Network Working Group
Internet-Draft
Intended status: Standards-Track
Expires: October 16, 2012

V. Manral
Hewlett Packard Co.
D. Eastlake
Huawei R&D USA
D. Ward
A. Banerjee
Cisco Systems
April 09, 2012

TRILL: Bidirectional Forwarding Detection (BFD) Support
draft-ietf-trill-rbridge-bfd-04

Abstract

This document specifies use of the BFD (Bidirectional Forwarding Detection) protocol in RBridge campuses based on the Rbridge Channel extension to the the TRILL (TRansparent Interconnection of Lots of Links) protocol.

BFD is a widely deployed OAM (Operations, Administration, and Maintenance) mechanism in IP and MPLS networks, using UDP and ACH encapsulation respectively. This document specifies the BFD encapsulation over TRILL.

Requirements Language

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].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on July 27, 2012.

Table of Contents

1. Introduction

Faster convergence is a critical feature of TRILL networks. The TRILL IS-IS Hellos used between RBridges provide a basic neighbor and continuity check for TRILL links. However, failure detection by non-receipt of such Hellos is based on the holding time parameter that is commonly set to a value of tens of seconds and, in any case, has a minimum expressible value of one second.

Some applications, including voice over IP, may wish, with high probability, to detect interruptions in continuity within a much shorter time period. In some cases physical layer failures can be detected very rapidly but this is not always possible, such as when there is a failure between two bridges that are in turn between two RBridges. There are also many subtle failures possible at higher levels. For example, some forms of failure could affect unicast frames while still letting multicast frames through; since all TRILL IS-IS Hellos are multicast such a failure cannot be detected with Hellos. Thus, a low overhead method for frequently testing continuity for the TRILL Data between neighbor RBridges is necessary for some applications. BFD protocol provides a low-overhead, short-duration detection of failures in the path between forwarding engines.

This document describes a TRILL encapsulation for BFD packets for networks that do not use IP addressing or for ones where it is not desireable.

2. Terminology

ACH: Associated Header

BFD: Bi-directional Forwarding Detection

IP: Internet Protocol

IS-IS: Intermediate-System to Intermediate-System

MPLS: Multi Protocol Label Switching

PPP: Point-to-Point Protocol

OAM: Operations, Administration, and Maintenance

3. BFD over TRILL

TRILL supports neighbor BFD Echo and one-hop and multi-hop BFD Control, as specified below, over the Rbridge Channel facility. Multi-destination BFD is beyond the scope of this document, although there is work being done in that area [MultiBFD]. The Rbridge Channel facility is specified in [TRILLChannel].

BFD over TRILL support is similar to BFD over IP support [RFC5881] except where differences are explicitly mentioned. When running BFD over TRILL both Single Hop as well as in Multi Hop sessions are supported.

Asynchronous mode is supported, and the demand mode is not supported for TRILL. BFD over TRILL supports the Echo function, however this can be used for only Single hop sessions.

The TRILL Header Hop count in the BFD packets is sent out with the maximum value of 63. To prevent spoofing attacks, the TRILL Hop count of a received session is checked [RFC5082]. For a single Hop session if the Hop count is less than 63, and the Rbridge Channel Header MH flag is zero, the packet is discarded. For Multi Hop sessions the Hop count check can be disabled if the MH flag is one.

Like in BFD for IP the format of the Echo Packet content is not defined.

New Rbridge Channel types for BFD TRILL Control frame and BFD echo packet are specified.

Authentication mechanisms as supported in BFD are also supported for BFD running over TRILL.

4. Sessions and Initialization

Within an RBridge campus, there will be only a single TRILL BFD Control session between two RBridges over a given interface visible to TRILL. This BFD session must be bound to this interface. As such, both sides of a session MUST take the "Active" role (sending initial BFD Control packets with a zero value of Your Discriminator), and any BFD packet from the remote machine with a zero value of Your Discriminator MUST be associated with the session bound to the remote system and interface.

Note that TRILL BFD provides OAM facilities for the TRILL Data plane. This is above whatever protocol is in use on a particular link, such as a PPP [TrillPPP] link or an Ethernet link. Link technology

specific OAM protocols may be used on a link between neighbor RBridges, for example Continuity Fault Management [802.1Q] if the link is Ethernet. But such link layer OAM and coordination between it and TRILL data plaen layer OAM, such as TRILL BFD, is beyond the scope of this document.

If lower level mechanisms, such as link aggregation [802.1AX], are in use that present a single logical interface to TRILL IS-IS, only a single TRILL BFD session can be established to any other RBridge over this logical interface. However, lower layer OAM could be aware of and/or run separately on each of the components of an aggregation.

5. Relationship to MPLS OAM

TRILL BFD uses the TRILL Rbridge Channel [TRILLChannel] similar to the way that MPLS OAM protocols use the MPLS Generic Associated Channel [RFC5586]. However, the RBridges that implement TRILL are IS-IS [IS-IS] based routers, not label switched routers; thus TRILL BFD is closer to IPv4/IPv6 BFD than to MPLS BFD.

TRILL BFD optionally includes support of BFD Echo that is not specified for MPLS BFD, due to the one way nature of the basic MPLS service.

6. TRILL BFD Control Protocol

TRILL BFD Control frames are unicast TRILL Rbridge Channel frames [TRILLChannel]. The Rbridge Channel Protocol value is given in Section 10.

The protocol specific data associated with the TRILL BFD Control protocol is as shown below. See [RFC5880] for further information on these fields.

TRILL BFD Control Protocol Data:

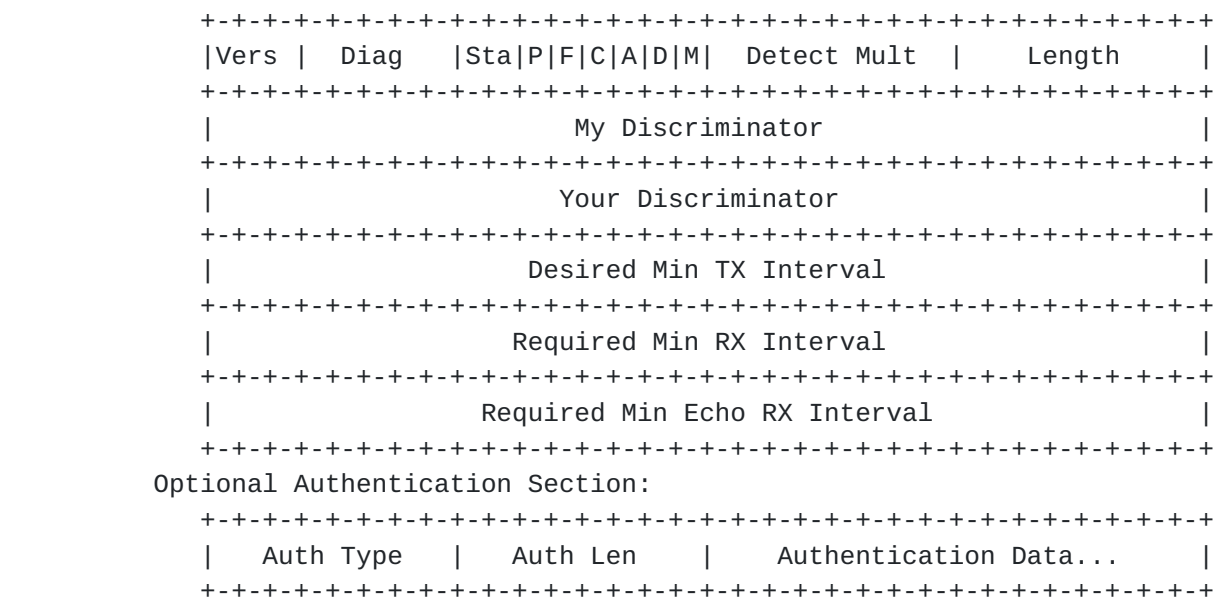

Optional Authentication Section:

7. One-Hop TRILL BFD Control

One-hop TRILL BFD Control is typically used to rapidly detect link and RBridge failures. TRILL BFD frames over one hop for such purposes SHOULD be sent with priority 7.

For neighbor RBridges RB1 and RB2, each RBridge sends one-hop TRILL BFD Control frames to the other only if TRILL IS-IS has detected bi-directional connectivity , that is, the adjacency is in the Two-Way or report Report state [RFC6327] and both RBridges indicate support of TRILL BFD is enabled. The BFD Enabled TLV is used to indicate this as specified in [RFC6213].

8. BFD Control Frame Processing

The following tests SHOULD be performed on received TRILL BFD Control frames before generic BFD processing.

Is the M bit in the TRILL Header non-zero? If so, discard the frame. TRILL support of multi-destination BFD Control is beyond the scope of this document, although work is being done in the Area [MultiBFD].

If the Channel Header MH flag is zero, indicating one-hop, test that the TRILL Header hop count received was 0x3F (i.e., is 0x3E if it has already been decremented) and if it is any other value discard the frame. If the MH Channel flag is one, indicating multi-hop, test that the TRILL Header hop count received was not less than a

configurable value that defaults to 0x30. If it is less, discard the frame.

9. TRILL BFD Echo Protocol

A TRILL BFD Echo frame is a unicast Rbridge Channel frame, as specified in [TRILLoam], which should be forwarded back by an immediate neighbor because both the ingress and egress nicknames are set to a nickname of the originating RBridge. Normal TRILL Data frame forwarding will cause the frame to be returned. The TRILL OAM protocol number for BFD Echo is given in Section 4.

TRILL BFD Echo frames SHOULD only be sent on a link if

A TRILL BFD Control session has been established,

TRILL BFD Echo support is indicated by the potentially echo responding RBridge.

The adjacency is in the Report state [RFC6327], and

The TRILL BFD Echo originating RBridge wishes to make use of this optional feature.

Since the originating RBridge is the RBridge that will be processing a returned Echo frame, the entire TRILL BFD Echo protocol specific data area is considered opaque and left to the discretion of the originating RBridge. Nevertheless, it is RECOMMENDED that this data include information by which the originating RBridge can authenticate the returned BFD Echo frame and confirm the neighbor that echoed the frame back. For example, it could include its own SystemID, the neighbor's SystemID, a session identifier and a sequence count as well as a Message Authentication Code.

9.1. BFD Echo Frame Processing

The following tests MUST be performed on returned TRILL BFD Echo frames before other processing. The RBridge Channel document requires that the information in the TRILL Header be given to the BFD protocol.

Is the M-bit in the TRILL Header non-zero? If so, discard the frame. TRILL support of multi-destination BFD Echo is beyond the scope of this document.

The TRILL BFD Echo frame should have gone exactly two hops so test that the TRILL Header hop count as received was 0x3E (i.e., 0x3D if

it has already been decremented) and if it is any other value discard the frame. The Rbridge Channel Header in the frame should have the MH bit equal to one and if it is zero, the frame is discarded.

10. Management and Operations Considerations

The TRILL BFD parameters on an RBridge are configurable. The default values are the same as in the IP BFD case [RFC5881], except where specified in this document such as for Hop Count.

It is required that the operator of an RBridge campus configure the rates at which TRILL BFD frames are transmitted on a link to avoid congestion (e.g., link, I/O, CPU) and false failure detection.

11. Security Considerations

Consistent with TRILL's goal of being able to operate with minimum configuration, the default for BFD security between neighbor RBridges is based on that state of IS-IS shared secret authentication for Hello between those RBridges. However, if such BFD security is configured then its configuration is independent of that for IS-IS security.

If IS-IS authentication is not in effect between neighbor RBridges then, by default, TRILL BFD between those RBridges is also unsecured. If such IS-IS authentication is in effect then, unless configured otherwise, TRILL BFD Control frames sent between those RBridges use BFD Meticulous Keyed SHA1 authentication [RF5880] with keying material derived as shown below

```
HMAC-SHA256 ( ( "TRILL BFD Control" | originatorMAC ),
              IS-IS-shared-key )
```

where HMAC-SHA256 is described in [FIPS180], "TRILL BFD Control" is the seventeen byte US ASCII [RFC20] string indicated that is then concatenated with the 6-byte MAC of the originating port. The MAC is included to minimize exposure of the same key to improve resistance to cryptanalysis. IS-IS-key is the secret keying material being used for IS-IS authentication on the link. In the Authentication Section of the BFD Control frame OAM protocol specific data, Auth Type would be 5, Auth Len would be 28, and Auth Key ID is zero. The RBridges MAY be configured to use other BFD security modes or keying material or configured to use no security.

Authentication for TRILL BFD Echo is a local implementation issue as BFD Echo frames are authenticated by their sender when received in the form of Echo responses. However, if TRILL IS-IS and BFD Control are being authenticated to a neighbor and BFD Echo is in use, BFD Echo frames to be returned by that neighbor SHOULD be authenticated and such authenticate SHOULD use different keying material from other types of authentication. For example, it could use keying material derived as follows:

```
HMAC-SHA256 ( ( "TRILL BFD Echo" | originatorMAC ),
              IS-IS-shared-key )
```

12. IANA Considerations

IANA is request to allocate two Rbridge Channel Protocol numbers from the range allocated by Standards Actions, as follows:

Protocol	Number
BFD Control	TBD (2 suggested)
BFD Echo	TBD (3 suggested)

13. Acknowledgements

The authors would like to thank Dave Katz the author of [RFC5880] and [RFC5881].

14. References

14.1. Normative References

[RFC20] Cerf, V., "ASCII format for network interchange", RFC 20, October 1969.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007.

[RFC5586] Bocci, M., Vigoureux, M., and S. Bryant, "MPLS Generic Associated Channel", RFC 5586, June 2009.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

[RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", RFC 6213, April 2011.

[FIPS180] US National Institute of Science and Technology, "Secure Hash Standard (SHS)", FIP PUB 180-4, 2011.

[RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.

[RFC6327] Eastlake, D., R. Perlman, A. Ghanwani, D. Dutt, V. Manral, "RBridges: Adjacency", RFC 6327, July 2011.

[TRILLChannel] Eastlake, D., V. Manral, Y. Li, S. Aldrin, D. Ward, "RBridges: RBridge Channel Support in TRILL", draft-ietf-trill-rbridge-channel, work in progress.

14.2. Informative References

[802.1AX] IEEE, "IEEE Standard for Local and metropolitan area networks / Link Aggregation", 802.1AX-2008, 1 January 2008.

[802.1Q] IEEE 802.1, "IEEE Standard for Local and metropolitan area networks - Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, May 2011.

[RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011.

[IS-IS] ISO/IEC 10589:2002, "Intermediate system to Intermediate system routeing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode NetworkService (ISO 8473)," ISO/IEC 10589:2002.

[MultiBFD] Katz, D. and D. Ward, "draft-ietf-bfd-multipoint", work in progress.

Authors' Addresses

Vishwas Manral
Hewlett Packard Co.
19111 Pruneridge Ave.
Cupertino, CA 95089
USA

Phone: 1-408-447-0000
Email: vishwas.manral@hp.com

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

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

Dave Ward
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95138
USA

Email: dward@cisco.com

Ayan Banerjee
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95138
USA

Phone: 1-408-525-8781
Email: ayabaner@cisco.com