

Internet Engineering Task Force (IETF)
Request for Comments: 7179
Updates: 6325
Category: Standards Track
ISSN: 2070-1721

D. Eastlake 3rd
Huawei
A. Ghanwani
Dell
V. Manral
Ionos Corp.
Y. Li
Huawei
C. Bestler
Nexenta Systems
May 2014

Transparent Interconnection of Lots of Links (TRILL): Header Extension

Abstract

The IETF Transparent Interconnection of Lots of Links (TRILL) base protocol (RFC 6325) specifies minimal hooks to safely support TRILL Header extensions. This document specifies an initial extension providing additional flag bits and specifies some of those bits. It updates RFC 6325.

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/rfc7179>.

Copyright Notice

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

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

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	3
2. TRILL Header Extensions	3
2.1. RBridge Extended Flag Handling Requirements	5
2.2. No Critical Surprises	5
2.3. Extended Header Flags	6
2.3.1. Critical Summary Bits	7
2.4. Conflict of Extensions	8
3. Specific Extended Header Flags	9
3.1. RBridge Channel Alert Extended Flags	9
4. Additions to IS-IS	9
5. IANA Considerations	10
6. Security Considerations	10
7. Acknowledgements	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11

1. Introduction

The base IETF Transparent Interconnection of Lots of Links (TRILL) protocol [RFC6325] provides a TRILL Header extension feature and describes minimal hooks to safely support header extensions. (This feature is called "options" in Section 3.8 of [RFC6325].) But, except for the first two bits, the TRILL base protocol document does not specify the structure of extensions to the TRILL Header nor the details of any particular extension.

This document is consistent with [RFC6325] and provides further details. It specifies an initial extension word providing additional flag bits and specifies some of those bits. Additional extensions, including TLV-encoded options, may be specified in later documents, for example, [Options] and [Options2].

Section 2 below describes some general principles of TRILL Header extensions and an initial extension. Section 3 specifies a pair of flags in this initial extension.

1.1. Conventions Used in This Document

The terminology and acronyms defined in [RFC6325] are used herein with the same meaning. Devices implementing the TRILL protocol are referred to as RBridges (Routing Bridges) or TRILL Switches.

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. TRILL Header Extensions

The base TRILL protocol includes a feature for extension of the TRILL Header (see [RFC6325], Sections 3.5 and 3.8). The 5-bit Op-Length header field gives the length of the extensions to the TRILL Header in units of 4 octets, which allows up to 124 octets of header extension. If Op-Length is zero, there are no header extensions present; else, the extension area follows immediately after the Ingress RBridge Nickname field of the TRILL Header. The first 32-bit word of the optional extensions area consists of an extended flags area and critical summary bits as specified in this document.

As described below, provision is made for

- o hop-by-hop flags, which might affect any RBridge that receives a TRILL Data frame with such a flag set,
- o ingress-to-egress flags, which would only necessarily affect the RBridge(s) where a TRILL frame is decapsulated,
- o flags affecting an as-yet-unspecified class of RBridges, for example, border RBridges in a TRILL campus extended to support multi-level IS-IS (Intermediate System to Intermediate System) [MultiLevel], and
- o both "critical" and "non-critical" flags.

Any RBridge receiving a frame with a critical hop-by-hop extension that it does not implement MUST discard the frame because it is unsafe to process the frame without understanding such a critical extension.

Any egress RBridge receiving a frame with a critical ingress-to-egress extension it does not implement MUST drop the frame if it is a unicast frame (TRILL Header M bit = 0); if it is a multi-destination TRILL Data frame (M=1), then it MUST NOT be egressed at that RBridge, but the egress RBridge still forwards such a frame on the distribution tree.

Non-critical extensions can be safely ignored.

Any extended flag indicating a significant change in the structure or interpretation of later parts of the frame that, if the extended flag were ignored, could cause a failure of service or violation of security policy MUST be a critical extension. If such an extended flag affects any fields that transit RBridges will examine, it MUST be a hop-by-hop critical extended flag.

Note: Most RBridge implementations are expected to be optimized for simple and common cases of frame forwarding and processing. Although the hard limit on the header extensions area length, the 32-bit alignment of the extension area, and the presence of critical extension summary bits, as described below, are intended to assist in the efficient hardware processing of frames with a TRILL Header extensions area, nevertheless the inclusion of extensions may cause frame processing using a "slow path" with inferior performance to "fast path" processing. Limited slow path throughput of such frames could cause some of them to be discarded.

2.1. RBridge Extended Flag Handling Requirements

All RBridges MUST check whether there are any critical flags set that are necessarily applicable to their processing of the frame. To assist in this task, critical summary bits are provided that cover not only the extended flags specified herein but will cover any further extensions that may be specified in future documents, for example, [Options] and [Options2]. If an RBridge does not implement all critical flags in a TRILL Data frame, it MUST treat the frame as having an unimplemented critical extension as described in Section 2. A transit or egress RBridge may assume that the critical summary bits are correct.

In addition, a transit RBridge:

- o MAY set or clear hop-by-hop flags as specified for such flags;
- o MUST adjust the length of the extensions area, including changing Op-Length in the TRILL Header, as appropriate if it adds or removes the extended header flags word;
- o MUST, if it adds the word of extended header flags or changes any critical flags, correctly set the critical summary bits in the extended header flags word;
- o MUST NOT remove the extended header flags word unless it is all zero (either on arrival or after permitted modifications); and
- o MUST NOT set or clear ingress-to-egress or reserved extended header flags except as specifically permitted in the specification of such flags.

2.2. No Critical Surprises

RBridges advertise the extended header flags they support in IS-IS PDUs (Protocol Data Units) [RFC7176]. Unless an RBridge advertises support for a critical extended header flag, it will not normally receive frames with that flag set. An RBridge is not required to support any extensions.

An RBridge SHOULD NOT set a critical extended flag in a frame unless,

- o for a critical hop-by-hop extended header flag, it has determined that the next hop RBridge or RBridges that will accept the frame support that flag,

"SHOULD NOT" is specified above since there may be cases where it is acceptable for those frames, particularly for the multi-destination case, to be discarded or not egressed by any R Bridges that do not implement the extended flag.

Bits	Description

0-2	Crit.: Critical summary bits. 0 CHbHS: Critical Hop-by-Hop extension(s) are present. 1 CItES: Critical Ingress-to-Egress extension(s) are present. 2 CRSVS: Critical Reserved extension(s) are present.
3-7	CHbH: Critical Hop-by-Hop extended flag bits.
8-13	NCHbH: Non-critical Hop-by-Hop extended flag bits.
14-16	CRSV: Critical Reserved extended flag bits.
17-20	NCRSV: Non-critical Reserved extended flag bits.
21-26	CItE: Critical Ingress-to-Egress extended flag bits.
27-31	NCItE: Non-critical Ingress-to-Egress extended flag bits.

2.3.1. Critical Summary Bits

The top three bits of the extended header flags area, bits 0, 1, and 2 above, are called the critical summary bits. They summarize the presence of critical extensions as follows:

CHbHS: If the CHbHS (Critical Hop-by-Hop Summary) bit is one, one or more critical hop-by-hop extensions are present. These might be critical hop-by-hop extended header flags or critical hop-by-hop extensions after the first word in the extensions area. Transit RBridges that do not support all of the critical hop-by-hop extensions present, for example, an RBridge that supported no critical hop-by-hop extensions, MUST drop the frame. If the CHbHS bit is zero, the frame is safe, from the point of view of extensions processing, for a transit RBridge to forward, regardless of what extensions that RBridge does or does not support.

CItES: If the CItES (Critical Ingress-to-Egress Summary) bit is a one, one or more critical ingress-to-egress extensions are present. These might be critical ingress-to-egress extended header flags or critical ingress-to-egress extensions after the first word in the extensions area. If the CItES bit is zero, no such extensions are present. If either CHbHS or CItES is non-zero, egress RBridges that do not support all critical extensions present, for example, an RBridge that supports no critical extensions, MUST drop the frame. If both CHbHS and CItES are zero, the frame is safe, from the point of view of extensions, for an egress RBridge to process, regardless of what extensions that RBridge does or does not support.

CRSVS: If the CRSVS (Critical Reserved Summary) bit is a one, one or more critical extensions are present that are reserved to apply to a class of RBridges to be specified in the future, for example, border RBridges in a TRILL campus extended to support multi-level IS-IS. This class will be a subset of transit RBridges. RBridges in this class MUST drop frames with the CRSVS bit set unless they implement all critical hop-by-hop and all critical reserved extensions present in the frame.

The critical summary bits enable simple and efficient processing of TRILL Data frames by egress RBridges that support no critical extensions, by transit RBridges that support no critical hop-by-hop extensions, and by RBridges in the reserved class that support no critical hop-by-hop or reserved extensions. Such RBridges need only check whether Op-Length is non-zero and, if it is, check the top one, two, or three bits just after the fixed portion of the TRILL Header. Based on those three bits, such RBridges can decide whether to discard or forward/process the frame.

2.4. Conflict of Extensions

Defining TRILL extensions including extended header flags that conflict with each other would be undesirable. Should conflicting extensions appear in the same packet, the results would be unpredictable if different implementations processed them in different orders. While rules could be defined to specify how to predictably process conflicting extensions, such rules would also limit implementation flexibility and could impose substantial processing burdens.

Conflicting extensions SHOULD NOT be defined, but if they are, careful thought should be given as to whether and how to specify the handling of conflicting extensions.

3. Specific Extended Header Flags

The table below shows the state of TRILL Header extended flag assignments. See Section 5 for IANA Considerations.

Bits	Purpose	Section
0-2	Critical Summary Bits	2.3.1
3-6	available critical hop-by-hop flags	
7	Critical Channel Alert flag	3.1
8	Non-critical Channel Alert flag	3.1
9-13	available non-critical hop-by-hop flags	
14-16	available critical reserved flags	
17-20	available non-critical reserved flags	
21-26	available critical ingress-to-egress flags	
27-31	available non-critical ingress-to-egress flags	

Table 1: Extended Header Flags Area

3.1. RBridge Channel Alert Extended Flags

The RBridge Channel Alert extended header flags indicate that the frame is an RBridge Channel frame [RFC7178] that requests processing at each hop.

If the Critical Channel Alert flag (bit 7) is a one and the RBridge does not implement the RBridge Channel feature or the particular RBridge Channel protocol involved [RFC7178] or the frame does not actually appear to be an RBridge Channel message, then the frame is discarded. This permits implementation, for example, of a channel message requiring strict source routing or the like, with assurance that it will be discarded rather than deviate from the directed path.

If the frame is not discarded as described above, then the presence of either the Critical or Non-critical Channel Alert flag alerts transit RBridges to the presence of an RBridge Channel message [RFC7178] that may require special handling. The non-critical alert flag supports, for example, an RBridge Channel protocol message including a "record route" function where not recording transit RBridges that do not support this function is acceptable.

4. Additions to IS-IS

RBridges use IS-IS Link State PDUs (LSPs) to inform other RBridges which extended header flags they support. The IS-IS PDU(s), TLV(s), or sub-TLV(s) used to encode and advertise this information are specified in a separate document [RFC7176].

5. IANA Considerations

IANA has created a "TRILL Extended Header Flags" subregistry within the TRILL Parameters registry. The "TRILL Extended Header Flags" subregistry is initially populated as specified in Table 1 in Section 3. References in that table to sections of this document have been replaced in the IANA subregistry by references to this document as an RFC.

New TRILL extended header flags are allocated by IETF Review [RFC5226].

To indicate support of extended header flags, IANA has assigned the following bits in the TRILL-VER and PORT-TRILL-VER Sub-TLV Capability Flag registries created by [RFC7176]:

- o Bits 3-13 of the PORT-TRILL-VER Sub-TLV Capability Flags have been assigned to indicate support of TRILL hop-by-hop extended header flags 3-13.
- o Bits 14-31 of the TRILL-VER Sub-TLV Capability Flags have been assigned to indicate support of TRILL extended header flags 14-31.

6. Security Considerations

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

For security considerations related to extended header flags, see the document where the flag is specified.

It is important that the critical summary bits in the extended header flags word be set properly. If set when critical extensions of the appropriate category are not present, frames may be unnecessarily discarded. If not set when critical extensions are present, frames may be mishandled or corrupted, and intended security policies may be violated.

The RBridge Channel Alert extended header flags have the following security considerations. Implementations should keep in mind that they might be erroneously set in a frame. If either RBridge Channel Alert flag is found set in a frame that is not an RBridge Channel message [RFC7178], the flag MAY be cleared and should have no effect except, possibly, delaying processing of the frame. If either RBridge Channel Alert flag is erroneously omitted from a frame, desired per-hop processing for the frame may not occur.

7. Acknowledgements

The following, listed in alphabetic order, are thanked for their valuable contributions: Ben Campbell, Adrian Farrel, Barry Leiba, and Thomas Narten.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.
- [RFC7176] Eastlake 3rd, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 7176, May 2014.
- [RFC7178] Eastlake 3rd, D., Manral, V., Li, Y., Aldrin, S., and D. Ward, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Support", RFC 7178, May 2014.

8.2. Informative References

- [MultiLevel] Perlman, R., Eastlake 3rd, D., Ghanwani, A., and H. Zhai, "Flexible Multilevel TRILL (Transparent Interconnection of Lots of Links)", Work in Progress, January 2014.
- [Options] Eastlake 3rd, D., Ghanwani, A., Manral, V., and C. Bestler, "RBridges: Further TRILL Header Extensions", Work in Progress, June 2012.
- [Options2] Eastlake 3rd, D., "RBridges: More Proposed TRILL Header Options", Work in Progress, October 2011.

Authors' Addresses

Donald Eastlake 3rd
Huawei R&D USA
155 Beaver Street
Milford, MA 01757
USA

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

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

EMail: anoop@alumni.duke.edu

Vishwas Manral
Ionos Corp.
4100 Moorpark Ave.
San Jose, CA 95117
USA

EMail: vishwas@ionosnetworks.com

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China

Phone: +86-25-56622310
EMail: liyizhou@huawei.com

Caitlin Bestler
Nexenta Systems
455 El Camino Real
Santa Clara, CA 95050
USA

EMail: caitlin.bestler@nexenta.com

