

Internet Engineering Task Force (IETF)
Request for Comments: 8362
Updates: 5340, 5838
Category: Standards Track
ISSN: 2070-1721

A. Lindem
A. Roy
Cisco Systems
D. Goethals
Nokia
V. Reddy Vallem

F. Baker
April 2018

OSPFv3 Link State Advertisement (LSA) Extensibility

Abstract

OSPFv3 requires functional extension beyond what can readily be done with the fixed-format Link State Advertisement (LSA) as described in RFC 5340. Without LSA extension, attributes associated with OSPFv3 links and advertised IPv6 prefixes must be advertised in separate LSAs and correlated to the fixed-format LSAs. This document extends the LSA format by encoding the existing OSPFv3 LSA information in Type-Length-Value (TLV) tuples and allowing advertisement of additional information with additional TLVs. Backward-compatibility mechanisms are also described.

This document updates RFC 5340, "OSPF for IPv6", and RFC 5838, "Support of Address Families in OSPFv3", by providing TLV-based encodings for the base OSPFv3 unicast support and OSPFv3 address family support.

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/rfc8362>.

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 (<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	3
1.1.	Requirements Notation	4
1.2.	OSPFv3 LSA Terminology	4
2.	OSPFv3 Extended LSA Types	4
3.	OSPFv3 Extended LSA TLVs	5
3.1.	Prefix Options Extensions	6
3.1.1.	N-bit Prefix Option	7
3.2.	Router-Link TLV	8
3.3.	Attached-Routers TLV	9
3.4.	Inter-Area-Prefix TLV	10
3.5.	Inter-Area-Router TLV	11
3.6.	External-Prefix TLV	12
3.7.	Intra-Area-Prefix TLV	13
3.8.	IPv6 Link-Local Address TLV	14
3.9.	IPv4 Link-Local Address TLV	14
3.10.	IPv6-Forwarding-Address Sub-TLV	15
3.11.	IPv4-Forwarding-Address Sub-TLV	15
3.12.	Route-Tag Sub-TLV	16
4.	OSPFv3 Extended LSAs	16
4.1.	OSPFv3 E-Router-LSA	16
4.2.	OSPFv3 E-Network-LSA	18
4.3.	OSPFv3 E-Inter-Area-Prefix-LSA	19
4.4.	OSPFv3 E-Inter-Area-Router-LSA	20
4.5.	OSPFv3 E-AS-External-LSA	21
4.6.	OSPFv3 E-NSSA-LSA	22
4.7.	OSPFv3 E-Link-LSA	22
4.8.	OSPFv3 E-Intra-Area-Prefix-LSA	24
5.	Malformed OSPFv3 Extended LSA Handling	25
6.	LSA Extension Backward Compatibility	25
6.1.	Full Extended LSA Migration	25
6.2.	Extended LSA Sparse-Mode Backward Compatibility	26

6.3. LSA TLV Processing Backward Compatibility	26
7. Security Considerations	27
8. IANA Considerations	27
8.1. OSPFv3 Extended LSA TLV Registry	27
8.2. OSPFv3 Extended LSA Sub-TLV Registry	28
9. References	29
9.1. Normative References	29
9.2. Informative References	30
Appendix A. Global Configuration Parameters	31
Appendix B. Area Configuration Parameters	31
Acknowledgments	32
Contributors	32
Authors' Addresses	33

1. Introduction

OSPFv3 requires functional extension beyond what can readily be done with the fixed-format Link State Advertisement (LSA) as described in RFC 5340 [OSPFV3]. Without LSA extension, attributes associated with OSPFv3 links and advertised IPv6 prefixes must be advertised in separate LSAs and correlated to the fixed-format LSAs. This document extends the LSA format by encoding the existing OSPFv3 LSA information in Type-Length-Value (TLV) tuples and allowing advertisement of additional information with additional TLVs. Backward-compatibility mechanisms are also described.

This document updates RFC 5340, "OSPF for IPv6", and RFC 5838, "Support of Address Families in OSPFv3", by providing TLV-based encodings for the base OSPFv3 support [OSPFV3] and OSPFv3 address family support [OSPFV3-AF].

A similar extension was previously proposed in support of multi-topology routing. Additional requirements for the OSPFv3 LSA extension include source/destination routing, route tagging, and others.

A final requirement is to limit the changes to OSPFv3 to those necessary for TLV-based LSAs. For the most part, the semantics of existing OSPFv3 LSAs are retained for their TLV-based successor LSAs described herein. Additionally, encoding details, e.g., the representation of IPv6 prefixes as described in Appendix A.4.1 in RFC 5340 [OSPFV3], have been retained. This requirement was included to increase the expedience of IETF adoption and deployment.

The following aspects of the OSPFv3 LSA extension are described:

1. Extended LSA Types
2. Extended LSA TLVs
3. Extended LSA Formats
4. Backward Compatibility

1.1. Requirements Notation

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.

1.2. OSPFv3 LSA Terminology

The TLV-based OSPFv3 LSAs described in this document will be referred to as Extended LSAs. The OSPFv3 fixed-format LSAs [OSPFV3] will be referred to as Legacy LSAs.

2. OSPFv3 Extended LSA Types


In order to provide backward compatibility, new LSA codes must be allocated. There are eight fixed-format LSAs defined in RFC 5340 [OSPFV3]. For ease of implementation and debugging, the LSA function codes are the same as the fixed-format LSAs only with 32, i.e., 0x20, added. The alternative to this mapping was to allocate a bit in the LS Type indicating the new LSA format. However, this would have used one half the LSA function code space for the migration of the eight original fixed-format LSAs. For backward compatibility, the U-bit MUST be set in the LS Type so that the LSAs will be flooded by OSPFv3 routers that do not understand them.

LSA function code	LS Type	Description
33	0xA021	E-Router-LSA
34	0xA022	E-Network-LSA
35	0xA023	E-Inter-Area-Prefix-LSA
36	0xA024	E-Inter-Area-Router-LSA
37	0xC025	E-AS-External-LSA
38	N/A	Unused (Not to be allocated)
39	0xA027	E-Type-7-LSA
40	0x8028	E-Link-LSA
41	0xA029	E-Intra-Area-Prefix-LSA

OSPFv3 Extended LSA Types

3. OSPFv3 Extended LSA TLVs

The format of the TLVs within the body of the Extended LSAs 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 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.

This document defines the following top-level TLV types:

- o 0 - Reserved
- o 1 - Router-Link TLV

- o 2 - Attached-Routers TLV
- o 3 - Inter-Area-Prefix TLV
- o 4 - Inter-Area-Router TLV
- o 5 - External-Prefix TLV
- o 6 - Intra-Area-Prefix TLV
- o 7 - IPv6 Link-Local Address TLV
- o 8 - IPv4 Link-Local Address TLV

Additionally, this document defines the following sub-TLV types:

- o 0 - Reserved
- o 1 - IPv6-Forwarding-Address sub-TLV
- o 2 - IPv4-Forwarding-Address sub-TLV
- o 3 - Route-Tag sub-TLV

In general, TLVs and sub-TLVs MAY occur in any order, and the specification should define whether the TLV or sub-TLV is required and the behavior when there are multiple occurrences of the TLV or sub-TLV. While this document only describes the usage of TLVs and sub-TLVs, sub-TLVs may be nested to any level as long as the sub-TLVs are fully specified in the specification for the subsuming sub-TLV.

For backward compatibility, an LSA is not considered malformed from a TLV perspective unless either a required TLV is missing or a specified TLV is less than the minimum required length. Refer to Section 6.3 for more information on TLV backward compatibility.

3.1. Prefix Options Extensions

The prefix options are extended from Appendix A.4.1.1 [OSPFV3]. The applicability of the LA-bit is expanded, and it SHOULD be set in Inter-Area-Prefix TLVs and MAY be set in External-Prefix TLVs when the advertised host IPv6 address, i.e., PrefixLength = 128 for the IPv6 Address Family or PrefixLength = 32 for the IPv4 Address Family [OSPFV3-AF], is an interface address. In RFC 5340, the LA-bit is only set in Intra-Area-Prefix-LSAs (Section 4.4.3.9 of [OSPFV3]). This will allow a stable address to be advertised without having to configure a separate loopback address in every OSPFv3 area.

3.1.1.1. N-bit Prefix Option

Additionally, the N-bit prefix option is defined. The figure below shows the position of the N-bit in the prefix options (value 0x20).

```

      0 1 2 3 4 5 6 7
+-----+-----+-----+-----+
|   |   | N|DN| P| x|LA|NU|
+-----+-----+-----+-----+

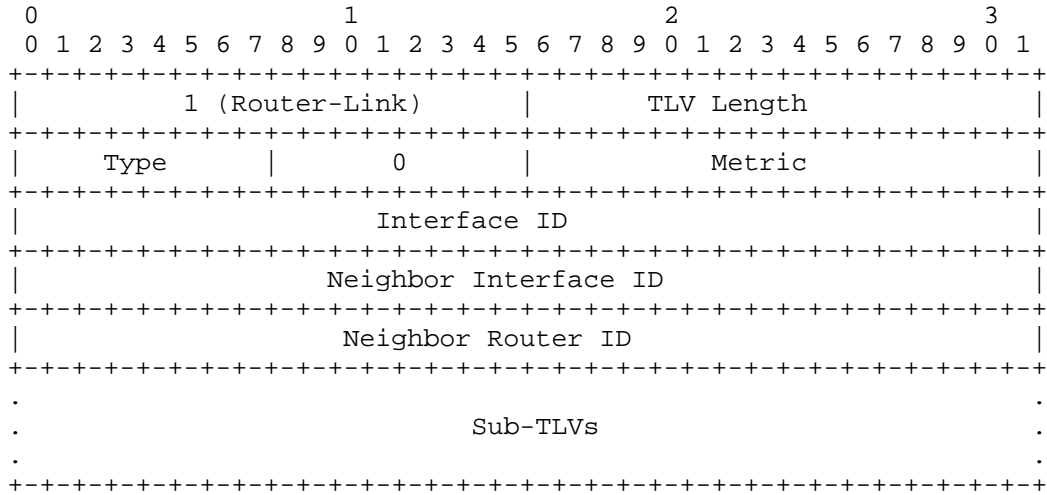
```

The Prefix Options Field

The N-bit is set in PrefixOptions for a host address (PrefixLength=128 for the IPv6 Address Family or PrefixLength=32 for the IPv4 Address Family [OSPFV3-AF]) that identifies the advertising router. While it is similar to the LA-bit, there are two differences. The advertising router MAY choose NOT to set the N-bit even when the above conditions are met. If the N-bit is set and the PrefixLength is NOT 128 for the IPv6 Address Family or 32 for the IPv4 Address Family [OSPFV3-AF], the N-bit MUST be ignored. Additionally, the N-bit is propagated in the PrefixOptions when an OSPFv3 Area Border Router (ABR) originates an Inter-Area-Prefix-LSA for an Intra-Area route that has the N-bit set in the PrefixOptions. Similarly, the N-bit is propagated in the PrefixOptions when an OSPFv3 Not-So-Stubby Area (NSSA) ABR originates an E-AS-External-LSA corresponding to an NSSA route as described in Section 3 of RFC 3101 [NSSA]. The N-bit is added to the Inter-Area-Prefix TLV (Section 3.4), External-Prefix TLV (Section 3.6), and Intra-Area-Prefix-TLV (Section 3.7). The N-bit is used as hint to identify the preferred address to reach the advertising OSPFv3 router. This would be in contrast to an anycast address [IPV6-ADDRESS-ARCH], which could also be a local address with the LA-bit set. It is useful for applications such as identifying the prefixes corresponding to Node Segment Identifiers (SIDs) in Segment Routing [SEGMENT-ROUTING]. There may be future applications requiring selection of a prefix associated with an OSPFv3 router.

3.2. Router-Link TLV

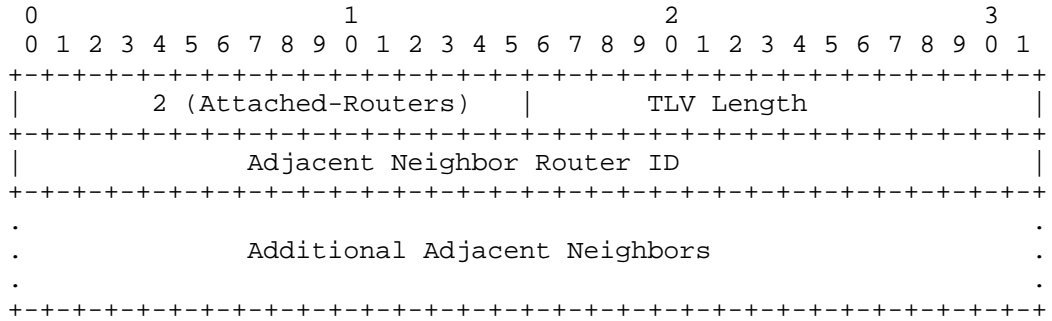
The Router-Link TLV defines a single router link, and the field definitions correspond directly to links in the OSPFv3 Router-LSA; see Appendix A.4.3 of [OSPFV3]. The Router-Link TLV is only applicable to the E-Router-LSA (Section 4.1). Inclusion in other Extended LSAs MUST be ignored.



Router-Link TLV

3.3. Attached-Routers TLV

The Attached-Routers TLV defines all the routers attached to an OSPFv3 multi-access network. The field definitions correspond directly to content of the OSPFv3 Network-LSA; see Appendix A.4.4 of [OSPFV3]. The Attached-Routers TLV is only applicable to the E-Network-LSA (Section 4.2). Inclusion in other Extended LSAs MUST be ignored.

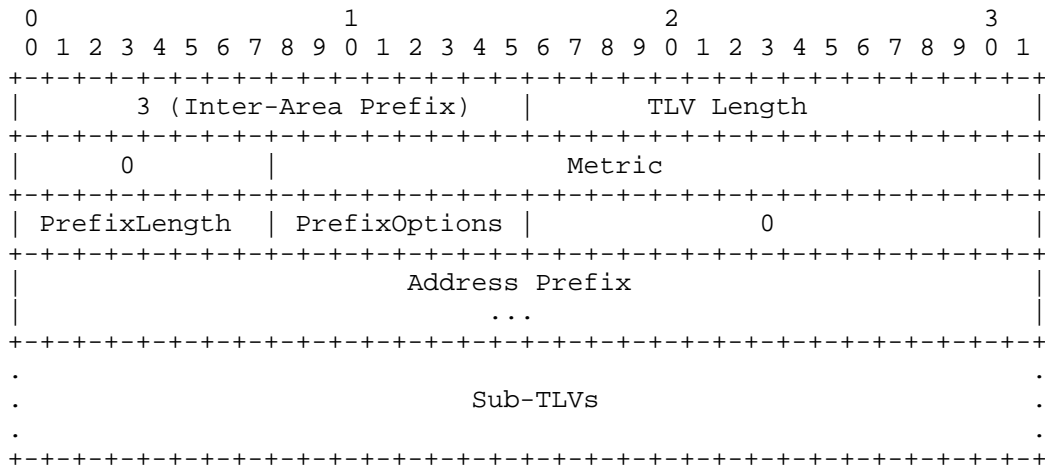


Attached-Routers TLV

There are two reasons for not having a separate TLV or sub-TLV for each adjacent neighbor. The first is to discourage using the E-Network-LSA for more than its current role of solely advertising the routers attached to a multi-access network. The router's metric as well as the attributes of individual attached routers should be advertised in their respective E-Router-LSAs. The second reason is that there is only a single E-Network-LSA per multi-access link with the Link State ID set to the Designated Router's Interface ID, and consequently, compact encoding has been chosen to decrease the likelihood that the size of the E-Network-LSA will require IPv6 fragmentation when advertised in an OSPFv3 Link State Update packet.

3.4. Inter-Area-Prefix TLV

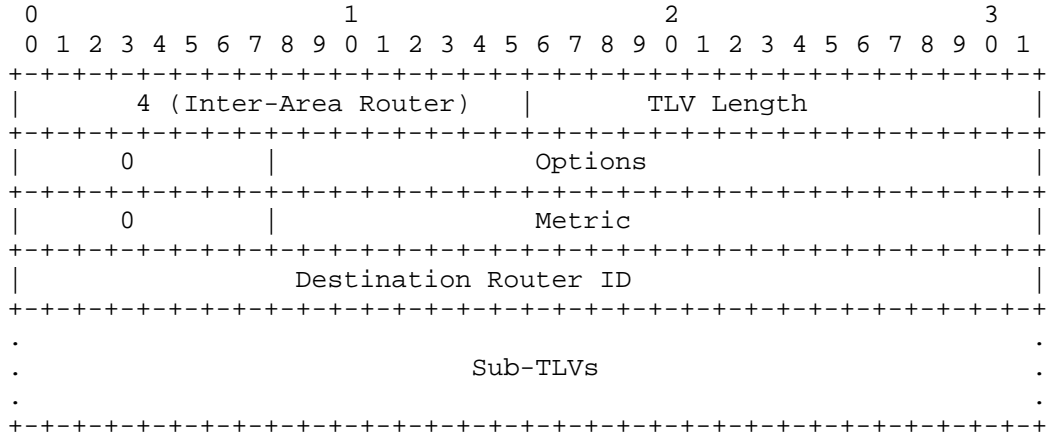
The Inter-Area-Prefix TLV defines a single OSPFV3 inter-area prefix. The field definitions correspond directly to the content of an OSPFv3 IPv6 Prefix, as defined in Appendix A.4.1 of [OSPFV3], and an OSPFv3 Inter-Area-Prefix-LSA, as defined in Appendix A.4.5 of [OSPFV3]. Additionally, the PrefixOptions are extended as described in Section 3.1. The Inter-Area-Prefix TLV is only applicable to the E-Inter-Area-Prefix-LSA (Section 4.3). Inclusion in other Extended LSAs MUST be ignored.



Inter-Area-Prefix TLV

3.5. Inter-Area-Router TLV

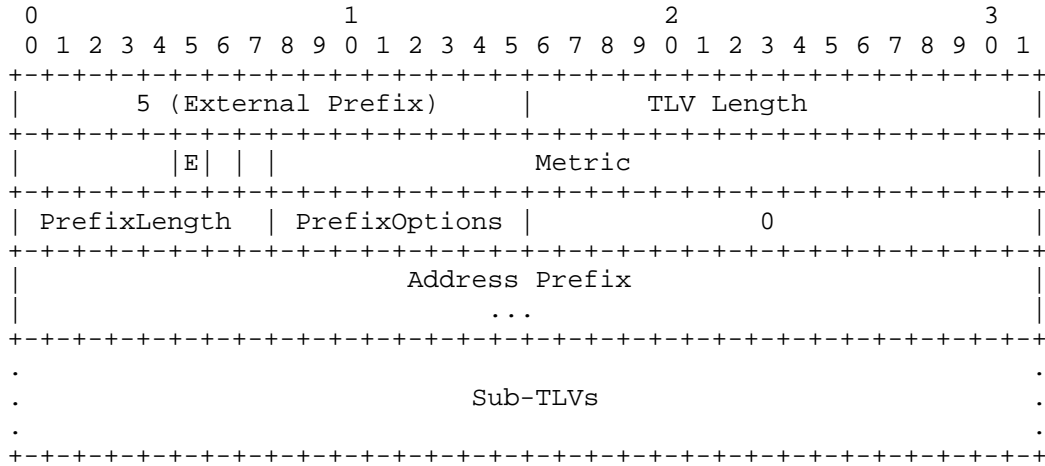
The Inter-Area-Router TLV defines a single OSPFv3 Autonomous System Boundary Router (ASBR) that is reachable in another area. The field definitions correspond directly to the content of an OSPFv3 Inter-Area-Router-LSA, as defined in Appendix A.4.6 of [OSPFV3]. The Inter-Area-Router TLV is only applicable to the E-Inter-Area-Router-LSA (Section 4.4). Inclusion in other Extended LSAs MUST be ignored.



Inter-Area-Router TLV

3.6. External-Prefix TLV

The External-Prefix TLV defines a single OSPFv3 external prefix. With the exception of omitted fields noted below, the field definitions correspond directly to the content of an OSPFv3 IPv6 Prefix, as defined in Appendix A.4.1 of [OSPFV3], and an OSPFv3 AS-External-LSA, as defined in Appendix A.4.7 of [OSPFV3]. The External-Prefix TLV is only applicable to the E-AS-External-LSA (Section 4.5) and the E-NSSA-LSA (Section 4.6). Additionally, the PrefixOptions are extended as described in Section 3.1. Inclusion in other Extended LSAs MUST be ignored.



External-Prefix TLV

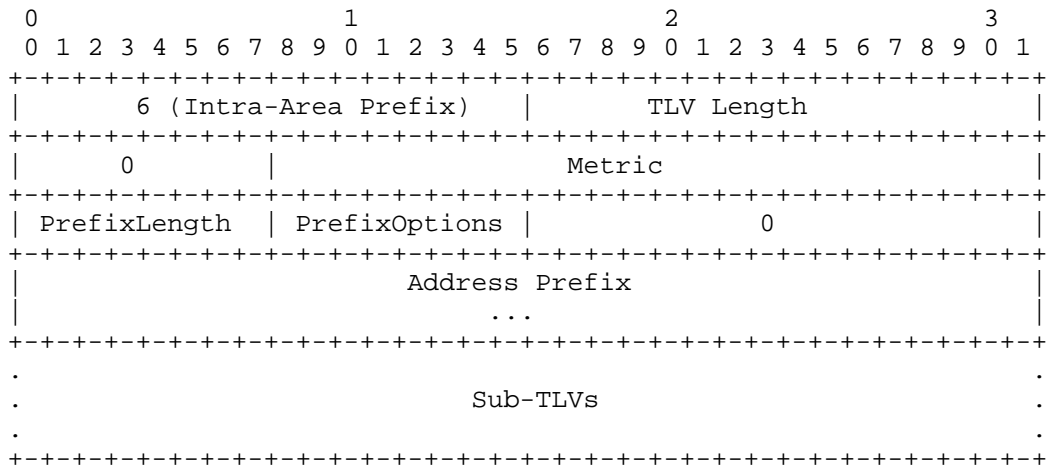
In the External-Prefix TLV, the optional IPv6/IPv4 Forwarding Address and External Route Tag are now sub-TLVs. Given the Referenced LS Type and Referenced Link State ID from the AS-External-LSA have never been used or even specified, they have been omitted from the External-Prefix TLV. If there were ever a requirement for a referenced LSA, it could be satisfied with a sub-TLV.

The following sub-TLVs are defined for optional inclusion in the External-Prefix TLV:

- o 1 - IPv6-Forwarding-Address sub-TLV (Section 3.10)
- o 2 - IPv4-Forwarding-Address sub-TLV (Section 3.11)
- o 3 - Route-Tag sub-TLV (Section 3.12)

3.7. Intra-Area-Prefix TLV

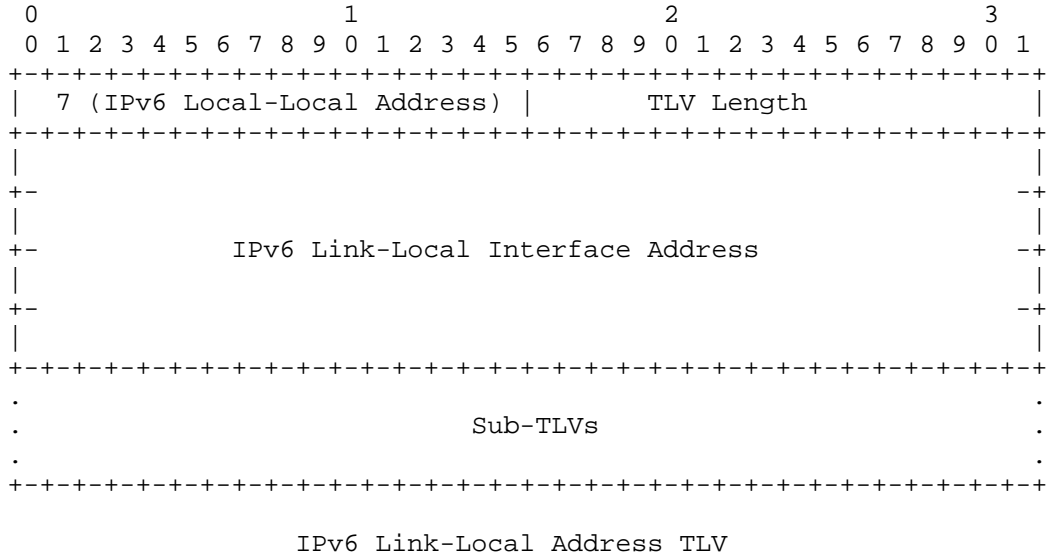
The Intra-Area-Prefix TLV defines a single OSPFv3 intra-area prefix. The field definitions correspond directly to the content of an OSPFv3 IPv6 Prefix, as defined in Appendix A.4.1 of [OSPFV3], and an OSPFv3 Link-LSA, as defined in Appendix A.4.9 of [OSPFV3]. The Intra-Area-Prefix TLV is only applicable to the E-Link-LSA (Section 4.7) and the E-Intra-Area-Prefix-LSA (Section 4.8). Additionally, the PrefixOptions are extended as described in Section 3.1. Inclusion in other Extended LSAs MUST be ignored.



Intra-Area-Prefix TLV

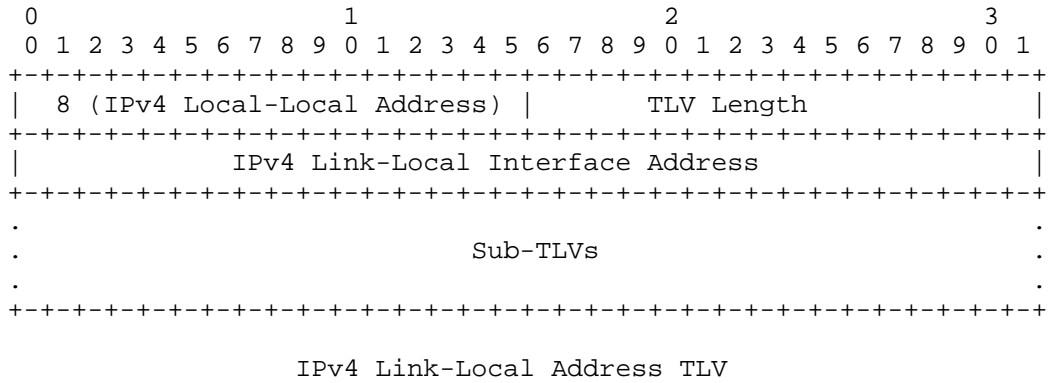
3.8. IPv6 Link-Local Address TLV

The IPv6 Link-Local Address TLV is to be used with IPv6 address families as defined in [OSPFV3-AF]. The IPv6 Link-Local Address TLV is only applicable to the E-Link-LSA (Section 4.7). Inclusion in other Extended LSAs MUST be ignored.



3.9. IPv4 Link-Local Address TLV

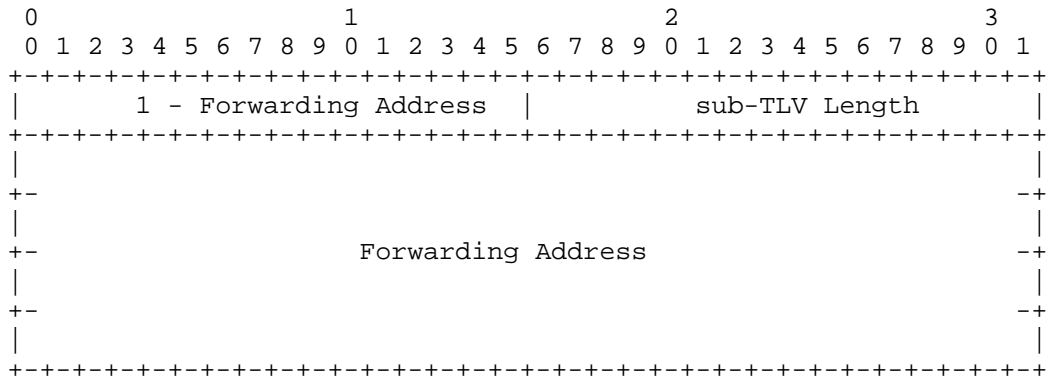
The IPv4 Link-Local Address TLV is to be used with IPv4 address families as defined in [OSPFV3-AF]. The IPv4 Link-Local Address TLV is only applicable to the E-Link-LSA (Section 4.7). Inclusion in other Extended LSAs MUST be ignored.



3.10. IPv6-Forwarding-Address Sub-TLV

The IPv6-Forwarding-Address TLV has identical semantics to the optional forwarding address in Appendix A.4.7 of [OSPFV3]. The IPv6-Forwarding-Address TLV is applicable to the External-Prefix TLV (Section 3.6). Specification as a sub-TLV of other TLVs is not defined herein. The sub-TLV is optional and the first specified instance is used as the forwarding address as defined in [OSPFV3]. Instances subsequent to the first MUST be ignored.

The IPv6-Forwarding-Address TLV is to be used with IPv6 address families as defined in [OSPFV3-AF]. It MUST be ignored for other address families. The IPv6-Forwarding-Address TLV length must meet a minimum length (16 octets), or it will be considered malformed as described in Section 6.3.




IPv6-Forwarding-Address TLV

3.11. IPv4-Forwarding-Address Sub-TLV

The IPv4-Forwarding-Address TLV has identical semantics to the optional forwarding address in Appendix A.4.7 of [OSPFV3]. The IPv4-Forwarding-Address TLV is applicable to the External-Prefix TLV (Section 3.6). Specification as a sub-TLV of other TLVs is not defined herein. The sub-TLV is optional, and the first specified instance is used as the forwarding address as defined in [OSPFV3]. Instances subsequent to the first MUST be ignored.

The IPv4-Forwarding-Address TLV is to be used with IPv4 address families as defined in [OSPFV3-AF]. It MUST be ignored for other address families. The IPv4-Forwarding-Address TLV length must meet a minimum length (4 octets), or it will be considered malformed as described in Section 6.3.




IPv4-Forwarding-Address TLV

3.12. Route-Tag Sub-TLV

The optional Route-Tag sub-TLV has identical semantics to the optional External Route Tag in Appendix A.4.7 of [OSPFV3]. The Route-Tag sub-TLV is applicable to the External-Prefix TLV (Section 3.6). Specification as a sub-TLV of other TLVs is not defined herein. The sub-TLV is optional, and the first specified instance is used as the Route Tag as defined in [OSPFV3]. Instances subsequent to the first MUST be ignored.

The Route-Tag TLV length must meet a minimum length (4 octets), or it will be considered malformed as described in Section 6.3.



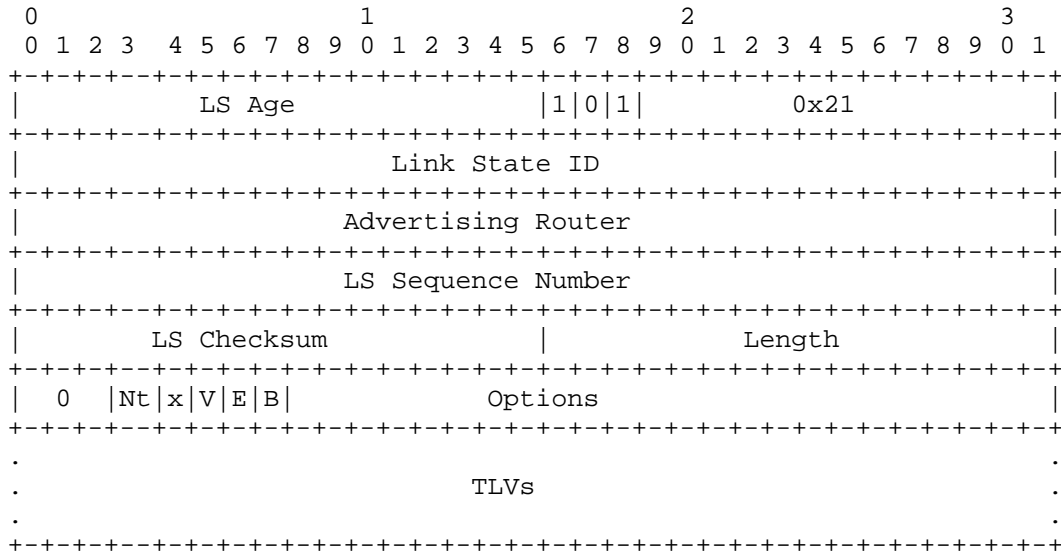
Route-Tag Sub-TLV

4. OSPFv3 Extended LSAs

This section specifies the OSPFv3 Extended LSA formats and encoding. The Extended OSPFv3 LSAs corresponded directly to the original OSPFv3 LSAs specified in [OSPFV3].

4.1. OSPFv3 E-Router-LSA

The E-Router-LSA has an LS Type of 0xA021 and has the same base information content as the Router-LSA defined in Appendix A.4.3 of [OSPFV3]. However, unlike the existing Router-LSA, it is fully extensible and represented as TLVs.

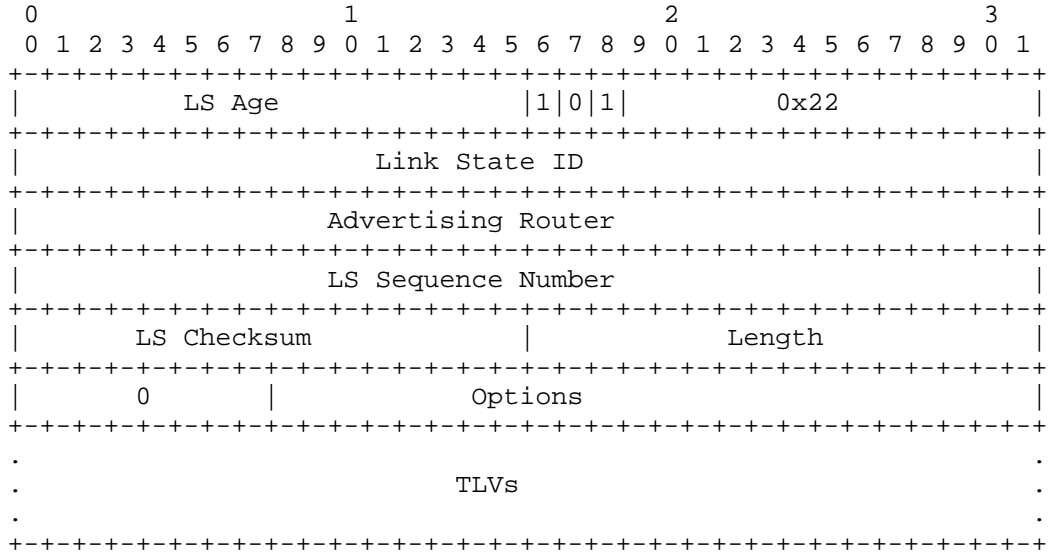


Extended Router-LSA

Other than having a different LS Type, all LSA Header fields are the same as defined for the Router-LSA. Initially, only the top-level Router-Link TLV (Section 3.2) is applicable, and an E-Router-LSA may include multiple Router-Link TLVs. Like the existing Router-LSA, the LSA length is used to determine the end of the LSA including any TLVs. Depending on the implementation, it is perfectly valid for an E-Router-LSA to not contain any Router-Link TLVs. However, this would imply that the OSPFv3 router doesn't have any adjacencies in the corresponding area and is forming an adjacency or adjacencies over an unnumbered link(s). Note that no E-Router-LSA stub link is advertised for an unnumbered link.

4.2. OSPFv3 E-Network-LSA

The E-Network-LSA has an LS Type of 0xA022 and has the same base information content as the Network-LSA defined in Appendix A.4.4 of [OSPFV3]. However, unlike the existing Network-LSA, it is fully extensible and represented as TLVs.




E-Network-LSA

Other than having a different LS Type, all LSA Header fields are the same as defined for the Network-LSA. Like the existing Network-LSA, the LSA length is used to determine the end of the LSA including any TLVs. Initially, only the top-level Attached-Routers TLV (Section 3.3) is applicable. If the Attached-Router TLV is not included in the E-Network-LSA, it is treated as malformed as described in Section 5. Instances of the Attached-Router TLV subsequent to the first MUST be ignored.

4.3. OSPFv3 E-Inter-Area-Prefix-LSA

The E-Inter-Area-Prefix-LSA has an LS Type of 0xA023 and has the same base information content as the Inter-Area-Prefix-LSA defined in Appendix A.4.5 of [OSPFV3]. However, unlike the existing Inter-Area-Prefix-LSA, it is fully extensible and represented as TLVs.




E-Inter-Area-Prefix-LSA

Other than having a different LS Type, all LSA Header fields are the same as defined for the Inter-Area-Prefix-LSA. In order to retain compatibility and semantics with the current OSPFv3 specification, each Inter-Area-Prefix LSA MUST contain a single Inter-Area-Prefix TLV. This will facilitate migration and avoid changes to functions such as incremental Shortest Path First (SPF) computation.

Like the existing Inter-Area-Prefix-LSA, the LSA length is used to determine the end of the LSA including any TLVs. Initially, only the top-level Inter-Area-Prefix TLV (Section 3.4) is applicable. If the Inter-Area-Prefix TLV is not included in the E-Inter-Area-Prefix-LSA, it is treated as malformed as described in Section 5. Instances of the Inter-Area-Prefix TLV subsequent to the first MUST be ignored.

4.4. OSPFv3 E-Inter-Area-Router-LSA

The E-Inter-Area-Router-LSA has an LS Type of 0xA024 and has the same base information content as the Inter-Area-Router-LSA defined in Appendix A.4.6 of [OSPFV3]. However, unlike the Inter-Area-Router-LSA, it is fully extensible and represented as TLVs.



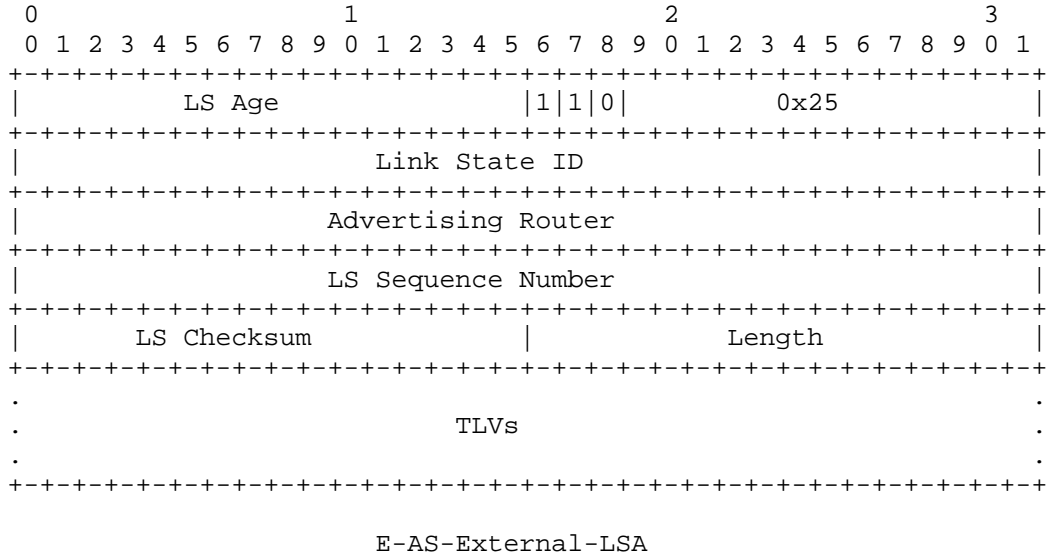
E-Inter-Area-Router-LSA

Other than having a different LS Type, all LSA Header fields are the same as defined for the Inter-Area-Router-LSA. In order to retain compatibility and semantics with the current OSPFv3 specification, each Inter-Area-Router-LSA MUST contain a single Inter-Area-Router TLV. This will facilitate migration and avoid changes to functions such as incremental SPF computation.

Like the existing Inter-Area-Router-LSA, the LSA length is used to determine the end of the LSA including any TLVs. Initially, only the top-level Inter-Area-Router TLV (Section 3.5) is applicable. If the Inter-Area-Router TLV is not included in the E-Inter-Area-Router-LSA, it is treated as malformed as described in Section 5. Instances of the Inter-Area-Router TLV subsequent to the first MUST be ignored.

4.5. OSPFv3 E-AS-External-LSA

The E-AS-External-LSA has an LS Type of 0xC025 and has the same base information content as the AS-External-LSA defined in Appendix A.4.7 of [OSPFV3]. However, unlike the existing AS-External-LSA, it is fully extensible and represented as TLVs.



Other than having a different LS Type, all LSA Header fields are the same as defined for the AS-External-LSA. In order to retain compatibility and semantics with the current OSPFv3 specification, each LSA MUST contain a single External-Prefix TLV. This will facilitate migration and avoid changes to OSPFv3 functions such as incremental SPF computation.

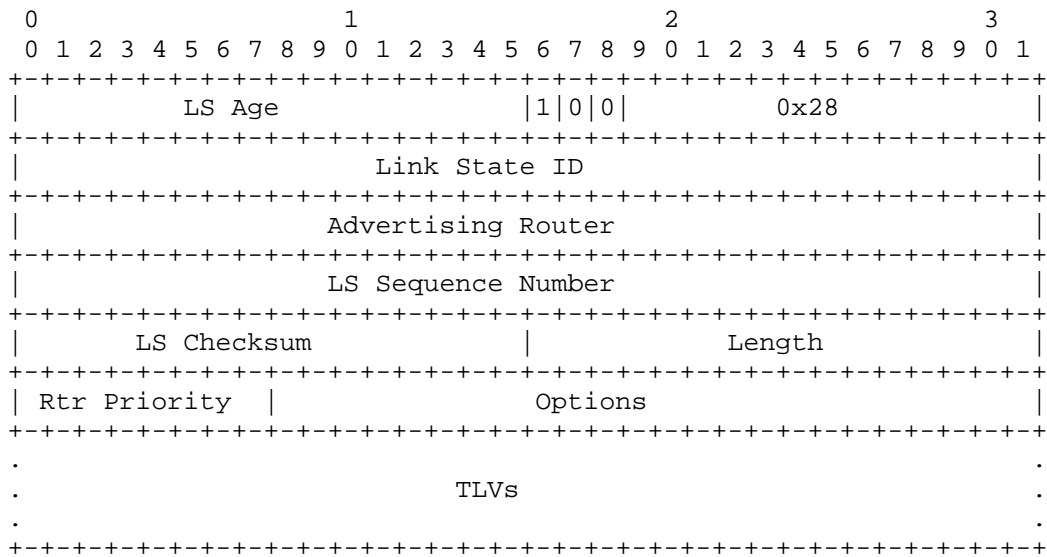
Like the existing AS-External-LSA, the LSA length is used to determine the end of the LSA including any TLVs. Initially, only the top-level External-Prefix TLV (Section 3.6) is applicable. If the External-Prefix TLV is not included in the E-External-AS-LSA, it is treated as malformed as described in Section 5. Instances of the External-Prefix TLV subsequent to the first MUST be ignored.

4.6. OSPFv3 E-NSSA-LSA

The E-NSSA-LSA will have the same format and TLVs as the Extended AS-External-LSA (Section 4.5). This is the same relationship that exists between the NSSA-LSA, as defined in Appendix A.4.8 of [OSPFV3], and the AS-External-LSA. The NSSA-LSA will have type 0xA027, which implies area flooding scope. Future requirements may dictate that supported TLVs differ between the E-AS-External-LSA and the E-NSSA-LSA. However, future requirements are beyond the scope of this document.

4.7. OSPFv3 E-Link-LSA

The E-Link-LSA has an LS Type of 0x8028 and will have the same base information content as the Link-LSA defined in Appendix A.4.9 of [OSPFV3]. However, unlike the existing Link-LSA, it is fully extensible and represented as TLVs.



E-Link-LSA

Other than having a different LS Type, all LSA Header fields are the same as defined for the Link-LSA.

Only the Intra-Area-Prefix TLV (Section 3.7), IPv6 Link-Local Address TLV (Section 3.8), and IPv4 Link-Local Address TLV (Section 3.9) are applicable to the E-Link-LSA. Like the Link-LSA, the E-Link-LSA

affords advertisement of multiple intra-area prefixes. Hence, multiple Intra-Area-Prefix TLVs (Section 3.7) may be specified, and the LSA length defines the end of the LSA including any TLVs.

A single instance of the IPv6 Link-Local Address TLV (Section 3.8) SHOULD be included in the E-Link-LSA. Instances following the first MUST be ignored. For IPv4 address families as defined in [OSPFV3-AF], this TLV MUST be ignored.

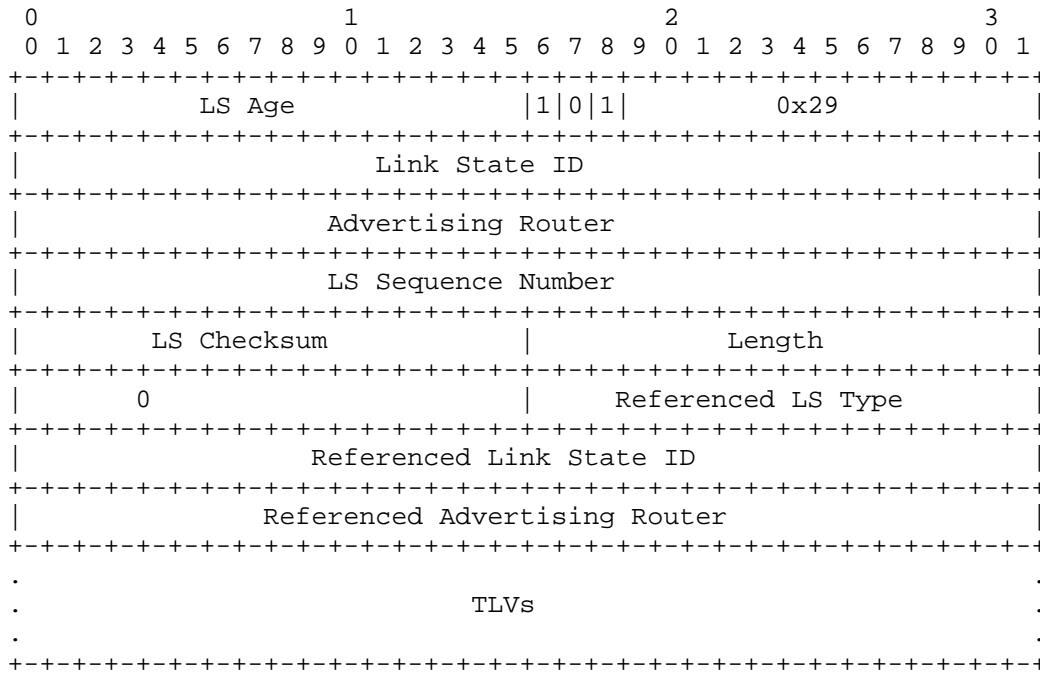
Similarly, only a single instance of the IPv4 Link-Local Address TLV (Section 3.9) SHOULD be included in the E-Link-LSA. Instances following the first MUST be ignored. For OSPFv3 IPv6 address families as defined in [OSPFV3-AF], this TLV SHOULD be ignored.

If the IPv4/IPv6 Link-Local Address TLV corresponding to the OSPFv3 Address Family is not included in the E-Link-LSA, it is treated as malformed as described in Section 5.

Future specifications may support advertisement of routing and topology information for multiple address families. However, this is beyond the scope of this document.

4.8. OSPFv3 E-Intra-Area-Prefix-LSA

The E-Intra-Area-Prefix-LSA has an LS Type of 0xA029 and has the same base information content as the Intra-Area-Prefix-LSA defined in Appendix A.4.10 of [OSPFV3] except for the Referenced LS Type. However, unlike the Intra-Area-Prefix-LSA, it is fully extensible and represented as TLVs. The Referenced LS Type MUST be either an E-Router-LSA (0xA021) or an E-Network-LSA (0xA022).



E-Intra-Area-Prefix-LSA

Other than having a different LS Type, all LSA Header fields are the same as defined for the Intra-Area-Prefix-LSA.

Like the Intra-Area-Prefix-LSA, the E-Intra-Area-Link-LSA affords advertisement of multiple intra-area prefixes. Hence, multiple Intra-Area-Prefix TLVs may be specified, and the LSA length defines the end of the LSA including any TLVs.

5. Malformed OSPFv3 Extended LSA Handling

Extended LSAs that have inconsistent length or other encoding errors, as described herein, MUST NOT be installed in the Link State Database, acknowledged, or flooded. Reception of malformed LSAs SHOULD be counted and/or logged for examination by the administrator of the OSPFv3 routing domain. Note that for the purposes of length validation, a TLV or sub-TLV should not be considered invalid unless the length exceeds the length of the LSA or does not meet the minimum length requirements for the TLV or sub-TLV. This allows for sub-TLVs to be added as described in Section 6.3.

Additionally, an LSA MUST be considered malformed if it does not include all of the required TLVs and sub-TLVs.

6. LSA Extension Backward Compatibility

In the context of this document, backward compatibility is solely related to the capability of an OSPFv3 router to receive, process, and originate the TLV-based LSAs defined herein. Unrecognized TLVs and sub-TLVs are ignored. Backward compatibility for future OSPFv3 extensions utilizing the TLV-based LSAs is out of scope and must be covered in the documents describing those extensions. Both full and, if applicable, partial deployment SHOULD be specified for future TLV-based OSPFv3 LSA extensions.

6.1. Full Extended LSA Migration

If `ExtendedLSASupport` is enabled (Appendix A), OSPFv3 Extended LSAs will be originated and used for the SPF computation. Individual OSPF Areas can be migrated separately with the Legacy AS-External-LSAs being originated and used for the SPF computation. This is accomplished by enabling `AreaExtendedLSASupport` (Appendix B).

An OSPFv3 routing domain or area may be non-disruptively migrated using separate OSPFv3 instances for the Extended LSAs. Initially, the OSPFv3 instances with `ExtendedLSASupport` will have a lower preference, i.e., higher administrative distance, than the OSPFv3 instances originating and using the Legacy LSAs. Once the routing domain or area is fully migrated and the OSPFv3 Routing Information Bases (RIBs) have been verified, the OSPFv3 instances using the Extended LSAs can be given preference. When this has been completed and the routing within the OSPF routing domain or area has been verified, the original OSPFv3 instance using Legacy LSAs can be removed.

6.2. Extended LSA Sparse-Mode Backward Compatibility

In this mode, OSPFv3 will use the Legacy LSAs for the SPF computation and will only originate Extended LSAs when LSA origination is required in support of additional functionality. Furthermore, those Extended LSAs will only include the top-level TLVs (e.g., Router-Link TLVs or Inter-Area TLVs), which are required for that new functionality. However, if a top-level TLV is advertised, it MUST include required sub-TLVs, or it will be considered malformed as described in Section 5. Hence, this mode of compatibility is known as "sparse-mode". The advantage of sparse-mode is that functionality utilizing the OSPFv3 Extended LSAs can be added to an existing OSPFv3 routing domain without the requirement for migration. In essence, this compatibility mode is very much like the approach taken for OSPFv2 [OSPF-PREFIX-LINK]. As with all the compatibility modes, backward compatibility for the functions utilizing the Extended LSAs must be described in the IETF documents describing those functions.

6.3. LSA TLV Processing Backward Compatibility

This section defines the general rules for processing LSA TLVs. To ensure compatibility of future TLV-based LSA extensions, all implementations MUST adhere to these rules:

1. Unrecognized TLVs and sub-TLVs are ignored when parsing or processing Extended LSAs.
2. Whether or not partial deployment of a given TLV is supported MUST be specified.
3. If partial deployment is not supported, mechanisms to ensure the corresponding feature is not deployed MUST be specified in the document defining the new TLV or sub-TLV.
4. If partial deployment is supported, backward compatibility and partial deployment MUST be specified in the document defining the new TLV or sub-TLV.
5. If a TLV or sub-TLV is recognized but the length is less than the minimum, then the LSA should be considered malformed, and it SHOULD NOT be acknowledged. Additionally, the occurrence SHOULD be logged with enough information to identify the LSA by type, Link State ID, originator, and sequence number and identify the TLV or sub-TLV in error. Ideally, the log entry would include the hexadecimal or binary representation of the LSA including the malformed TLV or sub-TLV.

6. Documents specifying future TLVs or Sub-TLVs MUST specify the requirements for usage of those TLVs or sub-TLVs.
7. Future TLVs or sub-TLVs must be optional. However, there may be requirements for sub-TLVs if an optional TLV is specified.

7. Security Considerations

In general, extensible OSPFv3 LSAs are subject to the same security concerns as those described in RFC 5340 [OSPFV3]. Additionally, implementations must assure that malformed TLV and sub-TLV permutations do not result in errors that cause hard OSPFv3 failures.

If there were ever a requirement to digitally sign OSPFv3 LSAs as described for OSPFv2 LSAs in RFC 2154 [OSPF-DIGITAL-SIGNATURE], the mechanisms described herein would greatly simplify the extension.

8. IANA Considerations

This specification defines nine OSPFv3 Extended LSA types as described in Section 2. These have been added to the existing OSPFv3 LSA Function Codes registry.

The specification defines a code point for the N-bit in the OSPFv3 Prefix-Options registry. The value 0x20 has been assigned.

This specification also creates two registries for OSPFv3 Extended LSA TLVs and sub-TLVs. The TLV and sub-TLV code points in these registries are common to all Extended LSAs, and their respective definitions must define where they are applicable.

8.1. OSPFv3 Extended LSA TLV Registry

The "OSPFv3 Extended LSA TLVs" registry defines top-level TLVs for Extended LSAs and has been placed in the existing OSPFv3 IANA registry.

Nine values have been allocated:

- o 0 - Reserved
- o 1 - Router-Link TLV
- o 2 - Attached-Routers TLV
- o 3 - Inter-Area-Prefix TLV
- o 4 - Inter-Area-Router TLV

- o 5 - External-Prefix TLV
- o 6 - Intra-Area-Prefix TLV
- o 7 - IPv6 Link-Local Address TLV
- o 8 - IPv4 Link-Local Address TLV

Types in the range 9-32767 are allocated via IETF Review or IESG Approval [RFC8126].

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

Types in the range 33024-45055 are to be assigned on a First Come First Served (FCFS) basis.

Types in the range 45056-65535 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 cover the range being assigned.

8.2. OSPFv3 Extended LSA Sub-TLV Registry

The "OSPFv3 Extended LSA Sub-TLVs" registry defines sub-TLVs at any level of nesting for Extended LSAs and has been placed in the existing OSPFv3 IANA registry.

Four values have been allocated:

- o 0 - Reserved
- o 1 - IPv6-Forwarding-Address sub-TLV
- o 2 - IPv4-Forwarding-Address sub-TLV
- o 3 - Route-Tag sub-TLV

Types in the range 4-32767 are allocated via IETF Review or IESG Approval.

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

Types in the range 33024-45055 are to be assigned on an FCFS basis.

Types in the range 45056-65535 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 cover the range being assigned.

9. References

9.1. Normative References

- [NSSA] 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>>.
- [OSPFV3] 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>>.
- [OSPFV3-AF] 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>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [TE] 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>>.

9.2. Informative References

[IPV6-ADDRESS-ARCH]

Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

[MT-OSPFV3]

Mirtorabi, S. and A. Roy, "Multi-topology routing in OSPFv3 (MT-OSPFv3)", Work in Progress, draft-ietf-ospf-mt-ospfv3-03, July 2007.

[OSPF-DIGITAL-SIGNATURE]

Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, DOI 10.17487/RFC2154, June 1997, <<https://www.rfc-editor.org/info/rfc2154>>.

[OSPF-PREFIX-LINK]

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>>.

[SEGMENT-ROUTING]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPFv3 Extensions for Segment Routing", Work in Progress, draft-ietf-ospf-ospfv3-segment-routing-extensions-11, January 2018.

Appendix A. Global Configuration Parameters

The global configurable parameter `ExtendedLSASupport` is added to the OSPFv3 protocol. If `ExtendedLSASupport` is enabled, the OSPFv3 router will originate OSPFv3 Extended LSAs and use the LSAs for the SPF computation. If `ExtendedLSASupport` is not enabled, a subset of OSPFv3 Extended LSAs may still be originated and used for other functions as described in Section 6.2.

Appendix B. Area Configuration Parameters

The area configurable parameter `AreaExtendedLSASupport` is added to the OSPFv3 protocol. If `AreaExtendedLSASupport` is enabled, the OSPFv3 router will originate link and area OSPFv3 Extended LSAs and use the LSAs for the SPF computation. Legacy AS-Scoped LSAs will still be originated and used for the AS-External-LSA computation. If `AreaExtendedLSASupport` is not enabled, a subset of OSPFv3 link and area Extended LSAs may still be originated and used for other functions as described in Section 6.2.

For regular areas, i.e., areas where AS-scoped LSAs are flooded, disabling `AreaExtendedLSASupport` for a regular OSPFv3 area (not a Stub or NSSA area) when `ExtendedLSASupport` is enabled is contradictory and SHOULD be prohibited by implementations.

Acknowledgments

OSPFv3 TLV-based LSAs were first proposed in "Multi-topology routing in OSPFv3 (MT-OSPFv3)" [MT-OSPFV3].

Thanks for Peter Psenak for significant contributions to the backward-compatibility mechanisms.

Thanks go to Michael Barnes, Mike Dubrovsky, Anton Smirnov, and Tony Przygienda for review of the draft versions and discussions of backward compatibility.

Thanks to Alan Davey for review and comments including the suggestion to separate the Extended LSA TLV definitions from the Extended LSAs definitions.

Thanks to David Lamparter for review and suggestions on backward compatibility.

Thanks to Karsten Thomann, Chris Bowers, Meng Zhang, and Nagendra Kumar for review and editorial comments.

Thanks to Alia Atlas for substantive Routing Area Director (AD) comments prior to IETF last call.

Thanks to Alvaro Retana and Suresh Krishnan for substantive comments during IESG Review.

Thanks to Mehmet Ersue for the Operations and Management (OPS) Directorate review.

Contributors

Sina Mirtorabi
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
United States of America

Email: sina@cisco.com

Authors' Addresses

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

Email: acee@cisco.com

Abhay Roy
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
United States of America

Email: akr@cisco.com

Dirk Goethals
Nokia
Copernicuslaan 50
Antwerp 2018
Belgium

Email: dirk.goethals@nokia.com

Veerendranatha Reddy Vallem
Bangalore
India

Email: vallem.veerendra@gmail.com

Fred Baker
Santa Barbara, California 93117
United States of America

Email: FredBaker.IETF@gmail.com

