

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2011

V. Manral, Ed.
IPInfusion Inc.
D. Eastlake
Huawei Inc.
D. Ward
Juniper Networks
A. Banerjee
Cisco Systems
March 13, 2011

**Rbridges: Bidirectional Forwarding Detection (BFD) support for TRILL
draft-manral-trill-bfd-encaps-01**

Abstract

This document specifies use of the BFD (Bidirectional Forwarding Detection) protocol in RBridge campuses based on the OAM (Operations, Administration, and Maintenance) Channel extension to the TRILL (Transparent Interconnection of Lots of Links) protocol.

BFD is a widely deployed OAM mechanism in IP and MPLS networks. However, in the present form a BFD packet cannot be sent over a TRILL network as it is either IP/ UDP encapsulated or encapsulated directly over MPLS or using ACH encapsulation. This document also defines BFD encapsulation over TRILL to address this shortcoming.

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 [RFC 2119](#) [[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 September 14, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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	4
2.	Terminology	4
3.	BFD over TRILL	4
4.	Sessions and Initialization	5
5.	Relationship to MPLS OAM	6
6.	TRILL BFD Control Protocol	6
7.	One-Hop TRILL BFD Control	7
8.	BFD Control Frame Processing	7
9.	TRILL BFD Echo Protocol	7
9.1.	BFD Echo Frame Processing	8
10.	Management and Operations Considerations	8
11.	Security Considerations	8
12.	IANA Considerations	9
13.	Acknowledgements	9
14.	References	9
14.1.	Normative References	9
14.2.	Informative References	9
	Authors' Addresses	10

1. Introduction

Faster convergence is a very 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 which 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 desirable.

2. Terminology

BFD: Bi-directional Forwarding Detection

OAM: Operations, Administration, and Maintenance

MPLS: Multi Protocol Label Switching

IS-IS: Intermediate-System to Intermediate-System

TTL: Time To Live

3. BFD over TRILL

TRILL supports neighbor BFD Echo and one-hop and multi-hop BFD Control, as specified below, over the TRILL OAM Channel facility. Multi-destination BFD is beyond the scope of this document. The OAM Channel facility is specified in [TRILLoam].

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

Asynchronous mode is supported, however 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 sent out with a value of 63. To prevent spoofing attacks, the TRILL Hop count of a received session is checked. For a single Hop session if the Hop count is less than 63 the packet is discarded if the GTSM mode [[RFC5082](#)] is set. For Multi Hop sessions the Hop count check can be disabled or the bfdTrillAcceptedHopCount value can be configured. If a packet is received with a hop count of less than bfdTrillAcceptedHopCount, the packet is discarded.

The format of the echo packet is not defined.

A new BFD TRILL header is defined.

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 [[802.1ag](#)] if the link is Ethernet. But such link layer OAM and coordination between it and TRILL data plane 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 OAM Channel [TRILLoam] in the same way that MPLS OAM protocols use the MPLS Generic Associated Channel [RFC5586]. However, the R Bridges that implement TRILL are 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 which is not specified for MPLS BFD.

6. TRILL BFD Control Protocol

TRILL BFD Control frames are unicast TRILL OAM Message Channel frames [TRILLoam]. The TRILL OAM Protocol value is given in [Section 4](#).

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:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Vers | Diag |Sta|P|F|C|A|D|M| Detect Mult | Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     My Discriminator                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Your Discriminator                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Desired Min TX Interval                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Required Min RX Interval                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Required Min Echo RX Interval                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Optional Authentication Section:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Auth Type | Auth Len | Authentication Data... |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


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 and both RBridges indicate support of TRILL BFD is enabled. The BFD Enabled TLV is used to indicate this as specified in [RFCbfdtlv]. The indication of TRILL BFD support with the BFD Enabled TLV overrides any indication of lack of support through failure to indicate support of the OAM-Channel TRILL Header extended flag.

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.

If the OAM 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 OAM 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 TRILL OAM Message Channel frame, as specified in [TRILloam], which should be bounced 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, 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 SHOULD be performed on returned TRILL BFD Echo frames before other processing. (In some implementations, the TRILL Header may not be available to the TRILL BFD Echo module in which case these check are not possible.)

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 TRILL OAM 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 at an RBridge are configurable... The default values are ... TBD.

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

This draft raises no new security considerations than those already mentioned in the BFD [[RFC5880](#)]. By keeping a separate flag for Single Hop and Multihop sessions it allows the TTL check to be performed thus preventing spoofing of packets.

However the same is possible even without the changes mentioned in this document. A device should rate limit the LSP ping packets redirected to the CPU so that the CPU is not overwhelmed.

12. IANA Considerations

IANA is request to allocate two TRILL OAM 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 a lot of folks. Names will be disclosed soon.

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4634] Eastlake, D. and T. Hansen, "US Secure Hash Algorithms (SHA and HMAC-SHA)", [RFC 4634](#), July 2006.
- [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.

14.2. Informative References

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

[802.1ag] IEEE, "IEEE Standard for Local and metropolitan area
networks / Virtual Bridged Local Area Networks / Connectivity
Fault Management", 802.1ag-2007, 17 December 2007.

[TrillPPP] Carlson, J., "PPP TRILL Protocol Control Protocol",
[draft-ietf-pppext-trill-protocol-02.txt](#), work in progress, May
2010.

Manral, et al.

Expires September 14, 2011

[Page 9]

Authors' Addresses

Vishwas Manral (editor)
IPInfusion Inc.
1188 E. Arques Ave.
Sunnyvale, CA 94085
USA

Phone: 408-400-1900
Email: vishwas@ipinfusion.com

Donald Eastlake 3rd
Huawei Inc.
155 Beaver Street
Milford, MA 01757
USA

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

Dave Ward
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089-1206
USA

Phone: 408-745-2000
Email: dward@juniper.net

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

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