

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 05, 2014

N. Akiya
C. Pignataro
D. Ward
Cisco Systems
July 04, 2013

Seamless Bidirectional Forwarding Detection (S-BFD) Alert Discriminator
and BFD Path Tracing
draft-akiya-bfd-seamless-alert-discrim-00

Abstract

This specification defines a concept of alert discriminator which operates over Seamless Bidirectional Forwarding Detection (S-BFD). New diagnostic codes, solely to be used together with alert discriminators, are also defined in this specification.

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 January 05, 2014.

Copyright Notice

Copyright (c) 2013 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. Overview	3
3. Alert Discriminator	3
4. Reflector BFD Session	4
5. Alert Discriminator Diagnostic Code	4
6. BFD Path Trace: Alert Discriminator Diagnostic Code 31	5
6.1. Initiator Procedures	5
6.1.1. Transmission S-BFD Control Packets	5
6.1.2. Reception of S-BFD Control Packets	6
6.2. Responder Procedures	6
6.2.1. Reception of S-BFD Control Packets	6
6.2.2. Transmission of S-BFD Control Packets	7
6.3. Possible Use Cases	7
7. Security Considerations	7
8. IANA Considerations	8
9. Acknowledgements	8
10. Contributing Authors	8
11. References	8
11.1. Normative References	8
11.2. Informative References	9
Authors' Addresses	9

1. Introduction

[RFC5880] defines the use of Bidirectional Forwarding Detection (BFD) protocol as a fast failure detection mechanism between nodes which are adjacent to each other or multiple hops away. [RFC5881] defines single hop BFD. Specifications such as [RFC5883] and [RFC5884] define multihop BFD.

When multihop BFD, IP based or MPLS based, declares a failure, responsibility of identifying the problematic point in the paths is often left to operators. ICMP echo request/reply (IP ping) [RFC0792] and LSP echo request/reply (LSP ping) [RFC4379] allow for tracing of hops to a specific target, and these are often used, manually or automatically, to attempt to isolate faults. However, when it comes to identifying the problematic point that caused BFD failure, there are couple of issues.

- o Usage of non-BFD packets can result in them being load balanced differently along the paths, causing those packets to traverse different paths than BFD packets to the target.
- o BFD is designed with simplicity and low-overhead as goals. Thus implementations often provide more preferable scale/performance capacities over IP/LSP ping, allowing for increased probability to identify short-lived transient issues.

Above points produced the desire to use BFD to trace hops to a specific target.

This specification defines a generic concept of alert discriminator which operates over Seamless Bidirectional Forwarding Detection (S-BFD) [I-D.akiya-bfd-seamless-base]. New diagnostic codes, solely to be used together with alert discriminators, are also defined in this specification. Finally, BFD path tracing is described as one of the use cases of defined mechanism.

It is worth noting that this specification does not reserve specific BFD discriminator value as the alert discriminator, but only defines the concept of alert discriminators.

2. Overview

A group of network nodes reserves a same BFD discriminator value as the alert discriminator. Alert discriminator operates as a BFD target identifier of alert type (3). A reflector BFD session is then responsible for monitoring incoming BFD control packets with alert discriminator as "your discriminator". Reflector BFD session, upon reception of BFD control packets with alert discriminator as "your discriminator", would examine BFD diagnostic code. Diagnostic code instructs how reflector BFD session is to behave. A network node is able to transmit S-BFD control packets with "your discriminator" as this alert discriminator and well known diagnostic code, to a particular target, and expect reflector BFD session on the target network node to behave accordingly.

3. Alert Discriminator

Alert discriminator is a BFD target identifier of type (3).

Value	BFD Target Identifier Type
-----	-----
3	Alert Discriminator

Uniqueness of alert discriminator is that same BFD discriminator value is reserved on group of network nodes as the alert discriminator.

For example, there are 4 network nodes in a network: A, B, C, D. 0x7F7F7F7F is chosen as the alert discriminator for this network. Nodes A, B, C and D will each reserve 0x7F7F7F7F as BFD target identifier type 3.

How alert discriminator value is to be chosen is outside the scope of this document.

4. Reflector BFD Session

One or more reflector BFD session(s) MUST be created on each network node which has reserved alert discriminator(s). Reflector BFD session MUST listen for incoming S-BFD control packets with "your discriminator" of BFD target identifier type 3, alert discriminators. Further procedures for a reflector BFD session processing incoming S-BFD control packets for BFD target identifier type 3 depends on specified BFD diagnostic code. Definition of BFD diagnostic code for alert discriminator usage and required reflector BFD session behavior for each are described in Section 5.

5. Alert Discriminator Diagnostic Code

[RFC5880] defines a field to describe diagnostic code in a BFD control packet, and defines set of diagnostic codes. This specification defines a new set of diagnostic codes to be used solely for S-BFD control packets using alert discriminators. New diagnostic codes specified in this document are only meaningful when used together with alert discriminators.

- o S-BFD control packets transmitted and received, destined for BFD target identifier of type 3, MUST NOT use diagnostic codes defined in [RFC5880] and MUST use diagnostic codes defined in this document.
- o [S-]BFD control packets transmitted and received, not destined for BFD target identifier of type 3, MUST use diagnostic codes defined in [RFC5880] and MUST NOT use diagnostic codes defined in this document.

Note that BFD diagnostic codes for alert discriminators are defined from highest possible values. Any future documents claiming alert discriminator diagnostic codes MUST use next available highest values from the reserved range. Alert discriminator diagnostic codes are defined as follow:

Value	Alert Discriminator Diagnostic Code Name
0-30	Reserved for future use
31	BFD path trace

When transmitted BFD control packet is targeted to a BFD target identifier of type 3, then BFD diagnostic code MUST NOT be zero. When receiving BFD control packet is targeted to a BFD target identifier of type 3, then packet with BFD diagnostic code of zero MUST be dropped.

Note that primary purpose of alert discriminator diagnostic codes are to provide hints to responder on why initiator is sending alert discriminator S-BFD packets.

6. BFD Path Trace: Alert Discriminator Diagnostic Code 31

BFD path trace, aka BFD traceroute, is performed through making use of the alert discriminator with alert discriminator diagnostic code 31.

6.1. Initiator Procedures

When a network node desires to trace hops to a BFD target, S-BFD control packets are transmitted with following contents.

6.1.1. Transmission S-BFD Control Packets

- o IP destination address or MPLS label stack MUST be set to describe the target.
- o "your discriminator" MUST be set to an alert discriminator.
- o BFD diagnostic code MUST be set to 31 (BFD path trace).
- o Poll (P) bit MUST be set.
- o Incrementing or decrementing IP/MPLS TTL.
- o Remaining packet contents are as per described in [I-D.akiya-bfd-seamless-ip].

When incrementing TTL is used towards the BFD target, TTL SHOULD start at value of 1. Completion of BFD path trace is reached when locally determined so (ex: no response from one of the nodes) or when one of following conditions are hit, and initiator MUST NOT transmit BFD path trace packets to further downstream network nodes:

- o Response S-BFD control packet has been received from intended BFD target.
- o In case IP address(es) of intended BFD target is unknown, two consecutive response S-BFD control packets (TTL+n and TTL+(n+1)) contain same IP source address.

When decrementing TTL is used, BFD path trace SHOULD start from the BFD target using TTL=N. How value of N is determined is outside the scope of this document. Completion of BFD path trace is reached when locally determined so or after performing BFD path trace operation to TTL=1.

Because there are no sequence numbers included in transmitted and received S-BFD control packets (without use of Authentication) for BFD path tracing, initiator SHOULD allow some delay between multiple BFD path tracing operations for a same target, if same "my discriminator" value is used on them. This is to ensure responses from multiple BFD path tracing operations do not conflict with each other, resulting in incorrectly recorded hops.

6.1.2. Reception of S-BFD Control Packets

If response S-BFD control packets do not contain "my discriminator" of alert discriminator, then packet MUST NOT be considered as response for BFD path tracing.

If response S-BFD control packets do not have Final (F) bit set, then packet MUST NOT be considered as response for BFD path tracing.

If response S-BFD control packets do not contain BFD diagnostic code 31, then packet MUST NOT be considered as response for BFD path tracing.

IP source address of valid response S-BFD control packets are recorded to form trace hops to the BFD target.

6.2. Responder Procedures

Reflector BFD session at the responder network node MUST operate with procedures described in [I-D.akiya-bfd-seamless-ip].

6.2.1. Reception of S-BFD Control Packets

Following conditions MUST be met for received S-BFD control packets targeted to BFD target identifier of type 3 to be considered for BFD path tracing:

- o BFD diagnostic code is 31 (BFD path trace).
- o Poll (P) bit is set.

6.2.2. Transmission of S-BFD Control Packets

Following procedures MUST be followed when transmitting a response S-BFD control packet for BFD path tracing:

- o BFD diagnostic code in response S-BFD packet MUST be set to 31 (BFD path trace).
- o Final (F) bit MUST be set.

6.3. Possible Use Cases

BFD path tracing may be desirable for following occasions.

- o When a BFD session is determined to have lost reachability to the target (ex: state transitions from UP to DOWN), immediately trigger BFD path trace to the target to attempt to isolate the fault.
- o While a particular BFD session is in UP state, occasionally trigger BFD path trace in the background to record the paths. Compare recorded paths to see how frequently paths are changing. If determined to be more frequent than expected, then log a warning to indicate potential network instability.
- o Just trigger BFD path trace, manually or automatically, as needed basis.

7. Security Considerations

Alert discriminator selected for a network should be kept from being disclosed to anybody or anything external to the network. This will prevent attacks from knowing the exact value for the alert discriminator. It is still possible for attacks to scan a range of BFD discriminator values to identify alert discriminator being used. Therefore, as described in [I-D.akiya-bfd-seamless-base], implementations MUST provide filtering capability based on source IP addresses.

In addition, same security considerations as [RFC5880], [RFC5881], [RFC5883], [RFC5884], [I-D.akiya-bfd-seamless-base] and [I-D.akiya-bfd-seamless-ip] apply to this document.

8. IANA Considerations

BFD Target Identifier types:

Value	BFD Target Identifier Type
-----	-----
3	Alert Discriminator

Alert Discriminator Diagnostic Code:

Value	Alert Discriminator Diagnostic Code Name
-----	-----
0-30	Reserved for future use
31	BFD path trace

9. Acknowledgements

TBD

10. Contributing Authors

Nagendra Kumar
Cisco Systems
Email: naikumar@cisco.com

Mallik Mudigonda
Cisco Systems
Email: mmudigon@cisco.com

Aswatnarayan Raghuram
AT&T
Email: ar2521@att.com

Glenward D. Hayden
AT&T
Email: gh1691@att.com

11. References

11.1. Normative References

[I-D.akiya-bfd-seamless-base]

Akiya, N., Pignataro, C., and D. Ward, "Seamless Bidirectional Forwarding Detection (BFD) with MPLS Label Verification Extension", draft-akiya-bfd-seamless-base-00 (work in progress), June 2013.

[I-D.akiya-bfd-seamless-ip]

Akiya, N., Pignataro, C., and D. Ward, "Seamless Bidirectional Forwarding Detection (BFD) for IP", draft-akiya-bfd-seamless-ip-00 (work in progress), June 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.

[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.

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

11.2. Informative References

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

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

Authors' Addresses

Nobo Akiya
Cisco Systems

Email: nobo@cisco.com

Carlos Pignataro
Cisco Systems

Email: cpignata@cisco.com

Dave Ward
Cisco Systems

Email: wardd@cisco.com