

Internet Engineering Task Force (IETF)
Request for Comments: 7357
Updates: 6325
Category: Standards Track
ISSN: 2070-1721

H. Zhai
F. Hu
ZTE
R. Perlman
Intel Labs
D. Eastlake 3rd
Huawei
O. Stokes
Extreme Networks
September 2014

Transparent Interconnection of Lots of Links (TRILL):
End Station Address Distribution Information (ESADI) Protocol

Abstract

The IETF TRILL (Transparent Interconnection of Lots of Links) protocol provides least-cost pair-wise data forwarding without configuration in multi-hop networks with arbitrary topologies and link technologies. TRILL supports multipathing of both unicast and multicast traffic. Devices that implement the TRILL protocol are called TRILL switches or RBridges (Routing Bridges).

ESADI (End Station Address Distribution Information) is an optional protocol by which a TRILL switch can communicate, in a Data Label (VLAN or fine-grained label) scoped way, end station address and reachability information to TRILL switches participating in ESADI for the relevant Data Label. This document updates RFC 6325, specifically the documentation of the ESADI protocol, and is not backwards compatible.

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/rfc7357>.

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	5
2. ESADI Protocol Overview	6
2.1. ESADI Virtual Link	10
2.2. ESADI Neighbor Determination	10
2.3. ESADI Payloads	11
3. ESADI DRB (Designated RBridge) Determination	11
4. ESADI PDU Processing	12
4.1. Unicasting ESADI PDUs	12
4.2. General Transmission of ESADI PDUs	13
4.3. General Receipt of ESADI PDUs	14
4.4. ESADI Reliable Flooding	14
5. End Station Addresses	15
5.1. Learning Confidence Level	15
5.2. Forgetting End Station Addresses	16
5.3. Duplicate MAC Address	16
6. ESADI-LSP Contents	18
6.1. ESADI Parameter Data	19
6.2. MAC-Reachability TLV	20
6.3. Default Authentication	21
7. IANA Considerations	21
7.1. ESADI Participation and Capability Flags	22
7.2. TRILL GENINFO TLV	23
8. Security Considerations	24
8.1. Privacy Considerations	25
9. Acknowledgements	26
10. References	26
10.1. Normative References	26
10.2. Informative References	28
Appendix A. Interoperability and Changes to RFC 6325	29
A.1. ESADI PDU Changes	29
A.2. Unicasting Changes	30
A.3. Message Timing Changes and Suggestions	30
A.4. Duplicate Address Reachability	30

1. Introduction

The TRILL (Transparent Interconnection of Lots of Links) protocol [RFC6325] provides least-cost pair-wise data forwarding without configuration in multi-hop networks with arbitrary topologies and link technologies, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. TRILL accomplishes this with the IS-IS (Intermediate System to Intermediate System) [IS-IS] [RFC1195] [RFC7176] link-state routing protocol using a header with a hop count. The design supports optimization of the distribution of multi-destination frames and two types of data labeling: VLANs (Virtual Local Area Networks) [RFC6325] and FGLs (fine-grained labels) [RFC7172]. Devices that implement TRILL are called TRILL switches or RBridges (Routing Bridges).

There are five ways a TRILL switch can learn end station addresses, as described in Section 4.8 of [RFC6325]. One of these is the ESADI (End Station Address Distribution Information) protocol, which is an optional Data Label scoped way by which TRILL switches can communicate with each other information such as end station addresses and their TRILL switch of attachment. A TRILL switch that is announcing interest in a Data Label MAY use the ESADI protocol to announce the end station address of some or all of its attached end stations in that Data Label to other TRILL switches that are running ESADI for that Data Label. (In the future, ESADI may also be used for other address and reachability information.)

By default, TRILL switches with connected end stations learn addresses from the data plane when ingressing and egressing native frames, although such learning can be disabled. The ESADI protocol's potential advantages over data plane learning include the following:

1. Security advantages:

- a) The ESADI protocol can be used to announce end stations with an authenticated enrollment (for example, enrollment authenticated by cryptographically based EAP (Extensible Authentication Protocol) [RFC3748] methods via [802.1X]).
- b) The ESADI protocol supports cryptographic authentication of its message payloads for more secure transmission.

2. Fast update advantages: The ESADI protocol provides a fast update of end station MAC (Media Access Control) addresses and their TRILL switch of attachment. If an end station is unplugged from one TRILL switch and plugged into another, ingressed frames with that end station's MAC address as their destination can be

black-holed. That is, they can be sent just to the older egress TRILL switch that the end station was connected to until cached address information at some remote ingress TRILL switch times out, possibly for tens of seconds [RFC6325].

MAC address reachability information, some ESADI parameters, and optional authentication information are carried in ESADI packets rather than in the TRILL IS-IS protocol. As specified below, ESADI is, for each Data Label, a virtual logical topology overlay in the TRILL topology. An advantage of using ESADI over using TRILL IS-IS is that the end station attachment information is not flooded to all TRILL switches but only to TRILL switches advertising ESADI participation for the Data Label in which those end stations occur.

1.1. Content and Precedence

This document updates [RFC6325], the TRILL base protocol specification, replacing the description of the TRILL ESADI protocol (Section 4.2.5 of [RFC6325], including all subsections), providing more detail on ESADI, updating other ESADI-related sections of [RFC6325], and prevailing over [RFC6325] in any case where they conflict. For this reason, familiarity with [RFC6325] is particularly assumed. These changes include a change to the format of ESADI-LSPs (ESADI Link State Protocol Data Units) that is not backwards compatible; this change is justified by the substantially increased amount of information that can be carried and in light of the very limited, if any, deployment of RFC 6325 ESADI. These changes are further discussed in Appendix A.

Section 2 of this document is the ESADI protocol overview. Section 3 specifies ESADI DRB (Designated RBridge) determination. Section 4 discusses the processing of ESADI PDUs. Section 5 discusses interaction with other modes of end station address learning. Section 6 describes the ESADI-LSP and its contents.

1.2. Terminology

This document uses the acronyms defined in [RFC6325], in addition to the following:

Data Label:	VLAN or FGL.
ESADI RBridge:	An RBridge that is participating in ESADI for one or more Data Labels.
FGL:	Fine-Grained Label [RFC7172].
LSP:	Link State PDU [IS-IS].

LSP number zero: A Link State PDU with fragment number equal to zero.

PDU: Protocol Data Unit.

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].

Capitalized IANA-related terms such as "IETF Review" are to be interpreted as described in [RFC5226].

2. ESADI Protocol Overview

ESADI is a Data Label scoped way for TRILL switches (also known as RBridges) to announce and learn end station addresses rapidly and securely. An RBridge that is announcing participation in ESADI for one or more Data Labels is called an ESADI RBridge.

ESADI is an optional protocol that is separate from the mandatory TRILL IS-IS implemented by all RBridges in a campus. There is a separate ESADI instance for each Data Label (VLAN or FGL) if ESADI is being used for that Data Label. In essence, for each such Data Label, there is a modified instance of the IS-IS reliable flooding mechanism in which ESADI RBridges may choose to participate. (These are not the instances specified in [RFC6822].) Multiple ESADI instances may share implementation components within an RBridge as long as that sharing preserves the independent operation of each instance of the ESADI protocol. For example, the ESADI link state database could be a single database with a field in each record indicating the Data Label to which it applies, or it could be a separate database per Data Label. However, the ESADI update process operates separately for each ESADI instance and independently from the TRILL IS-IS update process.

ESADI does no routing calculations, so there is no reason for pseudonodes in ESADI and none are created. (Pseudonodes [IS-IS] are a construct for optimizing routing calculations.) Furthermore, a relatively large amount of ESADI data will have to be distributed, under some circumstances, using ESADI mechanisms; this would require a large number of ESADI-LSP fragments. ESADI-LSP, ESADI-CSNP, and ESADI-PSNP (ESADI Link State PDU, Complete Sequence Number PDU, and Partial Sequence Number PDU) payloads are therefore formatted as Extended Level 1 Circuit Scope (E-L1CS) PDUs [RFC7356] (see also Section 6). This allows up to 2**16 fragments but does not support link state data associated with pseudonodes.

After the TRILL Header, ESADI packets have an inner Ethernet header with the Inner.MacDA of "All-Egress-RBridges" (formerly called "All-ESADI-RBridges"), an inner Data Label specifying the VLAN or FGL of interest, and the "L2-IS-IS" Ethertype followed by the ESADI payload, as shown in Figure 1.


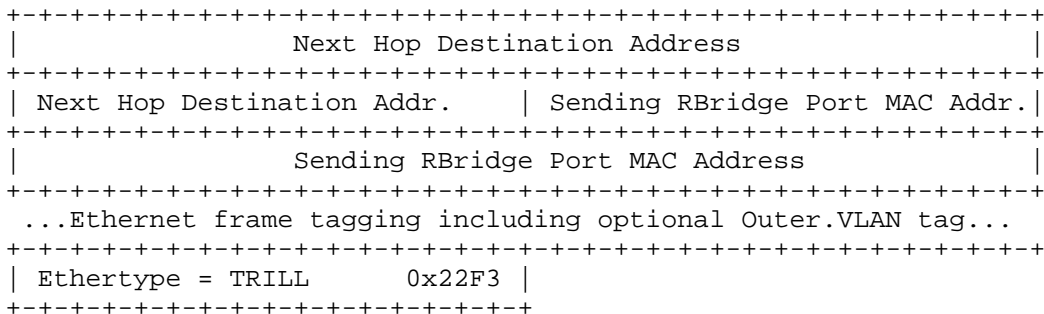



Figure 1: TRILL ESADI Packet Overview

TRILL ESADI packets sent on an Ethernet link are structured as shown in Figure 2. The outer VLAN tag will not be present if it was not included by the Ethernet port that sent the packet.

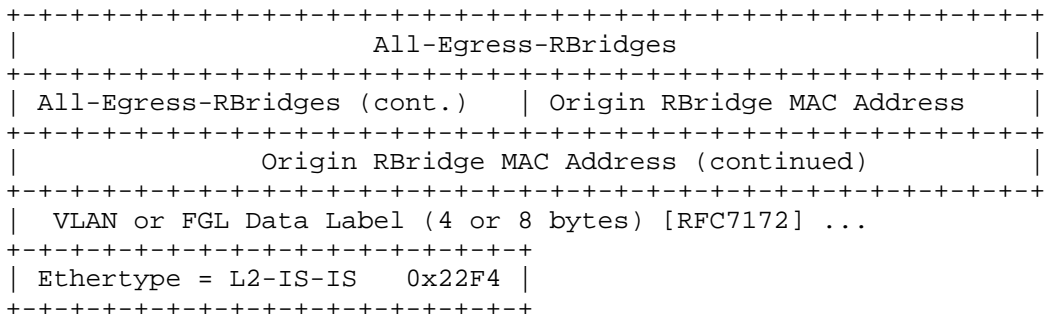
Outer Ethernet Header:




TRILL Header:



Inner Ethernet Header:



ESADI Payload (formatted as IS-IS):



Frame Check Sequence:




Figure 2: ESADI Ethernet Link Packet Format

The Next Hop Destination Address or Outer.MacDA is the All-RBridges multicast address if the ESADI PDU is being multicast. If it is being unicast, the Next Hop Destination Address is the unicast address of the next-hop RBridge. The VLAN for the Outer.VLAN

information, if present, will be the Designated VLAN for the link on which the packet is sent. The V and R fields will be zero while the M bit will be one, unless the ESADI PDU was unicast, in which case the M bit will be zero. The Data Label specified will be the VLAN or FGL to which the ESADI packet applies. The Origin RBridge MAC Address or Inner.MacSA MUST be a MAC address unique across the campus owned by the RBridge originating the ESADI packet -- for example, any of its port MAC addresses if it has any Ethernet ports -- and each ESADI RBridge MUST use the same Inner.MacSA for all of the ESADI packets it originates.

TRILL ESADI packets sent on a PPP link are structured as shown in Figure 3 [RFC6361].

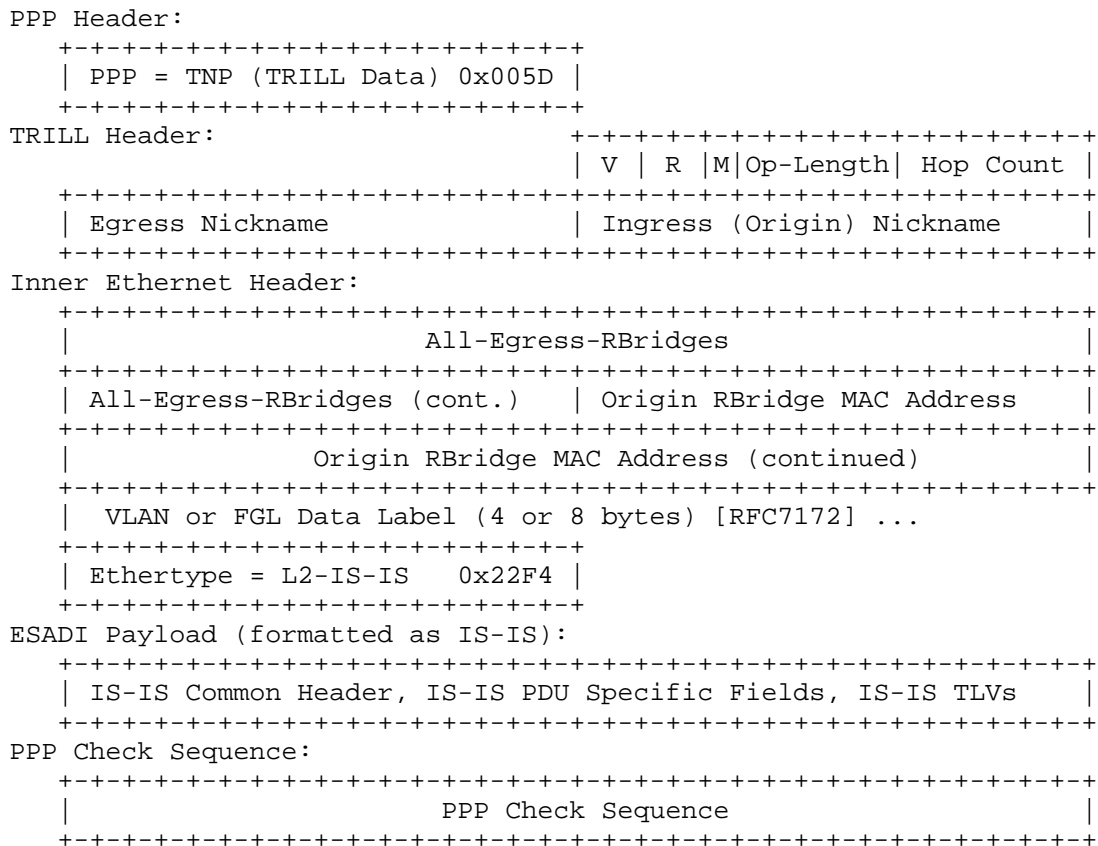


Figure 3: ESADI PPP Link Packet Format

2.1. ESADI Virtual Link

All RBridges forward ESADI packets as if they were ordinary TRILL Data packets. Because of this forwarding, it appears to an instance of the ESADI protocol at an RBridge that it is directly connected by a multi-access virtual link to all RBridges in the campus that are "data reachable" from it (see Section 2 of [RFC7180]) and are running ESADI for that Data Label. No "routing" calculation (least-cost path or distribution tree construction) ever has to be performed by ESADI. An ESADI RBridge merely transmits the ESADI packets it originates on this virtual link as described for TRILL Data packets in [RFC6325] and [RFC7172]. For multicast ESADI packets, it may use any distribution tree that it might use for an ordinary multi-destination TRILL Data packet. RBridges that do not implement the ESADI protocol, do not have it enabled, or are not participating in the ESADI protocol for the Data Label of an ESADI packet do not decapsulate or locally process the ESADI packet. Thus, ESADI packets are transparently tunneled through transit RBridges.

2.2. ESADI Neighbor Determination

The ESADI instance for Data Label X at an RBridge RB1 determines who its adjacent ESADI neighbors are by examining the TRILL IS-IS link state database for RBridges that are data reachable from RB1 (see Section 2 of [RFC7180]) and are announcing their participation in Data Label X ESADI. When an RBridge RB2 becomes data unreachable from RB1 or the relevant entries for RB2 are purged from the core IS-IS link state database, it is lost as a neighbor and also dropped from any ESADI instances from the point of view of RB1, and when RB2 is no longer announcing participation in Data Label X ESADI, it ceases to be a neighbor for any Data Label X ESADI instance. All these considerations are Data Label scoped. Because of these mechanisms whereby an ESADI instance at an ESADI RBridge can determine its ESADI adjacencies by examining the TRILL IS-IS link state database, there are no "Hellos" sent in ESADI and no adjacency information is carried in ESADI-LSPs.

A participation announcement in a VLAN scoped ESADI instance is generated by setting a flag bit in the Interested VLANs sub-TLV, and an announcement for an FGL scoped ESADI instance is generated by setting a flag bit in the Interested Labels sub-TLV [RFC7176] (see Section 7.1).

2.3. ESADI Payloads

TRILL ESADI packet payloads are structured like IS-IS Extended Level 1 Circuit Scope (E-L1CS) LSP, CSNP, and PSNP PDUs [RFC7356], except as indicated below, but are always TRILL encapsulated on the wire as if they were TRILL Data packets. The information distributed by the ESADI protocol includes a list of local end station MAC addresses connected to the originating RBridge and, for each such address, a 1-octet unsigned "Confidence" rating in the range 0-254 (see Section 6.2). It is entirely up to the originating RBridge which locally connected MAC addresses it wishes to advertise via ESADI and with what Confidence. It MAY advertise all, some, or none of such addresses. In addition, some ESADI parameters of the advertising RBridge (see Section 6.1) and, optionally, authentication information (see Section 6.3) are included. Future uses of ESADI may distribute other similar address and reachability information.

TRILL ESADI-LSPs MUST NOT contain a Data Label ID in their payload. The Data Label to which the ESADI data applies is the Data Label of the TRILL Data packet enclosing the ESADI payload. If a Data Label ID could occur within the payload, it might conflict with that TRILL Data packet Data Label and could conflict with any future Data Label mapping scheme that may be adopted [VLANmapping]. If a VLAN or FGL ID field within an ESADI-LSP PDU does include a value, that field's contents MUST be ignored.

3. ESADI DRB (Designated RBridge) Determination

Because ESADI does no adjacency announcement or routing, the ESADI-DRB never creates a pseudonode. However, a DRB [RFC7177] is still needed to issue ESADI-CSNP PDUs and respond to ESADI-PSNP PDUs for ESADI-LSP synchronization.

Generally speaking, the DRB election on the ESADI virtual link (see Section 2.1) operates similarly to the DRB election on a TRILL IS-IS broadcast link, as described in Section 4.2.1 ("DRB Election Details") of [RFC7177], with the following exceptions: in the Data Label X ESADI-DRB election at RB1 on an ESADI virtual link, the candidates are the local ESADI instance for Data Label X and all remote ESADI instances at RBridges that are (1) data reachable from RB1 [RFC7180] and (2) announcing in their TRILL IS-IS LSP that they are participating in ESADI for Data Label X. The winner is the instance with the highest ESADI Parameter 7-bit priority field with ties broken by the System ID, comparing fields as unsigned integers with the larger magnitude considered higher priority. "SNPA/MAC address" (Subnetwork Point of Attachment / MAC address) is not considered in this tiebreaking, and there is no "Port ID".

4. ESADI PDU Processing

Data Label X ESADI neighbors are usually not connected directly by a physical link but are always logically connected by a virtual link (see Section 2.1). There could be hundreds or thousands of ESADI RBridges (TRILL switches) on the virtual link. The only PDUs used in ESADI are the ESADI-LSP, ESADI-CSNP, and ESADI-PSNP PDUs. In particular, there are no Hello or MTU PDUs, because ESADI does not build a topology, does not do any routing calculations, and does not determine MTU. Instead, ESADI uses the distribution trees and the Sz campus minimum link MTU determined by the core TRILL IS-IS (see [RFC6325] and [RFC7180]).

4.1. Unicasting ESADI PDUs

For [IS-IS], PDU multicasting is normal on a local link and no effort is made to optimize to unicast, because on the typical physical link for which IS-IS was designed (commonly a piece of multi-access Ethernet cable), any frame made the link busy for that frame time. However, to ESADI instances, what appears to be a simple multi-access link is generally a set of multi-hop distribution trees that may or may not be pruned. Thus, transmitting a multicast frame on such a tree can impose a substantially greater load than transmitting a unicast frame. This load may be justified if there are likely to be multiple listeners but may not be justified if there is only one recipient of interest. For this reason, under some circumstances, ESADI PDUs MAY be TRILL unicast if it is confirmed that the destination RBridge supports receiving unicast ESADI PDUs (see Section 6.1).

The format of a unicast ESADI packet is the format of a multicast TRILL ESADI packet as described in Section 2 above, except as follows:

- o On an Ethernet link, in the outer Ethernet header the Outer.MacDA is the unicast address of the next-hop RBridge.
- o In the TRILL Header, the M bit is set to zero and the Egress Nickname is the nickname of the destination RBridge.

To support unicasting of ESADI PDUs, Section 4.6.2.2 of [RFC6325] is replaced with the following:

4.6.2.2. TRILL ESADI Packets

If $M = 1$, the egress nickname designates the distribution tree. The packet is forwarded as described in Section 4.6.2.5. In addition, if (1) the forwarding RBridge is interested in the specified VLAN or fine-grained label [RFC7172], (2) the forwarding RBridge implements the TRILL ESADI protocol, and (3) ESADI is enabled for the specified VLAN or fine-grained label, then the inner frame is decapsulated and provided to that local ESADI protocol.

If $M = 0$ and the egress nickname is not that of the receiving RBridge, the packet is forwarded as for known unicast TRILL Data frames as described in Section 4.6.2.4. If $M = 0$ and the egress nickname is that of the receiving RBridge, and the receiving RBridge supports unicast ESADI PDUs, then the ESADI packet is decapsulated and processed if it meets the three numbered conditions in the paragraph above; otherwise, it is discarded.

The references to "4.6.2.2", "4.6.2.4", and "4.6.2.5" above refer to those sections in [RFC6325].

4.2. General Transmission of ESADI PDUs

Following the usual [IS-IS] rules, an ESADI instance does not transmit any ESADI PDUs if it has no ESADI adjacencies. Such transmission would just be a waste of bandwidth.

The MTU available to ESADI payloads is at least 24 bytes less than that available to TRILL IS-IS because of the additional fields required ($2(\text{TRILL Ethertype}) + 6(\text{TRILL Header}) + 6(\text{Inner.MacDA}) + 6(\text{Inner.MacSA}) + 4/8(\text{Data Label})$ bytes). Thus, the inner ESADI payload, starting with the Intradomain Routing Protocol Discriminator byte, MUST NOT exceed Sz minus 24 for a VLAN ESADI instance or Sz minus 28 for an FGL ESADI instance; however, if a larger payload is received, it is processed normally (see [RFC6325] and [RFC7180] for discussions of Sz and MTU).

In all cases where this document says that an ESADI PDU is multicast, if the transmitting RBridge has only one neighbor and that neighbor advertises support for unicast, the PDU MAY be unicast (see Section 4.1).

A priority bit to indicate that an LSP fragment should be flooded with high priority is provided by [RFC7356]. This bit SHOULD be set on ESADI-LSP fragment zero because it is important that the ESADI Parameter APPsub-TLV get through promptly. This bit SHOULD NOT be set on other ESADI-LSP fragments to avoid giving undue priority to less urgent PDUs.

4.3. General Receipt of ESADI PDUs

In contrast with Layer 3 IS-IS PDU acceptance tests, which check the source inner and outer SNPA/MAC in order to verify that a PDU is from an adjacent TRILL switch, in TRILL ESADI adjacency is based on the system ID, so the system ID inside the PDU is all that is tested for.

If an ESADI instance believes that it has no ESADI neighbors, it ignores any ESADI PDUs it receives.

4.4. ESADI Reliable Flooding

The IS-IS reliable flooding mechanism (the Update Process) is modified for ESADI in the ways listed below. Except as otherwise stated, the ESADI update process works as described in [IS-IS], [RFC1195], and [RFC7356].

When an ESADI instance sees that it has a new ESADI neighbor, its self-originated ESADI-LSP fragments are scheduled to be sent and MAY be unicast to that neighbor if the neighbor is announcing in its LSP that it supports unicast ESADI (see Section 6.1). If all the other ESADI instances for the same Data Label send their self-originated ESADI-LSPs immediately, there may be a surge of traffic to that new neighbor. Therefore, the ESADI instances SHOULD wait an interval of time before sending their ESADI-LSP(s) to a new neighbor. The interval time value is up to the device implementation. One suggestion is that the interval time can be assigned a random value with a range based on the RBridge's nickname (or any one of its nicknames, if it holds more than one), such as $(2000 * \text{nickname} / 2^{**16})$ milliseconds, assuming "nickname" to be an unsigned quantity.

All the TRILL switches participating in an ESADI instance for some Data Label appear to ESADI to be adjacent. Thus, the originator of any active ESADI-LSP fragment always appears to be on link and, to spread the burden of such a response, could be the RBridge to respond to any ESADI-CSNP or PSNP request for that fragment. However, under very rare circumstances, it could be that some version of the LSP fragment with a higher sequence number is actually held by another ESADI RBridge on the link, so non-originators need to be able to respond eventually. Thus, when the receipt of a CSNP or PSNP causes the SRMflag (Send Routing Message flag [IS-IS]) to be set for an LSP

fragment, action is as specified in [IS-IS] for the originating ESADI RBridge of the fragment; however, at a non-originating ESADI RBridge, when changing the SRMflag from 0 to 1, the lastSent timestamp [IS-IS] is also set to the current time minus

$$\text{minimumLSPTransmissionInterval} * \text{Random (Jitter)} / 100$$

(where minimumLSPTransmissionInterval, Random, and Jitter are as in [IS-IS]). This will delay and jitter the transmission of the LSP fragment by non-originators. This gives the originator more time to send the fragment and provides more time for such an originator-transmitted copy to traverse the likely multi-hop path to non-originators and clear the SRMflag for the fragment at non-originators.

The multi-hop distribution tree method with Reverse Path Forwarding Check used for multicast distribution by TRILL will typically be less reliable than transmission over a single local broadcast link hop. For LSP synchronization robustness, in addition to sending ESADI-CSNPs as usual when it is the DRB, an ESADI RBridge SHOULD also transmit an ESADI-CSNP for an ESADI instance if all of the following conditions are met:

- o it sees one or more ESADI neighbors for that instance, and
- o it does not believe it is the DRB for the ESADI instance, and
- o it has not received or sent an ESADI-CSNP PDU for the instance for the average of the CSNP Time (see Section 6.1) of the DRB and its CSNP Time.

5. End Station Addresses

The subsections below discuss end station address considerations in the context of ESADI.

5.1. Learning Confidence Level

The Confidence level mechanism [RFC6325] allows an RBridge campus manager to cause certain address learning sources to prevail over others. MAC address information learned through a registration protocol, such as [802.1X] with a cryptographically based EAP [RFC3748] method, might be considered more reliable than information learned through the mere observation of data traffic. When such authenticated learned address information is transmitted via the ESADI protocol, the use of authentication in the TRILL ESADI-LSP packets could make tampering with it in transit very difficult. As a result, it might be reasonable to announce such authenticated

information via the ESADI protocol with a high Confidence, so it would be used in preference to any alternative learning from data observation.

5.2. Forgetting End Station Addresses

The end station addresses learned through the TRILL ESADI protocol should be forgotten through changes in ESADI-LSPs. The timeout of the learned end station address is up to the originating RBridge that decides when to remove such information from its ESADI-LSPs (or up to ESADI protocol timeouts if the originating RBridge becomes unreachable).

If RBridge R_{Bn} participating in the TRILL ESADI protocol for Data Label X no longer wishes to participate in ESADI, it ceases to participate by (1) clearing the ESADI Participation bit in the appropriate Interested VLANs or Interested Labels sub-TLV and (2) sending a final ESADI-LSP nulling out its ESADI-LSP information.

5.3. Duplicate MAC Address

With ESADI, it is possible to persistently see occurrences of the same MAC address in the same Data Label being advertised as reachable by two or more RBridges. The specification of how to handle this situation in [RFC6325] is updated by this document, by replacing the last sentence of the last paragraph of Section 4.2.6 of [RFC6325] as shown below to provide better traffic-spreading while avoiding possible address flip-flopping.

As background, assume some end station or set of end stations E_{Sn} have two or more ports with the same MAC address in the same Data Label with the ports connected to different RBridges (RB₁, RB₂, ...) by separate links. With ESADI, some other RBridge, RB₀, can persistently see that MAC address in that Data Label connected to multiple RBridges. When RB₀ ingresses a frame, say from E_{S0}, destined for that MAC and label, the current [RFC6325] text permits a wide range of behavior. In particular, [RFC6325] would permit RB₀ to use some rule, such as "always encapsulate to the egress with the lowest System ID", which would put all of this traffic through only one of the egress RBridges and one of the end station ports. With that behavior, there would be no load-spreading, even if there were multiple different ingress RBridges and/or different MAC addresses with the same reachability. [RFC6325] would also permit RB₀ to send different traffic to different egresses by doing ECMP (Equal Cost Multipath) at a flow level, which would likely result in return traffic for RB₀ to egress to E_{S0} from various of RB₁, RB₂, ... for the same MAC and label. The resulting address reachability flip-flopping perceived at RB₀ could cause problems.

This update to [RFC6325] avoids these potential difficulties by requiring that RB0 use one of the following two policies: (1) only encapsulate to one egress RBridge for any particular MAC and label, but select that egress pseudorandomly, based on the topology (including MAC reachability) or (2) if RB0 will not be disturbed by the returning TRILL Data packets showing the same MAC or by label flip-flopping between different ingresses, RB0 may use ECMP. Assuming multiple ingress RBridges and/or multiple MAC and label addresses, strategy 1 should result in load-spreading without address flip-flopping, while strategy 2 will produce better load-spreading than strategy 1 but with address flip-flopping from the point of view of RB0.

OLD [RFC6325] Section 4.2.6 text:

"... If confidences are also tied between the duplicates, for consistency it is suggested that RB2 direct all such frames (or all such frames in the same ECMP flow) toward the same egress RBridge; however, the use of other policies will not cause a network problem since transit RBridges do not examine the Inner.MacDA for known unicast frames."

NEW [RFC6325] Section 4.2.6 text:

"... If confidences are also tied between the duplicates, then RB2 MUST adopt one of the following two strategies:

1. In a pseudorandom way [RFC4086], select one of the egress RBridges that is least cost from RB2 and to which the destination MAC appears to be attached, and send all traffic for the destination MAC and VLAN (or FGL [RFC7172]) to that egress. This pseudorandom choice need only be changed when there is a change in campus topology or MAC attachment information. Such pseudorandom selection will, over a population of ingress RBridges, probabilistically spread traffic over the possible egress RBridges. Reasonable inputs to the pseudorandom selection are the ingress RBridge System ID and/or nickname, the VLAN or FGL, the destination MAC address, and a vector of the RBridges with connectivity to that MAC and VLAN or FGL. There is no need for different RBridges to use the same pseudorandom function.

As an example of such a pseudorandom function, if there are k egress RBridges ($RB_0, RB_1, \dots, RB_{(k-1)}$) all reporting attachment to address MAC_x in Data Label DL_y , then an ingress RBridge RB_{in} could select the one to which it will send all unicast TRILL Data packets addressed to MAC_x in DL_y based on the following:

$$FNV-32(RB_{in} \mid MAC_x \mid DL_y \mid RB_0 \mid RB_1 \mid \dots \mid RB_{(k-1)}) \bmod k$$

where the FNV (Fowler/Noll/Vo) algorithm is specified in [FNV], RB_x means the nickname for RBridge RB_x , " \mid " means concatenation, MAC_x is the destination MAC address, DL_y is the Data Label, and " $\bmod k$ " means the integer division remainder of the output of the FNV-32 function considered as a positive integer divided by k .

2. If RB_2 supports ECMP and will not be disturbed by return traffic from the same MAC and VLAN (or FGL [RFC7172]) coming from a variety of different RBridges, then it MAY send traffic using ECMP at the flow level to the egress RBridges that are least cost from RB_2 and to which the destination MAC appears to be attached."

6. ESADI-LSP Contents

The only PDUs used in ESADI are the ESADI-LSP, ESADI-CSNP, and ESADI-PSNP PDUs. Currently, the contents of an ESADI-LSP consist of zero or more MAC-Reachability TLVs, optionally an Authentication TLV, and exactly one ESADI parameter APPsub-TLV. Other similar data may be included in the future and, as in [IS-IS], an ESADI instance ignores any TLVs or sub-TLVs it does not understand. Because these PDUs are formatted as Extended Level 1 Circuit Scope (E-L1CS) PDUs [RFC7356], the Type and Length fields in the TLVs are 16-bit.

This section specifies the format for the ESADI Parameter APPsub-TLV, gives the reference for the ESADI MAC-Reachability TLV, and discusses default authentication configuration.

For robustness, the payload for an ESADI-LSP number zero and any ESADI-CSNP or ESADI-PSNP covering fragment zero MUST NOT exceed 1470 minus 24 bytes in length (1446 bytes) if it has an Inner.VLAN, or 1470 minus 28 bytes (1442 bytes) if it has an Inner.FGL. However, if an ESADI-LSP number zero or such an ESADI-CSNP or ESADI-PSNP is received that is longer, it is still processed normally. (As stated in Section 4.3.1 of [RFC6325], 1470 bytes was chosen to make it extremely unlikely that a TRILL control packet, even with reasonable additional headers, tags, and/or encapsulation, would encounter MTU problems on an inter-RBridge link.)

6.1. ESADI Parameter Data

Figure 4 presents the format of the ESADI parameter data. This APPsub-TLV MUST be included in a TRILL GENINFO TLV in ESADI-LSP number zero. If it is missing from ESADI-LSP number zero or if ESADI-LSP number zero is not known, priority for the sending RBridge defaults to 0x40 and CSNP Time defaults to 30. If there is more than one occurrence in ESADI-LSP number zero, the first occurrence will be used. Occurrences of the ESADI Parameter APPsub-TLV in non-zero ESADI-LSP fragments are ignored.

```

+-----+
| Type                                     | (2 bytes)
+-----+
| Length                                   | (2 bytes)
+-----+
|R| Priority                               | (1 byte)
+-----+
| CSNP Time                               | (1 byte)
+-----+
| Flags                                   | (1 byte)
+-----+
| Reserved for expansion                   | (variable)
+-----+
+-----+

```

Figure 4: ESADI Parameter APPsub-TLV

Type: Set to ESADI-PARAM sub-TLV (TRILL APPsub-TLV type 0x0001).
Two bytes, because this APPsub-TLV appears in an extended TLV [RFC7356].

Length: Variable, with a minimum of 3, but must fit within the ESADI packet. This field is encoded as an unsigned integer in network byte order [RFC7356].

R: A reserved bit that MUST be sent as zero and ignored on receipt.

Priority: Gives the originating RBridge's priority for being the DRB on the ESADI instance virtual link (see Section 3) for the Data Label in which the PDU containing the parameter data was sent. It is an unsigned 7-bit integer with the larger magnitude indicating higher priority. It defaults to 0x40 for an RBridge participating in ESADI for which it has not been configured.

CSNP Time: An unsigned byte that gives the amount of time in seconds during which the originating RBridge, if it is the DRB on the ESADI virtual link, will send at least three ESADI-CSNP PDUs. It defaults to 30 seconds for an RBridge participating in ESADI for which it has not been configured.

Flags: A byte of flags associated with the originating ESADI instance, as follows:

```

      0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+---+
| UN|           RESV           |
+---+---+---+---+---+---+---+---+

```

The UN flag indicates that the RBridge originating the ESADI-LSP, including this ESADI parameter data, will accept and properly process ESADI PDUs sent by TRILL unicast (see Section 4.1). The remaining RESV bits are reserved for future use and MUST be sent as zero and ignored on receipt.

Reserved for future expansion: Future versions of the ESADI Parameter APPsub-TLV may have additional information. A receiving ESADI RBridge ignores any additional data here, unless it implements such future expansion(s).

6.2. MAC-Reachability TLV

The primary information in TRILL ESADI-LSP PDUs consists of MAC-Reachability (MAC-RI) TLVs specified in [RFC6165]. These TLVs contain one or more unicast MAC addresses of end stations that are both on a port and in a VLAN for which the originating RBridge is Appointed Forwarder, along with the 1-octet unsigned Confidence in this information with a value in the range 0-254. If such a TLV is received containing a Confidence of 255, it is treated as if the Confidence was 254. (This is to assure that any received address information can be overridden by local address information statically configured with a Confidence of 255.)

The TLVs in TRILL ESADI PDUs, including the MAC-RI TLV, MUST NOT contain the Data Label ID. If a Data Label ID is present in the MAC-RI TLV, it is ignored. In the ESADI PDU, only the Inner.VLAN or Inner.FGL tag indicates the Data Label to which the ESADI-LSP applies.

6.3. Default Authentication

The Authentication TLV may be included in ESADI PDUs [RFC5310] [IS-IS]. The default for ESADI PDU authentication is based on the state of TRILL IS-IS shared secret authentication for TRILL IS-IS LSP PDUs. If TRILL IS-IS authentication and ESADI are implemented at a TRILL switch, then ESADI MUST be able to use the authentication algorithms implemented for TRILL IS-IS and implement the keying material derivation function given below. If ESADI authentication has been manually configured, that configuration is not restricted by the configuration of TRILL IS-IS security.

If TRILL IS-IS authentication is not in effect for LSP PDUs originated by a TRILL switch, then ESADI PDUs originated by that TRILL switch are by default also unsecured.

If such IS-IS LSP PDU authentication is in effect at a TRILL switch, then, unless configured otherwise, ESADI PDUs sent by that switch MUST use the same algorithm in their Authentication TLVs. The ESADI authentication keying material used is derived from the IS-IS LSP shared secret keying material as detailed below. However, such authentication MAY be configured to use some other keying material.

HMAC-SHA256 ("TRILL ESADI", IS-IS-LSP-shared-key)

In the algorithm above, HMAC-SHA256 is as described in [FIPS180] and [RFC6234], and "TRILL ESADI" is the 11-byte US ASCII [ASCII] string indicated. IS-IS-LSP-shared-key is secret keying material being used by the originating TRILL switch for IS-IS LSP authentication.

7. IANA Considerations

IANA allocation and registry considerations are given below. Three new sub-registries have been created in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry located at <http://www.iana.org/assignments/trill-parameters> -- two in Section 7.1 and one in Section 7.2 -- and various code points have been assigned.

7.1. ESADI Participation and Capability Flags

IANA Action 1:

IANA has created the following new sub-registry called "Interested VLANs Flag Bits" in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry.

Sub-registry: Interested VLANs Flag Bits

Registration Procedures: IETF Review

Note: These bits appear in the Interested VLANs record within the Interested VLANs and Spanning Tree Roots Sub-TLV (INT-VLAN) specified in [RFC7176].

References: [RFC7176], [RFC7357]

Bit	Mnemonic	Description	Reference
0	M4	IPv4 Multicast Router Attached	[RFC7176]
1	M6	IPv6 Multicast Router Attached	[RFC7176]
2	-	Unassigned	
3	ES	ESADI Participation	[RFC7357]
4-15	-	(used for a VLAN ID)	[RFC7176]
16-19	-	Unassigned	
20-31	-	(used for a VLAN ID)	[RFC7176]

The creation of this sub-registry (as immediately above) assigned bit 3 as the ESADI Participation bit in the Interested VLANs and Spanning Tree Roots sub-TLV. If The ESADI Participation bit is a one, it indicates that the originating RBridge is participating in ESADI for the indicated Data Label(s).

IANA Action 2:

IANA has created the following new sub-registry called "Interested Labels Flag Bits" in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry.

Sub-registry: Interested Labels Flag Bits

Registration Procedures: IETF Review

Note: These bits appear in the Interested Labels record within the Interested Labels and Spanning Tree Roots Sub-TLV (INT-LABEL) specified in [RFC7176].

References: [RFC7176], [RFC7357]

Bit	Mnemonic	Description	Reference
----	-----	-----	-----
0	M4	IPv4 Multicast Router Attached	[RFC7176]
1	M6	IPv6 Multicast Router Attached	[RFC7176]
2	BM	Bit Map	[RFC7176]
3	ES	ESADI Participation	[RFC7357]
4-7	-	Unassigned	

The creation of this sub-registry (as immediately above) assigned bit 3 as the ESADI Participation bit in the Interested Labels and Spanning Tree Roots sub-TLV. If The ESADI Participation bit is a one, it indicates that the originating RBridge is participating in ESADI for the indicated Data Label(s).

7.2. TRILL GENINFO TLV

IANA Action 3:

IANA has allocated the IS-IS Application Identifier 1 under the Generic Information TLV (#251) [RFC6823] for TRILL.

IANA Action 4:

IANA has created a sub-registry in the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry as follows:

Sub-registry: TRILL APPsub-TLV Types under IS-IS TLV 251
Application Identifier 1

Registration Procedures: IETF Review with additional requirements on the documentation of the use being registered as specified in Section 7.2 of [RFC7357].

Note: Types greater than 255 are only usable in contexts permitting a type larger than one byte, such as extended TLVs [RFC7356].

Reference: [RFC7357]

Type	Name	Reference
0	Reserved	[RFC7357]
1	ESADI-PARAM	[RFC7357]
2-254	Unassigned	[RFC7357]
255	Reserved	[RFC7357]
256-65534	Unassigned	[RFC7357]
65535	Reserved	[RFC7357]

TRILL APPsub-TLV Types 2 through 254 and 256 through 65534 are available for assignment by IETF Review. The RFC causing such an assignment will also include a discussion of security issues and of the rate of change of the information being advertised. TRILL APPsub-TLVs MUST NOT alter basic IS-IS protocol operation including the establishment of adjacencies, the update process, and the decision process for TRILL IS-IS [IS-IS] [RFC1195] [RFC7177]. The TRILL Generic Information TLV MUST NOT be used in an IS-IS instance zero [RFC6822] LSP but may be used in Flooding Scoped LSPs (FS-LSPs) [RFC7356].

The V, I, D, and S flags in the initial flags byte of a TRILL Generic Information TLV have the meanings specified in [RFC6823] but are not currently used, as TRILL operates as a Level 1 IS-IS area and no semantics are hereby assigned to the inclusion of an IPv4 and/or IPv6 address via the I and V flags. Thus, these I and V flags MUST be zero; however, if either or both are one, the space that should be taken by an IPv4 and/or IPv6 address, respectively, is skipped over and ignored. Furthermore, the use of multilevel IS-IS is an obvious extension for TRILL [MultiLevel], and future IETF Standards Actions may update or obsolete this specification to provide for the use of any or all of these flags in the TRILL GENINFO TLV.

The ESADI Parameters information, for which TRILL APPsub-TLV 1 is hereby assigned, is compact and slow changing (see Section 6.1).

For security considerations related to ESADI and the ESADI Parameter APPsub-TLV, see Section 8.

8. Security Considerations

ESADI PDUs can be authenticated through the inclusion of the Authentication TLV [RFC5310]. Defaults for such authentication are described in Section 6.3.

The ESADI-LSP data primarily announces MAC address reachability within a Data Label. Such reachability can, in some cases, be an authenticated registration (for example, a Layer 2 authenticated registration using cryptographically based EAP (Extensible

Authentication Protocol [RFC3748]) methods via [802.1X]). The combination of these techniques can cause ESADI MAC reachability information to be substantially more trustworthy than MAC reachability learned from observation of the data plane. Nevertheless, ESADI still involves trusting all other RBridges in the TRILL campus, at least those that have the keying material necessary to construct a valid Authentication TLV.

However, there may be cases where authenticated registration is used for end stations, because of a significant threat of forged packets on end station links, but it is not necessary to authenticate ESADI PDUs because that threat is not present for inter-RBridge trunks. For example, a TRILL campus with secure RBridges and inter-RBridge links configured as trunks but with some end stations connected via IEEE 802.11 wireless access links might use 802.11 authentication for the connection of such end stations but might not necessarily authenticate ESADI PDUs. Note that if the IS-IS LSPs in a TRILL campus are authenticated, perhaps due to a concern about forged packets, the ESADI PDUs will be authenticated by default as provided in Section 6.3.

MAC reachability learned from the data plane (the TRILL default) is overwritten by any future learning of the same type. ESADI advertisements are represented in the Data Label scoped link state database. Thus, ESADI makes visible any multiple attachments of the same MAC address within a Data Label to different RBridges (see Section 5.3). This may or may not be an error or misconfiguration, but ESADI at least makes it explicitly and persistently visible, which would not be the case with data plane learning.

For general TRILL security considerations, see [RFC6325].

8.1. Privacy Considerations

The address reachability information distributed by ESADI has substantial privacy considerations under many, but not all, circumstances.

For example, if ESADI were used in a TRILL campus with independent user end stations at the edge, the MAC addresses of such end stations could uniquely identify the users of those end stations. Their reachability would be sensitive information and, particularly if logged, could reveal such user information. On the other hand, if TRILL is being used to implement an Internet Exchange Point (IXP) to connect Internet Service Providers (ISPs), the MAC addresses being advertised in ESADI would typically be those of the ISP's directly connected IP router ports, since Layer 3 routers bound the TRILL campus, for which there would be few privacy concerns.

However, records of MAC attachment that include a modest amount of history, perhaps a few days' worth, can be useful in managing a network and troubleshooting network problems. It might, in some cases, also be legally required, or required for billing purposes or the like.

Network operators should seek a reasonable balance between these competing considerations, customized for the circumstances of their particular networks where ESADI is in use. They should not maintain logs of MAC reachability information for any longer than is clearly required.

9. Acknowledgements

The authors thank the following, listed in alphabetic order, for their suggestions and contributions:

David Black, Somnath Chatterjee, Adrian Farrel, Stephen Farrell, Sujay Gupta, Russ Housley, Pearl Liang, Kathleen Moriarty, Thomas Narten, Erik Nordmark, and Mingui Zhang.

10. References

10.1. Normative References

- [ASCII] American National Standards Institute (formerly United States of America Standards Institute), "USA Code for Information Interchange", ANSI X3.4-1968, 1968.

ANSI X3.4-1968 has been replaced by newer versions with slight modifications, but the 1968 version remains definitive for the Internet.
- [FIPS180] "Secure Hash Standard (SHS)", National Institute of Standards and Technology, Federal Information Processing Standard (FIPS) 180-4, March 2012, <<http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>>.
- [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.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, February 2009.
- [RFC6165] Banerjee, A. and D. Ward, "Extensions to IS-IS for Layer-2 Systems", RFC 6165, 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.
- [RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011.
- [RFC6823] Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", RFC 6823, December 2012.
- [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.
- [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.
- [RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, 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.
- [RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, September 2014.

10.2. Informative References

- [802.1X] IEEE 802.1, "IEEE Standard for Local and metropolitan area networks--Port-Based Network Access Control", IEEE Standard 802.1X-2010, February 2010.
- [FNV] Fowler, G., Noll, L., Vo, K., and D. Eastlake 3rd, "The FNV Non-Cryptographic Hash Algorithm", Work in Progress, April 2014.
- [MultiLevel] Perlman, R., Eastlake 3rd, D., Ghanwani, A., and H. Zhai, "Flexible Multilevel TRILL (Transparent Interconnection of Lots of Links)", Work in Progress, June 2014.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, May 2011.
- [RFC6822] Previdi, S., Ed., Ginsberg, L., Shand, M., Roy, A., and D. Ward, "IS-IS Multi-Instance", RFC 6822, December 2012.
- [VLANmapping] Perlman, R., Rijhsinghani, A., Eastlake 3rd, D., Banerjee, A., and D. Dutt, "TRILL: Campus Label and Priority Regions", Work in Progress, January 2014.

Appendix A. Interoperability and Changes to RFC 6325

This appendix summarizes the significant changes this document makes to the TRILL base protocol specification [RFC6325]. Although simultaneous use of [RFC6325] ESADI and ESADI as specified in this document in a TRILL campus is very unlikely due to non-deployment of [RFC6325] ESADI, this appendix also discusses, for each change, the interoperability considerations should such simultaneous use occur.

A.1. ESADI PDU Changes

The format of ESADI-LSP, ESADI-CSNP, and ESADI-PSNP PDU payloads is changed from the IS-IS Level 1 format [IS-IS] to the Extended Level 1 Circuit Scope format (E-L1CS) specified in [RFC7356]. This change is not backwards compatible with [RFC6325]. It is made in light of the information-carrying capacity of the E-L1CS format, which is 256 times greater than that of the base IS-IS format. It is anticipated that this greater information-carrying capacity will be needed, under some circumstances, to carry end station addressing information or other similar address and reachability information when it is added to ESADI in the future.

The PDU numbers used for the ESADI LSP, CSNP, and PSNP PDUs in [RFC6325] are 18, 24, and 26 [IS-IS]. With this document, the format changes, and the PDU numbers change to 10, 11, and 12 [RFC7356]. The use of different PDU numbers assures that a PDU will not be mis-parsed. Because of this, implementations of this document and implementations of [RFC6325] ESADI will discard each other's PDUs. Thus, address reachability or other information distributed through either type of ESADI implementation will only be communicated to other implementations of the same type, and the two communities will not communicate any information with each other.

Note that Rbridges can use the TRILL mandatory-to-implement, enabled-by-default data plane address learning in addition to ESADI. (Section 5 of this document and the material it references explain how to handle conflicts between different sources of address reachability information.) Simply leaving data plane address learning enabled would enable smooth incremental migration from [RFC6325] ESADI to the ESADI specification in this document, should that be necessary. The data plane address learning would fill in any gaps due to non-communication between the two types of ESADI implementations, although without the speed or security advantages of ESADI.

A.2. Unicasting Changes

Unicasting of ESADI PDUs is optionally supported, including replacing Section 4.6.2.2 of [RFC6325] with the new text given in Section 4.1 of this document. This unicast support is backwards compatible because it is only used when the recipient RBridge signals its support.

A.3. Message Timing Changes and Suggestions

The following timing-related ESADI message changes and suggestions are included in this document:

1. Provide for staggered delay for non-originators of ESADI-LSP fragments in response to requests for such fragments by CSNP and PSNP messages.
2. Suggest staggered timing of unicast ESADI-LSPs when a new ESADI RBridge appears on the ESADI virtual link.

These relate only to the timing of messages for congestion minimization. Should a message be lost, due to congestion or otherwise, it will be later retransmitted as a normal part of the robust flooding mechanism used by ESADI.

A.4. Duplicate Address Reachability

The handling of persistent reachability of the same MAC within the same Data Label from two or more RBridges is substantially modified, including the explicit replacement of some text in Section 4.2.6 of [RFC6325] (see Section 5.3 of this document). There is no problem with a mixture of ESADI implementations in a TRILL campus, some conforming to [RFC6325] and some conforming to this document, for handling this condition. The more implementations conform to the improved behavior specified in this document for this condition, the better the traffic-spreading will be, and the less likely address flip-flopping problems will be.

Authors' Addresses

Hongjun Zhai
ZTE Corporation
68 Zijinghua Road
Nanjing 200012
China
Phone: +86-25-52877345
EMail: zhai.hongjun@zte.com.cn

Fangwei Hu
ZTE Corporation
889 Bibo Road
Shanghai 201203
China
Phone: +86-21-68896273
EMail: hu.fangwei@zte.com.cn

Radia Perlman
EMC
2010 256th Ave. NE, #200
Bellevue, WA 98007
USA
EMail: Radia@alum.mit.edu

Donald Eastlake 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757
USA
Phone: +1-508-333-2270
EMail: d3e3e3@gmail.com

Olen Stokes
Extreme Networks
2121 RDU Center Drive, Suite 300
Morrisville, NC 27560
USA
EMail: ostokes@extremenetworks.com