

OSPF Working Group
Internet-Draft
Obsoletes: 6506 (if approved)
Intended status: Standards Track
Expires: December 25, 2013

M. Bhatia
Alcatel-Lucent
V. Manral
Hewlett Packard
A. Lindem
Ericsson
June 23, 2013

Supporting Authentication Trailer for OSPFv3
draft-acee-ospf-rfc6506bis-03.txt

Abstract

Currently, OSPF for IPv6 (OSPFv3) uses IPsec as the only mechanism for authenticating protocol packets. This behavior is different from authentication mechanisms present in other routing protocols (OSPFv2, Intermediate System to Intermediate System (IS-IS), RIP, and Routing Information Protocol Next Generation (RIPng)). In some environments, it has been found that IPsec is difficult to configure and maintain and thus cannot be used. This document defines an alternative mechanism to authenticate OSPFv3 protocol packets so that OSPFv3 does not only depend upon IPsec for authentication. This document obsoletes RFC 6506.

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 25, 2013.

Copyright Notice

Copyright (c) 2013 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	4
1.1. Requirements	5
1.2. Summary of Changes from RFC 6506	5
2. Proposed Solution	6
2.1. AT-Bit in Options Field	6
2.2. Basic Operation	7
2.3. IPv6 Source Address Protection	7
3. OSPFv3 Security Association	9
4. Authentication Procedure	12
4.1. Authentication Trailer	12
4.1.1. Sequence Number Wrap	13
4.2. OSPFv3 Header Checksum and LLS Data Block Checksum	14
4.3. Cryptographic Authentication Procedure	14
4.4. Cross-Protocol Attack Mitigation	15
4.5. Cryptographic Aspects	15
4.6. Message Verification	17
5. Migration and Backward Compatibility	20
6. Security Considerations	21
7. IANA Considerations	22
8. References	23
8.1. Normative References	23
8.2. Informative References	23
Appendix A. Acknowledgments	25
Authors' Addresses	26

1. Introduction

Unlike Open Shortest Path First version 2 (OSPFv2) [RFC2328], OSPF for IPv6 (OSPFv3) [RFC5340] does not include the AuType and Authentication fields in its headers for authenticating protocol packets. Instead, OSPFv3 relies on the IPsec protocols Authentication Header (AH) [RFC4302] and Encapsulating Security Payload (ESP) [RFC4303] to provide integrity, authentication, and/or confidentiality.

[RFC4552] describes how IPv6 AH and ESP extension headers can be used to provide authentication and/or confidentiality to OSPFv3.

However, there are some environments, e.g., Mobile Ad Hoc Networks (MANETs), where IPsec is difficult to configure and maintain, and this mechanism cannot be used.

[RFC4552] discusses, at length, the reasoning behind using manually configured keys, rather than some automated key management protocol such as Internet Key Exchange version 2 (IKEv2) [RFC5996]. The primary problem is the lack of a suitable key management mechanism, as OSPFv3 adjacencies are formed on a one-to-many basis and most key management mechanisms are designed for a one-to-one communication model. This forces the system administrator to use manually configured security associations (SAs) and cryptographic keys to provide the authentication and, if desired, confidentiality services.

Regarding replay protection, [RFC4552] states that:

Since it is not possible using the current standards to provide complete replay protection while using manual keying, the proposed solution will not provide protection against replay attacks.

Since there is no replay protection provided there are a number of vulnerabilities in OSPFv3 that have been discussed in [RFC6039].

Since there is no deterministic way to differentiate between encrypted and unencrypted ESP packets by simply examining the packet, it could be difficult for some implementations to prioritize certain OSPFv3 packet types, e.g., Hello packets, over the other types.

This document defines a new mechanism that works similarly to OSPFv2 [RFC5709] to provide authentication to the OSPFv3 packets and attempts to solve the problems related to replay protection and deterministically disambiguating different OSPFv3 packets as described above.

This document adds support for the Secure Hash Algorithms (SHAs)

defined in the US NIST Secure Hash Standard (SHS), which is specified by NIST FIPS 180-3. [FIPS-180-3] includes SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512. The Hashed Message Authentication Code (HMAC) authentication mode defined in NIST FIPS 198-1 [FIPS-198-1] is used.

It is believed that HMAC as defined in [RFC2104] is mathematically identical to [FIPS-198-1]; it is also believed that algorithms in [RFC6234] are mathematically identical to [FIPS-198-1].

1.1. Requirements

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 2119 [RFC2119].

1.2. Summary of Changes from RFC 6506

This document includes the following changes from RFC 6506 [RFC6506]:

1. Sections 2.2 and 4.2 explicitly state the Link-Local Signaling (LLS) block checksum calculation is omitted when an OSPFv3 authentication is used. The LLS block is included in the authentication digest calculation and computation of a checksum is unnecessary. Clarification of this issue was raised in an errata.
2. Section 4.5 includes a correction to the key preparation to use the protocol specific key (Ks) rather than the key (K) as the initial key (Ko). This problem was also raised in an errata.
3. Section 4.5 also includes a discussion of the choice of key length to be the hash length (L) rather than the block size (B). The discussion of this choice was included to clarify an issue raised in a rejected errata.
4. Section 4.1 and 4.6 indicate that sequence number checking is dependent on OSPFv3 packet type in order to account for packet prioritization as specified in [RFC4222]. This was an omission from RFC 6506.
5. Section 5 includes guidance on precisely the actions required for an OSPFv3 router providing a backward compatible transition mode.

2. Proposed Solution

To perform non-IPsec Cryptographic Authentication, OSPFv3 routers append a special data block, henceforth referred to as the Authentication Trailer, to the end of the OSPFv3 packets. The length of the Authentication Trailer is not included in the length of the OSPFv3 packet but is included in the IPv6 payload length, as shown in Figure 1.

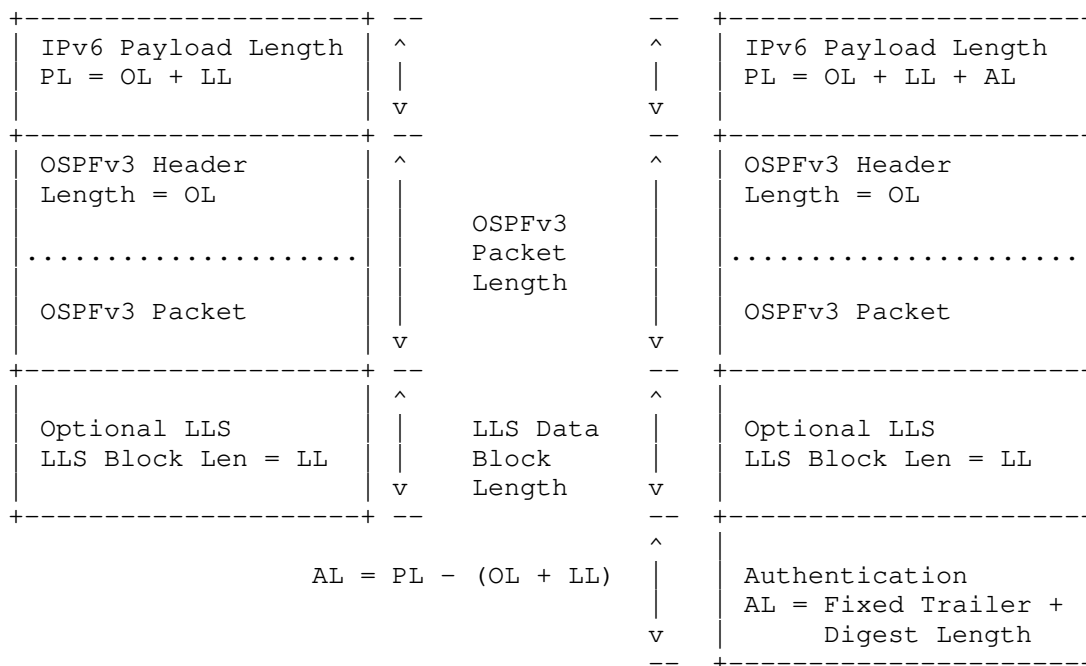


Figure 1: Authentication Trailer in OSPFv3

The presence of the Link-Local Signaling (LLS) [RFC5613] block is determined by the L-bit setting in the OSPFv3 Options field in OSPFv3 Hello and Database Description packets. If present, the LLS data block is included along with the OSPFv3 packet in the Cryptographic Authentication computation.

2.1. AT-Bit in Options Field

A new AT-bit (AT stands for Authentication Trailer) is introduced into the OSPFv3 Options field. OSPFv3 routers MUST set the AT-bit in OSPFv3 Hello and Database Description packets to indicate that all the packets on this link will include an Authentication Trailer. For OSPFv3 Hello and Database Description packets, the AT-bit indicates

the AT is present. For other OSPFv3 packet types, the OSPFv3 AT-bit setting from the OSPFv3 Hello/Database Description setting is preserved in the OSPFv3 neighbor data structure. OSPFv3 packet types that don't include an OSPFv3 Options field will use the setting from the neighbor data structure to determine whether or not the AT is expected.




Figure 2: OSPFv3 Options Field

The AT-bit, as shown in the figure above, MUST be set in all OSPFv3 Hello and Database Description packets that contain an Authentication Trailer.

2.2. Basic Operation

The procedure followed for computing the Authentication Trailer is much the same as described in [RFC5709] and [RFC2328]. One difference is that the LLS data block, if present, is included in the Cryptographic Authentication computation.

The way the authentication data is carried in the Authentication Trailer is very similar to how it is done in case of [RFC2328]. The only difference between the OSPFv2 Authentication Trailer and the OSPFv3 Authentication Trailer is that information in addition to the message digest is included. The additional information in the OSPFv3 Authentication Trailer is included in the message digest computation and is therefore protected by OSPFv3 Cryptographic Authentication as described herein.

Consistent with OSPFv2 Cryptographic Authentication [RFC2328] and Link-Local Signaling Cryptographic Authentication [RFC5613], checksum calculation and verification are omitted for both the OSPFv3 header checksum and the LLS Data Block when the OSPFv3 authentication mechanism described in this specification is used.

2.3. IPv6 Source Address Protection

While OSPFv3 always uses the Router ID to identify OSPFv3 neighbors, the IPv6 source address is learned from OSPFv3 Hello packets and copied into the neighbor data structure [RFC5340]. Hence, OSPFv3 is susceptible to Man-in-the-Middle attacks where the IPv6 source address is modified. To thwart such attacks, the IPv6 source address

will be included in the message digest calculation and protected by OSPFv3 authentication. Refer to Section 4.5 for details. This is different than the procedure specified in [RFC5709] but consistent with [MANUAL-KEY].

3. OSPFv3 Security Association

An OSPFv3 Security Association (SA) contains a set of parameters shared between any two legitimate OSPFv3 speakers.

Parameters associated with an OSPFv3 SA are as follows:

- o Security Association Identifier (SA ID)

This is a 16-bit unsigned integer used to uniquely identify an OSPFv3 SA, as manually configured by the network operator.

The receiver determines the active SA by looking at the SA ID field in the incoming protocol packet.

The sender, based on the active configuration, selects an SA to use and puts the correct Key ID value associated with the SA in the OSPFv3 protocol packet. If multiple valid and active OSPFv3 SAs exist for a given interface, the sender may use any of those SAs to protect the packet.

Using SA IDs makes changing keys while maintaining protocol operation convenient. Each SA ID specifies two independent parts, the authentication algorithm and the Authentication Key, as explained below.

Normally, an implementation would allow the network operator to configure a set of keys in a key chain, with each key in the chain having a fixed lifetime. The actual operation of these mechanisms is outside the scope of this document.

Note that each SA ID can indicate a key with a different authentication algorithm. This allows the introduction of new authentication mechanisms without disrupting existing OSPFv3 adjacencies.

- o Authentication Algorithm

This signifies the authentication algorithm to be used with this OSPFv3 SA. This information is never sent in clear text over the wire. Because this information is not sent on the wire, the implementer chooses an implementation-specific representation for this information.

Currently, the following algorithms are supported:

- * HMAC-SHA-1,

- * HMAC-SHA-256,
- * HMAC-SHA-384, and
- * HMAC-SHA-512.

- o Authentication Key

This value denotes the Cryptographic Authentication Key associated with this OSPFv3 SA. The length of this key is variable and depends upon the authentication algorithm specified by the OSPFv3 SA.

- o KeyStartAccept

The time that this OSPFv3 router will accept packets that have been created with this OSPFv3 SA.

- o KeyStartGenerate

The time that this OSPFv3 router will begin using this OSPFv3 SA for OSPFv3 packet generation.

- o KeyStopGenerate

The time that this OSPFv3 router will stop using this OSPFv3 SA for OSPFv3 packet generation.

- o KeyStopAccept

The time that this OSPFv3 router will stop accepting packets generated with this OSPFv3 SA.

In order to achieve smooth key transition, KeyStartAccept SHOULD be less than KeyStartGenerate, and KeyStopGenerate SHOULD be less than KeyStopAccept. If KeyStartGenerate or KeyStartAccept are left unspecified, the time will default to 0, and the key will be used immediately. If KeyStopGenerate or KeyStopAccept are left unspecified, the time will default to infinity, and the key's lifetime will be infinite. When a new key replaces an old, the KeyStartGenerate time for the new key MUST be less than or equal to the KeyStopGenerate time of the old key.

Key storage SHOULD persist across a system restart, warm or cold, to avoid operational issues. In the event that the last key associated with an interface expires, it is unacceptable to revert to an unauthenticated condition and not advisable to disrupt routing. Therefore, the router SHOULD send a "last Authentication Key

expiration" notification to the network operator and treat the key as having an infinite lifetime until the lifetime is extended, the key is deleted by the network operator, or a new key is configured.

4. Authentication Procedure

4.1. Authentication Trailer

The Authentication Trailer that is appended to the OSPFv3 protocol packet is described below:

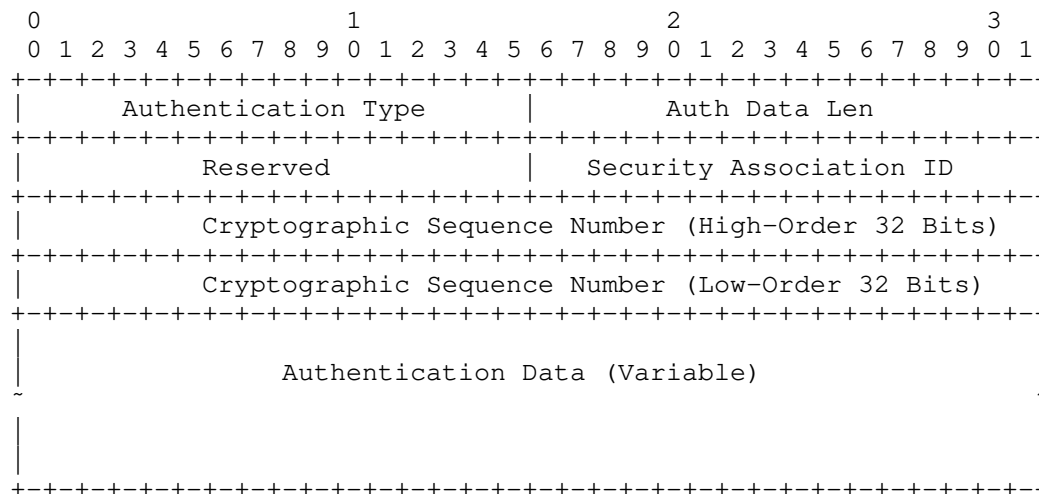


Figure 3: Authentication Trailer Format

The various fields in the Authentication Trailer are:

- o Authentication Type

16-bit field identifying the type of authentication. The following values are defined in this specification:

- 0 - Reserved.
- 1 - HMAC Cryptographic Authentication as described herein.

- o Auth Data Len

The length in octets of the Authentication Trailer (AT) including both the 16-octet fixed header and the variable length message digest.

- o Reserved

This field is reserved. It SHOULD be set to 0 when sending protocol packets and MUST be ignored when receiving protocol packets.

- o Security Association Identifier (SA ID)

16-bit field that maps to the authentication algorithm and the secret key used to create the message digest appended to the OSPFv3 protocol packet.

Though the SA ID implicitly implies the algorithm, the HMAC output size should not be used by implementers as an implicit hint because additional algorithms may be defined in the future that have the same output size.

- o Cryptographic Sequence Number

64-bit strictly increasing sequence number that is used to guard against replay attacks. The 64-bit sequence number MUST be incremented for every OSPFv3 packet sent by the OSPFv3 router. Upon reception, the sequence number MUST be greater than the sequence number in the last accepted OSPFv3 packet of the same packet type from the sending OSPFv3 neighbor. Otherwise, the OSPFv3 packet is considered a replayed packet and dropped. OSPFv3 packets of different types may arrive out of order if they are prioritized as recommended in [RFC4222].

OSPFv3 routers implementing this specification MUST use available mechanisms to preserve the sequence number's strictly increasing property for the deployed life of the OSPFv3 router (including cold restarts). One mechanism for accomplishing this would be to use the high-order 32 bits of the sequence number as a wrap/boot count that is incremented anytime the OSPFv3 router loses its sequence number state. Sequence number wrap is described in Section 4.1.1.

- o Authentication Data

Variable data that is carrying the digest for the protocol packet and optional LLS data block.

4.1.1. Sequence Number Wrap

When incrementing the sequence number for each transmitted OSPFv3 packet, the sequence number should be treated as an unsigned 64-bit value. If the lower-order 32-bit value wraps, the higher-order 32-bit value should be incremented and saved in non-volatile storage. If by some chance the OSPFv3 router is deployed long enough that there is a possibility that the 64-bit sequence number may wrap, all keys, independent of their key distribution mechanism, MUST be reset to avoid the possibility of replay attacks. Once the keys have been changed, the higher-order sequence number can be reset to 0 and saved

to non-volatile storage.

4.2. OSPFv3 Header Checksum and LLS Data Block Checksum

Both the checksum calculation and verification are omitted for the OSPFv3 header checksum and the LLS Data Block checksum [RFC5613] when the OSPFv3 authentication mechanism described in this specification is used. This implies:

- o For OSPFv3 packets to be transmitted, the OSPFv3 header checksum computation is omitted, and the OSPFv3 header checksum SHOULD be set to 0 prior to computation of the OSPFv3 Authentication Trailer message digest.
- o For OSPFv3 packets including an LLS Data Block to be transmitted, the OSPFv3 LLS Data Block checksum computation is omitted, and the OSPFv3 LLS Data Block checksum SHOULD be set to 0 prior to computation of the OSPFv3 Authentication Trailer message digest.
- o For received OSPFv3 packets including an OSPFv3 Authentication Trailer, OSPFv3 header checksum verification MUST be omitted. However, if the OSPFv3 packet does include a non-zero OSPFv3 header checksum, it will not be modified by the receiver and will simply be included in the OSPFv3 Authentication Trailer message digest verification.
- o For received OSPFv3 packets including an LLS Data Block and OSPFv3 Authentication Trailer, LLS Data Block checksum verification MUST be omitted. However, if the OSPFv3 packet does include an LLS Block with a non-zero checksum, it will not be modified by the receiver and will simply be included in the OSPFv3 Authentication Trailer message digest verification.

4.3. Cryptographic Authentication Procedure

As noted earlier, the SA ID maps to the authentication algorithm and the secret key used to generate and verify the message digest. This specification discusses the computation of OSPFv3 Cryptographic Authentication data when any of the NIST SHS family of algorithms is used in the Hashed Message Authentication Code (HMAC) mode.

The currently valid algorithms (including mode) for OSPFv3 Cryptographic Authentication include:

- o HMAC-SHA-1,
- o HMAC-SHA-256,

- o HMAC-SHA-384, and
- o HMAC-SHA-512.

Of the above, implementations of this specification MUST include support for at least HMAC-SHA-256 and SHOULD include support for HMAC-SHA-1 and MAY also include support for HMAC-SHA-384 and HMAC-SHA-512.

Implementations of this specification MUST use HMAC-SHA-256 as the default authentication algorithm.

4.4. Cross-Protocol Attack Mitigation

In order to prevent cross-protocol replay attacks for protocols sharing common keys, the two-octet OSPFv3 Cryptographic Protocol ID is appended to the Authentication Key prior to use. Other protocols using Cryptographic Authentication as specified herein MUST similarly append their respective Cryptographic Protocol IDs to their keys in this step. Refer to the IANA Considerations (Section 7).

4.5. Cryptographic Aspects

In the algorithm description below, the following nomenclature, which is consistent with [FIPS-198-1], is used:

H is the specific hashing algorithm (e.g., SHA-256).

K is the Authentication Key from the OSPFv3 Security Association.

Ks is a Protocol-Specific Authentication Key obtained by appending Authentication Key (K) with the two-octet OSPFv3 Cryptographic Protocol ID.

Ko is the cryptographic key used with the hash algorithm.

B is the block size of H, measured in octets rather than bits. Note that B is the internal block size, not the hash size.

For SHA-1 and SHA-256: B == 64

For SHA-384 and SHA-512: B == 128

L is the length of the hash, measured in octets rather than bits.

XOR is the exclusive-or operation.

Opad is the hexadecimal value 0x5c repeated B times.

Ipad is the hexadecimal value 0x36 repeated B times.

Apad is a value that is the same length as the hash output or message digest. The first 16 octets contain the IPv6 source address followed by the hexadecimal value 0x878FE1F3 repeated (L-16)/4 times. This implies that hash output is always a length of at least 16 octets.

1. Preparation of the Key

The OSPFv3 Cryptographic Protocol ID is appended to the Authentication Key (K) yielding a Protocol-Specific Authentication Key (Ks). In this application, Ko is always L octets long. While [RFC2104] supports a key that is up to B octets long, this application uses L as the Ks length consistent with [RFC4822], [RFC5310], and [RFC5709]. According to [FIPS-198-1], Section 3, keys greater than L octets do not significantly increase the function strength. Ks is computed as follows:

If the Protocol-Specific Authentication Key (Ks) is L octets long, then Ko is equal to Ks. If the Protocol-Specific Authentication Key (Ks) is more than L octets long, then Ko is set to H(Ks). If the Protocol-Specific Authentication Key (Ks) is less than L octets long, then Ko is set to the Protocol-Specific Authentication Key (Ks) with zeros appended to the end of the Protocol-Specific Authentication Key (Ks) such that Ko is L octets long.

2. First-Hash

First, the OSPFv3 packet's Authentication Data field in the Authentication Trailer is filled with the value Apad. This is very similar to the appendage described in [RFC2328], Section D.4.3, Items (6)(a) and (6)(d)).

Then, a First-Hash, also known as the inner hash, is computed as follows:

$$\text{First-Hash} = H(\text{Ko XOR Ipad} \parallel (\text{OSPFv3 Packet}))$$

When XORing Ko and Ipad, Ko will be padded with zeros to the length of Ipad.

Implementation Note: The First-Hash above includes the Authentication Trailer, as well as the OSPFv3 packet, as per [RFC2328], Section D.4.3, and, if present, the LLS data block [RFC5613].

The definition of Apad (above) ensures it is always the same length as the hash output. This is consistent with RFC 2328. Note that the "(OSPFv3 Packet)" referenced in the First-Hash function above includes both the optional LLS data block and the OSPFv3 Authentication Trailer.

The digest length for SHA-1 is 20 octets; for SHA-256, 32 octets; for SHA-384, 48 octets; and for SHA-512, 64 octets.

3. Second-Hash

Then a Second-Hash, also known as the outer hash, is computed as follows:

$$\text{Second-Hash} = H(\text{Ko XOR Opad} \parallel \text{First-Hash})$$

When XORing Ko and Opad, Ko will be padded with zeros to the length of Ipad.

4. Result

The resulting Second-Hash becomes the authentication data that is sent in the Authentication Trailer of the OSPFv3 packet. The length of the authentication data is always identical to the message digest size of the specific hash function H that is being used.

This also means that the use of hash functions with larger output sizes will also increase the size of the OSPFv3 packet as transmitted on the wire.

Implementation Note: [RFC2328], Appendix D specifies that the Authentication Trailer is not counted in the OSPF packet's own Length field but is included in the packet's IP Length field. Similar to this, the Authentication Trailer is not included in the OSPFv3 header length but is included in the IPv6 header payload length.

4.6. Message Verification

A router would determine that OSPFv3 is using an Authentication trailer by examining the AT-bit in the Options field in the OSPFv3 header for Hello and Database Description packets. The specification in the Hello and Database Description options indicates that other OSPFv3 packets will include the Authentication Trailer.

The Authentication Trailer (AT) is accessed using the OSPFv3 packet header length to access the data after the OSPFv3 packet and, if an

LLS data block [RFC5613] is present, using the LLS data block length to access the data after the LLS data block. The L-bit in the OSPFv3 options in Hello and Database Description packets is examined to determine if an LLS data block is present. If an LLS data block is present (as specified by the L-bit), it is included along with the OSPFv3 Hello or Database Description packet in the cryptographic authentication computation.

Due to the placement of the AT following the LLS data block and the fact that the LLS data block is included in the Cryptographic Authentication computation, OSPFv3 routers supporting this specification MUST minimally support examining the L-bit in the OSPFv3 options and using the length in the LLS data block to access the AT. It is RECOMMENDED that OSPFv3 routers supporting this specification fully support OSPFv3 Link-Local Signaling [RFC5613].

If usage of the Authentication Trailer (AT), as specified herein, is configured for an OSPFv3 link, OSPFv3 Hello and Database Description packets with the AT-bit clear in the options will be dropped. All OSPFv3 packet types will be dropped if AT is configured for the link and the IPv6 header length is less than the amount necessary to include an Authentication Trailer.

If the cryptographic sequence number in the AT is less than or equal to the last sequence number in the last OSPFv3 packet of the same type successfully received from the neighbor, the OSPFv3 packet MUST be dropped, and an error event SHOULD be logged. OSPFv3 packets of different types may arrive out of order if they are prioritized as recommended in [RFC4222].

Authentication-algorithm-dependent processing needs to be performed, using the algorithm specified by the appropriate OSPFv3 SA for the received packet.

Before an implementation performs any processing, it needs to save the values of the Authentication Data field from the Authentication Trailer appended to the OSPFv3 packet.

It should then set the Authentication Data field with Apad before the authentication data is computed (as described in Section 4.5). The calculated data is compared with the received authentication data in the Authentication Trailer. If the two do not match, the packet MUST be discarded and an error event SHOULD be logged.

After the OSPFv3 packet has been successfully authenticated, implementations MUST store the 64-bit cryptographic sequence number for each packet type received from the neighbor. The saved cryptographic sequence numbers will be used for replay checking for

subsequent packets received from the neighbor.

5. Migration and Backward Compatibility

All OSPFv3 routers participating on a link SHOULD be migrated to OSPFv3 Authentication at the same time. As with OSPFv2 authentication, a mismatch in the SA ID, Authentication Type, or message digest will result in failure to form an adjacency. For multi-access links, communities of OSPFv3 routers could be migrated using different Interface Instance IDs. However, at least one router would need to form adjacencies between both the OSPFv3 routers including and not including the Authentication Trailer. This would result in sub-optimal routing as well as added complexity and is only recommended in cases where authentication is desired on the link and migrating all the routers on the link at the same time isn't feasible.

In support of uninterrupted deployment, an OSPFv3 router implementing this specification MAY implement a transition mode where it includes the Authentication Trailer in transmitted packets but does not verify this information in received packets. This is provided as a transition aid for networks in the process of migrating to the authentication mechanism described in this specification. More specifically:

1. OSPFv3 routers in transition mode will include the OSPFv3 authentication trailer in transmitted packets and set the AT-Bit in the options field of transmitted Hello and Database Description packets. OSPFv3 routers receiving these packets and not having authentication configured will ignore the authentication trailer and AT-bit.
2. OSPFv3 routers in transition mode will also calculate and set the OSPFv3 header checksum and the LLS block checksum in transmitted packets so that they will not be dropped by OSPFv3 routers without authentication configured.
3. OSPFv3 routers in transition mode will authenticate received packets that have the AT-Bit set in the options field of Hello and Database Description packets or are from a neighbor that previously set the AT-Bit in the options field in Hello and Database Description packets.
4. OSPFv3 routers in transition mode will also accept packets without the options field AT-Bit set in Hello and Database Description packets. These packets will be assumed to be from OSPFv3 routers without authentication configured and they will not be authenticated. Additionally, the OSPFv3 header checksum and LLS block checksum will be validated.

6. Security Considerations

The document proposes extensions to OSPFv3 that would make it more secure than [RFC5340]. It does not provide confidentiality as a routing protocol contains information that does not need to be kept secret. It does, however, provide means to authenticate the sender of the packets that are of interest. It addresses all the security issues that have been identified in [RFC6039].

It should be noted that the authentication method described in this document is not being used to authenticate the specific originator of a packet but is rather being used to confirm that the packet has indeed been issued by a router that has access to the Authentication Key.

Deployments SHOULD use sufficiently long and random values for the Authentication Key so that guessing and other cryptographic attacks on the key are not feasible in their environments. Furthermore, it is RECOMMENDED that Authentication Keys incorporate at least 128 pseudo-random bits to minimize the risk of such attacks. In support of these recommendations, management systems SHOULD support hexadecimal input of Authentication Keys.

The mechanism described herein is not perfect and does not need to be perfect. Instead, this mechanism represents a significant increase in the effort required for an adversary to successfully attack the OSPFv3 protocol while not causing undue implementation, deployment, or operational complexity.

Refer to [RFC4552] for additional considerations on manual keying.

7. IANA Considerations

IANA has allocated the AT-bit (0x000400) in the "OSPFv3 Options (24 bits)" registry as described in Section 2.1.

IANA has created the "OSPFv3 Authentication Trailer Options" registry. This new registry initially includes the "OSPFv3 Authentication Types" registry, which defines valid values for the Authentication Type field in the OSPFv3 Authentication Trailer. The registration procedure is Standards Action.

Value/Range	Designation
0	Reserved
1	HMAC Cryptographic Authentication
2-65535	Unassigned

OSPFv3 Authentication Types

Finally, IANA has created the "Keying and Authentication for Routing Protocols (KARP) Parameters" category. This new category initially includes the "Authentication Cryptographic Protocol ID" registry, which provides unique protocol-specific values for cryptographic applications, such as but not limited to, prevention of cross-protocol replay attacks. Values can be assigned for both native IPv4/IPv6 protocols and UDP/TCP protocols. The registration procedure is Standards Action.

Value/Range	Designation
0	Reserved
1	OSPFv3
2-65535	Unassigned

Cryptographic Protocol ID

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [RFC5709] Bhatia, M., Manral, V., Fanto, M., White, R., Barnes, M., Li, T., and R. Atkinson, "OSPFv2 HMAC-SHA Cryptographic Authentication", RFC 5709, October 2009.
- [RFC6506] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 6506, February 2012.

8.2. Informative References

- [FIPS-180-3] US National Institute of Standards and Technology, "Secure Hash Standard (SHS)", FIPS PUB 180-3, October 2008.
- [FIPS-198-1] US National Institute of Standards and Technology, "The Keyed-Hash Message Authentication Code (HMAC)", FIPS PUB 198, July 2008.
- [MANUAL-KEY] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, "Security Extension for OSPFv2 when using Manual Key Management", Work in Progress, October 2011.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [RFC4222] Choudhury, G., "Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance", BCP 112, RFC 4222, October 2005.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)",

RFC 4303, December 2005.

- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, June 2006.
- [RFC4822] Atkinson, R. and M. Fanto, "RIPv2 Cryptographic Authentication", RFC 4822, February 2007.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, February 2009.
- [RFC5613] Zinin, A., Roy, A., Nguyen, L., Friedman, B., and D. Yeung, "OSPF Link-Local Signaling", RFC 5613, August 2009.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC6039] Manral, V., Bhatia, M., Jaeggli, J., and R. White, "Issues with Existing Cryptographic Protection Methods for Routing Protocols", RFC 6039, October 2010.
- [RFC6234] Eastlake, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, May 2011.

Appendix A. Acknowledgments

First and foremost, thanks to the US National Institute of Standards and Technology for their work on the SHA [FIPS-180-3] and HMAC [FIPS-198-1].

Thanks also need to go to the authors of the HMAC-SHA authentication RFCs including [RFC4822], [RFC5310], and [RFC5709]. The basic HMAC-SHA procedures were originally described by Ran Atkinson and Tony Li in [RFC4822].

Also, thanks to Ran Atkinson for help in the analysis of RFC 6506 errata.

Thanks to Srinivasan K L and Marek Karasek for their identification and submission of RFC 6506 errata.

Thanks to Sam Hartman for discussions on replay mitigation and the use of a 64-bit strictly increasing sequence number. Also, thanks to Sam for comments during IETF last call with respect to the OSPFv3 SA and sharing of key between protocols.

Thanks to Michael Barnes for numerous comments and strong input on the coverage of LLS by the Authentication Trailer (AT).

Thanks to Marek Karasek for providing the specifics with respect to backward compatible transition mode.

Thanks to Rajesh Shetty for numerous comments, including the suggestion to include an Authentication Type field in the Authentication Trailer for extendibility.

Thanks to Uma Chunduri for suggesting that we may want to protect the IPv6 source address even though OSPFv3 uses the Router ID for neighbor identification.

Thanks to Srinivasan KL, Shraddha H, Alan Davey, Russ White, Stan Ratliff, and Glen Kent for their support and review comments.

Thanks to Alia Atlas for comments made under the purview of the Routing Directorate review.

Thanks to Stephen Farrell for comments during the IESG review. Stephen was also involved in the discussion of cross-protocol attacks.

Authors' Addresses

Manav Bhatia
Alcatel-Lucent
Bangalore
India

Email: manav.bhatia@alcatel-lucent.com

Vishwas Manral
Hewlett Packard
USA

Email: vishwas.manral@hp.com

Acee Lindem
Ericsson
102 Carric Bend Court
Cary, NC 27519
USA

Email: acee.lindem@ericsson.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 14, 2014

A. Lindem
Ericsson
S. Mirtorabi
A. Roy
F. Baker
Cisco Systems
September 10, 2013

OSPFv3 LSA Extendibility
draft-acee-ospfv3-lsa-extend-02.txt

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 LSA. This document extends the LSA format by allowing the optional inclusion of Type-Length-Value (TLV) tuples in the LSAs. Backward compatibility mechanisms are also described.

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 March 14, 2014.

Copyright Notice

Copyright (c) 2013 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	4
1.1. Requirements notation	4
1.2. Acknowledgments	4
2. OSPFv3 Extended LSA Types	6
3. OSPFv3 Extended LSA TLV	7
4. OSPFv3 E-Router-LSA	8
5. OSPFv3 E-Network-LSA	10
6. OSPFv3 E-Inter-Area-Prefix-LSA	12
7. OSPFv3 E-Inter-Area-Router-LSA	14
8. OSPFv3 E-AS-External-LSA	16
9. OSPFv3 E-NSSA-LSA	18
10. OSPFv3 E-Link-LSA	19
11. OSPFv3 E-Intra-Area-Prefix-LSA	22
12. LSA Extension Backward Compatibility	23
12.1. Extended LSA Mixed-Mode Backward Compatibility	24
12.2. LSA TLV Processing Backward Compatibility	24
13. Security Considerations	25
14. IANA Considerations	26
15. References	27
15.1. Normative References	27
15.2. Informative References	27
Appendix A. Configurable Constants	28
Authors' Addresses	29

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 LSA. This document extends the LSA format by allowing the optional inclusion of Type-Length-Value (TLV) tuples in the LSAs. Backward compatibility mechanisms are also described.

A similar extension was previously proposed in support of multi-topology routing. Additional requirements for 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 LSA are retained for their TLV-based successor LSAs described herein. Additionally, encoding details, e.g., the representation of IPv6 prefixes as described in section 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 OSPFv3 LSA extension are described:

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

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].

1.2. 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, and Anton Smirnov for review of the draft versions and discussions of backward compatibility.

The RFC text was produced using Marshall Rose's xml2rfc tool.

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 was to allocate a bit in the LSA 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 will be set in 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 TLV

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 Type/Length/Value (TLV) tuples. The format of each TLV is:

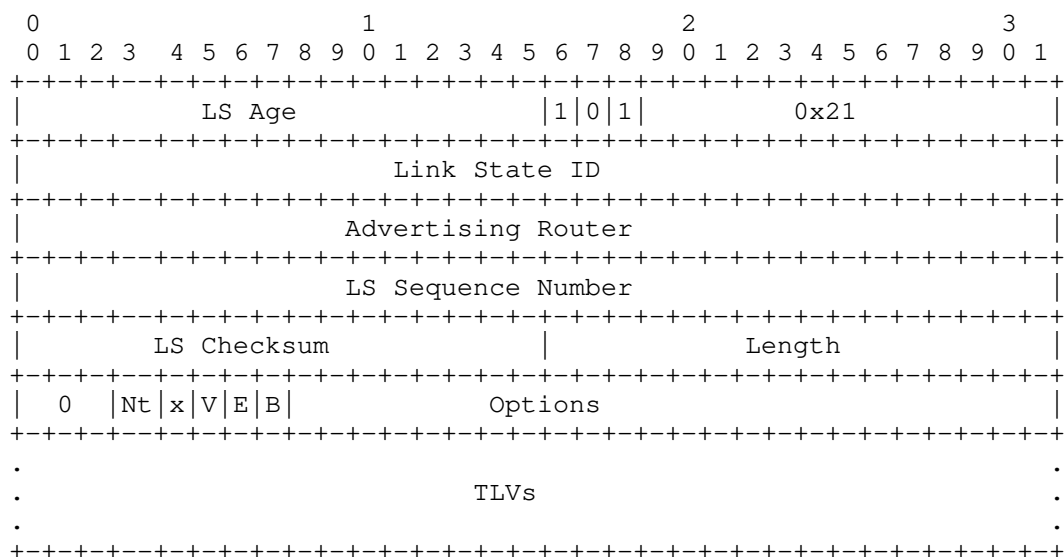


TLV Format

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. Unrecognized types are ignored.

4. OSPFv3 E-Router-LSA

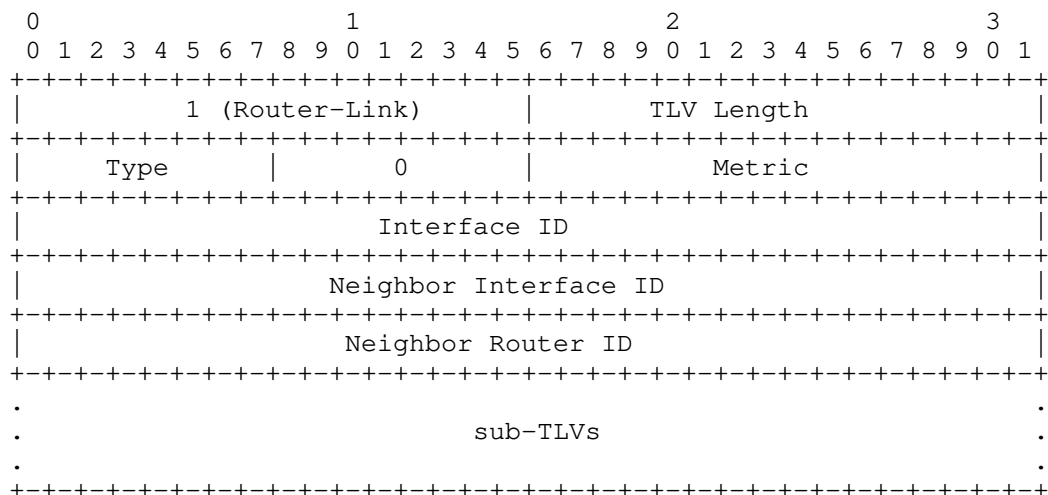
The E-Router-LSA has an LS Type of 0xA021 and has the same base information content as the Router-LSA, section 4.4.3.2 in [OSPFV3]. However, unlike the existing Router-LSA, it is fully extendable and represented as TLVs.



Extended Router-LSA

All LSA Header fields are the same as defined for the Router-LSA. The following top-level TLVs are defined:

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

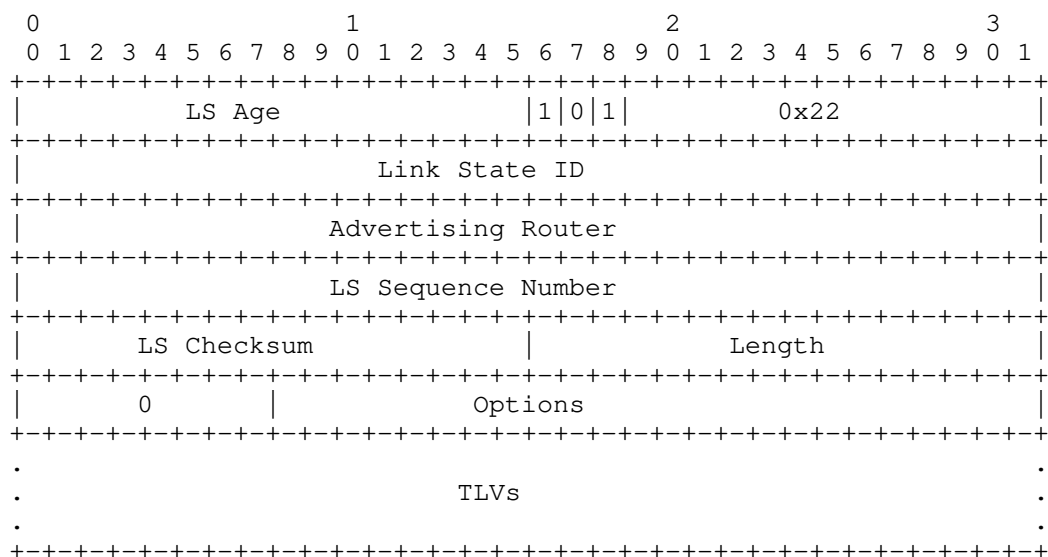


Router-Link TLV

Like the existing Router-LSA, the LSA length is used to determine the end of the LSA including TLVs. The Router-Link TLV is only applicable to the E-Router-LSA. Inclusion in other Extended LSAs MUST be ignored.

5. OSPFv3 E-Network-LSA

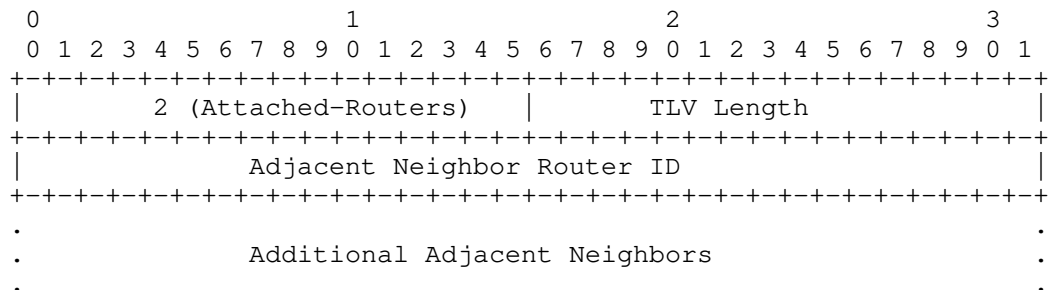
The E-Network-LSA has an LS Type of 0xA022 and has the same base information content as the Network-LSA, section 4.4.3.3 in [OSPFV3]. However, unlike the existing Network-LSA, it is fully extendable and represented as TLVs.



E-Network-LSA

All LSA Header fields are the same as defined for the Network-LSA. The following top-level TLVs are defined:

- o 2 - Attached-Routers TLV



+--+

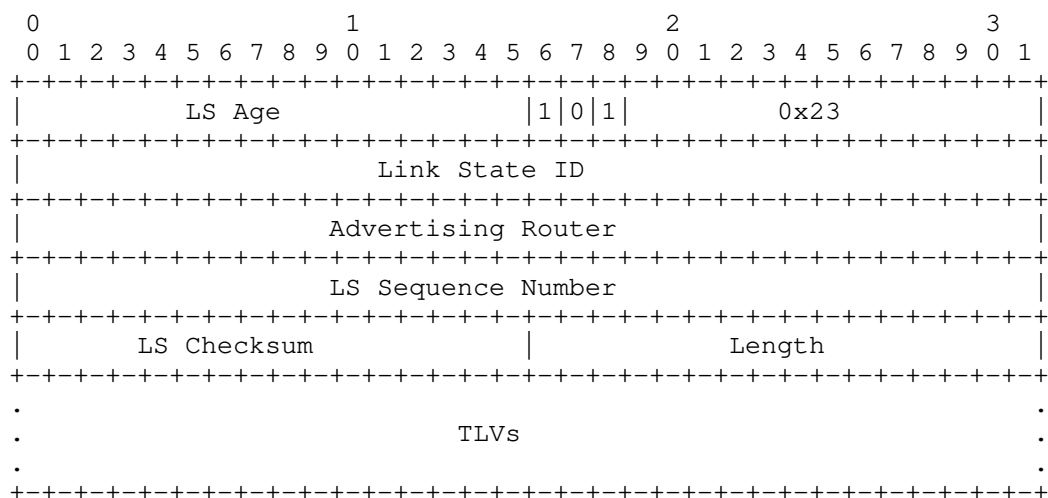
Attached-Routers TLV

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 her 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 of the size of the E-Network-LSA requiring IPv6 fragmentation when advertised in an OSPFv3 Link State Update packet.

Like the existing Network-LSA, the LSA length is used to determine the end of the LSA including TLVs. The Attached-Routers TLV is only applicable to the E-Network-LSA. Inclusion in other Extended LSAs MUST be ignored.

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

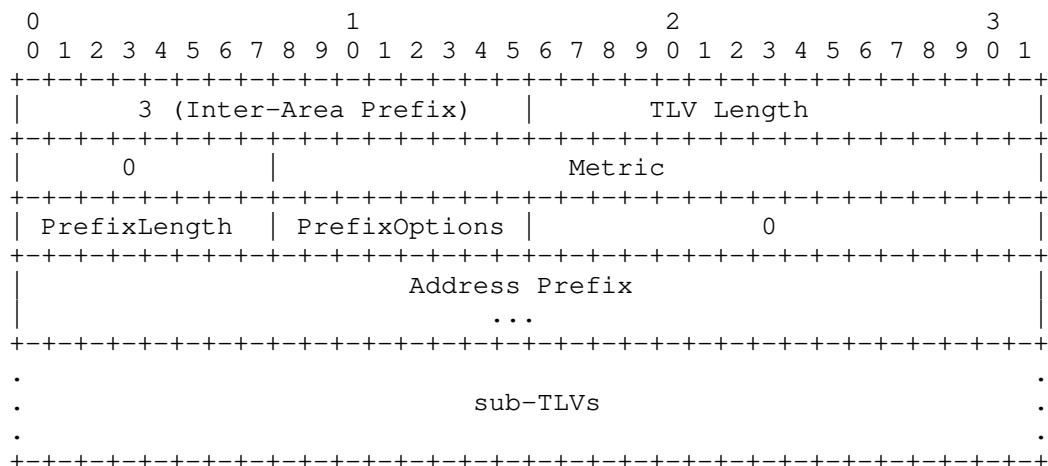
The E-Inter-Area-Prefix-LSA has an LS Type of 0xA023 and has the same base information content as the Inter-Area-Prefix-LSA, section 4.4.3.4 in [OSPFV3]. However, unlike the existing Inter-Area-Prefix-LSA, it is fully extendable and represented as TLVs.



E-Inter-Area-Prefix-LSA

All LSA Header fields are the same as defined for the Network-LSA. The following top-level TLVs are defined:

- o 3 - Inter-Area Prefix TLV



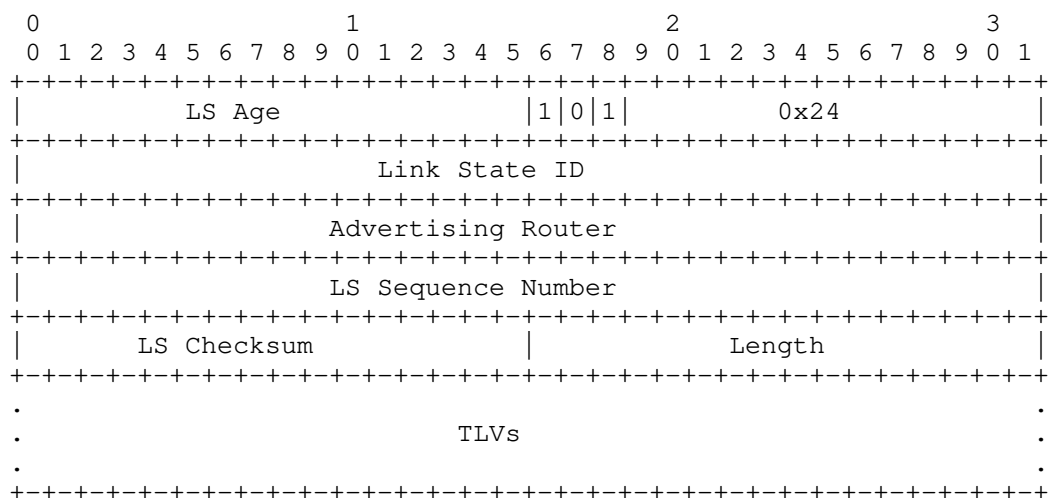
Inter-Area Prefix TLV

In order to retain compatibility and semantics with the current OSPFv3 specification, each LSA MUST contain a single Inter-Area Prefix TLV. This will facilitate migration and avoid changes to functions such as incremental SPF computation.

Like the existing Inter-Area-Prefix-LSA, the LSA length is used to determine the end of the LSA including TLV. The Inter-Area-Prefix TLV is only applicable to the E-Inter-Area-Prefix-LSA. Inclusion in other Extended LSAs MUST be ignored.

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

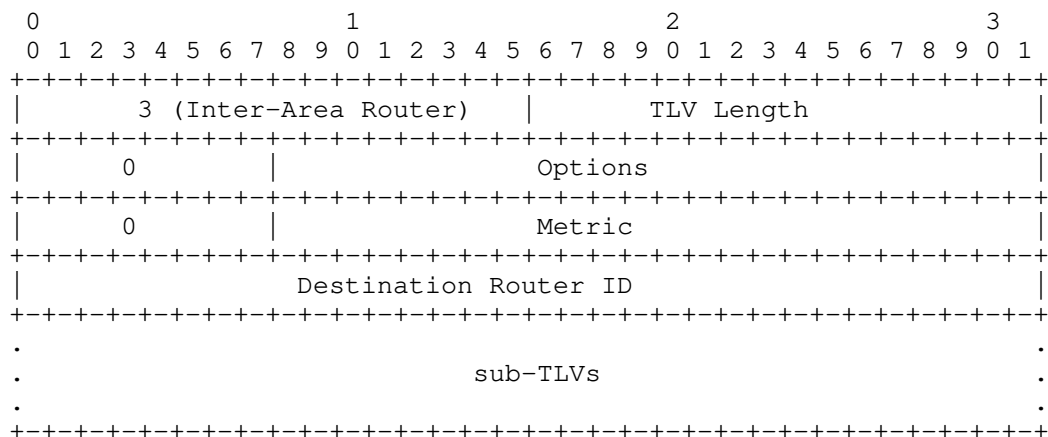
The E-Inter-Area-Router-LSA has an LS Type of 0xA024 and has the same base information content as the Inter-Area-Router-LSA, section 4.4.3.5 in [OSPFV3]. However, unlike the Inter-Area-Router-LSA, it is fully extendable and represented as TLVs.



E-Inter-Area-Router-LSA

All LSA Header fields are the same as defined for the Inter-Area-Router-LSA. The following top-level TLVs are defined:

- o 4 - Inter-Area Router TLV



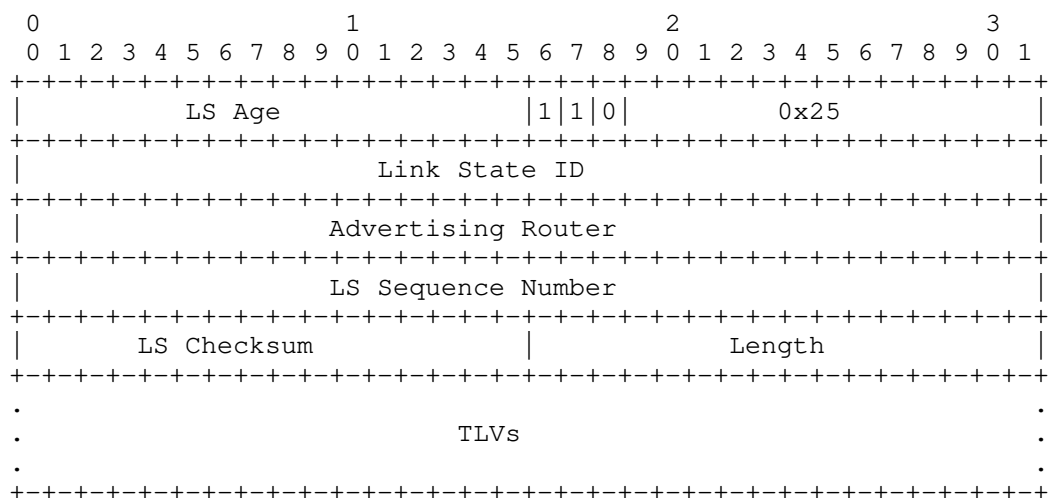
Inter-Area Router TLV

In order to retain compatibility and semantics with the current OSPFv3 specification, each 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 sub-TLVs. The Inter-Area-Router TLV is only applicable to the E-Inter-Area-Router-LSA. Inclusion in other Extended LSAs MUST be ignored.

8. OSPFv3 E-AS-External-LSA

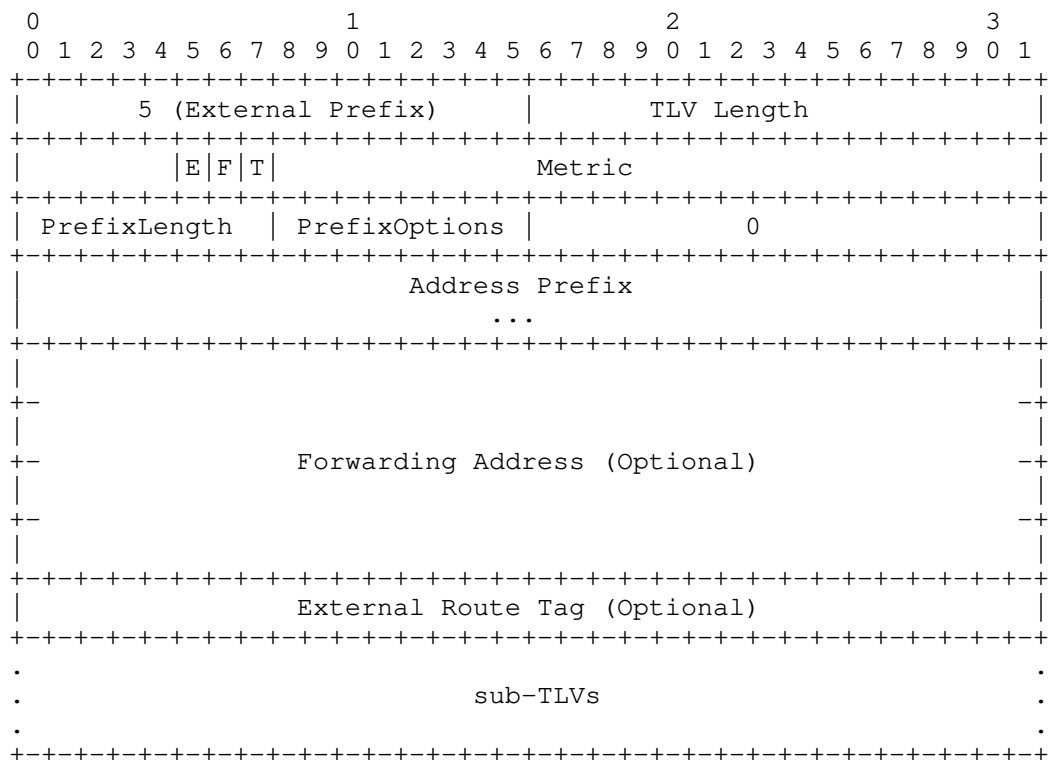
The E-AS-External-LSA has an LS Type of 0xC025 and has the same base information content as the AS-External-LSA, section 4.4.3.6 in [OSPFV3]. However, unlike the existing AS-External-LSA, it is fully extendable and represented as TLVs.



E-AS-External-LSA

All LSA Header fields are the same as defined for the AS-External-LSA. The following top-level TLVs are defined:

- o 5 - External Prefix TLV



External Prefix TLV

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 functions such as incremental SPF computation. 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.

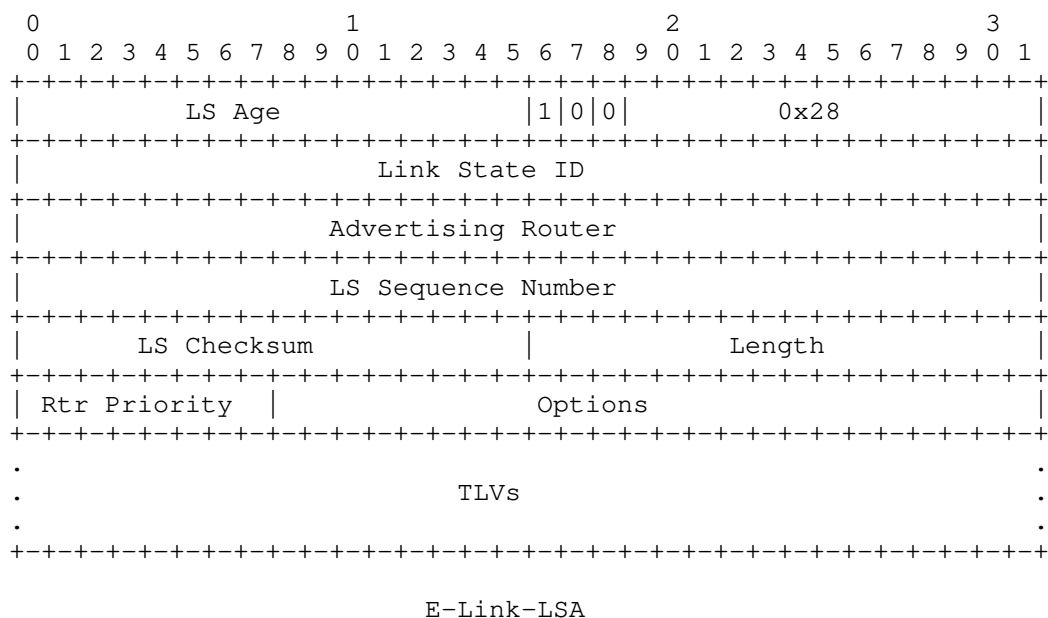
Like the existing AS-External-LSA, the LSA length is used to determine the end of the LSA including sub-TLVs. The External-Prefix TLV is only applicable to the E-AS-External-LSA and the E-NSSA-LSA. Inclusion in other Extended LSAs MUST be ignored.

9. OSPFv3 E-NSSA-LSA

The E-NSSA-LSA will have the same format and TLVs as the Extended AS-External-LSA Section 8. This is the same relationship as exists between the NSSA-LSA, section 4.4.3.7 in [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.

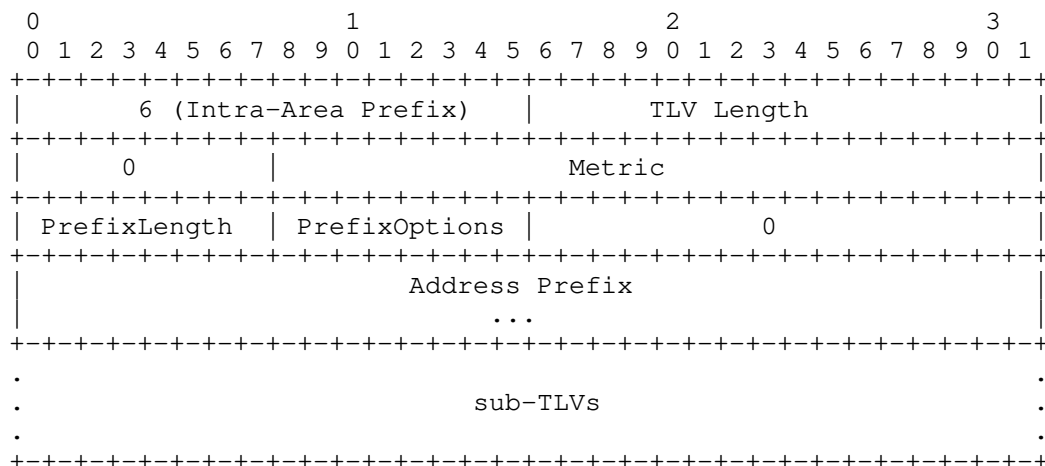
10. OSPFv3 E-Link-LSA

The E-Link-LSA has an LS Type of 0x8028 and will have the same base information content as the Link-LSA, section 4.4.3.8 in [OSPFV3]. However, unlike the existing Link-LFA, it is extendable and represented as TLVs.



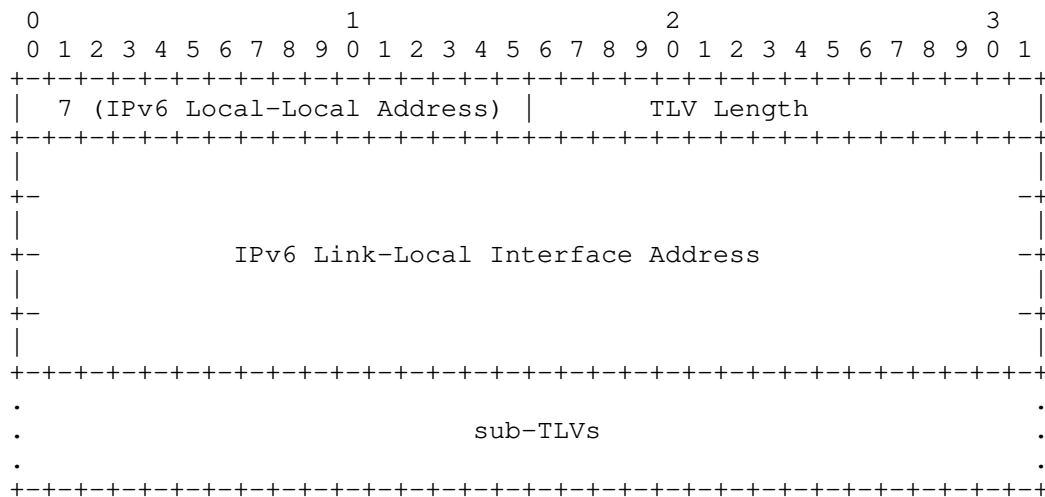
The following top-level TLVs are defined:

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



Intra-Area Prefix TLV

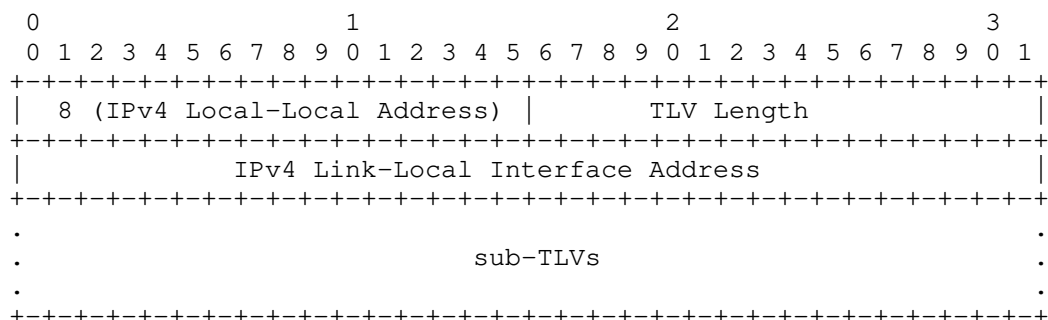
Like the Link-LSA, the E-Link-LSA affords advertisement of multiple intra-area prefixes. Hence, multiple Intra-Area Prefix TLVs may be specified and the LSA length defines the end of the LSA including all TLVs. The Intra-Area-Prefix TLV is only applicable to the E-Link-LSA and the E-Intra-Area-Prefix-LSA. Inclusion in other Extended LSAs MUST be ignored.



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. Inclusion in other Extended LSAs MUST be ignored. Only a single instance of the IPv6 Link-Local Address family SHOULD be included in the E-Link-LSA. Instances preceding the first MUST be ignored. For IPv4 address families as defined in [OSPFV3-AF], this TLV SHOULD be ignored. Future specifications may support advertisement of routing and topology information for multiple address families. However, this is beyond the scope of this document.

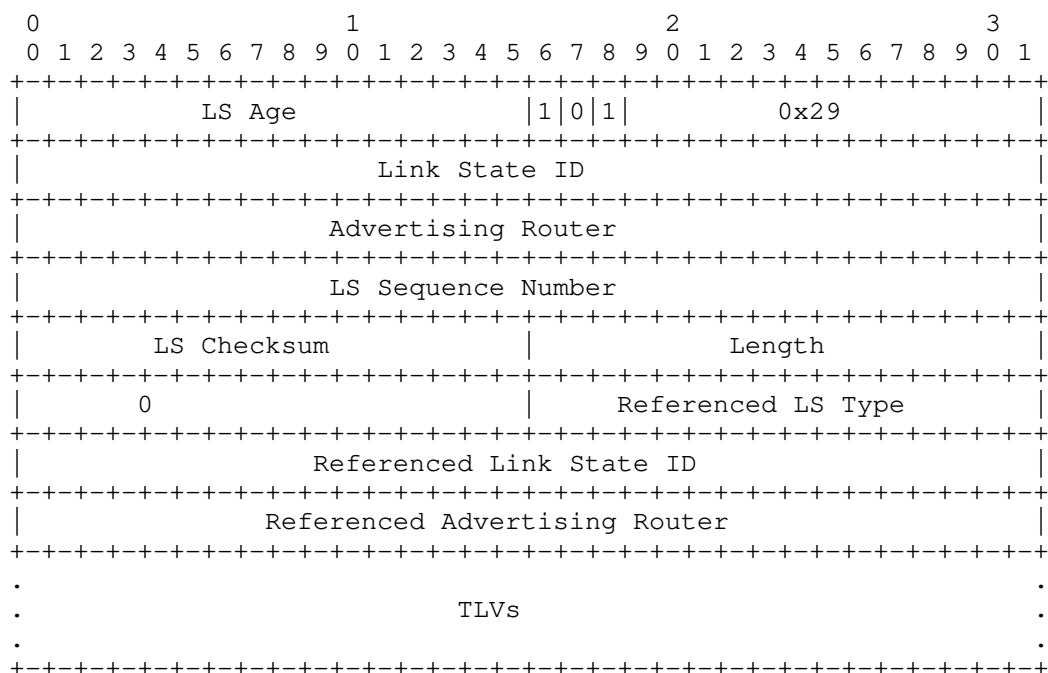


IPv4 Link-Local Address TLV

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. Inclusion in other Extended LSAs MUST be ignored. Only a single instance of the IPv4 Link-Local Address family SHOULD be included in the E-Link-LSA. Instances preceding the first MUST be ignored. For IPv6 address families as defined in [OSPFV3-AF]. Future specifications may support advertisement of routing and topology information for multiple address families. However, this is beyond the scope of this document.

11. 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, section 4.4.3.9 in [OSPFV3]. However, unlike the Intra-Area-Prefix-LSA, it is fully extendable and represented as TLVs.



E-Intra-Area-Prefix-LSA

All LSA Header fields are the same as defined for the Intra-Area-Prefix-LSA. The following top-level TLVs are defined:

- o 6 - Intra-Area-Prefix TLV (defined in Section 10)

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 all TLVs.

12.1. Extended LSA Mixed-Mode Backward Compatibility

An implementation MAY support configuration allowing a mixture of OSPFv3 routers supporting and not supporting TLV-based LSAs in the same OSPFv3 routing domain. In these deployments, the OSPFv3 routers configured with a value of MixedMode or MixedModeDegraded for ExtendedLSASupport, (Appendix A), MUST originate both the TLV-based and non-TLV-based versions of the OSPFv3 LSAs described herein. For the purposes of Shortest Path First (SPF) computation, if the configured value is MixedMode, the TLV-based LSAs MUST be used by OSPFv3 routers supporting this specification. If MixedModeDegraded is configured, the non-TLV-based versions of the OSPFv3 LSAs are used for SPF computation. OSPFv3 routers configured for mixed mode operation also MUST form adjacencies with OSPFv3 Routers sending OSPFv3 Hello and Database Description packets with the options field EL-bit clear. In this manner, OSPFv3 routing domains utilizing the new encodings can be gradually migrated with a worst-case cost of approximately doubling the number of LSAs in the routing domain.

12.2. 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 are 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.

13. Security Considerations

In general, extendible 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 which 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.

14. IANA Considerations

This specification defines nine OSPFv3 Extended LSA types as described in Section 2.

This specification also creates two registries OSPFv3 Extended-LSAs 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.

The OSPFv3 Extend-LSA TLV registry will define top-level TLVs for Extended-LSAs and should be placed in the existing OSPFv3 IANA registry. New values can be allocated via IETF Consensus or IESG Approval.

Nine initial values are 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

The OSPFv3 Extend-LSA sub-TLV registry will define sub-TLVs at any level of nesting for Extended-LSAs and should be placed in the existing OSPFv3 IANA registry. New values can be allocated via IETF Consensus or IESG Approval.

One initial value is allocated:

- o 0 - Reserved

15. References

15.1. Normative References

- [OSPFV3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [OSPFV3-AF] Lindem, A., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, April 2010.
- [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.

15.2. Informative References

- [MT-OSPFV3] Mirtorabi, S. and A. Roy, "Multi-topology routing in OSPFv3 (MT-OSPFV3)", draft-ietf-ospf-mt-ospfv3-04.txt (work in progress).
- [OSPF-DIGITAL-SIGNATURE] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.

Appendix A. Configurable Constants

An additional global configurable constant will be added to the OSPFv3 protocol.

ExtendedLSASupport

This is an enumeration type indicating the extent to which the OSPFv3 instance supports the TLV format described herein for Extended LSAs. The valid value for the enumeration are:

- * None - Non-extended LSAs will not be originated or used in the SPF calculation.
- * Normal - Extended LSAs will be originated and adjacencies will not be formed with OSPFv3 routers not supporting this specification.
- * MixedMode - Both extended and non-extended LSAs will be originated. OSPFv3 adjacencies will be formed with OSPFv3 routers not supporting this specification. The extended LSAs are used for the SPF computation.
- * MixedModeDegraded - Both extended and non-extended LSAs will be originated. OSPFv3 adjacencies will be formed with OSPFv3 routers not supporting this specification. The non-extended LSAs are used for the SPF computation.

Authors' Addresses

Acee Lindem
Ericsson
301 Midenhall Way
Cary, NC 27513
USA

Email: acee.lindem@ericsson.com

Sina Mirtorabi
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
USA

Email: sina@cisco.com

Abhay Roy
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
USA

Email: akr@cisco.com

Fred Baker
Cisco Systems
Santa Barbara, CA 93117
USA

Email: fred@cisco.com

OSPF Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 22, 2015

A. Atlas
S. Hegde
C. Bowers
Juniper Networks
J. Tantsura
Ericsson
Z. Li
Huawei Technologies
July 21, 2014

OSPF Extensions to Support Maximally Redundant Trees
draft-atlas-ospf-mrt-03

Abstract

This document specifies extensions to OSPF to support the distributed computation of Maximally Redundant Trees (MRT). Some example uses of the MRTs include IP/LDP Fast-Reroute and global protection or live-live for multicast traffic. The extensions indicate what MRT profile(s) each router supports. Different MRT profiles can be defined to support different uses and to allow transitioning of capabilities. An extension is introduced to flood MRT-Ineligible links, due to administrative policy.

The need for a mechanism to allow routers to advertise a worst-case FIB compute/install time is well understood for controlling convergence. This specification introduces the Controlled Convergence TLV to be carried in the Router Information LSA.

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 January 22, 2015.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. Terminology	3
4. Overview of OSPF Extensions for MRT	4
4.1. Supporting MRT Profiles	4
4.2. GADAG Root Selection	5
4.3. Triggering an MRT Computation	5
5. MRT Capability Advertisement	6
5.1. Advertising MRT Capability in OSPFv2	6
5.2. Advertising MRT Capability in OSPFv3	7
5.3. MRT Profile TLV in Router Information LSA	8
6. Advertising MRT-ineligible links for MRT	9
6.1. MRT-Ineligible Link sub-TLV	10
7. Worst-Case Network Convergence Time	10
8. Backwards Compatibility	11
8.1. Handling MRT Capability Changes	12
9. Implementation Status	12
10. Security Considerations	12
11. IANA Considerations	12
12. References	13
12.1. Normative References	13
12.2. Informative References	13
Authors' Addresses	14

1. Introduction

This document describes the OSPF extensions necessary to support the architecture that defines how IP/LDP Fast-Reroute can use MRTs [I-D.ietf-rtgwg-mrt-frr-architecture]. At least one common standardized algorithm (such as the lowpoint algorithm explained and fully documented in [I-D.ietf-rtgwg-mrt-frr-algorithm]) is required

so that the routers supporting MRT computation consistently compute the same MRTs. MRT can also be used to protect multicast traffic via either global protection or local protection. [I-D.atlas-rtgwg-mrt-mc-arch]

IP/LDP Fast-Reroute using MRTs can provide 100% coverage for link and node failures in an arbitrary network topology where the failure doesn't split the network. It can also be deployed incrementally inside an OSPF area; an MRT Island is formed of connected supporting routers and the MRTs are computed inside that island.

In the default MRT profile, a supporting router both computes the MRTs and creates the necessary transit forwarding state necessary to provide the two additional forwarding topologies, known as MRT-Blue and MRT-Red.

2. Requirements Language

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 [RFC2119]

3. Terminology

For ease of reading, some of the terminology defined in [I-D.ietf-rtgwg-mrt-frr-architecture] is repeated here.

network graph: A graph that reflects the network topology where all links connect exactly two nodes and broadcast links have been transformed into the standard pseudo-node representation.

Redundant Trees (RT): A pair of trees where the path from any node X to the root R along the first tree is node-disjoint with the path from the same node X to the root along the second tree. These can be computed in 2-connected graphs.

Maximally Redundant Trees (MRT): A pair of trees where the path from any node X to the root R along the first tree and the path from the same node X to the root along the second tree share the minimum number of nodes and the minimum number of links. Each such shared node is a cut-vertex. Any shared links are cut-links. Any RT is an MRT but many MRTs are not RTs.

MRT Island: From the computing router, the set of routers that support a particular MRT profile and are connected via MRT-eligible links.

GADAG: Generalized Almost Directed Acyclic Graph - a graph that is the combination of the ADAGs of all blocks. Transforming a network graph into a GADAG is part of the MRT algorithm.

MRT-Red: MRT-Red is used to describe one of the two MRTs; it is used to describe the associated forwarding topology and MT-ID. Specifically, MRT-Red is the decreasing MRT where links in the GADAG are taken in the direction from a higher topologically ordered node to a lower one.

MRT-Blue: MRT-Blue is used to describe one of the two MRTs; it is used to describe the associated forwarding topology and MT-ID. Specifically, MRT-Blue is the increasing MRT where links in the GADAG are taken in the direction from a lower topologically ordered node to a higher one.

4. Overview of OSPF Extensions for MRT

There are two separate aspects that need to be advertised in OSPF. Both derive from the need for all routers supporting an MRT profile to compute the same pair of MRTs to each destination. By executing the same algorithm on the same network graph, distributed routers will compute the same MRTs. Convergence considerations are discussed in [I-D.ietf-rtgwg-mrt-frr-architecture].

The first aspect that must be advertised is which MRT profile(s) are supported and the associated GADAG Root Selection Priority. The second aspect that must be advertised is any links that are not eligible, due to administrative policy, to be part of the MRTs. This must be advertised consistently across the area so that all routers in the MRT Island use the same network graph.

4.1. Supporting MRT Profiles

An MRT Profile defines the exact MRT Algorithm, the MRT-Red LDP MT-ID, the MRT-Blue LDP MT-ID, and the forwarding mechanisms supported for the transit MRT-Red and MRT-Blue forwarding topologies. Finally, the MRT Profile defines exact behavioral rules such as:

- o how reconvergence is handled,
- o inter-area forwarding behavior,

A router that advertises support for an MRT Profile MUST provide the specified forwarding mechanism for its MRT-Red and MRT-Blue forwarding topologies. A router that advertises support for an MRT Profile MUST implement an algorithm that produces the same set of MRT-Red and MRT-Blue next-hops for its MRT-Red and MRT-Blue

topologies as is provided by the algorithm specified in the MRT Profile.

A router MAY indicate support for multiple MRT Profiles. A router computes its local MRT Island for each separate MRT Profile that the router supports. Supporting multiple MRT Profiles also provides a mechanism for transitioning from one profile to another. Different uses of MRT forwarding topologies may behave better on different MRT profiles.

The default MRT Profile is defined in [I-D.ietf-rtgwg-mrt-frr-architecture]. Its behavior is intended to support IP/LDP unicast and multicast fast-reroute.

4.2. GADAG Root Selection

One aspect of the MRT algorithms is that the selection of the GADAG root can affect the alternates and the traffic through that GADAG root. Therefore, it is important to provide an operator with control over which router will play the role of GADAG root. A measure of the centrality of a node may help determine how good a choice a particular node is.

The GADAG Root Selection Policy (defined as part of an MRT profile) may make use of the GADAG Root Selection Priority value advertised in the MRT Profile TLV of the Router Information LSA. For example, the GADAG Root Selection Policy for the default MRT profile is the following: Among the routers in the MRT Island and with the highest priority advertised, an implementation MUST pick the router with the highest Router ID to be the GADAG root.

4.3. Triggering an MRT Computation

When an MRT Computation is triggered, it is triggered for a given MRT Profile in a given area. First, the associated MRT Island is determined. Then, the GADAG Root is selected. Finally, the actual MRT algorithm is run to compute the transit MRT-Red and MRT-Blue topologies. Additionally, the router MAY choose to compute MRT-FRR alternates or make other use of the MRT computation results.

Prefixes can be attached and detached and have their associated MRT-Red and MRT-Blue next-hops computed without requiring a new MRT computation.


5. MRT Capability Advertisement

A router that supports MRT indicates this by setting a newly defined M bit in the Router LSA. If the router provides no other information via a separate MRT Profile TLV, then the router supports the default MRT Profile with a GADAG Root Selection Priority of 128.

In addition, a router can advertise a newly-defined MRT Profile TLV within the scope of the OSPF router information LSA [RFC4970]. This TLV also includes the GADAG Root Selection Priority.

5.1. Advertising MRT Capability in OSPFv2

A new M-bit is defined in the Router-LSA (defined in [RFC2328] and updated in [RFC4915]), as pictured below.

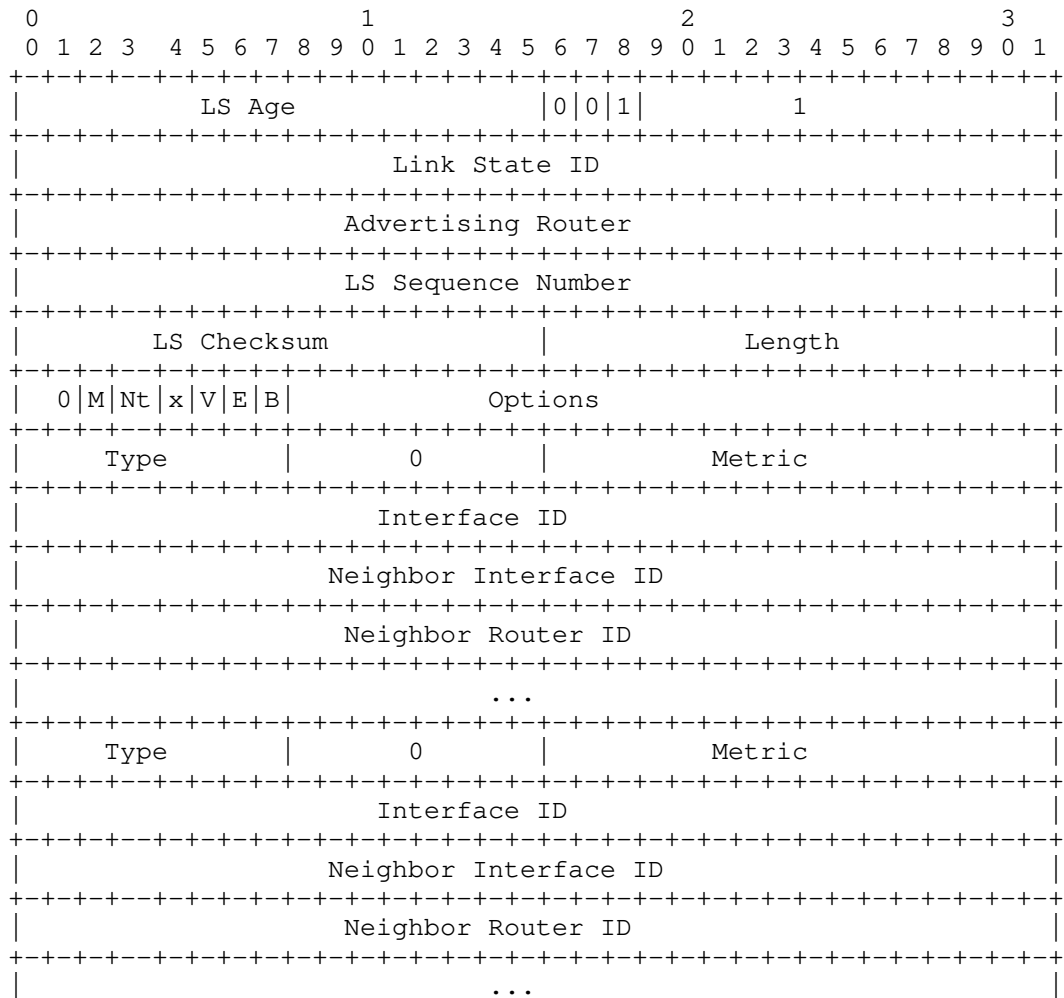


M-bit in OSPFv2 Router LSA

M bit: When set, the router supports MRT. If no separate MRT Profile TLV is advertised in the associated Router Information LSA, then the router supports the default MRT Profile and has a GADAG Root Selection Priority of 128.

5.2. Advertising MRT Capability in OSPFv3

Similarly, the M-bit is defined in the OSPFv3 Router LSA as shown below. Since there can be multiple router LSAs, the M-bit needs to be set on all of them.



M-bit in OSPFv3 Router LSA

M bit: When set, the router supports MRT. If no separate MRT Profile TLV is advertised in the associated Router Information LSA, then the router supports the default MRT Profile and has a GADAG Root Selection Priority of 128.

5.3. MRT Profile TLV in Router Information LSA

A router may advertise an MRT Profile TLV to indicate support for multiple MRT Profiles, for a non-default MRT Profile, and/or to indicate a non-default GADAG Root Selection Priority. The MRT Profile TLV is advertised within the OSPF router information LSA

which is defined for both OSPFv2 and OSPFv3 in [RFC4970]. The RI LSA MUST have area scope.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Type										Length																													
Profile ID(1)										GADAG Priority										Reserved																			
																																						
Profile ID(n)										GADAG Priority										Reserved																			

TYPE: TBA-MRT-OSPF-1 (To Be Allocated by IANA)

LENGTH: 4 * (number of Profiles)

Profile ID : 1 byte

GADAG Priority: 1 byte

MRT Profile TLV in Router Information LSA

Each Profile ID listed indicates support for a given MRT Profile, as defined in [I-D.ietf-rtgwg-mrt-frr-architecture]. A Profile ID value of 0 corresponds to the default MRT profile.

The GADAG Priority is the GADAG Root Selection Priority associated with the advertising router in the MRT Island for the associated MRT Profile, as indicated by the Profile ID. If multiple MRT Profiles are supported, then the length of this TLV varies. The ordering of the profiles inside the TLV is not significant. Multiple appearances of this TLV is not an error.

Lack of support for the default MRT profile is indicated by the presence of an MRT Profile TLV with a non-zero Profile ID value, and the absence of an MRT Profile TLV with a zero Profile ID value.

6. Advertising MRT-ineligible links for MRT

Due to administrative policy, some otherwise eligible links in the network topology may need to be excluded from the network graph upon which the MRT algorithm is run. Since the same network graph must be used across the area, it is critical for OSPF to flood which links to exclude from the MRT calculation. This is done by introducing a new MRT-Ineligible Link sub-TLV. For OSPFv2, this sub-TLV is carried in

the Extended Link TLV defined in [I-D.ietf-ospf-segment-routing-extensions]. For OSPFv3, this sub-TLV is carried in the Router-Link TLV defined in [I-D.ietf-ospf-ospfv3-lsa-extend].

If a link is marked by administrative policy as MRT-Ineligible, then a router MUST flood the OSPFv2 Extended Link TLV (or OSPFv3 Router-Link TLV) for that link, including the MRT-Ineligible Link sub-TLV. The OSPFv2 Extended Link TLV and OSPFv3 Router-Link TLV have area wide scope.

6.1. MRT-Ineligible Link sub-TLV

0										1										2										3											
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
Type										Length																															

TYPE: TBA-MRT-OSPF-2 in OSPFv2 Extended Link TLV
 TBA-MRT-OSPF-3 in OSPFv3 Router-Link TLV
 (To Be Allocated by IANA)
 LENGTH: 0

MRT-Ineligible Link sub-TLV

This zero-length sub-TLV can appear in the OSPFv2 Extended Link TLV or the OSPFv3 Router-Link TLV. Its presence indicates that the associated link MUST NOT be used in the MRT calculation for all profiles.

7. Worst-Case Network Convergence Time

As part of converging the network after a single failure, Section 12.2 of [I-D.ietf-rtgwg-mrt-frr-architecture] describes the need to wait for a configured or advertised period for all routers to be using their new SPTs. Similarly, any work on avoiding micro-forwarding loops during convergence[RFC5715] requires determining the maximum among all routers in the area of the worst-case route computation and FIB installation time. More details on the specific reasoning and need for flooding it are given in [I-D.atlas-bryant-shand-lf-timers].

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+-----																																							

TYPE: TBA-MRT-OSPF-4 (To Be Allocated by IANA)

LENGTH: 4

FIB compute/install time: This is the worst-case time the router may take to compute and install all OSPF routes in the area after a change to a stable network. The value is in milliseconds.

Controlled Convergence TLV in Router Information LSA

The Controlled Convergence TLV is carried in the Router Information LSA and flooded with area-wide scope. The FIB compute/install time value sent by a router SHOULD be an estimate taking into account network scale or real-time measurements, or both. Advertisements SHOULD be dampened to avoid frequent communication of small changes in the FIB compute/install time.

A router receiving the Controlled Convergence sub-TLV SHOULD estimate the network convergence time as the maximum of the FIB compute/install times advertised by the routers in an area, including itself. In order to account for routers that do not advertise the Controlled Convergence sub-TLV, a router MAY use a locally configured minimum network convergence time as a lower bound on the computed network convergence time. A router MAY use a locally configured maximum network convergence time as an upper bound on the computed network convergence time.

8. Backwards Compatibility

The MRT capability bit, the MRT Profile, the MRT-Ineligible Link, and the OSPFv3 MRT-Ineligible Link TLVs are defined in this document. They should not introduce any interoperability issues. Routers that do not support the MRT capability bit in the router LSA SHOULD silently ignore it. Routers that do not support the new MRT-related TLVs SHOULD silently ignore them.

8.1. Handling MRT Capability Changes

When a router changes from supporting MRT to not supporting MRT, it is possible that Router Information LSAs with MRT-related TLVs remain in the neighbors' database briefly. Such MRT-related TLVs SHOULD be ignored when the associated Router LSA from that router does not have the MRT capability set in its Router LSA.

When a router changes from not supporting MRT to supporting MRT, it will flood its Router LSA(s) with the M-bit set and may send an updated Router Information LSA. If a Router LSA is received with the M-bit newly set, an MRT computation SHOULD be scheduled but MAY be delayed up to 60 seconds to allow reception of updated related Router Information LSAs. In general, when changes in MRT-related information is received, an MRT computation SHOULD be triggered.

The rationale behind using the M bit in router LSA is to handle the MRT capability changes gracefully in case of version upgrade/downgrade. The M bit in router LSA ensures the latest "MRT capability" information is available for computation when there is a downgrade to the version that doesn't support MRT and RI LSA.

9. Implementation Status

[RFC Editor: please remove this section prior to publication.]

Please see [I-D.ietf-rtgwg-mrt-frr-architecture] for details on implementation status.

10. Security Considerations

This OSPF extension is not believed to introduce new security concerns. It relies upon the security architecture already provided for Router LSAs and Router Information LSAs.

11. IANA Considerations

IANA is requested to allocate values for the following OSPF Router Information TLV Types [RFC4970]: MRT Profile TLV (TBA-MRT-OSPF-1), and Controlled Convergence TLV (TBA-MRT-OSPF-4).

IANA is requested to allocate a value from the OSPF Extended Link LSA sub-TLV registry [I-D.ietf-ospf-segment-routing-extensions] for the MRT-Ineligible Link sub-TLV (TBA-MRT-OSPF-2).

IANA is requested to allocate a value from the OSPFv3 Extended-LSA sub-TLV registry [I-D.ietf-ospf-ospfv3-lsa-extend] for the MRT-Ineligible Link sub-TLV (TBA-MRT-OSPF-3).

12. References

12.1. Normative References

- [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-03 (work in progress), May 2014.
- [I-D.ietf-ospf-segment-routing-extensions]
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-01 (work in progress), July 2014.
- [I-D.ietf-rtgwg-mrt-frr-algorithm]
Enyedi, G., Csaszar, A., Atlas, A., Bowers, C., and A. Gopalan, "Algorithms for computing Maximally Redundant Trees for IP/LDP Fast-Reroute", draft-rtgwg-mrt-frr-algorithm-01 (work in progress), July 2014.
- [I-D.ietf-rtgwg-mrt-frr-architecture]
Atlas, A., Kebler, R., Bowers, C., Enyedi, G., Csaszar, A., Tantsura, J., Konstantynowicz, M., and R. White, "An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees", draft-rtgwg-mrt-frr-architecture-04 (work in progress), July 2014.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 4970, July 2007.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

12.2. Informative References

- [I-D.atlas-bryant-shand-lf-timers]
K, A. and S. Bryant, "Synchronisation of Loop Free Timer Values", draft-atlas-bryant-shand-lf-timers-04 (work in progress), February 2008.

- [I-D.atlas-rtgwg-mrt-mc-arch]
Atlas, A., Kebler, R., Wijnands, I., Csaszar, A., and G. Envedi, "An Architecture for Multicast Protection Using Maximally Redundant Trees", draft-atlas-rtgwg-mrt-mc-arch-02 (work in progress), July 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3137] Retana, A., Nguyen, L., White, R., Zinin, A., and D. McPherson, "OSPF Stub Router Advertisement", RFC 3137, June 2001.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, June 2007.
- [RFC5715] Shand, M. and S. Bryant, "A Framework for Loop-Free Convergence", RFC 5715, January 2010.

Authors' Addresses

Alia Atlas
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA

Email: akatlas@juniper.net

Shraddha Hegde
Juniper Networks
Embassy Business Park
Bangalore, KA 560093
India

Email: shraddha@juniper.net

Chris Bowers
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

Email: cbowers@juniper.net

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
USA

Email: jeff.tantsura@ericsson.com

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com