Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 9, 2014

N. Akiya
C. Pignataro
N. Kumar
Cisco Systems

June 7, 2014

Seamless Bidirectional Forwarding Detection (BFD) for Segment Routing (SR)
draft-akiya-bfd-seamless-sr-02

Abstract

Note: this document needs to be updated to align with changes in the S-BFD base document.

This specification defines procedures to use Seamless Bidirectional Forwarding Detection (S-BFD) in a Segment Routing (SR) based environment.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 http://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 December 9, 2014.

Copyright Notice

Table of Contents

1. Introduction

One application for Seamless Bidirectional Forwarding Detection (S-BFD) [I-D.akiya-bfd-seamless-base] is to perform full reachability validations, partial reachability validations and adjacency segment ID verifications on a Segment Routing (SR) based environment.

This specification defines procedures to use Seamless BFD in a SR based environment.

2. BFD Target Identifier Types

BFD target identifier type of value 2 is used for SR. Note that BFD target identifier type of value 2, which specifies segment routing node segment ID, is not tied to a specific routing protocol. If definitions and procedures need routing protocol specifics, then IGP specific SR types will be defined.

3. Reserved BFD Discriminators

With SR technology, BFD target identifier type 2 is used. Node segment IDs are used as BFD discriminators. BFD discriminator values corresponding to all or subset of local node segment IDs are to be allocated from the discriminator pool for Seamless BFD.

Example:

- BFD Target Identifier Type 2: Node segment ID 0x03E9A0FF maps to BFD discriminator 0x03E9A0FF.

4. BFD Target Identifier Table

With SR BFD target identifier type, only locally reserved BFD discriminators and corresponding information are to be in this table. No inter-node communications are needed to exchange BFD discriminator and BFD target identifier mappings.

5. Full Reachability Validations

5.1. Initiator Behavior

Any SR network node can attempt to perform a full reachability validation to any BFD target identifier of type 2 (node segment ID) on other network nodes, as long as destination BFD target identifier is provisioned to use this mechanism. Transmitted BFD control packet by the initiator is to have "your discriminator" corresponding to destination BFD target identifier of type 2.

Initiator is to use following procedures to construct BFD control packets to perform SR full reachability validations:

- MUST set "your discriminator" to target node segment ID.
- MUST use explicit label switching packet format described in [I-D.akiya-bfd-seamless-base].

5.2. Responder Behavior

To respond to received BFD control packet which was targeted to local BFD target identifier of type 2 (Segment Routing Node Segment ID), response BFD control packet is targeted to IP address taken from received "source IP address". Responder MUST validate obtained IP address is in valid format (ex: not Martian address). Responder MUST consult local routing table to ensure obtained IP address is reachable. Responder MAY impose node segment ID, corresponding to obtained IP address, on the response BFD control packet.

6. Partial Reachability Validations

Procedures described in [I-D.akiya-bfd-seamless-base] applies.

7. MPLS Label Verifications

With target identifier type 2, SR based, when a network node wants to test an adjacency segment ID, then adjacency segment ID (label value + EXP) being tested is encoded as lower 23 bits of localhost IP destination address. When passive BFD session receives a SR BFD control packet with lower 23 bits of IP destination address non-zero, then response will contain adjacency segment ID (label value + EXP) corresponding to incoming interface as lower 23 bits of localhost IP destination address.

Simple ASCII art is provided to illustrate the MPLS label verification concept on a SR network.

```
            md=50/yd=R3/DIP=127...R2R3
Active  [1] - - - - - - - - - - - - - - - >  Passive
BFD     < - - - - - - - - - - - - - - - [2]   BFD
Session     md=R3/yd=50/DIP=127...R3R2      Session

                            (adj SID R2R3)->
  R1 ------------------ R2 ------------------ R3
                            <-(adj SID R3R2)
```

If a response BFD control packet is received, then initiator can conclude that a packet has reached intended node correctly. With information embedded in last 23 bits of response BFD control packet from responder, initiator has the ability to perform further verifications on how responded node received BFD control packet.

8. Provisioning Active BFD Sessions for SR Networks

Many factors will influence how to provision active BFD sessions on which network nodes. This section provides some provisioning suggestions of active BFD sessions on SR networks. However, they are only suggestions. Less provisioning of active BFD sessions may be required in some cases, or further active BFD sessions may be required in other cases.

Traffic engineered segment routing

- Segment routing eliminates hop-by-hop signaling to create traffic engineered paths, as described in [I-D.previdi-filsfils-isis-segment-routing]. When traffic engineered segment routing path is instantiated on an ingress

node, with stack of segment IDs, absence of hop-by-hop signaling results in less confidence in reachability to egress as well as traversal of strictly routed segments. S-BFD can perform rapid verification of both prior to allowing service over instantiated traffic engineered segment routing paths. In addition, S-BFD can provide continuity check on both aspects, as detection time and coverage of S-BFD is much superior than IGP failure detection and convergence time.

Single node segment ID data forwarding

- In order to protect all data passing through local network using single node segment ID, active BFD sessions can be instantiated on each network node to verify full reachability to all node segment IDs.

Centralized controller initiated S-BFD

- Centralized controller based segment routing network monitoring techniques, such as the one described in [I-D.geib-spring-oam-usecase], are powerful. One aspect that is lacking from such techniques is the guarantee that monitor packet did indeed reach certain network node (i.e. u-turned at expected network node). Related aspect is the lack of guarantee that monitor packet over adjacency segment ID did indeed result in traversal of expected adjacency. Since S-BFD can fill in the missing holes, also running S-BFD in parallel from the central controller device will even strengthen the technique.

9. Security Considerations

Security considerations for BFD are discussed in [RFC5880] and security considerations for S-BFD are discussed in [I-D.akiya-bfd-seamless-base].

10. IANA Considerations

None

11. Acknowledgements

Authors would like to thank Marc Binderberger from Cisco Systems for providing valuable comments.

12. Contributing Authors

Dave Ward
Cisco Systems
Email: wardd@cisco.com

Tarek Saad
Cisco Systems
Email: tsaad@cisco.com

Siva Sivabalan
Cisco Systems
Email: msiva@cisco.com

13. References

13.1. Normative References

[I-D.akiya-bfd-seamless-base]
Akiya, N., Pignataro, C., Ward, D., Bhatia, M., and J. Networks, "Seamless Bidirectional Forwarding Detection (S-BFD)", draft-akiya-bfd-seamless-base-03 (work in progress), April 2014.

[I-D.previdi-filsfils-isis-segment-routing]
Previdi, S., Filsfils, C., Bashandy, A., Horneffer, M., Decraene, B., Litkowski, S., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., and J. Tantsura, "Segment Routing with IS-IS Routing Protocol", draft-previdi-filsfils-isis-segment-routing-02 (work in progress), March 2013.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

13.2. Informative References

[I-D.geib-spring-oam-usecase]
Geib, R. and C. Filsfils, "Use case for a scalable and topology aware MPLS data plane monitoring system", draft-geib-spring-oam-usecase-01 (work in progress), February 2014.

Authors' Addresses

Nobo Akiya
Cisco Systems

Email: nobo@cisco.com

Carlos Pignataro
Cisco Systems

Email: cpignata@cisco.com

Nagendra Kumar
Cisco Systems

Email: naikumar@cisco.com