data packets that are received, RBi MUST NOT egress it out of the RBv ports where it is not the DF for the frame's Inner.VLAN (or for the VLAN corresponding to the Inner.Label if the packet is an FGL one). Otherwise, whether or not to egress it out of such ports is further subject to the filtering check result of the frame's ingress nickname on these ports (see Section 5.3).

5.3. Ingress Nickname Filtering

As shown in Figure 3, CE1 may send multi-destination traffic in VLAN x to the TRILL campus via a member RBridge (say RB1). The traffic is then TRILL-encapsulated by RB1 and delivered through the TRILL campus to multi-destination receivers. RB2 may receive the traffic and egress it back to CE1 if it is the DF for VLAN x on the port to LAALP1. The traffic then loops back to CE1 (see Section 3.2 of [RFC7379]).

To fix the above issue, this document requires an ingress nickname filtering check. The idea is to check the ingress nickname of a multi-destination TRILL Data packet before egressing a copy of it out of an RBv port. If the ingress nickname matches the pseudo-nickname of the RBv (associated with the port), the filtering check should fail and the copy MUST NOT be egressed out of that RBv port. Otherwise, the copy is egressed out of that port if it has also

passed other checks, such as the Appointed Forwarder check described in Section 4.6.2.5 of [RFC6325] and the DF check described in Section 5.2.

Note that this ingress nickname filtering check has no effect on the multi-destination native frames that are received on access ports and replicated to other local ports (including RBv ports), since there is no ingress nickname associated with such frames. Furthermore, for the RBridge regular access ports, there is no pseudo-nickname associated with them, so no ingress nickname filtering check is required on those ports.

More details of data packet processing on RBv ports are given in the next section.

6. TRILL Traffic Processing

This section provides more details of native frame and TRILL Data packet processing as it relates to the RBv's pseudo-nickname.

6.1. Ingressing Native Frames

When RB1 receives a unicast native frame from one of its ports that has end-station service enabled, it processes the frame as described in Section 4.6.1.1 of [RFC6325], with the following exception:

- o If the port is an RBv port, RB1 uses the RBv's pseudo-nickname instead of one of its regular nickname(s) as the ingress nickname when doing TRILL encapsulation on the frame.

When RB1 receives a native multi-destination (broadcast, unknown unicast, or multicast) frame from one of its access ports (including regular access ports and RBv ports), it processes the frame as described in Section 4.6.1.2 of [RFC6325], with the following exceptions:

- o If the incoming port is an RBv port, RB1 uses the RBv's pseudo-nickname instead of one of its regular nickname(s) as the ingress nickname when doing TRILL encapsulation on the frame.

- o For the copies of the frame replicated locally to RBv ports, there are two cases, as follows:
 - If the outgoing port(s) is associated with the same pseudo-nickname as that of the incoming port but not with the same LAALP as the incoming port, the copies are forwarded out of that outgoing port(s) after passing the Appointed Forwarder check for the frame's VLAN. That is to say, the copies are processed on such port(s), as discussed in Section 4.6.1.2 of [RFC6325].
 - Else, the Designated Forwarder (DF) check is also made on the outgoing ports for the frame's VLAN after the Appointed Forwarder check, and the copies are not output through any ports that failed the DF check (i.e., RB1 is not the DF for the frame's VLAN on the ports). Otherwise, the copies are forwarded out of the outgoing ports that pass both the Appointed Forwarder check and the DF check (see Section 5.2).

For any such frames received, the MAC address information learned by observing it, together with the LAALP ID of the incoming port, SHOULD be shared with other member RBridges in the group (see Section 7).

6.2. Egressing TRILL Data Packets

This section describes egress processing of the TRILL Data packets received on an RBv member RBridge (say RBn). Section 6.2.1 describes the egress processing of unicast TRILL Data packets, and Section 6.2.2 specifies the egressing of multi-destination TRILL Data packets.

6.2.1. Unicast TRILL Data Packets

When receiving a unicast TRILL Data packet, RBn checks the egress nickname in the TRILL Header of the packet. If the egress nickname is one of RBn's regular nicknames, the packet is processed as defined in Section 4.6.2.4 of [RFC6325].

If the egress nickname is the pseudo-nickname of a local RBv, RBn is responsible for learning the source MAC address, unless data-plane learning has been disabled. The learned {Inner.MacSA, Data Label, ingress nickname} triplet SHOULD be shared within the AAE group as described in Section 7.

The packet is then decapsulated to its native form. The Inner.MacDA and Data Label are looked up in RBn's local forwarding tables, and one of the three following cases will occur. RBn uses the first case that applies and ignores the remaining cases:

- o If the destination end station identified by the Inner.MacDA and Data Label is on a local link, the native frame is sent onto that link with the VLAN from the Inner.VLAN or VLAN corresponding to the Inner.Label if the packet is FGL.
- o Else if RBn can reach the destination through another member RBridge (say RBk), it tunnels the native frame to RBk by re-encapsulating it into a unicast TRILL Data packet and sends it to RBk. RBn uses RBk's regular nickname instead of the pseudo-nickname as the egress nickname for the re-encapsulation, and the ingress nickname remains unchanged (somewhat similar to Section 2.4.2.1 of [RFC7780]). If the Hop Count value of the packet is too small for it to reach RBk safely, RBn SHOULD increase that value properly in doing the re-encapsulation. (NOTE: When receiving that re-encapsulated TRILL Data packet, as the egress nickname of the packet is RBk's regular nickname rather than the pseudo-nickname of a local RBv, RBk will process it per Section 4.6.2.4 of [RFC6325] and will not re-forward it to another RBridge.)
- o Else, RBn does not know how to reach the destination; it sends the native frame out of all the local ports on which it is Appointed Forwarder for the Inner.VLAN (or Appointed Forwarder for the VLAN into which the Inner.Label maps on that port for an FGL TRILL Data packet [RFC7172]).

6.2.2. Multi-Destination TRILL Data Packets

When RB1 receives a multi-destination TRILL Data Packet, it checks and processes the packet as described in Section 4.6.2.5 of [RFC6325], with the following exception:

- o On each RBv port where RBn is the Appointed Forwarder for the packet's Inner.VLAN (or for the VLAN to which the packet's Inner.Label maps on that port if it is an FGL TRILL Data packet), the DF check (see Section 5.2) and the ingress nickname filtering check (see Section 5.3) are further performed. For such an RBv port, if either the DF check or the filtering check fails, the frame MUST NOT be egressed out of that port. Otherwise, it can be egressed out of that port.

7. MAC Information Synchronization in Edge Group

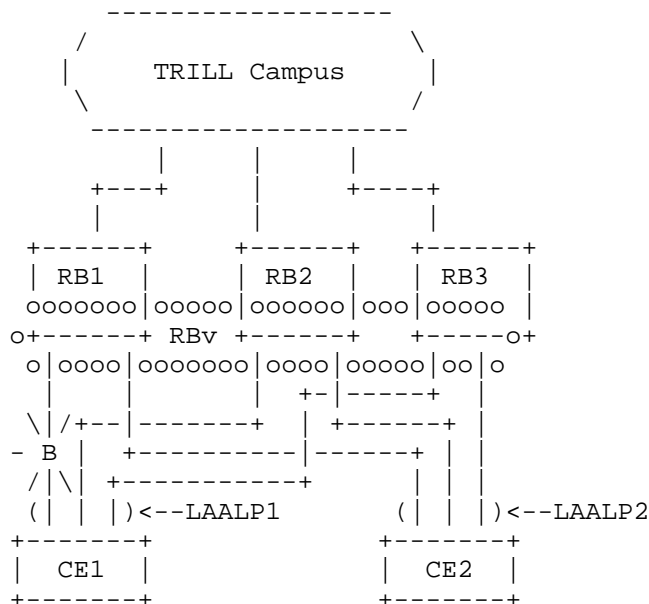
An edge RBridge, say RB1 in LAALP1, may have learned a correspondence between a {Data Label and MAC address} and nickname for a remote host (say h1) when h1 sends a packet to CE1. The returning traffic from CE1 may go to another member RBridge of LAALP1 (for example, RB2). RB2 may not have that correspondence stored. Therefore, it has to do the flooding for unknown unicast. Such flooding is unnecessary, since the returning traffic is almost always expected and RB1 had learned the address correspondence. To avoid the unnecessary flooding, RB1 SHOULD share the correspondence with other RBridges of LAALP1. RB1 synchronizes the correspondence by using the MAC-Reachability (MAC-RI) sub-TLV [RFC6165] in its ESADI-LSPs [RFC7357].

On the other hand, RB2 has learned the MAC address and Data Label of CE1 when CE1 sends a frame to h1 through RB2. The returning traffic from h1 may go to RB1. RB1 may not have CE1's MAC address and Data Label stored even though it is in the same LAALP for CE1 as RB2. Therefore, it has to flood the traffic out of all its access ports where it is Appointed Forwarder for the VLAN (see Section 6.2.1) or the VLAN the FGL maps to on that port if the packet is FGL. Such flooding is unnecessary, since the returning traffic is almost always expected and RB2 had learned CE1's MAC and Data Label information. To avoid that unnecessary flooding, RB2 SHOULD share the MAC address and Data Label with other RBridges of LAALP1. RB2 synchronizes the MAC address and Data Label by enclosing the relative MAC-RI TLV within a pair of boundary TRILL APPsub-TLVs for LAALP1 (see Section 9.3) in its ESADI-LSP [RFC7357]. After receiving the enclosed MAC-RI TLVs, the member RBridges of LAALP1 (i.e., LAALP1 related RBridges) treat the MAC address and Data Label as if it were learned by them locally on their member port of LAALP1; the LAALP1 unrelated RBridges just ignore LAALP1's boundary APPsub-TLVs and treat the MAC address and Data Label as specified in [RFC7357]. Furthermore, in order to make the LAALP1 unrelated RBridges know that the MAC and Data Label are reachable through the RBv that provides service to LAALP1, the Topology-ID/Nickname field of the MAC-RI TLV SHOULD carry the pseudo-nickname of the RBv, rather than a zero value or one of the originating RBridge's (i.e., RB2's) regular nicknames.

8. Member Link Failure in an RBv

As shown in Figure 4, suppose that the link RB1-CE1 fails. Although a new RBv will be formed by RB2 and RB3 to provide active-active service for LAALP1 (see Section 5), the unicast traffic to CE1 might still be forwarded to RB1 before the remote RBridge learns that CE1 is attached to the new RBv. That traffic might be disrupted by the link failure. Section 8.1 discusses failure protection in this scenario.

However, multi-destination TRILL Data packets can reach all member RBridges of the new RBv and be egressed to CE1 by either RB2 or RB3 (i.e., the new DF for the traffic's Inner.VLAN or the VLAN the packet's Inner.Label maps to in the new RBv). Although there might be a transient hang time between failure and the establishment of the new RBv, special actions to protect against downlink failure for such multi-destination packets are not needed.



B - Failed Link or Link Bundle

Figure 4: A Multi-Homed CE with a Failed Link

8.1. Link Protection for Unicast Frame Egressing

When the link CE1-RB1 fails, RB1 loses its direct connection to CE1. The MAC entry through the failed link to CE1 is removed from RB1's local forwarding table immediately. Another MAC entry learned from another member RBridge of LAALP1 (for example, RB2, since it is still a member RBridge of LAALP1) is installed into RB1's forwarding table (see Section 9.3). In that new entry, RB2 (identified by one of its regular nicknames) is the egress RBridge for CE1's MAC address. Then, when a TRILL Data packet to CE1 is delivered to RB1, it can be tunneled to RB2 after being re-encapsulated (the ingress nickname remains unchanged and the egress nickname is replaced by RB2's regular nickname) based on the above installed MAC entry (see bullet 2 in Section 6.2.1). RB2 then receives the frame and egresses it to CE1.

After failure recovery, RB1 learns that it can reach CE1 via link CE1-RB1 again by observing CE1's native frames or from the MAC information synchronization by member RBridge(s) of LAALP1 as described in Section 7. It then restores the MAC entry to its previous one and downloads it to its data-plane "fast path" logic.

9. TLV Extensions for Edge RBridge Group

The following subsections specify the APPsub-TLVs needed to support pseudo-nickname edge groups.

9.1. PN-LAALP-Membership APPsub-TLV

This APPsub-TLV is used by an edge RBridge to announce its associated pseudo-nickname LAALP information. It is defined as a sub-TLV of the TRILL GENINFO TLV [RFC7357] and is distributed in E-L1FS FS-LSPs [RFC7780]. It has the following format:

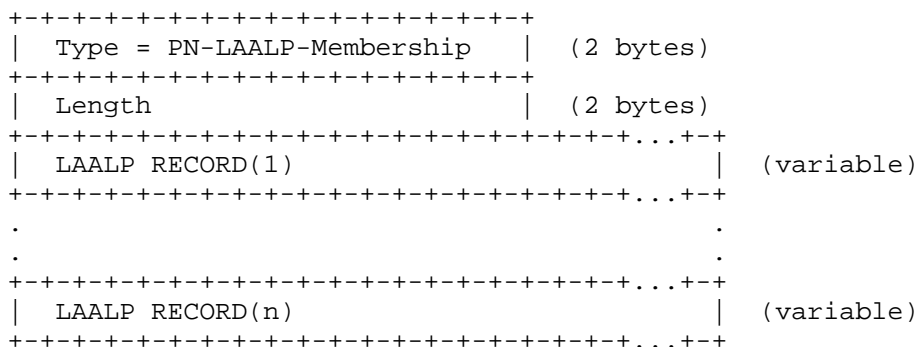



Figure 5: PN-LAALP-Membership Advertisement APPsub-TLV

where each LAALP RECORD has the following form:




- o PN-LAALP-Membership (2 bytes): Defines the type of this sub-TLV, 2.
- o Length (2 bytes): The sum of the lengths of the LAALP RECORDs.
- o OE (1 bit): A flag indicating whether or not the LAALP wants to occupy an RBv by itself; 1 for occupying by itself (or Occupying Exclusively (OE)). By default, it is set to 0 on transmit. This bit is used for edge RBridge group auto-discovery (see Section 4.1). For any one LAALP, the values of this flag might conflict in the LSPs advertised by different member RBridges of that LAALP. In that case, the flag for that LAALP is considered to be 1.
- o RESV (7 bits): MUST be transmitted as zero and ignored on receipt.
- o Size (1 byte): Size of the remaining part of the LAALP RECORD (2 plus the length of the LAALP ID).
- o Reusing Pseudo-Nickname (2 bytes): Suggested pseudo-nickname of the AAE group serving the LAALP. If the LAALP is not served by any AAE group, this field MUST be set to zero. It is used by the originating RBridge to help the vDRB to reuse the previous pseudo-nickname of an AAE group (see Section 4.2).
- o LAALP ID (variable): The ID of the LAALP. See Section 9.4.

On receipt of such an APPsub-TLV, if RBn is not an LAALP related edge RBridge, it ignores the sub-TLV; otherwise, it parses the sub-TLV. When new LAALPs are found or old ones are withdrawn compared to its old copy, and they are also configured on RBn, RBn performs the "Member RBridges Auto-Discovery" procedure described in Section 4.

9.2. PN-RBv APPsub-TLV

The PN-RBv APPsub-TLV is used by a Designated RBridge of a virtual RBridge (vDRB) to dictate the pseudo-nickname for the LAALPs served by the RBv. It is defined as a sub-TLV of the TRILL GENINFO TLV [RFC7357] and is distributed in E-L1FS FS-LSPs [RFC7780]. It has the following format:



- o PN-RBv (2 bytes): Defines the type of this sub-TLV, 3.
- o Length (2 bytes): 3+n*k bytes, where there are n LAALP IDs, each of size k bytes. k is found in the LAALP ID Size field below. If Length is not 3 plus an integer times k, the sub-TLV is corrupt and MUST be ignored.
- o RBv's Pseudo-Nickname (2 bytes): The appointed pseudo-nickname for the RBv that serves the LAALPs listed in the following fields.
- o LAALP ID Size (1 byte): The size of each of the following LAALP IDs in this sub-TLV. 8 if the LAALPs listed are MC-LAGs or DRNI (Section 6.3.2 of [802.1AX]). The value in this field is the k value that appears in the formula for Length above.
- o LAALP ID (LAALP ID Size bytes): The ID of the LAALP. See Section 9.4.

This sub-TLV may occur multiple times with the same RBv pseudo-nickname; this means that all of the LAALPs listed are identified by that pseudo-nickname. For example, if there are LAALP IDs of different length, then the LAALP IDs of each size would have to be listed in a separate sub-TLV.

Because a PN-RBv APPsub-TLV is distributed as part of the application link state by using the E-L1FS FS-LSP [RFC7780], creation, changes to contents, or withdrawal of a PN-RBv APPsub-TLV is accomplished by the Designated RBridge updating and flooding an E-L1FS PDU.

On receipt of such a sub-TLV, if RBn is not an LAALP related edge RBridge, it ignores the sub-TLV. Otherwise, if RBn is also a member RBridge of the RBv identified by the list of LAALPs, it associates the pseudo-nickname with the ports of these LAALPs and downloads the association to data-plane fast path logic. At the same time, RBn claims the RBv's pseudo-nickname across the campus and announces the RBv as its child on the corresponding tree or trees using the Affinity sub-TLV [RFC7176] [RFC7783].

9.3. PN-MAC-RI-LAALP Boundary APPsub-TLVs

In this document, two APPsub-TLVs are used as boundary APPsub-TLVs for an edge RBridge to enclose the MAC-RI TLV(s) containing the MAC address information learned from the local port of an LAALP when this RBridge wants to share the information with other edge RBridges. They are defined as TRILL APPsub-TLVs [RFC7357]. The PN-MAC-RI-LAALP-INFO-START APPsub-TLV has the following format:

```
+-----+
|Type=PN-MAC-RI-LAALP-INFO-START| (2 bytes)
+-----+
| Length | (2 bytes)
+-----+...+-----+
| LAALP ID | (variable)
+-----+...+-----+
```

- o PN-MAC-RI-LAALP-INFO-START (2 bytes): Defines the type of this sub-TLV, 4.
- o Length (2 bytes): The size of the following LAALP ID. 8 if the LAALP listed is an MC-LAG or DRNI.
- o LAALP ID (variable): The ID of the LAALP (see Section 9.4).

The PN-MAC-RI-LAALP-INFO-END APPsub-TLV is defined as follows:

```

+-----+
| Type=PN-MAC-RI-LAALP-INFO-END | (2 bytes)
+-----+
| Length                          | (2 bytes)
+-----+

```

- o PN-MAC-RI-LAALP-INFO-END (2 bytes): Defines the type of this sub-TLV, 5.
- o Length (2 bytes): 0.

This pair of APPsub-TLVs can be carried multiple times in an ESADI-LSP and in multiple ESADI-LSPs. When an LAALP related edge RBridge (say RBn) wants to share with other edge RBridges the MAC addresses learned on its local ports of different LAALPs, it uses one or more pairs of such APPsub-TLVs for each such LAALP in its ESADI-LSPs. Each encloses the MAC-RI TLVs containing the MAC addresses learned from a specific LAALP. Furthermore, if the LAALP is served by a local RBv, the value of the Topology-ID/Nickname field in the relative MAC-RI TLVs SHOULD be the pseudo-nickname of the RBv, rather than one of RBn's regular nicknames or a zero value. Then, on receipt of such a MAC-RI TLV, remote RBridges know that the contained MAC addresses are reachable through the RBv.

On receipt of such boundary APPsub-TLVs, when the edge RBridge is not an LAALP related one or cannot recognize such sub-TLVs, it ignores them and continues to parse the enclosed MAC-RI TLVs per [RFC7357]. Otherwise, the recipient parses the boundary APPsub-TLVs. The PN-MAC-RI-LAALP-INFO-START / PN-MAC-RI-LAALP-INFO-END pair MUST occur within one TRILL GENINFO TLV. If an END is encountered without any previous START in the ESADI-LSP, the END APPsub-TLV is ignored. After encountering a START, if the end of the ESADI-LSP is reached without encountering an END, then the end of the ESADI-LSP is treated as if it were a PN-MAC-RI-LAALP-INFO-END. The boundary APPsub-TLVs and TLVs between them are handled as follows:

- 1) If the edge RBridge is configured with the contained LAALP and the LAALP is also enabled locally, it treats all the MAC addresses contained in the following MC-RI TLVs enclosed by the corresponding pair of boundary APPsub-TLVs as if they were learned from its local port of that LAALP;
- 2) Else, it ignores these boundary APPsub-TLVs and continues to parse the following MAC-RI TLVs per [RFC7357] until another pair of boundary APPsub-TLVs is encountered.

9.4. LAALP IDs

The LAALP ID identifies an AAE RBridge group in the TRILL campus and thus MUST be unique across the campus. In all of the APPsub-TLVs specified above, the length of the LAALP ID can be determined from a size field. If that length is 8 bytes, the LAALP ID is an MC-LAG or DRNI identifier as specified in Section 6.3.2 of [802.1AX]. The meaning and structure of LAALP IDs of other lengths are reserved and may be specified in future documents.

10. OAM Packets

Attention must be paid when generating Operations, Administration, and Maintenance (OAM) packets. To ensure that the response messages can return to the originating member RBridge of an RBv, a pseudo-nickname cannot be used as the ingress nickname in TRILL OAM messages, except in the response to an OAM message that has that RBv's pseudo-nickname as the egress nickname. For example, assume that RB1 is a member RBridge of RBvi. RB1 cannot use RBvi's pseudo-nickname as the ingress nickname when originating OAM messages; otherwise, the responses to the messages may be delivered to another member RBridge of RBvi rather than RB1. But when RB1 responds to the OAM message with RBvi's pseudo-nickname as the egress nickname, it can use that pseudo-nickname as the ingress nickname in the response message.

Since RBridges cannot use OAM messages for the learning of MAC addresses (Section 3.2.1 of [RFC7174]), it will not lead to MAC address flip-flopping at a remote RBridge, even though RB1 uses its regular nicknames as ingress nicknames in its TRILL OAM messages, and at the same time RB1 uses RBvi's pseudo-nickname in its TRILL Data packets.

11. Configuration Consistency

The VLAN membership of all the RBridge ports in an LAALP MUST be the same. Any inconsistencies in VLAN membership may result in packet loss or non-shortest paths.

Take Figure 1 as an example. Suppose that RB1 configures VLAN1 and VLAN2 for the CE1-RB1 link, while RB2 only configures VLAN1 for the CE1-RB2 link. Both RB1 and RB2 use the same ingress nickname RBv for all frames originating from CE1. Hence, a remote RBridge (say RBx) will learn that CE1's MAC address in VLAN2 is originating from the RBv. As a result, on the return path, RBx may deliver VLAN2 traffic to RB2. However, RB2 does not have VLAN2 configured on the CE1-RB2 link, and hence the frame may be dropped or has to be redirected to RB1 if RB2 knows that RB1 can reach CE1 in VLAN2.

How LAALP implementations maintain consistent VLAN configuration on the TRILL switch LAALP ports is out of scope for the TRILL protocol. However, considering the consequences that might be caused by inconsistencies, TRILL switches MUST disable the ports connected to an LAALP with an inconsistent VLAN configuration.

It is important that if any VLAN in an LAALP is being mapped by edge RBridges to an FGL [RFC7172] the mapping MUST be the same for all edge RBridge ports in the LAALP. Otherwise, for example, unicast FGL TRILL Data packets from remote RBridges may get mapped into different VLANs, depending on which edge RBridge receives and egresses them.

It is important that RBridges in an AAE group not be configured to assert the OE-flag if any RBridge in the group does not implement it. Since, as stated in [RFC7379], the RBridges in an AAE edge group are expected to be from the same vendor, due to the proprietary nature of deployed LAALPs, this will normally follow automatically from all of the RBridges in an AAE edge group supporting, or not supporting, OE.

12. Security Considerations

Authenticity for contents transported in IS-IS PDUs is enforced using regular IS-IS security mechanisms [IS-IS] [RFC5310].

For security considerations pertaining to extensions transported by TRILL ESADI, see the Security Considerations section in [RFC7357].

Since currently deployed LAALPs [RFC7379] are proprietary, security over membership in, and internal management of, active-active edge groups is proprietary. If authentication is not used, a rogue RBridge that insinuates itself into an active-active edge group can disrupt end-station traffic flowing into or out of that group. For example, if there are N RBridges in the group, it could typically control 1/Nth of the traffic flowing out of that group and a similar amount of unicast traffic flowing into that group. For multi-destination traffic flowing into that group, it could control all that was in a VLAN for which it was the DF and can exercise substantial control over the DF election by changing its own System ID.

For general TRILL security considerations, see [RFC6325].

13. IANA Considerations

IANA has allocated four code points from the range below 255 for the four TRILL APPsub-TLVs specified in Section 9 and added them to the "TRILL APPsub-TLV Types under IS-IS TLV 251 Application Identifier 1" registry, as follows:

| Type | Name | Reference |
|------|----------------------------|-----------|
| 2 | PN-LAALP-Membership | RFC 7781 |
| 3 | PN-RBv | RFC 7781 |
| 4 | PN-MAC-RI-LAALP-INFO-START | RFC 7781 |
| 5 | PN-MAC-RI-LAALP-INFO-END | RFC 7781 |

14. References

14.1. Normative 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.
- [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>>.
- [RFC6165] Banerjee, A. and D. Ward, "Extensions to IS-IS for Layer-2 Systems", RFC 6165, DOI 10.17487/RFC6165, April 2011, <<http://www.rfc-editor.org/info/rfc6165>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<http://www.rfc-editor.org/info/rfc6234>>.
- [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>>.
- [RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, DOI 10.17487/RFC7356, September 2014, <<http://www.rfc-editor.org/info/rfc7356>>.
- [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>>.
- [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>>.
- [RFC7783] Senevirathne, T., Pathangi, J., and J. Hudson, "Coordinated Multicast Trees (CMT) for Transparent Interconnection of Lots of Links (TRILL)", RFC 7783, DOI 10.17487/RFC7783, February 2016, <<http://www.rfc-editor.org/info/rfc7783>>.

14.2. Informative 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.
- [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>>.
- [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>>.
- [RFC7782] Zhang, M., Perlman, R., Zhai, H., Durrani, M., and S. Gupta, "Transparent Interconnection of Lots of Links (TRILL) Active-Active Edge Using Multiple MAC Attachments", RFC 7782, DOI 10.17487/RFC7782, February 2016, <<http://www.rfc-editor.org/info/rfc7782>>.

Acknowledgments

We would like to thank Mingjiang Chen for his contributions to this document. Additionally, we would like to thank Erik Nordmark, Les Ginsberg, Ayan Banerjee, Dinesh Dutt, Anoop Ghanwani, Janardhanan Pathangi, Jon Hudson, and Fangwei Hu for their good questions and comments.

Contributors

Weiguo Hao
Huawei Technologies
101 Software Avenue
Nanjing 210012
China

Phone: +86-25-56623144
Email: haoweiguo@huawei.com

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

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

Authors' Addresses

Hongjun Zhai
Jinling Institute of Technology
99 Hongjing Avenue, Jiangning District
Nanjing, Jiangsu 211169
China

Email: honjun.zhai@tom.com

Tissa Senevirathne
Consultant

Email: tsenevir@gmail.com

Radia Perlman
EMC
2010 256th Avenue NE, #200
Bellevue, WA 98007
United States

Email: Radia@alum.mit.edu

Mingui Zhang
Huawei Technologies
No. 156 Beiqing Rd., Haidian District
Beijing 100095
China

Email: zhangmingui@huawei.com

Yizhou Li
Huawei Technologies
101 Software Avenue
Nanjing 210012
China

Phone: +86-25-56625409

Email: liyizhou@huawei.com