

TRILL Working Group
INTERNET-DRAFT
Intended status: Proposed Standard
Updates: [6325](#)

Hongjun Zhai
Fangwei Hu
ZTE
Radia Perlman
Intel Labs
Donald Eastlake
Huawei
Olen Stokes
Extreme Networks
November 26, 2013

Expires: May 25, 2014

**TRILL (Transparent Interconnection of Lots of Links):
ESADI (End Station Address Distribution Information) Protocol**
[<draft-ietf-trill-esadi-04.txt>](#)

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 multi-pathing 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 addresses and other 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 Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Distribution of this document is unlimited. Comments should be sent to the TRILL working group mailing list: trill@ietf.org.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>. The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Table of Contents

1. Introduction	4
1.1 Content and Precedence	5
1.2 Terminology	5
2. ESADI Protocol Overview	7
2.1 ESADI Virtual Link	9
2.2 ESADI Neighbor Determination	10
2.3 ESADI Payloads	10
3. ESADI DRB Determination	12
4. ESADI PDU processing	13
4.1 Unicasting ESADI PDUs	13
4.2 General Transmission of ESADI PDUs	14
4.3 General Receipt of ESADI PDUs	14
4.4 Details of Receiving and Sending ESADI PDUs	15
4.4.1 ESADI-CSNP Receipt	15
4.4.2 ESADI-PSNP Receipt	16
4.4.3 ESADI-LSP Receipt	16
4.4.4 Passage of Time	16
4.4.5 Neighbor Appearance	16
5. End Station Addresses	18
5.1 Learning Confidence Level	18
5.2 Forgetting End Station Addresses	18
5.3 Duplicate MAC Address	18
6. ESADI-LSP Contents	21
6.1 ESADI Parameter Data	21
6.2 MAC Reachability TLV	22
6.3 Default Authentication	23
7. IANA Considerations	24
7.1 ESADI Participation and Capability Flags	24
7.2 TRILL GENINFO TLV	25
8. Security Considerations	27
9. Acknowledgements	27
Normative references	28
Informative References	29
Appendix A: Changes to [RFC6325]	30
Appendix Z: Change History	31
Authors' Addresses	33

1. Introduction

The IETF 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 multi-pathing of both unicast and multicast traffic. TRILL accomplishes this with the IS-IS (Intermediate System to Intermediate System) [[IS-IS](#)] [[RFC1195](#)] [[rfc6326bis](#)] 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, [[RFCfg1](#)]). 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 \[RFC6325\]](#). The ESADI (End Station Address Distribution Information) protocol is an optional Data Label scoped way 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 additional types of 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: (1a) 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](#)]). (1b) 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, frames ingressed for that end station's MAC address 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 optionally 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 specification, essentially replacing the description of the ESADI protocol, providing more detail, and prevailing over [[RFC6325](#)] where they conflict. The changes are summarized in [Appendix A](#). These changes are not backwards compatible because, among other things, they change the format of ESADI-LSPs.

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

1.2 Terminology

This document uses the acronyms defined in [[RFC6325](#)] and the following:

Data Label - VLAN or FGL.

FGL - Fine Grained Label [[RFCfgl](#)].

LSP - Link State PDU.

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

PDU - Protocol Data Unit.

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 Considerations terms such as "IETF Review" as to be interpreted

as described in [[RFC5226](#)].

2. ESADI Protocol Overview

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

ESADI is a separate optional protocol from the mandatory TRILL IS-IS implemented by all R Bridges in a campus. There is a separate ESADI instance for each Data Label (VLAN or FGL). In essence, for each Data Label, there is a modified instance of the IS-IS reliable flooding mechanism in which ESADI R Bridges may choose to participate. (These are not the instances specified in [RFC6822].) It is an implementation decision how independent the multiple ESADI instances at an R Bridge are. For example, the ESADI link state database could be in a single database with a field in each record indicating the Data Label to which it applies or could be a separate database per Data Label. But the update process operates separately for each ESADI instance and independently from the TRILL IS-IS update process.

ESADI does no routing so there is no reason for pseudo-nodes in ESADI and none are created. Furthermore, there may be a requirement for a relatively large amount of data to be distributed through ESADI which might take a large number of ESADI-LSP fragments. ESADI-LSP, ESADI-CSNP, and ESADI-PSNP payloads are therefore formatted as Extended Level 1 Circuit Scope PDUs [FS-LSP] (see also [Section 6](#)). This allows up to 2^{16} fragments but does not support link state data associated with pseudo-nodes.

After the TRILL header, ESADI packets have an inner Ethernet header with the Inner.MacDA of "All-Egress-R Bridges" (formerly called "All-ESADI-R Bridges"), 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 below. The outer VLAN tag will not be present if it was stripped by an Ethernet port out of which the packet was sent.

Outer Ethernet Header:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Next Hop Destination Address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Next Hop Destination Addr.      | Sending RBridge Port MAC Addr. |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Sending RBridge Port MAC Address                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
...Ethernet frame tagging including optional Outer.VLAN tag...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ethertype = TRILL      0x22F3 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TRILL Header:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| V | R | M | Op-Length | Hop Count |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Egress Nickname           | Ingress (Origin) Nickname       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Inner Ethernet Header:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               All-Egress-RBridges                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| All-Egress-RBridges cont.    | Origin RBridge MAC Address    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Origin RBridge MAC Address continued                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| VLAN or FGL Data Label (4 or 8 bytes) ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ethertype = L2-IS-IS    0x22F4 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

ESADI Payload (formatted as IS-IS):

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| IS-IS Common Header, IS-IS PDU Specific Fields, IS-IS TLVs |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Frame Check Sequence:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               FCS (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 field will be

one unless the ESADI PDU was unicast, in which case the M field 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 RBridge MUST use the same Inner.MacSA for all of the ESADI packets that RBridge originates.

TRILL ESADI packets sent on a PPP link are structured as shown below [[RFC6361](#)].

PPP Header:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| PPP = TNP (TRILL data) 0x005D |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

TRILL Header:

```
+-----+-----+-----+-----+-----+-----+
| V | R |M|Op-Length| Hop Count |
+-----+-----+-----+-----+-----+-----+
| Egress Nickname           | Ingress (Origin) Nickname   |
+-----+-----+-----+-----+-----+-----+
```

Inner Ethernet Header:

```
+-----+-----+-----+-----+-----+-----+
|                               All-Egress-RBridges                               |
+-----+-----+-----+-----+-----+-----+
| All-Egress-RBridges cont.    | Origin RBridge MAC Address    |
+-----+-----+-----+-----+-----+-----+
|                               Origin RBridge MAC Address continued              |
+-----+-----+-----+-----+-----+-----+
| VLAN or FGL Data Label (4 or 8 bytes) ...                                     |
+-----+-----+-----+-----+-----+-----+
| Ethertype = L2-IS-IS   0x22F4 |
+-----+-----+-----+-----+-----+-----+
```

ESADI Payload (formatted as IS-IS):

```
+-----+-----+-----+-----+-----+-----+
| IS-IS Common Header, IS-IS PDU Specific Fields, IS-IS TLVs |
+-----+-----+-----+-----+-----+-----+
```

PPP Check Sequence:

```
+-----+-----+-----+-----+-----+-----+
|                               PPP Check Sequence                               |
+-----+-----+-----+-----+-----+-----+
```

Figure 3: ESADI PPP Link Packet Format

2.1 ESADI Virtual Link

All transit 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 [\[ClearCorrect\]](#)) and are running ESADI for that Data Label. No "routing" computation or routing decisions ever have 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 [\[RFCfg1\]](#). 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 for the Data Label of an ESADI packet do not decapsulate or locally process the TRILL 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 [\[ClearCorrect\]](#)) 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, and when RB2 is no longer announcing participation in Data Label X ESADI, it ceases to be a neighbor for the Data Label X ESADI instance. All these considerations being Data Label scoped. Because of these mechanisms, there are no "Hellos" sent in ESADI.

Participation announcement in a VLAN scoped ESADI instance is through setting a flag bit in the Interested VLANs sub-TLV and announcement for an FGL scoped ESADI instance is through setting a flag bit in the Interested Labels sub-TLV [\[rfc6326bis\]](#). (See [Section 7.1](#))

2.3 ESADI Payloads

TRILL ESADI packet payloads are structured like IS-IS Extended Level 1 Circuit Scoped LSP, CSNP, and PSNP PDUs [\[FS-LSP\]](#), 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

one 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 types of 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 is ignored.

3. ESADI DRB Determination

Because ESADI does no adjacency announcement or routing, the ESADI-DRB never creates a pseudonode. But a DRB (Designated RBridge [[rfc6327bis](#)]) is still needed for ESADI-LSP synchronization by issuing ESADI-CSNP PDUs and responding to ESADI-PSNP PDUs.

Generally speaking, the DRB election on the ESADI virtual link (see [Section 2.1](#)) operates similarly to a TRILL IS-IS broadcast link [[rfc6327bis](#)] 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 (1) are data reachable from RB1 [[ClearCorrect](#)], and (2) are 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 System ID, comparing fields as unsigned integers with the larger magnitude considered higher priority. "SNPA/MAC address" is not considered in this tie breaking 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 R Bridges (TRILL switches) on the virtual link. There are only ESADI-LSP, ESADI-CSNP and ESADI-PSNP PDUs used in ESADI. In particular, there are no Hello or MTU PDUs because ESADI does not build a topology, does not do any routing, and uses the campus Sz MTU.

4.1 Unicasting ESADI PDUs

In 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. But 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-LSP and ESADI-PSNP PDUs MAY be TRILL unicast if it is confirmed that the destination R Bridge supports receiving unicast ESADI PDUs (see [Section 6.1](#)).

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 the forwarding R Bridge is (1) interested in the specified VLAN or Fine Grained Label, (2) implements the TRILL ESADI protocol, and (3) ESADI is enabled for that VLAN or Fine Grained Label, 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 R Bridge, the packet is forwarded as for known unicast TRILL Data in [Section 4.6.2.4](#). If M=0 and the egress nickname is that of the receiving R Bridge and the receiving R Bridge 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

An ESADI instance SHOULD NOT transmit any ESADI PDUs if it has no ESADI neighbors. They 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(TRILL Ethertype) + 6(TRILL Header) + 6(Inner.MacDA) + 6(Inner.MacSA) + 4/8(Inner.VLAN/Inner.FGL) 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 [[ClearCorrect](#)] for discussions of Sz and MTU.)

The format of a unicast ESADI packet is the format of multicast TRILL ESADI packet, 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.

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.

[FS-LSP] provides a priority bit to indicate that an LSP fragment should be flooded with high priority. This bit SHOULD be set on ESADI-LSP fragment zero and SHOULD NOT be set on any other ESADI-LSP fragment.

4.3 General Receipt of ESADI PDUs

Because ESADI neighbor adjacency is in terms of System ID, all PDU acceptance tests that in TRILL/IS-IS check that the PDU is from an adjacent router instead check that the System ID is that of an ESADI neighbor and do not check either the source Inner or Outer SNPA/MAC.

If an ESADI instance believes that it has no ESADI neighbors, it

ignores any ESADI PDUs it receives.

4.4 Details of Receiving and Sending ESADI PDUs

Event	Section
Receive ESADI-CSNP	See Section 4.4.1
Receive ESADI-PSNP	See Section 4.4.2
Receive ESADI-LSP	See Section 4.4.3
Passage of Time	See Section 4.4.4
Neighbor Appearance	See Section 4.4.5

4.4.1 ESADI-CSNP Receipt

If an ESADI RBridge RB1 believes it is DRB on a virtual link for Data Label X, it ignores an ESADI-CSNP it receives.

If RB1 believes it is not DRB:

When RB1 receives an ESADI-CSNP from RB2 and detects that it has ESADI-LSPs that RB2 is missing, it sets the transmission flag only for its own ESADI-LSPs that RB2 is missing. Missing ESADI-LSPs originated by other ESADI RBridges will be detected by those other ESADI RBridges because all data reachable ESADI RBridges participating for Data Label X are adjacent.

When RB1 receives an ESADI-CSNP from RB2 and detects that it is missing ESADI-LSPs originated by RBridges reachable from RB1 that RB2 has, it generates one or more ESADI-PSNP PDUs. Generally, ESADI-PSNPs are multicast and may be fragmented as in IS-IS; however, if RB1 is missing ESADI-LSPs from RBx and RBx is advertising unicast ESADI PDU support, RB1 MAY construct one or more ESADI-PSNP fragments listing only RBx ESADP-LSPs and unicast

those ESADI-PSNPs to RBx.

4.4.2 ESADI-PSNP Receipt

When RBx receives an ESADI-PSNP PDU, if RBx is the originator of any ESADI-LSPs requested by the ESADI-PSNP those ESADI-LSPs will be multicast on the virtual link.

4.4.3 ESADI-LSP Receipt

Processing of a received ESADI-LSP is as in [[FS-LSP](#)].

In the case of receiving an ESADI-LSP with a smaller sequence number than the copy stored in the local EASDI Link State Database the local ESADI instance will also schedule multicasting its stored copy.

4.4.4 Passage of Time

If an ESADI instance believes it is DRB, it multicasts an ESADI-CSNP periodically (thrice per CSNP Time, see [Section 6.1](#)) to keep the link state database synchronized among its neighbors on the virtual link.

The multi-hop TRILL multi-destination packet distribution with Reverse Path Forwarding Check will typically be less reliable than the single hop link-local LSP synchronization of TRILL IS-IS. Therefore, for LSP synchronization robustness, in addition to sending ESADI-CSNPs when it is 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 DRB for the ESADI instance, and
- o it has not received or sent an ESADI-CSNP PDU for the instance for the CSNP Time (see [Section 6.1](#)) of the DRB.

4.4.5 Neighbor Appearance

When an ESADI instance sees that it has a new neighbor, its self-originated EASDI-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. So the ESADI instances SHOULD wait an interval of time before sending their ESADI-LSP to a new neighbor. The interval time value is up to the device implementation and is subject to the usual IS-IS timing jitter. 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 $(2 * \text{nickname} / 0xFFC0)$ seconds assuming "nickname" and 0xFFC0 to be unsigned quantities.

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 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 time out 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_N 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 with 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 replacing the last sentence

of the last paragraph of [Section 4.2 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 have two or more ports with the same MAC&label with each port connected to different RBridges (RB1, RB2, ...) by separate links. (Label is a VLAN or FGL.) With ESADI, some other RBridge, RB0, can persistently see that MAC&label connected to multiple RBridges. When RB0 ingresses a frame destined for that MAC&label, the current [\[RFC6325\]](#) text permits a wide range of behavior. In particular, it would permit RB0 to use some rule such as always send to the egress with the lowest System ID, which would put all of this traffic through one of the egress RBridges and one of the end station ports. There would be no load spreading even if there were multiple different ingress RBridges and/or different MAC addresses with the same reachability. It also would also permit RB0 to send different traffic to different egresses by doing ECMP at a flow level, which would likely result in return traffic to RB0 from RB1, RB2, ... for the same MAC&label. The resulting address flip-flopping could cause problems. This update to [\[RFC6325\]](#) avoids these potential difficulties by requiring RB0 to use one of the following two policies: (1) use only one egress for a particular MAC&label but to select that egress pseudo-randomly based on the topology including MAC reachability or (2) if it will not be disturbed by the returning TRILL Data packets showing the same MAC&label flip-flopping between different ingresses, it may use ECMP. Assuming multiple ingress RBridges and/or multiple MAC addresses, strategy 1 should result in load spreading without address flip-flopping while strategy 2 will produce more uniform load spreading with address flip-flopping from the point of view of RB0.

OLD [\[RFC6325\]](#) 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 [\[RFC6325\]](#) text:

"...

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

1. In a pseudo-random 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) to that egress. This pseudo-random choice need only be changed when there is a

change in campus topology or MAC attachment information. Such pseudo-random selection will, over a population of ingress RBridges, probabilistically spread traffic over the possible egress RBridges. Reasonable inputs to the pseudo-random 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. There is no need for different RBridges to use the same pseudo-random function.

2. If RB2 supports ECMP and will not be disturbed by return traffic from the same MAC and VLAN (or FGL) coming from different ingress RBridges, then it MAY send traffic using ECMP at the flow level to the egress RBridges that are least cost from RB2 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 consists of zero or more MAC Reachability TLVs, optionally an Authentication TLV, and exactly one ESADI parameter APPsub-TLV. Other 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 PDUs [[FS-LSP](#)], the Type and Length fields in the TLVs are 16-bit.

This section specifies the format for the ESADI parameter data 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. But if an ESADI-LSP number zero or such an ESADI-CSNP or ESADI-PSNP is received that is longer, it is still processed normally.

6.1 ESADI Parameter Data

The figure below 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 data APPsub-TLV in non-zero ESADI-LSP fragments are ignored.

```

+---+---+---+---+---+---+---+---+---+---+
| Type                                     | (2 byte)
+---+---+---+---+---+---+---+---+---+---+
| Length                                   | (2 byte)
+---+---+---+---+---+---+---+---+---+---+
|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 subTLV (TRILL APPsub-TLV type 0x0001). Two bytes because this APPsub-TLV appears in an Extended TLV [[FS-LSP](#)].

Length: Set to 3 to 65,535 but must fit within the ESADI packet.

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

Priority: The Priority field gives the originating RBridge's priority for being 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 seven-bit integer with larger magnitude indication 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 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. 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 Parameters 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 one octet unsigned Confidence in this information with a value in the range 0-254. If such a TLV is received with 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. 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, by default, ESADI PDUs originated by that TRILL switch are 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 above HMAC-SHA256 is as described in [[FIPS180](#)] [[RFC6234](#)] and "TRILL ESADI" is the eleven 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.

7.1 ESADI Participation and Capability Flags

IANA is requested to allocate bit TBD [3 recommended] as the "ESADI Participation" bit in the Interested VLANs sub-TLV and the Interested Labels sub-TLVs [[rfc6326bis](#)]. If The ESADI Participation bit is a one, it indicates that the originating RBridge is participating in ESADI for the indicated VLAN(s) or FGL(s). In addition, IANA is requested to create two sub-registries in the TRILL Parameters Registry for such bits as follows:

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).

References: [[rfc6326bis](#)], [This document]

Bit	Mnemonic	Description	Reference
---	-----	-----	-----
0	M4	IPv4 Multicast Router Attached	[rfc6326bis]
1	M6	IPv6 Multicast Router Attached	[rfc6326bis]
2	-	available for allocation	
3	ES	ESADI Participation	This document
4-15	-	(used for a VLAN ID)	[rfc6326bis]
16-19	-	available for allocation	
20-31	-	(used for a VLAN ID)	[rfc6326bis]

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).

References: [[rfc6326bis](#)], [this document]

Bit	Mnemonic	Description	Reference
---	-----	-----	-----
0	M4	IPv4 Multicast Router Attached	[rfc6326bis]
1	M6	IPv6 Multicast Router Attached	[rfc6326bis]
2	BM	Bit Map	[rfc6326bis]
3	ES	ESADI Participation	This document
4-7	-	available for allocation	

7.2 TRILL GENINFO TLV

IANA is requested to allocate the IS-IS Application Identifier TBD [1 suggested] under the Generic Information TLV (#251) [[RFC6823](#)] for TRILL and to create a subregistry in the TRILL Parameters Registry as follows:

Sub-Registry: TRILL APPsub-TLV Types under IS-IS TLV #251
Application Identifier #TBD

Registration Procedures: IETF Review

Note: Types greater than 255 are only usable in contexts permitting a type larger than one byte, such as Extended TLVs [[FS-LSP](#)].

Reference: <this RFC>

Type	Name	Reference
-----	-----	-----
0	Reserved	<this RFC>
1	ESADI-PARAM	<this RFC>
2-254	Available	<this RFC>
255	Reserved	<this RFC>
256-65534	Available	<this RFC>
65535	Reserved	<this RFC>

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](#)] [[rfc6327bis](#)]. The TRILL Generic Information TLV MUST NOT be used in an IS-IS instance zero [[RFC6822](#)].

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 flags MUST be zero; however, use of multi-level 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 parameters APPsub-TLV, see [Section 8](#).

8. Security Considerations

ESADI PDUs can be authenticated through the inclusion of the Authentication TLV as described in [Section 6.3](#). The ESADI-LSP data primarily announces MAC&label reachability. 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.

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 data label scoped link state database. Thus ESADI makes visible any multiple attachments of the same MAC&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](#)].

9. Acknowledgements

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

Somnath Chatterjee, Thomas Narten, and Erik Nordmark

This document was produced with raw nroff. All macros used were defined in the source file.

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)", United States of American, National Institute of Science 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] - International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, Nov 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.
- [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.
- [ClearCorrect] - Eastlake, D., Zhang, M., Ghanwani, A., Manral, V.,

A. Benerjee, "TRILL: Clarifications, Corrections, and Updates",

H. Zhai, et al

[Page 28]

[draft-ietf-trill-clear-correct](#), in RFC Editor's queue.

- [FS-LSP] - Ginsberg, L., S. Previdi, Y. Yang, "IS-IS Flooding Scope LSPs", [draft-ietf-isis-fs-lsp](#), work in progress.
- [rfc6326bis] - Eastlake, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [draft-ietf-isis-rfc6326bis](#), work in progress.
- [rfc6327bis] - Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, H., and V. Manral, "Routing Bridges (Rbridges): Adjacency", [draft-ietf-trill-rfc6327bis](#), work in progress.
- [RFCfgl] - Eastlake, D., M. Zhang, P. Agarwal, R. Perlman, D. Dutt, "TRILL (Transparent Interconnection of Lots of Links): Fine-Grained Labeling", [draft-ietf-trill-fine-labeling](#), in RFC Editor's queue.

Informative References

- [802.1X] - IEEE 802.1, "IEEE Standard for Local and metropolitan area networks / Port-Based Network Access Control", IEEE Std 802.1X-2010, 5 February 2010.
- [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.
- [MultiLevel] - Perlman, R., D. Eastlake, A. Ghanwani, H. Zhai, "Multilevel TRILL (Transparent Interconnection of Lots of Links)", [draft-perlman-trill-rbridge-multilevel](#), work in progress.
- [VLANmapping] - Perlman, R., D. Dutt, A. Banerjee, A. Rijhsinghani, and D. Eastlake, "Rbridges: Campus VLAN and Priority Regions", [draft-ietf-trill-rbridge-vlan-mapping](#), work in progress.

Appendix A: Changes to [[RFC6325](#)]

Below is a list of the main changes this document makes to the TRILL base protocol specification [[RFC6325](#)]:

1. The format of ESADI-LSP, ESDAI-CSNP, and ESADI-PSNP PDU payloads is changed from the base IS-IS format to the Extended Level 1 Circuit Scoped format in [[FS-LSP](#)].
2. Unicasting of ESADI PDUs is supported.
3. The handling of persistent reachability of the same MAC within the same Data Label from two or more RBridge is substantially modified (see [Section 5.3](#)).

Appendix Z: Change History

RFC Editor: Please delete this section before publication.

Z.1 From -00 to -01

1. Add [Section 6.3](#) "Default Authentication".
2. Add "Acknowledgements" Section.
3. Change requirement from "MAY" to "SHOULD" for an ESADI RBridge that is not DRB to send an ESADI-CSNP if it does not receive an ESADI-CSNP in long enough.
4. Default CSNP Time was listed as 30 in one place and 40 in another. Change to uniformly specify 30.
5. Update references to [RFC 6326](#) to reference the 6326bis draft.
6. Relax allocation criteria for TRILL APPsub-TLV type code points from Standard Action to IETF Review.
7. Numerous Editorial changes.

Z.2 From -01 to -02

1. Extend to cover FGL and well as VLAN and introduce the term "Data Label" to cover both.
2. Expand number of LSP fragments to 2^{16} .
3. Simplify neighbor detection to no longer require possession of ESADI LSP zero.
4. Add update to last sentence of [Section 4.2 of \[RFC6325\]](#).
5. Update references for publication of RFCs 6822 and 6823.
6. Additional minor changes.

Z.3 From -02 to -03

1. Replace instances of "IS-IS and data unreachable" with just "data unreachable" as data unreachability implies IS-IS unreachability [[ClearCorrect](#)].
2. With ESADI, there is just one virtual link on which all participating TRILL switches are adjacent. Thus, all of the useful ESADI-LSPs in an ESADI link state database are originated by a station on this virtual link. To avoid overworking the ESADI DRB on the link, ESADI-LSPs sent by a reachable TRILL switch in response to an ESADI-PSNP should be sent by the TRILL switch originating those EASDI-LSPs.
3. Re-organize material on sending and receiving ESADI PDUs into more smaller subsections that cover all the different circumstances.
4. Substantially expand Security Considerations section.
5. Numerous editorial changes.

Z.4 From -03 to -04

1. Change to using Extended Level 1 Circuit Scope [[FS-LSP](#)] for EASDI-LSP, ESADI-CSNP, and ESDAI-PSNP PDUs.
2. Update references to [RFC 6327](#) to the rfc6327bis draft.
3. Sort Informational References RFCs in numeric order.'
4. Add [Appendix A](#): summary of changes to [[RFC6325](#)].
5. Minor editing changes.

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
Intel Labs
2200 Mission College Blvd.
Santa Clara, CA 95054-1549 USA

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

Donald Eastlake
Huawei R&D USA
155 Beaver Street
Milford, MA 01757 USA

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

Olen Stokes
Extreme Networks
Pamlico Building One, Suite 100
3306 East NC Hwy 54
RTP, NC 27709 USA

Email: ostokes@extremenetworks.com

Copyright and IPR Provisions

Copyright (c) 2013 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. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of [RFC 5378](#). No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under [RFC 5378](#), shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

