Link State Routing
Internet-Draft
Intended status: Standards Track
Expires: September 25, 2021

K. Talaulikar
P. Psenak
Cisco Systems, Inc.
A. Fu
Bloomberg
M. Rajesh
Juniper Networks
March 24, 2021

OSPF Strict-Mode for BFD
draft-ietf-lsr-ospf-bfd-strict-mode-03

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 of strict-mode for BFD session establishment for OSPF adjacency. If both OSPF neighbors advertise the strict-mode for BFD, adjacency formation will be blocked until a BFD session has been successfully established.

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 https://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 September 25, 2021.

Copyright Notice

Table of Contents

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 and trigger the rerouting of traffic around the failure faster than with OSPF hello packet monitoring.

The use of BFD for monitoring routing protocols adjacencies is described in [RFC5882]. When BFD monitoring is enabled for OSPF adjacencies, 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. Note that it is possible in some failure scenarios for the network to be in a state such that an OSPF adjacency can be established but a BFD session cannot be established and maintained. In certain other scenarios, a degraded or poor quality link will allow OSPF adjacency formation to succeed but the BFD session establishment will fail or the BFD session will flap. In this case, traffic that gets forwarded over such a link may

experience packet drops while the failure of the BFD session establishment would not enable fast routing convergence if the link were to go down or flap.

To avoid the routing churn associated with these scenarios, it would be beneficial to not 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 in which not all OSPF routers support BFD and are enabled for BFD 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".

This document specifies the OSPF protocol extensions using link-local signaling (LLS) [RFC5613] for a router to indicate to its neighbor the willingness to establish its adjacency using the strict-mode for BFD. It also introduces an extension for OSPFv3 link-local signaling of the interface IPv4 address when used for an IPv4 address-family (AF) instance to enable discovery of the IPv4 addresses for BFD session setup.

A similar functionality for IS-IS is specified [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 field. This bit is defined for the LLS block included in Hello and Database Description (DD) packets and indicates that BFD is enabled on the link and that the router requests strict-mode for BFD. Section 7 describes the position of the B-bit.

A router MUST include the LLS block with the LLS Type 1 Extended Options and Flags TLV with the B-bit set in its Hello and DD packets when strict-mode for BFD 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. It has the following format:

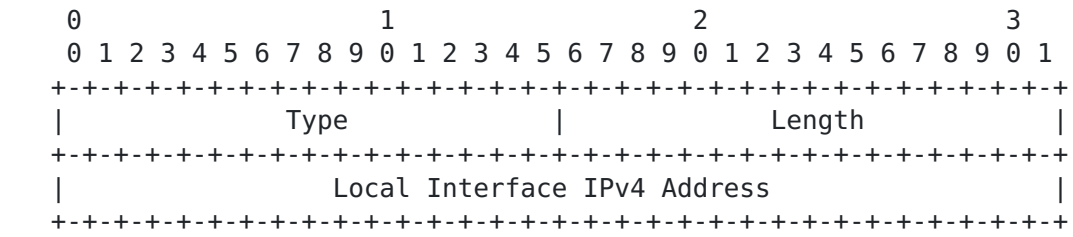

where:

Type: 21

Length: 4 octet

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

4. Procedures

A router supporting strict-mode for BFD advertises this capability through its Hello packets as described in Section 2. When a router supporting strict-mode for BFD discovers a new neighbor router that also supports strict-mode for BFD, then it will establish a BFD session first with that neighbor 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 as below when strict-mode for BFD is used:

Init (without strict-mode for BFD)

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 strict-mode for BFD)

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 not already completed (e.g., in the event of transition from 2-way state). Neighbors in Init state or higher will be listed in the Hello packets associated with the interface if they either have a corresponding BFD session established or have not advertised strict-mode for BFD in the Hello packet LLS Extended Options and Flags.

Whenever the neighbor state transitions to Down state, the removal of the BFD session associated with that neighbor SHOULD be requested by OSPF and subsequent BFD session establishment SHOULD similarly be requested by OSPF upon transitioning into Init state. This may result in the deletion and creation of the BFD session 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 strict-mode for BFD is enabled on the router and the specific neighbor indicates strict-mode for BFD capability via its Hello LLS options. When BFD is enabled, but the strict-mode for operation has not be signaled by both neighbors, then an implementation SHOULD start the BFD session establishment only in 2-Way state or higher 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 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 strict-mode for BFD) on an OSPF router would result in it not setting the B-bit in its subsequent Hello LLS options. Disabling strict-mode for BFD has no effect on the BFD operations and would not result in bringing down of any established BFD session. Disabling BFD would result in the BFD session being brought down due to Admin reason [RFC5882] and hence 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. If the adjacency is already up (i.e., in its terminal state of Full or 2-way with non-DR routers on a multi-access interface) with a neighbor that also supports strict-mode for BFD, then an implementation SHOULD NOT bring this

adjacency down but instead use the strict-mode for BFD operation after the next transition into Init state. However, if the adjacency is not up, then an implementation MAY bring such an adjacency down so it can use the strict-mode for BFD for its adjacency establishment.

4.1. OSPFv3 IPv4 Address-Family Specifics

Multiple AF support 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 AF, it is required that BFD is used to monitor and verify the IPv4 data plane connectivity between the routers on the link and, hence, the BFD session is setup 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 is received. This results in the setup of the BFD session either after the adjacency is established or later in the adjacency formation sequence.

To enable operation in strict-mode for BFD, 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 of the OSPFv3 Hello packets for IPv4 AF instances makes this possible. Implementations that support strict-mode for BFD operation for OSPFv3 IPv4 AF instances MUST include the Local Interface IPv4 Address TLV in the LLS block of their Hello packets whenever the B-bit is also set in the LLS Options and Flags field. A receiver MUST ignore the B-bit (i.e., not operate in BFD strict mode) when the Local Interface IPv4 Address TLV is not present in OSPFv3 Hello message for IPv4 AF OSPFv3 instances.

4.2. Graceful Restart Considerations

An implementation needs to handle scenarios where both graceful restart (GR) and the strict-mode for BFD operation are deployed together. The GR aspects discussed in [RFC5882] also apply with strict-mode for BFD operation. Additionally, in strict-mode for BFD operation, since the OSPF adjacency formation is delayed until the BFD session establishment, 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 & Management Considerations

An implementation SHOULD report the BFD session status along with the OSPF Init adjacency state when operating in strict-mode for BFD and perform logging operations on state transitions to 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 links or those with packet loss, BFD sessions may flap frequently. In such scenarios, OSPF strict-mode for BFD may be deployed in conjunction with a BFD dampening or hold-down mechanism to avoid frequent adjacency flaps that cause routing churn.

6. Backward Compatibility

An implementation MUST support OSPF adjacency formation and operations with a neighbor router that does not advertise the strict-mode for BFD capability - both when that neighbor router does not support BFD and when it does support BFD but not in the strict-mode of operation as described in this document. Implementations MAY provide an option to specifically enable BFD operation only in the strict-mode. In this case, an OSPF adjacency with a neighbor that does not support strict-mode for BFD would not be established successfully. Implementations MAY provide an option to disable strict-mode for BFD 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 of Hello packets from its neighbors and continue to establish BFD sessions, if enabled, without delaying the OSPF adjacency formation. Since the router that does not support this specification would not have set the B-bit in the LLS block of its own Hello packets, its neighbor routers that support this specification would not use strict-mode for BFD with such OSPF routers. As a result, the behavior would be the same as before this specification. Therefore, there are no backward compatibility issues or implementations considerations beyond what is specified herein.

7. IANA Considerations

This specification updates Link Local Signaling TLV Identifiers registry.

Following values have been assigned via early allocation:

o B-bit from "LLS Type 1 Extended Options and Flags" registry at bit position 0x00000010.

o Type 21 - 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 failure to detect bidirectional forwarding failures. If authentication is being used in the OSPF routing domain [RFC5709][RFC7474], then the Cryptographic Authentication TLV [RFC5613] SHOULD also be used to protect the contents of the LLS block.

9. Acknowledgements

The authors would like to acknowledge the review and inputs from Acee Lindem, Manish Gupta and Balaji Ganesh.

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

10. References

10.1. 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>.

10.2. 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>.

[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>.

Authors' Addresses

Ketan Talaulikar
Cisco Systems, Inc.
India

Email: ketant@cisco.com

Peter Psenak
Cisco Systems, Inc.
Apollo Business Center
Mlynske nivy 43
Bratislava 821 09
Slovakia

Email: ppsenak@cisco.com

Albert Fu
Bloomberg
USA

Email: afu14@bloomberg.net

Rajesh M
Juniper Networks
India

Email: mrajesh@juniper.net