Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 09, 2013

N. Akiya
C. Pignataro
N. Kumar
Cisco Systems

June 07, 2013

Seamless Bidirectional Forwarding Detection (BFD) for Segment Routing (SR)
draft-akiya-bfd-seamless-sr-00

Abstract

This specification defines procedures to use Seamless Bidirectional Forwarding Detection (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 09, 2013.

Copyright Notice

Table of Contents

1. Introduction

One application for Seamless Bidirectional Forwarding Detection (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. BFD discriminator values corresponding to all or subset of local node segment IDs are to be reserved on corresponding network node. Node segment IDs are used as BFD discriminators. Corresponding BFD discriminators MUST be reserved and those BFD discriminators MUST NOT be used for other BFD sessions.

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

- SR TE LSP has path-protection and no local repairs on transit nodes: Active BFD sessions should be instantiated on the LSP ingress. Instantiated active BFD sessions should perform full

reachability validation to all node segment IDs that are immediate nexthop of all adjacency segment IDs used in the LSP. This verifies that strict switching based on adjacency segment IDs is being switched to correct downstream node segment. If multiple links exist on one or more of adjacency points being validated, MPLS label verification technique should also be provisioned to ensure correct link is being traversed. Lastly, full reachability validation should be performed from LSP ingress to LSP egress to verify end-to-end reachability. Fate of the LSP is tied to all active BFD sessions instantiated on LSP ingress.

- o SR TE LSP has local repairs on transit nodes: Active BFD sessions should be instantiated on each local repair points, using combination of full reachability validation technique and MPLS label verification technique. These active sessions are programmed to be one of the triggers of local repair procedures. Lastly, full reachability validation should be performed from LSP ingress to LSP egress to verify end-to-end reachability, but this should be provisioned with more relaxed failure detection count than other active BFD sessions instantiated on transit repair points. Fate of the LSP is tied only to the active BFD session verifying end-to-end reachability on LSP ingress.

Single node segment ID data forwarding

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

- o Additionally, it may be beneficial to provision active BFD sessions on other network nodes (non-edge) for local repair purposes. These network nodes can also instantiate active BFD sessions to desired identifier (edge or non-edge).

9. Security Considerations

Same security considerations as [RFC5880], [RFC5881], [RFC5883], [RFC5884], [RFC5885] and [I-D.akiya-bfd-seamless-base] apply to this document.

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.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.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

13.2. Informative References

[I-D.ietf-bfd-on-lags]
Bhatia, M., Chen, M., Boutros, S., Binderberger, M., and J. Haas, "Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces", draft-ietf-bfd-on-lags-00 (work in progress), May 2013.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

[RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.

[RFC5885] Nadeau, T. and C. Pignataro, "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.

[RFC6428] Allan, D., Swallow Ed. , G., and J. Drake Ed. , "Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile", RFC 6428, November 2011.

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