

Internet Engineering Task Force (IETF)
Request for Comments: 9384
Category: Standards Track
ISSN: 2070-1721

J. Haas
Juniper Networks
March 2023

A BGP Cease NOTIFICATION Subcode for Bidirectional Forwarding Detection (BFD)

Abstract

The Bidirectional Forwarding Detection (BFD) protocol (RFC 5880) is used to detect loss of connectivity between two forwarding engines, typically with low latency. BFD is leveraged by routing protocols, including the Border Gateway Protocol (BGP), to bring down routing protocol connections more quickly than the original protocol timers.

This document defines a subcode for the BGP Cease NOTIFICATION message (Section 6.7 of RFC 4271) for use when a BGP connection is being closed due to a BFD session going down.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9384>.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
2. Requirements Language
3. BFD Cease NOTIFICATION Subcode
4. Operational Considerations
5. Security Considerations
6. IANA Considerations
7. References
 - 7.1. Normative References
 - 7.2. Informative References

Acknowledgments

1. Introduction

The Bidirectional Forwarding Detection (BFD) protocol [RFC5880] is used to detect loss of connectivity between two forwarding engines, typically with low latency. BFD is utilized as a service for various clients, including routing protocols, to provide an advisory mechanism for those clients to take appropriate actions when a BFD session goes down [RFC5882]. This is typically used by the clients to trigger closure of their connections more quickly than the original protocol timers might allow.

Border Gateway Protocol version 4 (BGP-4) [RFC4271] terminates its connections upon Hold Timer expiration when the speaker does not receive a BGP message within the negotiated Hold Time interval. As per Sections 4.2 and 4.4 of [RFC4271], the minimum Hold Time interval is at least three seconds, unless KEEPALIVE processing has been disabled by negotiating the distinguished Hold Time of zero.

If a BGP speaker desires to have its connections terminate more quickly than the negotiated BGP Hold Timer can accommodate upon loss of connectivity with a neighbor, the BFD protocol can be relied upon by BGP speakers to supply that faster detection. When the BFD session state changes to Down, the BGP speaker terminates the connection with a Cease NOTIFICATION message sent to the neighbor, if possible, and then closes the TCP connection for the session.

This document defines a subcode, "BFD Down", to be sent with the Cease NOTIFICATION message that indicates the reason for this type of connection termination.

2. 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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. BFD Cease NOTIFICATION Subcode

The value 10 has been allocated by IANA for the "BFD Down" Cease NOTIFICATION message subcode.

When a BGP connection is terminated due to a BFD session going into the Down state, the BGP speaker SHOULD send a NOTIFICATION message with the error code "Cease" and the error subcode "BFD Down".

4. Operational Considerations

A BFD session may go into the Down state when there is only a partial loss of connectivity between two BGP speakers. Operators using BFD for their BGP connections make choices regarding what BFD timers are used based upon a variety of criteria -- for example, stability vs. fast failure.

In the event of a BGP connection being terminated due to a "BFD Down" event from partial loss of connectivity as detected by BFD, the remote BGP speaker might be able to receive a BGP Cease NOTIFICATION message with the "BFD Down" subcode. The receiving BGP speaker will then have an understanding that the connection is being terminated because of a BFD-detected issue and not an issue with the BGP speaker.

When there is a total loss of connectivity between two BGP speakers,

it may not have been possible for the Cease NOTIFICATION message to have been sent. Even so, BGP speakers SHOULD provide this reason as part of their operational state. Examples include `bgpPeerLastError` per the BGP MIB [RFC4273] and "last-error" per [BGP-YANG].

When the procedures in [RFC8538] for sending a NOTIFICATION message with a "Cease" code and "Hard Reset" subcode are required, and the BGP connection is being terminated because BFD has gone into the Down state, the "BFD Down" subcode SHOULD be encapsulated in the Hard Reset's data portion of the NOTIFICATION message.

5. Security Considerations

Similar to [RFC4486], this document defines a subcode for the BGP Cease NOTIFICATION message that provides information to aid network operators in correlating network events and diagnosing BGP peering issues. This subcode is purely informational and has no impact on the BGP Finite State Machine beyond that already documented by [RFC4271], Sections 6.6 and 6.7.

6. IANA Considerations

IANA has assigned the value 10 from the "BGP Cease NOTIFICATION message subcodes" registry (<https://www.iana.org/assignments/bgp-parameters/>), with the name "BFD Down" and a reference to this document.

7. References

7.1. Normative References

- [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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [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>>.
- [RFC5882] Katz, D. and D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, DOI 10.17487/RFC5882, June 2010, <<https://www.rfc-editor.org/info/rfc5882>>.
- [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>>.
- [RFC8538] Patel, K., Fernando, R., Scudder, J., and J. Haas, "Notification Message Support for BGP Graceful Restart", RFC 8538, DOI 10.17487/RFC8538, March 2019, <<https://www.rfc-editor.org/info/rfc8538>>.

7.2. Informative References

- [BGP-YANG] Jethanandani, M., Patel, K., Hares, S., and J. Haas, "YANG Model for Border Gateway Protocol (BGP-4)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-model-16, 1 March 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-model-16>>.

[RFC4273] Haas, J., Ed. and S. Hares, Ed., "Definitions of Managed Objects for BGP-4", RFC 4273, DOI 10.17487/RFC4273, January 2006, <<https://www.rfc-editor.org/info/rfc4273>>.

[RFC4486] Chen, E. and V. Gillet, "Subcodes for BGP Cease Notification Message", RFC 4486, DOI 10.17487/RFC4486, April 2006, <<https://www.rfc-editor.org/info/rfc4486>>.

Acknowledgments

Thanks to Jeff Tantsura and Dale Carder for their comments on this document.

Mohamed Boucadair provided feedback as part of the Routing Directorate review of this document.

In 2006, Bruno Rijsman had written a proposal that was substantively similar to this document: draft-rijsman-bfd-down-subcode. That draft did not progress in the Inter-Domain Routing (IDR) Working Group at that time. The author of this document was unaware of Bruno's prior work when creating this proposal.

Author's Address

Jeffrey Haas
Juniper Networks
Email: jhaas@juniper.net