

Internet Engineering Task Force (IETF)
Request for Comments: 8770
Updates: 6987
Category: Standards Track
ISSN: 2070-1721

K. Patel
Arrcus
P. Pillay-Esnault
PPE Consulting
M. Bhardwaj
S. Bayraktar
Cisco Systems
April 2020

Host Router Support for OSPFv2

Abstract

The Open Shortest Path First Version 2 (OSPFv2) protocol does not have a mechanism for a node to repel transit traffic if it is on the shortest path. This document defines a bit called the Host-bit (H-bit). This bit enables a router to advertise that it is a non-transit router. This document also describes the changes needed to support the H-bit in the domain. In addition, this document updates RFC 6987 to advertise Type 2 External and Not-So-Stubby Area (NSSA) Link State Advertisements (LSAs) (RFC 3101) with a high cost in order to repel traffic effectively.

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/rfc8770>.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
2. Requirements Language
3. Host-Bit Support
4. SPF Modifications
5. Autodiscovery and Backward Compatibility
6. OSPF AS-External-LSAs / NSSA-LSAs with Type 2 Metrics
7. IANA Considerations

| | |
|------|-------------------------|
| 8. | Security Considerations |
| 9. | References |
| 9.1. | Normative References |
| 9.2. | Informative References |
| | Acknowledgements |
| | Authors' Addresses |

1. Introduction

The OSPFv2 protocol specifies a Shortest Path First (SPF) algorithm that identifies transit vertices based on their adjacencies. Therefore, OSPFv2 does not have a mechanism to prevent traffic transiting a participating node if it is a transit vertex in the only existing or shortest path to the destination. The use of metrics to make the node undesirable can help to repel traffic only if an alternative better route exists.

A mechanism to move traffic away from the shortest path is particularly useful for a number of use cases:

1. Graceful isolation of a router, to avoid blackhole scenarios when there is a reload and possible long reconvergence times.
2. Closet switches that are not usually used for transit traffic but need to participate in the topology.
3. Overloaded routers that could use such a capability to temporarily repel traffic until they stabilize.
4. BGP route reflectors, known as virtual Route Reflectors, that are not in the forwarding path but are in central locations such as data centers. Such route reflectors are typically used for route distribution and are not capable of forwarding transit traffic. However, they need to learn the OSPF topology to perform SPF computation for optimal routes and reachability resolution for their clients [BGP-ORR].

This document describes the functionality provided by the Host-bit (H-bit); this functionality prevents other OSPFv2 routers from using the host router by excluding it in path calculations for transit traffic in OSPFv2 routing domains. If the H-bit is set, then the calculation of the shortest-path tree for an area, as described in Section 16.1 of [RFC2328], is modified by including a check to verify that transit vertices DO NOT have the H-bit set (see Section 4). Furthermore, in order to repel traffic effectively, this document updates [RFC6987] so that Type 2 External and Not-So-Stubby Area (NSSA) Link State Advertisements (LSAs) [RFC3101] are advertised with a high cost (see Section 6). OSPFv3 [RFC5340] defines an option bit, known as the R-bit, for router-LSAs; the H-bit supports similar functionality.

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. Host-Bit Support

This document defines a new router-LSA bit, known as the Host-bit or the H-bit. An OSPFv2 router advertising a router-LSA with the H-bit set indicates that it MUST NOT be used as a transit router (see Section 4) by other OSPFv2 routers in the area that support the H-bit functionality.

If the H-bit is not set, then backward compatibility is achieved, as the behavior will be the same as in [RFC2328].

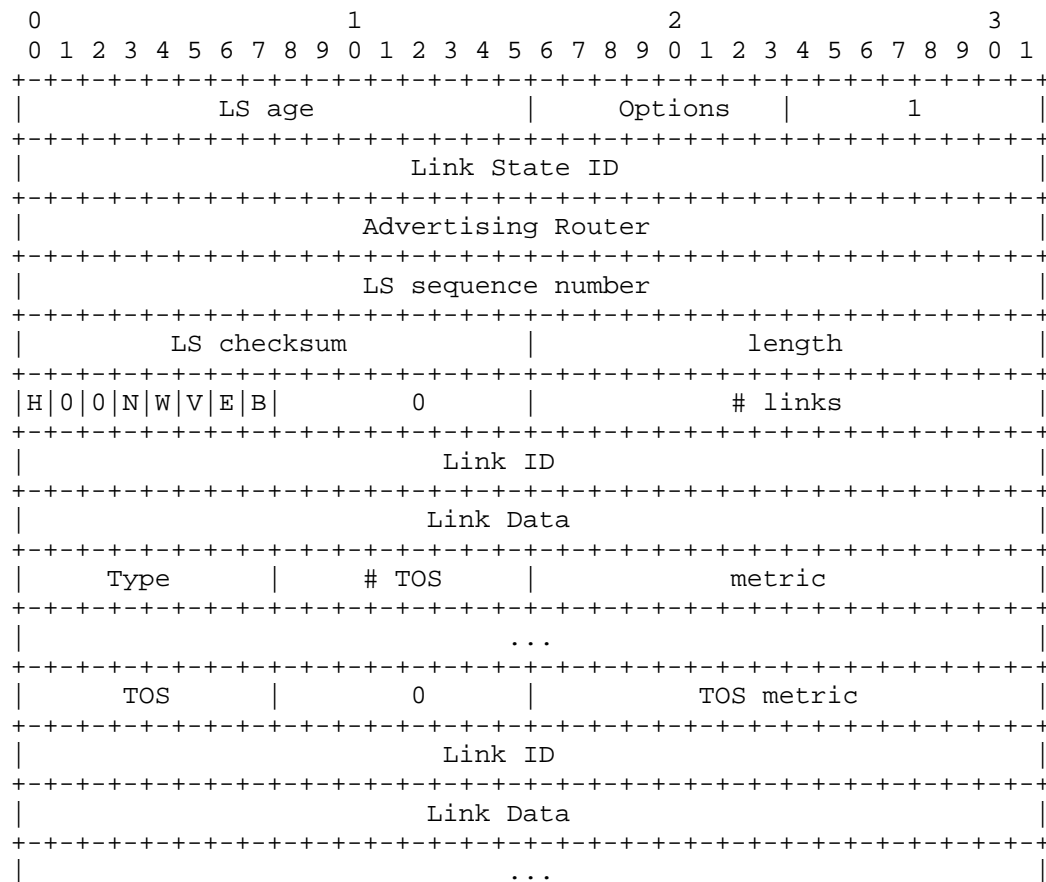


Figure 1: OSPF Router-LSA

Bit H is the high-order bit of the OSPF flags, as shown below.




Figure 2: OSPF Router-LSA Option Bits

When the H-bit is set, the OSPFv2 router is a host (non-transit) router and is incapable of forwarding transit traffic. In this mode, the other OSPFv2 routers in the area MUST NOT use the host router for transit traffic but may send traffic to its local destinations.

An OSPFv2 router originating a router-LSA with the H-bit set MUST advertise all its non-stub links with a link cost of MaxLinkMetric [RFC6987].

When the H-bit is set, an Area Border Router (ABR) MUST advertise the same H-bit setting in its self-originated router-LSAs for all attached areas. The consistency of the setting will prevent inter-area traffic transiting through the router by suppressing advertisements of prefixes from other routers in the area in its summary-LSAs. Only IPv4 prefixes associated with its local interfaces MUST be advertised in summary-LSAs to provide reachability to end hosts attached to a router with the H-bit set.

When the H-bit is set, the host router cannot act as an Autonomous System Border Router (ASBR). Indeed, ASBRs are transit routers to

prefixes that are typically imported through redistribution of prefixes from other routing protocols. Therefore, non-local IPv4 prefixes, e.g., those imported from other routing protocols, SHOULD NOT be advertised in AS-external-LSAs if the H-bit is set. Some use cases, such as an overloaded router or a router being gracefully isolated, may benefit from continued advertisements of non-local prefixes. In these cases, the Type 2 metric in AS-external-LSAs MUST be set to LSInfinity [RFC2328] to repel traffic (see Section 6 of this document).

4. SPF Modifications

The SPF calculation described in Section 16.1 of [RFC2328] is modified to ensure that the routers originating router-LSAs with the H-bit set will not be used for transit traffic. Step (2) is modified to include a check on the H-bit, as shown below. (Please note that all of the sub-procedures of Step (2) remain unchanged and are not included in the excerpt below.)

- (2) Call the vertex just added to the tree "vertex V". Examine the LSA associated with vertex V. This is a lookup in Area A's link state database based on the Vertex ID. If this is a router-LSA, and the H-bit of the router-LSA is set, and vertex V is not the root, then the router should not be used for transit and Step (3) should be executed immediately. If this is a router-LSA and bit V of the router-LSA (see Appendix A.4.2) is set, set Area A's TransitCapability to TRUE. In any case, each link described by the LSA gives the cost to an adjacent vertex. For each described link (say it joins vertex V to vertex W):

5. Autodiscovery and Backward Compatibility

To reduce the possibility of any routing loops due to partial deployment, this document defines an OSPF Router Information (RI) LSA capability bit [RFC7770]. See Section 7 (Table 2).

The RI LSA MUST be area-scoped.

Autodiscovery via announcement of the OSPF Host Router capability (Section 7) ensures that the H-bit functionality and its associated SPF changes MUST only take effect if all the routers in a given OSPF area support this functionality.

In normal operation, it is possible that the RI LSA will fail to reach all routers in an area in a timely manner. For example, if a new router without H-bit support joins an area that previously had only H-bit-capable routers with the H-bit set, then it may take some time for the RI LSA to propagate to all routers. While it is propagating, the routers in the area will gradually detect the presence of a router that does not support the capability and will revert back to the normal SPF calculation. During the propagation time, the area as a whole is unsure of the status of the new router; this type of situation can cause temporary transient loops.

The following recommendations will mitigate transient routing loops:

- * Implementations are RECOMMENDED to provide a configuration parameter to manually override enforcement of the H-bit functionality in partial deployments where the topology guarantees that OSPFv2 routers not supporting the H-bit do not compute routes resulting in routing loops.
- * All routers with the H-bit set MUST advertise all of the router's non-stub links with a metric equal to MaxLinkMetric [RFC6987] in its LSAs in order to prevent OSPFv2 routers (unless a last-resort

path) that do not support the H-bit from attempting to use the non-stub links for transit traffic.

- * All routers supporting the H-bit MUST check the RI LSAs of all nodes in the area to verify that all nodes support the H-bit before actively using the H-bit feature. If any router does not advertise the OSPF Host Router capability (Section 7), then the SPF modifications described in Section 4 MUST NOT be used in the area.

6. OSPF AS-External-LSAs / NSSA-LSAs with Type 2 Metrics

When calculating the path to a prefix in an OSPF AS-external-LSA or NSSA-LSA [RFC3101] with a Type 2 metric, the advertised Type 2 metric is taken as more significant than the OSPF intra-area or inter-area path. Hence, advertising the links with MaxLinkMetric as specified in [RFC6987] does not discourage transit traffic when calculating AS-external or NSSA routes with Type 2 metrics.

Consequently, this document updates [RFC6987] so that the Type 2 metric in any self-originated AS-external-LSAs or NSSA-LSAs is advertised as LSInfinity-1 [RFC2328]. If the H-bit is set, then the Type 2 metric MUST be set to LSInfinity.

7. IANA Considerations

IANA has registered the following value in the "OSPFv2 Router Properties Registry".

| Value | Description | Reference |
|-------|--------------|-----------|
| 0x80 | Host (H-bit) | RFC 8770 |

Table 1: H-Bit

IANA has registered the following in the "OSPF Router Informational Capability Bits" registry.

| Bit Number | Capability Name | Reference |
|------------|------------------|-----------|
| 7 | OSPF Host Router | RFC 8770 |

Table 2: OSPF Host Router Capability Bit

8. Security Considerations

This document introduces the H-bit, which is a capability feature that restricts the use of a router for transit, while only its local destinations are reachable. This is a subset of the operations of a normal router and therefore should not introduce new security considerations beyond those already known in OSPFv2 [RFC2328]. The feature introduces the advertisement of host router capability information to all OSPFv2 routers in an area. This information can be leveraged for discovery and verification that all routers in the area support the capability before the feature is turned on. In the event that a rogue or buggy router incorrectly advertises its capability, possible scenarios are as follows:

- * The router does not have the capability but sends the H-bit set in its LSAs. In this case, a routing loop is possible. However, this is mitigated by the fact that this router should be avoided anyway. Moreover, the link metrics cost (MaxLinkMetric) of this

router will mitigate this situation. In any case, a router advertising the H-bit capability without its link metrics cost equal to MaxLinkMetric could be a rogue router and should be avoided.

- * The router has the capability but sends the H-bit clear in its LSAs. In this case, the router merely prevents the support of other H-bit routers in the area and prevents all the routers from running the modified SPF. Any impacts are also mitigated in this scenario, as other H-bit routers in the area also advertise the MaxLinkMetric cost, so they will still be avoided unless they are the last-resort path.
- * The rogue router is on the only transit path for some destinations and sends the H-bit set (for no good/valid reason) in its LSAs, and effectively partitions the network. This case is indistinguishable from the normal case where an operator may consciously decide to set the H-bit to perform maintenance on a router that is on the only transit path. The OSPF protocol will continue to function within the partitioned domains.

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>>.
- [RFC6987] Retana, A., Nguyen, L., Zinin, A., White, R., and D. McPherson, "OSPF Stub Router Advertisement", RFC 6987, DOI 10.17487/RFC6987, September 2013, <<https://www.rfc-editor.org/info/rfc6987>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <<https://www.rfc-editor.org/info/rfc7770>>.
- [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

- [BGP-ORR] Raszuk, R., Ed., Cassar, C., Aman, E., Decraene, B., and K. Wang, "BGP Optimal Route Reflection (BGP-ORR)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-optimal-route-reflection-20, 8 January 2020, <<https://tools.ietf.org/html/draft-ietf-idr-bgp-optimal-route-reflection-20>>.
- [RFC3101] Murphy, P., "The OSPF Not-So-Stubby Area (NSSA) Option", RFC 3101, DOI 10.17487/RFC3101, January 2003, <<https://www.rfc-editor.org/info/rfc3101>>.
- [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>>.

Acknowledgements

The authors would like to acknowledge Hasmit Grover for discovering the limitation in [RFC6987], and Acee Lindem, Abhay Roy, David Ward, Burjiz Pithawala, and Michael Barnes for their comments.

Authors' Addresses

Keyur Patel
Arrcus

Email: keyur@arrcus.com

Padma Pillay-Esnault
PPE Consulting

Email: padma.ietf@gmail.com

Manish Bhardwaj
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America

Email: manbhard@cisco.com

Serpil Bayraktar
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
United States of America

Email: serpil@cisco.com