

Network Working Group
INTERNET-DRAFT
Intended status: Proposed Standard

Donald Eastlake 3rd
Stellar Switches
Ayan Banerjee
Dinesh Dutt
Cisco
Radia Perlman
Intel
Anoop Ghanwani
Brocade
June 27, 2010

Expires: December 26, 2010

TRILL Use of IS-IS
<[draft-eastlake-isis-trill-01.txt](#)>

Abstract

The IETF has standardized the 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 Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Distribution of this document is unlimited. Comments should be sent to the ISIS and TRILL working group mailing lists: <isis-wg@ietf.org> and <rbridge@postel.org>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Table of Contents

1. Introduction.....	3
1.1 Conventions used in this document.....	3
2. TLV and sub-TLV Extensions to IS-IS for TRILL.....	4
2.1 The Group Address TLV.....	4
2.1.1 The Group MAC Address sub-TLV.....	4
2.2 Multi-Topology Aware Port Capability TLV.....	6
2.2.1 The Special VLANs and Flags sub-TLV.....	6
2.2.2 Enabled-VLANs sub-TLV.....	8
2.2.3 Appointed Forwarders sub-TLV.....	9
2.3 Sub-TLVs for the Router Capability TLV.....	10
2.3.1 The TRILL Version sub-TLV.....	10
2.3.2 The Nickname sub-TLV.....	10
2.3.3 The Trees sub-TLV.....	11
2.3.4 The Tree Identifiers Sub-TLV.....	12
2.3.5 The Trees Used Identifiers Sub-TLV.....	13
2.3.6 Interested VLANs and Spanning Tree Roots sub-TLV....	13
2.3.7 The VLAN Group sub-TLV.....	15
2.4 MTU sub-TLV of the Extended Reachability TLV.....	16
2.5 TRILL Neighbor TLV.....	17
3. The MTU PDUs.....	19
4. Use of Existing PDUs and TLVs.....	20
4.1 TRILL IIH PDUs.....	20
4.2 Area Address.....	20
4.3 Protocols Supported.....	20
5. Acknowledgements.....	21
6. IANA Considerations.....	22
6.1 Allocations From Existing Registries.....	22
6.2 New Sub-Registries Created and Their Initial Contents.	23
7. Security Considerations.....	24
8. References.....	25
8.1 Normative References.....	25
8.2 Informative References.....	25
Change History.....	26
Changes from -00 to -01.....	26

1. Introduction

The IETF has standardized the TRILL protocol, which provides transparent Layer 2 forwarding using encapsulation with a hop count and [\[IS-IS\]](#) link state routing [\[RFCtrill\]](#). 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 are compatible with IEEE 802.1 customer bridges and can incrementally replace such bridges.

This document specifies the data formats and code points for the IS-IS [\[IS-IS\]](#) [\[RFC1195\]](#) extensions to support TRILL.

1.1 Conventions used in this document

The terminology and acronyms defined in [\[RFCtrill\]](#) are used herein with the same meaning.

Additional acronyms used in this document:

IIH - IS-IS Hello

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

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 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 6](#).

2.1 The Group Address TLV

The Group Address (GADDR) TLV, IS-IS TLV type 142 [TBD], 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 advertising 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 6.2](#)).

GADDR has the following format:

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

- o Type: TLV Type, set to GADDR-TLV 142 [TBD].
- 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 [\[RFCtrill\]](#). 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 TLV

The Multi Topology aware Port Capability (MT-PORT-CAP) TLV is IS-IS TLV type 143 [TBD] and has the format shown below. The sub-TLVs that it carries are a new series of sub-TLVs.

```

+---+---+---+---+
|Type=MT PORTCAP|          (1 bytes)
+---+---+---+---+
| Length          |          (1 bytes)
+---+---+---+---+---+---+---+---+---+---+
|  RESV  |   Topology-ID   | (2 bytes)
+---+---+---+---+---+---+---+---+---+---+
|               sub-TLVs...
+---+---+---+---+---+---+---+---+---+---+

```

- o Type: TLV Type, set to MT-PORT-CAP TLV 143 [TBD].
- o Length: Variable, minimum 2.
- o RESV: 4 reserved bits that MUST be sent as zero and ignored on receipt.
- o Topology-ID: In TRILL this value is set to zero and ignored on receipt. When this TLV is used in other protocols, it may be a 12-bit field containing the ID of the topology being announced or zero if base topology information is being carried.
- o sub-TLVs: The TLV value contains sub-TLVs formatted as described in [\[RFC5305\]](#).

The MT-PORT-CAP TLV is carried only within a Hello PDU. The sub-sections of this section below specify sub-TLVs to be carried in the MT-PORT-CAP TLV.

2.2.1 The Special VLANs and Flags sub-TLV

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


```

+---+---+---+---+---+
|Type=VLAN Flags|          (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 [TBD].
- o Length: 8.
- o Port ID: An ID for the port on which the enclosing TRILL IIH PDU is being sent as specified in [\[RFCtrill\] Section 4.4.2](#).
- o Sender nickname: If the sending IS is holding any nicknames as discussed in [\[RFCtrill\] 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 need a nickname for their first hop to insert as the ingress nickname to correctly format a TRILL encapsulated data frame. See [\[RFCtrill\] 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 [\[RFCtrill\] Section 4.4.5](#).
- o Desig.VLAN: The 12-bit ID of the designated VLAN for the link as specified in [\[RFCtrill\] Section 4.2.4.2](#).
- o AF, AC, VM, BY, and TR: These flag bits have the following meaning each, if set to one, as specified in the listed section of [\[RFCtrill\]](#):

RFCtrill		
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 [[RFCtrill](#)] [Section 4.4.2](#). It has the following format:

```

+---+---+---+---+---+
|Type=EnabledVLAN|          (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 [TBD].
- o Length: Variable, minimum 3.
- o RESV: 4 reserved bits that MUST be sent as zero and ignored on receipt.
- o Start VLAN ID, VLAN bit-map: The third and subsequent bytes of the value provide a bit map of enabled VLANs starting at the VLAN ID

indicated in the lower 12 bits of the first two bytes. The highest

order bit of the third byte indicates the VLAN equal to the starting ID while the lowest order bit of the third byte indicates that ID plus 7. Etc.

If this sub-TLV is present more than once in a Hello PDU, a VLAN is enabled for end station service on the port where the Hello was sent if this is indicated in any occurrence of this sub-TLV in the Hello.

2.2.3 Appointed Forwarders sub-TLV

The DRB on a link uses the Appointed Forwarders sub-TLV to inform other ISs 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 [RFCtrill]. It has the following format:

```

+---+---+---+---+---+
|Type=ApptFwrdrs|          (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 Appointed Forwarders sub-TLV 3 [TBD].
- 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. A VLAN's ID appears as both the start and end VLAN to specify that single VLAN. As specified in Section 4.4 of [[RFCtrill](#)], 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 below of this [Section 2.3](#) 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 [[RFCtrill](#)].

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

- o Type: Router Capability sub-TLV Type, set to 5 (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 [[RFCtrill](#)] [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 [\[RFCtrill\]](#).
- o Tree Root Priority: This is an unsigned 16-bit integer priority to be a tree root as specified in Section 4.5 of [\[RFCtrill\]](#).
- o Nickname: This is an unsigned 16-bit integer as specified in Section 3.7 of [\[RFCtrill\]](#).

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 [\[RFCtrill\]](#). 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 [\[RFCtrill\]](#).
- o Maximum trees able to compute: An unsigned 16-bit integer as specified in Section 4.5 of [\[RFCtrill\]](#).
- o Number of trees to use: An unsigned 16-bit integer as specified in Section 4.5 of [\[RFCtrill\]](#).

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 [\[RFCtrill\]](#). 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 (...) |

```

+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

D. Eastlake, et al

[Page 12]

- 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 first sub-TLV. Subsequent sub-TLVs will have the starting number of the ordered list. 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 [TBD] (TREE-USE-IDs) and the trees listed are those that the originating IS wishes to use as specified in [\[RFCtrill\]](#) [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 [\[RFCtrill\]](#) 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,

D. Eastlake, et al

[Page 13]

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, and
 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 a 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 [\[RFCtrill\] 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 so used, 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 [\[RFCtrill\]](#) [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 [\[RFCtrill\]](#). It is initialized to zero at an IS when the 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 [\[RFCtrill\]](#) [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 differend 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 give. 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 [\[RFCtrill\]](#) [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 11 (VLAN-GROUP).
- o Length: 4 + 2*n, where n may be 0.
- 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.

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 [\[RFCtrill\] 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 6 [TBD].
- 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 [RFCtrill].

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 [RFCtrill]. The structure of the TRILL Neighbor TLV is as follows:

```

+---+---+---+---+---+
|Type=TRILLNeigh|          (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 [TBD].
- o Length: $1 + 9*n$, where n is the number of neighbor records.
- o S: Smallest flag. If this bit is a one, then the list of neighbors includes the neighbor with the smallest MAC address.
- o L: Largest flag. If this bit is a one, then the list of neighbors includes the neighbor with the largest MAC address.

- o RESV: These seven bits are reserved for future use and MUST be set

to zero on transmission 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 [[RFCtrill](#)] [Section 4.3.1](#)).
- o MTU: This field is set to the largest successfully tested MTU size for this neighbor or 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 [[RFCtrill](#)] [Section 4.4.2.1](#), all MAC address may fit into one TLV, in which case both the S and L flags would be set to one in that TLV. Otherwise, the highest MAC address in a TRILL Neighbor TLV with the L flag zero will also appear as a MAC address in some other TRILL Neighbor TLV (possibly in a different TRILL IIH PDU). And the lowest MAC address in a TRILL Neighbor TLV with the S flag zero will also appear in some other TRILL Neighbor TLV (possibly in a different TRILL IIH PDU).

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 [\[RFCtrill\] 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 [[RFCtrill](#)] specifies the contents of the TRILL IIH and how its use differs in TRILL from Layer 3 LAN IIH PDU use.

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 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 2.6](#)).

4.2 Area Address

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

```
+-----+-----+
| 0x01 (Area Address Type) | 0x02 (Length of Value) |
+-----+-----+
| 0x01 (Length of Address) | 0x00 (zero Area Address) |
+-----+-----+
```

4.3 Protocols Supported

NLPID 0xC0 has been assigned to TRILL [[RFCnlpid](#)]. A Protocols Supported TLV (#129, [[RFC1195](#)]) including that value SHOULD appear in TRILL IIH PDUs and fragment zero LSP PDUs.

5. Acknowledgements

The authors gratefully acknowledge the contributions and review by the following: Mike Shand, Dino Farinacci, Les Ginsberg, Dave Ward, and Russ White, particularly the detailed and helpful comments by Mike Shand.

6. IANA Considerations

IANA will allocate the existing registry code points listed in [Section 6.1](#) and create the new registries and their initial contents as listed in [Section 6.2](#).

6.1 Allocations From Existing Registries

This document creates two new IS-IS PDUs, namely the MTU-PROBE-PDU, and MTU-ACK-PDU, as described in [Section 3](#). IANA will assign a new PDU type to these PDUs and reflect them in the PDU registry. [suggested values below]

MTU-PROBE-PDU	Level-1 PDU Number: 23
MTU-ACK-PDU	Level-1 PDU Number: 28

This document specifies three new IS-IS TLV types, namely the Group Address TLV (GADDR, type 142), the Multi-Topology Aware Port Capability TLV (MT-PORT-CAP, type 143), and the TRILL Neighbor TLV (type 145), that need to be reflected in the IS-IS TLV code-point registry. The PDUs in which these TLVs are permitted for TRILL is 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, as follows:

- 1+ = MUST occur at least once and MAY occur multiple times. If absent, the PDU is ignored.
- * = MAY occur 0, 1, or more times.

	Section	TLV#	IIH	LSP	SNP	NUMBER
GADDR-TLV	2.1	142	-	X	-	*
MT-PORT-CAP-TLV	2.2	143	X	-	-	1+
TRILL Neighbor TLV	2.5	145	X	-	-	*

This document specifies eight new sub-TLVs from existing sub-TLV sequences, namely 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. The final "NUMBER" 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.
- * = MAY occur 0, 1, or more times.

	Section	sub-TLV#	Router Capability	Extended IS Reach	NUMBER
TRILL-VER	2.3.1	5	X	-	0-1
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	-	*
VLAN-GROUP	2.3.7	11	X	-	*
MTU	2.4	6	-	X	0-1

6.2 New Sub-Registries Created and Their Initial Contents

This document creates new sub-TLV IS-IS sub-registries for sub-TLVs within the Group Address (GADDR) and Multi-topology Port Capability (MT-PORT-CAP) TLVs and specifies four sub-TLVs within these two new registries, namely Group MAC Address (GMAC-ADDR), VLAN and Flags, Enabled-VLANs, and AppointedForwarders. 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. The final "NUMBER" column indicates the permitted number of occurrences of the sub-TLV cumulatively within all the occurrences of their TLV in that TLV's carrying PDU (or set of PDUs in the case of LSP) as follows:

- 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.

	Section	sub-TLV#	Group Address	MT Port Capability	NUMBER
GMAC-ADDR	2.1.1	1	X	-	*
VLAN-FLAGS	2.2.1	1	-	X	1
Enabled-VLANs	2.2.2	2	-	X	*
AppointedFwrdrs	2.2.3	3	-	X	*

The assignment of additional sub-TLV type values in the Group Address and MT Port Capability TLVs requires IETF Review as specified in [\[RFC5226\]](#) except that types values 0x00 and 0xFF require an IETF Standards action for assignment.

7. Security Considerations

For general TRILL protocol security considerations, see [[RFCtrill](#)].

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](#)].

8. References

Normative and informative references for this document are given below.

8.1 Normative References

- [IS-IS] - 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.
- [RFCtrill] - Perlman, R., D. Eastlake, D. Dutt, S. Gai, and A. Ghanwani, "Rbridges: Base Protocol Specification", [draft-ietf-trill-rbridge-protocol-16.txt](#), in RFC Editor's queue.
- [RFCnlpid] - Eastlake, D., "IANA Considerations for Network Layer Protocol Identifiers", [draft-eastlake-nlpid-iana-considerations-04.txt](#), in RFC Editor's queue.

8.2 Informative References

- [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.

Change History

RFC Editor Note: Delete this section before publication.

Changes from -00 to -01

Numerous clarification and editorial changes based mostly on excellent comments provided by Mike Shand.

Addition of [Section 4.3](#) on Protocols Supported TLV.

Authors' Addresses

Donald E. Eastlake 3rd
Stellar Switches
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 G. 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 Communications Systems
1745 Technology Drive
San Jose, CA 95110 USA

Phone: +1-408-333-7149
Email: anoop@brocade.com

Copyright and IPR Provisions

Copyright (c) 2010 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 BSD License. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of [RFC 5378](#). No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under [RFC 5378](#), shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

