

Internet Engineering Task Force (IETF)
Request for Comments: 7684
Category: Standards Track
ISSN: 2070-1721

P. Psenak
Cisco Systems
H. Gredler
Independent
R. Shakir
Jive Communications, Inc.
W. Henderickx
Alcatel-Lucent
J. Tantsura
Ericsson
A. Lindem
Cisco Systems
November 2015

OSPFv2 Prefix/Link Attribute Advertisement

Abstract

OSPFv2 requires functional extension beyond what can readily be done with the fixed-format Link State Advertisements (LSAs) as described in RFC 2328. This document defines OSPFv2 Opaque LSAs based on Type-Length-Value (TLV) tuples that can be used to associate additional attributes with prefixes or links. Depending on the application, these prefixes and links may or may not be advertised in the fixed-format LSAs. The OSPFv2 Opaque LSAs are optional and fully backward compatible.

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 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7684>.

Copyright Notice

Copyright (c) 2015 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 (<http://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

1. Introduction	3
1.1. Requirements Notation	3
2. OSPFv2 Extended Prefix Opaque LSA	3
2.1. OSPFv2 Extended Prefix TLV	5
3. OSPFv2 Extended Link Opaque LSA	8
3.1. OSPFv2 Extended Link TLV	9
4. Backward Compatibility	10
5. Security Considerations	10
6. IANA Considerations	11
6.1. OSPFv2 Extended Prefix Opaque LSA TLVs Registry	11
6.2. OSPFv2 Extended Prefix TLV Sub-TLVs Registry	12
6.3. OSPFv2 Extended Prefix TLV Flags Registry	12
6.4. OSPFv2 Extended Link Opaque LSA TLVs Registry	12
6.5. OSPFv2 Extended Link TLV Sub-TLVs Registry	13
7. References	13
7.1. Normative References	13
7.2. Informative References	14
Acknowledgements	14
Authors' Addresses	15

1. Introduction

OSPFv2 requires functional extension beyond what can readily be done with the fixed-format Link State Advertisements (LSAs) as described in RFC 2328 [OSPFV2]. This document defines OSPFv2 Opaque LSAs based on Type-Length-Value (TLV) tuples that can be used to associate additional attributes with prefixes or links. Depending on the application, these prefixes and links may or may not be advertised in the fixed-format LSAs. The OSPFv2 Opaque LSAs are optional and fully backward compatible. This is in contrast to the approach taken in OSPFv3 [OSPFV3-EXTEND] where the existing LSAs will be replaced by TLV-based extended LSAs.

New requirements such as source/destination routing, route tagging, and segment routing necessitate this extension.

This specification defines the following OSPFv2 Opaque LSAs:

1. OSPFv2 Extended Prefix Opaque LSA - Allows advertisement of additional attributes for prefixes advertised in Router-LSAs, Network-LSAs, Summary-LSAs (IP network), NSSA-LSAs, and AS-external-LSAs [OSPFV2][RFC3101].
2. OSPFv2 Extended Link Opaque LSA - Allows advertisement of additional attributes for links advertised in Router-LSAs.

Additionally, the following TLVs are defined:

1. OSPFv2 Extended Prefix TLV - Top-level TLV advertising attributes for a prefix in the OSPFv2 Extended Prefix Opaque LSA.
2. OSPFv2 Extended Link TLV - Top-level TLV advertising attributes for a link in the OSPFv2 Extended Link Opaque LSA.

1.1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [KEYWORDS].

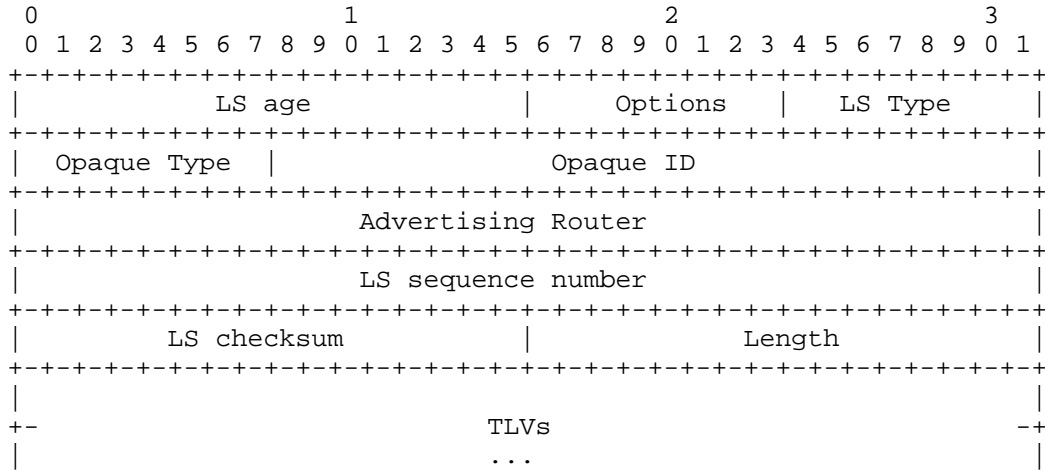
2. OSPFv2 Extended Prefix Opaque LSA

The OSPFv2 Extended Prefix Opaque LSA is used to advertise additional prefix attributes. Opaque LSAs are described in [OPAQUE].

Multiple OSPFv2 Extended Prefix Opaque LSAs can be advertised by an OSPFv2 router. The flooding scope of the OSPFv2 Extended Prefix Opaque LSA depends on the scope of the advertised prefixes and is

under the control of the advertising router. In some cases (e.g., mapping server deployment [SEGMENT-ROUTING]), the LSA flooding scope may be greater than the scope of the corresponding prefixes.

The format of the OSPFv2 Extended Prefix Opaque LSA is as follows:

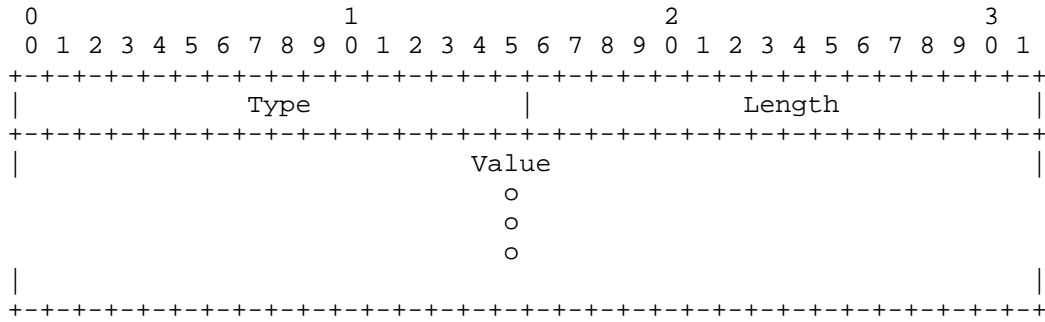


OSPFv2 Extended Prefix Opaque LSA

The Opaque Type used by the OSPFv2 Extended Prefix Opaque LSA is 7. The Opaque Type is used to differentiate the various types of OSPFv2 Opaque LSAs and is described in Section 3 of [OPAQUE]. The LS Type may be 10 or 11, indicating that the Opaque LSA flooding scope is area-local (10) or AS-wide (11) [OPAQUE]. The LSA Length field [OSPFV2] represents the total length (in octets) of the Opaque LSA, including the LSA header and all TLVs (including padding).

The Opaque ID field is an arbitrary value used to maintain multiple OSPFv2 Extended Prefix Opaque LSAs. For OSPFv2 Extended Prefix Opaque LSAs, the Opaque ID has no semantic significance other than to differentiate OSPFv2 Extended Prefix Opaque LSAs originated by the same OSPFv2 router. If multiple OSPFv2 Extended Prefix Opaque LSAs include the same prefix, the attributes from the Opaque LSA with the lowest Opaque ID SHOULD be used.

The format of the TLVs within the body of the OSPFv2 Extended Prefix Opaque LSA is the same as the format used by the Traffic Engineering Extensions to OSPFv2 [TE]. The variable TLV section consists of one or more nested TLV tuples. Nested TLVs are also referred to as sub-TLVs. The format of each TLV is:

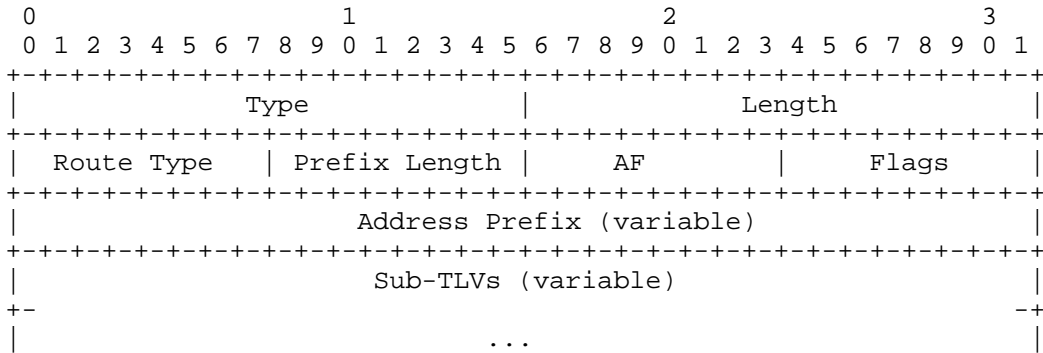


TLV Format

The Length field defines the length of the value portion in octets (thus, a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the Length field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). Nested TLVs are also 32-bit aligned. For example, a 1-byte value would have the Length field set to 1, and 3 octets of padding would be added to the end of the value portion of the TLV. The padding is composed of zeros.

2.1. OSPFv2 Extended Prefix TLV

The OSPFv2 Extended Prefix TLV is used to advertise additional attributes associated with the prefix. Multiple OSPFv2 Extended Prefix TLVs MAY be advertised in each OSPFv2 Extended Prefix Opaque LSA. However, since the Opaque LSA type defines the flooding scope, the LSA flooding scope MUST satisfy the application-specific requirements for all the prefixes included in a single OSPFv2 Extended Prefix Opaque LSA. The OSPFv2 Extended Prefix TLV has the following format:



OSPFv2 Extended Prefix TLV

Type

The TLV type. The value is 1 for this TLV type.

Length

Variable, dependent on sub-TLVs.

Route Type

The type of the OSPFv2 route. If the route type is 0 (Unspecified), the information inside the OSPFv2 External Prefix TLV applies to the prefix regardless of prefix's route type. This is useful when prefix-specific attributes are advertised by an external entity that is not aware of the route type associated with the prefix. Supported types are:

- 0 - Unspecified
- 1 - Intra-Area
- 3 - Inter-Area
- 5 - Autonomous System (AS) External
- 7 - Not-So-Stubby Area (NSSA) External

These route types correspond directly to the OSPFv2 LSAs types as defined in the "OSPFv2 Link State (LS) Type" registry in <http://www.iana.org/assignments/ospfv2-parameters>. Specification of route types other than those defined will prevent correlation with existing OSPFv2 LSAs and is beyond the scope of this specification.

Prefix Length

Length of prefix in bits.

AF

Address family for the prefix. Currently, the only supported value is 0 for IPv4 unicast. The inclusion of address family in this TLV allows for future extension.

Flags

This one-octet field contains flags applicable to the prefix. Supported Flags include:

- 0x80 - A-Flag (Attach Flag): An Area Border Router (ABR) generating an OSPFv2 Extended Prefix TLV for an inter-area prefix that is locally connected or attached in another connected area SHOULD set this flag.

0x40 - N-Flag (Node Flag): Set when the prefix identifies the advertising router, i.e., the prefix is a host prefix advertising a globally reachable address typically associated with a loopback address. The advertising router MAY choose to not set this flag even when the above conditions are met. If the flag is set and the prefix length is not a host prefix, then the flag MUST be ignored. The flag is preserved when the OSPFv2 Extended Prefix Opaque LSA is propagated between areas.

Address Prefix

For the address family IPv4 unicast, the prefix itself is encoded as a 32-bit value. The default route is represented by a prefix of length 0. Prefix encoding for other address families is beyond the scope of this specification.

If this TLV is advertised multiple times for the same prefix in the same OSPFv2 Extended Prefix Opaque LSA, only the first instance of the TLV is used by receiving OSPFv2 routers. This situation SHOULD be logged as an error.

If this TLV is advertised multiple times for the same prefix in different OSPFv2 Extended Prefix Opaque LSAs originated by the same OSPFv2 router, the OSPFv2 advertising router is re-originating OSPFv2 Extended Prefix Opaque LSAs for multiple prefixes and is most likely repacking Extended-Prefix-TLVs in OSPFv2 Extended Prefix Opaque LSAs. In this case, the Extended-Prefix-TLV in the OSPFv2 Extended Prefix Opaque LSA with the smallest Opaque ID is used by receiving OSPFv2 routers. This situation may be logged as a warning.

It is RECOMMENDED that OSPFv2 routers advertising OSPFv2 Extended Prefix TLVs in different OSPFv2 Extended Prefix Opaque LSAs re-originate these LSAs in ascending order of Opaque ID to minimize the disruption.

If this TLV is advertised multiple times for the same prefix in different OSPFv2 Extended Prefix Opaque LSAs originated by different OSPFv2 routers, the application using the information is required to determine which OSPFv2 Extended Prefix Opaque LSA is used. For example, the application could prefer the LSA providing the best path to the prefix.

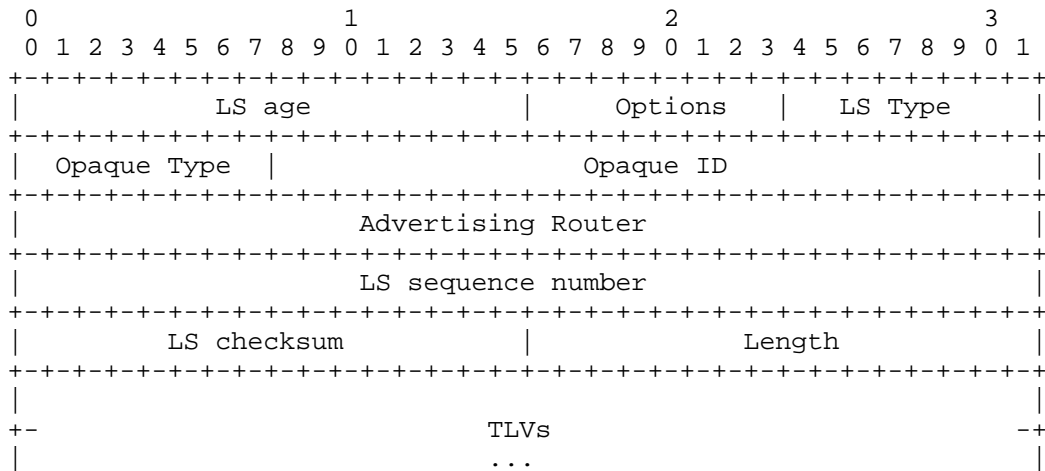
This document creates a registry for OSPFv2 Extended Prefix sub-TLVs in Section 6.

3. OSPFv2 Extended Link Opaque LSA

The OSPFv2 Extended Link Opaque LSA is used to advertise additional link attributes. Opaque LSAs are described in [OPAQUE].

The OSPFv2 Extended Link Opaque LSA has an area flooding scope. Multiple OSPFv2 Extended Link Opaque LSAs can be advertised by a single router in an area.

The format of the OSPFv2 Extended Link Opaque LSA is as follows:



OSPFv2 Extended Link Opaque LSA

The Opaque Type used by the OSPFv2 Extended Link Opaque LSA is 8. The LS Type is 10, indicating that the Opaque LSA flooding scope is area-local [OPAQUE]. The Opaque Type is used to differentiate the various types of OSPFv2 Opaque LSAs and is described in Section 3 of [OPAQUE]. The LSA Length field [OSPFV2] represents the total length (in octets) of the Opaque LSA, including the LSA header and all TLVs (including padding).

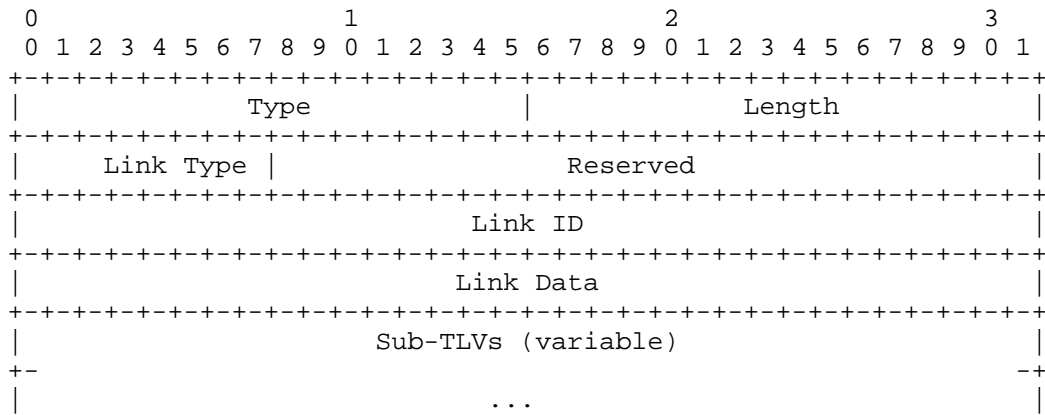
The Opaque ID field is an arbitrary value used to maintain multiple OSPFv2 Extended Prefix Opaque LSAs. For OSPFv2 Extended Link Opaque LSAs, the Opaque ID has no semantic significance other than to differentiate OSPFv2 Extended Link Opaque LSAs originated by the same OSPFv2 router. If multiple OSPFv2 Extended Link Opaque LSAs include the same link, the attributes from the Opaque LSA with the lowest Opaque ID will be used.

The format of the TLVs within the body of the OSPFv2 Extended Link Opaque LSA is the same as described in Section 2.

3.1. OSPFv2 Extended Link TLV

The OSPFv2 Extended Link TLV is used to advertise various attributes of the link. It describes a single link and is constructed of a set of sub-TLVs. There are no ordering requirements for the sub-TLVs. Only one OSPFv2 Extended Link TLV SHALL be advertised in each OSPFv2 Extended Link Opaque LSA, allowing for fine granularity changes in the topology.

The OSPFv2 Extended Link TLV has following format:



OSPFv2 Extended Link TLV

Type

The TLV type. The value is 1 for this TLV type.

Length

Variable, dependent on sub-TLVs.

Link Type

Link Type is defined in Section A.4.2 of [OSPFV2] and in the "OSPFv2 Router LSA Link Type (Value 1)" registry at <<http://www.iana.org/assignments/ospfv2-parameters>>. Specification of link types other than those defined will prevent correlation with existing OSPFv2 Router-LSA links and is beyond the scope this specification.

Link ID

Link ID is defined in Section A.4.2 of [OSPFV2].

Link Data

Link Data is defined in Section A.4.2 of [OSPFV2].

If this TLV is advertised multiple times in the same OSPFv2 Extended Link Opaque LSA, only the first instance of the TLV is used by receiving OSPFv2 routers. This situation SHOULD be logged as an error.

If this TLV is advertised multiple times for the same link in different OSPFv2 Extended Link Opaque LSAs originated by the same OSPFv2 router, the OSPFv2 Extended Link TLV in the OSPFv2 Extended Link Opaque LSA with the smallest Opaque ID is used by receiving OSPFv2 routers. This situation may be logged as a warning.

It is RECOMMENDED that OSPFv2 routers advertising OSPFv2 Extended Link TLVs in different OSPFv2 Extended Link Opaque LSAs re-originate these LSAs in ascending order of Opaque ID to minimize the disruption.

This document creates a registry for OSPFv2 Extended Link sub-TLVs in Section 6.

4. Backward Compatibility

Since Opaque OSPFv2 LSAs are optional and backward compatible [OPAQUE], the extensions described herein are fully backward compatible. However, future OSPFv2 applications utilizing these extensions MUST address backward compatibility of the corresponding functionality.

5. Security Considerations

In general, new LSAs defined in this document are subject to the same security concerns as those described in [OSPFV2] and [OPAQUE].

OSPFv2 applications utilizing these OSPFv2 extensions must define the security considerations relating to those applications in the specifications corresponding to those applications.

Additionally, implementations must assure that malformed TLV and sub-TLV permutations are detected and do not provide a vulnerability for attackers to crash the OSPFv2 router or routing process. Malformed LSAs MUST NOT be stored in the Link State Database (LSDB), acknowledged, or reflooded. Reception of malformed LSAs SHOULD be counted and/or logged for further analysis. In this context, a malformed LSA is one that cannot be parsed due to a TLV or sub-TLV overrunning the end of the subsuming LSA, TLV, or sub-TLV or where there is data remaining to be parsed but the length of the remaining data is less than the size of a TLV header.

6. IANA Considerations

This specification updates the "Opaque Link-State Advertisements (LSA) Option Types" registry with the following values:

- o 7 - OSPFv2 Extended Prefix Opaque LSA
- o 8 - OSPFv2 Extended Link Opaque LSA

This specification also creates five new registries:

- o OSPFv2 Extended Prefix Opaque LSA TLVs
- o OSPFv2 Extended Prefix TLV Sub-TLVs
- o OSPFv2 Extended Prefix TLV Flags
- o OSPFv2 Extended Link Opaque LSA TLVs
- o OSPFv2 Extended Link TLV Sub-TLVs

6.1. OSPFv2 Extended Prefix Opaque LSA TLVs Registry

The "OSPFv2 Extend Prefix Opaque LSA TLVs" registry defines top-level TLVs for OSPFv2 Extended Prefix Opaque LSAs and has been added to the "Open Shortest Path First v2 (OSPFv2) Parameters" registry. New values can be allocated via IETF Review or IESG Approval [RFC5226].

The following initial values have been allocated:

- o 0 - Reserved
- o 1 - OSPFv2 Extended Prefix TLV

Types in the range 32768-33023 are for Experimental Use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

Types in the range 33024-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA considerations covering the range being assigned.

6.2. OSPFv2 Extended Prefix TLV Sub-TLVs Registry

The "OSPFv2 Extended Prefix TLV Sub-TLVs" registry defines sub-TLVs at any level of nesting for OSPFv2 Extended Prefix TLVs and has been added to the "Open Shortest Path First v2 (OSPFv2) Parameters" registry. New values can be allocated via IETF Review or IESG Approval.

The following initial value has been allocated:

- o 0 - Reserved

Types in the range 32768-33023 are for Experimental Use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

Types in the range 33024-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA considerations covering the range being assigned.

6.3. OSPFv2 Extended Prefix TLV Flags Registry

The "OSPFv2 Extended Prefix TLV Flags" registry defines the bits in the 8-bit OSPFv2 Extended Prefix TLV Flags (Section 2.1). This specification defines the A (0x80) and N (0x40) bits. This registry has been added to the "Open Shortest Path First v2 (OSPFv2) Parameters" registry. New values can be allocated via IETF Review or IESG Approval.

6.4. OSPFv2 Extended Link Opaque LSA TLVs Registry

The "OSPFv2 Extended Link Opaque LSA TLVs" registry defines top-level TLVs for OSPFv2 Extended Link Opaque LSAs and has been added to the "Open Shortest Path First v2 (OSPFv2) Parameters" registry. New values can be allocated via IETF Review or IESG Approval.

The following initial values have been allocated:

- o 0 - Reserved
- o 1 - OSPFv2 Extended Link TLV

Types in the range 32768-33023 are for Experimental Use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

Types in the range 33024-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA considerations covering the range being assigned.

6.5. OSPFv2 Extended Link TLV Sub-TLVs Registry

The "OSPFv2 Extended Link TLV Sub-TLVs" registry defines sub-TLVs at any level of nesting for OSPFv2 Extended Link TLVs and has been added to the "Open Shortest Path First v2 (OSPFv2) Parameters" registry. New values can be allocated via IETF Review or IESG Approval.

The following initial value has been allocated:

- o 0 - Reserved

Types in the range 32768-33023 are for Experimental Use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

Types in the range 33024-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA considerations covering the range being assigned.

7. References

7.1. Normative References

- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [OPAQUE] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [OSPFV2] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.
- [TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.

7.2. Informative References

[OSPFv3-EXTEND]

Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", Work in Progress, draft-ietf-ospf-ospfv3-lsa-extend-08, October 2015.

[RFC3101] Murphy, P., "The OSPF Not-So-Stubby Area (NSSA) Option", RFC 3101, DOI 10.17487/RFC3101, January 2003, <<http://www.rfc-editor.org/info/rfc3101>>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

[SEGMENT-ROUTING]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", Work in Progress, draft-ietf-ospf-segment-routing-extensions-05, June 2015.

Acknowledgements

We would like to thank Anton Smirnov for his contribution.

Thanks to Tony Przygienda for his review and comments.

Thanks to Wim Henderickx, Greg Harkins, Peter Psenak, Eric Wu, Shraddha Hegde, and Csaba Mate for their responses to the implementation survey.

Thanks to Tom Petch and Chris Bowers for review and comments.

Thanks to Alia Atlas and Alvaro Retana for their AD review and comments.

Thanks to Carlos Pignataro and Ron Bonica for Operations Directorate review and comments.

Thanks to Suresh Krishnan for the Gen-ART review and comments.

Thanks to Ben Campbell, Kathleen Moriarty, and Barry Leiba for IESG review and comments.

Authors' Addresses

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

Hannes Gredler
Independent
Email: hannes@gredler.at

Rob Shakir
Jive Communications, Inc.
1275 W 1600 N, Suite 100
Orem, UT 84057
United States
Email: rjs@rob.sh

Wim Henderickx
Alcatel-Lucent
Copernicuslaan
Antwerp, 2018 94089
Belgium
Email: wim.henderickx@alcatel-lucent.com

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
United States
Email: jeff.tantsura@ericsson.com

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
United States
Email: acee@cisco.com

