

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

Hongjun Zhai
Fangwei Hu
ZTE
Radia Perlman
Intel Labs
Donald Eastlake
Huawei
June 25, 2012

Expires: December 24, 2012

**TRILL (Transparent Interconnection of Lots of Links):
The ESADI (End Station Address Distribution Information) Protocol**
[<draft-ietf-trill-esadi-00.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 for multi-pathing of both unicast and multicast traffic. Devices that implement the TRILL protocol are called R Bridges (Routing Bridges) or TRILL Switches.

The ESADI (End Station Address Distribution Information) protocol is a VLAN (Virtual Local Area Network) scoped way that TRILL switches can communicate end station addresses to each other. An R Bridge announcing VLAN-x connectivity (normally a VLAN-x appointed forwarder) and running the TRILL ESADI protocol can receive remote address information and/or transmit local address information for VLAN-x to other such R Bridges. This document updates [RFC 6325](#), specifically the documentation of the ESADI protocol.

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: rbridge@postel.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.....	6
3. ESADI DRB Determination.....	10
4. ESADI PDU processing.....	11
4.1 Sending of ESADI PDUs.....	12
4.2 Receipt of ESADI PDUs.....	13
5. End Station Addresses.....	14
5.1 Learning Confidence Level.....	14
5.2 Forgetting End Station Addresses.....	14
6. ESADI-LSP Contents.....	15
6.1 ESADI Parameter Data.....	15
6.2 MAC Reachability TLV.....	16
7. IANA Considerations.....	17
7.1 ESADI Participation and Capability Flags.....	17
7.2 TRILL GENAPP TLV.....	17
8. Security Considerations.....	19
9. References.....	20
9.1 Normative references.....	20
9.2 Informative References.....	21

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 by using the IS-IS (Intermediate System to Intermediate System) [[IS-IS](#)] [[RFC1195](#)] [[RFC6326](#)] link state routing protocol and encapsulating traffic using a header that includes a hop count. The design supports VLANs (Virtual Local Area Networks) and optimization of the distribution of multi-destination frames based on VLANs and IP multicast groups. Devices that implement TRILL are called RBridges (Routing Bridges) or TRILL switches.

There are five ways an RBridge can learn end station addresses as described in [Section 4.8 of \[RFC6325\]](#). The ESADI (End Station Address Distribution Information) protocol is an optional VLAN scoped way RBridges can communicate locally learned end station addresses with each other. An RBridge that is announcing connectivity to VLAN-x (normally a VLAN-x appointed forwarder [[RFC6439](#)]) MAY use the ESADI protocol to announce the end station address of some or all of its attached VLAN-x end nodes to other RBridges that are running ESADI for VLAN-x.

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

1. Security advantages: 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](#)]). In addition, 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 stations MAC (Media Access Control) addresses. If an end station is unplugged from one RBridge and plugged into another, frames addressed to that older RBridge can be black holed. They can be sent just to the older RBridge that the end station was connected to until cached address information at some remote RBridge times out, possibly for tens of seconds or more [[RFC6325](#)].

MAC address reachability information and some ESADI parameters are

carried in ESADI frames rather than in the TRILL IS-IS protocol. As described below, ESADI is, for each VLAN, a virtual logical topology

overlay in the TRILL topology. An advantage of using ESADI is that the end station attachment information is not flooded to all RBridges through TRILL IS-IS but only to participating RBridges advertising ESADI support for the VLAN in which those end stations occur.

1.1 Content and Precedence

This document updates [[RFC6325](#)], the TRILL basic specification, essentially replacing the description of the ESADI protocol, and prevails over [[RFC6325](#)] where they conflict.

[Section 2](#) is the ESADI protocol overview. [Section 3](#) specifies ESADI DRB state. [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 contents.

1.2 Terminology

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

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

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

2. ESADI Protocol Overview

ESADI is a VLAN scoped way that R Bridges can announce and learn end station addresses rapidly and securely. An R Bridge that is announcing itself as connected to one or more VLANs (usually because it is an Appointed Forwarder for those VLANs [[RFC6439](#)]) and participates in the ESADI protocol is called an ESADI R Bridge.

ESADI is a separate protocol from the TRILL IS-IS implemented by all R Bridges in a campus. There is a separate ESADI instance for each VLAN. In essence, for each VLAN, 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 being specified in [[MultiInstance](#)].) It is an implementation decision how independent the multiple ESADI instances at an R Bridge are. For example, the ESADI link state could be in a single database with a field in each record indicating the VLAN to which it applies or could be a separate database per VLAN. But the update process operates separately for each ESADI instance and independently from the TRILL IS-IS update process.

After the TRILL header, ESADI frames have an inner Ethernet header with the Inner.MacDA of "All-Egress-R Bridges" (formerly called "All-ESADI-R Bridges"), an Inner.VLAN tag specifying the VLAN of interest, and the "L2-IS-IS" Ethertype followed by the ESADI payload as shown in Figure 1.

TRILL ESADI frame Structure




Figure 1.

TRILL ESADI frames sent on an Ethernet link are structured as shown below. The outer VLAN tag will not be present if it was stripped by the port out of which the frame was sent.

Outer Ethernet Header:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           All-RBridges Multicast Address           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| All-RBridges Multicast Addr. | Sending RBridge Port MAC Addr. |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Sending RBridge Port MAC Address           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ethertype = C-Tag      0x8100 | Outer.VLAN Tag Information |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ethertype = TRILL      0x22F3 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TRILL Header:

```

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

```

Inner Ethernet Header:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Next Hop Destination Address           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Next Hop Dest. Address cont. | Origin RBridge MAC Address |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Origin RBridge MAC Address continued           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ethertype = C-Tag      0x8100 | Inner.VLAN Tag Information |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| 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 Frame Format

The VLAN specified in the Outer.VLAN information will always be the Designated VLAN for the link on which the frame is sent. The V and R fields will be zero while the M field will be one unless the RBridge supports unicasting ESADI PDUs, in which case the M field MAY be zero. The VLAN specified in the Inner.VLAN information will be the VLAN to which the ESADI frame applies. The Origin RBridge MAC Address or Inner.MacSA MUST be a globally unique MAC address owned by the RBridge originating the ESADI frame, for example, any of its port MAC

addresses, and each RBridge MUST use the same Inner.MacSA for all of the ESADI frames that RBridge originates.

TRILL ESADI frames sent on a PPP link are structured as shown below.

PPP Header:

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

TRILL Header:

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

Inner Ethernet Header:

```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Next Hop Destination Address          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Next Hop Dest. Address cont. | Origin RBridge MAC Address |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Origin RBridge MAC Address continued          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Ethertype = C-Tag      0x8100 | Inner.VLAN Tag Information |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| 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 Frame Format

All transit RBridges forward ESADI frames as if they were ordinary TRILL Data frames. 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 other RBridges in the campus running ESADI for that VLAN. No "routing" computation or routing decisions ever have to be performed by ESADI. An ESADI RBridge merely transmits the ESADI frames it originates on this virtual link as described for TRILL Data frames in [\[RFC6325\]](#). For multicast ESADI frames, which is the normal case, it may use any distribution tree that it might use for a normal multi-destination TRILL Data frame. RBridges that do not implement the ESADI protocol, do not have it enabled, or are not participating for the Inner.VLAN of an ESADI frame do not decapsulate or locally process any multicast TRILL ESADI frames they receive. Thus the ESADI frames are

transparently tunneled through transit RBridges.

TRILL ESADI frame payloads are structured like IS-IS PDUs, except as indicated below, but are always TRILL encapsulated on the wire as if they were TRILL Data frames. The ESADI instance for VLAN-x at an RBridge RB1 determines who its ESADI potential neighbors are by logically examining the TRILL IS-IS link state database for RBridges that are data and IS-IS reachable from RB1 (see Section 2 of [\[ClearCorrect\]](#)) and are announcing their participation in VLAN-x ESADI. When an RBridge RB2 becomes IS-IS or data unreachable from RB1 or any of the relevant entries for RB2 are purged from the core IS-IS link state database, it is lost as a potential neighbor and also dropped from any ESADI instances. And when RB2 is no longer announcing participation in VLAN-x ESADI, it ceases to be a potential neighbor for the VLAN-x ESADI instance. RB2 becomes an actual ESADI adjacency for RB1 when it is a potential neighbor and RB1 holds an ESADI-LSP zero for RB2, all these considerations being VLAN scoped. Because of these mechanisms, there are no "Hellos" sent in ESADI.

The information distributed by the ESADI protocol is 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 6.1](#)). 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. Future uses of ESADI may distribute additional types of information.

TRILL ESADI-LSPs MUST NOT contain a VLAN ID in their payload. The VLAN ID to which the ESADI data applies is the Inner.VLAN of the TRILL Data frame enclosing the ESADI payload. If a VLAN ID could occur within the payload, it might conflict with the Inner.VLAN and could conflict with any future VLAN mapping scheme that may be adopted [\[VLANmapping\]](#). If a VLAN ID field within an ESADI-LSP PDU does include a VLAN ID, its contents is ignored.

3. ESADI DRB Determination

Generally speaking, the DRB state on the ESADI link operates similarly to a TRILL IS-IS broadcast link with the following exception:

In the VLAN-x ESADI-DRB election at RB1 on an ESADI virtual link, comparing with [\[RFC6327\]](#), the candidates are the local ESADI instance for VLAN-x and all remote ESADI instances at RBridges that (1) are data and IS-IS reachable from RB1 [\[ClearCorrect\]](#), (2) are announcing in their TRILL IS-IS LSP that they are participating in ESADI for VLAN-x, and (3) for which RB1 is holding an ESADI-LSP zero with an ESADI Parameters APPsub-TLV. 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. In particular "SNPA/MAC address" is not considered and there is no "Port ID".

Because ESADI does no routing, the ESADI-DRB does not create a pseudonode.

4. ESADI PDU processing

VLAN-x ESADI neighbors are usually not connected directly by a physical link, but are always logically connected by a virtual link. There could be hundreds or thousands of ESADI RBridges 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 and does not do any routing.

In Layer 3 IS-IS, PDU multicasting is normally on a local link and no effort is made to optimize to unicast because under the original conditions when IS-IS was designed (commonly a piece of multi-access Ethernet cable) any frame made the link busy for that frame time. But in ESADI 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 imposes a 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.

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

"4.6.2.2. TRILL ESADI Frames

If M=1, the egress nickname designates the distribution tree. The frame is forwarded as described in [Section 4.6.2.5](#). In addition, if the forwarding RBridge is (1) interested in the specified VLAN, for example it is Appointed Forwarder for that VLAN on at least one port, (2) implements the TRILL ESADI protocol, and (3) ESADI is enabled for that VLAN, 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 frame 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 RBridge and the receiving RBridge supports unicast ESADI PDUs, then the ESADI frame 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 references to those sections in [\[RFC6325\]](#).

[Section 4.1](#) describes the sending of ESADI PDUs. [Section 4.2](#) covers the receipt of ESADI PDUs.

4.1 Sending of ESADI PDUs

The MTU available to instances of ESADI is at least 24 bytes less than that available to TRILL IS-IS because of the additional fields required ($2(\text{TRILL Ethertype}) + 6(\text{TRILL Header}) + 6(\text{Inner.MacDA}) + 6(\text{Inner.MacSA}) + 4(\text{Inner.VLAN})$). Thus the inner ESADI payload, starting with the Intradomain Routing Protocol Discriminator byte, MUST NOT exceed Sz minus 24; however, if a larger payload is received, it is processed normally.

Once an ESADI instance is operationally up for VLAN-x, it multicasts its self-originated ESADI-LSP number zero on the virtual link to announce its ESADI parameters. When the other ESADI instances receive the ESADI-LSP number zero and find a new neighbor, their self-originated 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. If all the other ESADI instances send their self-originated ESADI-LSPs immediately, there may be a surge of traffic to that new neighbor. So the other ESADI instances should wait an interval time before sending the ESADI-LSP 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 ESADI priority when implementation.

If the ESADI instance believes it is DRB, it multicasts an ESADI-CSNP periodically to keep the Link State Database synchronized among its neighbors on the virtual link. After receiving an ESADI-PSNP PDU, the DRB will multicast the ESADI-LSPs requested by the ESADI-PSNP on the virtual link.

For robustness, if an ESADI instance has two or more ESADI neighbors and is not DRB and it receives no ESADI-CSNP PDUs for at least the CSNP Time (see [Section 6.1](#)) of the DRB, it MAY transmit an ESADI-CSNP.

In the case of receiving an ESADI-LSP with a smaller sequence number than the copy stored in the local ESADI Link State Database, the local ESADI instance will also schedule to transmit the stored copy and MAY unicast it to the sender of the received ESADI-LSP if it is confirmed that the sender supports receiving unicast ESADI PDUs.

The format of a unicast ESADI frame is the format of TRILL ESADI frame, in [Section 4.2 in \[RFC6325\]](#), 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.

4.2 Receipt of ESADI PDUs

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

Because all ESADI instance for VLAN-x are adjacent, when RB1 receives an ESADI-CSNP from RB2 and detects that it has ESADI-LSPs that RB1 is missing, it sets the transmission flag only for its own ESADI-LSPs that RB1 is missing. Missing ESADI-LSPs originated by other ESADI instances will be detected by those other instances.

When receiving an ESADI-PSNP PDU, if the local ESADI instance is DRB, ESADI-LSP PDU requested by the ESADI-PSNP will be multicast on the virtual link.

5. End Station Addresses

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 frames. When such authenticated learned address information is transmitted via the ESADI protocol, the use of authentication in the TRILL ESADI-LSP frames 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 TRILL ESADI protocol should be forgotten by its ESADI-LSP. The time out of the learned end station address is up to the originating RBridge to decide when to remove such information from its ESADI-LSPs (or up to ESADI protocol timeouts if the originating RBridge becomes inaccessible).

If RBridge RBn participating in the TRILL ESADI protocol for VLAN-x is no longer appointed forwarder for VLAN-x on any port where it is providing end station service, it ceases to participate in ESADI after sending a final ESADI-LSP nulling out its ESADI-LSP information. (All other RBridges that are VLAN-x forwarder on at least one of their ports notice that the link state data for RBn has changed to show that it no longer has a link in VLAN-x. In response, they forget all end station address information they have learned from decapsulating VLAN-x frames that show RBn as the ingress RBridge.)

6. ESADI-LSP Contents

The only PDUs used in ESADI are the Level 1 ESADI-LSP, ESADI-CSNP, and ESADI-PSNP PDUs. The content of an ESADI-LSP consists of zero or more MAC Reachability TLVs, optionally an Authentication TLV, and exactly one ESADI parameter APPsub-TLV in ESADI-LSP zero. This section specifies the format for ESADI parameter data APPsub-TLV and gives the reference for the ESADI MAC Reachability TLV. In the future, there may be other TLVs or sub-TLVs carried in ESADI-LSPs.

ESADI-LSP number zero MUST NOT exceed 1470 minus 24 bytes in length (1446 bytes) but if received 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 GENAPP TLV in ESADI-LSP number zero. If it is missing from ESADI-LSP number zero, priority is assumed to be zero and CSNP time 40. If there is more than one occurrence in ESADI-LSP zero, the first occurrence will be used. Occurrences of the ESADI parameter data APPsub-TLV in non-zero ESADI-LSP fragments are ignored.

```

+-+--+--+--+--+
| Type                | (1 byte)
+-+--+--+--+--+
| Length              | (1 byte)
+-+--+--+--+--+
|R| Priority          | (1 byte)
+-+--+--+--+--+
| CSNP Time          | (1 byte)
+-+--+--+--+--+
| Reserved for expansion | (variable)
+-+--+--+--+...
```

Figure 4. ESADI Parameter APPsub-TLV

Type: set to ESADI-PARAM subTLV (TRILL APPsub-TLV type 1).

Length: Set to 2 to 255.

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

Priority: The Priority field gives the ESADI instance's priority for being DRB on the TRILL ESADI virtual link for the VLAN 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 link, will send at least 3 EASDI-CSNP PDUs. It defaults to 30 seconds for an RBridge participating in ESADI for which it has not been configured.

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 as 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.

To avoid conflict with the Inner.VLAN ID, the TLVs in TRILL ESADI PDUs, including the MAC-RI TLV, MUST NOT contain the VLAN ID. If a VLAN-ID is present in the MAC-RI TLV, it is ignored. In the encapsulated TRILL ESADI frame, only the Inner.VLAN tag indicates the VLAN to which the ESADI-LSP applies.

7. IANA Considerations

IANA allocation considerations are given below.

7.1 ESADI Participation and Capability Flags

IANA is requested to allocate an "ESADI Participation" and the "capability of receiving unicast ESADI PDU" bit in the Interested VLANs and Spanning Tree Roots sub-TLV [RFC6326]. (bit 2 and 3 respectively in the Interested VLANs field recommended) If bit 2 is a one, it indicates that the originating RBridge is participating in ESADI for the indicated VLAN or VLANs. If bit 3 is a one, it indicates that the originating RBridge has the capability of receiving and processing unicast ESADI PDUs.

7.2 TRILL GENAPP TLV

IANA is requested to allocate an IS-IS Application Identifier under the Generic Information TLV (#251) for TRILL [RFCgenapp] and to create a subregistry in the TRILL Parameters Registry for "TRILL APPsub-TLVs under IS-IS TLV #251 Application Identifier #TBD". The initial contents of this subregistry are as follows:

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

TRILL APPsub-TLV Types 2 through 254 are available for allocation by Standard Action, as modified by [RFC4020]. The standards track RFC causing such an allocation 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 TRILL IS-IS protocol operation including the establishment of IS-IS adjacencies, the update process, and the decision process for TRILL IS-IS [IS-IS] [RFC1195] [RFC6327]. The TRILL Generic Information TLV MUST NOT be used in an IS-IS instance zero [MultiInstance].

The V, I, D, and S flags in the initial flags byte of a TRILL Generic Information TLV have the meanings specified in [RFCgenapp] but are not currently used as TRILL operates as a Level 1 IS-IS area and no

semantics is 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 GENAPP TLV.

The ESADI Parameters information, for which 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

For general TRILL Security Considerations, see [[RFC6325](#)].

More TBD

9. References

Normative and informative references for this document are below.

9.1 Normative references

- [IS-IS] - International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, 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.
- [RFC4020] - Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", [BCP 100](#), [RFC 4020](#), February 2005.
- [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.
- [RFC6326] - Eastlake, D., Banerjee, A., Dutt, D., Perlman, R., and A. Ghanwani, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [RFC 6326](#), July 2011.
- [RFC6327] - Eastlake 3rd, D., Perlman, R., Ghanwani, A., Dutt, D., and V. Manral, "Routing Bridges (RBridges): Adjacency", [RFC 6327](#), July 2011.
- [RFC6439] - Perlman, R., Eastlake, D., Li, Y., Banerjee, A., and F. Hu, "Routing Bridges (RBridges): Appointed Forwarders", [RFC 6439](#), November 2011.
- [RFCgenapp] - Ginsberg, L., S. Previdi, M. Shand, "Advertising Generic Information in IS-IS", [draft-ietf-isis-genapp-04.txt](#), in RFC Editor's queue.
- [ClearCorrect] - [draft-ietf-trill-clear-correct](#), work in progress.

9.2 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.
- [MultiInstance] - Previdi, S., L. Ginsberg, M. Shand, A. Roy, D. Ward, [draft-ietf-isis-mi](#), work in progress.
- [MultiLevel] - Perlman, R., D. Eastlake, A. Ghanwani, H. Zhai, [draft-perlman-trill-rbridge-multilevel](#), work in progress.
- [RFC3748] - Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, Ed., "Extensible Authentication Protocol (EAP)", [RFC 3748](#), June 2004.
- [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.

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

Copyright and IPR Provisions

Copyright (c) 2012 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.

