

Network Working Group
INTERNET-DRAFT
Intended status: Proposed Standard
Expires: April 14, 2013

Donald Eastlake
Yizhou Li
Huawei
October 15, 2012

Interface Addresses TLV
<[draft-eastlake-isis-ia-tlv-00.txt](#)>

Abstract

This document specifies an IS-IS (Intermediate System to Intermediate System) TLV that enables the reporting by an Intermediate System of sets of addresses of different types such that all of the addresses in each set designate the same interface. 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) or the IPv6 Neighbor Discovery (ND) protocol in some cases.

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.

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.....	3
1.1 Conventions Used in This Document.....	3
2. The Interface Addresses TLV.....	4
3 IA-TLV sub-TLVs.....	7
3.1 AFN Size sub-TLV.....	7
3.2 Data Label sub-TLV.....	8
3.3 Nickname sub-TLV.....	9
4. Example.....	10
5. IANA Considerations.....	11
6. Security Considerations.....	12
7. Normative References.....	13
8. Informative References.....	13

1. Introduction

This document specifies an IS-IS (Intermediate System to Intermediate System [[ISO-10589](#)] [[RFC1195](#)]) TLV that enables the reporting by an Intermediate System in its LSP (Link State PDUs) of sets of addresses of different types such that all of the addresses in each set designate the same interface. For example, an EUI-48 MAC (Extended Unique Identifier 48-bit, Media Access Control [[RFC5342](#)]) address, IPv4 address, and IPv6 address can be reported as all three 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, [[RFC826](#)]), Reverse Address Resolution Protocol (RARP, [[RFC903](#)]) or the IPv6 Neighbor Discovery (ND [[RFC4861](#)]) protocols in some cases.

Although, in some IETF protocols, address field types are represented by EtherType [[802](#)] or Hardware Type [[RFC5494](#)] only Address Family Number is used in this 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. The Interface Addresses TLV

The Interface Addresses TLV is used to indicate that a set of addresses indicate the same interface. These addresses can be in different address families. For example, it can be used to declare that an interface has a particular IPv4 address, IPv6 address, and EUI-48 MAC address.

```

+---+---+---+---+---+
| Type = TBD          | (1 byte)
+---+---+---+---+---+
| Length              | (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
|E|RESV | Topology-ID | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+
| Addr Set End       | (1 byte)
+---+---+---+---+---+
| Confidence         | (1 byte)
+---+---+---+---+---+
| Template Size      | (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+
| AFN 1              | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+
| AFN 2              | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+
| ...
+---+---+---+---+---+---+---+---+---+---+---+---+
| AFN K              | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+...-+
| Address Set 1      (size determined by Template) |
+---+---+---+---+---+---+---+---+---+---+---+---+...-+
| Address Set 2      (size determined by Template) |
+---+---+---+---+---+---+---+---+---+---+---+---+...-+
| ...
+---+---+---+---+---+---+---+---+---+---+---+---+...-+
| Address Set N      (size determined by Template) |
+---+---+---+---+---+---+---+---+---+---+---+---+...-+
| optional sub-TLVs ...
+---+---+---+---+---+---+---+---+---+---+---+---+...

```

Figure 1. The Interface Addresses TLV

- o Type: Interface Addresses TLV type, set to TBD[#247 suggested] (IA-TLV).
- o Length: Variable, minimum 7. If length is 6 or less, the TLV MUST be ignored.

- o E: The E (rEachability) flag is set to one to indicate that the interfaces whose addresses are being given in the TLV are

reachable through the Intermediate System that is advertising the TLV.

- o RESV: Reserved. MUST be sent as zero and ignored on receipt.
- o Topology-ID: This field carries a topology ID [[RFC5120](#)] or zero if topologies are not in use.
- o 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](#)]
- o Confidence: This 8-bit quantity indicates the confidence level in the addresses being transported. The semantics of the Confidence are further defined by the technology using the IA-TLV.
- o Template Size: A byte containing the unsigned integer number K of AFNs (Address Family Numbers) in the template specifying the structure of each Address Set occurring later in the TLV. The minimum valid value is 1 and there must be room in the TLV for the AFNs. If Template Size is zero or too big, the TLV MUST be ignored.
- o AFN: A two-byte Address Family Number. The number of AFNs present is given in the Template Size field. 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 IPv4 and EUI-48, in that order, then each Address set present will consist of a 4-byte IPv4 address followed by a 6-byte MAC 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.
- o Address Set: Each address set consists of a sequence of Template Size number of addresses of the types given by the AFN sequence template earlier in the TLV in the same order as those AFNs. 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 TLV is ignored.
- o 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-TLVs [[RFC5305](#)]. These sub-TLVs are from a new

sequence of sub-TLVs. Any such sub-TLVs that are not known to the receiving IS are ignored. Should this not be possible, for example there is only one remaining byte or an apparent sub-TLV extends beyond the end of the TLV, the containing IA-TLV is considered corrupt and is ignored. Several sub-TLV types are specified in [Section 3](#).

This Reachable Interface Addresses TLV occurs only within LSP PDUs and may occur multiple times in the same or different LSP PDUs with the same or different templates.

Different IA-TLVs within the same or different LSP PDUs from the same IS 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-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-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 IS 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 IS can be assumed to know the size of IPv4 and IPv6 addresses (AFNs 1 and 2) and the size of the additional AFNs allocated by the IANA Considerations below. Should an IS wish to include an AFN that some receiving IS in the campus may not know, it SHOULD include an AFN-Size sub-TLV as described below. If an IA-TLV is received with one or more AFNs in its template for which the receiving IS does not know the length and for which an AFN-Size sub-TLV is not present, that IA-TLV will be ignored.

3 IA-TLV sub-TLVs

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

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

3.1 AFN Size sub-TLV

Using this sub-TLV, the originating IS 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 IS appears in the template. If an AFN Size sub-TLV is present for each such AFN, the at least the IA-TLV can be parses the Address Sets and make use of any address types present that it does understand.
2. If an AFN occurs in the template that represents a variable length address, this sub-TLV gives its size for all occurrences in that IA-TLV.

```

+---+---+---+---+---+
|subType = AFNsz|           (1 byte)
+---+---+---+---+---+
| Length          |           (1 byte)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| AFN Size Record(s) |           (3 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Where each AFN Size Record is structured as follows:

```

+---+---+---+---+---+---+---+---+---+
| AFN |           | (2 bytes)
+---+---+---+---+---+---+---+---+---+
| AdrSize |           (1 byte)
+---+---+---+---+---+

```

- o Type: AFN-Size sub-TLV type, set to 1 (AFNsz).
- o Length: 3*n where n is the number of AFN Size Records present. If n is not a multiple of 3, the sub-TLV MUST be ignored.
- o AFN Size Record(s): Zero or more 3-byte records, each giving the

size of an address type identified by an AFN,

- o AFN: The AFN whose length is being specified by the AFN Size Record.
- o AdrSize: The length of the address specified by the AFN field.

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

The declaration of a size for an AFN address type controls for the occurrences of the AFN in the enclosing IA-TLV and for and subsequent IA-TLVs in the enclosing LSP PDU until the next occurrence in the LSP PDU of an AFN Size sub-TLV for that AFN. Thus, an AFN Size sub-TLV for a fixed size AFN need only be included in the first IA-TLV in a PDU. But one must be included in or before first IA-TLV using an AFN in each PDU where that AFN appears in the template if needed. Otherwise Address Sets will not be parseable by a receiving IS. If multiple AFN Size Records occur for the same AFN in the same AFN Size sub-TLV the last size given controls.

An AFN Size sub-TLV for any AFN known to the receiving IS (which always includes AFN 1 and 2 and the AFNs specified in [Section 5](#)) is compared with the size known to the IS and if they differ, the enclosing IA-TLV is ignored.

[3.2 Data Label sub-TLV](#)

```

+---+---+---+---+
|Type==DATA-LABEL|           (1 byte)
+---+---+---+---+
| Length          |           (1 byte)
+---+---+---+---+---+---+---+...
| Data Label      |           (variable)
+---+---+---+---+---+---+---+...
```

- o Type: Data Label sub-TLV type, set to 2 (DATA-LABEL).
- o Length: variable, minimum 1. If Length is zero, the sub-TLV MUST be ignored.
- o Data Label: A data label under which all the interfaces listed in the TLV can be reached. Exact meaning for different lengths depends on the technology in use. If absent, no label is specified. If more than one different Data Label sub-TLV occurs in an Interface Addresses TLV, it indicates that the interfaces listed can be reach via all the labels given.

For TRILL use, if Length=2, the Data Label is a VLAN-ID which appears right justified in two bytes with four leading bits that

MUST be sent as zero and ignored on receipt.

D. Eastlake, et al.

[Page 8]

3.3 Nickname sub-TLV

```

+---+---+---+---+
|Type=NICKNAME  |          (1 byte)
+---+---+---+---+
| Length        |          (1 byte)
+---+---+---+---+---+---+---+---+...
| Nickname      |          (variable)
+---+---+---+---+---+---+---+---+...

```

- o Type: Data Label sub-TLV type, set to 3 (NICKNAME).
- o Length: variable, minimum 1. If Length is zero, the sub-TLV MUST be ignored.
- o Nickname: The nickname of an Intermediate System through which all the interfaces listed in the TLV can be reached. Exact meaning for different lengths depends on the technology in use. If absent, no nickname is specified. If more than one different Nickname sub-TLV occurs in an Interface Addresses TLV, it indicates that the interfaces listed can be reach via all the nicknames given. Occurrence of one or more Nickname sub-TLVs in an Interface Addresses TLV does not change the effect of the E flag bit in the TLV initial fixed fields; the E flag having the value one still indicates that the interfaces are reachable through the Intermediate System advertising the TLV.

4. Example

TBD

5. IANA Considerations

IANA is requested to allocate a new TLV number for IA-TLV [#247 suggested because #147 is the MAC Reachability (MAC-RI) TLV].

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

Number	Description	References
TBD(26)	EUI-48	RFC 5342 , this document
TBD(27)	EUI-64	RFC 5342 , this document
TBD(28)	IPv6/64	This document.

[[[Curiously enough, the AFN RFC references seem to always use "Address Family Identifier" although IANA uses "Address Family Number". Furthermore, neither of the two RFCs pointed to by the IANA Address Family Number Registry actually has IANA Considerations for Address Family Numbers although the IANA Registry has shows policies for different ranges of AFNs. Conceivably, such IANA Considerations should appear here.]]]

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-64 or EUI-48 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 [[RFC5342](#)].

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

Name: Sub-TLVs for TLV TBD[#247]
 Procedure: Expert Review
 Reference: This document

Type	Description	Reference
----	-----	-----
0	Reserved	
1	AFN Size	This document
2	Nickname	This document
3	Data Label	This document
4-254	Available	This document
255	Reserved	

[[[Should administrative tag sub-TLVs (see [RFC 5130](#)) be allowed?]]]

6. Security Considerations

This document raises no new security issues for IS-IS. IS-IS security may be used to secure IS-IS PDUs containing the IA-TLV. See [[RFC5304](#)] and [[RFC5310](#)].

7. Normative References

- [ISO-10589] - ISO/IEC 10589:2002, Second Edition, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 2002.
- [RFC1195] - Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", 1990.
- [RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5120] - Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), February 2008.
- [RFC5304] - Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", [RFC 5304](#), October 2008.
- [RFC5305] - Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", 2008.
- [RFC5310] - Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), February 2009.

8. Informative References

- [802] - IEEE, "Standard for Local and Metropolitan Area Networks: Overview and Architecture", IEEE Std. 802-2001, 8 March 2002.
- [RFC826] - Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48-bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, [RFC 826](#), November 1982.
- [RFC903] - Finlayson, R., Mann, T., Mogul, J., and M. Theimer, "A Reverse Address Resolution Protocol", STD 38, [RFC 903](#), June 1984.
- [RFC4861] - Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC5342] - Eastlake 3rd, D., "IANA Considerations and IETF Protocol

Usage for IEEE 802 Parameters", [BCP 141](#), [RFC 5342](#), September

D. Eastlake, et al.

[Page 13]

2008.

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

Acknowledgements

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

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

Authors' Addresses

Donald Eastlake
Huawei R&D USA
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: +86-25-56624558
Email: liyizhou@huawei.com

Copyright, Disclaimer, and Additional 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.

