Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 9, 2015

P. Psenak
Cisco Systems
H. Gredler
Juniper Networks, Inc.
R. Shakir
British Telcom
W. Henderickx
Alcatel-Lucent
J. Tantsura
Ericsson
A. Lindem
Cisco Systems

June 7, 2015

OSPFv2 Prefix/Link Attribute Advertisement
draft-ietf-ospf-prefix-link-attr-06.txt

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 OSPF opaque LSAs based on Type-Length-Value (TLV) tuples that can be used to associate additional attributes with prefixes or links. Dependent on the application, these prefixes and links may or not be advertised in the fixed-format LSAs. The OSPF opaque LSAs are optional and fully backward compatible.

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 http://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 December 9, 2015.

Table of Contents

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 OSPF opaque LSAs based on Type-Length-Value (TLV) tuples that can be used to associate additional attributes with prefixes or links. Dependent on the application, these prefixes and links may or not be advertised in the fixed-format LSAs. The OSPF opaque LSAs are optional and fully backward compatible. This is in contrast to the approach taken in OSPFv3 [I-D.ietf-ospf-ospfv3-lsa-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, Network-Summary-LSAs, NSSA-LSAs, and AS-External-LSAs [OSPFV2]

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 [RFC-KEYWORDS].

2. OSPFv2 Extended Prefix Opaque LSA

The OSPFv2 Extended Prefix Opaque LSA will be 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), 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:

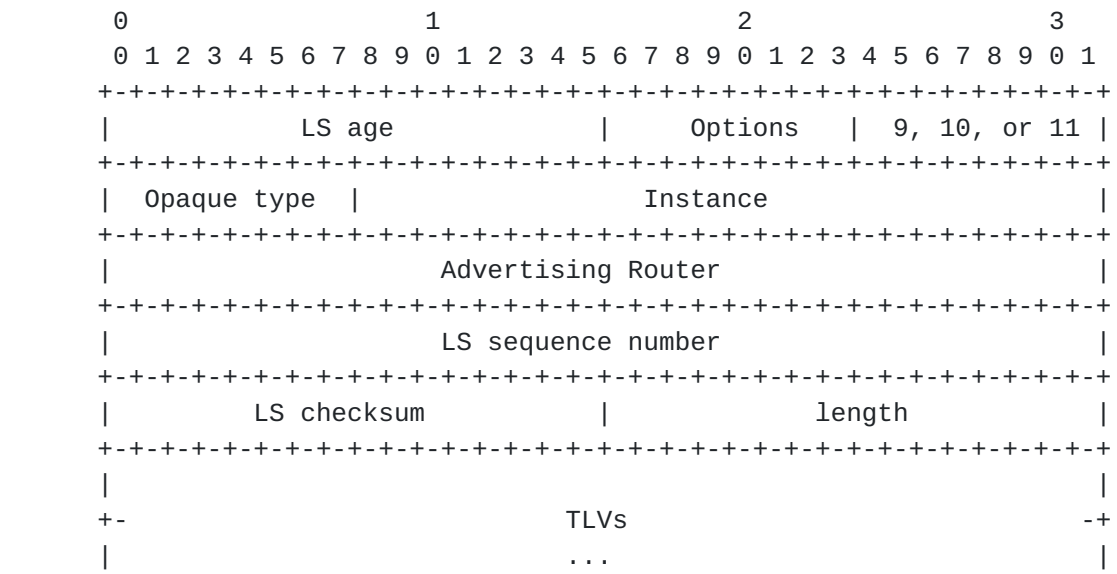

OSPFv2 Extended Prefix Opaque LSA

The opaque type used by OSPFv2 Extended Prefix Opaque LSA is 7.

The Instance field is an arbitrary value used to maintain multiple Extended Prefix Opaque LSAs. A maximum of 16777216 Extended Prefix Opaque LSAs may be sourced by a single OSPF instance.

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 OSPF [TE]. The variable TLV section consists of one or more nested Type/Length/Value (TLV) tuples. Nested TLVs are also referred to as sub-TLVs. The format of each TLV is:

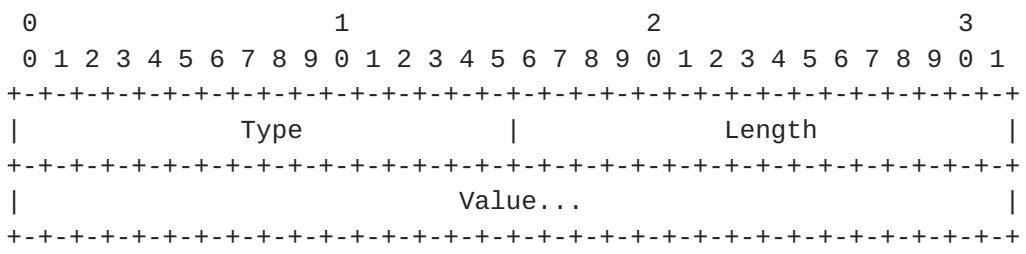

TLV Format

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.

2.1. OSPFv2 Extended Prefix TLV

The OSPF Extended Prefix TLV is used to advertise additional attributes associated with the prefix. Multiple OSPF 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 OSPF Extended Prefix TLV has the following format:

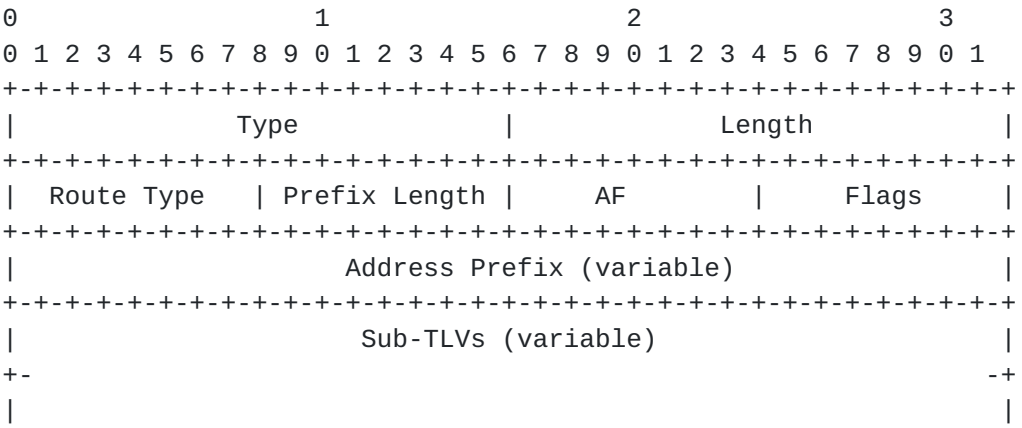

OSPFv2 Extended Prefix TLV

Type
The TLV type. Suggested value is 1.

Length
Variable dependent on sub-TLVs.

Route Type
Route type: type of the OSPF route. If the route type is 0 (Unspecified), the information inside the OSPF 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, which is not aware of the route-type associated with the prefix. Supported types are:

0 - Unspecified

1 - Intra-Area

3 - Inter-Area

5 - AS External

7 - NSSA External

Prefix Length
Length in of the prefix in bits.

AF
Address family for the prefix. Currently, the only supported value is 0 for IPv4 unicast.

Flags: 1 octet field. The following flags are defined:

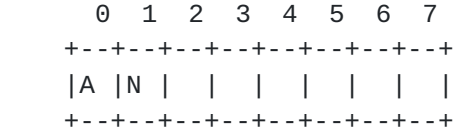

where:

A-Flag: Attach flag. An ABR generating Extended Prefix TLV for inter-area prefix that is locally connected or attached in other connected area SHOULD set this flag.

N-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 OSPFv2 Extended Prefix Opaque LSA is propagated between areas.

Address Prefix
The prefix itself encoded as an even multiple of 32-bit words, padded with zeroed bits as necessary. This encoding consumes ((PrefixLength + 31) / 32) 32-bit words. The default route is represented by a prefix of length 0.

If this TLV is advertised multiple times for the same prefix in the same OSPFv2 Extended Prefix Opaque LSA, only the first instance 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 OSPF router, the OSPF advertising router is re-originating Extended Prefix Opaque LSAs for multiple prefixes and is most likely repacking Extended-Prefix-TLVs in Extended Prefix Opaque LSAs. In this case, the Extended-Prefix-TLV in the Extended Prefix Opaque LSA with the smallest Instance is used by receiving OSPFv2 Routers. This situation MAY be logged as a warning.

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

If this TLV is advertised multiple times for the same prefix in different OSPFv2 Extended Prefix Opaque LSAs originated by the different OSPF 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 OSPF Extended Prefix sub-TLVs in Section 7.

3. OSPFv2 Extended Link Opaque LSA

The OSPFv2 Extended Link Opaque LSA will be 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:

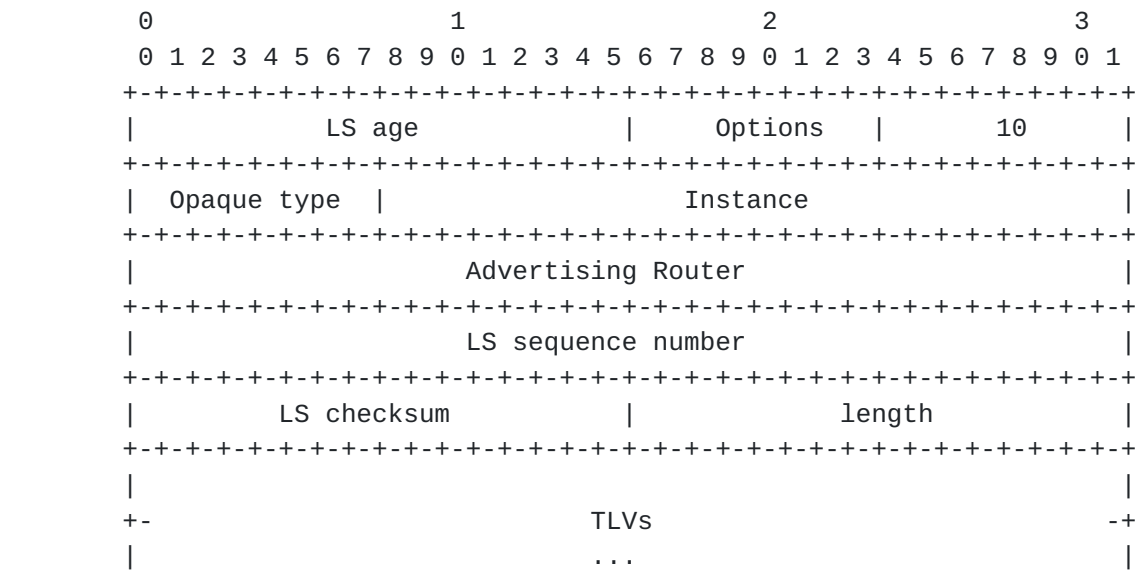

OSPFv2 Extended Link Opaque LSA

The Opaque type used by OSPFv2 Extended Link Opaque LSA is 8.

The Instance field is an arbitrary value used to maintain multiple Extended Link Opaque LSAs. A maximum of 16777216 Extended Link Opaque LSAs may be sourced by a single OSPF instance.

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 Extended Link TLV SHALL be advertised in each Extended Link Opaque LSA, allowing for fine granularity changes in the topology.

The Extended Link TLV has following format:

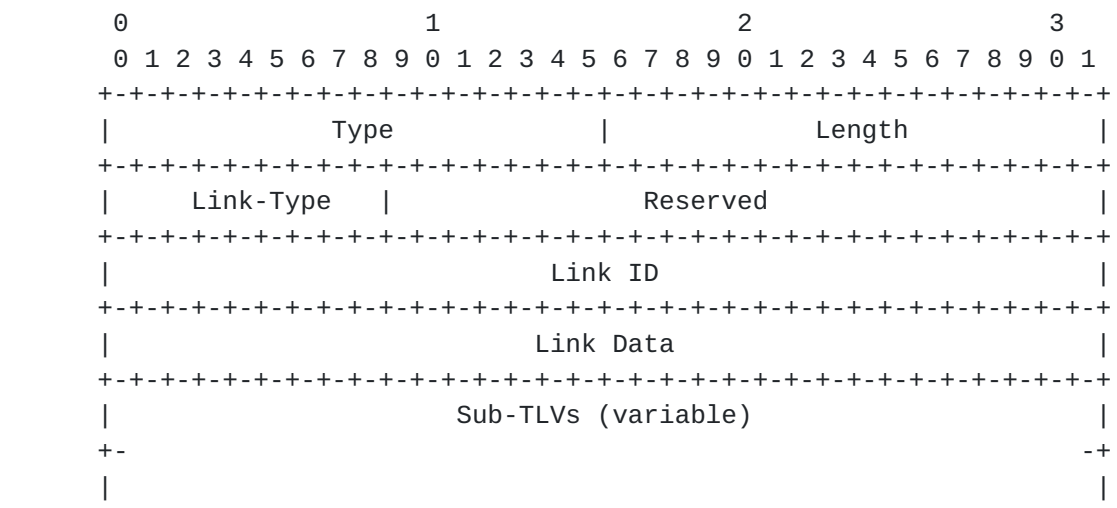

OSPFv2 Extended Link TLV

Type
The TLV type. Suggested value is 1.

Length
Variable dependent on sub-TLVs.

Link-Type
Link-Type is defined in section A.4.2 of [OSPFV2].

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 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 OSPF router, the Extended Link TLV in the Extended Link Opaque LSA with the smallest Instance is used by receiving OSPFv2 Routers. This situation MAY be logged as a warning.

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

This document creates a registry for OSPF Extended Link sub-TLVs in Section 7.

4. Backward Compatibility

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

5. Implementation Status

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in RFC 6982. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to RFC 6982, "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

5.1. Implementation Survey Results

An implementation survey with seven questions related to the implementer's support of OSPFv2 Prefix/Link Attributes was sent to the OSPF WG list and several known implementers. This section contains responses from four implementers who completed the survey. No external means were used to verify the accuracy of the information submitted by the respondents. The respondents are considered experts on the products they reported on. Additionally, responses were omitted from implementers who indicated that they have not implemented the function yet.

Four vendors and one open source entity replied to the survey. These included Alcatel-Lucent, Cisco, Huawei, Juniper, and FreeRouter (http://freerouter.nop.hu). Cisco and Alcatel-Lucent also did

interoperability testing. FreeRouter did interoperability testing with Cisco. The Cisco, Alcatel-Lucent, and FreeRouter implementations are in released software versions. The Huawei and Juniper implementation software releases are pending. For prefix attributes, the recent change incorporating the A-Flag is pending implementation for all four vendors. The FreeRouter implementation includes support for the A-Flag. Implementation of the N-flag is pending for the Huawei and Juniper implementations. Otherwise, all the survey respondents have full implementations. For all four vendors and the FreeRouter implementation, segment routing [SEGMENT-ROUTING] was an application making use of the extensions. Additionally, Cisco has implemented Topology-Independent Loop-Free Alternatives (TI-LFA) [TI-LFA] and Bit Indexed Egress Replication (BIER) advertisement [BIER].

Alcatel-Lucent's support of this specification is included in SR OS, Release 13.0.R4. Cisco's support is included in IOS-XR 5.3.2. The FreeRouter implementation is available in the FreeRouter 15.6.4 distribution. Huawei and Juniper will respectively provide support in future versions Versatile Routing Platform (VRP) and JUniper Network Operating System (JUNOS).

6. Security Considerations

In general, new LSAs defined in this document are subject to the same security concerns as those described in [OSPFV2]. Additionally, implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors that cause hard OSPF failures.

7. IANA Considerations

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

- 7 (IANA Early Allocation [RFC7120]) - OSPFv2 Extended Prefix Opaque LSA

- 8 (IANA Early Allocation [RFC7120]) - OSPFv2 Extended Link Opaque LSA

This specification also creates four new registries:

- OSPF Extended Prefix Opaque LSA TLVs

- OSPF Extended Prefix TLV Sub-TLVs

- OSPF Extended Link Opaque LSA TLVs

- OSPF Extended Link TLV Sub-TLVs

7.1. OSPF Extended Prefix Opaque LSA TLV Registry

The OSPF Extend Prefix Opaque LSA TLV registry will define top-level TLVs for Extended Prefix Opaque LSA and should be placed in the existing OSPF IANA registry. New values can be allocated via IETF Consensus or IESG Approval.

The following initial values are allocated:

- 0 - Reserved

- 1 - OSPF 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 that covers the range being assigned.

7.2. OSPF Extended Prefix TLV Sub-TLV Registry

The OSPF Extended Prefix TLV sub-TLV registry will define sub-TLVs at any level of nesting for Extended Prefix TLV and should be placed in the existing OSPF IANA registry. New values can be allocated via IETF Consensus or IESG Approval.

The following initial values are allocated:

- 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 that covers the range being assigned.

7.3. OSPF Extended Link Opaque LSA TLV Registry

The OSPF Extended Link Opaque LSA TLV registry will define top-level TLVs for Extended Link Opaque LSAs and should be placed in the existing OSPF IANA registry. New values can be allocated via IETF Consensus or IESG Approval.

Following initial values are 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 am IETF specification that specifies IANA Considerations that covers the range being assigned.

7.4. OSPF Extended Link TLV Sub-TLV Registry

The OSPF Extended Link TLV sub-TLV registry will define sub-TLVs at any level of nesting for Extended Link TLV and should be placed in the existing OSPF IANA registry. New values can be allocated via IETF Consensus or IESG Approval.

The following initial values are 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 that covers the range being assigned.

8. Acknowledgments

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.

9. References

9.1. Normative References

[OPAQUE] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

[OSPFV2] Moy, J., "OSPF Version 2", RFC 2328, April 1998.

[RFC-KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[TE] Katz, D., Yeung, D., and K. Kompella, "Traffic Engineering Extensions to OSPF", RFC 3630, September 2003.

9.2. Informative References

[BIER] Psenak, P., Kumar, N., Wijnands, I., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPF Extensions for BIER", draft-ietf-bier-ospf-bier-extensions-00.txt (work in progress), April 2015.

[I-D.ietf-ospf-ospfv3-lsa-extend] Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", draft-ietf-ospf-ospfv3-lsa-extend-06 (work in progress), February 2015.

[RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 7120, January 2014.

[SEGMENT-ROUTING] Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", draft-ietf-ospf-segment-routing-extensions-04.txt (work in progress), February 2015.

[TI-LFA] Francois, P., Filsfils, C., Bashandy, A., Decraene, B., and S. Litkowski, "Topology Independent Fast Reroute using Segment Routing", draft-francois-spring-segment-routing-ti-lfa-01.txt (work in progress), October 2014.

Authors' Addresses

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

Email: ppsenak@cisco.com

Hannes Gredler
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

Email: hannes@juniper.net

Rob Shakir
British Telcom
London
UK

Email: rob.shakir@bt.com

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
USA

Email: jeff.tantsura@ericsson.com

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
USA

Email: acee@cisco.com