IDR WorkGroup
Internet-Draft
Intended status: Standards Track
Expires: October 5, 2019

M. Zheng
A. Lindem
Cisco Systems

April 3, 2019

BGP BFD Strict-Mode
draft-merciaz-idr-bgp-bfd-strict-mode-01

Abstract

This document specifies extensions to RFC4271 BGP-4 that enable a BGP speaker to signal additional Bidirectional Forwarding Detection (BFD) extensions using an optional parameter BFD capability. This BFD 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, Otherwise, a BGP speaker and its peer should not require a BFD session for BGP session establishment.

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 October 5, 2019.

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 neighbor session until s 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 Capability

The BGP Capability [RFC5492] for BFD parameters will allow a BGP speaker's BFD capabilities including its support for BFD strict-mode. This capability is defined as follows:

Capability code: TBD

Capability length: 1 octet

Capability value: Consists of 1 octet BFD flags as follows:

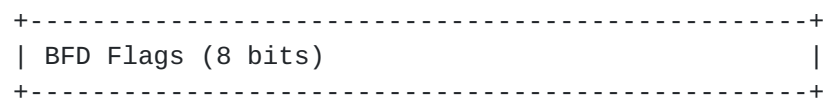

The use and meaning of the fields are as follows:

BFD Flags: This field contains bit flags relating to BFD.

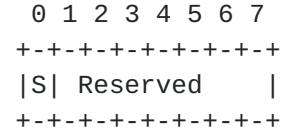

The most significant bit is defined as state of Strict-Mode ("Strict-Mode", or "S") bit, which can be used by a BGP speaker to signal its support for BFD Strict-mode. When set (value 1), this bit indicates that the BGP speaker has the BFD "Strict-mode" enabled. If both local BGP speaker and its peer have BFD strict-mode enabled, then BGP session establishment will be prevented until a BFD session is established between the peering addresses. A BGP speaker with BFD

strict-mode enabled MUST advertise the BFD capability with "S" bit set.

The remaining bits are reserved and SHOULD be set to zero by the sender and MUST be ignored by the receiver.

4. Operation

A BGP speaker that supports capabilities advertisement sends an OPEN message to its BGP peer, the message MAY include an Optional Parameter, called Capabilities. The parameter lists the capabilities supported by the speaker. By following BGP capabilities advertisement procedures defined in [RFC5492], BFD capability advertisement for strict-mode is advertised to BGP peers.

If both BGP speakers advertise the BFD capability with the strict-mode bit set, then the BGP state machine will be held in OpenConfirm state [RFC4271] until a BFD session is established or the BGP session is terminated for some other reason (e.g., the BGP Hold time expires.

A BGP speaker which supports capabilities advertisement and has BFD strict-mode enabled MUST include the BGP BFD capability with the "S" Bit set in the BGP capabilities it advertises.

A BGP speaker which supports BFD capability advertisement, examines the list of capabilities present in the Capabilities BFD Parameter that the speaker receives from its peer. If both the local and remote BGP speakers BFD strict-mode enabled, then the BGP session will be held in OpenConfirm state until a BFD session is established between the two BGP speaker or the BGP session terminates for some other reason, e.g., the BGP hold timer expires. If either peer has not advertised the BFD Capability with strict-mode enabled, 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].

5. Backward Compatibility

The BFD capability will not introduce any backward compatibility will not result in any backward compatibility issues as long as the procedures in [RFC5492] and Section 4 are followed. As per [RFC5492], a BGP speaker which does not support the BFD capability will ignore the BFD capability. If a BGP speaker advertising the capability receives the Unsupported Capability NOTIFICATION message and terminates the BGP session, the BGP speaker advertising the BFD capability SHOULD simply attempt to reestablish the BGP session with the BFD capability omitted.

6. Security Considerations

This specification doesn't change the basic security model inherent in [RFC4271]. However, it does introduce a new indirect attack vector for BGP since it is now dependent on BFD.

7. IANA Considerations

This document defines a new BGP capability - BFD Capability. The Capability Code for BFD Capability is TBD.

IANA is requested to establish a "BGP BFD Capability Flags" registry within the "Border Gateway Protocol (BGP) Parameters" grouping. The Registration Procedure should be Standards Action, the initial values as follows:

Bit Position	Name	Short Name	Reference
0	Strict-Mode	S	this document
1-7	Unassigned		this document

8. Acknowledgement

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

9. 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>.

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

[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>.

Authors' Addresses

Mercia Zheng
Cisco Systems
821 Alder Drive,
MILPITAS, CALIFORNIA 95035
UNITED STATES

Email: merciaz@cisco.com

Acee Lindem
Cisco Systems
821 Alder Drive,
MILPITAS, CALIFORNIA 95035
UNITED STATES

Email: acee@cisco.com