INTERNET-DRAFT
Intended status: Proposed Standard

Expires: January 14, 2014

Donald Eastlake
Yizhou Li
Huawei
Radia Perlman
Intel

July 15, 2013

TRILL: Interface Addresses APPsub-TLV
<draft-eastlake-trill-ia-appsubtlv-00.txt>

Abstract

This document specifies a TRILL (Transparent Interconnection of Lots of Links) IS-IS application sub-TLV that enables the reporting by a TRILL switch sets of addresses such that all of the addresses in each set designate the same interface (port). For example, an EUI-48 MAC (Extended Unique Identifier 48-bit, Media Access Control) address, IPv4 address, and IPv6 address can be reported as all corresponding to the same interface. Such information could be used, for example, to synthesize responses to or by-pass the need for the Address Resolution Protocol (ARP), the IPv6 Neighbor Discovery (ND) protocol, of the flooding of unknown MAC addresses, in some cases.

Status of This Memo

Table of Contents

1. Introduction

This document specifies a TRILL (Transparent Interconnection of Lots of Links) [RFC6325] IS-IS application sub-TLV (APPsub-TLV [RFC6823]) that enables the convenient representation of sets of addresses such that all of the addresses in each set designate the same end station interface (port). For example, an EUI-48 MAC (Extended Unique Identifier 48-bit, Media Access Control [RFC5342bis]) address, IPv4 address, and IPv6 address can be reported as all three corresponding to the same interface.

This APPsub-TLV is used inside the TRILL GENINFO TLV as specified in [ESADI]. It is expected to be used in Directory Assisted TRILL Edge services [DirectoryFramework].

Although, in some IETF protocols, address field types are represented by EtherType [RFC5342bis] or Hardware Type [RFC5494] only Address Family Number is used in this APPsub-TLV.

1.1 Conventions Used in This Document

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. Format of the Interface Addresses APPsub-TLV

The Interface Addresses APPsub-TLV is used to indicate that a set of addresses indicate the same end-station interface and to associate that interface with the TRILL switch by which the interface is reachable. These addresses can be in different address families. For example, it can be used to declare that an end-station interface with a particular IPv4 address, IPv6 address, and EUI-48 MAC address is reachable from a particular TRILL switch.

The Template field indicates certain well known sets of addresses or gives a number of AFNs. When AFNs are listed, the set of AFNs provides an explicit template for the type and order of addresses in each Address Set.

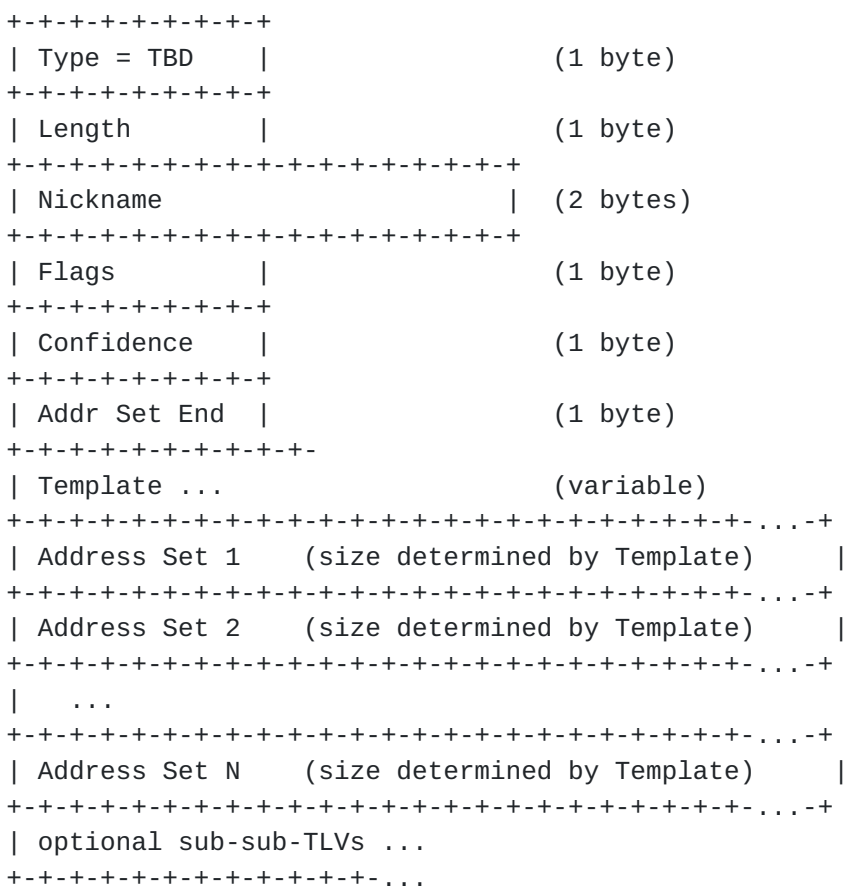

Figure 1. The Interface Addresses APPsub-TLV

- Type: Interface Addresses TRILL APPsub-TLV type, set to TBD[#2 suggested] (IA-SUBTLV).

- Length: Variable, minimum 5. If length is 4 or less, the APPsub-TLV MUST be ignored.

o Nickname: The nickname of the RBridge by which the address sets

are reachable.

- Flags: A byte of flags as follows:

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|D|L|   Resv    |
+-+-+-+-+-+-+-+-+
```

  D: If D is one, the APPsub-TLV contains Push Directory information.

  L: If L is one, the APPsub-TLV contains information learned locally be observing ingressed frames. (Both D and L can one in the same APPsub-TLV.)

  Resv: Additional reserved flag bits that MUST be sent as zero and ignored on receipt.

- Confidence: This 8-bit quantity indicates the confidence level in the addresses being transported [RFC6325].

- Addr Set End: The unsigned offset of the byte, within the TLV value part, of the last byte of the last Address Set. This will be the byte just before the first sub-TLV if any sub-TLVs are present. [RFC5305]

- Template: The initial byte of this field is the unsigned integer K. It K has a value from 1 to 63, it indicates that this initial byte is followed by a list of K AFNs (Address Family Numbers) in the template specifying the structure and order of each Address Set occurring later in the TLV. The minimum valid value is 1. If K is 64 to 255, it indicates that the Template for each Address Set is a specific well known Template. If the Template includes explicit AFNs, they look like the following.

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| AFN 1                         |  (2 bytes)
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| AFN 2                         |  (2 bytes)
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| AFN K                         |  (2 bytes)
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

  For K in the 64 to 255 range, some values indicate combinations of a specific number of 48-bit MAC addresses, IPv4 addresses, and IPv6 addresses in that order. If M is the number of MAC addresses

(limited to 1 or 2), v4 is the number of IPv4 addresses (limited

to 0, 1, or 2) and v6 is the number of IPv6 addresses (limited to 0 through 4 inclusive), the value of K is

K = 63 + M + 2*v4 + 6*v6

That equation specifies values of K from 64 through 93. Values from 94 through 255 are available for assignment by IETF Review.

- AFN: A two-byte Address Family Number. The number of AFNs present is given in first byte of the Template field if that value is less than 64. This sequence specifies the structure of the Address Sets occurring later in the TLV. For example, if Template Size is 2 and the two AFNs present are the AFNs for EUI-48 and IPv4, in that order, then each Address set present will consist of a 6-byte MAC address followed by a 4-byte IPv4 address. If any AFNs are present that are unknown to the receiving IS and the length of the corresponding address is not provided by a sub-TLV as specified below, the receiving IS will be unable to parse the Address Sets and MUST ignore the enclosing TLV.

- Address Set: Each address set consists of a sequence of addresses of the types given by the Template earlier in the TLV. No alignment, other than to a byte boundary, is guaranteed. The addresses in each Address Set are contiguous with no unused bytes between them and the Address Sets are contiguous with no unused bytes between Address Sets. The Address Sets must fit within the TLV. If the product of the size of an Address Set and the number of Address Sets is so large that this is not true, the APPsub-TLV is ignored.

- sub-sub-TLVs: If the Address Sets indicated by Addr Sets End do not completely fill the Length of the TLV, the remaining bytes are parsed as sub-sub-TLVs [RFC5305]. Any such sub-sub-TLVs that are not known to the receiving RBridge are ignored. Should this not be possible, for example there is only one remaining byte or an apparent sub-sub-TLV extends beyond the end of the TLV, the containing IA-APPsub-TLV is considered corrupt and is ignored. Several sub-sub-TLV types are specified in Section 3.

Different IA-APPsub-TLVs within the same or different ESADI-LSPs or Pull Directory responses from the same RBridge may have different Templates. The same AFN may occur more than once in a Template and the same address may occur in more than one address set. For example, an EUI-48 MAC address interface might have three IPv6 addresses. This could be represented by an IA-APPsub-TLV whose Template specifically provided for one EUI-48 address and three IPv6 addresses, which might be an efficient format if there were multiple interfaces with that pattern. Alternatively, a Template with one EUI-48 and one IPv6

address could be used in an IA-APPsub-TLV with three address sets each having the same EUI-48 address but different IPv6 addresses,

which might be the most efficient format if only one interface had multiple IPv6 addresses and other interfaces had only one IPv6 address.

In order to be able to parse the Address Sets, a receiving RBridge must know at least the size of the address each AFN in the Template specifies; however, the presence of the Addr Set End field means that the sub-TLVs, if any, can always be located by a receiving IS. An RBridge can be assumed to know the size of EUI-48, IPv4, and IPv6 addresses (AFNs 16389, 1, and 2) and the size of the additional AFNs allocated by the IANA Considerations below. Should an RBridge wish to include an AFN that some receiving RBridge in the campus may not know, it SHOULD include an AFN-Size sub-sub-TLV as described below. If an IA-APPsub-TLV is received with one or more AFNs in its template for which the receiving RBridge does not know the length and for which an AFN-Size sub-sub-TLV is not present, that IA-APPsub-TLV will be ignored.

3. IA-APPsub-TLV sub-sub-TLVs

IA-APPsub-TLVs may have trailing sub-sub-TLVs [RFC5305] as specified below. These sub-sub-TLVs occur after the Address Sets and the amount of space available for sub-sub-TLVs is determined from the overall IA-APPsub-TLV length and the value of the Addr Set End byte.

There is no ordering restriction on sub-sub-TLVs. Unless otherwise specified each sub-sub-TLV type can occur zero, one, or many times in an IA-APPsub-TLV.

3.1 AFN Size sub-sub-TLV

Using this sub-TLV, the originating RBridge can specify the size of an address type. This is useful under two circumstances:

1. One or more AFNs that are unknown to the receiving RBridge appears in the template. If an AFN Size sub-sub-TLV is present for each such AFN, then at least the IA-APPsub-TLV can be parsed.

2. If an AFN occurs in the Template that represents a variable length address, this sub-sub-TLV gives its size for all occurrences in that IA-APPsubTLV.

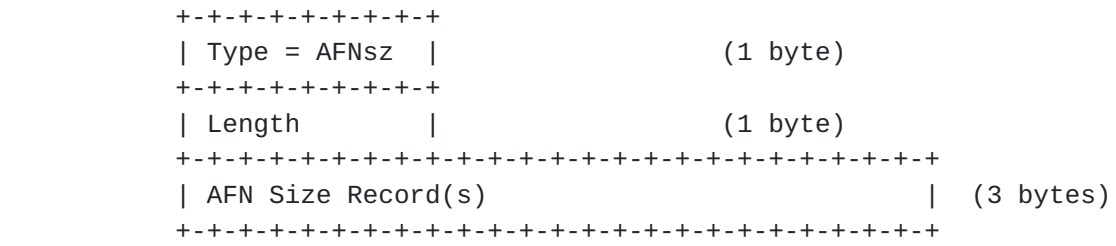

Where each AFN Size Record is structured as follows:

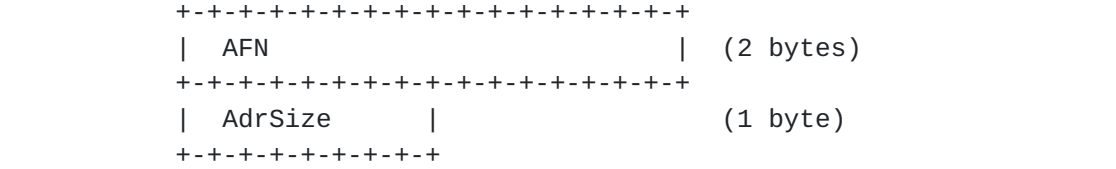

- Type: AFN-Size sub-sub-TLV type, set to 1 (AFNsz).

- Length: 3*n where n is the number of AFN Size Records present. If n is not a multiple of 3, the sub-sub-TLV MUST be ignored.

- AFN Size Record(s): Zero or more 3-byte records, each giving the size of an address type identified by an AFN,

- AFN: The AFN whose length is being specified by the AFN Size Record.

- AdrSize: The length of the address specified by the AFN field.

This sub-sub-TLV may occur multiple times in an enclosing IA-APPsub-TLV.

An AFN Size sub-sub-TLV for any AFN known to the receiving RBridge (which always includes AFN 1, 2, and 16389 and the AFNs specified in xxx) is compared with the size known to the RBridge and if they differ, the IA-APPsub-TLV is ignored.

3.2 Fixed Address sub-sub-TLV

There may be cases where, in an Interface Addresses TLV, the same address would appear across every address set in the TLV. To avoid having a larger template and wasted space in all Address Sets, this sub-sub-TLV can be used to indicate such a fixed address

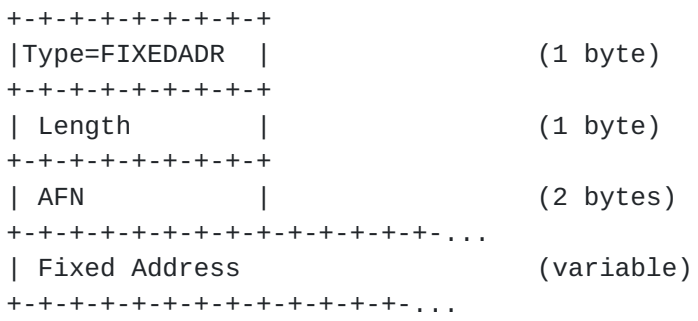

- Type: Data Label sub-sub-TLV type, set to 2 (FIXEDADR).

- Length: variable, minimum 3. If Length is 2 or less, the sub-sub-TLV MUST be ignored.

- AFN: Address Family Number of the Fixed Address.

- Fixed Address: The address of the type indicated by the preceding AFN field that is considered to be part of every Address Set in the IA-APPsub-TLV.

3.3 Data Label sub-sub-TLV

When used with Push or Pull Directories, the Data Label is indicated by the Data Label of the ESADI instance (Push) or RBridge Channel

message (Pull) in which the IA APPsub-TLV appears and any occurrence

of this sub-sub-TLV is ignored. However, the IA APPsub-TLV might be used in other contexts where this sub-sub-TLV indicates the Data Label of the Address Sets and multiple occurrences of this sub-sub-TLV indicate that the Address Sets exist in all of the Data Labels.

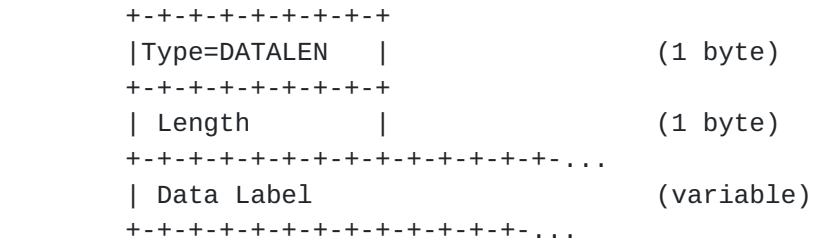

- Type: Data Label sub-TLV type, set to 3 (DATALEN).

- Length: 2 or 3

- Data Label: If length is 2, the bottom 12 bits of the Data Label are a VLAN ID and the top 4 bits are reserved (MUST be sent as zero and ignored on receipt). If the length is 3, the three Data Label bytes contain an FGL [RFCfgl].

3.4 Topology sub-sub-TLV

The presence of this sub-sub-TLV indicates that the Address Sets are in the topology give. If it occurs multiple times, then the Address Sets are in all of the topologies listed.

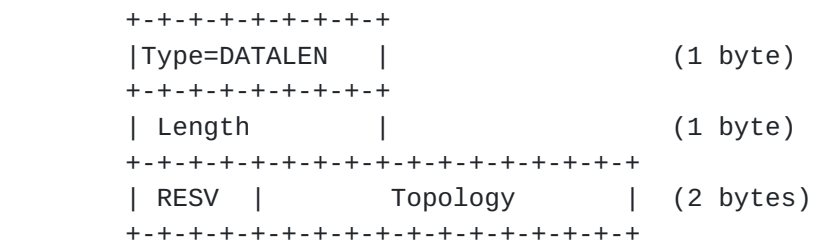

- Type: Data Label sub-TLV type, set to 3 (DATALEN).

- Length: 2.

RESV: Four reserved bits. MUST be sent as zero and ignored on receipt.

- Topology: The 12-bit topology number [RFC5120].

4. Security Considerations

TBD...

5. IANA Considerations

5.1 Additional AFN Number Allocation

IANA is requested to allocate three new AFN numbers as follows:

Number	Description	References
TBD(29)	OUI	[RFC5342bis], this document
TBD(30)	MAC/24	This document.
TBD(31)	IPv6/64	This document.

The OUI AFN is provided so that MAC addresses can be abbreviated if they have the same upper 24 bits. In particular, if there is an OUI provided as a Fixed Address sub-sub-TLV (see Section 5.2.2) then, whenever a MAC/24 address appears within an Address Set (as indicated by the Template), the OUI is used as the first 24 bits of the actual MAC address for the Address Set.

MAC/24 is a 24-bit suffix intended to be pre-fixed by an OUI as in the previous paragraph. In absence of an OUI specified as a Fixed Address in the same APPsub-TLV, an Address Set containing an MAC/24 address cannot be used.

IPv6/64 is an 8-byte quantity that is the first 64 bits of an IPv6 address. If present, there will normally be an EUI-48 or EUI-64 address in the address set to provide the lower 64 bits of the IPv6 address. For this purpose, an EUI-48 is expanded to 64 bits as described in [RFC5342bis].

5.2 IA APPsub-TLV Sub-Sub-TLVs SubRegistry

IANA is requested to establish a new subregistry for sub-sub-TLVs of the Interface Addresses APPsub-TLV with initial contents as shown below.

Name: Interface Addresses APPsub-TLV Sub-Sub-TLVs

Procedure: IETF Review

Reference: This document

Type	Description	Reference
0	Reserved	
1	AFN Size	This document
2	Fixed Address	This document
3	Data Label	This document
4	Topology	This document
5-254	Available	This document
255	Reserved	

Acknowledgments

The authors gratefully acknowledge the contributions and review by the following:

Linda Dunbar

The document was prepared in raw nroff. All macros used were defined within the source file.

Normative References

[RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

[RFC5120] - Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

[RFC5305] - Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

[RFC5342bis] - Eastlake 3rd, D., "IANA Considerations and IETF Protocol Usage for IEEE 802 Parameters", BCP 141, RFC 5342, September 2008.

[RFC5494] - Arkko, J. and C. Pignataro, "IANA Allocation Guidelines for the Address Resolution Protocol (ARP)", RFC 5494, April 2009.

[RFC6325] - Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.

[RFC6823] - Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", RFC 6823, December 2012.

[RFCfgl] - D. Eastlake, M. Zhang, P. Agarwal, R. Perlman, D. Dutt, "TRILL: Fine-Grained Labeling", draft-ietf-trill-fine-labeling-07.txt, in RFC Editor's queue.

[ESADI] - Zhai, H., F. Hu, R. Perlman, D. Eastlake, O. Stokes, "TRILL (Transparent Interconnection of Lots of Links): The ESADI (End Station Address Distribution Information) Protocol", draft-ietf-trill-esadi, work in progress.

Informational References

[DirectoryFramework] - Dunbar, L., D. Eastlkae, R. Perlman, I. Gashinsky, "TRILL Edge Directory Assistance Framework", draft-ietf-trill-directory-framework, work in progress.

[ARP reduction] - Shah, et. al., "ARP Broadcast Reduction for Large Data Centers", Oct 2010.

Authors' Addresses

Donald Eastlake
Huawei Technologies
155 Beaver Street
Milford, MA 01757 USA

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

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012 China

Phone: +86-25-56622310
Email: liyizhou@huawei.com

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

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