BFD Working Group
Internet-Draft
Intended status: Informational
Expires: January 14, 2021

W. Cheng
R. Wang
China Mobile
X. Min
A. Liu
ZTE Corp.
R. Rahman
Cisco Systems
July 13, 2020

Unaffiliated BFD Echo Function
draft-cw-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 BFD echo 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 January 14, 2021.

Copyright Notice

Copyright (c) 2020 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 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

To minimize the impact of device 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 path between adjacent forwarding engines. The faults can be interface, data link, and even forwarding engine faults. 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, also supports echo function to reduce the device requirement for BFD. When the echo function is activated, the local system sends a BFD echo packet and the remote system loops back the packet through the forwarding path. If several consecutive echo packets are not received, the session is declared to be Down.

When using BFD echo function, it is not clear whether the devices using echo function need to support the full BFD procotol, including maintaining the state machine of BFD session as described in [RFC5880] and [RFC5881]. According to different understanding, there are two typical scenarios as below:

1. Full BFD procotol capability with affiliated echo function: this scenario requires both the local device and the neighboring device to support BFD protocol.

2. Only BFD echo function without full BFD procotol capability: this scenario requires only the local device to support sending BFD packets.

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

Unaffiliated BFD echo function described in this document reuses the BFD echo function as described in [RFC5880] and [RFC5881], but independent of BFD asynchronous mode, that means it doesn't need BFD protocol capability of state machine, but only BFD echo function to a deployed device supporting BFD detection. When using unaffiliated BFD echo function, just the local device works on BFD protocol and the BFD peer doesn't, which only loopback the received BFD echo packets as usual data packets without enabling BFD protocol.

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.

2. Unaffiliated BFD Echo Behavior

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 capablity, 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. A BFD echo session can be established at the device that supports BFD, and the device will send the BFD echo packets with the IP address destined for itself, whereas the other peer device just loopback the received BFD echo packets.

After receiving a BFD echo packet, the device that does not support BFD protocol immediately loops back the packet by normal IP forwarding, implementing quick link failure detection. As shown in Figure 1, device A supports BFD, whereas device B does not support BFD. To rapidly detect any faults with the IP link between device A and device B, a BFD echo session can be provisioned and created at device A, and device A starts sending BFD echo packets, which should include a BFD echo session demultiplexing field, such as BFD discriminator defined in [RFC5880]. After receiving the BFD echo packets sent from device A, device B immediately loops back them, this allows device A to rapidly detect a connectivity loss to device B.

```
   Device A                                   Device B
                      BFD echo session
   BFD Enabled                                BFD Echo packets loopback
   +--------+                                 +---------+
   |   A    |---------------------------------|    B    |
   |        |Inf 1                       Inf 1|         |
   +--------+10.1.1.1/24          10.1.1.2/24+---------+
   BFD is supported.                BFD is not supported.
```

Figure 1: Unaffiliated BFD Echo deployment scenario

3. Discussion

Unaffiliated BFD echo function is reasonable and useful. Firstly, unaffiliated BFD echo can use BFD protocol capability in the local BFD-supported device, while using IP forwarding capability in the peer non-BFD-supported 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.

4. 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.

5. IANA Considerations

This document has no IANA action requested.

6. Acknowledgements

TBD.

7. References

7.1. 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>.

7.2. 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

Aihua Liu
ZTE Corp.
Shenzhen
CN

Email: liu.aihua@zte.com.cn

Reshad Rahman
Cisco Systems
Kanata
CA

Email: rrahman@cisco.com

BFD Working Group X. Min
Internet-Draft ZTE Corp.
Intended status: Standards Track February 28, 2020
Expires: August 31, 2020

Bidirectional Forwarding Detection (BFD) Padding Alteration
draft-xiao-bfd-padding-alteration-00

Abstract

This document describes the procedures of the Bidirectional Forwarding Detection (BFD) protocol to alter BFD padding, using BFD in Asynchronous mode.

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 August 31, 2020.

Copyright Notice

Copyright (c) 2020 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 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

"BFD Encapsulated in Large Packets" [I-D.ietf-bfd-large-packets] defines a new BFD variable bfd.PaddedPduSize, which can be configured to the value the BFD transport protocol payload size is increased to, and the contents of this additional payload MUST be zero. In other words, the BFD protocol can be used to verify path MTU with BFD padding.

Considering that there are implementations that do not support Large BFD Packets, [I-D.ietf-bfd-large-packets] specifies that when an implementation is incapable of processing Large BFD Packets, the BFD state machine will behave as if it were not receiving Control packets. The implementation would follow standard BFD procedures with regards to not having received Control packets.

"Bidirectional Forwarding Detection (BFD)" [RFC5880] specifies the Poll Sequence as an exchange of BFD Control packets that is used in some circumstances to ensure that the remote system is aware of parameter changes.

This document describes the procedures of the Bidirectional Forwarding Detection (BFD) protocol to alter BFD padding, using BFD in Asynchronous mode. The procedure defined in this document refers to but differs from the Poll Sequence defined in Section 6.5 of [RFC5880].

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. Turn On BFD Padding

There are two scenarios when turning on BFD padding. One is to turn on BFD padding when the BFD session is Up and running with the Small BFD Packets. Another one is to turn on BFD padding before the BFD session is Up.

2.1. Switch From the Small BFD Packets

Consider a scenario when the BFD session is Up and running with the Small BFD Packets, which don't include any padding. If the BFD variable bfd.PaddedPduSize is set to a new non-zero value, i.e., switch from the Small BFD Packets to the Large BFD Packets whose size equals to the provisioned new value, then the procedure described in this section is applied, and the procedure is called Padding Poll Sequence.

A Padding Poll Sequence consists of a system sending periodic Small BFD Control packets and additional Padding Poll, namely Large BFD Control packets with the Poll (P) bit set, which is used to notify the BFD peer system, that it intends to switch from the current Small BFD Packets, to the Large BFD Packets whose size equals to the provisioned new value. When the BFD peer system receives a Padding Poll, it immediately responds a BFD Control packet with the Final (F) bit set, independent of any periodic BFD Control packets it is sending. When the system sending the Padding Poll sequence receives a replied packet with Final, the Padding Poll Sequence is considered to be successfully terminated. Typically, the successfully terminated sequence consists of a single packet in each direction, but considering packet losses or relatively long packet latencies, multiple Padding Poll packets may be sent before the sequence successfully terminates, the maximum number of Padding Poll packets and the interval between sequential Padding Poll packets are implementation specific local matters. When the system sending the Padding Poll sequence doesn't receive any replied packet with Final, over the whole process of sending multiple Poll packets, the Padding Poll Sequence is considered to be unsuccessfully terminated.

If a Padding Poll Sequence is successfully terminated, the system requesting the Padding Poll Sequence would switch to the periodic transmission of Large BFD Packets whose size equals to the provisioned new value, with the Poll (P) bit clear.

If a Padding Poll Sequence is unsuccessfully terminated, the system requesting the Padding Poll Sequence would report an error to its management, and continue the periodic transmission of Small BFD Control packets with the Poll (P) bit clear.

2.2. Switch From the Start of a BFD Session

If the BFD variable bfd.PaddedPduSize is set to a non-zero value indicating padding before the BFD session is Up, then the procedure described in this section is applied.

At first the BFD session would be Up and running with the Small BFD Packets which don't include any padding, following the BFD session establishing procedures defined in [RFC5880], and then switch from the Small BFD Packets to the Large BFD Packets, so that the Padding Poll Sequence described in Section 2.1 is applied.

If BFD variable bfd.PaddedPduSize is set on both sides of the tested path, then the Padding Poll Sequence would happen on both directions separately.

3. Increase BFD Padding

While the BFD session is Up and running with the Large BFD Packets which include padding, if the BFD variable bfd.PaddedPduSize is set to a new value indicating larger padding, i.e., switch from the Large BFD Packets to the Large BFD Packets whose size equals to the provisioned new value, then the Padding Poll Sequence described in Section 2.1 is applied. Specifically, Padding Poll would use Larger BFD Packets with the Poll (P) bit set.

4. Decrease BFD Padding

While the BFD session is Up and running with the Large BFD Packets which include padding, if the BFD variable bfd.PaddedPduSize is set to a new value indicating smaller padding, i.e., switch from the Larger BFD Packets to the Large BFD Packets whose size equals to the provisioned new value, then the Poll Sequence described in Section 6.5 of [RFC5880] is applied. Specifically, the Poll Sequence MUST be performed by setting the Poll (P) bit on those scheduled periodic transmissions, and additional packets MUST NOT be sent.

5. Turn Off BFD Padding

While the BFD session is Up and running with the Large BFD Packets which include padding, if the BFD padding is turned off, i.e., switch from the Large BFD Packets to the Small BFD Packets, then the Poll Sequence described in Section 6.5 of [RFC5880] is also applied.

6. Security Considerations

This document does not raise any additional security issues beyond those of the specifications referred to in the list of normative references.

7. IANA Considerations

This document has no IANA action requested.

8. Acknowledgements

TBA.

9. Normative References

[I-D.ietf-bfd-large-packets]
Haas, J. and A. Fu, "BFD Encapsulated in Large Packets", draft-ietf-bfd-large-packets-02 (work in progress), November 2019.

[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>.

[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>.

Author's Address

Xiao Min
ZTE Corp.
Nanjing
China

Phone: +86 18061680168
Email: xiao.min2@zte.com.cn