

R Bridges: More Proposed TRILL Header Options
<[draft-eastlake-trill-rbridge-more-options-03.txt](#)>

Abstract

The TRILL base protocol standard, [RFC 6325](#), specifies minimal hooks for options and [draft-ietf-trill-rbridge-options](#) specifies the format for options and a proposed initial set of options. This draft is a holding location for additional proposed options. It is not intended that this draft will ever progress to be an RFC.

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 <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/1id-abstracts.html>

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

Table of Contents

<u>1.</u>	<u>Introduction.....</u>	<u>3</u>
<u>1.1</u>	<u>Terminology.....</u>	<u>3</u>
<u>2.</u>	<u>Extended Flag Options.....</u>	<u>4</u>
<u>3.</u>	<u>TLV Options.....</u>	<u>5</u>
<u>3.1</u>	<u>Additional Flags TLV Option.....</u>	<u>5</u>
<u>3.2</u>	<u>Port ID TLV Option.....</u>	<u>5</u>
<u>3.3</u>	<u>Extended Options TLV.....</u>	<u>6</u>
<u>3.3.1</u>	<u>Extended Options IANA Considerations.....</u>	<u>7</u>
<u>3.4</u>	<u>Vendor Options TLV.....</u>	<u>8</u>
<u>3.5</u>	<u>Authentication TLV Option.....</u>	<u>8</u>
<u>3.5.1</u>	<u>Authentication Option Computations.....</u>	<u>9</u>
<u>3.5.2</u>	<u>Authentication Option Fields and Flags.....</u>	<u>9</u>
<u>4.</u>	<u>Acknowledgement.....</u>	<u>10</u>
<u>5.</u>	<u>Security Considerations.....</u>	<u>10</u>
<u>6.</u>	<u>IANA Considerations.....</u>	<u>10</u>
<u>7.</u>	<u>Normative References.....</u>	<u>11</u>
<u>8.</u>	<u>Informative References.....</u>	<u>11</u>

1. Introduction

The IETF has standardized the TRILL protocol [[RFC6325](#)] a solution for transparent shortest-path frame routing in multi-hop Ethernet networks with arbitrary topologies, using an existing link-state routing protocol and encapsulation with a hop count. That standard specifies minimal hooks for options and [[RFCopt](#)] specifies the format for options and a proposed initial set of options.

This draft is a holding location for other proposed TRILL header options. The options described here may or may not ever be adopted as extensions to the TRILL protocol specification. Most of the descriptions herein need at least some further work, may be missing IANA or Security Considerations, or have other problems. It is not intended that this draft will ever progress to be an RFC.

1.1 Terminology

The terminology and acronyms of [[RFC6325](#)] are 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. Extended Flag Options

Options that appear bit encoded in the TRILL Header options area (extended header flag options) are specified in Section 2.3.1 of [\[RFCopt\]](#) and a specific bit option is specified in [Section 3](#) of [\[RFCopt\]](#).

No additional extended flag options are specified in this version of this document.

3. TLV Options

Options that are Type, Length, Value (TLV) encoded in the TRILL Header options area (TLV options) are specified in [Section 2.3.2](#) and 2.3.3 of [\[RFCopt\]](#). Two specific TLV options are specified in [Section 4](#) of [\[RFCopt\]](#).

A number of potential additional TRILL Header TLV options are specified below.

3.1 Additional Flags TLV Option

This option provides a means of adding a variety of additional flags to the TRILL Header beyond the extended header flag options available in the first four octets of the options area.

The value of a flags option consists of additional flags, eight per octet, numbered from the high-order to the low-order bit. Thus flag 1 is the 0x80 bit of the first octet, flag 8 is the 0x01 bit of that octet, flag 9 is the 0x80 bit of the second octet, etc. The number of additional flags that can be defined is bounded only by the options space that can be available. All flags not present, because they would be in value octets beyond those specified by the option Length, are considered zero.

This option can appear once in a frame for each possible combination of the ingress-to-egress, hop-by-hop, non-critical, critical, mutable and immutable flags (all combinations except for mutable critical hop-by-hop, which is a meaningless). To simplify canonicalization for security, this option MUST NOT be included if all of the flag bits would be zero and the value MUST NOT have any trailing zero octets. Thus its Length MUST be at least 1 and at least the last octet of the value present MUST be non-zero.

The option fields and flags are as follows:

- o Type is 0xTBD.
- o Length is variable with a minimum value of 1.
- o IE, NC, MT can be any valid combination producing several variations of this option.

3.2 Port ID TLV Option

The primary purpose of the Port ID option is to normally avoid the unicast Inner.MacDA address lookup at the egress RBridge to find the port on which to send a decapsulated native frame.

This option provides a 2-octet logical destination port and a 2-octet logical source port that, in some ways, could be considered extensions to the 6-octet inner destination and source MAC addresses in a frame. These logical port designators are local to the destination and source R Bridges respectively. They may be any values that those R Bridges find convenient to efficiently map to their physical ports except that the value 0x0000 is used to indicate that a logical port designator is unknown and the value 0xFFFF is reserved and MUST NOT appear in a port ID option. Should an R Bridge implementing this option receive a frame with a destination port ID of 0xFFFF it handles it as if the destination port ID was 0x0000.

R Bridges that implement this option learn the Port ID for a remote MAC address from the source Port ID field in the Port ID option, if present, in frames they decapsulate in the same way they can learn the ingress R Bridge and VLAN. This information is timed out or replaced in the same manner as remote MAC address information learned from the data plane. Such R Bridges include their local Port ID in the source field of a Port ID option when encapsulating a frame when inclusion of this option is indicated by their local policy.

For known unicast TRILL data frames, one would expect ingress R Bridges implementing this option to include the option if they are sending to egress R Bridges that also implement the option. For multi-destination TRILL data frames, inclusion of a Port ID option with a source port ID may make sense but the destination port ID is meaningless and ignored by egress R Bridges.

The option fields and flags are as follows:

- o Type is 0xTBD.
- o Length MUST be 6. The data is the 2-octet destination port ID followed by the 2-octet source port ID followed by two reserved octets. The reserved octets MUST be set to zero when this option is inserted by an ingress R Bridge, be copied without change but otherwise ignored by transit R Bridges, and be ignored by egress R Bridges.
- o IE and NC MUST be one. This is an ingress-to-egress non-critical option.
- o MT MUST be zero. This is an immutable option.

[3.3](#) Extended Options TLV

This option provides an extension mechanism for shorter options to avoid using up top-level Types in TLV option encoding. These extended options could be referred to as sub-TLV encoded.

The value part of an Extended Options TLV consists of extended

options each of which is sub-TLV encoded as follows:

```

| 0  1  2  3  4  5  6  7  8| 9 10 11 12 13 14 15|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Extended Type           | Length | Value...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |         |

```

The Extended Type is an unsigned 12-bit number whose high order part is the first octet and whose low order part is the high order four bits of the second octet.

The Length field is the low order four bits of the second octet. It gives the length of the extended option value, that is, the number of additional octets after the extended type and length.

There is no need for IE, NC, or MT flags in an extended option sub-TLV because each particular extended option inherits its IE, NC, and MT status from the enclosing Extended Options TLV.

The length specified in an Extended Options TLV MUST be exactly the sum of the total lengths of the Extended Options that its value area contains, that is, the sum of the enclosed Extended Option sub-TLV lengths plus two times the number of Extended Option sub-TLVs for their initial two octets.

If multiple Extended Options are present in an Extended Options TLV, there is no space between them. Thus, while an Extended Option is an integer number of octets, it has no other special alignment. Transit RBRidges MUST NOT re-order the extended options within an ingress-to-egress Extended Options TLV.

The option fields and flags are as follows:

- o Type is 0xTBD.
- o Length minimum 2.
- o IE, NC, and MT can be any valid combination producing several variations of this option.

[3.3.1](#) Extended Options IANA Considerations

The Extended Types 0x000 and 0xFFF are reserved and require IETF standards action for allocation. The values 0x001 through 0xFFE are available for assignment based on RFC publication. The assigning RFC MUST specify, for each extended option type, the allowed values of the IE, NC, and MT bits in the Extended Options TLV that will enclose occurrences of that specific Extended Option.

3.4 Vendor Options TLV

This option is provided to supply a standard method for the specification of vendor specific options in a TRILL Header. Vendor identity and specific options are encoded into the value associated with the option. Since vendor specific options could require any valid combination of the IE, NC, and MT bits, they inherit they inherit these from enclosing Vendor Options TLV.

Each vendor specific option starts with three bytes of OUI followed by a forth byte that has, in the bottom four bits, the length of the additional value of the vendor specific option in bytes. The top four bits of this fourth byte are available for vendor use, for example, to distinguish different options.

Transit RBridges MUST NOT re-order the vendor specific options within an ingress-to-egress Vendor Options TLV. The Vendor Options fields and flags are as follows:

- o Type is 0xTBD.
- o Length is variable with a minimum of 4.
- o IE, NC, and MT can be any valid combination producing several variations of this option.

3.5 Authentication TLV Option

[The write-up of this option is not complete.]

TRILL provides an authentication option that builds on the IS-IS security keying [[RFC5304](#)] [[RFC5310](#)] and can be applied to frames with a TRILL Header.

The first octet of the option value is the same algorithm selection code as for the IS-IS Authentication TLV (IS-IS TLV #). The value length for the option is variable and depends on the algorithm in the same way as the value in the IS-IS security TLV. Algorithm zero indicates a plain text password which must be configured in code which generates and checks this TLV and is NOT RECOMMENDED. Thus far, other algorithms have indicated HMAC signing of a canonical form of the message using a shared secret that must likewise be configured.

This option can appear up to twice in a frame, once for ingress-to-egress authentication and once for hop-by-hop authentication.

3.5.1 Authentication Option Computations

For algorithms which depend on the value of the frame (i.e., all strong authentication algorithms), the frame must be canonicalized before the authentication code is computed or verified. This canonicalization is logically done by copying the frame starting with the TRILL Header and modifying the copy as follows:

1. Set the TRILL Header Hop Count to zero.
2. Clear the octets of the Security Option after the algorithm selection code.
3. For all mutable options, setting the option Length to zero and remove any value octets.
4. If an ingress-to-egress authentication code is being computed, since hop-by-hop options can be added or deleted in transit, all hop-by-hop options must be removed from the frame copy. When a critical hop-by-hop option is removed, any required adjustments must be made in the remainder of the frame, as if it was about to be forwarded to an RBridge that did not support that hop-by-hop critical option.
5. Adjust the TRILL Header Op-Length downward as necessary to make it correct for the other adjustments made in the frame copy.

The authentication code is then calculated using this copy and either inserted into the authentication option in the real frame for transmission or compared against the authentication code in the real frame for verification.

3.5.2 Authentication Option Fields and Flags

The security option fields and flags are as follows:

- o Type is 0xTBD.
- o Length MUST be at least 1.
- o IE is variable. There may be an ingress-to-egress or hop-by-hop security option in a frame or both.
- o NC and MT MUST be zero. This is a critical, immutable option.

4. Acknowledgement

The Port ID option was suggested as part of the TRILL Header by Silvano Gai.

5. Security Considerations

-- TBD --

6. IANA Considerations

-- See IANA Considerations sub-sections above. --

7. Normative References

- [RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997
- [RFC6325] - Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (R Bridges): Base Protocol Specification", [RFC 6325](#), July 2011.
- [RFCopt] - D. Eastlake, A. Ghanwani, V. Manral, and C. Bestler, "R Bridges: TRILL Header Options", [draft-ietf-trill-rbridge-options](#), work in progress, June 2010.

8. Informative References

- [RFC5304] - Li, T. and R. Atkinson, "Intermediate System to Intermediate System (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.

Authors' Addresses

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

tel: +1-508-333-2270
email: d3e3e3@gmail.com

Copyright and IPR Provisions

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. 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.

