

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: November 5, 2016

N. Akiya
Big Switch Networks
C. Pignataro
D. Ward
Cisco
May 4, 2016

**Seamless Bidirectional Forwarding Detection (S-BFD) for
IPv4, IPv6 and MPLS
draft-ietf-bfd-seamless-ip-05**

Abstract

This document defines procedures to use Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6 and MPLS environments.

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 November 5, 2016.

Copyright Notice

Copyright (c) 2016 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 (<http://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	2
2.	S-BFD UDP Port	2
3.	S-BFD Echo UDP Port	3
4.	S-BFD Control Packet Demultiplexing	3
5.	Initiator Procedures	3
5.1.	Details of S-BFD Control Packet Sent by SBFDDInitiator . .	4
5.1.1.	Target vs. Remote Entity (S-BFD Discriminator) . . .	4
6.	Responder Procedures	5
6.1.	Details of S-BFD Control Packet Sent by SBFDDReflector . .	5
7.	Security Considerations	5
8.	IANA Considerations	6
9.	Acknowledgements	6
10.	Contributors	6
11.	References	7
11.1.	Normative References	7
11.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Seamless Bidirectional Forwarding Detection (S-BFD), [[I-D.ietf-bfd-seamless-base](#)], defines a generalized mechanism to allow network nodes to seamlessly perform continuity checks to remote entities. This document defines necessary procedures to use S-BFD on IPv4, IPv6 and MPLS environments.

The reader is expected to be familiar with the IP [[RFC0791](#)] [[RFC2460](#)], BFD [[RFC5880](#)], MPLS BFD [[RFC5884](#)], and S-BFD [[I-D.ietf-bfd-seamless-base](#)] terminologies and protocol constructs.

[2.](#) S-BFD UDP Port

A new UDP port is defined for the use of the S-BFD on IPv4, IPv6 and MPLS environments: 7784.

On S-BFD control packets from the SBFDDInitiator to the SBFDDReflector, the SBFDDReflector session MUST listen for incoming S-BFD control packets on the port 7784. SBFDDInitiator sessions MUST transmit S-BFD control packets with destination port 7784. The source port of the

S-BFD control packets transmitted by SBFDInitiator sessions can be any but MUST NOT be 7784. The same UDP source port number MUST be used for all S-BFD control packets associated with a particular SBFDInitiator session. The source port number is unique among all SBFDInitiator sessions on the system.

On S-BFD control packets from the SBFDReflector to the SBFDInitiator, the SBFDInitiator MUST listen for reflected S-BFD control packets at its source port.

3. S-BFD Echo UDP Port

The BFD Echo port defined by [\[RFC5881\]](#), port 3785, is used for the S-BFD Echo function on IPv4, IPv6 and MPLS environments. SBFDInitiator sessions MUST transmit S-BFD echo packets with destination port 3785. The setting of the UDP source port [\[RFC5881\]](#) and the procedures [\[I-D.ietf-bfd-seamless-base\]](#) for the S-BFD Echo function are outside the scope of this document.

4. S-BFD Control Packet Demultiplexing

The S-BFD Control Packet demultiplexing follows the procedure specified in Section 7.1. of [\[I-D.ietf-bfd-seamless-base\]](#). Received S-BFD control packet MUST be demultiplexed with the destination UDP port field.

This procedure for an S-BFD packet is executed on both the initiator and the reflector. If the port is 7784 (i.e., S-BFD packet for S-BFDReflector), then the packet MUST be looked up to locate a corresponding SBFDReflector session based on the value from the "your discriminator" field in the table describing S-BFD discriminators. If the port is not 7784, but the packet is demultiplexed to be for an SBFDInitiator, then the packet MUST be looked up to locate a corresponding SBFDInitiator session based on the value from the "your discriminator" field in the table describing BFD discriminators. In that case, then the destination IP address of the packet SHOULD be validated to be for itself. If the packet demultiplexes to a classical BFD session, then the procedures from [\[RFC5880\]](#) apply.

5. Initiator Procedures

S-BFD control packets are transmitted with IP header, UDP header and BFD control header ([\[RFC5880\]](#)). When S-BFD control packets are explicitly label switched (i.e. not IP routed which happen to go over an LSP, but explicitly sent on a specific LSP), the former is prepended with a label stack. Note that this document does not make a distinction between a single-hop S-BFD scenario and a multi-hop S-BFD scenario, both scenarios are supported.

Necessary values in the BFD control headers are described in [\[I-D.ietf-bfd-seamless-base\]](#). [Section 5.1](#) describes necessary values in the MPLS header, IP header and UDP header when an SBFDDiscriminator on the initiator is sending S-BFD control packets.

5.1. Details of S-BFD Control Packet Sent by SBFDDiscriminator

- o Specifications common to both IP routed S-BFD control packets and explicitly label switched S-BFD control packets:
 - * Source IP address field of the IP header MUST be set to a local IP address that is expected to be routable by the target (i.e. not IPv6 link-local address when the target is multiple hops away).
 - * UDP destination port MUST be set to a well-known UDP destination port assigned for S-BFD: 7784.
 - * UDP source port MUST NOT be set to 7784.
- o Specifications for IP routed S-BFD control packets:
 - * Destination IP address field of the IP header MUST set to an IP address of the target.
 - * TTL field of the IP header SHOULD be set to 255.
- o Specifications for explicitly label switched S-BFD control packets:
 - * S-BFD control packets MUST have the label stack that is expected to reach the target.
 - * TTL field of the top most label SHOULD be 255.
 - * The destination IP address MUST be chosen from the 127/8 range for IPv4 and from the 0:0:0:0:0:FFFF:7F00/104 range for IPv6, as with [\[RFC5884\]](#).
 - * TTL field of the IP header MUST be set to 1.

5.1.1. Target vs. Remote Entity (S-BFD Discriminator)

Typically, an S-BFD control packet will have "your discriminator" field corresponding to an S-BFD discriminator of the remote entity located on the target network node defined by the destination IP address or the label stack. It is, however, possible for an SBFDDiscriminator to carefully set the "your discriminator" and TTL

fields to perform a continuity test in the direction towards a target, but destined to a transit network node and not to the target itself.

[Section 5.1](#) intentionally uses the word "target", instead of "remote entity", to accommodate this possible S-BFD usage through TTL expiry. This also requires S-BFD control packets not be dropped by the responder node due to TTL expiry. Thus implementations on the responder MUST allow received S-BFD control packets taking TTL expiry exception path to reach corresponding reflector BFD session.

6. Responder Procedures

S-BFD control packets are IP routed back to the initiator, and will have IP header, UDP header and BFD control header. If an SBFDRReflector receives an S-BFD control packet with UDP source port as 7784, the packet MUST be discarded. Necessary values in the BFD control header are described in [[I-D.ietf-bfd-seamless-base](#)]. [Section 6.1](#) describes necessary values in the IP header and UDP header when an SBFDRReflector on the responder is sending S-BFD control packets.

6.1. Details of S-BFD Control Packet Sent by SBFDRReflector

- o Destination IP address field of the IP header MUST be copied from source IP address field of received S-BFD control packet.
- o Source IP address field of the IP header MUST be set to a local IP address that is expected to be visible by the initiator (i.e. not IPv6 link-local address when the initiator is multiple hops away). The source IP address SHOULD be copied from the destination IP address field of the received S-BFD control packet, except when it is from the 127/8 range for IPv4 or from the 0:0:0:0:0:FFFF:7F00/104 range for IPv6.
- o TTL field of the IP header MUST be set to 255.
- o UDP destination port MUST be copied from received UDP source port.
- o UDP source port MUST be copied from received UDP destination port.

7. Security Considerations

Security considerations for S-BFD are discussed in [[I-D.ietf-bfd-seamless-base](#)]. Additionally, implementing the following measures will strengthen security aspects of the mechanism described by this document:

- o Implementations MUST provide filtering capability based on source IP addresses of received S-BFD control packets: [[RFC2827](#)].
- o Implementations MUST NOT act on received S-BFD control packets containing source Martian IP addresses (i.e., address that, by application of the current forwarding tables, would not have its return traffic routed back to the sender.)
- o Implementations MUST ensure that response S-BFD control packets generated to the initiator by the SBFDRreflector have a reachable target (ex: destination IP address).

8. IANA Considerations

A new value 7784 was allocated from the "Service Name and Transport Protocol Port Number Registry". The allocated registry entry is:

```
Service Name (REQUIRED)
  s-bfd
Transport Protocol(s) (REQUIRED)
  udp
Assignee (REQUIRED)
  IESG <iesg@ietf.org>
Contact (REQUIRED)
  BFD Chairs <bfd-chairs@ietf.org>
Description (REQUIRED)
  Seamless Bidirectional Forwarding Detection (S-BFD)
Reference (REQUIRED)
  RFC.this (RFC Editor, please update at publication)
Port Number (OPTIONAL)
  7784
```

9. Acknowledgements

The authors would like to thank the BFD WG members for helping to shape the contents of this document. In particular, significant contributions were made by following people: Marc Binderberger, Jeffrey Haas, Santosh Pallagatti, Greg Mirsky, Sam Aldrin, Vengada Prasad Govindan, Mallik Mudigonda and Srihari Raghavan.

10. Contributors

The following are key contributors to this document:

```
Tarek Saad, Cisco Systems, Inc.
Siva Sivabalan, Cisco Systems, Inc.
Nagendra Kumar, Cisco Systems, Inc.
```


11. References

11.1. Normative References

- [I-D.ietf-bfd-seamless-base]
Akiya, N., Pignataro, C., Ward, D., Bhatia, M., and J. Networks, "Seamless Bidirectional Forwarding Detection (S-BFD)", [draft-ietf-bfd-seamless-base-09](#) (work in progress), April 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), DOI 10.17487/RFC5880, June 2010, <<http://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, <<http://www.rfc-editor.org/info/rfc5881>>.

11.2. Informative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), DOI 10.17487/RFC0791, September 1981, <<http://www.rfc-editor.org/info/rfc791>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [BCP 38](#), [RFC 2827](#), DOI 10.17487/RFC2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", [RFC 5884](#), DOI 10.17487/RFC5884, June 2010, <<http://www.rfc-editor.org/info/rfc5884>>.

Authors' Addresses

Nobo Akiya
Big Switch Networks

Email: nobo.akiya.dev@gmail.com

Carlos Pignataro
Cisco Systems, Inc.

Email: cpignata@cisco.com

Dave Ward
Cisco Systems, Inc.

Email: wardd@cisco.com