ï»¿

Internet Engineering Task Force (IETF)
Request for Comments: 9084
Category: Standards Track
ISSN: 2070-1721

A. Wang
China Telecom
A. Lindem
Cisco Systems, Inc.
J. Dong
Huawei Technologies
P. Psenak
K. Talaulikar, Ed.
Cisco Systems, Inc.

August 2021

OSPF Prefix Originator Extensions

Abstract

This document defines OSPF extensions to include information associated with the node originating a prefix along with the prefix advertisement. These extensions do not change the core OSPF route computation functionality but provide useful information for network analysis, troubleshooting, and use cases like traffic engineering.

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/rfc9084.

Table of Contents

1. Introduction

Prefix attributes are advertised in OSPFv2 [RFC2328] using the Extended Prefix Opaque Link State Advertisement (LSA) [RFC7684] and in OSPFv3 [RFC5340] using the various Extended Prefix LSA types [RFC8362].

The procedures for identification of the originating router for a prefix in OSPF vary by the type of the prefix and, currently, it is not always possible to produce an accurate result. For intra-area prefixes, the originating router is identified by the Advertising Router field of the area-scoped LSA used for those prefix advertisements. However, for inter-area prefixes advertised by an Area Border Router (ABR), the Advertising Router field of their area-scoped LSAs is set to the ABR itself and the information about the router originating the prefix advertisement is lost in the process of prefix propagation across areas. For Autonomous System (AS) external prefixes, the originating router may be considered as the Autonomous System Border Router (ASBR) and is identified by the Advertising Router field of the AS-scoped LSA used. However, the actual originating router for the prefix may be a remote router outside the OSPF domain. Similarly, when an ABR performs translation of Not-So-Stubby Area (NSSA) [RFC3101] LSAs to AS-external LSAs, the information associated with the NSSA ASBR (or the router outside the OSPF domain) is not propagated across the OSPF domain.

While typically the originator of information in OSPF is identified by its OSPF Router ID, it does not necessarily represent a reachable address for the router since the OSPF Router ID is a 32-bit number that is unique in the OSPF domain. For OSPFv2, a common practice is to use one of the IPv4 addresses of the node (e.g., a loopback interface) as the OSPF Router ID. However, this cannot always be assumed and this approach does not apply to IPv6 addresses with OSPFv3. The IPv4/IPv6 Router Address, as respectively defined in [RFC3630] and [RFC5329] for OSPFv2 and OSPFv3, provides an address to reach the advertising router.

The primary use case for the extensions proposed in this document is to be able to identify the originator of a prefix in the network. In cases where multiple prefixes are advertised by a given router, it is also useful to be able to associate all these prefixes with a single router even when prefixes are advertised outside of the area in which they originated. It also helps to determine when the same prefix is being originated by multiple routers across areas.

This document proposes extensions to the OSPF protocol for the inclusion of information associated with the router originating the prefix along with the prefix advertisement. These extensions do not change the core OSPF route computation functionality. They provide useful information for topology analysis and traffic engineering, especially on a controller when this information is advertised as an attribute of the prefixes via mechanisms such as Border Gateway Protocol - Link State (BGP-LS) [RFC7752] [RFC9085].

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. Protocol Extensions

This document defines the Prefix Source OSPF Router-ID and the Prefix Source Router Address Sub-TLVs. They are used, respectively, to include the Router ID of, and a reachable address of, the router that originates the prefix as a prefix attribute.

2.1. Prefix Source OSPF Router-ID Sub-TLV

For OSPFv2, the Prefix Source OSPF Router-ID Sub-TLV is an optional sub-TLV of the OSPFv2 Extended Prefix TLV [RFC7684]. For OSPFv3, the Prefix Source OSPF Router-ID Sub-TLV is an optional sub-TLV of the Intra-Area-Prefix TLV, Inter-Area-Prefix TLV, and External-Prefix TLV [RFC8362] when originating either an IPv4 [RFC5838] or an IPv6 prefix advertisement.

The Prefix Source OSPF Router-ID Sub-TLV has the following format:

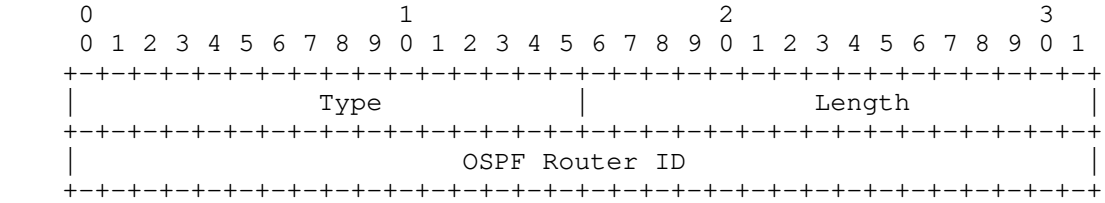

Figure 1: Prefix Source OSPF Router-ID Sub-TLV Format

Where:

Type: 4 for OSPFv2 and 27 for OSPFv3

Length: 4

OSPF Router ID: the OSPF Router ID of the OSPF router that originated the prefix advertisement in the OSPF domain

The parent TLV of a prefix advertisement MAY include more than one Prefix Source OSPF Router-ID Sub-TLV, one corresponding to each of the Equal-Cost Multipath (ECMP) nodes that originated the advertised prefix.

For intra-area prefix advertisements, the Prefix Source OSPF Router-ID Sub-TLV MUST be considered invalid and ignored if the OSPF Router ID field is not the same as the Advertising Router field in the containing LSA. Similar validation cannot be reliably performed for inter-area and external prefix advertisements.

A received Prefix Source OSPF Router-ID Sub-TLV with the OSPF Router ID field set to 0 MUST be considered invalid and ignored. Additionally, reception of such sub-TLVs SHOULD be logged as an error (subject to rate limiting).

2.2. Prefix Source Router Address Sub-TLV

For OSPFv2, the Prefix Source Router Address Sub-TLV is an optional sub-TLV of the OSPFv2 Extended Prefix TLV [RFC7684]. For OSPFv3, the Prefix Source Router Address Sub-TLV is an optional sub-TLV of the Intra-Area-Prefix TLV, Inter-Area-Prefix TLV, and External-Prefix TLV [RFC8362] when originating either an IPv4 [RFC5838] or an IPv6 prefix advertisement.

The Prefix Source Router Address Sub-TLV has the following format:

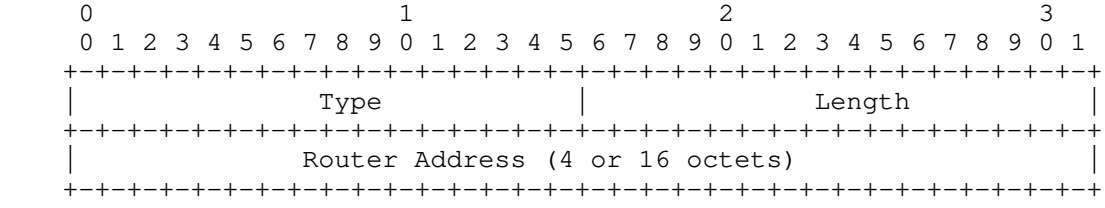

Figure 2: Prefix Source Router Address Sub-TLV Format

Where:

Type: 5 for OSPFv2 and 28 for OSPFv3

Length: 4 or 16

Router Address: A reachable IPv4 or IPv6 router address for the

router that originated the IPv4 or IPv6 prefix advertisement, respectively. Such an address would be semantically equivalent to what may be advertised in the OSPFv2 Router Address TLV [RFC3630] or in the OSPFv3 Router IPv6 Address TLV [RFC5329].

The parent TLV of a prefix advertisement MAY include more than one Prefix Source Router Address Sub-TLV, one corresponding to each of the Equal-Cost Multipath (ECMP) nodes that originated the advertised prefix.

A received Prefix Source Router Address Sub-TLV that has an invalid length (i.e., not consistent with the prefix's address family) MUST be considered invalid and ignored. Additionally, reception of such sub-TLVs SHOULD be logged as an error (subject to rate limiting).

3. Elements of Procedure

This section describes the procedure for the advertisement of the Prefix Source OSPF Router-ID and Prefix Source Router Address Sub-TLVs along with the prefix advertisement.

The OSPF Router ID of the Prefix Source OSPF Router-ID is set to the OSPF Router ID of the node originating the prefix in the OSPF domain.

If the originating node is advertising an OSPFv2 Router Address TLV [RFC3630] or an OSPFv3 Router IPv6 Address TLV [RFC5329], then the same address MUST be used in the Router Address field of the Prefix Source Router Address Sub-TLV. When the originating node is not advertising such an address, implementations can select a unique and reachable local address (for example, advertised with the N-Flag set [RFC7684] or N-bit set [RFC8362]) on the originating node to advertise in the Router Address field.

When an ABR generates inter-area prefix advertisements into its non-backbone areas corresponding to an inter-area prefix advertisement from the backbone area, the only way to determine the originating node information is based on the Prefix Source OSPF Router-ID and Prefix Source Router Address Sub-TLVs present in the inter-area prefix advertisement originated into the backbone area by an ABR from another non-backbone area. The ABR performs its prefix calculation to determine the set of nodes that contribute to ECMP paths for the prefix. It MUST only use the prefix originator information from this set of nodes. The ABR MUST NOT include the Prefix Source OSPF Router-ID or the Prefix Source Router Address Sub-TLVs when it is unable to determine the information for the originating nodes contributing ECMP paths.

Implementations may support the propagation of the originating node information along with a redistributed prefix into the OSPF domain from another routing domain. The details of such mechanisms are outside the scope of this document. Such implementations may also provide control on whether the Router Address in the Prefix Source Router Address Sub-TLV is set as the ASBR node address or as the address of the actual node outside the OSPF domain that owns the prefix.

When translating NSSA prefix advertisements [RFC3101] to AS external prefix advertisements, the NSSA ABR follows the same procedures as an ABR generating inter-area prefix advertisements for the propagation of the originating node information.

4. Security Considerations

Since this document extends the OSPFv2 Extended Prefix LSA, the security considerations for [RFC7684] are applicable. Similarly, since this document extends the OSPFv3 E-Intra-Area-Prefix-LSA, E-Inter-Area-Prefix-LSA, E-AS-External-LSA, and E-NSSA-LSA, the security considerations for [RFC8362] are applicable. The new sub-TLVs introduced in this document are optional and do not affect the OSPF route computation and therefore do not affect the security aspects of OSPF protocol operations.

A rogue node that can inject prefix advertisements may use the extensions introduced in this document to advertise bogus prefix source information.

5. Operational Considerations

Consideration should be given to the operational impact of the increase in the size of the OSPF Link-State Database as a result of the protocol extensions in this document. Based on deployment design and requirements, a subset of prefixes may be identified for which originating node information is required to be included in prefix advertisements.

The propagation of prefix source node information for prefix advertisements advertised across an OSPF area or domain boundaries will expose information outside of an area or domain where it would normally be hidden or abstracted by the base OSPF protocol. Based on deployment design and requirements, the propagation of node information across area or domain boundaries may be limited to a subset of prefixes in the ABRs or ASBRs, respectively.

The identification of the node that is originating a specific prefix in the network may aid in the debugging of issues related to prefix reachability within an OSPF network.

6. IANA Considerations

Per this document, IANA has allocated the following codepoints from the "OSPFv2 Extended Prefix TLV Sub-TLVs" registry under the "Open Shortest Path First v2 (OSPFv2) Parameters" registry.

Value	Description	Reference
4	Prefix Source OSPF Router-ID	RFC 9084
5	Prefix Source Router Address	RFC 9084

Table 1: Codepoints in OSPFv2 Extended Prefix TLV Sub-TLVs

Per this document, IANA has allocated the following codepoints from the "OSPFv3 Extended-LSA Sub-TLVs" registry under the "Open Shortest Path First v3 (OSPFv3) Parameters" registry.

Value	Description	Reference
27	Prefix Source OSPF Router-ID	RFC 9084
28	Prefix Source Router Address	RFC 9084

Table 2: Codepoints in OSPFv3 Extended-LSA Sub-TLVs

7. References

7.1. 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>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering

(TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <https://www.rfc-editor.org/info/rfc3630>.

[RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <https://www.rfc-editor.org/info/rfc5329>.

[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>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <https://www.rfc-editor.org/info/rfc7684>.

[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>.

[RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <https://www.rfc-editor.org/info/rfc8362>.

7.2. Informative References

[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>.

[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>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <https://www.rfc-editor.org/info/rfc9085>.

Acknowledgement

Many thanks to Les Ginsberg for his suggestions on this document. Also, thanks to Jeff Tantsura, Rob Shakir, Gunter Van de Velde, Goethals Dirk, Smita Selot, Shaofu Peng, John E. Drake, and Baalajee S. for their review and valuable comments. The authors would also like to thank Alvaro Retana for his detailed review and suggestions for the improvement of this document.

Authors' Addresses

Aijun Wang
China Telecom
Beiqijia Town
Changping District
Beijing
102209
China

Email: wangaj3@chinatelecom.cn

Acee Lindem
Cisco Systems, Inc.
301 Midenhall Way
Cary, NC 27513
United States of America

Email: acee@cisco.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China

Email: jie.dong@huawei.com

Peter Psenak
Cisco Systems, Inc.
Eurovea Centre, Central 3
Pribinova Street 10
81109 Bratislava
Slovakia

Email: ppsenak@cisco.com

Ketan Talaulikar (editor)
Cisco Systems, Inc.
India

Email: ketant@cisco.com