BFD Working Group
Internet Draft
Intended status: StandardTrack
Expires: April 23, 2026

W. Jiang
China Mobile
C. Lin
New H3C Technologies
X. Min
ZTE Corp.

October 20, 2025

Multiple Hop Unaffiliate BFD
draft-jiang-bfd-multi-hop-unaffiliate-03

Abstract

The Bidirectional Forwarding Detection (BFD) is a fault detection protocol designed to rapidly identify communication failure between two forwarding engines. This document suggests utilizing BFD Echo when the local system supports BFD, but the neighboring system does not. BFD Control packets and their processing procedures can be executed over the BFD Echo port, where the neighboring system solely loops packets back to the local system.

This document serves as an update to RFC 5880 and draft-ietf-bfd-unaffiliate-echo-10.

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 https://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 April 23, 2026.

Copyright Notice

Table of Contents

1. Introduction

The continuous advancement of network technology has led to widespread adoption of tunneling techniques such as GRE, VxLAN, and SR-Policy for secure data transmission and precise traffic management. These tunneling methods encapsulate the data to be transmitted as the tunnel payload and envelop it with a tunnel header specifying the destination. Upon reaching the tunnel's destination, the original message is extracted from the encapsulation for further processing.

With the increasing use of tunneling technology, it is critical for network devices to rapidly identify communication faults within multi-hop tunnels and take appropriate measures to rectify these faults to ensure service continuity. Upon detecting tunnel faults, one may choose to utilize the tunnel's backup path, select an alternative tunnel, or transition to a traditional IP path for transmission.

For example, in the context of utilizing SRv6-Policy for traffic steering, it is crucial for swift recognition of SRv6-Policy tunnel malfunctions and prompt failover to a backup path or SRv6-BE. BFD [RFC5880] serves as a low-overhead, short-duration method to detect faults on the communication path between adjacent forwarding engines, offering Asynchronous and Demand modes to accommodate various deployment scenarios. BFD also supports an Echo function to reduce device requirements. Activation of the Echo function involves the local system sending BFD Echo packets, which are looped back by the remote system through the forwarding path. If multiple consecutive BFD Echo packets are not received, the BFD session is declared to be Down.

[draft-ietf-bfd-unaffiliated-echo] outlines the detailed process of utilizing echo-BFD for fault detection in single-hop scenarios. This document delves into the use of BFD for detecting faults at both ends of the tunnel, specifically focusing on multi-hop path fault detection.

1.1. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Overview

The advancement of network technology has enabled the effective use of tunneling techniques such as GRE, VxLAN, and SR-Policy, providing secure data transmission and precise control over traffic. These techniques involve encapsulating the data within a tunnel payload and adding a tunnel header that specifies the tunnel's destination. The original message is then transmitted to the tunnel's destination as the payload. Upon reaching the destination, the tunnel encapsulation is removed, allowing the original message to be processed further.

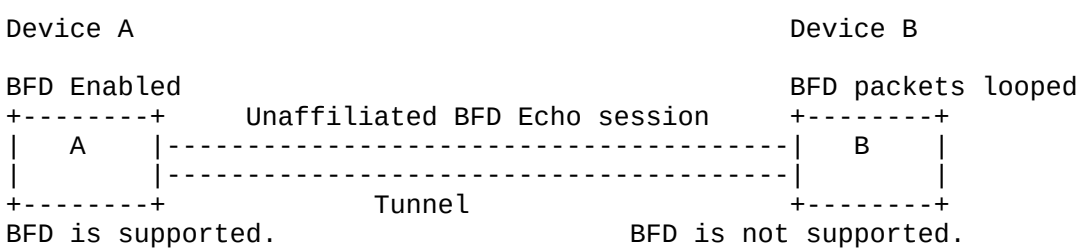

Figure 1: Unaffiliated BFD Echo diagram

As shown in Figure 1, device A supports BFD, while device B does not. Device A sends Unaffiliated BFD Echo packets through the tunnel, and upon receiving these packets, device B, which is multiple hops away from the BFD peer, immediately loops them back through the tunnel. This process enables device A to swiftly detect tunnel connectivity issues. It's important to note that device B does not intercept or parse any BFD protocol field within the received Unaffiliated BFD Echo packet.

In the context of tunnel forwarding, as the original packet is encapsulated within a tunnel for transmission, the outer TTL value of the encapsulated packet is decremented by 1 at each hop as it traverses multiple intermediate points. When the tunnel encapsulation is removed at the tunnel endpoint to restore the original packet, the TTL value from the outer tunnel encapsulation is copied back into the TTL field of the original packet header.

In the scenario where the tunnel destination device performs loopback with BFD packets, the original packet is re-encapsulated based on the local forwarding path and then sent back to the tunnel source device through the tunnel. All Unaffiliated BFD Echo packets for the session must be sent with a Time to Live (TTL) or Hop Limit value of 255.

Regarding the modification to the requirements outlined in [draft-ietf-bfd-unaffiliated-echo], when receiving echo BFD packets with TTL or Hop Limit values other than 254, the original document required them to be dropped. The modified requirement states that no such check is performed, and the packets are processed normally as long as the TTL value is not decremented to 0, while it is also recommended to perform the check based on the typical maximum number of hops in the tunnel scenario.

Overall, the Unaffiliated BFD Echo packet reuses the format of the BFD Control packet defined in [RFC5880], with the specified fields populated as described. These include parameters such as My Discriminator, Your Discriminator, Desired Min TX Interval, Required Min RX Interval, Required Min Echo RX Interval, and Detect Mult, each set according to the guidelines provided.

3. Procedure

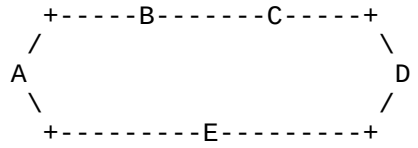

Figure 2: Multiple Hop BFD Echo Example

Based on the provided information, the process of using BFD Echo packets for tunnel detection involves the initiation of BFD Echo packets from device A and their transmission through a tunnel to the endpoint device D. Device D then loops the BFD Echo packets back to device A, resulting in detection of the tunnel’s availability for normal packet transmission. If no looped back BFD Echo packets are received within multiple intervals, tunnel fault status is reported to the upper-layer protocol for appropriate action.

The uncertainty of the transmission path can make it difficult to check the TTL values of looped-back packets. Consequently, the check on the TTL value of 254 as mentioned in [draft-ietf-bfd-unaffiliated-echo] may not be directly applicable. Instead, TTL checking should be performed based on the maximum hop count of the tunnel.

The removal of encapsulation at the tunnel’s endpoint is crucial to ensure that BFD Echo packets are looped back only after reaching the endpoint, rather than being prematurely processed by intermediate nodes. Optionally, a strict path control method can be employed, which involves specifying a strict forwarding path from the tunnel’s

starting point, ensuring the original BFD Echo packet serves as the payload of the tunnel, makes a round trip within the tunnel, and ultimately returns to the starting point with the original packet's TTL remaining unchanged throughout.

The BFD endpoint can flexibly decide whether to perform TTL checking based on the actual circumstances. In a scenario where a specific path is strictly specified, it may check if the TTL value of the looped back Echo BFD packet is 255. In other cases, the TTL value of the looped-back Echo BFD packet should be checked based on the maximum possible hops of the tunnel.

This comprehensive process addresses various considerations involved in using BFD Echo packets for tunnel detection and illustrates the complex nature of managing and monitoring tunnel connectivity.

4. Security Considerations

All Security Considerations from [RFC5880] and [RFC5881] apply.

In order to mitigate the potential reflector attack by the remote attackers, or infinite loop of the Unaffiliated BFD Echo packets, it's RECOMMENDED to put two requirements, also known as Generalized TTL Security Mechanism (GTSM) [RFC5082], on the device looping Unaffiliated BFD Echo packets, the TTL value should not be reset. Instead, the TTL value from the outer tunnel header should be copied into the inner packet, and then the packet is sent back to the tunnel source device.

If the tunnel source specifies the round-trip path of the BFD Echo packet using a strict path, then the BFD Echo packet remains unchanged as the payload, being transmitted to the tunnel destination and then back to the tunnel source. To ensure that the BFD Echo packet is transmitted to the tunnel destination, when specifying a strict path, the tunnel destination should be included as one of the nodes in the strict path.

5. IANA Considerations

This document has no IANA action requested.

6. References

6.1. Normative References

TBD

Authors' Addresses

Wenying Jiang
China Mobile
Beijing
China

Email: jiangwenying@chinamobile.com

Changwang Lin
New H3C Technologies
Beijing
China

Email: linchangwang.04414@h3c.com

Xiao Min
ZTE Corp.
Nanjing
China
Phone: +86 18061680168
Email: xiao.min2@zte.com.cn