

OSPF  
Internet-Draft  
Updates: [2328](#) (if approved)  
Intended status: Standards Track  
Expires: February 28, 2019

K. Patel  
Arrcus  
P. Pillay-Esnault  
Huawei Technologies  
M. Bhardwaj  
S. Bayraktar  
Cisco Systems  
August 27, 2018

H-bit Support for OSPFv2  
draft-ietf-ospf-ospfv2-hbit-06

## Abstract

OSPFv3 defines an option bit for router-LSAs known as the R-bit in [RFC5340](#). If the R-bit is clear, an OSPFv3 router can participate in OSPF topology flooding, however it will not be used as a transit router. In such cases, other routers in the OSPFv3 routing domain only install routes to allow local traffic delivery. This document defines the H-bit functionality to prevent other OSPFv2 routers from using the router for transit traffic in OSPFv2 routing domains as described in [RFC 2328](#). This document updates [RFC 2328](#).

## 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 February 28, 2019.

## Copyright Notice

Copyright (c) 2018 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

Internet-Draft

August 2018

(<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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

## Table of Contents

|                       |   |                   |
|-----------------------|---|-------------------|
| <a href="#">1.</a>    | Introduction . . . . .  | <a href="#">2</a> |
| <a href="#">2.</a>    | Requirements Language . . . . .                               | <a href="#">3</a> |
| <a href="#">3.</a>    | H-bit Support . . . . .                                       | <a href="#">3</a> |
| <a href="#">4.</a>    | SPF Modifications . . . . .                                   | <a href="#">5</a> |
| <a href="#">5.</a>    | Auto Discovery and Backward Compatibility . . . . .           | <a href="#">5</a> |
| <a href="#">6.</a>    | OSPF AS-External-LSAs/NSSA LSAs with Type 2 Metrics . . . . . | <a href="#">6</a> |
| <a href="#">7.</a>    | IANA Considerations . . . . .                                 | <a href="#">6</a> |
| <a href="#">8.</a>    | Security Considerations . . . . .                             | <a href="#">7</a> |
| <a href="#">9.</a>    | Acknowledgements . . . . .                                    | <a href="#">7</a> |
| <a href="#">10.</a>   | References . . . . .  | <a href="#">7</a> |
| <a href="#">10.1.</a> | Normative References . . . . .                                | <a href="#">7</a> |
| <a href="#">10.2.</a> | Informative References . . . . .                              | <a href="#">7</a> |
|                       | Authors' Addresses . . . . .                                  | <a href="#">8</a> |

## [1.](#) Introduction

OSPFv3 [[RFC5340](#)] defines an option bit for router-LSAs known as the R-bit. If the R-bit is clear, an OSPFv3 router can participate in OSPFv3 topology flooding without acting as a transit router. In such cases, other routers in the OSPFv3 routing domain only install routes used for local traffic.

This functionality is particularly useful for BGP Route Reflectors, known as virtual Route Reflectors (vRRs), that are not in the forwarding path but are in central locations such as data centers. Such Route Reflectors typically are used for route distribution and

Internet-Draft

August 2018

are not capable of forwarding transit traffic. However, they need to learn the OSPF topology for:

1. SPF computation for Optimal Route Reflection functionality as defined in [[I-D.ietf-idr-bgp-optimal-route-reflection](#)]
2. Reachability resolution for its Route Reflector Clients.

This document defines the R-bit functionality equivalent for OSPFv2 defined in [[RFC2328](#)] by introducing a new router-LSA bit known as the "H-bit". This document updates [appendix A.4.2 of RFC 2328](#).

## [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.](#) H-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 to other OSPFv2 routers in the area supporting the functionality that it MUST NOT be used as a transit router. The bit value usage of the H-bit is reversed from the R-bit defined in OSPFv3 [[RFC5340](#)] to support backward compatibility. The modified OSPFv2 router-LSA format is:

Internet-Draft

August 2018

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|               LS age               | Options | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Link State ID               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Advertising Router               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               LS sequence number               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               LS checksum               | length |
+-----+-----+-----+-----+-----+-----+-----+-----+
| H|0|0|N|W|V|E|B| 0 | # links |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Link ID               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Link Data               |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type | # TOS | metric |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               ...               |
+-----+-----+-----+-----+-----+-----+-----+-----+
| TOS | 0 | TOS metric |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Link ID               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Link Data               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

bit H

When set, an OSPFv2 router is a non-transit router and is incapable of forwarding transit traffic.

When the H-bit is set, an OSPFv2 router is a non-transit router and should not be used to forward transit traffic. In this mode, the other OSPFv2 routers in the area SHOULD NOT use the originating OSPFv2 router for transit traffic, but MAY use the OSPFv2 router for local traffic destined to that OSPFv2 router.

An OSPFv2 router originating a router-LSA with the H-bit set SHOULD advertise all its non-local router links with a link cost of MaxLinkMetric as defined in [Section 3 of \[RFC6987\]](#). This is to increase the applicability of the H-bit to partial deployments where it is the responsibility of the operator to ensure that OSPFv2

routers not supporting the H-bit do not install routes causing routing loops.

When the H-bit is set, IPv4 prefixes associated with local interfaces in other areas MAY be advertised in summary LSAs. Non-local IPv4 prefixes, e.g., those advertised by other routers and installed during the SPF computation, MAY be advertised in summary-LSAs if configured by policy. Likewise, when the H-bit is set, only IPv4 prefixes associated with local interfaces MAY be advertised in AS-external LSAs. Non-local IPv4 prefixes, e.g., those exported from other routing protocols, MUST NOT be advertised in AS-external-LSAs. Finally, when the H-bit is set, an Area Border Router (ABR) MUST advertise a consistent H-bit setting in its self-originated router-LSAs for all attached areas.

#### [4.](#) SPF Modifications

The SPF calculation described in [section 16.1 \[RFC2328\]](#) will be 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 as follows:

- 2) Call the vertex just added to the tree vertex V. Examine the LSA associated with vertex V. This is a lookup in the 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 Section 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.](#) Auto Discovery and Backward Compatibility

To avoid the possibility of any routing loops due to partial deployment, this document defines a OSPF Router-Information LSA functional capability bit known as the Host Support capability.

Auto Discovery via announcement of the Host Support Functional Capability ensures that the H-bit functionality and its associated SPF changes SHOULD only take effect if all the routers in a given OSPF area support this functionality.

Implementations are encouraged 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. More precisely, the advertisement of MaxLinkMetric for the router's non-local links will prevent OSPFv2 routers not supporting the H-bit from attempting to use it for transit traffic.

## [6.](#) OSPF AS-External-LSAs/NSSA LSAs with Type 2 Metrics

When calculating the path to an OSPF AS-External-LSA or NSSA-LSA 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. Consequently, OSPF routers implementing [\[RFC6987\]](#) or this specification should advertise a Type-2 metric of LSInfinity for any self-originated AS-External-LSAs or NSSA-LSAs in situations when the OSPF router is acting as a stub router [\[RFC6987\]](#) or implementing this specification.

## [7.](#) IANA Considerations

IANA is requested to create the OSPF Router-LSA bit registry with the following assignments:

| Value | Description                            | Reference                 |
|-------|--|---------------------------|
| 0x01  | Area Border Router (B-bit)             | <a href="#">[RFC2328]</a> |
| 0x02  | AS Boundary Router (E-bit)             | <a href="#">[RFC2328]</a> |
| 0x04  | Virtual Link Endpoint (V-bit)          | <a href="#">[RFC2328]</a> |
| 0x08  | Historic (W-bit)                       | <a href="#">[RFC1584]</a> |
| 0x10  | Unconditional NSSA Translator (Nt-bit) | <a href="#">[RFC3101]</a> |
| 0x20  | Unassigned                             |                           |
| 0x40  | Unassigned                             |                           |
| 0x80  | Host (H-bit)                           | This Document             |

This document also defines a new Router Functional Capability [\[RFC7770\]](#) known as the Host Support Functional Capability. This document requests IANA to allocate the value of this capability from the Router Functional Capability Bits TLV.

## [8.](#) Security Considerations

This document introduces no new security considerations beyond those already specified in [\[RFC6987\]](#), [\[RFC2328\]](#), and [\[RFC5340\]](#).

## [9.](#) Acknowledgements

The authors would like to acknowledge Hasmit Grover for discovery of

the limitation in [RFC6987], Acee Lindem, Abhay Roy, David Ward, Burjiz Pithawala and Michael Barnes for their comments.

## 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>>.
- [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>>.
- [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>>.

### 10.2. Informative References

- [I-D.ietf-idr-bgp-optimal-route-reflection]  
Raszuk, R., Cassar, C., Aman, E., Decraene, B., and K. Wang, "BGP Optimal Route Reflection (BGP-ORR)", [draft-ietf-idr-bgp-optimal-route-reflection-16](#) (work in progress), April 2018.

- [RFC1584] Moy, J., "Multicast Extensions to OSPF", [RFC 1584](#),



DOI 10.17487/RFC1584, March 1994,  
<<https://www.rfc-editor.org/info/rfc1584>>.

[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>>.

#### Authors' Addresses

Keyur Patel  
Arrcus

Email: [keyur@arrcus.com](mailto:keyur@arrcus.com)

Padma Pillay-Esnault  
Huawei Technologies  
2330 Central Expressway  
Santa Clara, CA 95050  
USA

Email: [padma@huawei.com](mailto:padma@huawei.com)

Manish Bhardwaj  
Cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
USA

Email: [manbhard@cisco.com](mailto:manbhard@cisco.com)

Serpil Bayraktar  
Cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
USA

Email: [serpil@cisco.com](mailto:serpil@cisco.com)