BFD Working Group
Internet-Draft
Updates: 5880 (if approved)
Intended status: Standards Track
Expires: May 6, 2021

W. Cheng
R. Wang
China Mobile
X. Min
ZTE Corp.
R. Rahman
Cisco Systems
R. Boddireddy
Juniper Networks

November 2, 2020

Unaffiliated BFD Echo Function
draft-ietf-bfd-unaffiliated-echo-01

Abstract

Bidirectional Forwarding Detection (BFD) is a fault detection protocol that can quickly determine a communication failure between two forwarding engines. This document proposes a use of the BFD Echo function where the local system supports BFD but the neighboring system does not support BFD.

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 May 6, 2021.

Table of Contents

1. Introduction

To minimize the impact of device/link faults on services and improve network availability, a network device must be able to quickly detect faults in communication with adjacent devices. Measures can then be taken to promptly rectify the faults to ensure service continuity.

BFD [RFC5880] is a low-overhead, short-duration method to detect faults on the communication path between adjacent forwarding engines. The faults can be on interface, data link, and even forwarding engine. It is a single, unified mechanism to monitor any media and protocol layers in real time.

BFD defines Asynchronous mode to satisfy various deployment scenarios, and also supports Echo function to reduce the device requirement for BFD. When the Echo function is activated, the local system sends BFD Echo packets and the remote system loops back the received Echo packets through the forwarding path. If several consecutive BFD Echo packets are not received by the local system, then the BFD session is declared to be Down.

When using BFD Echo function, there are two typical scenarios as below:

o Full BFD protocol capability with affiliated Echo function: this scenario requires both the local device and the neighboring device to support full BFD protocol.

o Only BFD Echo function without full BFD protocol capability: this scenario requires only the local device to support sending and demultiplexing BFD Control packets.

The two typical scenarios are both reasonable and useful, and the latter is referred to as Unaffiliated BFD Echo function in this document.

Section 6.2.2 of [BBF-TR-146] describes one use case of the Unaffiliated BFD Echo function, and at least one more use case is known in the field BFD deployment.

This document describes the use of the Unaffiliated BFD Echo function over IPv4 and IPv6 for single IP hop.

2. Updates to RFC 5880

The Unaffiliated BFD Echo function described in this document reuses the BFD Echo function as described in [RFC5880] and [RFC5881], but does not require BFD asynchronous mode. When using the Unaffiliated BFD Echo function, only the local system has the BFD protocol enabled, the remote system just loops back the received BFD Echo packets as regular data packets.

With that said, this document updates [RFC5880] with respect to its descriptions on the BFD Echo function as follows.

o [RFC5880] states in the 4th paragraph of Section 3.2:

An adjunct to both modes is the Echo function. When the Echo function is active, a stream of BFD Echo packets is transmitted in such a way as to have the other system loop them back through its forwarding path. If a number of packets of the echoed data stream are not received, the session is declared to be down. The Echo function may be used with either Asynchronous or Demand mode. Since the Echo function is handling the task of detection, the rate of periodic transmission of Control packets may be reduced (in the case of Asynchronous mode) or eliminated completely (in the case of Demand mode).

* This paragraph is now updated to:

An adjunct or complement to both modes is the Echo function. When the Echo function is active, a stream of BFD Echo packets is

transmitted in such a way as to have the other system loop them back through its forwarding path. If a number of packets of the echoed data stream are not received, the session is declared to be down. The Echo function may be used with either Asynchronous or Demand mode. Since the Echo function is handling the task of detection, the rate of periodic transmission of Control packets may be reduced (in the case of Asynchronous mode) or eliminated completely (in the case of Demand mode). The Echo function may also be used independently, with neither Asynchronous nor Demand mode.

o [RFC5880] states in the 3rd and 9th paragraphs of Section 6.1:

Once the BFD session is Up, a system can choose to start the Echo function if it desires and the other system signals that it will allow it. The rate of transmission of Control packets is typically kept low when the Echo function is active.

If the session goes Down, the transmission of Echo packets (if any) ceases, and the transmission of Control packets goes back to the slow rate.

* The two paragraphs are now updated to:

When a system is running with Asynchronous mode, once the BFD session is Up, it can choose to start the Echo function if it desires and the other system signals that it will allow it. The rate of transmission of Control packets is typically kept low when the Echo function is active.

In Asynchronous mode, if the session goes Down, the transmission of Echo packets (if any) ceases, and the transmission of Control packets goes back to the slow rate.

o [RFC5880] states in the 2nd paragraph of Section 6.4:

When a system is using the Echo function, it is advantageous to choose a sedate reception rate for Control packets, since liveness detection is being handled by the Echo packets. This can be controlled by manipulating the Required Min RX Interval field (see section 6.8.3).

* This paragraph is now updated to:

When a system is using the Echo function with Asynchronous mode, it is advantageous to choose a sedate reception rate for Control packets, since liveness detection is being handled by the Echo

packets. This can be controlled by manipulating the Required Min RX Interval field (see section 6.8.3).

o [RFC5880] states in the 2nd paragraph of Section 6.8:

When a system is said to have "the Echo function active" it means that the system is sending BFD Echo packets, implying that the session is Up and the other system has signaled its willingness to loop back Echo packets.

* This paragraph is now updated to:

When a system in Asynchronous or Demand mode is said to have "the Echo function active" it means that the system is sending BFD Echo packets, implying that the session is Up and the other system has signaled its willingness to loop back Echo packets.

o [RFC5880] states in the 7th paragraph of Section 6.8.3:

When the Echo function is active, a system SHOULD set bfd.RequiredMinRxInterval to a value of not less than one second (1,000,000 microseconds). This is intended to keep received BFD Control traffic at a negligible level, since the actual detection function is being performed using BFD Echo packets.

* This paragraph is now updated to:

When the Echo function is active with Asynchronous mode, a system SHOULD set bfd.RequiredMinRxInterval to a value of not less than one second (1,000,000 microseconds). This is intended to keep received BFD Control traffic at a negligible level, since the actual detection function is being performed using BFD Echo packets.

o [RFC5880] states in the 1st and 2nd paragraphs of Section 6.8.9:

BFD Echo packets MUST NOT be transmitted when bfd.SessionState is not Up. BFD Echo packets MUST NOT be transmitted unless the last BFD Control packet received from the remote system contains a nonzero value in Required Min Echo RX Interval.

BFD Echo packets MAY be transmitted when bfd.SessionState is Up. The interval between transmitted BFD Echo packets MUST NOT be less than the value advertised by the remote system in Required Min Echo RX Interval, except as follows:

A 25% jitter MAY be applied to the rate of transmission, such that the actual interval MAY be between 75% and 100% of the

advertised value. A single BFD Echo packet MAY be transmitted between normally scheduled Echo transmission intervals.

* The two paragraphs are now updated to:

When a system is using the Echo function with either Asynchronous or Demand mode, BFD Echo packets MUST NOT be transmitted when bfd.SessionState is not Up, and BFD Echo packets MUST NOT be transmitted unless the last BFD Control packet received from the remote system contains a nonzero value in Required Min Echo RX Interval.

When a system is using the Echo function with either Asynchronous or Demand mode, BFD Echo packets MAY be transmitted when bfd.SessionState is Up, and the interval between transmitted BFD Echo packets MUST NOT be less than the value advertised by the remote system in Required Min Echo RX Interval, except as follows:

A 25% jitter MAY be applied to the rate of transmission, such that the actual interval MAY be between 75% and 100% of the advertised value. A single BFD Echo packet MAY be transmitted between normally scheduled Echo transmission intervals.

3. Unaffiliated BFD Echo Procedures

As shown in Figure 1, device A supports BFD, whereas device B does not support BFD. To rapidly detect any IP forwarding faults between device A and device B, a BFD Echo session MUST be created at device A, and the BFD Echo session is RECOMMENDED to follow the BFD state machine defined in Section 6.2 of [RFC5880], except that the received state is not sent but echoed from the remote system. In this case, although BFD Echo packets are transmitted with destination UDP port 3785 as defined in [RFC5881], the BFD Echo packets sent by device A are BFD Control packets too, the looped BFD Echo packets back from device B would drive BFD state change at device A, substituting the BFD Control packets sent from the BFD peer.

Once a BFD Echo session is created at device A, it starts sending BFD Echo packets, which SHOULD include a BFD Echo session demultiplexing field, such as BFD Your Discriminator defined in [RFC5880] (BFD My Discriminator can be set to 0 to avoid confusion), except that device A can use IP source address or UDP source port to demultiplex BFD Echo session, or there is only one BFD Echo session running at device A. Device A would send BFD Echo packets with IP destination address destined for itself, such as the IP address of interface 1 of device A. All BFD Echo packets for the session MUST be sent with a Time to Live (TTL) or Hop Limit value of 255.

Considering the BFD peer wouldn't advertise Required Min Echo RX Interval as defined in [RFC5880], the transmit interval for sending BFD Echo packets MUST be provisioned at device A, how to make sure the BFD peer is willing and able to loop back BFD Echo packets sent with the provisioned transmit interval is outside the scope of this document. Considering the BFD peer wouldn't advertise Detect Mult as defined in [RFC5880], the Detect Mult for calculating the Detection Time MUST be provisioned at device A, the Detection Time in device A is equal to the provisioned Detect Mult multiplied by the provisioned transmit interval.

After receiving the BFD Echo packets sent from device A, the one-hop-away BFD peer device B immediately loops them back by normal IP forwarding, this allows device A to rapidly detect a connectivity loss to device B.

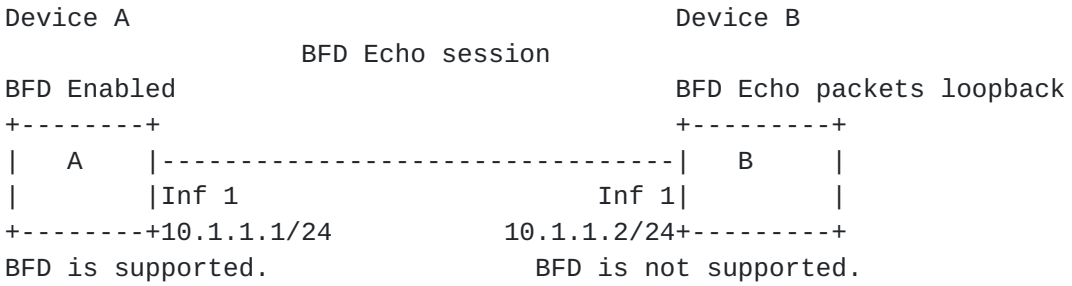

Figure 1: Unaffiliated BFD Echo deployment scenario

4. Unaffilicated BFD Echo Applicability

With the more and more application of BFD detection, there are some scenarios the BFD Echo function is deployed. And due to the different capabilities of the devices deploying BFD Echo function, it's required to apply Unaffiliated BFD Echo to the devices that couldn't afford the overhead of the full BFD protocol capability, such as the servers running virtual machines or some Internet of Things (IoT) devices. Unaffiliated BFD Echo can be used when two devices are connected and only one of them supports BFD protocol capability.

Unaffiliated BFD Echo function is reasonable and useful. Firstly, Unaffiliated BFD Echo can use BFD protocol capability at the local BFD-supported device, while using IP forwarding capability at the peer BFD-unsupported device, so Unaffiliated BFD Echo can support fast detecting and manage BFD sessions very effectively. Secondly, it is scalable when using Unaffiliated BFD Echo to adapt to different capabilities of devices.

5. Security Considerations

Unicast Reverse Path Forwarding (uRPF), as specified in [RFC3704] and [RFC8704], is a security feature that prevents the IP address spoofing attacks which is commonly used in DoS, DDoS. uRPF has two modes called strict mode and loose mode. uRPF strict mode means that the router will perform checks for all incoming packets on a certain interface: whether the router has a matching entry for the source IP in the routing table and whether the router uses the same interface to reach this source IP as where the router received this packet on. Note that the use of BFD Echo function would prevent the use of uRPF in strict mode.

6. IANA Considerations

This document has no IANA action requested.

7. Acknowledgements

The authors would like to acknowledge Ketan Talaulikar, Greg Mirsky and Santosh Pallagatti for their careful review and very helpful comments.

8. Contributors

Liu Aihua
ZTE
Email: liu.aihua@zte.com.cn

Qian Xin
ZTE
Email: qian.xin2@zte.com.cn

Zhao Yanhua
ZTE
Email: zhao.yanhua3@zte.com.cn

9. References

9.1. Normative References

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <https://www.rfc-editor.org/info/rfc5881>.

9.2. Informative References

[BBF-TR-146] Broadband Forum, "BBF Technical Report - Subscriber Sessions Issue 1", 2013, <https://www.broadband-forum.org/technical/download/TR-146.pdf>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <https://www.rfc-editor.org/info/rfc3704>.

[RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <https://www.rfc-editor.org/info/rfc8704>.

Authors' Addresses

Weiqiang Cheng
China Mobile
Beijing
CN

Email: chengweiqiang@chinamobile.com

Ruixue Wang
China Mobile
Beijing
CN

Email: wangruixue@chinamobile.com

Xiao Min
ZTE Corp.
Nanjing
CN

Email: xiao.min2@zte.com.cn

Reshad Rahman
Cisco Systems
Kanata
CA

Email: rrahman@cisco.com

Raj Chetan Boddireddy
Juniper Networks

Email: rchetan@juniper.net