

Internet Engineering Task Force (IETF)
Request for Comments: 7176
Obsoletes: 6326
Category: Standards Track
ISSN: 2070-1721

D. Eastlake 3rd
Huawei
T. Senevirathne
Cisco
A. Ghanwani
Dell
D. Dutt
Cumulus Networks
A. Banerjee
Insieme Networks
May 2014

Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS

Abstract

The IETF Transparent Interconnection of Lots of Links (TRILL) protocol provides optimal pair-wise data frame forwarding without configuration in multi-hop networks with arbitrary topology and link technology; it also provides support for multipathing of both unicast and multicast traffic. This document specifies the data formats and code points for the IS-IS extensions to support TRILL. These data formats and code points may also be used by technologies other than TRILL. This document obsoletes RFC 6326.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7176>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | |
|--|----|
| 1. Introduction | 3 |
| 1.1. Conventions Used in This Document | 4 |
| 2. TLV and Sub-TLV Extensions to IS-IS for TRILL | 4 |
| 2.1. Group Address TLV | 5 |
| 2.1.1. Group MAC Address Sub-TLV | 5 |
| 2.1.2. Group IPv4 Address Sub-TLV | 7 |
| 2.1.3. Group IPv6 Address Sub-TLV | 8 |
| 2.1.4. Group Labeled MAC Address Sub-TLV | 9 |
| 2.1.5. Group Labeled IPv4 Address Sub-TLV | 10 |
| 2.1.6. Group Labeled IPv6 Address Sub-TLV | 11 |
| 2.2. Multi-Topology-Aware Port Capability Sub-TLVs | 12 |
| 2.2.1. Special VLANs and Flags Sub-TLV | 12 |
| 2.2.2. Enabled-VLANs Sub-TLV | 13 |
| 2.2.3. Appointed Forwarders Sub-TLV | 14 |
| 2.2.4. Port TRILL Version Sub-TLV | 15 |
| 2.2.5. VLANs Appointed Sub-TLV | 17 |
| 2.3. Sub-TLVs of the Router Capability and MT-Capability TLVs .. | 17 |
| 2.3.1. TRILL Version Sub-TLV | 18 |
| 2.3.2. Nickname Sub-TLV | 19 |
| 2.3.3. Trees Sub-TLV | 20 |
| 2.3.4. Tree Identifiers Sub-TLV | 20 |
| 2.3.5. Trees Used Identifiers Sub-TLV | 21 |
| 2.3.6. Interested VLANs and Spanning Tree Roots Sub-TLV .. | 22 |
| 2.3.7. VLAN Group Sub-TLV | 24 |
| 2.3.8. Interested Labels and Spanning Tree Roots Sub-TLV .. | 25 |
| 2.3.9. RBridge Channel Protocols Sub-TLV | 27 |
| 2.3.10. Affinity Sub-TLV | 29 |
| 2.3.11. Label Group Sub-TLV | 30 |
| 2.4. MTU Sub-TLV for Extended Reachability and MT-ISN TLVs | 31 |
| 2.5. TRILL Neighbor TLV | 31 |
| 3. MTU PDUs | 33 |

| | |
|---|----|
| 4. Use of Existing PDUs and TLVs | 35 |
| 4.1. TRILL IIH PDUs | 35 |
| 4.2. Area Address | 35 |
| 4.3. Protocols Supported | 35 |
| 4.4. Link State PDUs (LSPs) | 35 |
| 4.5. Originating LSP Buffer Size | 36 |
| 5. IANA Considerations | 36 |
| 5.1. TLVs | 36 |
| 5.2. Sub-TLVs | 36 |
| 5.3. PDUs | 38 |
| 5.4. Reserved and Capability Bits | 38 |
| 5.5. TRILL Neighbor Record Flags | 39 |
| 6. Security Considerations | 39 |
| 7. Changes from RFC 6326 | 39 |
| 8. References | 41 |
| 8.1. Normative References | 41 |
| 8.2. Informative References | 43 |
| 9. Acknowledgements | 44 |

1. Introduction

The IETF Transparent Interconnection of Lots of Links (TRILL) protocol [RFC6325] [RFC7177] provides transparent forwarding in multi-hop networks with arbitrary topology and link technologies using a header with a hop count and link-state routing. TRILL provides optimal pair-wise forwarding without configuration, safe forwarding even during periods of temporary loops, and support for multipathing of both unicast and multicast traffic. Intermediate Systems (ISs) implementing TRILL are called Routing Bridges (RBridges) or TRILL Switches.

This document, in conjunction with [RFC6165], specifies the data formats and code points for the IS-IS [ISO-10589] [RFC1195] extensions to support TRILL. These data formats and code points may also be used by technologies other than TRILL.

This document obsoletes [RFC6326], which generally corresponded to the base TRILL protocol [RFC6325]. There has been substantial development of TRILL since the publication of those documents. The main changes from [RFC6326] are summarized below, and a full list is given in Section 7.

1. Added multicast group announcements by IPv4 and IPv6 address.
2. Added facilities for announcing capabilities supported.
3. Added a tree affinity sub-TLV whereby ISs can request distribution tree association.

4. Added multi-topology support.
5. Added control-plane support for TRILL Data frame fine-grained labels. This support is independent of the data-plane representation.
6. Fixed the verified erratum [Err2869] in [RFC6326].

Changes herein to TLVs and sub-TLVs specified in [RFC6326] are backward compatible.

1.1. Conventions Used in This Document

The terminology and acronyms defined in [RFC6325] are used herein with the same meaning.

Additional acronyms and phrases used in this document are:

BVL - Bit Vector Length

BVO - Bit Vector Offset

IIH - IS-IS Hello

IS - Intermediate System. For this document, all relevant intermediate systems are RBridges [RFC6325].

MAC - Media Access Control

MT - Multi-Topology

NLPID - Network Layer Protocol Identifier

SNPA - Subnetwork Point of Attachment (MAC Address)

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. TLV and Sub-TLV Extensions to IS-IS for TRILL

This section, in conjunction with [RFC6165], specifies the data formats and code points for the TLVs and sub-TLVs for IS-IS to support the IETF TRILL protocol. Information as to the number of occurrences allowed, such as for a TLV in a PDU or set of PDUs or for a sub-TLV in a TLV, is summarized in Section 5.

2.1. Group Address TLV

The Group Address (GADDR) TLV, IS-IS TLV type 142, is carried in an LSP PDU and carries sub-TLVs that in turn advertise multicast group listeners. The sub-TLVs that advertise listeners are specified below. The sub-TLVs under GADDR constitute a new series of sub-TLV types (see Section 5.2).

GADDR has the following format:

```


+++++
|Type=GADDR-TLV |                (1 byte)
+++++
|  Length      |                (1 byte)
+++++
|      sub-TLVs...
+++++

```

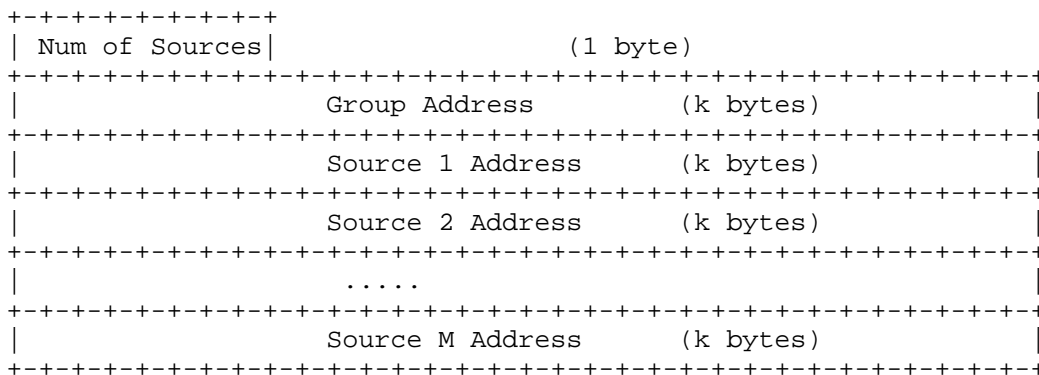
- o Type: TLV type, set to GADDR-TLV 142.
- o Length: variable depending on the sub-TLVs carried.
- o sub-TLVs: The Group Address TLV value consists of sub-TLVs formatted as described in [RFC5305].

2.1.1. Group MAC Address Sub-TLV

The Group MAC Address (GMAC-ADDR) sub-TLV is sub-TLV type number 1 within the GADDR TLV. In TRILL, it is used to advertise multicast listeners by MAC address as specified in Section 4.5.5 of [RFC6325]. It has the following format:



where each group record is of the following form with k=6:



- o Type: GADDR sub-TLV type, set to 1 (GMAC-ADDR).
- o Length: $5 + m + k*n = 5 + m + 6*n$, where m is the number of group records and n is the sum of the number of group and source addresses.
- o RESV: Reserved. 4-bit fields that 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 VLAN ID: This carries the 12-bit VLAN identifier for all subsequent MAC addresses in this sub-TLV, or the value zero if no VLAN is specified.
- o Num Group Recs: A 1-byte unsigned integer that is the number of group records in this sub-TLV.
- o GROUP RECORDS: Each group record carries the number of sources. If this field is zero, it indicates a listener for (*,G), that is, a listener not restricted by source. It then has a 6-byte (48-bit) multicast MAC address followed by 6-byte source MAC addresses. If the sources do not fit in a single sub-TLV, the same group address may be repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GMAC-ADDR sub-TLV is carried only within a GADDR TLV.

2.1.2. Group IPv4 Address Sub-TLV

The Group IPv4 Address (GIP-ADDR) sub-TLV is IS-IS sub-TLV type 2 within the GADDR TLV. It has the same format as the Group MAC Address sub-TLV described in Section 2.1.1 except that $k=4$. The fields are as follows:

- o Type: sub-TLV type, set to 2 (GIP-ADDR).
- o Length: $5 + m + k*n = 5 + m + 4*n$, where m is the number of group records and n is the sum of the number of group and source addresses.
- o Topology-ID: This field carries a topology ID [RFC5120] or zero if topologies are not in use.
- o RESV: Must be sent as zero on transmission and is ignored on receipt.
- o VLAN ID: This carries a 12-bit VLAN identifier that is valid for all subsequent addresses in this sub-TLV, or the value zero if no VLAN is specified.
- o Num Group Recs: A 1-byte unsigned integer that is the number of group records in this sub-TLV.

- o GROUP RECORDS: Each group record carries the number of sources. If this field is zero, it indicates a listener for (*,G), that is, a listener not restricted by source. It then has a 4-byte (32-bit) IPv4 Group Address followed by 4-byte source IPv4 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GIP-ADDR sub-TLV is carried only within a GADDR TLV.

2.1.3. Group IPv6 Address Sub-TLV


The Group IPv6 Address (GIPV6-ADDR) sub-TLV is IS-IS sub-TLV type 3 within the GADDR TLV. It has the same format as the Group MAC Address sub-TLV described in Section 2.1.1 except that $k=16$. The fields are as follows:

- o Type: sub-TLV type, set to 3 (GIPV6-ADDR).
- o Length: $5 + m + k*n = 5 + m + 16*n$, where m is the number of group records and n is the sum of the number of group and source addresses.
- o Topology-Id: This field carries a topology ID [RFC5120] or zero if topologies are not in use.
- o RESV: Must be sent as zero on transmission and is ignored on receipt.
- o VLAN ID: This carries a 12-bit VLAN identifier that is valid for all subsequent addresses in this sub-TLV, or the value zero if no VLAN is specified.
- o Num Group Recs: A 1-byte unsigned integer that is the number of group records in this sub-TLV.
- o GROUP RECORDS: Each group record carries the number of sources. If this field is zero, it indicates a listener for (*,G), that is, a listener not restricted by source. It then has a 16-byte (128-bit) IPv6 Group Address followed by 16-byte source IPv6 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

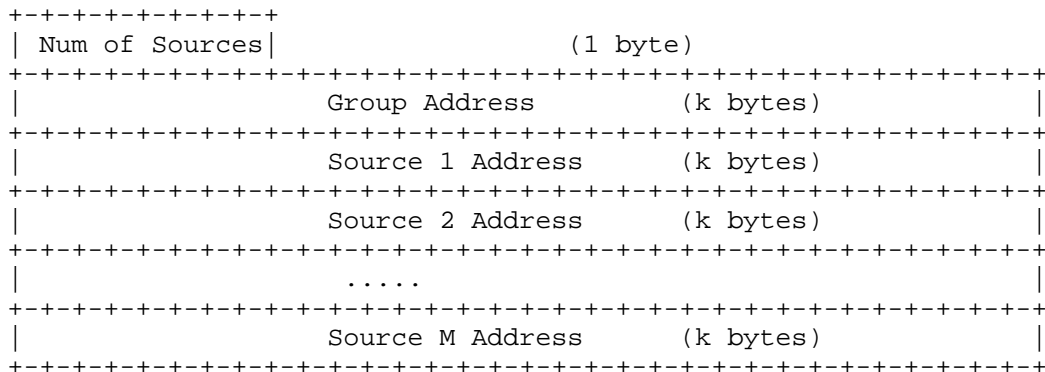
The GIPV6-ADDR sub-TLV is carried only within a GADDR TLV.

2.1.4. Group Labeled MAC Address Sub-TLV

The GMAC-ADDR sub-TLV of the Group Address (GADDR) TLV specified in Section 2.1.1 provides for a VLAN ID. The Group Labeled MAC Address sub-TLV, below, extends this to a fine-grained label.



where each group record is of the following form with k=6:



- o Type: GADDR sub-TLV type, set to 4 (GLMAC-ADDR).
- o Length: $6 + m + k*n = 6 + m + 6*n$, where m is the number of group records and n is the sum of the number of group and source addresses.

- o RESV: Reserved. 4-bit field that 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 Label: This carries the fine-grained label [RFC7172] identifier for all subsequent MAC addresses in this sub-TLV, or the value zero if no label is specified.
- o Num Group Recs: A 1-byte unsigned integer that is the number of group records in this sub-TLV.
- o GROUP RECORDS: Each group record carries the number of sources. If this field is zero, it indicates a listener for (*,G), that is, a listener not restricted by source. It then has a 6-byte (48-bit) multicast address followed by 6-byte source MAC addresses. If the sources do not fit in a single sub-TLV, the same group address may be repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GLMAC-ADDR sub-TLV is carried only within a GADDR TLV.

2.1.5. Group Labeled IPv4 Address Sub-TLV

The Group Labeled IPv4 Address (GLIP-ADDR) sub-TLV is IS-IS sub-TLV type 5 within the GADDR TLV. It has the same format as the Group Labeled MAC Address sub-TLV described in Section 2.1.4 except that $k=4$. The fields are as follows:

- o Type: sub-TLV type, set to 5 (GLIP-ADDR).
- o Length: $6 + m + k*n = 6 + m + 4*n$, where m is the number of group records and n is the sum of the number of group and source addresses.
- o Topology-ID: This field carries a topology ID [RFC5120] or zero if topologies are not in use.
- o RESV: Must be sent as zero on transmission and is ignored on receipt.
- o Label: This carries the fine-grained label [RFC7172] identifier for all subsequent IPv4 addresses in this sub-TLV, or the value zero if no label is specified.
- o Num Group Recs: A 1-byte unsigned integer that is the number of group records in this sub-TLV.

- o GROUP RECORDS: Each group record carries the number of sources. If this field is zero, it indicates a listener for (*,G), that is, a listener not restricted by source. It then has a 4-byte (32-bit) IPv4 Group Address followed by 4-byte source IPv4 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GLIP-ADDR sub-TLV is carried only within a GADDR TLV.

2.1.6. Group Labeled IPv6 Address Sub-TLV

The Group Labeled IPv6 Address (GLIPV6-ADDR) sub-TLV is IS-IS sub-TLV type 6 within the GADDR TLV. It has the same format as the Group Labeled MAC Address sub-TLV described in Section 2.1.4 except that $k=16$. The fields are as follows:

- o Type: sub-TLV type, set to 6 (GLIPV6-ADDR).
- o Length: $6 + m + k*n = 6 + m + 16*n$, where m is the number of group records and n is the sum of the number of group and source addresses.
- o Topology-Id: This field carries a topology ID [RFC5120] or zero if topologies are not in use.
- o RESV: Must be sent as zero on transmission and is ignored on receipt.
- o Label: This carries the fine-grained label [RFC7172] identifier for all subsequent IPv6 addresses in this sub-TLV, or the value zero if no label is specified.
- o Num Group Recs: A 1-byte unsigned integer that is the number of group records in this sub-TLV.
- o GROUP RECORDS: Each group record carries the number of sources. If this field is zero, it indicates a listener for (*,G), that is, a listener not restricted by source. It then has a 16-byte (128-bit) IPv6 Group Address followed by 16-byte source IPv6 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GLIPV6-ADDR sub-TLV is carried only within a GADDR TLV.

2.2. Multi-Topology-Aware Port Capability Sub-TLVs

TRILL makes use of the Multi-Topology-Aware Port Capability TLV (MT-Port-Cap-TLV) as specified in [RFC6165]. The following subsections specify the sub-TLVs transported by the MT-Port-Cap-TLV for TRILL.

2.2.1. Special VLANs and Flags Sub-TLV

In TRILL, a Special VLANs and Flags (VLAN-FLAGS) sub-TLV is carried in every IIH PDU. It has the following format:

```

+---+---+---+---+---+---+---+---+---+---+
|   Type                               | (1 byte)
+---+---+---+---+---+---+---+---+---+---+
|   Length                              | (1 byte)
+---+---+---+---+---+---+---+---+---+---+
|   Port ID                             | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+
|   Sender Nickname                     | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+
|AF|AC|VM|BY|   Outer.VLAN              | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+
|TR|R |R |R |   Designated-VLAN        | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+
    
```

- o Type: sub-TLV type, set to MT-Port-Cap-TLV VLAN-FLAGS sub-TLV 1.
- o Length: 8.
- o Port ID: An ID for the port on which the enclosing TRILL IIH PDU is being sent as specified in [RFC6325], Section 4.4.2.
- o Sender Nickname: If the sending IS is holding any nicknames as discussed in [RFC6325], Section 3.7, one MUST be included here. Otherwise, the field is set to zero. This field is to support intelligent end stations that determine the egress IS (RBridge) for unicast data through a directory service or the like and that need a nickname for their first hop to insert as the ingress nickname to correctly format a TRILL Data frame (see [RFC6325], Section 4.6.2, point 8). It is also referenced in connection with the VLANs Appointed Sub-TLV (see Section 2.2.5) and can be used as the egress on one-hop RBridge Channel messages [RFC7178], for example, those use for BFD over TRILL [RFC7175].
- o Outer.VLAN: A copy of the 12-bit outer VLAN ID of the TRILL IIH frame containing this sub-TLV, as specified in [RFC6325], Section 4.4.5.


- o Designated-VLAN: The 12-bit ID of the Designated VLAN for the link, as specified in [RFC6325], Section 4.2.4.2.
- o AF, AC, VM, BY, and TR: These flag bits have the following meanings when set to one, as specified in the listed section of [RFC6325]:

| RFC 6325 | | |
|----------|---------|---|
| Bit | Section | Meaning if bit is one |
| ----- | | |
| AF | 4.4.2 | Originating IS believes it is Appointed Forwarder for the VLAN and port on which the containing IIH PDU was sent. |
| AC | 4.9.1 | Originating port configured as an access port (TRILL traffic disabled). |
| VM | 4.4.5 | VLAN mapping detected on this link. |
| BY | 4.4.2 | Bypass pseudonode. |
| TR | 4.9.1 | Originating port configured as a trunk port (end-station service disabled). |

- o R: Reserved bit. MUST be sent as zero and ignored on receipt.

2.2.2. Enabled-VLANs Sub-TLV

The optional Enabled-VLANs sub-TLV specifies the VLANs enabled at the port of the originating IS on which the containing Hello was sent, as specified in [RFC6325], Section 4.4.2. It has the following format:



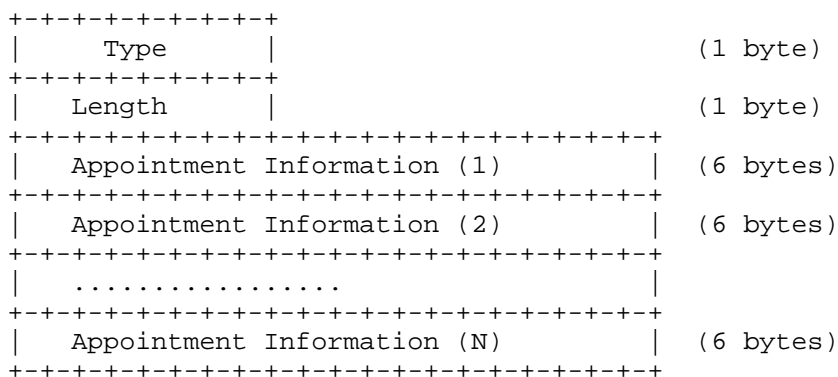
- o Type: sub-TLV type, set to MT-Port-Cap-TLV Enabled-VLANs sub-TLV 2.
- o Length: Variable, minimum 3.

- o RESV: 4 reserved bits that MUST be sent as zero and ignored on receipt.
- o Start VLAN ID: The 12-bit VLAN ID that is represented by the high-order bit of the first byte of the VLAN bit-map.
- o VLAN bit-map: The highest-order bit indicates the VLAN equal to the start VLAN ID, the next highest bit indicates the VLAN equal to start VLAN ID + 1, continuing to the end of the VLAN bit-map field.


If this sub-TLV occurs more than once in a Hello, the set of enabled VLANs is the union of the sets of VLANs indicated by each of the Enabled-VLAN sub-TLVs in the Hello.

2.2.3. Appointed Forwarders Sub-TLV

The Designated RBridge (DRB) on a link uses the Appointed Forwarders sub-TLV to inform other ISSs on the link that they are the designated VLAN-x forwarder for one or more ranges of VLAN IDs as specified in [RFC6439]. It has the following format:



where each appointment is of the form:

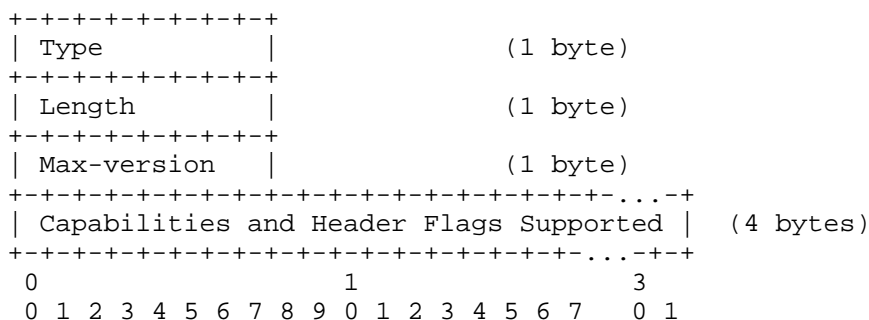


- o Type: sub-TLV type, set to MT-Port-Cap-TLV AppointedFwrdrsr sub-TLV 3.
- o Length: 6*n bytes, where there are n appointments.
- o Appointee Nickname: The nickname of the IS being appointed a forwarder.
- o RESV: 4 bits that MUST be sent as zero and ignored on receipt.
- o Start.VLAN, End.VLAN: This VLAN ID range is inclusive. Setting both Start.VLAN and VLAN.end to the same value indicates a range of one VLAN ID. If Start.VLAN is not equal to VLAN.end and Start.VLAN is 0x000, the sub-TLV is interpreted as if Start.VLAN was 0x001. If Start.VLAN is not equal to VLAN.end and VLAN.end is 0xFFFF, the sub-TLV is interpreted as if VLAN.end was 0xFFFE. If VLAN.end is less than Start.VLAN, the sub-TLV is ignored. If both Start.VLAN and VLAN.end are 0x000 or both are 0xFFFF, the sub-TLV is ignored. The values 0x000 or 0xFFFF are not valid VLAN IDs, and a port cannot be enabled for them.

An IS's nickname may occur as Appointed Forwarder for multiple VLAN ranges by occurrences of this sub-TLV within the same or different MT Port Capability TLVs within an IIH PDU. See [RFC6439].

2.2.4. Port TRILL Version Sub-TLV

The Port TRILL Version (PORT-TRILL-VER) sub-TLV indicates the maximum version of the TRILL standard supported and the support of optional hop-by-hop capabilities. By implication, lower versions are also supported. If this sub-TLV is missing from an IIH, it is assumed that the originating IS only supports the base version (version zero) of the protocol [RFC6325] and supports no optional capabilities indicated by this sub-TLV.




- o Type: MT-Port-Cap-TLV sub-TLV type, set to 7 (PORT-TRILL-VER).
- o Length: 5.
- o Max-version: A one-byte unsigned integer set to the maximum version supported.
- o Capabilities and Header Flags Supported: A bit vector of 32 bits numbered 0 through 31 in network order. Bits 3 through 13 indicate that the corresponding TRILL Header hop-by-hop extended flags [RFC7179] are supported. Bits 0 through 2 and 14 to 31 are reserved to indicate support of optional capabilities. A one bit indicates that the flag or capability is supported by the sending IS. Bits in this field MUST be set to zero except as permitted for a capability being advertised or if a hop-by-hop extended header flag is supported.

This sub-TLV, if present, MUST occur in an MT-Port-Cap-TLV in a TRILL IIH. If there is more than one occurrence, the minimum of the supported versions is assumed to be correct and a capability or header flag is assumed to be supported only if indicated by all occurrences. The flags and capabilities for which support can be indicated in this sub-TLV are disjoint from those in the TRILL-VER sub-TLV (Section 2.3.1) so they cannot conflict. The flags and capabilities indicated in this sub-TLV relate to hop-by-hop processing that can differ between the ports of an IS (RBridge) and thus must be advertised in IIHs. For example, a capability requiring cryptographic hardware assist might be supported on some ports and not others. However, the TRILL version is the same as that in the PORT-TRILL-VER sub-TLV. An IS, if it is adjacent to the sending IS of TRILL version sub-TLV(s), uses the TRILL version it received in PORT-TRILL-VER sub-TLV(s) in preference to that received in TRILL-VER sub-TLV(s).

2.2.5. VLANs Appointed Sub-TLV

The optional VLANs Appointed sub-TLV specifies, for the port of the originating IS on which the containing Hello was sent, the VLANs for which it is Appointed Forwarder. This sub-TLV has the following format:



- o Type: sub-TLV type, set to MT-Port-Cap-TLV VLANS-Appointed sub-TLV 8.
- o Length: Variable, minimum 3.
- o RESV: 4 reserved bits that MUST be sent as zero and ignored on receipt.
- o Start VLAN ID: The 12-bit VLAN ID that is represented by the high-order bit of the first byte of the VLAN bit-map.
- o VLAN bit-map: The highest-order bit indicates the VLAN equal to the start VLAN ID, the next highest bit indicates the VLAN equal to start VLAN ID + 1, continuing to the end of the VLAN bit-map field.


If this sub-TLV occurs more than once in a Hello, the originating IS is declaring that it believes itself to be Appointed Forwarder on the port on which the enclosing IIH was sent for the union of the sets of VLANs indicated by each of the VLANs-Appointed sub-TLVs in the Hello.

2.3. Sub-TLVs of the Router Capability and MT-Capability TLVs

The Router Capability TLV is specified in [RFC4971] and the MT-Capability TLV in [RFC6329]. All of the following sub-sections specify sub-TLVs that can be carried in the Router Capability TLV (#242) and the MT-Capability TLV (#144) with the same sub-TLV number for both TLVs. These TLVs are in turn carried only by LSPs.

2.3.1. TRILL Version Sub-TLV

The TRILL Version (TRILL-VER) sub-TLV indicates the maximum version of the TRILL standard supported and the support of optional capabilities by the originating IS. By implication, lower versions are also supported. If this sub-TLV is missing, it is assumed that the originating IS only supports the base version (version zero) of the protocol [RFC6325], and no optional capabilities indicated by this sub-TLV are supported.



- o Type: Router Capability sub-TLV type, set to 13 (TRILL-VER).
- o Length: 5.
- o Max-version: A one-byte unsigned integer set to the maximum version supported.
- o Capabilities and Header Flags Supported: A bit vector of 32 bits numbered 0 through 31 in network order. Bits 14 through 31 indicate that the corresponding TRILL Header extended flags [RFC7179] are supported. Bits 0 through 13 are reserved to indicate support of optional capabilities. A one bit indicates that the originating IS supports the flag or capability. For example, support of multi-level TRILL IS-IS [MultiLevel]. Bits in this field MUST be set to zero except as permitted for a capability being advertised or an extended header flag supported.


This sub-TLV, if present in a Router Capability TLV, MUST occur in the LSP number zero for the originating IS. If found in a Router Capability TLV in other fragments, it is ignored. If there is more than one occurrence in LSP number zero, the minimum of the supported versions is assumed to be correct, and an extended header flag or capability is assumed to be supported only if indicated by all occurrences. The flags and capabilities for which support can be indicated in this sub-TLV are disjoint from those in the PORT-TRILL-VER sub-TLV (Section 2.2.4) so they cannot conflict. However, the

TRILL version is the same as that in the PORT-TRILL-VER sub-TLV, and an IS that is adjacent to the originating IS of TRILL-VER sub-TLV(s) uses the TRILL version it received in PORT-TRILL-VER sub-TLV(s) in preference to that received in TRILL-VER sub-TLV(s).


For multi-topology-aware TRILL Switches, the TRILL version and capabilities announced for the base topology are assumed to apply to all topologies for which a separate TRILL version announcement does not occur in an MT-Capability TLV. Such announcements for non-zero topologies need not occur in fragment zero.

2.3.2. Nickname Sub-TLV

The Nickname (NICKNAME) Router Capability sub-TLV carries information about the nicknames of the originating IS, along with information about its priority to hold those nicknames and the priority for each nickname to be a tree root as specified in [RFC6325], Section 3.7.3. Multiple instances of this sub-TLV may occur.



where each nickname record is of the form:

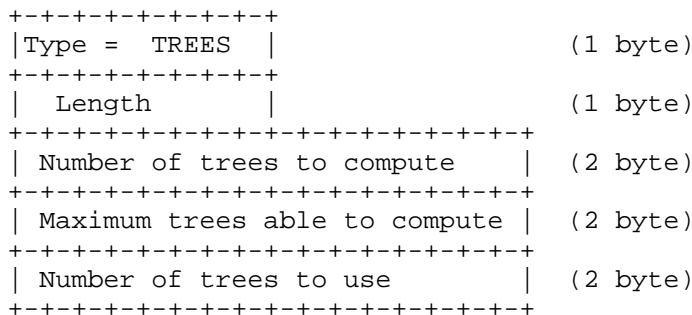


- o Type: Router Capability and MT-Capability sub-TLV type, set to 6 (NICKNAME).
- o Length: 5*n, where n is the number of nickname records present.

- o Nickname.Pri: An 8-bit unsigned integer priority to hold a nickname as specified in Section 3.7.3 of [RFC6325].
- o Tree Root Priority: This is an unsigned 16-bit integer priority to be a tree root as specified in Section 4.5 of [RFC6325].
- o Nickname: This is an unsigned 16-bit integer as specified in Section 3.7 of [RFC6325].

2.3.3. Trees Sub-TLV

Each IS providing TRILL service uses the TREES sub-TLV to announce three numbers related to the computation of distribution trees as specified in Section 4.5 of [RFC6325]. Its format is as follows:

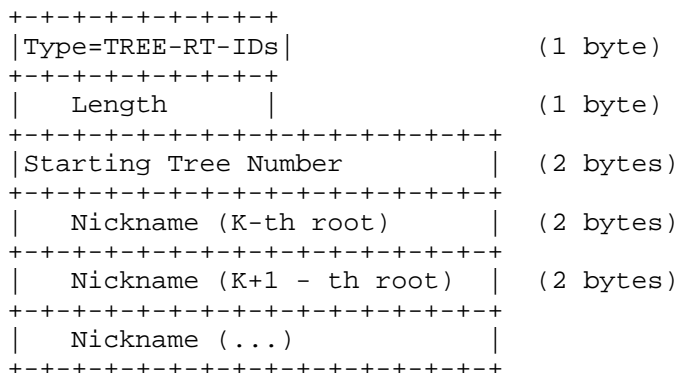


- o Type: Router Capability and MT-Capability sub-TLV type, set to 7 (TREES).
- o Length: 6.
- o Number of trees to compute: An unsigned 16-bit integer as specified in Section 4.5 of [RFC6325].
- o Maximum trees able to compute: An unsigned 16-bit integer as specified in Section 4.5 of [RFC6325].
- o Number of trees to use: An unsigned 16-bit integer as specified in Section 4.5 of [RFC6325].

2.3.4. Tree Identifiers Sub-TLV

The Tree Identifiers (TREE-RT-IDs) sub-TLV is an ordered list of nicknames. When originated by the IS that has the highest priority to be a tree root, it lists the distribution trees that the other ISs are required to compute as specified in Section 4.5 of [RFC6325]. If

this information is spread across multiple sub-TLVs, the starting tree number is used to allow the ordered lists to be correctly concatenated. The sub-TLV format is as follows:



- o Type: Router Capability and MT-Capability sub-TLV type, set to 8 (TREE-RT-IDs).
- o Length: $2 + 2*n$, where n is the number of nicknames listed.
- o Starting Tree Number: This identifies the starting tree number of the nicknames that are trees for the domain. This is set to 1 for the sub-TLV containing the first list. Other Tree-Identifiers sub-TLVs will have the number of the starting list they contain. In the event the same tree identifier can be computed from two such sub-TLVs and they are different, then it is assumed that this is a transient condition that will get cleared. During this transient time, such a tree SHOULD NOT be computed unless such computation is indicated by all relevant sub-TLVs present.
- o Nickname: The nickname at which a distribution tree is rooted.

2.3.5. Trees Used Identifiers Sub-TLV

This Router Capability sub-TLV has the same structure as the Tree Identifiers sub-TLV specified in Section 2.3.4. The only difference is that its sub-TLV type is set to 9 (TREE-USE-IDs), and the trees listed are those that the originating IS wishes to use as specified in [RFC6325], Section 4.5.

2.3.6. Interested VLANs and Spanning Tree Roots Sub-TLV

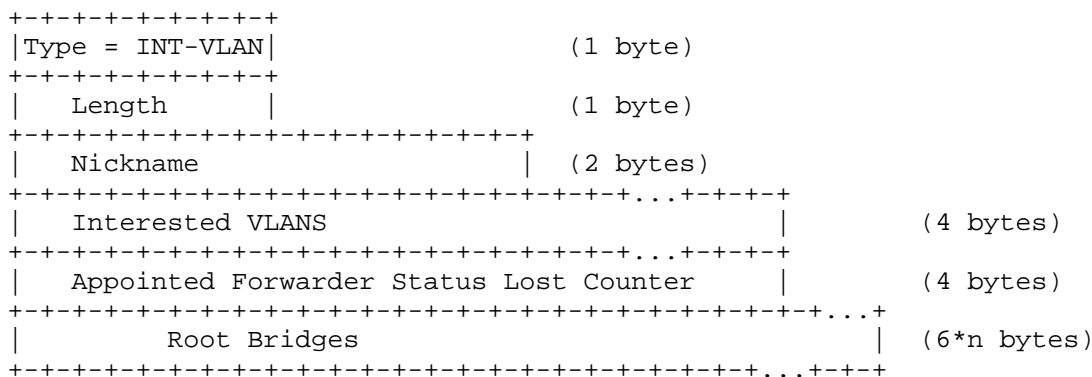
The value of this sub-TLV consists of a VLAN range and information in common to all of the VLANs in the range for the originating IS. This information consists of flags, a variable length list of spanning tree root bridge IDs, and an Appointed Forwarder status lost counter, all as specified in the sections of [RFC6325] listed with the respective information items below.

In the set of LSPs originated by an IS, the union of the VLAN ranges in all occurrences of this sub-TLV MUST be the set of VLANs for which the originating IS is Appointed Forwarder on at least one port, and the VLAN ranges in multiple VLANs sub-TLVs for an IS MUST NOT overlap unless the information provided about a VLAN is the same in every instance. However, as a transient state, these conditions may be violated. If a VLAN is not listed in any INT-VLAN sub-TLV for an IS, that IS is assumed to be uninterested in receiving traffic for that VLAN. If a VLAN appears in more than one INT-VLAN sub-TLV for an IS with different information in the different instances, the following apply:


- If those sub-TLVs provide different nicknames, it is unspecified which nickname takes precedence.
- The largest Appointed Forwarder status lost counter, using serial number arithmetic [RFC1982], is used.
- The originating IS is assumed to be attached to a multicast IPv4 router for that VLAN if any of the INT-VLAN sub-TLVs assert that it is so connected and similarly for IPv6 multicast router attachment.
- The root bridge lists from all of the instances of the VLAN for the originating IS are merged.

To minimize such occurrences, wherever possible, an implementation SHOULD advertise the update to an interested VLAN and Spanning Tree Roots sub-TLV in the same LSP fragment as the advertisement that it replaces. Where this is not possible, the two affected LSP fragments should be flooded as an atomic action. An IS that receives an update to an existing interested VLAN and Spanning Tree Roots sub-TLV can minimize the potential disruption associated with the update by employing a hold-down timer prior to processing the update so as to allow for the receipt of multiple LSP fragments associated with the same update prior to beginning processing.

The sub-TLV layout is as follows:



- o Type: Router Capability and MT-Capability sub-TLV type, set to 10 (INT-VLAN).
- o Length: 10 + 6*n, where n is the number of root bridge IDs.
- o Nickname: As specified in [RFC6325], Section 4.2.4.4, this field may be used to associate a nickname held by the originating IS with the VLAN range indicated. When not used in this way, it is set to zero.
- o Interested VLANs: The Interested VLANs field is formatted as shown below.



- M4, M6: These bits indicate, respectively, that there is an IPv4 or IPv6 multicast router on a link for which the originating IS is Appointed Forwarder for every VLAN in the indicated range as specified in [RFC6325], Section 4.2.4.4, item 5.1.
- R, RESV: These reserved bits MUST be sent as zero and are ignored on receipt.
- VLAN.start and VLAN.end: This VLAN ID range is inclusive. Setting both VLAN.start and VLAN.end to the same value indicates a range of one VLAN ID. If VLAN.start is not equal to VLAN.end and VLAN.start is 0x000, the sub-TLV is interpreted as if VLAN.start was 0x001. If VLAN.start is not equal to

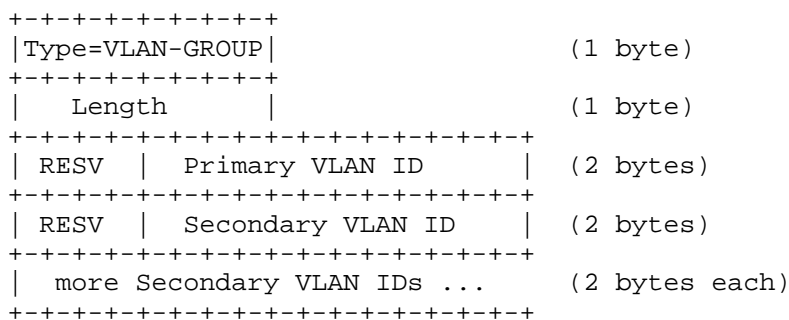
VLAN.end and VLAN.end is 0xFFFF, the sub-TLV is interpreted as if VLAN.end was 0xFFFE. If VLAN.end is less than VLAN.start, the sub-TLV is ignored. If both VLAN.start and VLAN.end are 0x000 or both are 0xFFFF, the sub-TLV is ignored. The values 0x000 or 0xFFFF are not valid VLAN IDs, and a port cannot be enabled for them.

- o Appointed Forwarder Status Lost Counter: This is a count of how many times a port that was Appointed Forwarder for the VLANs in the range given has lost the status of being an Appointed Forwarder for some port as discussed in Section 4.8.3 of [RFC6325]. It is initialized to zero at an IS when the zeroth LSP sequence number is initialized. No special action need be taken at rollover; the counter just wraps around.
- o Root Bridges: The list of zero or more spanning tree root bridge IDs is the set of root bridge IDs seen for all ports for which the IS is Appointed Forwarder for the VLANs in the specified range as discussed in [RFC6325], Section 4.9.3.2. While, of course, at most one spanning tree root could be seen on any particular port, there may be multiple ports in the same VLANs connected to different bridged LANs with different spanning tree roots.

An INT-VLAN sub-TLV asserts that the information provided (multicast router attachment, Appointed Forwarder status lost counter, and root bridges) is the same for all VLANs in the range specified. If this is not the case, the range MUST be split into subranges meeting this criteria. It is always safe to use sub-TLVs with a "range" of one VLAN ID, but this may be too verbose.

2.3.7. VLAN Group Sub-TLV

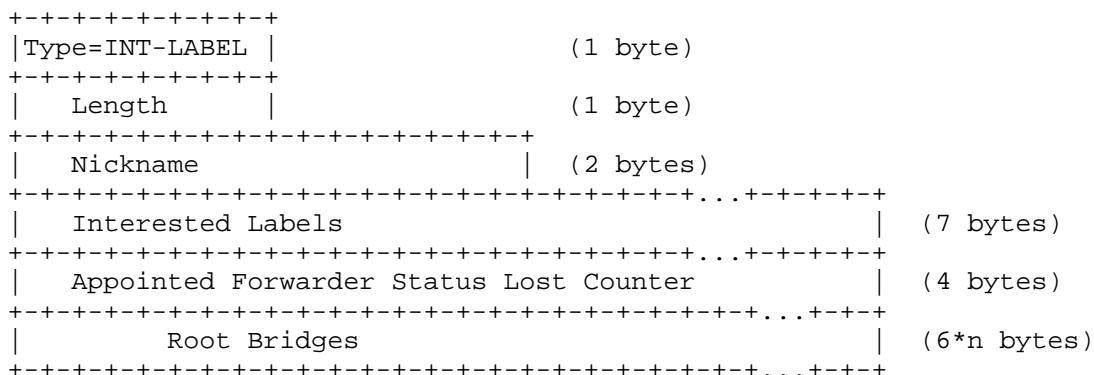
The VLAN Group sub-TLV consists of two or more VLAN IDs as specified in [RFC6325], Section 4.8.4. This sub-TLV indicates that shared VLAN learning is occurring at the originating IS between the listed VLANs. It is structured as follows:



- o Type: Router Capability and MT-Capability sub-TLV type, set to 14 (VLAN-GROUP).
- o Length: 4 + 2*n, where n is the number of secondary VLAN ID fields beyond the first. n MAY be zero.
- o RESV: a 4-bit field that MUST be sent as zero and ignored on receipt.
- o Primary VLAN ID: This identifies the primary VLAN ID.
- o Secondary VLAN ID: This identifies a secondary VLAN in the VLAN Group.
- o more Secondary VLAN IDs: zero or more byte pairs, each with the top 4 bits as a RESV field and the low 12 bits as a VLAN ID.

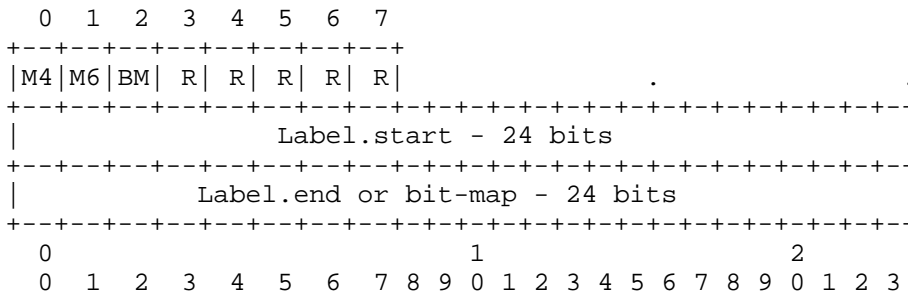
2.3.8. Interested Labels and Spanning Tree Roots Sub-TLV

An IS that can handle fine-grained labeling [RFC7172] announces its fine-grained label connectivity and related information in the Interested Labels and Spanning Tree Roots sub-TLV (INT-LABEL). It is a variation of the Interested VLANs and Spanning Tree Roots sub-TLV (INT-VLAN) and is structured as follows.



- o Type: Router Capability and MT-Capability sub-TLV type, set to 15 (INT-LABEL).
- o Length: 11 + 6*n, where n is the number of root bridge IDs.
- o Nickname: This field may be used to associate a nickname held by the originating IS with the Interested Labels indicated. When not used in this way, it is set to zero.

- o Interested Labels: The Interested Labels field is seven bytes long and formatted as shown below.



- M4, M6: These bits indicate, respectively, that there is an IPv4 or IPv6 multicast router on a link to which the originating IS is Appointed Forwarder for the VLAN corresponding to every label in the indicated range.
- BM: If the BM (bit-map) bit is zero, the last three bytes of the Interested Labels is a Label.end label number. If the BM bit is one, those bytes are a bit-map as described below.
- R: These reserved bits MUST be sent as zero and are ignored on receipt.
- Label.start and Label.end: If the BM bit is zero, this fine-grained label [RFC7172] ID range is inclusive. These fields are treated as unsigned integers. Setting them both to the same label ID value indicates a range of one label ID. If Label.end is less than Label.start, the sub-TLV is ignored.
- Label.start and bit-map: If the BM bit is one, the fine-grained labels that the IS is interested in are indicated by a 24-bit bit-map. The interested labels are the Label.start number plus the bit number of each one bit in the bit-map. So, if bit zero of the bit-map is a one, the IS is interested in the label with value Label.start, and if bit 23 of the bit-map is a one, the IS is interested in the label with value Label.start+23.

- o Appointed Forwarder Status Lost Counter: This is a count of how many times a port that was Appointed Forwarder for a VLAN mapping to the fine-grained label in the range or bit-map given has lost the status of being an Appointed Forwarder as discussed in Section 4.8.3 of [RFC6325]. It is initialized to zero at an IS when the zeroth LSP sequence number is initialized. No special action need be taken at rollover; the counter just wraps around.

- o Root Bridges: The list of zero or more spanning tree root bridge IDs is the set of root bridge IDs seen for all ports for which the IS is Appointed Forwarder for a VLAN mapping to the fine-grained label in the specified range or bit-map. (See [RFC6325], Section 4.9.3.2.) While, of course, at most one spanning tree root could be seen on any particular port, there may be multiple relevant ports connected to different bridged LANs with different spanning tree roots.

An INT-LABEL sub-TLV asserts that the information provided (multicast router attachment, Appointed Forwarder status lost counter, and root bridges) is the same for all labels specified. If this is not the case, the sub-TLV MUST be split into subranges and/or separate bit maps meeting this criteria. It is always safe to use sub-TLVs with a "range" of one VLAN ID, but this may be too verbose.

2.3.9. RBridge Channel Protocols Sub-TLV

An IS announces the RBridge Channel protocols [RFC7178] it supports through use of this sub-TLV.

```

+++++
|Type=RBCHANNELS|                               (1 byte)
+++++
| Length |                                       (1 byte)
+++++-----+
| Zero or more bit vectors | (variable)
+++++

```

- o Type: Router Capability and MT-Capability RBridge Channel Protocols sub-TLV, set to 16 (RBCHANNELS).
- o Length: variable.
- o Bit Vectors: Zero or more byte-aligned bit vectors where a one bit indicates support of a particular RBridge Channel protocol. Each byte-aligned bit vector is formatted as follows:

```

| 0 1 2 3 4 5 6 7| 8 9 10 11 12 13 14 15|
+++++-----+
| Bit Vector Length | Bit Vector Offset |
+++++-----+
| bits
+++++

```

The Bit Vector Length (BVL) is a seven-bit unsigned integer field giving the number of bytes of bit vector. The Bit Vector Offset (BVO) is a nine-bit unsigned integer field.

The bits in each bit vector are numbered in network order, the high-order bit of the first byte of bits being bit $0 + 8*BVO$, the low-order bit of that byte being $7 + 8*BVO$, the high order bit of the second byte being $8 + 8*BVO$, and so on for BVL bytes. A bit vector of RBridge Channel protocols supported MUST NOT extend beyond the end of the value in the sub-TLV in which it occurs. If it does, it is ignored. If multiple byte-aligned bit vectors are present in one such sub-TLV, their representations are contiguous, the BVL field for the next starting immediately after the last byte of bits for the previous bit vector. The one or more bit vectors present MUST exactly fill the sub-TLV value. If there are one or two bytes of value left over, they are ignored; if more than two, an attempt is made to parse them as one or more bit vectors.

If different bit vectors overlap in the protocol number space they refer to and they have inconsistent bit values for a channel protocol, support for the protocol is assumed if any of these bit vectors has a 1 for that protocol.

The absence of any occurrences of this sub-TLV in the LSP for an IS implies that the IS does not support the RBridge Channel facility. To avoid wasted space, trailing bit vector zero bytes SHOULD be eliminated by reducing BVL, any null bit vectors (ones with BVL equal to zero) eliminated, and generally the most compact encoding used. For example, support for channel protocols 1 and 32 could be encoded as

```
BVL = 5
BVO = 0
 0b01000000
 0b00000000
 0b00000000
 0b00000000
 0b10000000
```

or as

```
BVL = 1
BVO = 0
 0b01000000
BLV = 1
BVO = 4
 0b1000000
```

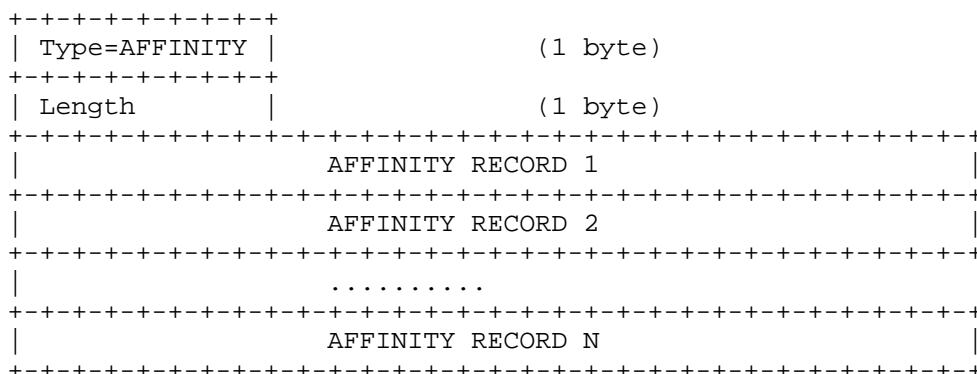
The first takes 7 bytes while the second takes only 6; thus, the second would be preferred.

In multi-topology-aware RBridges, RBridge Channel protocols for which support is announced in the base topology are assumed to be supported in all topologies for which there is no separate announcement for RBridge Channel protocol support.

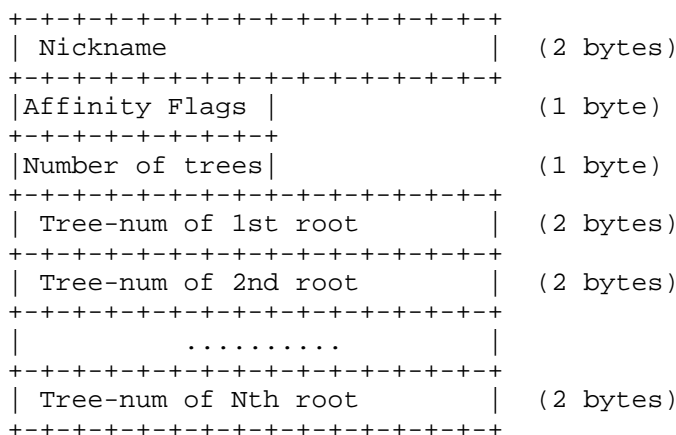
2.3.10. Affinity Sub-TLV

Association of an IS to a multi-destination distribution tree through a specific path is accomplished by using the Affinity sub-TLV. The announcement of an Affinity sub-TLV by RB1 with the nickname of RB2 as the first part of an Affinity Record in the sub-TLV value is a request by RB1 that all ISs in the campus connect RB2 as a child of RB1 when calculating any of the trees listed in that Affinity Record. Examples of use include [Affinity] and [Resilient].

The structure of the Affinity sub-TLV is shown below.



where each AFFINITY RECORD is structured as follows:




- o Type: Router Capability and MT-Capability sub-TLV type, set to 17 (AFFINITY).
- o Length: size of all Affinity Records included, where an Affinity Record listing n tree roots is 4+2*n bytes long.
- o Nickname: 16-bit nickname of the IS whose associations to the multi-destination trees listed in the Affinity Record are through the originating IS.
- o Affinity Flags: 8 bits reserved for future needs to provide additional information about the affinity being announced. MUST be sent as zero and ignored on receipt.
- o Number of trees: A one-byte unsigned integer giving the number of trees for which affinity is being announced by this Affinity Record.
- o Tree-num of roots: The tree numbers of the distribution trees this Affinity Record is announcing.

There is no need for a field giving the number of Affinity Records as this can be determined by processing those records.

2.3.11 Label Group Sub-TLV

The Label Group sub-TLV consists of two or more fine-grained label [RFC7172] IDs. This sub-TLV indicates that shared label MAC address learning is occurring at the announcing IS between the listed labels. It is structured as follows:




- o Type: Router Capability and MT-Capability sub-TLV type, set to 18 (LABEL-GROUP).
- o Length: 6 + 3*n, where n is the number of secondary VLAN ID fields beyond the first. n MAY be zero.

- o Primary Label ID: This identifies the primary Label ID.
- o Secondary Label ID: This identifies a secondary Label ID in the Label Group.
- o more Secondary Label IDs: zero or more byte triples, each with a Label ID.

2.4. MTU Sub-TLV for Extended Reachability and MT-ISN TLVs

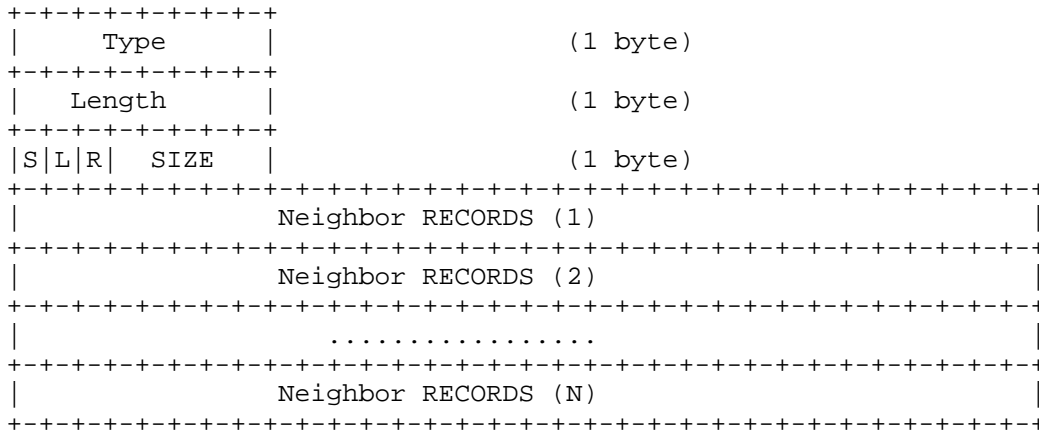
The MTU sub-TLV is used to optionally announce the MTU of a link as specified in [RFC6325], Section 4.2.4.4. It occurs within the Extended Reachability (#22) and MT-ISN (Intermediate System Neighbors) (#222) TLVs.




- o Type: Extended Reachability and MT-ISN sub-TLV type, set to MTU sub-TLV 28.
- o Length: 3.
- o F: Failed. This bit is a one if MTU testing failed on this link at the required campus-wide MTU.
- o RESV: 7 bits that MUST be sent as zero and ignored on receipt.
- o MTU: This field is set to the largest successfully tested MTU size for this link or zero if it has not been tested, as specified in Section 4.3.2 of [RFC6325].

2.5. TRILL Neighbor TLV

The TRILL Neighbor TLV is used in TRILL broadcast link IIH PDUs (see Section 4.1 below) in place of the IS Neighbor TLV, as specified in Section 4.4.2.1 of [RFC6325] and in [RFC7177]. The structure of the TRILL Neighbor TLV is as follows:



The information present for each neighbor is as follows:



- o Type: TLV type, set to TRILL Neighbor TLV 145.
- o Length: 1 + (SIZE+3)*n, where n is the number of neighbor records, which may be zero.
- o S: Smallest flag. If this bit is a one, then the list of neighbors includes the neighbor with the smallest MAC address considered as an unsigned integer.
- o L: Largest flag. If this bit is a one, then the list of neighbors includes the neighbor with the largest MAC address considered as an unsigned integer.
- o R, RESV: These bits are reserved and MUST be sent as zero and ignored on receipt.
- o SIZE: The SNPA size as an unsigned integer in bytes except that 6 is encoded as zero. An actual size of zero is meaningless and cannot be encoded. The meaning of the value 6 in this field is reserved, and TRILL Neighbor TLVs received with a SIZE of 6 are ignored. The SIZE is inherent to the technology of a link and is fixed for all TRILL Neighbor TLVs on that link but may vary

between different links in the campus if those links are different technologies, for example, 6 for EUI-48 SNPAs or 8 for EUI-64 SNPAs [RFC7042]. (The SNPA size on the various links in a TRILL campus is independent of the System ID size.)

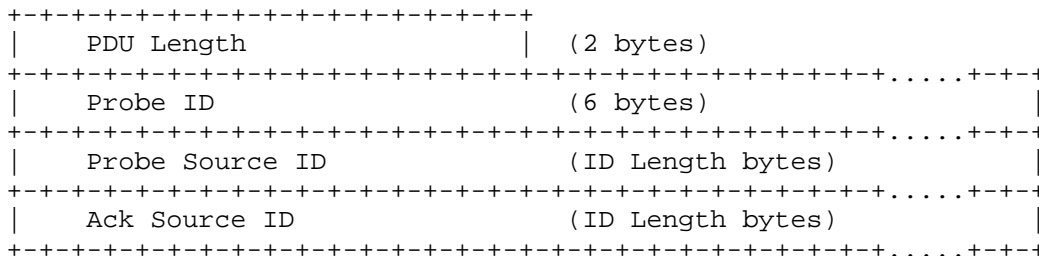
- o F: Failed. This bit is a one if MTU testing to this neighbor failed at the required campus-wide MTU (see [RFC6325], Section 4.3.1).
- o O: OOMF. This bit is a one if the IS sending the enclosing TRILL Neighbor TLV is willing to offer the Overload Originated Multi-destination Frame (OOMF) service [RFC7180] to the IS whose port has the SNPA in the enclosing Neighbor RECORD.
- o MTU: This field is set to the largest successfully tested MTU size for this neighbor or to zero if it has not been tested.
- o SNPA (MAC Address): Subnetwork Point of Attachment of the neighbor.

As specified in [RFC7177] and Section 4.4.2.1 of [RFC6325], all MAC addresses may fit into one TLV, in which case both the S and L flags would be set to one in that TLV. If the MAC addresses don't fit into one TLV, the highest MAC address in a TRILL Neighbor TLV with the L flag zero MUST also appear as a MAC address in some other TRILL Neighbor TLV (possibly in a different TRILL IIH PDU). Also, the lowest MAC address in a TRILL Neighbor TLV with the S flag zero MUST also appear in some other TRILL Neighbor TLV (possibly in a different TRILL IIH PDU). If an IS believes it has no neighbors, it MUST send a TRILL Neighbor TLV with an empty list of neighbor RECORDS, which will have both the S and L bits on.

3. MTU PDUs

The IS-IS MTU-probe and MTU-ack PDUs are used to optionally determine the MTU on a link between ISs as specified in Section 4.3.2 of [RFC6325] and in [RFC7177].

The MTU PDUs have the IS-IS PDU common header (up through the Maximum Area Addresses byte) with PDU Type numbers as indicated in Section 5. They also have a common fixed MTU PDU header as shown below that is $8 + 2 * (\text{ID Length})$ bytes long, 20 bytes in the case of the usual 6-bytes System IDs.



As with other IS-IS PDUs, the PDU Length gives the length of the entire IS-IS packet starting with and including the IS-IS common header.

The Probe ID field is an opaque 48-bit quantity set by the IS issuing an MTU-probe and copied by the responding IS into the corresponding MTU-ack. For example, an IS creating an MTU-probe could compose this quantity from a port identifier and probe sequence number relative to that port.

The Probe Source ID is set by an IS issuing an MTU-probe to its System ID and copied by the responding IS into the corresponding MTU-ack. The Ack Source ID is set to zero in MTU-probe PDUs and ignored on receipt. An IS issuing an MTU-ack sets the Ack Source ID field to its System ID. The System ID length is usually 6 bytes but could be a different value as indicated by the ID Length field in the IS-IS PDU Header.

The TLV area follows the MTU PDU header area. This area MAY contain an Authentication TLV and MUST be padded with the Padding TLV to the exact size being tested. Since the minimum size of the Padding TLV is 2 bytes, it would be impossible to pad to exact size if the total length of the required information-bearing fixed fields and TLVs added up to 1 byte less than the desired length. However, the length of the fixed fields and substantive TLVs for MTU PDUs is expected to be quite small compared with their minimum length (minimum 1470-byte MTU on an IEEE 802.3 link, for example), so this should not be a problem.

4. Use of Existing PDUs and TLVs

The sub-sections below provide details of TRILL use of existing PDUs and TLVs.

4.1. TRILL IIH PDUs

The TRILL IIH PDU is the variation of the IIH PDU used by the TRILL protocol. Section 4.4 of the TRILL standard [RFC6325] and [RFC7177] specify the contents of the TRILL IIH and how its use in TRILL differs from Layer 3 LAN IIH PDU use. The adjacency state machinery for TRILL neighbors is specified in detail in [RFC7177].

In a TRILL IIH PDU, the IS-IS common header and the fixed PDU Header are the same as a Level 1 IIH PDU.

The IS-IS Neighbor TLV (6) is not used in a TRILL IIH and is ignored if it appears there. Instead, TRILL LAN IIH PDUs use the TRILL Neighbor TLV (see Section 2.5).

4.2. Area Address

TRILL uses a fixed zero Area Address as specified in [RFC6325], Section 4.2.3. This is encoded in a 4-byte Area Address TLV (TLV #1) as follows:

```

+++++
| 0x01, Area Address Type | (1 byte)
+++++
| 0x02, Length of Value | (1 byte)
+++++
| 0x01, Length of Address | (1 byte)
+++++
| 0x00, zero Area Address | (1 byte)
+++++

```

4.3. Protocols Supported

NLPID (Network Layer Protocol ID) 0xC0 has been assigned to TRILL [RFC6328]. A Protocols Supported TLV (#129, [RFC1195]) including that value appears in TRILL IIH PDUs and LSP number zero PDUs.

4.4. Link State PDUs (LSPs)

An LSP number zero MUST NOT be originated larger than 1470 bytes, but a larger LSP number zero successfully received MUST be processed and forwarded normally.

4.5. Originating LSP Buffer Size

The originatingLSPBufferSize TLV (#14) MUST be in LSP number zero; however, if found in other LSP fragments, it is processed normally. Should there be more than one originatingLSPBufferSize TLV for an IS, the minimum size, but not less than 1470, is used.

5. IANA Considerations

This section gives IANA considerations for the TLVs, sub-TLVs, and PDUs specified herein. A number of new code points are assigned, and those that were assigned by [RFC6326] are included here for convenience. IANA has replaced all [RFC6326] references in the IANA registries with references to this document.

5.1. TLVs

This document specifies two IS-IS TLV types -- namely, the Group Address TLV (GADDR-TLV; type 142) and the TRILL Neighbor TLV (type 145). The PDUs in which these TLVs are permitted for TRILL are shown in the table below along with the section of this document where they are discussed. The final "NUMBER" column indicates the permitted number of occurrences of the TLV in their PDU, or set of PDUs in the case of LSPs, which in these two cases is "*" indicating that the TLV MAY occur 0, 1, or more times.

IANA has registered these two code points in the IANA IS-IS TLV registry (ignoring the "Section" and "NUMBER" columns, which are irrelevant to that registry).

| | Section | TLV | IIH | LSP | SNP | Purge | NUMBER |
|--------------------|---------|-----|-----|-----|-----|-------|--------|
| | ===== | === | === | === | === | ===== | ===== |
| GADDR-TLV | 2.1 | 142 | n | y | n | n | * |
| TRILL Neighbor TLV | 2.5 | 145 | y | n | n | n | * |

5.2. Sub-TLVs

This document specifies a number of sub-TLVs. The TLVs in which these sub-TLVs occur are shown in the second table below along with the section of this document where they are discussed. The TLVs within which these sub-TLVs can occur are determined by the presence of an "X" in the relevant column; the column headers are described in the first table below. In some cases, the column header corresponds to two different TLVs in which the sub-TLV can occur.

| Column Head | TLV | RFC | TLV Name |
|-------------|-------|-------|--------------------------|
| ===== | ===== | ===== | ===== |
| Grp. Adr. | 142 | 7176 | Group Address |
| MT Port | 143 | 6165 | MT-Port-Cap-TLV |
| MT Cap. | 242 | 4971 | Router CAPABILITY |
| | 144 | 6329 | MT-Capability |
| Ext. Reach | 22 | 5305 | Extended IS Reachability |
| | 222 | 5120 | MT-ISN |

The final "NUMBER" column below indicates the permitted number of occurrences of the sub-TLV cumulatively within all occurrences of their TLV(s) in those TLVs' carrying a PDU (or set of PDUs in the case of LSPs), as follows:

- 0-1 = MAY occur zero or one times.
- 1 = MUST occur exactly once. If absent, the PDU is ignored. If it occurs more than once, results are unspecified.
- * = MAY occur 0, 1, or more times.

The values in the "Section" and "NUMBER" columns are irrelevant to the IANA sub-registries.

| Name | Section | sub-TLV# | Grp. Adr. | MT Port | MT Cap. | Ext. Reach | NUMBER |
|-----------------|---------|----------|-----------|---------|---------|------------|--------|
| ===== | ===== | ===== | ===== | ===== | ===== | ===== | ===== |
| GMAC-ADDR | 2.1.1 | 1 | X | - | - | - | * |
| GIP-ADDR | 2.1.2 | 2 | X | - | - | - | * |
| GIPV6-ADDR | 2.1.3 | 3 | X | - | - | - | * |
| GLMAC-ADDR | 2.1.4 | 4 | X | - | - | - | * |
| GLIP-ADDR | 2.1.5 | 5 | X | - | - | - | * |
| GLIPV6-ADDR | 2.1.6 | 6 | X | - | - | - | * |
| VLAN-FLAGS | 2.2.1 | 1 | - | X | - | - | 1 |
| Enabled-VLANs | 2.2.2 | 2 | - | X | - | - | * |
| AppointedFwrdrs | 2.2.3 | 3 | - | X | - | - | * |
| PORT-TRILL-VER | 2.2.4 | 7 | - | X | - | - | 0-1 |
| VLANs-Appointed | 2.2.5 | 8 | - | X | - | - | * |
| NICKNAME | 2.3.2 | 6 | - | - | X | - | * |
| TREES | 2.3.3 | 7 | - | - | X | - | 0-1 |
| TREE-RT-IDs | 2.3.4 | 8 | - | - | X | - | * |
| TREE-USE-IDs | 2.3.5 | 9 | - | - | X | - | * |
| INT-VLAN | 2.3.6 | 10 | - | - | X | - | * |
| TRILL-VER | 2.3.1 | 13 | - | - | X | - | 0-1 |
| VLAN-GROUP | 2.3.7 | 14 | - | - | X | - | * |
| INT-LABEL | 2.3.8 | 15 | - | - | X | - | * |
| RBCHANNELS | 2.3.9 | 16 | - | - | X | - | * |

| | | | | | | | |
|-------------|---------|----------|-----------|---------|---------|------------|--------|
| AFFINITY | 2.3.10 | 17 | - | - | X | - | * |
| LABEL-GROUP | 2.3.11 | 18 | - | - | X | - | * |
| MTU | 2.4 | 28 | - | - | - | X | 0-1 |
| ===== | | | | | | | |
| Name | Section | sub-TLV# | Grp. Adr. | MT Port | MT Cap. | Ext. Reach | NUMBER |

IANA has entered the newly assigned sub-TLV numbers in the above table in the relevant existing sub-TLV registries, as determined by which column has an X for that sub-TLV. For the sub-TLVs from NICKNAME through and including VLAN-GROUP, which previously existed only in the registry of sub-TLVs under TLV 242, IANA has added each sub-TLV with the same sub-TLV number to the existing registry for sub-TLVs under TLV 144.

5.3. PDUs

The IS-IS PDUs registry remains as established in [RFC6326] except that the references to [RFC6326] are updated to reference this document.

5.4. Reserved and Capability Bits

Any reserved bits (R), bits in reserved fields (RESV), or capabilities bits in the PORT-TRILL-VER and TRILL-VER sub-TLVs, which are specified herein as "MUST be sent as zero and ignored on receipt" or the like, are allocated based on IETF Review [RFC5226].

Two sub-registries have been created within the TRILL Parameters Registry as follows:

Sub-Registry Name: TRILL-VER Sub-TLV Capability Flags
 Registration Procedures: IETF Review
 Reference: (This document)

| Bit | Description | Reference |
|-------|-------------------------------|------------|
| ===== | ===== | ===== |
| 0 | Affinity sub-TLV support. | [Affinity] |
| 1 | FGL-safe | [RFC7172] |
| 2-13 | Unassigned | |
| 14-31 | Extended header flag support. | [RFC7179] |

Sub-Registry Name: PORT-TRILL-VER Sub-TLV Capability Flags
 Registration Procedures: IETF Review
 Reference: (This document)

| Bit | Description | Reference |
|-------|-----------------------------------|-----------|
| 0 | Hello reduction support. | [RFC7180] |
| 1-2 | Unassigned | |
| 3-13 | Hop-by-hop extended flag support. | [RFC7179] |
| 14-31 | Unassigned | |

5.5. TRILL Neighbor Record Flags

A sub-registry has been created within the TRILL Parameters Registry as follows:

Sub-Registry Name: TRILL Neighbor TLV NEIGHBOR RECORD Flags
 Registration Procedures: Standards Action
 Reference: (This document)

| Bit | Short Name | Description | Reference |
|-----|------------|-----------------------|-----------------------------|
| 0 | Fail | Failed MTU test | [RFC6325][RFC7176][RFC7177] |
| 1 | OOMF | Offering OOMF service | [RFC7180] |
| 2-7 | - | Unassigned | |

6. Security Considerations

For general TRILL protocol security considerations, see the TRILL base protocol standard [RFC6325].

This document raises no new security issues for IS-IS. IS-IS security may be used to secure the IS-IS messages discussed here. See [RFC5304] and [RFC5310]. Even when IS-IS authentication is used, replays of Hello packets can create denial-of-service conditions; see [RFC6039] for details. These issues are similar in scope to those discussed in Section 6.2 of [RFC6325], and the same mitigations may apply.

7. Changes from RFC 6326

Non-editorial changes from [RFC6326] are summarized in the list below:

1. Added five sub-TLVs under the Group Address (GADDR) TLV covering VLAN-labeled IPv4 and IPv6 addresses and fine-grained-labeled MAC, IPv4, and IPv6 addresses (Sections 2.1.2, 2.1.3, 2.1.4, 2.1.5, and 2.1.6).

2. Added the PORT-TRILL-VER sub-TLV (Section 2.2.4).
3. Added the VLANs-Appointed sub-TLV (Section 2.2.5).
4. Changed the TRILL-VER sub-TLV as listed below.
 - a. Added 4 bytes of TRILL Header extended flags and capabilities supported information.
 - b. Required that the TRILL-VER sub-TLV appear in LSP number zero.

The above changes to TRILL-VER are backward compatible because the [RFC6326]-conformant implementations of TRILL thus far have only supported version zero and not supported any optional capabilities or extended flags, the level of support indicated by the absence of the TRILL-VER sub-TLV. Thus, if an [RFC6326]-conformant implementation of TRILL rejects this sub-TLV due to the changes specified in this document, it will, at worst, decide that support of version zero and no extended flags or capabilities is indicated, which is the best an [RFC6326]-conformant implementation of TRILL can do anyway. Similarly, a TRILL implementation that supports TRILL-VER as specified herein and rejects TRILL-VER sub-TLVs in an [RFC6326]-conformant TRILL implementation because they are not in LSP number zero will decide that the implementation supports only version zero with no extended flag or capabilities support, which will be correct (Section 2.3.1).

5. Clarified the use of invalid VLAN IDs (0x000 and 0xFFFF) in the Appointed Forwarders sub-TLV and the Interested VLANs and Spanning Tree Roots sub-TLV (Sections 2.2.3 and 2.3.6).
6. Added the Interested Labels and Spanning Tree Roots sub-TLV to indicate attachment of an IS to a fine-grained label [RFC7172] analogous to the existing Interested VLANs and Spanning Tree Roots sub-TLV for VLANs (Section 2.3.8).
7. Added the RBridge Channel Protocols sub-TLV so ISs can announce the RBridge Channel protocols they support (Section 2.3.9).
8. Permitted specification of the length of the link SNPA field in TRILL Neighbor TLVs. This change is backward compatible because the size of 6 bytes is specially encoded as zero, the previous value of the bits in the new SIZE field (Section 2.5).

9. Made the size of the MTU PDU Header Probe Source ID and Ack Source ID fields be the ID Length from the IS-IS PDU Header rather than the fixed value 6 (Section 3).
10. For robustness, required that LSP number zero PDUs be originated as no larger than 1470 bytes but processed regardless of size (Section 4.4).
11. Required that the originatingLSPBufferSize TLV, if present, appear in LSP number zero (Section 4.5).
12. Created sub-registries for and specified the IANA Considerations policy for reserved and capability bits in the TRILL version sub-TLVs (Section 5.4).
13. Added the distribution tree Affinity sub-TLV so ISs can request distribution tree attachments (Section 2.3.10).
14. Added the LABEL-GROUP sub-TLV analogous to the VLAN-GROUP sub-TLV (Section 2.3.11).
15. Added multi-topology to permit sub-TLVs previously only in the Router Capability TLV to also appear in the MT-Capability TLV and to permit the MTU sub-TLV previously limited to the Extended Reachability TLV to also appear in the MT-ISN TLV.
16. Added a sub-registry for Neighbor TLV Neighbor RECORD flag bits (Section 5.5).
17. Explicitly stated that if the number of sources in a GADDR-TLV sub-TLV is zero, it indicates a listener for (*,G), that is, a listener not restricted by source (Section 2.1).

8. References

8.1. Normative References

[ISO-10589]

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)", Second
Edition, November 2002.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and
dual environments", RFC 1195, December 1990.

- [RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4971] Vasseur, JP., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", RFC 4971, July 2007.
- [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.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 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.
- [RFC6328] Eastlake 3rd, D., "IANA Considerations for Network Layer Protocol Identifiers", BCP 164, RFC 6328, July 2011.
- [RFC6329] Fedyk, D., Ed., Ashwood-Smith, P., Ed., Allan, D., Bragg, A., and P. Unbehagen, "IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging", RFC 6329, April 2012.
- [RFC6439] Perlman, R., Eastlake, D., Li, Y., Banerjee, A., and F. Hu, "Routing Bridges (RBridges): Appointed Forwarders", RFC 6439, November 2011.
- [RFC7172] Eastlake 3rd, D., Zhang, M., Agarwal, P., Perlman, R., and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL): Fine-Grained Labeling", RFC 7172, May 2014.
- [RFC7177] Eastlake 3rd, D., Perlman, R., Ghanwani, A., Yang, Y., and V. Manral, "Transparent Interconnection of Lots of Links (TRILL): Adjacency", RFC 7177, May 2014.

- [RFC7178] Eastlake 3rd, D., Manral, V., Li, Y., Aldrin, S., and D. Ward, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Support", RFC 7178, May 2014.
- [RFC7179] Eastlake 3rd, D., Ghanwani, A., Manral, V., Li, Y., and C. Bestler, "Transparent Interconnection of Lots of Links (TRILL): Header Extension", RFC 7179, May 2014.
- [RFC7180] Eastlake 3rd, D., Zhang, M., Ghanwani, A., Manral, V., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", RFC 7180, May 2014.

8.2. Informative References

- [Err2869] RFC Errata, Errata ID 2869, RFC 6326, <<http://www.rfc-editor.org>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, October 2008.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, February 2009.
- [RFC6039] Manral, V., Bhatia, M., Jaeggli, J., and R. White, "Issues with Existing Cryptographic Protection Methods for Routing Protocols", RFC 6039, October 2010.
- [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.
- [RFC7042] Eastlake 3rd, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", BCP 141, RFC 7042, October 2013.
- [RFC7175] Manral, V., Eastlake 3rd, D., Ward, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Bidirectional Forwarding Detection (BFD) Support", RFC 7175, May 2014.
- [Affinity] Senevirathne, T., Pathangi, J., and J. Hudson, "Coordinated Multicast Trees (CMT) for TRILL", Work in Progress, April 2014.

[MultiLevel]

Perlman, R., Eastlake 3rd, D., Ghanwani, A., and H. Zhai, "Flexible Multilevel TRILL (Transparent Interconnection of Lots of Links)", Work in Progress, January 2014.

[Resilient]

Zhang, M. Senevirathne, T., Pathangi, J., Banerjee, A., and A. Ghanwani, "TRILL Resilient Distribution Trees", Work in Progress, December 2013.

9. Acknowledgements

The authors gratefully acknowledge the contributions and reviews by the following individuals: Ross Callon, Spencer Dawkins, Adrian Farrel, Alexey Melnikov, Radia Perlman, Carlos Pignataro, and Joe Touch.

The authors also acknowledge the contributions to [RFC6326] by the following individuals: Mike Shand, Stewart Bryant, Dino Farinacci, Les Ginsberg, Sam Hartman, Dan Romascanu, Dave Ward, and Russ White. In particular, thanks to Mike Shand for his detailed and helpful comments.

Authors' Addresses

Donald Eastlake 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757
USA

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

Tissa Senevirathne
Cisco Systems
375 East Tasman Drive,
San Jose, CA 95134
USA

Phone: +1-408-853-2291
EMail: tsenevir@cisco.com

Anoop Ghanwani
Dell
5450 Great America Parkway
Santa Clara, CA 95054
USA

EMail: anoop@alumni.duke.edu

Dinesh Dutt
Cumulus Networks
1089 West Evelyn Avenue
Sunnyvale, CA 94086
USA

EMail: ddutt.ietf@hobbesdutt.com

Ayan Banerjee
Insieme Networks
210 West Tasman Drive
San Jose, CA 95134
USA

EMail: ayabaner@gmail.com