

Internet Engineering Task Force (IETF)
Request for Comments: 7177
Obsoletes: 6327
Updates: 6325
Category: Standards Track
ISSN: 2070-1721

D. Eastlake 3rd
Huawei
R. Perlman
Intel Labs
A. Ghanwani
Dell
H. Yang
Cisco
V. Manral
Ionos Corp.
May 2014

Transparent Interconnection of Lots of Links (TRILL): Adjacency

Abstract

The IETF Transparent Interconnection of Lots of Links (TRILL) protocol supports arbitrary link technologies between TRILL switches, including point-to-point links and multi-access Local Area Network (LAN) links that can have multiple TRILL switches and end stations attached. TRILL uses Intermediate System to Intermediate System (IS-IS) routing. This document specifies the establishment, reporting, and termination of IS-IS adjacencies between TRILL switches, also known as RBridges (Routing Bridges). It also concerns four other link-local aspects of TRILL: Designated RBridge (DRB) selection, MTU (Maximum Transmission Unit) testing, pseudonode creation, and BFD (Bidirectional Forwarding Detection) session bootstrapping in connection with adjacency. State diagrams are included where appropriate. This document obsoletes RFC 6327 and 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/rfc7177>.

Copyright Notice

Copyright (c) 2014 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. Content and Precedence	5
1.2. Terminology and Acronyms	5
2. The TRILL Hello Environment and Purposes	6
2.1. RBridge Interconnection by Ethernet	6
2.2. Handling Native Frames	7
2.3. Zero or Minimal Configuration	8
2.4. MTU Robustness	8
2.5. Purposes of the TRILL Hello Protocol	9
3. Adjacency State Machinery	10
3.1. TRILL Hellos, Ports, and VLANs	10
3.2. Adjacency Table Entries and States	11
3.3. Adjacency and Hello Events	12
3.4. Adjacency State Diagram and Table	15
3.5. Multiple Parallel Links	17
3.6. Insufficient Space in Adjacency Table	17
4. LAN Ports and DRB State	17
4.1. Port Table Entries and DRB Election State	18
4.2. DRB Election Events	19
4.2.1. DRB Election Details	20
4.2.2. Change in DRB	20
4.2.3. Change in Designated VLAN	21
4.3. Port State Table and Diagram	21
5. MTU Matching	22
6. BFD-Enabled TLV and BFD Session Bootstrapping	24
7. Pseudonodes	25
8. More TRILL Hello Details	26
8.1. Contents of TRILL Hellos	27
8.2. Transmitting TRILL Hellos	27
8.2.1. TRILL Neighbor TLVs	28
8.3. Receiving TRILL Hellos	29
9. Multiple Ports on the Same Broadcast Link	29
10. IANA Considerations	30
11. Security Considerations	30
Appendix A. Changes from RFC 6327	31
Appendix B. Changes to RFC 6325	31
Normative References.....	32
Informative References.....	33
Acknowledgements.....	34

1. Introduction

The IETF Transparent Interconnection of Lots of Links (TRILL) protocol [RFC6325] provides optimal pair-wise data frame forwarding without configuration, safe forwarding even during transient loops, and support for transmission of both unicast and multicast traffic taking advantage of multiple paths in multi-hop networks with arbitrary topology. End stations are connected to TRILL switches with Ethernet, but TRILL switches can be interconnected with arbitrary link technology. TRILL accomplishes this by using [IS-IS] (Intermediate System to Intermediate System) link-state routing [RFC1195] [RFC7176] and a header in TRILL Data packets that includes a hop count. The design supports data labeling by Virtual Local Area Networks (VLANs) and fine-grained labels [RFC7172] as well as optimization of the distribution of multi-destination frames based on data label and multicast groups. Devices that implement TRILL are called TRILL switches or RBridges (Routing Bridges).

This document provides detailed specifications for five of the link-local aspects of the TRILL protocol used on broadcast links (also called LAN or multi-access links) and for the three of these aspects that are also used on point-to-point (P2P) links. It includes state diagrams and implementation details where appropriate. Alternative implementations that interoperate on the wire are permitted.

The scope of this document is limited to the following aspects of the TRILL protocol; their applicability, along with the most relevant section of this document, are as shown here:

LAN	P2P	Section	Link-Local Aspect
---	---	-----	-----
X	X	3	Adjacency formation and dissolution
X		4	DRB (Designated RBridge) election
X	X	5	MTU (Maximum Transmission Unit) matching
X	X	6	1-hop BFD (Bidirectional Forwarding Detection) for adjacency
X		7	Creation and use of pseudonodes [IS-IS]

Table 1: LAN/P2P Applicability

There is no DRB (also known as DIS (Designated Intermediate System)) election and no pseudonode creation on links configured as point-to-point.

For other aspects of the TRILL base protocol, see [RFC6325], [RFC6439], [RFC7180], and [ESADI].

This document obsoletes [RFC6327]. See Appendix A for a summary of changes from [RFC6327]. This document updates [RFC6325] as described in Appendix B.

1.1. Content and Precedence

In cases of conflict between this document and [RFC6325], this document prevails.

Section 2 below explains the rationale for the differences between the TRILL Hello protocol and the Layer 3 IS-IS Hello protocol in light of the environment for which the TRILL protocol is designed. It also describes the purposes of the TRILL Hello protocol.

Section 3 describes the adjacency state machine, its states, and its relevant events.

Section 4 describes the Designated RBridge (DRB) election state machine for RBridge LAN ports, its states, and its relevant events.

Section 5 describes MTU testing and matching on a TRILL link.

Section 6 discusses one-hop BFD session bootstrapping in connection with adjacency.

Section 7 discusses pseudonode creation and use on LAN links.

Section 8 provides more details on the reception and transmission of TRILL Hellos.

Section 9 discusses the case of multiple ports from one RBridge on the same link.

1.2. Terminology and Acronyms

This document uses the acronyms defined in [RFC6325], supplemented by the following additional acronyms:

BFD - Bidirectional Forwarding Detection [RFC7175].

SNPA - Subnetwork Point of Attachment [IS-IS].

TRILL switch - an alternative name for an RBridge.

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].

2. The TRILL Hello Environment and Purposes

[IS-IS] has subnetwork-independent functions and subnetwork-dependent functions. Currently, Layer 3 use of IS-IS supports two types of subnetworks: (1) point-to-point link subnetworks between routers and (2) general broadcast (LAN) subnetworks. Because of the differences between the environment of Layer 3 routers and the environment of TRILL RBridges, instead of the subnetwork-dependent functions used at Layer 3, which are specified in Sections 8.2 and 8.4 of [IS-IS], the TRILL protocol uses modified subnetwork-dependent functions for point-to-point subnetworks and broadcast (LAN) subnetworks. The differences between the TRILL and Layer 3 environments are described in Sections 2.1 through 2.4 followed by a summation, in Section 2.5, of the purposes of the TRILL Hello protocol.

2.1. RBridge Interconnection by Ethernet

TRILL supports the interconnection of RBridges by multi-access LAN links such as Ethernet. Because this includes general bridged LANs [802.1Q], the links between RBridges may contain devices or services that can restrict VLAN connectivity, such as [802.1Q] bridges or carrier Ethernet services. In addition, RBridge Ethernet ports, like [802.1Q] ports, can be configured to restrict input/output on a VLAN basis.

For this reason, TRILL Data and TRILL IS-IS packets are sent on Ethernet links in a Designated VLAN that is assumed to provide connectivity between all RBridges on the link. The Designated VLAN is dictated for a LAN link by the elected Designated RBridge on that link (DRB, equivalent to the Designated Intermediate System at Layer 3). On an RBridge Ethernet port configured as point-to-point, TRILL Data and IS-IS packets are sent in that port's Desired Designated VLAN, regardless of the state of any other ports on the link. Connectivity on an Ethernet link configured as point-to-point generally depends on both ends being configured with the same Desired Designated VLAN. Because TRILL Data packets flow between RBridges on an Ethernet link only in the link's Designated VLAN, adjacency for routing calculations is based only on connectivity characteristics in that VLAN.

(Non-Ethernet links, such as PPP [RFC6361] generally do not have any Outer.VLAN labeling, so the Designated VLAN for such links has no effect.)

2.2. Handling Native Frames

This section discusses the handling of "native" frames as defined in Section 1.4 of [RFC6325]. As such, this section is not applicable to point-to-point links between TRILL switches or any link where all the TRILL switch ports on the link have been configured as "trunk ports" by setting the end-station service disable bit for the port (see Section 4.9.1 of [RFC6325]).

Layer 3 data packets, such as IP packets, are already "tamed" when they are originated by an end station: they include a hop count and Layer 3 source and destination address fields. Furthermore, for ordinary data packets, there is no requirement to preserve any outer Layer 2 addressing, and, if the packets are unicast, they are explicitly addressed to their first-hop router.

In contrast, TRILL switches must accept, transport, and deliver "untamed" native frames: native frames that lack a hop count field usable by TRILL and have Layer 2 MAC (Media Access Control) addresses that indicate their source and destination. These Layer 2 addresses must be preserved for delivery to the native frame's Layer 2 destination. One resulting difference is that RBridge ports providing native frame service must receive in promiscuous MAC address mode, while Layer 3 router ports typically receive in a selective MAC address mode.

TRILL handles these requirements by having, on the link where an end station originates a native frame, one RBridge "ingress" such a locally originated native frame by adding a TRILL Header that includes a hop count, thus converting it to a TRILL Data packet. This augmented packet is then routed to one RBridge on the link having the destination end station for the frame (or one RBridge on each such link if it is a multi-destination frame). Such final RBridges perform an "egress" function, removing the TRILL Header and delivering the original frame to its destination(s). (For the purposes of TRILL, a Layer 3 router is an end station.)

Care must be taken to avoid a loop that would involve egressing a native frame and then re-ingressing it because, while it is in native form, it would not be protected by a hop count and could loop forever. Such a loop could, for multi-destination frames, even involve multiplication of the number of frames each time around and would likely saturate all links involved within milliseconds. For TRILL, safety against such loops for a link is more important than transient loss of data connectivity on that link.

The primary TRILL defense mechanism against such loops, which is mandatory, is to assure that, as far as practically possible, there is only a single RBridge that is in charge of ingressing and egressing native frames from and to a link where TRILL is offering end-station service. This is the Designated RBridge and Appointed Forwarder mechanism initially specified in the TRILL base protocol [RFC6325], discussed in Section 2.5 below, and further specified in both Section 4 below and [RFC6439].

2.3. Zero or Minimal Configuration

TRILL provides connectivity and least-cost paths with zero configuration. For additional services, it strives to require only minimal configuration; however, services that require configuration when offered by [802.1Q] bridges, such as non-default VLANs or priority, will require configuration. This differs from Layer 3 routing where routers typically need to be configured as to the subnetworks connected to each port, etc., to provide service.

2.4. MTU Robustness

TRILL IS-IS needs to be robust against links with reasonably restricted MTUs, including links that accommodate only the classic Ethernet frame size, despite the addition of reasonable headers such as VLAN tags. Such robustness is particularly required for TRILL Hellos to assure correct adjacency and the election of a unique DRB on LAN links.

TRILL will also be used inside data centers where it is common for all or most of the links and switches to support frames substantially larger than the classic Ethernet maximum size. For example, they may have an MTU adequate to comfortably handle Fiber Channel over Ethernet frames, for which T11 recommends a 2,500-byte MTU [FCoE], or even 9K byte jumbo frames. It would be beneficial for a TRILL campus with such a large MTU to be able to safely make use of it.

These needs are met by a mandatory maximum on the size of TRILL Hellos and by the optional use of MTU testing as described below.

2.5. Purposes of the TRILL Hello Protocol

There are three purposes for the TRILL Hello protocol. They are listed below, along with a reference to the section of this document in which each is discussed:

1. To determine which RBridge neighbors have acceptable connectivity to be reported as part of the topology (Section 3)
2. To elect a unique Designated RBridge on broadcast (LAN) links (Section 4)
3. To determine the MTU with which it is possible to safely communicate with each RBridge neighbor (Section 5)

In Layer 3 IS-IS, all three of these functions are combined. Hellos may be padded to the maximum length (see [RFC3719], Section 6) so that a router neighbor is not discovered if it is impossible to communicate with it using maximum-sized Layer 3 IS-IS packets. Also, even if Hellos from a neighbor R2 are received by R1, if connectivity to R2 is not 2-way (i.e., R2 does not list R1 in R2's Hello), then R1 does not consider R2 as a Designated Intermediate System (Designated Router) candidate. Because of this logic, it is possible at Layer 3 for multiple Designated Routers to be elected on a LAN, with each representing the LAN as a pseudonode. It appears to the topology as if the LAN is now two or more separate LANs. Although this is surprising, this does not cause problems for Layer 3.

In contrast, this behavior is not acceptable for TRILL, since in TRILL it is important that all RBridges on a link know about each other, and on broadcast (LAN) links that they choose a single RBridge to be the DRB to control the native frame ingress and egress. Otherwise, multiple RBridges might ingress/egress the same native frame, forming loops that are not protected by the hop count in the TRILL Header as discussed above.

The TRILL Hello protocol is best understood by focusing separately on each of these three functions listed above, which we do in Sections 3, 4, and 5.

One other issue with TRILL LAN Hellos is to ensure that subsets of the information can appear in any single message, and be processable, in the spirit of IS-IS Link State PDUs (LSPs) and Complete Sequence Number PDUs (CSNPs). LAN TRILL Hello packets, even though they are not padded, can become very large. An example where this might be the case is when some sort of backbone technology interconnects hundreds of TRILL sites over what would appear to TRILL to be a giant Ethernet, where the RBridges connected to that cloud will perceive

that backbone to be a single link with hundreds of neighbors. Thus, the TRILL LAN Hello uses a different Neighbor TLV [RFC7176] that lists neighbors seen for a range of MAC (SNPA) addresses.

3. Adjacency State Machinery

Each RBridge port has associated with it a port state, as discussed in Section 4, and a table of zero or more adjacencies (if the port is configured as point-to-point, zero, or one) as discussed in this section. The states such adjacencies can have, the events that cause adjacency state changes, the actions associated with those state changes, a state table, and a state diagram are given below.

3.1. TRILL Hellos, Ports, and VLANs

The determination of adjacencies on links is made using TRILL Hellos (see Section 8), an optional MTU test (see Section 5), and, optionally, BFD (see Section 6) and/or other connectivity tests. If the MAC (SNPA) addresses of more than one RBridge port on a broadcast link are the same, all but one of such ports are put in the Suspended state (see Section 4) and do not participate in the link, except to monitor whether they should stay suspended. If the two ports on a point-to-point link have MAC (SNPA) addresses, it does not affect TRILL operation if they are the same. (PPP ports, for example, do not have MAC addresses [RFC6361].)

The following items **MUST** be the same for all TRILL Hellos issued by an RBridge on a particular Ethernet port, regardless of the VLAN in which the Hello is sent:

- Source MAC address,
- Priority to be the DRB,
- Desired Designated VLAN,
- Port ID, and,
- if included, BFD-Enabled TLV [RFC6213] and PORT-TRILL-VER sub-TLV [RFC7176].

Of course, the priority, Desired Designated VLAN, and possibly the inclusion or value of the PORT-TRILL-VER sub-TLV, and/or BFD-Enabled TLV can change on occasion, but then the new value(s) must similarly be used in all TRILL Hellos on the LAN port, regardless of VLAN.

On broadcast links:

Because bridges acting as glue on an Ethernet broadcast link might be configured in such a way that some VLANs are partitioned, it is necessary for RBridges to transmit Hellos on Ethernet links with multiple VLAN tags. The conceptually simplest solution may have been to have RBridges transmit up to 4,094 times as many Hellos, one with each legal VLAN ID enabled at each port, but this would obviously have deleterious performance implications. So, the TRILL protocol specifies that if RB1 knows it is not the DRB, it transmits its Hellos on only a limited set of VLANs. Only an RBridge that believes itself to be the DRB on a broadcast Ethernet link "sprays" its TRILL Hellos on all of its enabled VLANs at the port. And in both cases, an RBridge can be configured to send Hellos on only a subset of those VLANs. The details are given in [RFC6325], Section 4.4.3.

On point-to-point links:

If the link technology is VLAN sensitive, such as Ethernet, an RBridge sends TRILL Hellos only in the Desired Designated VLAN for which it is configured.

3.2. Adjacency Table Entries and States

Every adjacency is in one of four states, whether it is one of the adjacencies on a broadcast link or the one possible adjacency on a point-to-point link. An RBridge participates in LSP synchronization at a port as long as it has one or more adjacencies out of that port that are in the 2-Way or Report state.

Down: This is a virtual state for convenience in creating state diagrams and tables. It indicates that the adjacency is nonexistent, and there is no entry in the adjacency table for it.

Detect: A neighbor RBridge has been detected through receipt of a TRILL Hello, but either 2-way connectivity has not been confirmed or the detection was on an Ethernet link in a VLAN other than the Designated VLAN.

2-Way: 2-way connectivity to the neighbor has been found and, if the link is Ethernet, it was found on the Designated VLAN, but some enabled test, such as the link MTU meeting the minimum campus requirement or BFD confirming link connectivity, has not yet succeeded.

Report: There is 2-way connectivity to the neighbor (on the Designated VLAN if an Ethernet link); all enabled tests have succeeded, including, if enabled, MTU and/or BFD testing. This state will cause adjacency to be reported in an LSP (with appropriate provision for a pseudonode, if any, as described in Section 7).

For an adjacency in any of the three non-Down states (Detect, 2-Way, or Report), there will be an adjacency table entry. That entry will give the state of the adjacency and will also include the information listed below.

- o The address, if any, of the neighbor, the Port ID, and the System ID in the received Hellos. Together, these three quantities uniquely identify the adjacency on a broadcast link.
- o One or more Hello holding timers. For a point-to-point adjacency, there is a single Hello holding timer. For a broadcast LAN adjacency, there are exactly two Hello holding timers: a Designated VLAN holding timer and a non-Designated VLAN holding timer. Each timer consists of a 16-bit unsigned integer number of seconds.
- o If the adjacency is on a broadcast link, the 7-bit unsigned priority of the neighbor to be the DRB.
- o The 5 bytes of data from the PORT-TRILL-VER received in the most recent TRILL Hello from the neighbor RBridge.
- o The VLAN that the neighbor RBridge wants to be the Designated VLAN on the link, called the Desired Designated VLAN.
- o For an adjacency table at an RBridge that supports BFD, a flag indicating whether the last received TRILL Hello from the neighbor RBridge contained a BFD-Enabled TLV (see Section 6).

3.3. Adjacency and Hello Events

The following events can change the state of an adjacency:

- A0. Receive a TRILL Hello for a broadcast LAN adjacency whose source MAC address (SNPA) is equal to that of the port on which it is received. This is a special event that cannot occur on a port configured as point-to-point and is handled as described immediately after this list of events. It does not appear in the state transition table or diagram.

- A1. Receive a TRILL Hello (other than an A0 event) such that:
- If received on an Ethernet port, it was received in the Designated VLAN.
 - If received for a broadcast LAN adjacency, it contains a TRILL Neighbor TLV that explicitly lists the receiving port's (SNPA) address.
 - If received for a point-to-point adjacency, it contains a Three-Way Handshake TLV with the receiver's System ID and Extended Circuit ID.
- A2. Event A2 is not possible for a port configured as point-to-point. Receive a TRILL Hello (other than an A0 event) such that either
- The port is Ethernet and the Hello was not on the Designated VLAN (any TRILL Neighbor TLV in such a Hello is ignored), or
 - The Hello does not contain a TRILL Neighbor TLV covering an address range that includes the receiver's (SNPA) address.
- A3. Receive a TRILL Hello (other than an A0 event) such that:
- If received on an Ethernet port, it was received in the Designated VLAN.
 - If received for a broadcast LAN adjacency, it contains one or more TRILL Neighbor TLVs covering an address range that includes the receiver's (SNPA) address and none of which list the receiver.
 - If received for a point-to-point adjacency, it contains a Three-Way Handshake TLV with either the System ID or Extended Circuit ID or both not equal to that of the receiver.
- A4. Either
- (1) the Hello holding timer expires on a point-to-point adjacency, or
 - (2) on a broadcast LAN adjacency,
 - (2a) both Hello timers expire simultaneously or
 - (2b) one Hello timer expires when the other Hello timer is already in the expired state.

- A5. For a broadcast LAN adjacency, the Designated VLAN Hello holding timer expires, but the non-Designated VLAN Hello holding timer still has time left until it expires. This event cannot occur for a point-to-point adjacency.
- A6. MTU if enabled, BFD if enabled, and all other enabled connectivity tests successful.
- A7. MTU if enabled, BFD if enabled, and all other enabled connectivity tests were successful but one or more now fail.
- A8. The RBridge port goes operationally down.

For the special A0 event, the Hello is examined to determine if it has a higher priority than the port on which it is received such that the sending port should be the DRB as described in Section 4.2.1. If the Hello is of lower priority than the receiving port, it is discarded with no further action. If it is of higher priority than the receiving port, then any adjacencies for the receiving port are discarded (transitioned to the Down state), and the port is suspended as described in Section 4.2.

The receipt of a TRILL Hello that is not an event A0 causes the following actions (except where the Hello would have created a new adjacency table entry but both the adjacency table is full and the Hello is too low priority to displace an existing entry as described in Section 3.6). The Designated VLAN referred to is the Designated VLAN dictated by the DRB determined without taking the received TRILL LAN Hello into account (see Section 4) for a broadcast LAN and the local Desired Designated VLAN for a port configured as point-to-point.

- o If the receipt of a Hello creates a new adjacency table entry, the neighbor RBridge MAC (SNPA) address (if any), Port ID, and System ID are set from the Hello.
- o For a point-to-point adjacency, the Hello holding timer is set from the Holding Time field of the Hello. For a broadcast link adjacency, the appropriate Hello holding timer for that adjacency, depending on whether or not the Hello was received in the Designated VLAN, is set to the Holding Time field of the Hello and if the receipt of the LAN Hello is creating a new adjacency table entry, the other timer is set to expired.
- o For a broadcast link adjacency, the priority of the neighbor RBridge to be the DRB is set to the priority field of the LAN Hello.

- o For a broadcast link adjacency, the VLAN that the neighbor RBridge wants to be the Designated VLAN on the link is set from the Hello.
- o The 5 bytes of PORT-TRILL-VER data are set from that sub-TLV in the Hello or set to zero if that sub-TLV does not occur in the Hello.
- o For a broadcast link, if the creation of a new adjacency table entry or the priority update above changes the results of the DRB election on the link, the appropriate RBridge port event (D2 or D3) occurs, after the above actions, as described in Section 4.2.
- o For a broadcast link adjacency, if there is no change in the DRB, but the neighbor Hello is from the DRB and has a changed Designated VLAN from the previous Hello received from the DRB, the result is a change in Designated VLAN for the link as specified in Section 4.2.3.

An event A4 resulting in the adjacency transitioning to the Down state may also result in an event D3 as described in Section 4.2.

Concerning events A6 and A7, if none of MTU, BFD, or other testing is enabled, A6 is considered to occur immediately upon the adjacency entering the 2-Way state, and A7 cannot occur.

See further TRILL Hello receipt details in Section 8.

3.4. Adjacency State Diagram and Table

The table below shows the transitions between the states defined above, based on the events defined above:

Event	Down	Detect	2-Way	Report
A1	2-Way	2-Way	2-Way	Report
A2	Detect	Detect	2-Way	Report
A3	Detect	Detect	Detect	Detect
A4	N/A	Down	Down	Down
A5	N/A	Detect	Detect	Detect
A6	N/A	N/A	Report	Report
A7	N/A	N/A	2-Way	2-Way
A8	Down	Down	Down	Down

Table 2: Adjacency State Table

"N/A" indicates that the event to the left is not applicable in the state at the top of the column. These events affect only a single adjacency. The special A0 event transitions all adjacencies to Down, as explained immediately after the list of adjacency events in Section 3.3.

The diagram below presents the same information as that in the state table:

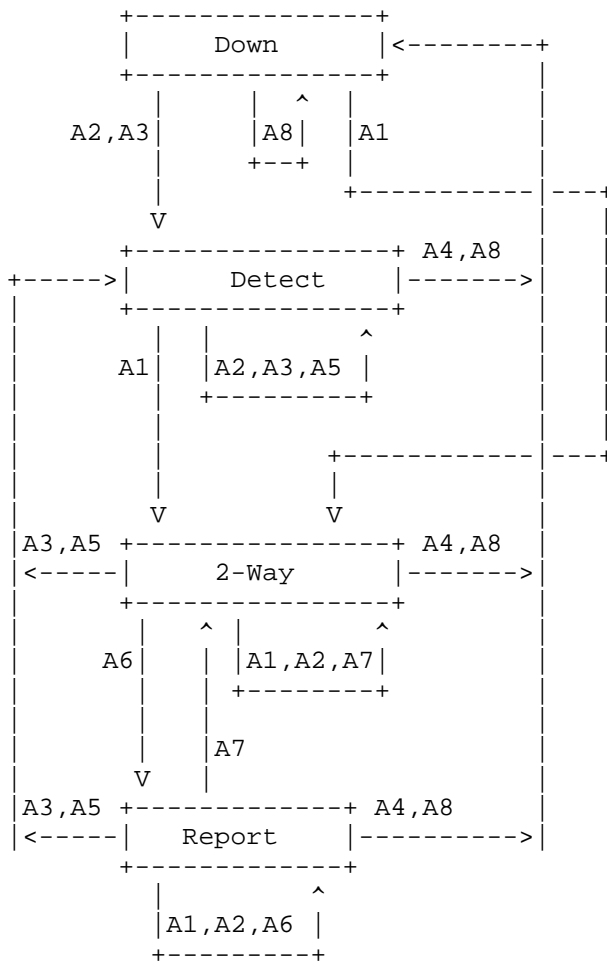


Figure 1: Adjacency State Diagram

3.5. Multiple Parallel Links

There can be multiple parallel adjacencies between neighbor RBridges that are visible to TRILL. (Multiple low-level links that have been bonded together by technologies such as link aggregation [802.1AX] appear to TRILL as a single link over which only a single TRILL adjacency can be established.)

Any such links that have pseudonodes (see Section 7) are distinguished in the topology; such adjacencies, if they are in the Report state, appear in LSPs as per Section 7. However, there can be multiple parallel adjacencies without pseudonodes because they are point-to-point adjacencies or LAN adjacencies for which a pseudonode is not being created. Such parallel, non-pseudonode adjacencies in the Report state appear in LSPs as a single adjacency. The cost of such an adjacency MAY be adjusted downwards to account for the parallel paths. Multipathing across such parallel connections can be freely done for unicast TRILL Data traffic on a per-flow basis but is restricted for multi-destination traffic, as described in Section 4.5.2 (point 3) of [RFC6325] and Appendix C of [RFC6325].

3.6. Insufficient Space in Adjacency Table

If the receipt of a TRILL Hello would create a new adjacency table entry (that is, would transition an adjacency out of the Down state), there may be no space for the new entry. For ports that are configured as point-to-point and can thus only have zero or one adjacency not in the Down state, it is RECOMMENDED that space be reserved for one adjacency so that this condition cannot occur.

When there is adjacency table space exhaustion, the DRB election priority (see Section 4.2.1) of the new entry that would be created is compared with the DRB election priority for the existing entries. If the new entry is higher priority than the lowest priority existing entry, it replaces the lowest priority existing entry, which is transitioned to the Down state.

4. LAN Ports and DRB State

This section specifies the DRB election process in TRILL at a broadcast (LAN) link port. Since there is no such election when a port is configured as point-to-point, this section does not apply in that case.

The information at an RBridge associated with each of its broadcast LAN ports includes the following:

- o Enablement bit, which defaults to enabled.
- o The 5 bytes of PORT-TRILL-VER sub-TLV data used in TRILL Hellos sent on the port.
- o SNPA address (commonly a 48-bit MAC address) of the port.
- o Port ID, used in TRILL Hellos sent on the port.
- o The Holding Time, used in TRILL Hellos sent on the port.
- o The priority to be the DRB, used in TRILL LAN Hellos sent on the port.
- o BFD support. If the port supports BFD, a BFD Enabled flag that controls whether or not a BFD-Enabled TLV is included in TRILL Hellos sent on the port.
- o The DRB state of the port, determined as specified below.
- o A 16-bit unsigned Suspension Timer, measured in seconds.
- o The Desired Designated VLAN. The VLAN this RBridge wants to be the Designated VLAN for the link out of this port, used in TRILL Hellos sent on the port if the link is Ethernet.
- o A table of zero or more adjacencies (see Section 3).

4.1. Port Table Entries and DRB Election State

The TRILL equivalent of the DIS (Designated Intermediate System) on a broadcast link is the DRB or Designated RBridge. The DRB election state machinery is described below.

Each RBridge port that is not configured as point-to-point is in one of the following four DRB states:

Down: The port is operationally down. It might be administratively disabled or down at the link layer. In this state, there will be no adjacencies for the port, and no TRILL Hellos or other TRILL IS-IS PDUs or TRILL Data packets are accepted or transmitted.

Suspended: Operation of the port is suspended because there is a higher priority port on the link with the same MAC (SNPA) address. This is the same as the Down state, with the exception that TRILL

Helloes are accepted for the sole purpose of determining whether to change the value of the Suspension Timer for the port as described below.

DRB: The port is the DRB and can receive and transmit TRILL Data packets.

Not DRB: The port is deferring to another port on the link, which it believes is the DRB, but can still receive and transmit TRILL Data packets.

4.2. DRB Election Events

The following events can change the DRB state of a port. Note that this is only applicable to broadcast links. There is no DRB state or election at a port configured to be point-to-point.

- D1. The port becomes enabled or the Suspension Timer expires while the port is in the Suspended state.
- D2. The adjacency table for the port changes, and there are now entries for one or more other RBridge ports on the link that appear to be higher priority to be the DRB than the local port.
- D3. The port is not Down or Suspended, and the adjacency table for the port changes, so there are now no entries for other RBridge ports on the link that appear to be higher priority to be the DRB than the local port.
- D4. A TRILL LAN Hello is received that has the same MAC address (SNPA) as the receiving port and higher priority to be the DRB as described for event A0.
- D5. The port becomes operationally down.

Event D1 is considered to occur on RBridge boot if the port is administratively and link-layer enabled.

Event D4 causes the port to enter the Suspended state and all adjacencies for the port to be discarded (transitioned to the Down state). If the port was in some state other than Suspended, the Suspension Timer is set to the Holding Time in the Hello that causes event D4. If it was in the Suspended state, the Suspension Timer is set to the maximum of its current value and the Holding Time in the Hello that causes event D4.

4.2.1. DRB Election Details

Events D2 and D3 constitute losing and winning the DRB election at the port, respectively.

The candidates for election are the local RBridge and all RBridges with which there is an adjacency on the port in an adjacency state other than the Down state. The winner is the RBridge with highest priority to be the DRB, as determined from the 7-bit priority field in that RBridge's Hellos received and the local port's priority to be the DRB field, with MAC (SNPA) address as a tiebreaker, Port ID as a secondary tiebreaker, and System ID as a tertiary tiebreaker. These fields are compared as unsigned integers, with the larger magnitude being considered higher priority.

Resorting to the secondary and tertiary tiebreakers should only be necessary in rare circumstances when multiple ports have the same priority and MAC (SNPA) address and some of them are not yet suspended. For example, RB1, which has low priority to be the DRB on the link, could receive Hellos from two other ports on the link that have the same MAC address as each other and are higher priority to be the DRB. One of these two ports with the same MAC address will be suspended and cease sending Hellos, and the Hello from it received by RB1 will eventually time out. But, in the meantime, RB1 can use the tiebreakers to determine which port is the DRB and thus which port's Hello to believe for such purposes as setting the Designated VLAN on the link.

4.2.2. Change in DRB

Events D2 and D3 result from a change in the apparent DRB on the link. Unnecessary DRB changes should be avoided, especially on links offering native frame service, as a DRB change will generally cause a transient interruption to native frame service.

If a change in the DRB on the link changes the Designated VLAN on an Ethernet link, the actions specified in Section 4.2.3 are taken.

If an RBridge changes in either direction between being the DRB and not being the DRB at a port, this will generally change the VLANs on which that RBridge sends Hellos through that port, as specified in Section 4.4.3 of [RFC6325].

4.2.3. Change in Designated VLAN

Unnecessary changes in the Designated VLAN on an Ethernet link should be avoided because a change in the Designated VLAN can cause a transient interruption to adjacency and thus to TRILL Data forwarding on the link. When practical, all RBridge ports on a link should be configured with the same Desired Designated VLAN so that if the winner of the DRB election changes for any reason, the Designated VLAN will remain the same.

If an RBridge detects a change in Designated VLAN on an Ethernet link, then, for all adjacency table entries for a port to that link, the RBridge takes the following steps, in the order given.

- o The non-Designated VLAN Hello holding timer is set to the maximum of its time to expiration and the current time to expiration of the Designated VLAN Hello holding timer.
- o The Designated VLAN Hello holding timer is then set to expired (if necessary), and an event A5 occurs for the adjacency (see Section 3.3).

If the Designated VLAN for a link changes, this will generally change the VLANs on which Hellos are sent by an RBridge port on that link as specified in Section 4.4.3 of [RFC6325].

4.3. Port State Table and Diagram

The table below shows the transitions between the DRB states defined above, based on the events defined above:

Event	Down	Suspended	DRB	Not DRB
D1	DRB	DRB	N/A	N/A
D2	N/A	N/A	Not DRB	Not DRB
D3	N/A	N/A	DRB	DRB
D4	N/A	Suspended	Suspended	Suspended
D5	Down	Down	Down	Down

Table 3: Port State Table

"N/A" indicates that the event to the left is not applicable in the state at the top of the column.

The diagram below presents the same information as in the state table:

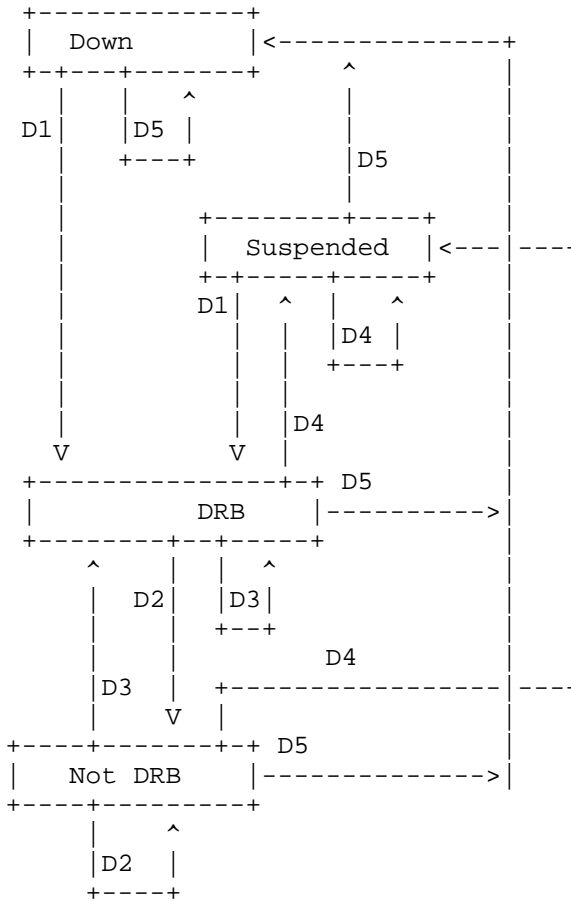


Figure 2: Port State Diagram

5. MTU Matching

The purpose of MTU testing is to ensure that the links used in the campus topology can pass TRILL IS-IS packets, particularly LSP PDUs, at the TRILL campus MTU. The LSP PDUs generated at a TRILL switch could, as part of the flooding process, be sent over any adjacency in the campus. To assure correct operation of IS-IS, an LSP PDU must be able to reach every RBridge in the IS-IS reachable campus; this might be impossible if the PDU exceeded the MTU of an adjacency that was part of the campus topology.

An RBridge, RB1, determines the desired campus link MTU by calculating the minimum of its `originatingLlLSPBufferSize` and the `originatingLlLSPBufferSize` of other RBridges in the campus, as advertised in the link-state database, but not less than 1,470 bytes. Although `originatingLlLSPBufferSize` in Layer 3 [IS-IS] is limited to the range 512 to 1,492 bytes inclusive, in TRILL it is limited to the range 1,470 to 65,535 bytes inclusive. (See Section 5 of [RFC7180].)

Although MTU testing is optional, it is mandatory for an RBridge to respond to an MTU-probe PDU with an MTU-ack PDU [RFC6325] [RFC7176]. The use of multicast or unicast for MTU-probe and MTU-ack is an implementation choice. However, the burden on the link is generally minimized by the following: (1) multicasting MTU-probes when a response from all other RBridges on the link is desired, such as when initializing or reconfirming MTU, (2) unicasting MTU-probes when a response from a single RBridge is desired, such as one that has just been detected on the link, and (3) unicasting all MTU-ack packets.

RB1 can test the MTU size to RB2 as described in Section 4.3.2 of [RFC6325]. For this purpose, MTU testing is only done in the Designated VLAN. An adjacency that fails the MTU test at the campus MTU will not enter the Report state, or, if the adjacency is in that state, it leaves that state. Thus, an adjacency failing the MTU test at the campus minimum MTU will not be reported by the RBridge performing the test. Since inclusion in least-cost route computation requires the adjacency to be reported by both ends, as long as the RBridge at either end of the adjacency notices the MTU failure, it will not be so used.

If RB1 tests MTU size, it reports the largest size for which the MTU test succeeds or a flag indicating that it fails at the campus MTU. This report always appears with the neighbor in RB1's TRILL Neighbor TLV. RB1 MAY also report this with the adjacency in an Extended Reachability TLV in RB1's LSP. RB1 MAY choose to test MTU sizes greater than the desired campus MTU as well as the desired campus MTU.

Most types of TRILL IS-IS packets, such as LSPs, can make use of the campus MTU. The exceptions are TRILL Hellos, which must be kept small for loop safety, and the MTU PDUs, whose size must be adjusted appropriately for the tests being performed.

6. BFD-Enabled TLV and BFD Session Bootstrapping

When the adjacency between R Bridges reaches the 2-Way state, TRILL Hellos will already have been exchanged. If an R Bridge supports BFD [RFC7175], it will have learned whether the other R Bridge has BFD enabled by whether or not a BFD-Enabled TLV [RFC6213] was included in its Hellos. In addition, TRILL Hellos include a nickname of the sending R Bridge [RFC7176] so that information will be available to the receiving R Bridge.

The BFD-Enabled TLVs in TRILL Hellos will look like the following:

```

+-----+
| Type=148          | (1 byte)
+-----+
| Length=3*n       | (1 byte)
+-----+
| RESV |           | MT ID=0           | (2 bytes)
+-----+
| NLPID=0xC0       | (1 byte)
+-----+
| possible additional | (3*(n-1) bytes)
| topology/NLPID pairs
+-----+

```

Figure 3: BFD-Enabled TLV Example/Format

Type = 148 for BFD Enabled [RFC6213].

Length will be 3 times the number of topology and protocol pairs in the TLV.

MT ID is a topology ID [RFC5120] that will be zero unless multi-topology is being supported [MT].

NLPID is a Network Layer Protocol ID [RFC6328] and will be 0xC0 for TRILL, but additional topology and protocol pairs could conceivably be listed.

An R Bridge port initiates a one-hop BFD session with another R Bridge if the following conditions are met: (1) it has BFD enabled, (2) it has an adjacency to another R Bridge in the 2-Way or Report state, and (3) the Hellos it receives indicate that the other R Bridge also has BFD enabled. Either (a) BFD was enabled on both R Bridge ports when the adjacency changed to the 2-Way or Report state, (b) the adjacency was already in the 2-Way or Report state and BFD was enabled on one R Bridge port when BFD had been enabled on the other, or (c) BFD was simultaneously enabled on both R Bridge ports.

In such a BFD session, BFD is encapsulated as specified in [RFC7175]. The egress nickname to be used will have been learned from received Hellos. On a point-to-point link, the Any-RBridge nickname [RFC7180] can also be used as egress, since support of that nickname is required by support of RBridge Channel [RFC7178] and support of RBridge Channel is required for support of BFD over TRILL.

The rare case of transient nickname conflict (due to the network operator configuring a conflict, new connectivity to a previously isolated RBridge, or the like) can cause transient failure of an ongoing BFD session. This can be avoided in the one-hop point-to-point case by using the Any-RBridge egress nickname. In cases where Any-RBridge cannot be used as the egress nickname and a transient nickname conflict is detected for the intended destination of a BFD session, initiation of the session SHOULD be delayed until the conflict is resolved.

If a one-hop BFD session is initiated when the adjacency is in the 2-Way state, the adjacency MUST NOT advance to the Report state until BFD and any other enabled connectivity tests (including MTU, if enabled) have succeeded, as specified in Section 3.

If a one-hop BFD session is established when the adjacency is in the Report state, due to enablement at the RBridges, then, to minimize unnecessary topology changes, the adjacency MUST remain in the Report state unless and until the BFD session (or some other enabled connectivity test) fails.

7. Pseudonodes

This section only applies to broadcast links, as there is no DRB and there cannot be a pseudonode [IS-IS] for a link configured as point-to-point. The Designated RBridge (DRB), determined as described above, controls whether a pseudonode will be used on a link.

If the DRB sets the bypass pseudonode bit in its TRILL LAN Hellos, the RBridges on the link (including the DRB) just directly report all their adjacencies on the LAN that are in the Report state. If the DRB does not set the bypass pseudonode bit in its TRILL Hellos, then (1) the DRB reports in its LSP its adjacency to the pseudonode, (2) the DRB sends LSPs on behalf of the pseudonode in which it reports adjacency to all other RBridges on the link where it sees that adjacency in the Report state, and (3) all other RBridges on the link report their adjacency to the pseudonode if they see their adjacency to the DRB as being in the Report state and do not report any other adjacencies on the link. Setting the bypass pseudonode bit has no effect on how LSPs are flooded on a link. It only affects what LSPs are generated.

It is anticipated that many links between RBridges will actually be point-to-point even in cases where the link technology supports operation as a multi-access broadcast link, in which case using a pseudonode merely adds to the complexity. For example, if RB1 and RB2 are the only RBridges on the link, and RB1 is the DRB, then if RB1 creates a pseudonode -- for example, RB1.25 -- that is used, there are then 3 LSPs: RB1.25, RB1, and RB2, where RB1.25 reports connectivity to RB1 and RB2, and RB1 and RB2 each just say they are connected to RB1.25. However, if DRB RB1 sets the bypass pseudonode bit in its Hellos, then there will be only 2 LSPs: RB1 and RB2, each reporting connectivity to each other.

A DRB SHOULD set the bypass pseudonode bit in its Hellos if it has not seen at least two simultaneous adjacencies in the Report state since it last rebooted or was reset by network management.

8. More TRILL Hello Details

This section provides further details on the receipt, transmission, and content of TRILL Hellos. Unless otherwise stated, it applies to both LAN and point-to-point Hellos.

TRILL Hellos, like all TRILL IS-IS packets, are primarily distinguished from Layer 3 IS-IS packets on Ethernet by being sent to the All-IS-IS-RBridges multicast address (01-80-C2-00-00-41). TRILL IS-IS packets on Ethernet also have the L2-IS-IS Ethertype (0x22F4) and are Ethertype encoded.

Although future extensions to TRILL may include the use of Level 2 IS-IS, [RFC6325] specifies TRILL using a single Level 1 Area using the fixed Area Address zero (see Section 4.2 of [RFC7176]).

IS-IS Layer 3 routers are frequently connected to other Layer 3 routers that are part of a different routing domain. In that case, the externalDomain flag (see [IS-IS]) is normally set for the port through which such a connection is made. The setting of this flag to "true" causes no IS-IS PDUs to be sent out of the port and any IS-IS PDUs received to be discarded, including Hellos. RBridges operate in a different environment where all neighbor RBridges merge into a single campus. For loop safety, RBridges do not implement the externalDomain flag or implement it with the fixed value "false". They send and can receive TRILL Hellos on every port that is not disabled.

8.1. Contents of TRILL Hellos

The table below lists mandatory (M) and optional (O) content TLVs for TRILL Hellos that are particularly relevant to this document, distinguishing between TRILL LAN Hellos and TRILL P2P Hellos. A "-" indicates that an occurrence would be ignored. There are additional TLVs and sub-TLVs that can occur in TRILL Hellos [RFC7176].

LAN	P2P	Number	Content Item
---	---	-----	-----
M	M	1	Area Addresses TLV with Area Address zero only
M	M	1	MT Port Capabilities TLV containing a VLAN-FLAGS sub-TLV
O	O	0-n	Other MT Port Capability TLVs
M	-	0-n	TRILL Neighbor TLV (see Section 8.2.1)
-	M	1	Three-Way Handshake TLV
O	O	0-n	Protocols Supported TLV -- MUST list the TRILL NLPID (0xC0) [RFC6328]
O	O	0-1	bfd-Enabled TLV
O	O	0-1	Authentication TLV
-	-	0-n	Padding TLV -- SHOULD NOT be included

Table 4: TRILL Hello Contents

A TRILL Hello MAY also contain any TLV permitted in a Layer 3 IS-IS Hello. As with all IS-IS PDUs, TLVs that are unsupported/unknown in TRILL Hellos are ignored.

8.2. Transmitting TRILL Hellos

TRILL Hellos are sent with the same timing as Layer 3 IS-IS Hellos [IS-IS]; however, no Hellos are sent if a port is in the Suspended or Down state or if the port is disabled.

TRILL Hello PDUs SHOULD NOT be padded and MUST NOT be sent if they exceed 1,470 bytes; however, a received TRILL Hello longer than 1,470 bytes is processed normally.

TRILL Hello PDU headers MUST conform to the following:

- o Maximum Area Addresses equal to 1.
- o Circuit Type equal to 1.

See Section 8.1 for mandatory Hello TLV contents and some optional Hello TLV contents.

8.2.1. TRILL Neighbor TLVs

A TRILL Neighbor TLV SHOULD NOT be included in TRILL point-to-point Hellos, as it MUST be ignored in that context and wastes space.

TRILL Neighbor TLVs sent in a LAN Hello on an Ethernet link MUST show the neighbor information, as sensed by the transmitting RBridge, for the VLAN on which the Hello is sent. Since implementations conformant to this document maintain such information on a per-VLAN basis only for the Designated VLAN, such implementations only send the TRILL Neighbor TLV in TRILL LAN Hellos in the Designated VLAN.

It is RECOMMENDED that, if there is sufficient room, a TRILL Neighbor TLV or TLVs, as described in Section 4.4.2.1 of [RFC6325], covering the entire range of MAC addresses and listing all adjacencies with a non-zero Designated VLAN Hello Holding Time, or an empty list of neighbors if there are no such adjacencies, be in TRILL Hellos sent on the Designated VLAN. If this is not possible, then TRILL Neighbor TLVs covering sub-ranges of MAC addresses should be sent so that the entire range is covered reasonably promptly. Delays in sending TRILL Neighbor TLVs will delay the advancement of adjacencies to the Report state and the discovery of some link failures. Rapid (for example, sub-second) detection of link or node failures is best addressed with a protocol designed for that purpose, such as BFD (see Section 6).

To ensure that any RBridge RB2 can definitively determine whether RB1 can hear RB2, RB1's neighbor list MUST eventually cover every possible range of IDs, that is, within a period that depends on RB1's policy and not necessarily within any specific period such as its Holding Time. In other words, if X1 is the smallest ID reported in one of RB1's neighbor lists, and the "smallest" flag is not set, then X1 MUST appear in a different neighbor list as well, as the largest ID reported in that fragment. Or lists may overlap, as long as there is no gap, such that some range, say, between Xi and Xj, would never appear in any list.

8.3. Receiving TRILL Hellos

Assuming that a packet is labeled as TRILL IS-IS -- for example, on Ethernet it has the L2-IS-IS Ethertype and the All-IS-IS-RBridges destination multicast address or is so marked by the appropriate code point on other link types such as PPP [RFC6361] or a pseudowire [RFC7173] -- it will be examined to see if it appears to be an IS-IS PDU. If so, and it appears to be a TRILL Hello PDU, the following tests are performed:

- o The type of Hello PDU (LAN or P2P) is compared with the port configuration. If a LAN Hello is received on a port configured to be point-to-point or a P2P Hello is received on a port not configured to be point-to-point, it is discarded.
- o If the Circuit Type field is not 1, the PDU is discarded.
- o If the PDU does not contain an Area Address TLV or it contains an Area Address TLV that is not the single Area Address zero, it is discarded.
- o If the Hello includes a Protocols Supported TLV that does not list the TRILL NLPID (0xC0), it is discarded. It is acceptable if there is no Protocols Supported TLV present.
- o If the Hello does not contain an MT Port Capabilities TLV containing a VLAN-FLAGS sub-TLV [RFC7176], it is discarded.
- o If the maximumAreaAddresses field of the PDU is not 1, it is discarded.
- o If IS-IS authentication is in use on the link and either the PDU has no Authentication TLV or validation of the PDU's Authentication TLV fails, it is discarded.

If none of the rules in the list above cause the packet to be discarded and the packet is parseable, it is assumed to be a well-formed TRILL Hello received on the link. It is treated as an event A0, A1, A2, or A3, based on the criteria listed in Section 3.3.

9. Multiple Ports on the Same Broadcast Link

It is possible for an RBridge RB1 to have multiple ports on the same broadcast (LAN) link that are not in the Suspended state. It is important for RB1 to recognize which of its ports are on the same link. RB1 can detect this condition based on receiving TRILL Hello messages with the same LAN ID on multiple ports.

The DRB election is port-based (see Section 4), and only the Hellos from the elected port can perform certain functions such as dictating the Designated VLAN or whether a pseudonode will be used; however, the election also designates the RBridge with that port as the DRB for the link. An RBridge may choose to load split some tasks among its ports on the link if it has more than one. Section 4.4.4 of [RFC6325] describes when it is safe to do so.

10. IANA Considerations

This document serves as a reference for 'Fail' (Failed MTU test), value 0, in the "TRILL Neighbor TLV NEIGHBOR RECORD Flags" registry. IANA has updated that reference to point to this RFC.

11. Security Considerations

This memo provides improved documentation of some aspects of the TRILL base protocol standard, particularly five aspects of the TRILL adjacency establishment and Hello protocol as listed in Section 1. It does not change the security considerations of the TRILL base protocol as detailed in Section 6 of [RFC6325].

See [RFC7175] for security considerations for BFD, whose use in connection with TRILL adjacency is discussed in this document, particularly Section 6.

Appendix A. Changes from RFC 6327

This document has the following changes from [RFC6327]. It obsoletes [RFC6327].

1. This document extends the TRILL Hello size limit, MTU testing, and state machine to point-to-point links.
2. This document incorporates the updates to [RFC6327] from [RFC7180].
3. The bulk of [RFC6327] was written from the point of view that links between TRILL switches would only be Ethernet. In fact, they could be any technology, such as PPP [RFC6361], pseudowire [RFC7173], or IP [TrillIP]. This replacement document generalizes [RFC6327] to cover such link types.
4. This document includes a specification of one-hop BFD session establishment in connection with adjacency.
5. Numerous editorial changes were incorporated.

Appendix B. Changes to RFC 6325

Section 2 of this document replaces Section 4.4.1 of [RFC6325]. Section 8 of this document replaces Section 4.4.2 of [RFC6325], except for Section 4.4.2.1. The changes in [RFC6325] made by this document include

- Prohibiting the sending of TRILL Hellos out of a port while it is in the Suspended state, and the specification of the Suspended state. ([RFC6325] specifies that Hellos be sent with the same timing as [IS-IS].)
- Permitting the inclusion of the Three-Way-Handshake TLV, BFD-Enabled TLV, and other TLVs in TRILL Hellos when these were omitted in TRILL Hello contents lists in Section 4.4.2 of [RFC6325].
- Extending the TRILL Hello protocol to support point-to-point and non-Ethernet links.

Normative References

- [IS-IS] ISO/IEC 10589:2002, Second Edition, "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)", 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.
- [RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", RFC 6213, April 2011.
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.
- [RFC6328] Eastlake 3rd, D., "IANA Considerations for Network Layer Protocol Identifiers", BCP 164, RFC 6328, July 2011.
- [RFC7175] Manral, V., Eastlake 3rd, D., Ward, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Bidirectional Forwarding Detection (BFD) Support", RFC 7175, May 2014.
- [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, May 2014.
- [RFC7178] Eastlake 3rd, D., Manral, V., Li, Y., Aldrin, S., and D. Ward, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Support", RFC 7178, May 2014.
- [RFC7180] Eastlake 3rd, D., Zhang, M., Ghanwani, A., Manral, V., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", RFC 7180, May 2014.

Informative References

- [802.1AX] "IEEE Standard for Local and metropolitan area networks--Link Aggregation", 802.1AX-2008, January 2008.
- [802.1Q] "IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", 802.1Q-2011, August 2011.
- [ESADI] Zhai, H., Hu, F., Perlman, R., Eastlake 3rd, D., and O. Stokes, "TRILL: ESADI (End Station Address Distribution Information) Protocol", Work in Progress, April 2014.
- [FCoE] Excerpt of discussion of "FCoE Max Size" generated from T11/09-251v1, 04/27/2009, "FCoE frame or FCoE PDU", <www.t11.org>.
- [MT] Eastlake, D., Zhang, M., Banerjee, A., and V. Manral, "TRILL: Multi-Topology", Work in Progress, September 2013.
- [RFC3719] Parker, J., Ed., "Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)", RFC 3719, February 2004.
- [RFC6327] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Dutt, D., and V. Manral, "Routing Bridges (RBridges): Adjacency", RFC 6327, July 2011.
- [RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011.
- [RFC6439] Perlman, R., Eastlake, D., Li, Y., Banerjee, A., and F. Hu, "Routing Bridges (RBridges): Appointed Forwarders", RFC 6439, November 2011.
- [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, May 2014.
- [RFC7173] Yong, L., Eastlake 3rd, D., Aldrin, S., and J. Hudson, "Transparent Interconnection of Lots of Links (TRILL) Transport Using Pseudowires", RFC 7173, May 2014.
- [TrillIP] Wasserman, M., Eastlake, D., and D. Zhang, "Transparent Interconnection of Lots of Links (TRILL) over IP", Work in Progress, March 2014.

Acknowledgements

The suggestions and comments of the following are gratefully acknowledged:

Stewart Bryant, Elwyn Davies, Adrian Farrel, Brian Haberman, Jon Hudson, Arnel Lim, and Gayle Noble.

The authors of [RFC6327] included Dinesh Dutt. The acknowledgements section of [RFC6327] includes the following, listed in alphabetic order: Jari Arkko, Ayan Banerjee, Les Ginsberg, Sujay Gupta, David Harrington, Pete McCann, Erik Nordmark, and Mike Shand.

Authors' Addresses

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

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

Radia Perlman
Intel Labs
2200 Mission College Blvd.
Santa Clara, CA 95054
USA

Phone: +1-408-765-8080
EMail: Radia@alum.mit.edu

Anoop Ghanwani
Dell
5450 Great America Parkway
Santa Clara, CA 95054
USA

EMail: anoop@alumni.duke.edu

Howard Yang
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

EMail: howardy@cisco.com

Vishwas Manral
Ionos Corp.
4100 Moorpark Ave.
San Jose, CA 95117
USA

EMail: vishwas@ionosnetworks.com