Source Packet Routing in Networking
Internet-Draft
Intended status: Standards Track
Expires: July 24, 2021

M. Chen
X. Chen
Huawei

January 20, 2021

Segment Routing Policy for Unaffiliated BFD Echo Function
draft-chen-spring-sr-policy-for-ubfd-00

Abstract

This document describes how to leverage Segment Routing (SR) Policy to make sure that the Unaffiliated BFD (U-BFD) Echo packets must be transmitted to the remote system before being looped back to the local system. This enables that U-BFD works not only for one hop scenario but for multiple hops scenario as well.

In addition, this document also defines a way to explicitly specify the loop back path of the Echo packets. This is useful in the case where the forward and reverse path of the Echo packets are required to follow the same path.

Requirements Language

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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 July 24, 2021.

Table of Contents

1. Introduction

BFD Echo function was originally defined in [RFC5880] and [RFC5881], where the remote system is required to loop the BFD Echo packets back to the local system. To support BFD Echo Function, some negotiations between the local system and remote system are needed, and both the local and remote system need to maintain the BFD session state.

Unaffiliated BFD Echo Function (U-BFD) is defined in [I-D.ietf-bfd-unaffiliated-echo]. Where the destination IP address of the BFD Echo packets is set to one of the IP addresses of the local system. Therefore, the Echo packets can be automatically looped back (through normal IP forwarding) by the remote system to the local system. With U-BFD, the remote system does not need to support any BFD related functions and maintain any session states. This further simplifies the BFD Echo Function process at the remote system hence increases the saleability.

But, the U-BFD works when there is only one hop between the local system and remote system. Otherwise, the Echo packets will be

prematurely looped back by an intermediate node to the local system, therefore the Echo packets will not be transmitted to the remote system. This may result in false negative issue. Take the following figure (Figure 1) as an example, if the U-BFD is expected to monitor the path between node A and node C, node A (as the local system) sets the destination IP to itself and sends the Echo packets to node B. Since node B has the route to node A, the Echo packets will be directly forwarded back to node A. If there is a failure on the path between node B and node C, obviously, the U-BFD session cannot detect it.

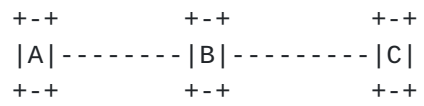

Figure 1, Multi-hop Scenario

In addition, in some scenarios, for example, mobile backhaul network, where the forward and reverse direction of a path are required to along the same path. When apply BFD in mobile backhaul network, it also expects that the BFD control packets in both directions follow the same path, otherwise, it may result in false positive issue. Take the following figure (Figure 2) as an example, there are two paths (A-B-C, A-D-C) between node A and node C. Assuming that it expects to monitor the path A-B-C by using BFD, where node A is the local system and node C is the remote system. If node C chooses path C-D-A to send the control packets, when a failure occurs on path C-D-A, node A (the local system) will not receive the BFD packets and hence consider that path A-B-C is failed. But actually path A-B-C is working.

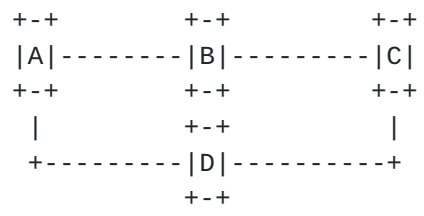

Figure 2, Multi-hop, Multi-path Scenario

To solve the above issues, there needs a way to make sure that U-BFD Echo packets must be transmitted to the remote system before being looped back. And when looping back the U-BFD Echo packets, the remote system should send the U-BFD Echo packets along a specified path.

Since Segment Routing (SR) Policy [I-D.ietf-spring-segment-routing-policy] allows a headend node to steer a packet flow along any path. This document leverages the SR

policy to make sure that the U-BFD Echo packets must be transmitted to the remote system before being looped back. This enables that U-BFD Echo Function works not only for one hop scenario but for multiple hops scenario. By using SR policy, the loop back path of the Echo packets can be specified as well. This is useful in the case where the forward and loop back path of the Echo packets are required to follow the same path.

2. SR Policy for U-BFD

As defined in [I-D.ietf-spring-segment-routing-policy], an SR Policy is identified through the tuple <headend, color, endpoint>. To support U-BFD, the endpoint of the Policy MUST be set to the same as the headend that is specified as an IPv4 or IPv6 address of the headend node. To make sure that the U-BFD Echo packets must be transmitted to the remote system, the SR Policy MUST have a candidate path that is associated with a Segment-List. The Segment-List MUST include a SID that identifies the remote system. To specify the loop back path, a series of SIDs or a Binding SID (BSID) that is associated with the loop back path MUST be included in the Segment-List.

Based on the topology in Figure 2, below are some examples that show how the SR Polices for U-BFD can be instantiated at the headend node.

1. The forward direction forwarding is based on the SR Policy, the loop back direction forwarding is based on IP forwarding. This way, only the SIDs of the forward path should be included in the Segment-List, and at least the SID of the remote system should be included. The SR Policy will make sure that the U-BFD Echo packets are transmitted to the remote system. When receives the Echo packets, the remote system will decapsulate the Echo packets and then forward them back to the local system according to IP forwarding.

```
SR policy POL1 <headend = A, color = 1, endpoint = A>
  Candidate-path CP1 <protocol-origin = 20, originator =
                      100:1.1.1.1, discriminator = 1>
    Preference 200
    Weight W1,
    SID-List <B,C>
```

2. Both the forward and loop back direction forwarding are based on the SR Policy. A BSID (carried in the Segment-List) is used to identify the loop back path <B, A>. Using BSID can decrease the SID stack depth.

```
SR policy POL2 <headend = A, color = 1, endpoint = A>
  Candidate-path CP1 <protocol-origin = 20, originator =
                      100:1.1.1.1, discriminator = 2>
    Preference 200
    Weight W1,
    SID-List <B, C, BSID>
```

3. Both the forward and loop back direction forwarding are based on the SR Policy. The forward path <B, C> and loop back path <B, A> are explicitly included in the Segment-List. This is suitable for the case where the hops of forward and loop back path are not too much, the ingress (local system) has the capability to handle the whole SIDs of the forward and loop back path.

```
SR policy POL2 <headend = A, color = 1, endpoint = A>
  Candidate-path CP1 <protocol-origin = 20, originator =
                      100:1.1.1.1, discriminator = 3>
    Preference 200
    Weight W1,
    SID-List <B, C, B, A>
```

3. IANA Considerations

This document makes no request of IANA.

4. Security Considerations

This document does not introduce additional security requirements and mechanisms other than the ones described in [I-D.ietf-bfd-unaffiliated-echo] and [I-D.ietf-spring-segment-routing-policy].

5. Acknowledgements

6. Contributors

The following people have substantially contributed to this document:

Pingwei Fan
Huawei

EMail: fanpingwei@huawei.com

7. References

7.1. Normative References

[I-D.ietf-bfd-unaffiliated-echo]
Cheng, W., Wang, R., Min, X., Rahman, R., and R. Boddireddy, "Unaffiliated BFD Echo Function", draft-ietf-bfd-unaffiliated-echo-01 (work in progress), November 2020.

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-09 (work in progress), November 2020.

[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>.

[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

[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>.

Authors' Addresses

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Xinjun Chen
Huawei

Email: ifocus.chen@huawei.com