

Network Working Group
Internet Draft
Intended status: Standards Track
Expiration Date: January 2008

Sina Mirtorabi
Force10 Networks
Abhay Roy
Cisco Systems
July 2007

Multi-topology routing in OSPFv3 (MT-OSPFv3)
draft-ietf-ospf-mt-ospfv3-03.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 29, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document describes an extensible mechanism to support multiple topologies (MT) in OSPFv3. These topologies can be used within the same address family in order to compute different paths for different classes of service, or belong to different address families allowing an integrated definition of address family with OSPFv3. The extension described in this document can further facilitate any future extensions of OSPFv3.

1. Introduction

Multi-topology routing as described in this document is similar in concept to M-ISIS [[M-ISIS](#)]. It is used to introduce an integrated definition of other address families in OSPFv3. Each address-family may also need to support multiple topologies, to compute different paths for different classes of service or in-band management network.

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 2119](#) [[RFC-KEYWORDS](#)].

2. Potential Solutions

In order to support multiple topologies at least two different solutions are possible. We discuss them briefly below, and describe issues with each of them.

2.1 Using Different Instance IDs

[INSTANCES] defines a range of Instance IDs for each address family. It is therefore possible to define multiple topologies by using different Instance IDs. However this approach is inefficient due to the overhead involved in having to manage multiple adjacencies, multiple link-state databases etc.

2.2 Using an integrated approach with existing LSAs

Another solution would be to define multiple topologies as an integrated approach within each instance. This can be done by redefining an unused field in the link description of Router LSA and define it as a multi-topology identifier (MT-ID). We will have to generate N link descriptions for a link with N topologies configured on it. This seems wasteful as the link description is replicated N times, further we have some backward compatibility issues.

Also, there is a need to identify prefixes carried for each topology, i.e. prefix-LSAs need to carry MT-IDs and there is no possibility to reuse the existing prefix-LSAs.

3. Proposed Solution

We propose to define new LSAs in order to achieve this. Not only does this allow an optimum definition of topologies within OSPFv3, it also does not have any backward compatibility issues as new LSAs will be ignored by old routers.

The flexible encoding proposed for the new LSAs can also facilitate any future extension in OSPFv3.

4. MT Capability

We define a Multi-topology capability bit in Options field.

```

      0               1               2
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
  +-+-+-+-+-+-+-+-+
  | | | | | | | | | | | | | MT|AF|* |* |DC| R| N|MC| E|V6|
  +-+-+-+-+-+-+-+-+

```

The Options field

MT-bit

This bit is set when a router supports MT-OSPFv3 as specified in this memo.

5. T-bit in LS type

We define a new T-bit (TLV based) in LS type field in order to extend the existing LSAs. This will define new LSA types homogeneously compared to the existing LSA types. The U-bit and the T-bit are set.

```

      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
  +-+-+-+-+-+-+-+-+
  |U |S2|S1|T |      LSA Function Code      |
  +-+-+-+-+-+-+-+-+

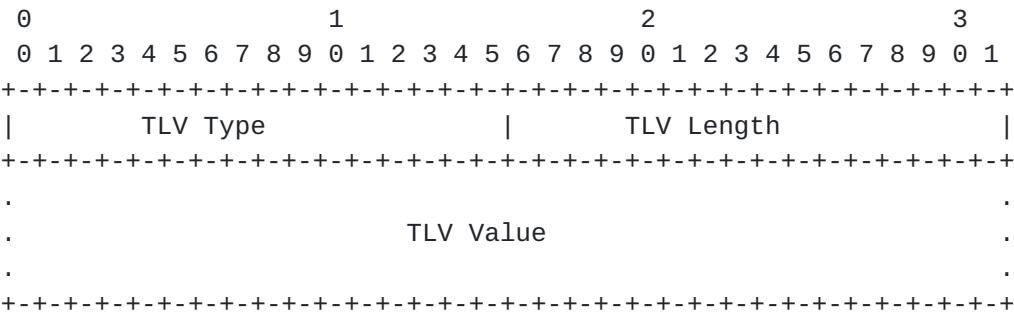
```

For the function codes defined in [OSPFv3] the LS types become:

LSA function code	LS Type	Description
1	0xB001	E-Router-LSA
2	0xB002	E-Network-LSA
3	0xB003	E-Inter-Area-Prefix-LSA
4	0xB004	E-Inter-Area-Router-LSA
5	0xD005	E-AS-External-LSA
6	0xB006	E-Group-membership-LSA
7	0xB007	E-Type-7-LSA
8	0x9008	E-Link-LSA
9	0xB009	E-Intra-Area-Prefix-LSA

6. OSPFv3 TLVs and sub-TLVs

All the Extended LSAs have a flexible TLV format encoding. OSPFv3 TLVs/sub-TLVs have a 16 bit Type and a 16 bit Length field followed by the TLV/sub-TLV value. The Length is set to the length of Value field in bytes. Any unknown TLV/sub-TLV should be ignored.



OSPFv3 TLV Format

7. Default Topology

In order to interact with non-MT capable routers we define default topology as the topology that is built by using the existing LSAs as specified in OSPFv3 [[OSPFv3](#)].

We define MT topologies as topologies which are other than the Default Topology. A MT topology will be defined by using the new LSAs as specified in this memo.

When all routers are MT capable, there is no need to generate existing LSAs as defined in [[OSPFv3](#)]. The new LSAs can be used even for Default Topology. A global configurable parameter RFC2740Compatibility (see [Appendix A](#)) is used to control the generation of existing or new LSAs.

8. MT-ID Fields

We define a 8 bit MT-ID field which is present in various LSA types. Each MT-ID is also accompanied with a MT-ID Metric field which carries a metric specific to one MT.

When a MT capable router participates in Default Topology, depending on RFC2740Compatibility (see [Appendix A](#)) it will generate existing LSAs or extended LSAs for the Default Topology.

MT-ID value 0 is reserved for carrying information about Default Topology in the new LSAs.

9. MT Control packet and IPv6 link local address

IPv6 link local address is MT independent and is used for MT-OSPFv3 control packets.

10. Forming adjacency in MT

Each interface can be configured to belong to a set of topologies. A single adjacency will be formed even if the interface is configured to participate in multiple topologies.

11. Advertising MT topology

When a router is configured with multiple topologies on a link, it advertises the list of MT-IDs and their corresponding metrics in E-router-LSA. When a MT-capable router participates in default topology, based on RFC2740Compatibility it may generate existing or Extended Router-LSAs.

Network LSAs are common to all MT. The DR will announce an adjacency to all attached routers independently of their MT-ID value. When RFC2740Compatibility is disabled on DR (and all routers should be MT capable) E-network-LSA will be used instead of network-LSA. This allow a smooth migration to extended LSAs.

12. Advertising MT prefix

When a MT-capable router participate in non-Default Topology, it generates extended prefix LSAs with MT-ID in which it participates. When a MT-capable router participates in default topology, based on RFC2740Compatibility it may generate existing or Extended prefix LSAs.

13. Advertising intra-area-prefix-LSA on multi-access link

On multi-access links, the DR is responsible to generate prefix-LSA on behalf of the LAN, this is done by considering the prefix advertised in link-LSAs.

If RFC2740Compatibility is disabled the DR will generate Extended prefix-LSAs. If RFC2740Compatibility is enabled we select a Multi-Topology DR (MT-DR) which generates the E-intra-area-prefix-LSA on behalf of the LAN.

MT-DR is elected by considering the highest Router ID among MT-capable routers (done by examining MT-bit of neighbors).

The E-intra-area-prefix-LSA generated by the MT-DR will have the Referenced LS type set to 2, Referenced Link State ID set to DR's Router ID and Referenced Advertising Router set to DR's Router ID. Note that MT-DR's role is to just generate the E-intra-area-prefix-LSA whereas DR is responsible for network LSA generation and helping in flooding on the multi-access link.

14. MT Area Boundary

Area boundaries for all topologies are the same but an interface can be configured to not participate in all topologies. This will make a router's B-bit setting topology independent whereas reachability to the ABR will be topology dependent.

15. MT SPF Computation

When a link participates in a topology, it's MT-ID value is carried in extended Router-LSA. A separate SPF is computed for each topology by considering only the link/metric for the corresponding topology.

Network LSAs are used by all topologies during the SPF computation. Nexthops computed during the MT SPF MUST belong to the same topology.

Similarly when processing prefix-LSAs or external-LSAs, only prefixes which contain a valid MT Metric for that MT SPF are considered reachable in that topology.

During SPF computation for the Default Topology, independently of RFC2740Compatibility value, extended LSA are used when present otherwise existing LSA are used. This allows a smooth transition across RFC2740Compatibility changes.

16. Forwarding in MT

Forwarding must make sure that only routes belonging to the single topology are used to forward the packet along its way from source to destination. Therefore user configuration **MUST** be consistently applied throughout the network so that an incoming packet be associated with the corresponding topology. It is outside of the scope of this document to consider different methods of associating an incoming packet to the corresponding topology routes.

17. MT-ID reserved value

The following MT-ID values are defined:

- | | |
|--------|--|
| 0 | - Reserved for IPv6 unicast topology |
| 1 | - Reserved for IPv4 in-band management purposes |
| 2 | - Reserved for IPv4 unicast topology |
| 3 | - Reserved for IPv6 multicast topology |
| 4 | - Reserved for IPv4 multicast topology |
| 5-31 | - Reserved for IETF consensus. |
| 32-255 | - Reserved for development, experimental and proprietary features [RFC3692]. |

18. IPv4 address family considerations

OSPFv3 runs on the top of IPv6 and uses IPv6 link local addresses for OSPFv3 control packets and next hop calculations. Although IPv6 link local addresses could be used as next hops for IPv4 address families, it is desirable to have IPv4 next hop addresses. For example, in IPv4 multicast having the nexthop address the same as the PIM neighbor address (IPv4 address) makes it easier to know to which upstream neighbor to send a PIM join when doing a RPF lookup for a source. It is also easier for troubleshooting purposes to have a next hop with the same semantics as the AF.

In order to achieve this, the link's IPv4 address will be advertised in the E-link-LSA, see [section 20.6](#).

We call direct interface address (DIA) the address that is reachable directly via the link provided that a layer 3 to layer 2 mapping is available. Note that there is no explicit need for the IPv4 link addresses to be on the same subnet. An implementation should resolve layer 3 to layer 2 mappings via ARP or ND for a DIA even if the IPv4 address is not on the same subnet as the router's interface IP address.

19. Backward compatibility and interaction with non-MT routers

An MT capable router will interact (in Default Topology) with non-MT capable routers by using the existing LSAs. MT information is carried in new LSAs which are ignored by non-MT routers therefore this document does not introduce any backward compatibility issues.

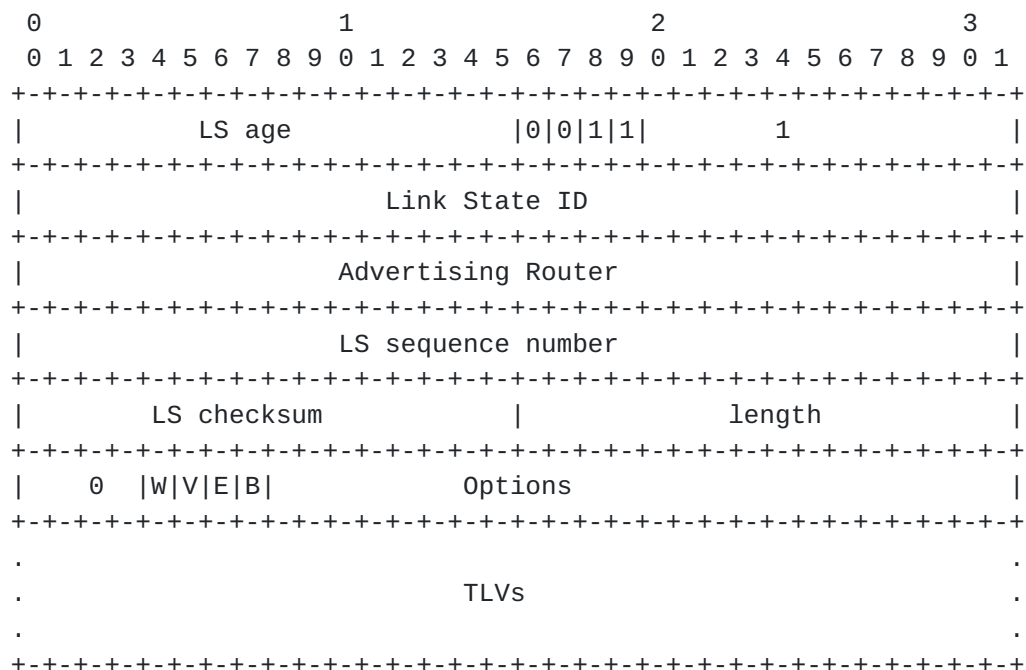
When all routers are MT capable, RFC2740Compatibility can be set to disable and only extended LSAs are used for Default Topology.

20. Extended LSA Formats

20.1 Extended Router-LSA

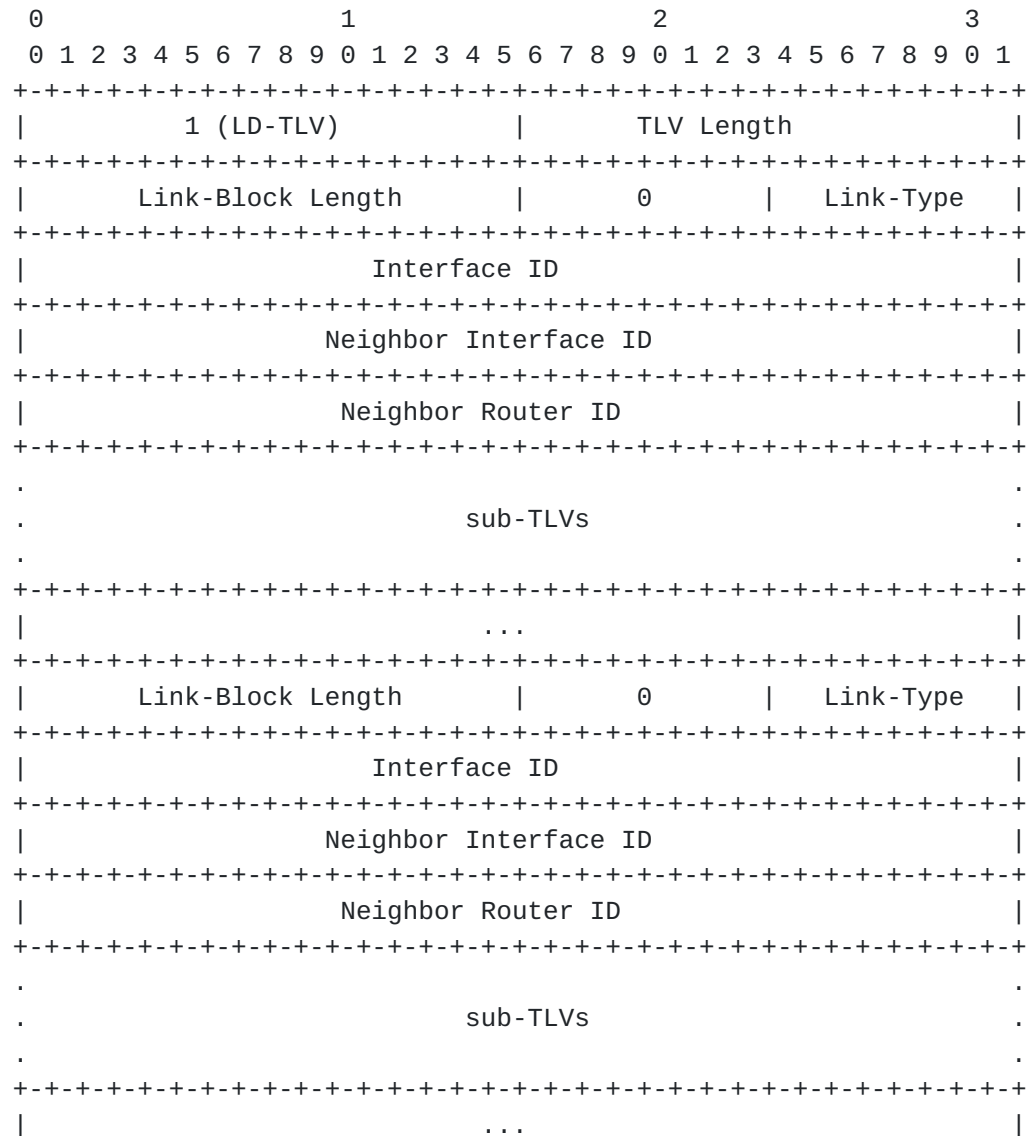
We define a new E-router-LSA with LS type equal to 0xB001. This LSA, extends router LSA by defining TLVs within the LSA payload. The LSA has a fixed portion followed by TLVs. Each TLV could further contains sub-TLVs.

The processing and generation of this LSA is the same as for router-LSA defined in [OSPFv3].



All fields are defined as in [\[OSPFv3\]](#).

We define a Link-description TLV (LD-TLV). This TLV extends the router-LSA payload by defining sub-TLVs within each link description.



Link-Block Length : This field is used as a pointer to the next link block. It defines the length of link block including Link-Block Length, Link-Type field, Interface ID, Neighbor Interface ID, Neighbor Router ID and sub-TLVs

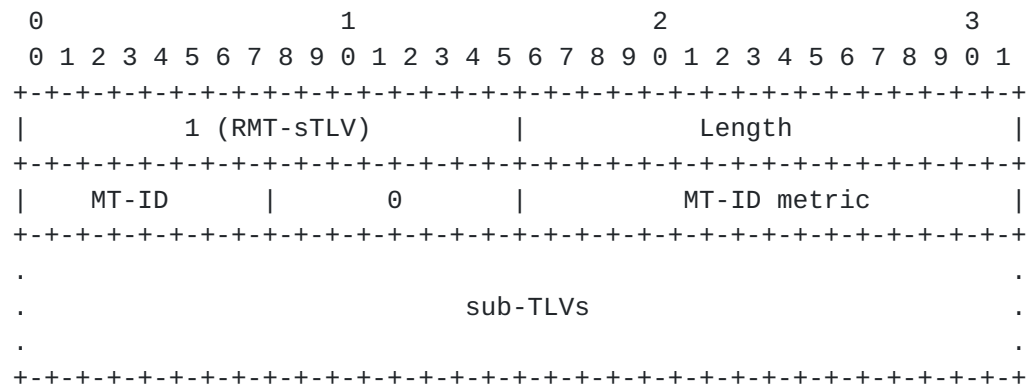
All other fields are defined as in [\[OSPFv3\]](#).

LD-TLV is the only top level TLV defined in this document. This TLV should not be repeated within an E-router-LSA fragment, instead

multiple link descriptions are included within the LD-TLV.
(Link-Block length indicates the next link description).

We define a Router Multi-Topology sub-TLV (RMT-sTLV) below. This sub-TLV could further contain sub-TLVs.

E-router-LSA must contain the LD-TLV and each link description must contain the RMT-sTLV.



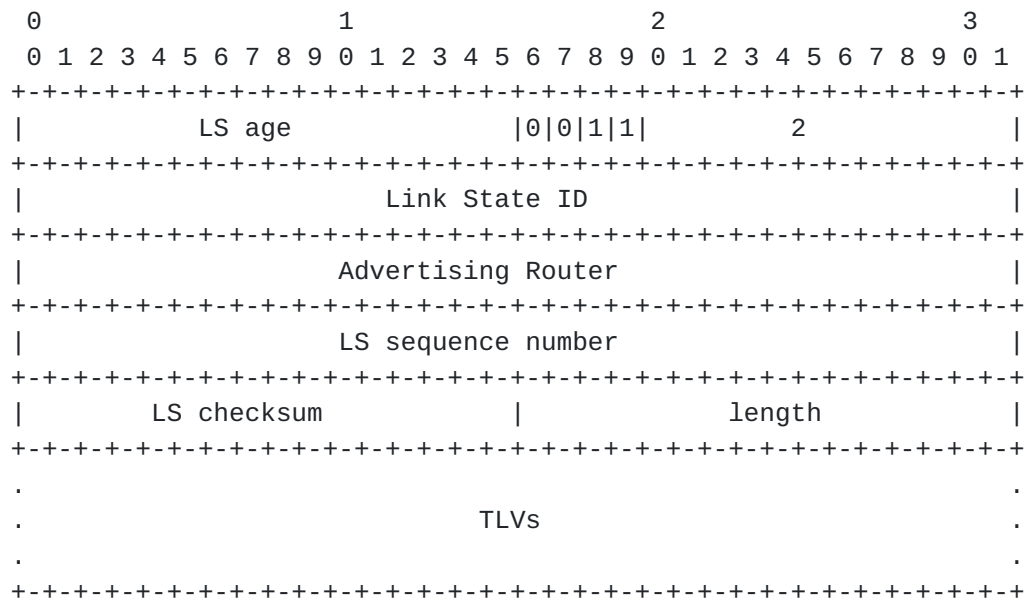
When a link participates in different topologies, it will include the RMT-sTLV with MT-IDs and MT-ID metrics for corresponding topologies.

20.2 Extended Network-LSA

The network LSA does not contain any MT information as the DR is shared by all topologies therefore the existing network LSA can be used independently of the router participation in a MT.

However we define an E-network-LSA in order to allow for any future extensions. The LS type is equal to 0xB002. This LSA extends network-LSA by defining TLVs within the LSA payload.

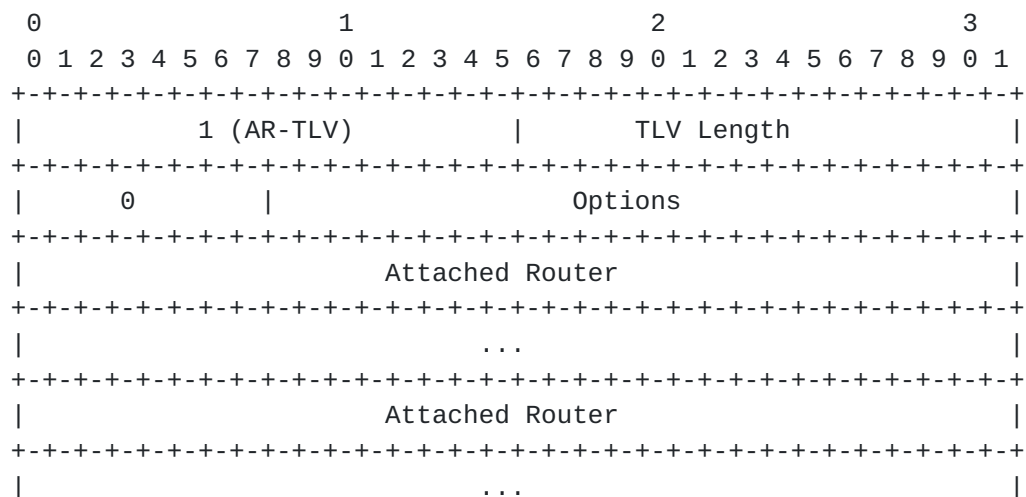
The processing and generation of this LSA is the same as for network-LSA defined in [\[OSPFv3\]](#).



All fields are defined as in [[OSPFv3](#)].

We define a Attached-Router TLV (AR-TLV). This TLV has similar contents as the network-LSA payload.

E-network-LSA must contain AR-TLV.



All fields are defined as in [[OSPFv3](#)].

20.3 Extended Inter-Area-Prefix-LSA

We define a new E-inter-area-prefix-LSA with LS type of 0xB003. This LSA, extends Inter-Area-Prefix-LSA by defining TLVs within the LSA payload. The LSA has a fixed portion followed by TLVs. Each TLV could further contains sub-TLVs. This LSA is originated by area border routers to describe routes to prefixes associated with a MT-ID that belong to other areas.

An implementation could decide to advertise one or more prefixes within one E-inter-area-prefix-LSA.

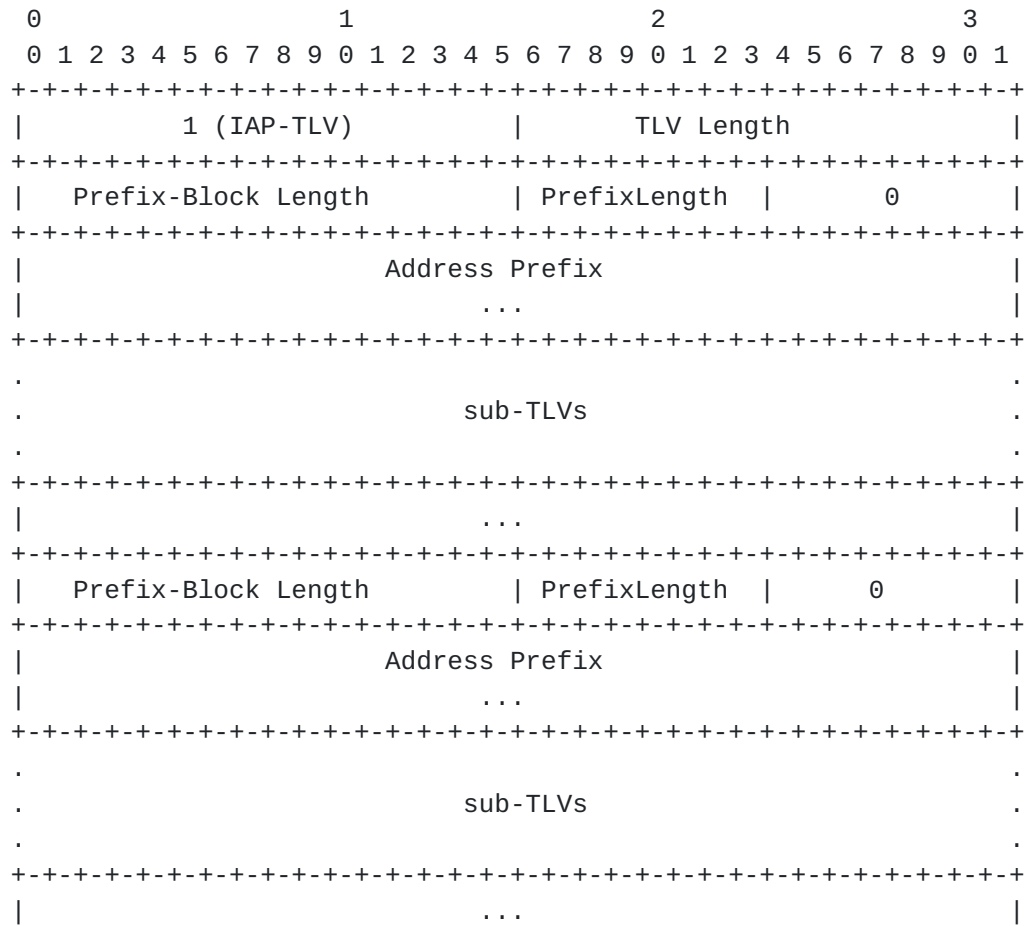
The processing and generation of this LSA is the same as for as inter-area-prefix-LSA as defined in [OSPFv3].

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS age               |0|0|1|1|               3       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link State ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Advertising Router               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS sequence number               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS checksum               |               length       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.               TLVs
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

We define an Inter-area Prefix TLV (IAP-TLV). This TLV extends the Inter-Area-Prefix-LSA payload by defining sub-TLVs within each each prefix.




Prefix-Block Length : This field is used as a pointer to the next prefix block. It define the length of prefix block including Prefix-Block Length, PrefixLength, Address Prefix and sub-TLVs.

All other fields are defined as in [\[OSPFv3\]](#).

IAP-TLV is the only top level TLV defined by this document. This TLV should not be repeated within an E-Inter-area-prefix-LSA, instead multiple prefix values are included within the IAP-TLV. (Prefix-Block Length indicates the next prefix block).

We define an Inter area Multi-Topology sub-TLV (IAMT-sTLV) below. This sub-TLV could further contain sub-TLVs.

Each prefix block must contain an Inter area Multi-Topology sub-TLV (IAMT-sTLV).




[20.4](#) Extended Inter-Area-Router-LSA

We define a new E-inter-area-router-LSA with LS type of 0xB004. This LSA, extends Inter-Area-Router-LSA by defining TLVs within the LSA payload. The LSA has a fixed portion followed by TLVs. Each TLV could further contains sub-TLVs. This LSA is originated by area border routers and describes routes to routers in other areas.

An implementation could decide to advertise information about one or more routers within one E-inter-area-router-LSA.

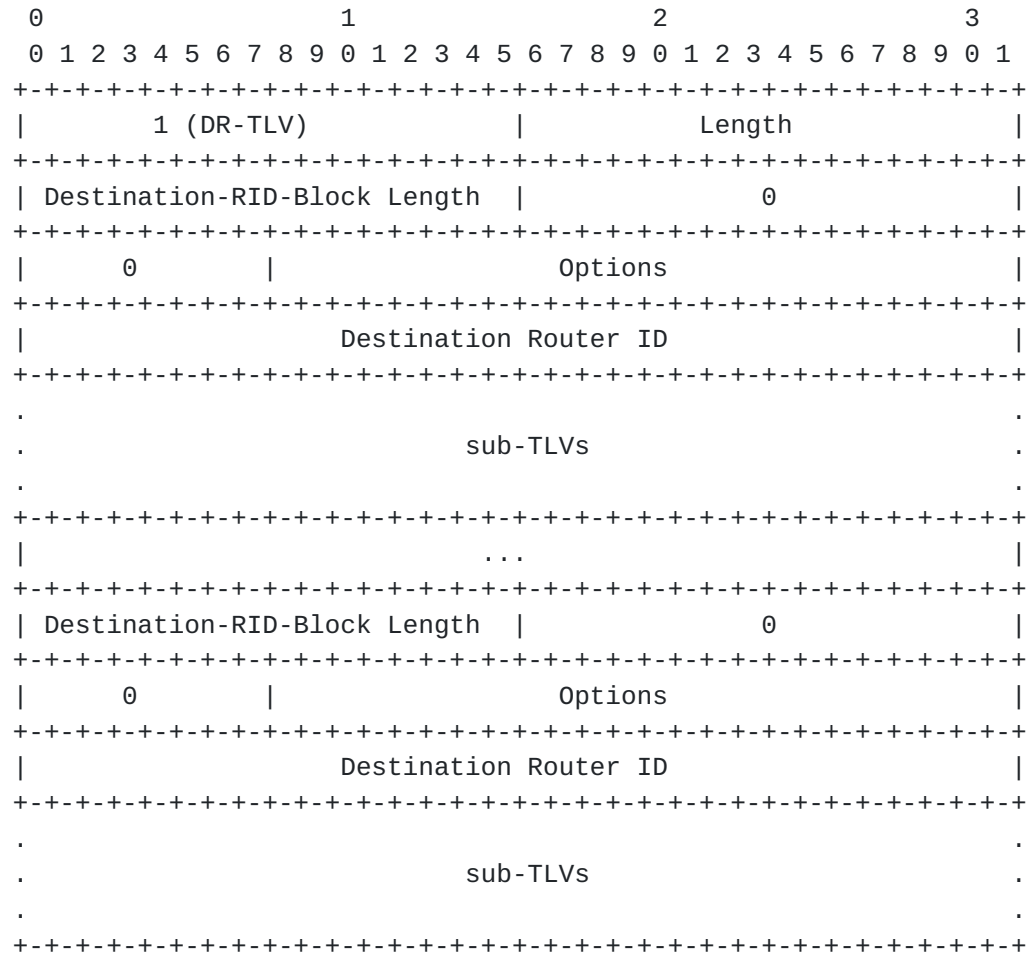
The processing and generation of this LSA is the same as for as inter-area-router-LSA as defined in [[OSPFv3](#)].



All fields are defined as in [[OSPFv3](#)].

We define a Dest-Router TLV (DR-TLV) below. This TLV extends the Inter-area-router-LSA payload by defining sub-TLVs within each Destination Router.

E-inter-area-router-LSA must contain the DR-TLV.



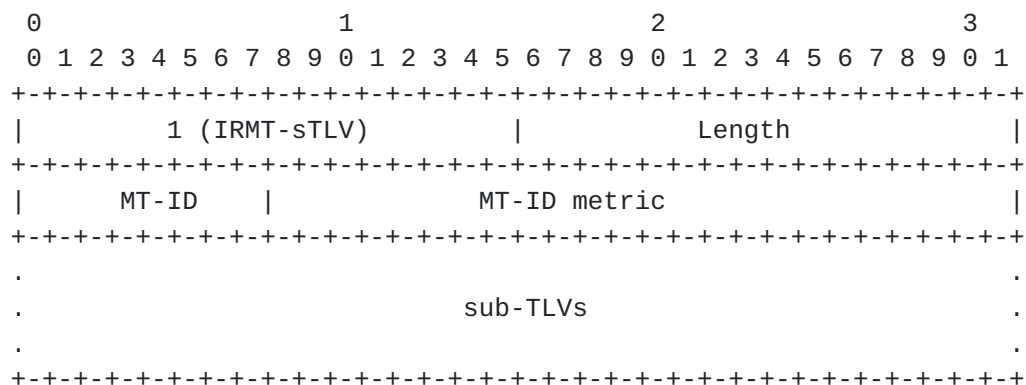
Destination-RID-Block Length : This field is used as a pointer to the next Destination Router ID block. It defines the length of Destination Router ID block including Destination-RID-Block Length, Options, Destination Router ID and sub-TLVs.

Destination Router ID: It is defined in [[OSPFv3](#)].

DR-TLV is the only top level TLV defined by this document. This TLV should not be repeated within an E-Inter-area-router-LSA, instead multiple destination router values are included within the DR-TLV (Destination-RID-Block Length indicates the next destination router block).

We define an Inter Router Multi-Topology sub-TLV (IRMT-sTLV) below.

DR-TLV must contain the IRMT-sTLV.



20.5 Extended AS-external-LSA

We define a new E-AS-external-LSA with LS type of 0xD005. This LSA extends AS-External-LSA by defining TLVs within the LSA payload. The LSA has a fixed portion followed by TLVs. Each TLV could further contains sub-TLVs. E-AS-External-LSA is originated by AS boundary routers, and describes destinations external to the AS associated to a MT-ID.

An implementation could decide to advertise one or more prefixes within one E-AS-external-LSA.

The processing and generation of this LSA is the same as for an AS-external-LSA as defined in [[OSPFv3](#)].

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS age               |0|1|0|1|               5       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Link State ID               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Advertising Router               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               LS sequence number               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.               TLVs
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

We define an External Prefix TLV (EP-TLV). This TLV extends the External-Prefix-LSA payload by defining sub-TLVs within each prefix.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               1 (EP-TLV)               | TLV Length               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Prefix-Block Length               | PrefixLength | 0               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Address Prefix               |
|               ...               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.               sub-TLVs
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               ...               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Prefix-Block Length               | PrefixLength | 0               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Address Prefix               |
|               ...               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.               sub-TLVs
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

|

...

|

Mirtorabi, Roy

[Page 17]

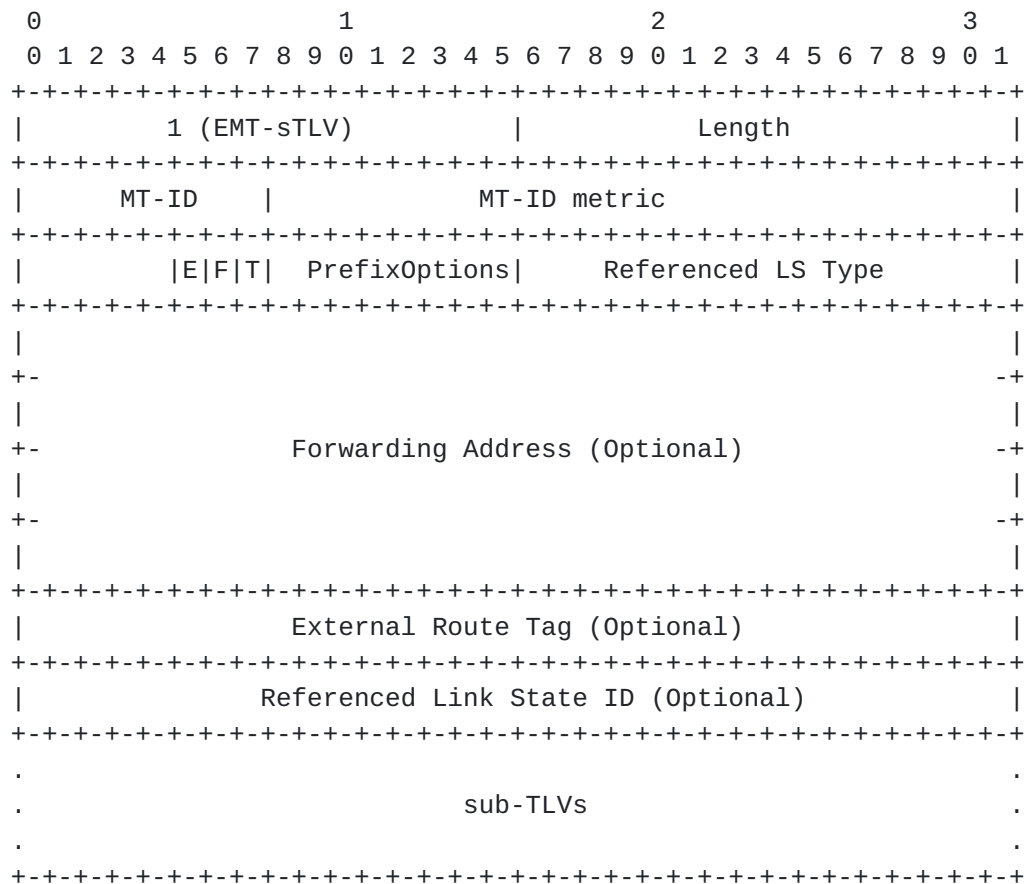
Prefix-Block Length : This field is used as a pointer to the next prefix block. It defines the length of prefix block including Prefix-Block Length, PrefixLength, Address Prefix and sub-TLVs.

All other fields are defined as in [\[OSPFv3\]](#).

EP-TLV is the only top level TLV defined by this document. This TLV should not be repeated within an E-AS-External-LSA, instead multiple prefix are included within the EP-TLV. (Prefix-Block Length indicates the next prefix block).

We define an External Multi-Topology sub-TLV (EMT-sTLV) below. This EMT-subTLV could further contain sub-TLVs.

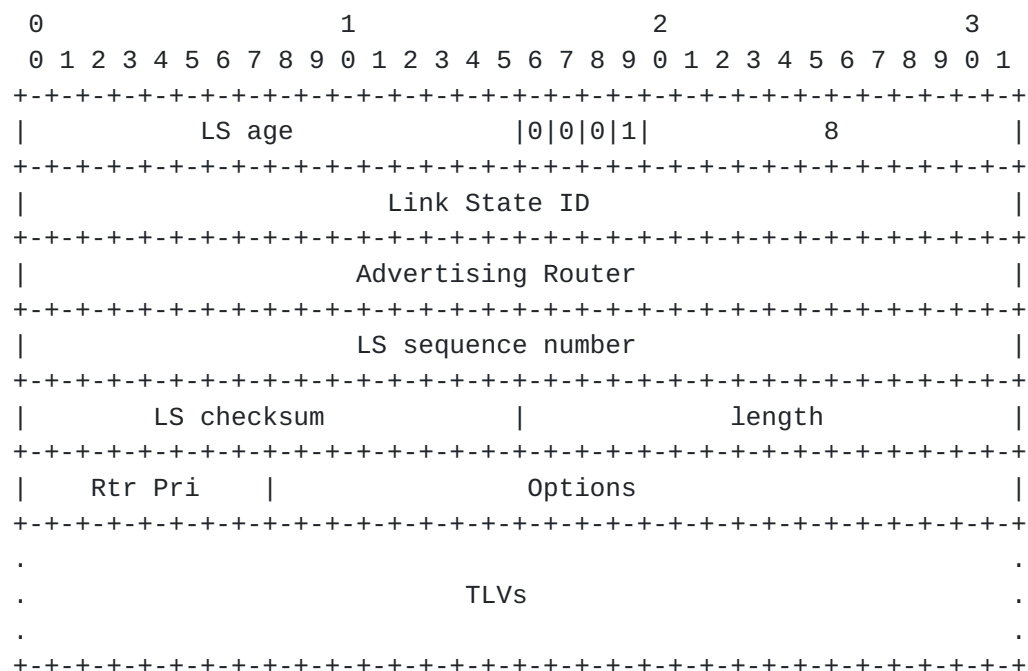
Each prefix block must contain an External Multi-Topology sub-TLV (EMT-sTLV).



20.6 Extended Link-LSA

We define a new E-link-LSA with LS type of 0x9008. This LSA extends Link-LSA by defining TLVs within the LSA payload. The LSA has a fixed portion followed by TLVs. Each TLV could further contains sub-TLVs. E-Link-LSA is generated for each link and carries each link's prefix in the corresponding topology. It also carries next hop IP information for the supported address families.

The processing and generation of this LSA is the same as for Link-Lsa as defined in [OSPFv3].



We define the following three TLVs

- o IPv6 Next hop TLV (NH6-TLV)
- o IPv4 Next hop TLV (NH4-TLV)
- o Prefix Multi-topology TLV (PMT-TLV)

Next hop TLVs carry IPv6/IPv4 information for next hop calculation. IPv6 next hop information MUST be a link local IPv6 address. Prefix-TLV carries router link's prefix on multi-access link. This information is used by DR in order to include those prefix in its E-intra-area prefix LSA.

NH6-TLV has the following format:


```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               1 (NH6-TLV)               |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                                 |
+-                                                                 -+
|                                                                 |
+-                                                                 -+
|               Link-local Interface Address               |
+-                                                                 -+
|                                                                 |
+-                                                                 -+
|                                                                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

This TLV MUST be present if the link participate in a MT belonging to IPv6 address family.

NH4-TLV has the following format:

```

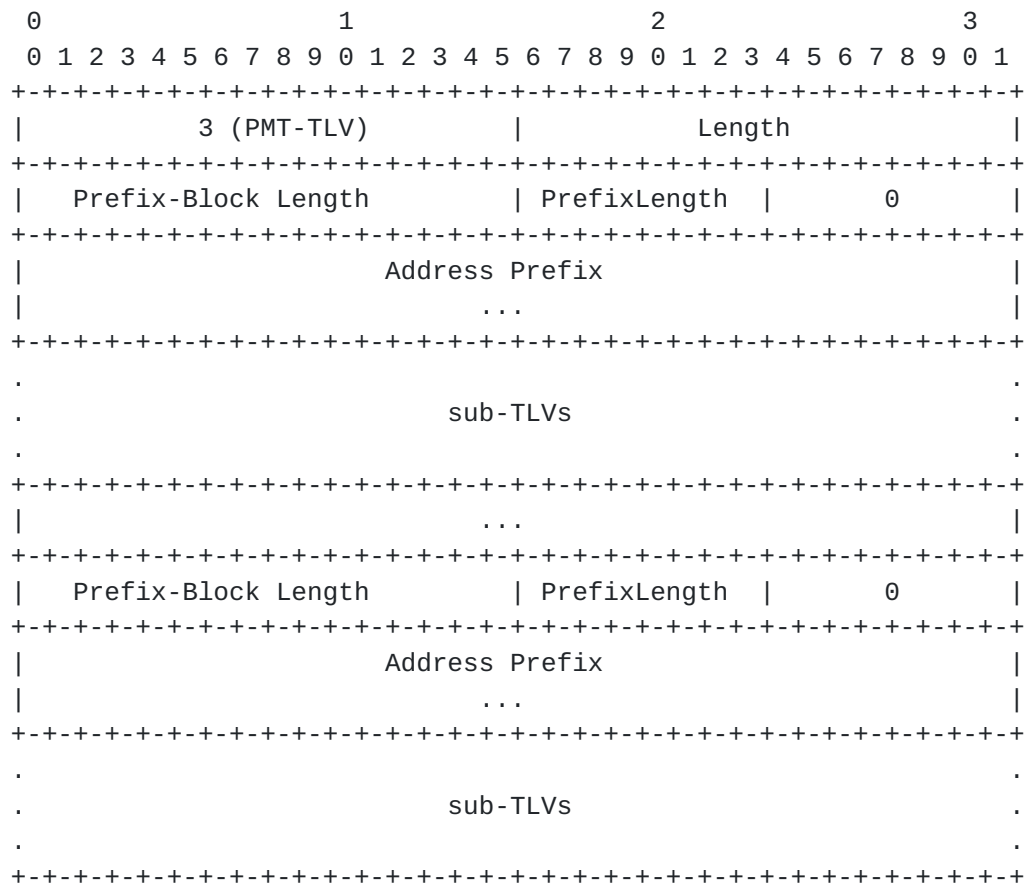
      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               2 (NH4-TLV)               |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               IPv4 address               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

This TLV MUST be present if the link participate in a MT belonging to IPv4 address family.

PMT-TLV has the following format:

PMT-TLV extends link-LSA by defining TLV under each address prefix.
This TLV should only be present on multi-access links.



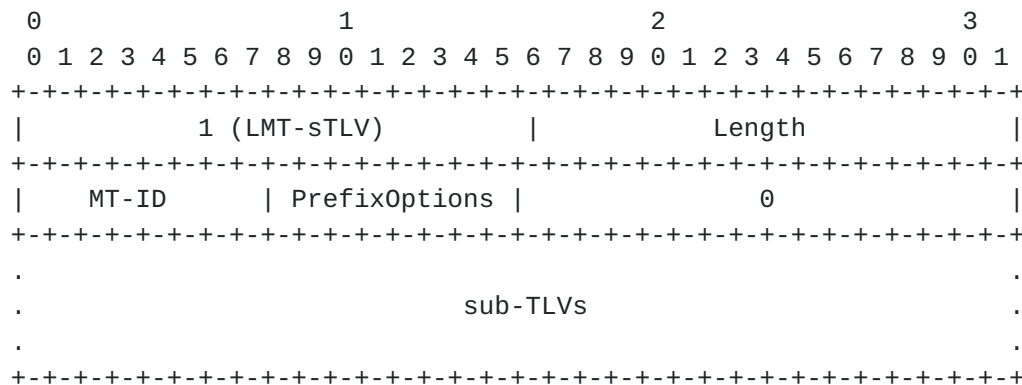
Prefix-Block Length : This field is used as a pointer to the next prefix block. It defines the length of prefix block including Prefix-Block Length, PrefixLength Address Prefix and sub-TLVs.

All other fields are defined as in [\[OSPFv3\]](#).

PMT-TLV should not be repeated within an E-Link-LSA, instead multiple prefix are included within the PMT-TLV. (Prefix-Block Length indicates the next prefix block).

We define a Link Multi-Topology sub-TLV (LMT-sTLV) below. This sub-TLV could further contain sub-TLVs.

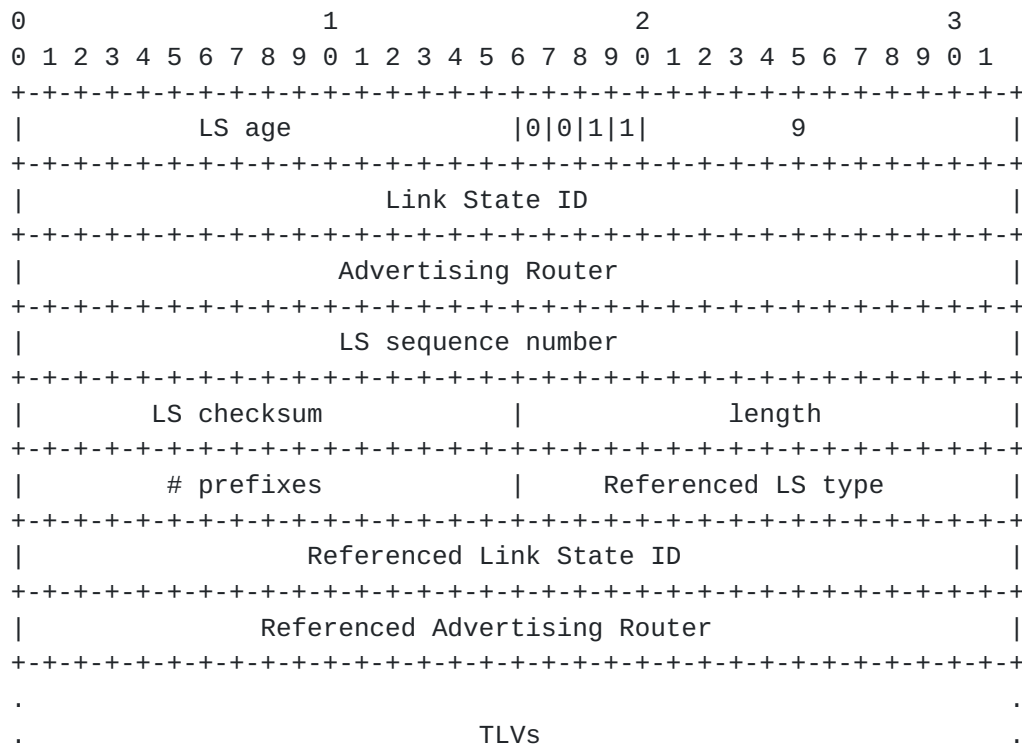
Each prefix block must contain the Link Multi-Topology sub-TLV (LMT-sTLV).



20.7 Extended Intra-Area-Prefix-LSA

We define a new E-intra-area-prefix-LSA with LS type of 0xB009. This LSA extends Intra-Area-Prefix-LSA by defining TLVs within the LSA payload. The LSA has a fixed portion followed by TLVs. Each TLV could further contains sub-TLVs. A router generates E-Intra-Area-Prefix-LSAs to advertise one or more prefixes associated with a topology.

The processing and generation of this LSA is the same as for intra-area-prefix-LSA defined in [[OSPFv3](#)].



+ - + - + - + - + - + - + - + - + - + - + - + - + - + - + - + - + - +

Mirtorabi, Roy

[Page 22]

[illegible]

All other fields are defined as in [[OSPFv3](#)].

We define an Intra-Area Multi-Topology sub-TLV (IMT-sTLV) below. This sub-TLV could further contain sub-TLVs.

Each prefix block must contain Intra-Area Multi-Topology sub-TLV (IMT-STLV).


```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           1 (IMT-sTLV)           |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   MT-ID       | PrefixOptions |           MT-ID metric       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.                               sub-TLVs                               .
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

21. Security Considerations

The technique described in this document does not introduce any new security issues to the OSPFv3 protocol.

22. Acknowledgements

The authors would like to thank Alex Zinin, Acee Lindem, Tom Henderson, Jeff Ahrenholz and Paul Wells for their comments on the document.

23. IANA Considerations

IANA is requested to create a new registry, "OSPFv3 multi-topology ID values" with the assignment listed in [Section 17](#) of this document and registration policies for future assignments.

24. References

Normative References

- [OSPFv3] R. Coltun, D. Ferguson and J. Moy, "OSPF for IPv6", [RFC 2740](#), December 1999.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), January 2004.
- [RFC-KEYWORDS] Bradner, S., "Key words for use in RFC's to Indicate Requirement Levels", [RFC 2119](#), March 1997.

Informative Reference

[INSTANCES] Mirtorabi, S. et al, "Support of address families in OSPFv3", Internet Draft, work in progress

[M-ISIS] T. Przygienda, N. Shen, N. Sheth, "M-ISIS: Multi Topology (MT) Routing in IS-IS", Internet Draft, work in progress

Appendix A. Global Parameter

RFC2740Compatibility

This parameter controls which LSAs are used for Default Topology. When set to "enabled", the Default Topology is described using existing LSAs [[OSPFv3](#)]. When set to "disabled" the Default Topology is described using Extended LSAs as specified in this memo. This parameter is set to "enabled" by default.

Authors' Addresses

Sina Mirtorabi
Force10 Networks
350 Holger Way
San Jose, CA 95134, USA

E-mail: sina@force10networks.com

Abhay Roy
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134, USA

E-mail: akr@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).