

Network Working Group
INTERNET-DRAFT
Intended status: Proposed Standard
Expires: April 11, 2014

Donald Eastlake
Weiguo Hao
Huawei
October 12, 2013

TRILL: Nickname and Label Properties
<[draft-eastlake-trill-nick-label-prop-02.txt](#)>

Abstract

There are a number of current and prospective requirements to indicate properties of nicknames, labels, and blocks thereof, for use with the TRILL (Transparent Interconnection of Lots of Links, [RFC 6325](#)) protocol. To meet that need, this document specifies IS-IS (Intermediate System to Intermediate System) sub-TLVs and some of their uses.

Status of This Memo

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

Distribution of this document is unlimited. Comments should be sent to the TRILL working group mailing list.

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

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>. The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Table of Contents

1.	Introduction.....	3
1.1	Conventions Used in This Document.....	3
2.	The Nickname Properties Sub-TLV.....	4
3.	The Label Properties Sub-TLV.....	7
4.	IANA Considerations.....	9
5.	Security Considerations.....	9
6.	Normative References.....	10
7.	Informative References.....	10
	Acknowledgements.....	11
	Authors' Addresses.....	12

1. Introduction

There are a number of current and prospective requirements to indicate properties of Nicknames, data Labels, and blocks thereof, for use with the TRILL (Transparent Interconnection of Lots of Links [[RFC6325](#)]) protocol. To meet that need, this document specifies two IS-IS (Intermediate System to Intermediate System [[ISO-10589](#)] [[RFC1195](#)]) sub-TLVs and some of their uses.

These sub-TLVs are used to flag properties of Nicknames and data Labels. Provision is made for flags associated with

- o individual Nicknames,
- o blocks of Nicknames,
- o individual Labels, and
- o blocks of Labels.

In addition, different sizes of Nicknames and data Labels can be accommodated.

The sub-TLVs specified in this document are used as follows:

- o They appear only in the IS-IS Router Capability and MT-Capability TLVs, which are TLVs number 242 and 144 and are specified in [[RFC4971](#)] and [[RFC6329](#)] respectively.
- o They can appear multiple times in the same or different Capability or MT-Capability TLVs.

1.1 Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

2. The Nickname Properties Sub-TLV

The structure of the Nickname properties (NICK-PROP) sub-TLV is as shown below.

```

+---+---+---+---+---+
| Type = TBD          |          (1 byte)
+---+---+---+---+---+
| Length              |          (1 byte)
+---+---+---+---+---+...
| NICK-PROP RECORD 1  |          (variable)
+---+---+---+---+---+...
| NICK-PROP RECORD 2  |          (variable)
+---+---+---+---+---+...
| ...
+---+---+---+---+---+...
| NICK-PROP RECORD K
+---+---+---+---+---+...

```

Figure 1. The Nickname Properties Sub-TLV

- o Type: NICK-PROP Sub-TLV type, set to TBD.
- o Length: Variable.
- o NICK-PROP RECORD: Variable length record as described below.

Each NICK-PROP RECORD is structured as follows.

```

    0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+---+---+---+---+---+---+---+---+---+---+---+---+
|NB|  NKSZ  |IN|OK|      RESV                      |
+---+---+---+---+---+---+---+---+---+---+---+---+
| Nickname Information                (variable)
+---+---+---+---+---+---+---+---+---+---+...

```

NB: If the NB (Nickname Block) flag is zero, the Nickname Information is a single Nickname. If the NB flag is one, the Nickname Information consists of an initial and final Nickname that are treated as unsigned integers and specify a block of Nicknames inclusively.

NKSZ: The NKSZ field specifies the size of each of the one or two Nickname values (see NB bit) that occur in the Nickname Information

The assigned value of NKSZ is as follows:

NKSZ

0	16-bit Nicknames
1-7	Available for assignment

IN: The IN flag is ignored unless the NB flag is zero and the Nickname Information is a Nickname owned by the RBridge advertising the enclosing NICK-PROP sub-TLV. When it is not being ignored and is equal to one, the IN flag indicates that the Nickname may be used by the advertising RBridge as the ingress Nickname for TRILL Data frames it creates.

TRILL switches may have multiple Nicknames [[RFC6325](#)] but there is no reason within the TRILL protocol for a TRILL switch to use more than one Nickname as the ingress Nickname for TRILL Data frames it creates. If a TRILL switch is not using all the Nicknames it holds as ingress Nicknames, it SHOULD use the NICK-PROP sub-TLV to indicate the Nickname (or Nicknames) it is using as ingress. This reduces the amount of Reverse Path Forwarding Check (RPFC) state in the campus. The amount of such state at each TRILL switch port is roughly proportional to the product the number of ingress Nicknames in use and the number of multi-destination distribution trees in use. If a TRILL switch does not advertise any ingress Nickname or Nicknames using the IN flag, it may use any Nickname it hold as an ingress Nickname.

OK: The OK flag is only effective if the NB flag is one and the NICK-PROP sub-TLV is advertised by the TRILL switch that is highest priority to be a tree root. TRILL switches that understand the OK flag and see one or more OK flags that are effective and are set to one dynamically select their Nickname as specified in [[RFC6325](#)] and [[clearcorrect](#)] except that they do so only from the block or blocks of nicknames advertised by NICK-PROP sub-TLV NICK-PROP RECORDS with an effective OK flag set to one. Such blocks are called the OK Nicknames blocks and a Nickname that is included in them is called an OK Nickname. If any Nickname value in the range from 0xFFC0 through 0xFFFF or equal to 0x0000 is advertised as part of an OK Nickname block they are ignored. For example, if 0xE000 through 0xFFFF was advertised as an OK Nickname block, it would be treated as if 0xE000 through 0xFFBF was advertised.

If the OK Nickname blocks change such that any TRILL switch is holding a Nickname that is no longer OK, that TRILL switch MUST allocate a new OK Nickname. To maximize network stability, all TRILL switches that might become highest priority tree root SHOULD advertise the same OK Nickname blocks.

Intended uses of this flag are to restrict Nicknames within part
of a network so as to support some methods to implement

[multilevel] or [multidatcenter] TRILL.

RESV: The remaining 10 bits are reserved and MUST be set to zero and ignored on receipt.

Nickname Information: If NB is zero, this information consists of a single Nickname. If NB is one, this information consists of an initial and final Nickname and represents a block of Nicknames. The Nicknames are treated as unsigned integers in network byte order. If the final Nickname of a block is less than the initial Nickname, the NICK-PROP RECORD is ignored. If the initial and final Nicknames are equal, then a block of size one is indicated. Otherwise a block of Nickname values with a size greater than one is indicated, starting with initial Nickname through and including the final Nickname. If the size of each Nickname value is not a multiple of 8 bits, the Nickname values are padded with initial reserved bits up to the next multiple of 8. These reserved bits MUST be sent as zero and ignored on receipt.

Each NICK-PROP RECORD must fit within the Length of the NICK-PROP sub-TLV. If there is a truncated NICK-PROP RECORD at the end of the sub-TLV, that RECORD is ignored.

For IANA Considerations in assigning values of NKSZ and bits in the RESV field, see [Section 4](#).

3. The Label Properties Sub-TLV

The structure of the Label properties (LABEL-PROP) sub-TLV is as shown below.

```

+-+--+--+--+--+--+
| Type = TBD      |           (1 byte)
+-+--+--+--+--+--+
| Length          |           (1 byte)
+-+--+--+--+--+--+...
| LABEL-PROP RECORD 1      (variable)
+-+--+--+--+--+--+...
| LABEL-PROP RECORD 2      (variable)
+-+--+--+--+--+--+...
| ...
+-+--+--+--+--+--+...
| LABEL-PROP RECORD K
+-+--+--+--+--+--+...
```

Figure 2. The Label Properties Sub-TLV

- o Type: LABEL-PROP Sub-TLV type, set to TBD.
- o Length: Variable.
- o LABEL-PROP RECORD: Variable length record as described below.

Each LABEL-PROP RECORD is structured as follows.

```

  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|LB|  LBSZ  | SCP |MG|   RESV   |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Label Information          (variable)
+-+--+--+--+--+--+--+--+--+--+...
```

LB: If the LB (Label Block) flag is zero, the Label Information is a single Label. If the LB flag is one, the Label Information consists of an initial and final Label that are treated as unsigned integers and specify a block of Labels inclusively.

LBSZ: The LBSZ field specifies the size and type of each of the one or two Label values (see LB bit) that occur in the Nickname Information

The assigned values of LBSZ are as follows:

LBSZ

- 0 12-bit VLAN ID
- 1 24-bit FGL Label [[FGL](#)]
- 2-7

SCP: The SCP or Scope field only has an effect if it is non-zero and the LABEL-PROP sub-TLV is advertised by the TRILL switch that is highest priority to be a tree root. It indicates the scope of propagation of TRILL Data frames having the data label or any of the data labels in the block indicated by the PROPERTY RECORD. The following values for SCP are currently specified:

SCP Meaning

- 0 No scope specified
- 1 Available for assignment
- 2 Throughout a campus
- 3 Local, within part of a campus [[TreeDistr](#)]

MG: The MG bit is used to indicate the management Label. It is only effective if the LB flag is zero and the LABEL-PROP sub-TLV is advertised by the TRILL switch that is highest priority to be a tree root. All TRILL switches that understand this bit MUST indicate interest in the listed Label unless this bit is set for more than one Label, in which case only the lowest valued such Label will be considered the management Label. The failure of a TRILL switch to indicate interest in this label will be ignored and tree distribution of TRILL data frames with this label will not be pruned.

RESV: The remaining 9 bits are reserved. See [Section 4](#) for IANA Considerations.

Label Information: If LB is zero, this information consists of a single Label. If LB is one, this information consists of an initial and final Label and represents a block of Labels. The Labels are treated as unsigned integers in network byte order. If the final Label for a block is less than the initial Label, the LABEL-PROP RECORD is ignored. If the initial and final Labels are equal, then a block of size one is indicated. Otherwise a block of Label values with a size greater than one is indicated, starting with initial Label through and including the final Label. If the size of each Label value is not a multiple of 8 bits, each Label value is padded with initial reserved bits up to the next multiple of 8. These reserved bits MUST be sent as zero and ignored on receipt. For example, 12-bit Labels are padded with four initial zeros.

4. IANA Considerations

TBD

5. Security Considerations

TBD

For general TRILL security considerations, see [[RFC6325](#)].

6. 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., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", [RFC 4971](#), July 2007
- [RFC6325] - Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", [RFC 6325](#), 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.
- [clearcorrect] - D. Eastlake, M. Zhang, A. Ghanwani, A. Banerjee, V. Manral, [draft-ietf-trill-clear-correct-06.txt](#), in RFC Editor's queue.
- [FGL] - D. Eastlake, M. Zhang, P. Agarwal, R. Perlman, D. Dutt, "TRILL (Transparent Interconnection of Lots of Links): Fine-Grained Labeling", [draft-ietf-trill-fine-labeling](#), in RFC Editor queue.
- [multilevel] - Perlman, R., D. Eastlake, A. Ghanwani, H. Zhai, "Multilevel TRILL (Transparent Interconnection of Lots of Links)", [draft-perlman-trill-rbridge-multilevel](#), work in progress.

7. Informative References

- [TreeDistr] - [draft-wu-trill-lsp-ext-tree-distr-opt](#), work in progress.

Acknowledgements

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

tbd

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

Authors' Addresses

Donald Eastlake
Huawei R&D USA
155 Beaver Street
Milford, MA 01757 USA

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

Weiguo Hao
Huawei Technologies
101 Software Avenue,
Nanjing 210012, China

Phone: +86-25-56623144
Email: haoweiguo@huawei.com

Copyright, Disclaimer, and Additional IPR Provisions

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of [RFC 5378](#). No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under [RFC 5378](#), shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

