
Stream: Internet Engineering Task Force (IETF)
RFC: [9747](#)
Updates: [5880](#)
Category: Standards Track
Published: March 2025
ISSN: 2070-1721
Authors: W. Cheng R. Wang X. Min, Ed. R. Rahman R. Boddireddy
China Mobile China Mobile ZTE Corp. Equinix Juniper Networks

RFC 9747

Unaffiliated Bidirectional Forwarding Detection (BFD) Echo

Abstract

This document specifies an extension to the Bidirectional Forwarding Detection (BFD) protocol that enables the use of the BFD Echo function without the need for an associated BFD control session. This "Unaffiliated BFD Echo" mechanism allows rapid detection of forwarding path failures in networks where establishing BFD control sessions is impractical or undesirable. By decoupling the Echo function from the control plane, network devices can utilize BFD's fast failure detection capabilities in a simplified manner, enhancing network resiliency and operational efficiency.

This document updates RFC 5880 by defining a new Unaffiliated BFD Echo mechanism.

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 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9747>.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	4
2. Unaffiliated BFD Echo Procedures	4
3. Updates to RFC 5880	6
4. Operational Considerations	10
5. Security Considerations	10
6. IANA Considerations	10
7. References	11
7.1. Normative References	11
7.2. Informative References	11
Acknowledgements	11
Contributors	12
Authors' Addresses	12

1. Introduction

To minimize the impact of device and link faults on services and to improve network availability in single-hop scenarios, a network device needs the capability to quickly detect communication faults with adjacent devices. Prompt detection allows for timely remedial actions to ensure service continuity.

BFD [[RFC5880](#)] provides a low-overhead, short-interval method for detecting faults on the communication path between adjacent forwarding engines, which may include interfaces, data links, and the forwarding engines themselves. BFD offers a unified mechanism to monitor any media and protocol layers in real time.

BFD defines two primary modes -- Asynchronous mode and Demand mode -- to accommodate various deployment scenarios. Additionally, it supports an Echo function that reduces the level of BFD support required in device implementations, as described in [Section 3.2](#) of [\[RFC5880\]](#). 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, as described in [Section 5](#) of [\[RFC5880\]](#) and [Section 4](#) of [\[RFC5881\]](#). If several consecutive BFD Echo packets are not received by the local system, the BFD session is declared Down.

There are two typical scenarios when using the BFD Echo function:

- Full BFD protocol capability with adjunct Echo function (Affiliated BFD Echo): This scenario requires both the local device and the adjacent device to support the full BFD protocol. This operation remains unchanged from [\[RFC5880\]](#).
- BFD Echo-Only method without full BFD protocol capability (Unaffiliated BFD Echo): This scenario requires only the local device to support sending and demultiplexing BFD Control packets. In this case, BFD Control packets are sent over the BFD Echo port, and the processing procedures for Asynchronous mode are used with the modifications specified in this document. Note that this method requires the local device to send packets with one of its own IP addresses as the destination address, upon receipt of which the adjacent device loops them back to the local device. Also note that this method monitors the connectivity to a device over a specific interface and does not verify the availability of a specific IP address at that device.

This document specifies the Unaffiliated BFD Echo scenario.

[Section 5](#) of [\[RFC5880\]](#) indicates that the payload of an Affiliated BFD Echo packet is a local matter; therefore, its contents are outside the scope of that specification. This document, however, specifies the contents of the Unaffiliated BFD Echo packet and the procedures for handling them. While this may appear to contravene [Section 5](#) of [\[RFC5880\]](#), the core behavior in that RFC states that the contents of BFD Echo packets are a local matter; this document is defining that "local matter". Regarding the selection of IP addresses, the rules stated in [Section 4](#) of [\[RFC5881\]](#) are applicable to the encapsulation of an Unaffiliated BFD Echo packet.

[Section 6.2.2](#) of [\[BBF-TR-146\]](#) describes a use case for the Unaffiliated BFD Echo.

This document updates [\[RFC5880\]](#) by defining a new method of BFD Echo-only operation which only impacts the sender of BFD Echo packets without requiring an implementation to support the BFD protocol at the loopback device, such that any IP forwarder can loop back the BFD Echo packets. It specifies the use of the Unaffiliated BFD Echo over IPv4 and IPv6 for a single IP hop. The reason why it cannot be used for multihop paths is that the Unaffiliated BFD Echo packets would be looped back by the first hop. A full description of the updates to [\[RFC5880\]](#) is provided in [Section 3](#).

1.1. Conventions Used in This Document

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. Unaffiliated BFD Echo Procedures

This section specifies the Unaffiliated BFD Echo procedures.




Figure 1: Unaffiliated BFD Echo

As shown in Figure 1, device A supports BFD, whereas device B is a regular IP forwarder that does not support BFD. Device A would send Unaffiliated BFD Echo packets, and after receiving the Unaffiliated 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. Note that device B would not intercept any received Unaffiliated BFD Echo packet or parse any BFD protocol field within the Unaffiliated BFD Echo packet.

An Unaffiliated BFD Echo session is not actually a BFD session because there is no coordination of BFD protocol state between the two link ends: the remote end does not support BFD and so cannot engage in a BFD session. The local end as an initiator may regard the Unaffiliated BFD Echo session as a BFD session from its own standpoint.

For the Unaffiliated Echo procedure, an Unaffiliated BFD Echo session is established on device A. The session **MUST** adhere to the BFD state machine specified in Section 6.2 of [RFC5880], with the exception that the received state is not derived from BFD Control packets originating from the

remote system, but rather from packets that are generated by the local system and looped back from the remote system. Consequently, the AdminDown state is not utilized in Unaffiliated BFD Echo.

BFD Control packets are transmitted and received as Unaffiliated BFD Echo packets, using UDP destination port 3785, as defined in [RFC5881]. The standard procedures for BFD Asynchronous sessions are applied to the looped BFD Control packets, including packet validation and authentication, in accordance with [RFC5880].

Once an Unaffiliated BFD Echo session is created on device A, it starts sending Unaffiliated BFD Echo packets. Unaffiliated BFD Echo packets with zeroed "Your Discriminator" field are demultiplexed to the proper session based on the source IP address or UDP source port, once the remote system loops back the local discriminator, all further received packets are demultiplexed based on the "Your Discriminator" field only, which is conformed to the procedure specified in Section 6.3 of [RFC5880]. An Unaffiliated BFD Echo packet follows the same encapsulation rules as for a BFD Echo packet as specified in Section 4 of [RFC5881]. All Unaffiliated BFD Echo packets for the session **MUST** be sent with a TTL or Hop Limit value of 255. Received packets **MUST** have a TTL or Hop Limit value of 254 (similar to Appendix A of [RFC5082] to verify against a configured number of hops); otherwise, the received packets **MUST** be dropped.

In the context of an Unaffiliated BFD Echo packet, the "Desired Min TX Interval" and "Required Min RX Interval" fields, as defined in [RFC5880], **MUST** be populated with a specific value to prevent the potential exposure of uninitialized memory. It is **RECOMMENDED** that these fields be set to a value of 1 second (1,000,000 microseconds). However, upon receipt, these values **MUST** be ignored and **MUST NOT** be used in the calculation of the Detection Time.

The "Required Min Echo RX Interval" field, as defined in [RFC5880], **MUST** be populated with a specific value to prevent the potential exposure of uninitialized memory. It is **RECOMMENDED** that this field be set to 0. However, this value **MUST** be ignored upon receipt. The transmission interval for Unaffiliated BFD Echo packets when in the Up state **MUST** be provisioned on device A.

The functionality of the Unaffiliated BFD Echo feature is dependent on device B performing IP forwarding. While this capability is typically expected to be supported on routers, it may not be enabled by default on hosts. The method for provisioning device B to loop back Unaffiliated BFD Echo packets is outside the scope of this document.

Similar to what's specified in [RFC5880], the Unaffiliated BFD Echo session begins with the periodic, slow transmission of Unaffiliated BFD Echo packets. The slow transmission rate should be no greater than one packet per second, until the session on device A is Up. After the session is Up, the provisioned transmission interval is used. When the Unaffiliated BFD Echo session on device A goes Down, the slow transmission rate is resumed. The "Detect Mult" defined in [RFC5880] **MUST** be set to a value provisioned on device A. When the `bfd.SessionState` is Up and a "Detect Mult" number of Unaffiliated BFD Echo packets have not arrived at device A as they should, the device A **MUST** set `bfd.SessionState` to Down and `bfd.LocalDiag` to 2 (Echo Function Failed)", as specified in Section 6.8.5 of [RFC5880].

In summary, the Unaffiliated BFD Echo packet reuses the format of the BFD Control packet defined in [RFC5880], and the fields within the Unaffiliated BFD Echo packet are populated as follows:

- My Discriminator: **MUST** be set to the provisioned local discriminator.
- Your Discriminator: **MUST** initially be set to 0, and then **MUST** be set to the value of "My Discriminator" looped back from the remote system.
- Desired Min TX Interval: **MUST** be set to a specific value, with a suggested value of 1 second (1,000,000 microseconds).
- Required Min RX Interval: **MUST** be set to a specific value, with a suggested value of 1 second (1,000,000 microseconds).
- Required Min Echo RX Interval: **MUST** be set to a specific value, with a suggested value of 0.
- Detect Mult: **MUST** be set to the provisioned maximum allowable number of consecutively lost Unaffiliated BFD Echo packets.

3. Updates to RFC 5880

The Unaffiliated BFD Echo described in this document reuses the BFD Echo function as described in [RFC5880] and [RFC5881], but does not require BFD Asynchronous or Demand mode. In the Unaffiliated BFD Echo operation, only the local system has the BFD protocol enabled, while the remote system simply loops back the received BFD Echo packets as ordinary data packets, without engaging in the BFD protocol.

This document updates [RFC5880] with respect to its descriptions on the BFD Echo function as follows.

The 4th paragraph of Section 3.2 of [RFC5880] is updated as below:

OLD TEXT

An adjunct to both modes is the Echo function.

NEW TEXT

An adjunct to both modes is the Echo function, which can also be running independently.

OLD TEXT

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).

NEW TEXT

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.

The 3rd and 9th paragraphs of [Section 6.1](#) of [\[RFC5880\]](#) are updated as below:

OLD TEXT

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.

NEW TEXT

When a system is running with Asynchronous or Demand 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 for Asynchronous mode or eliminated completely for Demand mode when the Echo function is active.

OLD TEXT

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.

NEW TEXT

In Asynchronous mode or Demand 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.

The 2nd paragraph of [Section 6.4](#) of [\[RFC5880\]](#) is updated as below:

OLD TEXT

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).

NEW TEXT

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). Note that a system operating in Demand mode would direct the remote system to cease the periodic transmission of BFD Control packets, by setting the Demand (D) bit in its BFD Control packets.

The 2nd paragraph of [Section 6.8](#) of [\[RFC5880\]](#) is updated as below:

OLD TEXT

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.

NEW TEXT

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.

The 7th paragraph of [Section 6.8.3](#) of [\[RFC5880\]](#) is updated as below:

OLD TEXT

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.

NEW TEXT

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. A system operating in Demand mode would not receive BFD Control traffic.

The 1st and 2nd paragraphs of [Section 6.8.9](#) of [RFC5880] are updated as below:

OLD TEXT

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.

NEW TEXT

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.

OLD TEXT

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...

NEW TEXT

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...

4. Operational Considerations

All operational considerations from [RFC5880] apply. Since this mechanism leverages existing BFD machinery, particularly periodic pacing of traffic based on configuration, there's no real possibility to create congestion. Moreover, creating congestion would be counterproductive to check the bidirectional connectivity.

Some devices that would benefit from the use of BFD may be unable to support the full BFD protocol. Examples of such devices include servers running virtual machines, or Internet of Things (IoT) devices. By using Unaffiliated BFD Echo, these devices only need to support a basic loopback function.

As specified in Section 2 of this document, some configuration is needed to make the Unaffiliated BFD Echo work, although the configuration won't go beyond the scope of [RFC5880]. At a BFD-enabled local system, the Unaffiliated BFD Echo session can coexist with another type of BFD session. In that scenario, the remote system for the Unaffiliated BFD Echo session must be different from the remote system for the other type of BFD session, and the local system's discriminators for different BFD sessions must be different. At the same time, it's not necessary for the local system to differentiate the Unaffiliated BFD Echo session from the other type of BFD session.

5. Security Considerations

All security considerations from [RFC5880] and [RFC5881] apply.

Unaffiliated BFD Echo requires the remote device to loop Unaffiliated BFD Echo packets. In order to provide this service, the remote device cannot make use of Unicast Strict Reverse Path Forwarding (RPF) [RFC3704], otherwise the Unaffiliated BFD Echo packets might not pass the RPF check at the remote device.

As described in Section 5 of [RFC5880], BFD Echo packets may be spoofed. Specifically for Unaffiliated BFD Echo, a DoS attacker may send spoofed Unaffiliated BFD Echo packets to the loopback device, so some form of authentication **SHOULD** be included. Considering the Unaffiliated BFD Echo packets in this document are also BFD Control packets, the "Authentication Section" as defined in [RFC5880] for a BFD Control packet is **RECOMMENDED** to be included within the Unaffiliated BFD Echo packet.

As stated in Section 2, in order to avoid unset values being a potential vector for disclosure of uninitialized memory, all fields of the Unaffiliated BFD Echo packet **MUST** be populated with a certain value, even if some of the fields are ignored on receipt.

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

7.2. Informative References

- [BBF-TR-146] Broadband Forum, "TR-146: Subscriber Sessions", Broadband Forum Technical Report, TR-146, Issue 1, May 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>>.
- [RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<https://www.rfc-editor.org/info/rfc5082>>.

Acknowledgements

The authors would like to acknowledge Ketan Talaulikar, Greg Mirsky, Santosh Pallagatti, Aijun Wang, Éric Vyncke, Adrian Farrel, Tim Wicinski, Dhruv Dhody, Stephen Farrell, Gunter Van de Velde, Gyan Mishra, Brian Trammell, Gorry Fairhurst, Mahesh Jethanandani, John Scudder, Murray Kucherawy, and Zaheduzzaman Sarker for their careful reviews and very helpful comments.

The authors would like to acknowledge Jeff Haas for his guidance, insightful review, and very helpful comments.

The authors would like to acknowledge Erik Auerswald for his insightful comments during the discussion of this document.

The authors would like to acknowledge Detao Zhao for the very helpful discussion.

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

Authors' Addresses

Weiqiang Cheng

China Mobile

Beijing

China

Email: chengweiqiang@chinamobile.com

Ruixue Wang

China Mobile

Beijing

China

Email: wangruixue@chinamobile.com

Xiao Min (EDITOR)

ZTE Corp.

Nanjing

China

Email: xiao.min2@zte.com.cn

Reshad Rahman

Equinix

Ottawa

Canada

Email: reshad@yahoo.com

Raj Chetan Boddireddy

Juniper Networks

Email: rchetan@juniper.net