

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

D. Eastlake
Huawei
A. Banerjee
D. Dutt
Cisco
R. Perlman
Intel
A. Ghanwani
Brocade
July 2011

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

Abstract

The IETF has standardized the Transparent Interconnection of Lots of Links (TRILL) protocol, which provides transparent Layer 2 forwarding using encapsulation with a hop count and IS-IS link state routing. This document specifies the data formats and code points for the IS-IS extensions to support TRILL.

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/rfc6326>.

Copyright Notice

Copyright (c) 2011 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	3
2. TLV and Sub-TLV Extensions to IS-IS for TRILL	3
2.1. The Group Address TLV	3
2.1.1. The Group MAC Address Sub-TLV	4
2.2. Multi-Topology-Aware Port Capability Sub-TLVs	5
2.2.1. The Special VLANs and Flags Sub-TLV	6
2.2.2. Enabled-VLANs Sub-TLV	7
2.2.3. Appointed Forwarders Sub-TLV	8
2.3. Sub-TLVs for the Router Capability TLV	9
2.3.1. The TRILL Version Sub-TLV	9
2.3.2. The Nickname Sub-TLV	10
2.3.3. The Trees Sub-TLV	11
2.3.4. The Tree Identifiers Sub-TLV	11
2.3.5. The Trees Used Identifiers Sub-TLV	12
2.3.6. Interested VLANs and Spanning Tree Roots Sub-TLV ...	12
2.3.7. The VLAN Group Sub-TLV	15
2.4. MTU Sub-TLV of the Extended Reachability TLV	15
2.5. TRILL Neighbor TLV	16
3. The MTU PDUs	18
4. Use of Existing PDUs and TLVs	19
4.1. TRILL IIH PDUs	19
4.2. Area Address	19
4.3. Protocols Supported	19
5. IANA Considerations	20
5.1. Allocations from Existing Registries	20
5.2. New Sub-Registries Created	21
6. Security Considerations	22
7. References	22
7.1. Normative References	22
7.2. Informative References	23
8. Acknowledgements	23
Appendix A. Initial IS-IS PDU Registry	24

1. Introduction

The IETF has standardized the TRILL protocol [RFC6325], which provides transparent Layer 2 forwarding using encapsulation 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 as well as supporting VLANs. Intermediate Systems (ISs) implementing TRILL can incrementally replace IEEE [802.1Q-2005] bridges.

This document, in conjunction with [RFC6165], specifies the data formats and code points for the IS-IS [ISO-10589] [RFC1195] extensions to support TRILL.

1.1. Conventions Used in This Document

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

Additional acronyms used in this document are:

IIH - IS-IS Hello

IS - Intermediate System (for this document, all relevant intermediate systems are R Bridges)

NLPID - Network Layer Protocol Identifier

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 added to IS-IS to support the TRILL standard. 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 provided in Section 5.

2.1. The Group Address TLV

The Group Address (GADDR) TLV, IS-IS TLV type 142, is carried only in an LSP PDU and carries sub-TLVs that in turn advertise multicast group listeners. Section 2.1.1 below specifies a sub-TLV that advertises listeners by MAC address. It is anticipated that

additional sub-TLVs for additional address types such as IP addresses will be specified in other documents. 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. The 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 as specified in Section 4.5.5 of [RFC6325]. It has the following format:

```

+-----+
|Type=GMAC-ADDR | (1 byte)
+-----+
| Length | (1 byte)
+-----+
| RESV | Topology-ID | (2 bytes)
+-----+
| RESV | VLAN ID | (2 bytes)
+-----+
|Num Group Recs | (1 byte)
+-----+
| GROUP RECORDS (1)
+-----+
| .....
+-----+
| GROUP RECORDS (N)
+-----+

```

where each group record is of the form:

```

+-----+
| Num of Sources|                               (1 byte)
+-----+
|               Group Address                (6 bytes)
+-----+
|               Source 1 Address              (6 bytes)
+-----+
|               Source 2 Address              (6 bytes)
+-----+
|               ....
+-----+
|               Source M Address              (6 bytes)
+-----+

```

- o Type: GADDR sub-TLV type, set to 1 (GMAC-ADDR).
- o Length: Variable, minimum 5.
- o RESV: Reserved. 4-bit fields that MUST be sent as zero and ignored on receipt.
- o Topology-ID: This field is not used in TRILL, where it is sent as zero and ignored on receipt, but is included for use by other technologies.
- 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 Number of Group Records: A 1-byte integer that is the number of group records in this sub-TLV.
- o Group Record: Each group record carries the number of sources. It then has a 48-bit multicast address followed by 48-bit 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.

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

TRILL makes use of the Multi-Topology-Aware Port Capability (MT-PORT-CAP) TLV as specified in [RFC6165]. The remainder of this section specifies the sub-TLVs that TRILL uses the MT-PORT-CAP TLV to transport.

2.2.1. The 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 |   Desig.VLAN   | (2 bytes)
+-----+-----+

```

- o Type: sub-TLV type, set to MT-PORT-CAP 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 encapsulated data frame. See [RFC6325], Section 4.6.2, point 8.
- o Outer.VLAN: A copy of the 12-bit outer VLAN ID of the TRILL IIH frame containing this sub-TLV when that frame was sent, as specified in [RFC6325], Section 4.4.5.
- o Desig.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 for end station service at the port of the originating IS on which the Hello was sent, as specified in [RFC6325], Section 4.4.2. It has the following format:

```

+++++
|   Type   |                               (1 byte)
+++++
|  Length  |                               (1 byte)
+++++
| RESV | Start VLAN ID | (2 bytes)
+++++
| VLAN bit-map....
+++++

```

- o Type: sub-TLV type, set to MT-PORT-CAP 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 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 Section 4.2.4 of [RFC6325]. It has the following format:

```

+-----+
|      Type      | (1 byte)
+-----+
|      Length     | (1 byte)
+-----+
| Appointment Information (1) | (6 bytes)
+-----+
| .....          |
+-----+
| Appointment Information (N) | (6 bytes)
+-----+

```

where each appointment is of the form:

```

+-----+
| Appointee Nickname | (2 bytes)
+-----+
| RESV | Start.VLAN | (2 bytes)
+-----+
| RESV | End.VLAN   | (2 bytes)
+-----+

```

- o Type: sub-TLV type, set to MT-PORT-CAP AppointedFwrdrs 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: These fields are the VLAN IDs of the appointment range, inclusive. To specify a single VLAN, the VLAN's ID appears as both the start and end VLAN. As specified in Section 4.4 of [RFC6325], appointing an IS forwarder on a port for a VLAN not enabled on that port has no effect.

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.

2.3. Sub-TLVs for the Router Capability TLV

The Router Capability TLV is specified in [RFC4971]. All of the sub-sections of this Section 2.3 below specify sub-TLVs that can be carried in the Router Capability TLV for TRILL.

2.3.1. The TRILL Version Sub-TLV

The TRILL Version (TRILL-VER) sub-TLV indicates the maximum version of the TRILL standard supported. By implication, lower versions are also supported. If this sub-TLV is missing, the originating IS only supports the base version of the protocol [RFC6325].

```

+-----+
| Type           | (1 byte)
+-----+
| Length         | (1 byte)
+-----+
| Max-version    | (1 byte)
+-----+

```

- o Type: Router Capability sub-TLV type, set to 13 (TRILL-VER).
- o Length: 1.
- o Max-version: Set to maximum version supported.

2.3.2. The 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 as specified in [RFC6325], Section 3.7.3. Multiple instances of this sub-TLV may be carried.

```

+-----+
|Type = NICKNAME|                               (1 byte)
+-----+
|  Length      |                               (1 byte)
+-----+
|               |
|      NICKNAME RECORDS (1)                    |
+-----+
|               |
|      .....                                    |
+-----+
|               |
|      NICKNAME RECORDS (N)                    |
+-----+

```

where each nickname record is of the form:

```

+-----+
| Nickname.Pri |                               (1 byte)
+-----+
|   Tree Root Priority   | (2 byte)
+-----+
|           Nickname     | (2 bytes)
+-----+

```

- o Type: Router 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. The 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:

```

+-----+-----+
|Type =  TREES  |                               (1 byte)
+-----+-----+
|  Length       |                               (1 byte)
+-----+-----+
| Number of trees to compute | (2 byte)
+-----+-----+
| Maximum trees able to compute | (2 byte)
+-----+-----+
| Number of trees to use       | (2 byte)
+-----+-----+

```

- o Type: Router 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. The 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 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:

```

+-----+
|Type=TREE-RT-IDs|          (1 byte)
+-----+
|   Length       |          (1 byte)
+-----+
|Starting Tree Number|      (2 bytes)
+-----+
|   Nickname (K-th root) |    (2 bytes)
+-----+
|   Nickname (K+1 - th root) | (2 bytes)
+-----+
|   Nickname (...)      |
+-----+

```

- o Type: Router 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 a 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. The 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 Router Capability 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 precisely 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 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:

```

+++++
|Type = INT-VLAN|                (1 byte)
+++++
|  Length      |                (1 byte)
+++++
|  Nickname     |                (2 bytes)
+++++
| Interested VLANs |                (4 bytes)
+++++
| Appointed Forwarder Status Lost Counter | (4 bytes)
+++++
|      Root Bridges      | (6*n bytes)
+++++

```

- o Type: Router 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.

0	1	2	3	4 - 15	16 - 19	20 - 31
+	+	+	+	+	+	+
M4	M6	R	R	VLAN.start	RESV	VLAN.end
+	+	+	+	+	+	+

- 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. A range of one VLAN ID is indicated by setting them both to that VLAN ID value.
- 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 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, only one spanning tree root could be seen on any particular port, there may be multiple ports in the same VLAN 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. The VLAN Group Sub-TLV

The VLAN Group Router Capability 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 announcing IS between the listed VLANs. It is structured as follows:

```

+-----+
|Type=VLAN-GROUP|          (1 byte)
+-----+
|  Length      |          (1 byte)
+-----+
| RESV  | Primary VLAN ID | (2 bytes)
+-----+
| RESV  | Secondary VLAN ID | (2 bytes)
+-----+
| more Secondary VLAN IDs ... (2 bytes each)
+-----+

```

- o Type: Router Capability sub-TLV type, set to 14 (VLAN-GROUP).
- o Length: $4 + 2*n$, where n is the number of secondary VLAN ID fields, which 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.

2.4. MTU Sub-TLV of the Extended Reachability TLV

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 TLV (type 22).

```

+-----+
| Type = MTU | (1 byte)
+-----+
| Length | (1 byte)
+-----+
| F | Reserved | (1 byte)
+-----+
| MTU | (2 bytes)
+-----+

```

- o Type: Extended Reachability 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 Reserved: 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 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 [RFC6327]. The structure of the TRILL Neighbor TLV is as follows:

```

+-----+
| Type | (1 byte)
+-----+
| Length | (1 byte)
+-----+
| S | L | RESV | (1 byte)
+-----+
| Neighbor RECORDS (1) |
+-----+
| ..... |
+-----+
| Neighbor RECORDS (N) |
+-----+

```


The information present for each neighbor is as follows:

```

+-----+
|F|  RESV      | (1 bytes)
+-----+
|      MTU      | (2 bytes)
+-----+
|  MAC Address  | (6 bytes)
+-----+

```

- o Type: TLV Type, set to TRILL Neighbor TLV 145.
- o Length: 1 + 9*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 RESV: These 7 bits are reserved use and MUST be sent as zero and ignored on receipt.
- 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 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 MAC Address: The MAC address of the neighbor as in the IS Neighbor TLV (6).

As specified in [RFC6327] 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 RBridge 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. The MTU PDUs

Two PDUs are added to IS-IS, the MTU-probe and MTU-ack PDUs. They are used to optionally determine the MTU on a link between ISs as specified in [RFC6325], Section 4.3.2.

The MTU PDUs have the IS-IS PDU common header (up through the Maximum Area Addresses byte) with two new PDU Type numbers, one each, as listed in Section 6. They also have a 20-byte common fixed MTU PDU header as shown below.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   PDU Length                               |   (2 bytes)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Probe ID                               (6 bytes)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Probe Source ID                         (6 bytes)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Ack Source ID                         (6 bytes)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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 arbitrary 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. An IS issuing an MTU-ack sets this field to its System ID.

The TLV area follows the MTU PDU header area. This area MAY contain an Authentication TLV and MUST be padded to the exact size being tested with the Padding TLV. 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 will be quite small compared with their minimum length (minimum 1470-byte MTU on an 802.3 link, for example), so this will 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 LAN IIH PDU used by the TRILL protocol. Section 4.4 of the TRILL standard [RFC6325] specifies 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 Section 4.4 of [RFC6325] and in [RFC6327].

In a TRILL IIH PDU, the IS-IS common header and the fixed PDU Header are the same as a Level 1 LAN IIH PDU. The Maximum Area Addresses octet in the common header MUST be set to 0x01.

The IS-IS Neighbor TLV (6) is not used in a TRILL IIH and is ignored if it appears there. Instead, TRILL 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 (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 0xC0 has been assigned to TRILL [RFC6328]. A Protocols Supported TLV (129, [RFC1195]) including that value MUST appear in TRILL IIH PDUs and LSP number zero PDUs.

5. IANA Considerations

IANA has allocated the existing registry code points listed in Section 5.1 and created two new registries with the initial contents as described in Section 5.2.

5.1. Allocations from Existing Registries

This document specifies two new 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 LSP, which in these two cases is "*" indicating that the TLV MAY occur 0, 1, or more times.

IANA 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	NUMBER
GADDR-TLV	2.1	142	-	X	-	*
TRILL Neighbor TLV	2.5	145	X	-	-	*

This document specifies eleven new sub-TLVs from existing sub-TLV sequences -- namely, VLAN-FLAGS, Enabled-VLANs, AppointedFwrdrs, TRILL Version (TRILL-VER), NICKNAME, TREES, TREE-RT-IDs, TREE-USE-IDs, INT-VLAN, VLAN-GROUP, and MTU. The TLVs in which these sub-TLVs occur are shown in the table below along with the section of this document where they are discussed.

Those sub-TLVs with an "X" in the column labeled "MT Port Capabil." are sub-TLVs of TLV 143 [RFC6165], the MT-PORT-CAP-TLV. Those sub-TLVs with an "X" in the column labeled "Router Capabil." are sub-TLVs of TLV 242, the IS-IS Router CAPABILITY TLV. Those sub-TLVs with an "X" in the column labeled "Extended IS Reach" are sub-TLVs of TLV 22, the Extended IS reachability TLV.

The final "NUM" column indicates the permitted number of occurrences of the sub-TLV cumulatively within all occurrences of their TLV in that TLV's carrying PDU (or set of PDUs in the case of LSP), as follows:

- 0-1 = MAY occur zero or one times. If it occurs more than once, results are unspecified.
- 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 "NUM" columns are irrelevant to the IANA sub-registries.

	Section	sub-TLV#	MT Port Capabil.	Router Capabil.	Extended IS Reach	NUM
VLAN-FLAGS	2.2.1	1	X	-	-	1
Enabled-VLANs	2.2.2	2	X	-	-	*
AppointedFwrdrs	2.2.3	3	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	-	*
MTU	2.4	28	-	-	X	0-1

5.2. New Sub-Registries Created

This document creates two new IS-IS PDUs -- namely, the MTU-PROBE-PDU and MTU-ACK-PDU, as described in Section 3. IANA assigned new PDU types to these PDUs and reflect them in a newly created PDU registry (see Appendix A).

MTU-PROBE-PDU PDU Number: 23
 MTU-ACK-PDU PDU Number: 28

IANA created a new sub-TLV IS-IS sub-registry for sub-TLVs within the Group Address (GADDR) TLV and specified an initial sub-TLV within that registry -- namely, the Group MAC Address (GMAC-ADDR) sub-TLV (1). The GMAC-ADDR sub-TLV may occur 0, 1, or more times in a GADDR TLV.

The initial sub-registry is shown below.

Registry Name: IS-IS Group Address Type Codes for TLV 10
 Reference: This document
 Registration Procedures: Expert Review [RFC5226]

Registry:					
Value	Group	Address	Type	Code	Reference
-----	-----				-----
0	Reserved				This document
1	GMAC-ADDR				This document
2-254	Unassigned				This document
255	Reserved				This document

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. References

7.1. 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.
- [RFC4971] Vasseur, JP. and N. Shen, "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", 2007.
- [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", 2008.

- [RFC6165] Banerjee, A. and D. Ward, "Extensions to IS-IS for Layer-2 Systems", RFC 6165, April 2011.
- [RFC6325] Perlman, R., Eastlake, D., Dutt, D., Gai, S., and A. Ghanwani, "RBridges: Base Protocol Specification", RFC 6325, July 2011.
- [RFC6327] Eastlake, D., Perlman, R., Ghanwani, A., Dutt, D., and V. Manral, "RBridges: Adjacency", RFC 6327, July 2011.
- [RFC6328] Eastlake, D., "IANA Considerations for Network Layer Protocol Identifiers", RFC 6328, July 2011.

7.2. Informative References

- [802.1Q-2005] "IEEE Standard for Local and metropolitan area networks / Virtual Bridged Local Area Networks", 802.1Q-2005, 19 May 2006.
- [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.

8. Acknowledgements

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

Appendix A. Initial IS-IS PDU Registry

The following is the suggested initial IS-IS PDU registry before MTU-PROBE-PDU and MTU-ACK-PDU, which should be added with this document as REFERENCE:

Registry Name: IS-IS PDUs

Reference: This document

Registration Procedures: IETF Review [RFC5226]

MNEMONIC	PDU#	REFERENCE
Unassigned	0-14	
L1-LAN-HELLO-PDU	15	[ISO-10589]
L2-LAN-HELLO-PDU	16	[ISO-10589]
P2P-HELLO-PDU	17	[ISO-10589]
L1-LSP-PDU	18	[ISO-10589]
Unassigned	19	
L2-LSP-PDU	20	[ISO-10589]
Unassigned	21-23	
L1-CSNP-PDU	24	[ISO-10589]
L2-CSNP-PDU	25	[ISO-10589]
L1-PSNP-PDU	26	[ISO-10589]
L2-PSNP-PDU	27	[ISO-10589]
Unassigned	28-31	

Authors' Addresses

Donald Eastlake
Huawei
155 Beaver Street
Milford, MA 01757 USA

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

Ayan Banerjee
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134 USA

EMail: ayabaner@cisco.com

Dinesh Dutt
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706 USA

Phone: +1-408-527-0955
EMail: ddutt@cisco.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

Anoop Ghanwani
Brocade
130 Holger Way
San Jose, CA 95134 USA

Phone: +1-408-333-7149
EMail: anoop@alumni.duke.edu

