Stream: Internet Engineering Task Force (IETF)
RFC: 9355
Updates: 2328
Category: Standards Track
Published: February 2023
ISSN: 2070-1721
Authors: K. Talaulikar, Ed. *Cisco Systems, Inc.* P. Psenak *Cisco Systems, Inc.* A. Fu *Bloomberg* M. Rajesh *Juniper Networks*

RFC 9355
OSPF Bidirectional Forwarding Detection (BFD) Strict-Mode

Abstract

This document specifies the extensions to OSPF that enable an OSPF router to signal the requirement for a Bidirectional Forwarding Detection (BFD) session prior to adjacency formation. Link-Local Signaling (LLS) is used to advertise the requirement for strict-mode BFD session establishment for an OSPF adjacency. If both OSPF neighbors advertise BFD strict-mode, adjacency formation will be blocked until a BFD session has been successfully established.

This document updates RFC 2328 by augmenting the OSPF neighbor state machine with a check for BFD session up before progression from Init to 2-Way state when operating in OSPF BFD strict-mode.

Table of Contents

1. Introduction
   1.1. Requirements Language
2. LLS B-Bit Flag
3. Local Interface IPv4 Address TLV
4. Procedures
   4.1. OSPFv3 IPv4 AF Specifics
   4.2. Graceful Restart Considerations
5. Operations and Management Considerations
6. Backward Compatibility
7. IANA Considerations
8. Security Considerations
9. References
   9.1. Normative References
   9.2. Informative References
Acknowledgements
Authors' Addresses

1. Introduction

Bidirectional Forwarding Detection (BFD) [RFC5880] enables routers to monitor data plane connectivity and to detect faults in the bidirectional path between them. BFD is leveraged by routing protocols like OSPFv2 [RFC2328] and OSPFv3 [RFC5340] to detect connectivity failures for established adjacencies faster than the OSPF Hello dead timer detection and to trigger rerouting of traffic around the failure. The use of BFD for monitoring routing protocol adjacencies is described in [RFC5882].

When BFD monitoring is enabled for OSPF adjacencies by the network operator, the BFD session is bootstrapped based on the neighbor address information discovered by the exchange of OSPF Hello packets. Faults in the bidirectional forwarding detected via BFD then result in the OSPF adjacency being brought down. A degraded or poor-quality link may result in intermittent packet drops. In such scenarios, implementations prior to the extensions specified in this document may still get an OSPF adjacency established over such a link; however, given the more aggressive monitoring intervals supported by BFD, a BFD session may not get established and/or may flap. The traffic forwarded over such a link would experience packet drops, and the failure of the BFD session establishment will not enable fast routing convergence. OSPF adjacency flaps may occur over such links when OSPF brings up the adjacency only for it to be brought down again by BFD.

To avoid the routing churn associated with these scenarios, it would be beneficial not to allow OSPF to establish an adjacency until a BFD session is successfully established and has stabilized. However, this would preclude the OSPF operation in an environment where not all OSPF routers support BFD and have it enabled on the link. A solution is to block OSPF adjacency establishment until a BFD session is established as long as both neighbors advertise such a requirement. Such a mode of OSPF BFD usage is referred to as "strict-mode". Strict-mode introduces signaling support in OSPF to achieve the blocking of adjacency formation until BFD session establishment occurs, as described in Section 4.1 of [RFC5882].

This document specifies the OSPF protocol extensions using Link-Local Signaling (LLS) [RFC5613] for a router to indicate to its neighbor the willingness to require BFD strict-mode for OSPF adjacency establishment (refer to Section 2). It also introduces an extension to OSPFv3 LLS of the interface IPv4 address (refer to Section 3) to be used for the BFD session setup when OSPFv3 is used for an IPv4 Address Family (AF) instance.

This document updates [RFC2328] by augmenting the OSPF neighbor state machine with a check for BFD session up before progression from Init to 2-Way state when operating in OSPF BFD strict-mode.

The extensions and procedures for OSPF BFD strict-mode also apply for adjacency over virtual links using BFD multi-hop [RFC5883] procedures.

A similar functionality for IS-IS is specified in [RFC6213].

1.1. 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.

2. LLS B-Bit Flag

This document defines the B-bit in the LLS Type 1 Extended Options and Flags. This bit is defined for the LLS block that is included in Hello and Database Description (DD) packets. The B-bit indicates that BFD is enabled on the link and that the router requests OSPF BFD strict-mode. Section 7 describes the position of the B-bit.

A router **MUST** include the LLS block with the B-bit set in the LLS Type 1 Extended Options and Flags in its Hello and DD packets when OSPF BFD strict-mode is enabled on the link.

3. Local Interface IPv4 Address TLV

The Local Interface IPv4 Address TLV is an LLS TLV defined for OSPFv3 IPv4 AF instance [RFC5838] protocol operation as described in Section 4.1.

It has the following format:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              Type             |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 Local Interface IPv4 Address                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type: 21

Length: 4 octets

Local Interface IPv4 Address: The primary IPv4 address of the local interface.

4. Procedures

A router supporting OSPF BFD strict-mode advertises this capability through its Hello packets as described in Section 2. When a router supporting OSPF BFD strict-mode discovers a new neighbor router that also supports OSPF BFD strict-mode, it will establish a BFD session with that neighbor first before bringing up the OSPF adjacency as described further in this section.

This document updates the OSPF neighbor state machine as described in [RFC2328]. Specifically, the operations related to the Init state are modified as described below when OSPF BFD strict-mode is used:

Init (without OSPF BFD strict-mode):
: In this state, a Hello packet has recently been received from the neighbor. However, bidirectional communication has not yet been established with the neighbor (i.e., the router itself did not appear in the neighbor's Hello packet). All neighbors in this state (or higher) are listed in the Hello packets sent from the associated interface.

Init (with OSPF BFD strict-mode):
: In this state, a Hello packet has recently been received from the neighbor. However, bidirectional communication has not yet been established with the neighbor (i.e., the router itself did not appear in the neighbor's Hello packet). BFD session establishment with the neighbor is requested if it's not already completed (e.g., in the event of transition from 2-Way state). Neighbors in Init state or higher will be listed in Hello packets associated with the interface if they either have a corresponding BFD session established or have not advertised OSPF BFD strict-mode in the LLS Type 1 Extended Options and Flags advertised in the Hello packet.

When the neighbor state transitions to Down state, the removal of the BFD session associated with that neighbor is requested by OSPF; subsequent BFD session establishment is similarly requested by OSPF upon transitioning into Init state. This may result in BFD session deletion and creation, respectively, when OSPF is the only client interested in the BFD session with the neighbor address.

An implementation **MUST NOT** wait for BFD session establishment in Init state unless OSPF BFD strict-mode is enabled by the operator on the interface and the specific neighbor indicates OSPF BFD strict-mode capability via the LLS Type 1 Extended Options and Flags advertised in the Hello packet. When BFD is enabled, but OSPF BFD strict-mode has not been signaled by both neighbors, an implementation **SHOULD** start BFD session establishment only in 2-Way or greater state. This makes it possible for an OSPF router to support BFD operation in both strict-mode and normal mode across different interfaces or even across different neighbors on the same multi-access interface.

Once the OSPF state machine has moved beyond the Init state, any change in the B-bit advertised in subsequent Hello packets **MUST NOT** result in any trigger in either the OSPF adjacency or the BFD session management (i.e., the B-bit is considered only when in Init state). Disabling BFD (or OSPF BFD strict-mode) on an OSPF interface would result in it not setting the B-bit in the LLS Type 1 Extended Options and Flags advertised in subsequent Hello packets. Disabling OSPF BFD strict-mode has no effect on BFD operations and would not result in the bringing down of any established BFD sessions. Disabling BFD would result in the BFD session being brought down due to AdminDown State (described in Section 3.2 of [RFC5882]); hence, it would not bring down the OSPF adjacency.

When BFD is enabled on an interface over which we already have an existing OSPF adjacency, it would result in the router setting the B-bit in its subsequent Hello packets and the initiation of BFD session establishment to the neighbor. If the adjacency is already up (i.e., in its terminal state of Full or 2-Way with routers that are not designated routers on a multi-access interface) with a neighbor that also supports OSPF BFD strict-mode, then an implementation **SHOULD NOT** bring this adjacency down into the Init state to avoid disruption to routing operations and instead use the OSPF BFD strict-mode wait only after a transition to Init state. However, if the adjacency is not up, then an implementation **MAY** bring such an adjacency down so it can use the OSPF BFD strict-mode for its adjacency establishment.

4.1. OSPFv3 IPv4 AF Specifics

Support for multiple AFs in OSPFv3 [RFC5838] requires the use of an IPv6 link-local address as the source address for Hello packets, even when forming adjacencies for IPv4 AF instances. In most deployments of OSPFv3 IPv4 AFs, it is required that BFD is used to monitor and verify IPv4 data plane connectivity between the routers on the link; hence, the BFD session is set up using IPv4 neighbor addresses. The IPv4 neighbor address on the interface is learned only later in the adjacency formation process when the neighbor's Link-LSA (Link State Advertisement) is received. This results in the setup of the BFD IPv4 session either after the adjacency is established or later in the adjacency formation sequence.

To operate in OSPF BFD strict-mode, it is necessary for an OSPF router to learn its neighbor's IPv4 link address during the Init state of adjacency formation (ideally, when it receives the first Hello). The use of the Local Interface IPv4 Address TLV (as defined in Section 3) in the LLS block advertised in OSPFv3 Hello packets for IPv4 AF instances makes this possible. Implementations that support OSPF BFD strict-mode for OSPFv3 IPv4 AF instances **MUST** include the Local Interface IPv4 Address TLV in the LLS block advertised in their Hello packets whenever the B-bit is also set in the LLS Type 1 Extended Options and Flags. A receiver **MUST** ignore the B-bit (i.e., not operate in strict-mode for BFD) when the Local Interface IPv4 Address TLV is not present in OSPFv3 Hello messages for OSPFv3 IPv4 AF instances.

4.2. Graceful Restart Considerations

An implementation needs to handle scenarios where both graceful restart (GR) and the OSPF BFD strict-mode are deployed together. The graceful restart aspects related to process restart scenarios discussed in Section 3.3 of [RFC5882] also apply with OSPF BFD strict-mode. Additionally, since the OSPF adjacency formation is delayed until the BFD session establishment in OSPF BFD strict-mode, the resultant delay in adjacency formation may affect or break the GR-based recovery. In such cases, it is **RECOMMENDED** that the GR timers are set such that they provide sufficient time to allow for normal BFD session establishment delays.

5. Operations and Management Considerations

An implementation **SHOULD** report the BFD session status along with the OSPF Init adjacency state when OSPF BFD strict-mode is enabled and support logging operations on neighbor state transitions that include the BFD events. This allows an operator to detect scenarios where an OSPF adjacency may be stuck waiting for BFD session establishment.

In network deployments with noisy or degraded links with intermittent packet loss, BFD sessions may flap, resulting in OSPF adjacency flaps. In turn, this may cause routing churn. The use of OSPF BFD strict-mode along with mechanisms such as hold-down (a delay in bringing up the initial OSPF adjacency following BFD session establishment) and/or dampening (a delay in bringing up the OSPF adjacency following failure detected by BFD) may help reduce the frequency of adjacency flaps and therefore reduce the associated routing churn. The details of these mechanisms are outside the scope of this document.

[RFC9129] specifies the base OSPF YANG module. The required configuration and operational elements for this feature are expected to be introduced as augmentation to this base OSPF YANG module.

6. Backward Compatibility

An implementation **MUST** support OSPF adjacency formation and operations with a neighbor router that does not advertise the OSPF BFD strict-mode capability: both when that neighbor router does not support BFD and when it does support BFD but does not signal the OSPF BFD strict-mode as described in this document. Implementations **MAY** provide a local configuration option to force BFD operation only in OSPF BFD strict-mode (i.e, adjacency will not come up unless BFD session is established). In this case, an OSPF adjacency with a neighbor that does not support OSPF BFD strict-mode would not be established successfully. Implementations **MAY** provide a local configuration option to enable BFD without the OSPF BFD strict-mode, which results in the router not advertising the B-bit and BFD operation being performed in the same way as prior to this specification.

The signaling specified in this document happens at a link-local level between routers on that link. A router that does not support this specification would ignore the B-bit in the LLS block advertised in Hello packets from its neighbors and continue to establish BFD sessions (if enabled) without delaying the OSPF adjacency formation. Since a router that does not support this specification would not have set the B-bit in the LLS block advertised in its own Hello packets, its neighbor routers supporting this specification would not use OSPF BFD strict-mode with such OSPF routers. As a result, the behavior would be the same as without this specification. Therefore, there are no backward compatibility issues or implementation considerations beyond what is specified herein.

7. IANA Considerations

This specification makes the following updates under the "Open Shortest Path First (OSPF) Link Local Signaling (LLS) - Type/Length/Value Identifiers (TLV)" parameters.

- In the "LLS Type 1 Extended Options and Flags" registry, the B-bit has been assigned the bit position 0x00000010.
- In the "Link Local Signaling TLV Identifiers (LLS Types)" registry, the Type 21 has been assigned to the Local Interface IPv4 Address TLV.

8. Security Considerations

The security considerations for "OSPF Link-Local Signaling" [RFC5613] also apply to the extension described in this document. Inappropriate use of the B-bit in the LLS block of an OSPF Hello message could prevent an OSPF adjacency from forming or lead to the failure of detecting bidirectional forwarding failures. If authentication is being used in the OSPF routing domain [RFC5709] [RFC7474], then the Cryptographic Authentication TLV [RFC5613] **MUST** also be used to protect the contents of the LLS block.

9. References

9.1. 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>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC5613] Zinin, A., Roy, A., Nguyen, L., Friedman, B., and D. Yeung, "OSPF Link-Local Signaling", RFC 5613, DOI 10.17487/RFC5613, August 2009, <https://www.rfc-editor.org/info/rfc5613>.

[RFC5838] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, DOI 10.17487/RFC5838, April 2010, <https://www.rfc-editor.org/info/rfc5838>.

[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>.

9.2. Informative References

[RFC5709] Bhatia, M., Manral, V., Fanto, M., White, R., Barnes, M., Li, T., and R. Atkinson, "OSPFv2 HMAC-SHA Cryptographic Authentication", RFC 5709, DOI 10.17487/RFC5709, October 2009, <https://www.rfc-editor.org/info/rfc5709>.

[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>.

[RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883, June 2010, <https://www.rfc-editor.org/info/rfc5883>.

[RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", RFC 6213, DOI 10.17487/RFC6213, April 2011, <https://www.rfc-editor.org/info/rfc6213>.

[RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <https://www.rfc-editor.org/info/rfc7474>.

[RFC9129] Yeung, D., Qu, Y., Zhang, Z., Chen, I., and A. Lindem, "YANG Data Model for the OSPF Protocol", RFC 9129, DOI 10.17487/RFC9129, October 2022, <https://www.rfc-editor.org/info/rfc9129>.

Acknowledgements

The authors would like to acknowledge the review and inputs from Acee Lindem, Manish Gupta, Balaji Ganesh, Les Ginsberg, Robert Raszuk, Gyan Mishra, Muthu Arul Mozhi Perumal, Russ Housley, and Wes Hardaker.

The authors would like to acknowledge Dylan van Oudheusden for highlighting the problems in using OSPF BFD strict-mode for BFD sessions for OSPFv3 IPv4 AF instances and Baalajee S for his suggestions on the approach to address it.

The authors would like to thank John Scudder for his AD review and suggestions to improve the document.

Authors' Addresses

Ketan Talaulikar (EDITOR)
Cisco Systems, Inc.
India
Email: ketant.ietf@gmail.com

Peter Psenak
Cisco Systems, Inc.
Apollo Business Center
Mlynske nivy 43
821 09 Bratislava
Slovakia
Email: ppsenak@cisco.com

Albert Fu
Bloomberg
United States of America
Email: afu14@bloomberg.net

Rajesh M
Juniper Networks
India
Email: mrajesh@juniper.net