Workgroup: IDR Workgroup
Internet-Draft:
draft-ietf-idr-bgp-bfd-strict-mode-07
Updates: RFC4271 (if approved)
Published: 5 January 2023
Intended Status: Standards Track
Expires: 9 July 2023

Authors: M. Zheng, Ciena
A. Lindem, Cisco Systems
J. Haas, Juniper Networks, Inc.
A. Fu, Bloomberg L.P.

BGP BFD Strict-Mode

Abstract

This document specifies extensions to RFC4271 BGP-4 that enable a BGP speaker to negotiate additional Bidirectional Forwarding Detection (BFD) extensions using a BGP capability. This BFD Strict-Mode Capability enables a BGP speaker to prevent a BGP session from being established until a BFD session is established. It is referred to as BGP BFD "strict-mode". BGP BFD strict-mode will be supported when both the local speaker and its remote peer are BFD strict-mode capable.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Bidirectional Forwarding Detection BFD [RFC5882] enables routers to monitor data plane connectivity and to detect faults in the bidirectional forwarding path between them. This capability is leveraged by routing protocols such as BGP [RFC4271] to rapidly react to topology changes in the face of path failures.

The BFD interaction with BGP is specified in Section 10.2 of [RFC5882]. When BFD is enabled for a BGP neighbor, faults in the bidirectional forwarding detected by BFD result in session termination. It is possible in some failure scenarios for the network to be in a state such that a BGP session may be established but a BFD session cannot be established. In some other scenarios, it may be possible to establish a BGP session, but a degraded or poor-quality link may result in the corresponding BFD session going up and down frequently.

To avoid situations which result in routing churn and to minimize the impact of network interruptions, it will be beneficial to disallow BGP to establish a session until BFD session is successfully established and has stabilized. We refer to this mode of operation as BGP BFD "strict-mode". However, always using "strict-mode" would preclude BGP operation in an environment where not all routers support BFD strict-mode or have BFD enabled. This document defines BGP "strict-mode" operation as preventing BGP session establishment until both the local and remove speakers have a stable BFD session. The document also specifies the BGP protocol extensions for BGP capability [RFC5492] for announcing BFD parameters including a BGP speaker's support for "strict-mode", i.e., requiring a BFD session for BGP session establishment.

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 Strict-Mode Capability

The BGP Strict-Mode Capability [RFC5492] will allow a BGP speaker's to advertise this capability. The capability is defined as follows:

Capability code: 74

Capability length: 0 octets

4. BGP BFD Strict-Mode Procedures

A BGP speaker which supports capabilities advertisement and has BFD strict-mode enabled MUST include the BFD Strict-Mode Capability.

A BGP speaker which supports the BFD Strict-Mode Capability, examines the list of capabilities present in the capabilities that the speaker receives from its peer. If both the local and remote BGP speakers include the BFD Strict-Mode Capability, the BGP finite state machine does not transition to the Established state from OpenConfirm state [RFC4271] until the BFD session is in the Up state (see below for AdminDown state). This means that a KEEPALIVE message is not sent nor is the KeepaliveTimer set.

If the BFD session does not transition to the Up state, and the HoldTimer has been negotiated to a non-zero value, the BGP FSM will close the session appropriately. If the HoldTimer has been negotiated to a zero value, the session should be closed after a time of X. This time X is referred as "BGP BFD Hold time". The proposed default BGP BFD Hold time value is 30 seconds. The BGP BFD Hold time value is configurable.

If BFD session is in the AdminDown state, then the BGP finite state machine will proceed normally without input from BFD. This means that BFD session "AdminDown" state WILL NOT prevent the BGP state transition from the OpenConfirm state to the Established state.

Once the BFD session has transitioned to the Up state, the BGP FSM may proceed to transition from the OpenConfirm state to state

Established state. Once in the Established state, a KEEPALIVE message is sent and a KeepaliveTimer for the BGP peer is started.

BGP strict-mode cannot be enabled unless BFD is configured for the BGP peer. If BFD is removed for the BGP peer, then BGP strict-mode will also be disabled.

Note that it is fully possible to have BFD enabled between the peers without BGP strict-mode.

If either BGP peer has not advertised the BFD Strict-Mode Capability, then a BFD session WILL NOT be required for the BGP session to reach Established state. This does not preclude usage of BFD after BGP session establishment [RFC5882].

If BFD strict-mode is enabled or disabled for a BGP peer and the BGP session state is not Established state, then the BGP will close the session.

If the BFD Stict-Mode is enabled or disabled for a BGP peer and the BGP session state is Established state, the local BFD strict-mode configuration will be modified but the session will remain in Established state.

Since BFD strict-mode is only applicable during BGP session establishment, this inconsistency would not have an impact on the Established session unless the remote BGP peer is waiting in OpenConfirm state. To avoid this situation, BFD strict-mode SHOULD be modified consistently on both the local and remote BGP peers.

4.1. Stability Considerations

The use of BGP BFD strict-mode along with mechanisms such as hold-down (a delay in the initial BGP Establishment state following BFD session establishment) and/or dampening (a delay in the BGP Establishment state following failure detected by BFD) may help reduce the frequency of BGP session flaps and therefore reduce the associated routing churn. The details of these mechanisms are outside the scope of this document.

5. Manageability Considerations

Auto-configuration is possible for the enabling BGP BFD strict-mode. However, the configuration automation is out of the scope of this document.

To simplify troubleshoorting and avoid inconsistencies, it is RECOMMENDED that BFD strict-mode configuration be consistent for both BGP peers.

6. Security Considerations

The mechanism defined in this document interacts with the BGP finite state machine when so configured. The security considerations of BFD thus, become considerations for BGP-4 [RFC4271] so used. Given that a BFD session is required for a BGP session, a Denial-of-Service (DoS) attack on BGP can now be mounted by preventing a BFD session between the BGP peers from being established or interrupting an existing BFD session. The use of the BFD Authentication mechanism defined in [RFC5880] is thus RECOMMENDED when used to protect BGP-4 [RFC4271].

7. IANA Considerations

This document defines the BGP BFD Strict-Mode Capability. The Capability Code 74 has been assigned from the First-Come-First-Served range (64-238) of the Capability Codes registry.

8. Acknowledgement

The authors would like to acknowledge the review and inputs from Shyam Sethuram, Mohammed Mirza, Bruno Decraene, Carlos Pignataro, and Enke Chen.

9. Normative References

[RFC2119] Bradner, S. and RFC Publisher, "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., Hares, S., Ed., and RFC Publisher, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC5492] Scudder, J., Chandra, R., and RFC Publisher, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <https://www.rfc-editor.org/info/rfc5492>.

[RFC5880] Katz, D., Ward, D., and RFC Publisher, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC5882] Katz, D., Ward, D., and RFC Publisher, "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. and RFC Publisher, "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>.

Authors' Addresses

Mercia Zheng
Ciena
3939 N. 1st Street
San Jose, CA 95134
United States

Email: merciaz.ietf@gmail.com

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
United States

Email: acee.ietf@gmail.com

Jeffrey Haas
Juniper Networks, Inc.
1133 Innovation Way
SUNNYVALE, CALIFORNIA 94089
United States

Email: jhaas@juniper.net

Albert Fu
Bloomberg L.P.

Email: afu14@bloomberg.net