

Network Working Group
Requests for Comments: 2740
Category: Standards Track

R. Coltun
Siara Systems
D. Ferguson
Juniper Networks
J. Moy
Sycamore Networks
December 1999

OSPF for IPv6

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

This document describes the modifications to OSPF to support version 6 of the Internet Protocol (IPv6). The fundamental mechanisms of OSPF (flooding, DR election, area support, SPF calculations, etc.) remain unchanged. However, some changes have been necessary, either due to changes in protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6.

Changes between OSPF for IPv4 and this document include the following. Addressing semantics have been removed from OSPF packets and the basic LSAs. New LSAs have been created to carry IPv6 addresses and prefixes. OSPF now runs on a per-link basis, instead of on a per-IP-subnet basis. Flooding scope for LSAs has been generalized. Authentication has been removed from the OSPF protocol itself, instead relying on IPv6's Authentication Header and Encapsulating Security Payload.

Most packets in OSPF for IPv6 are almost as compact as those in OSPF for IPv4, even with the larger IPv6 addresses. Most field- and packet-size limitations present in OSPF for IPv4 have been relaxed. In addition, option handling has been made more flexible.

All of OSPF for IPv4's optional capabilities, including on-demand circuit support, NSSA areas, and the multicast extensions to OSPF (MOSPF) are also supported in OSPF for IPv6.

Table of Contents

| | | |
|---------|--|----|
| 1 | Introduction | 4 |
| 1.1 | Terminology | 4 |
| 2 | Differences from OSPF for IPv4 | 4 |
| 2.1 | Protocol processing per-link, not per-subnet | 5 |
| 2.2 | Removal of addressing semantics | 5 |
| 2.3 | Addition of Flooding scope | 5 |
| 2.4 | Explicit support for multiple instances per link | 6 |
| 2.5 | Use of link-local addresses | 6 |
| 2.6 | Authentication changes | 7 |
| 2.7 | Packet format changes | 7 |
| 2.8 | LSA format changes | 8 |
| 2.9 | Handling unknown LSA types | 10 |
| 2.10 | Stub area support | 10 |
| 2.11 | Identifying neighbors by Router ID | 11 |
| 3 | Implementation details | 11 |
| 3.1 | Protocol data structures | 12 |
| 3.1.1 | The Area Data structure | 13 |
| 3.1.2 | The Interface Data structure | 13 |
| 3.1.3 | The Neighbor Data Structure | 14 |
| 3.2 | Protocol Packet Processing | 15 |
| 3.2.1 | Sending protocol packets | 15 |
| 3.2.1.1 | Sending Hello packets | 16 |
| 3.2.1.2 | Sending Database Description Packets | 17 |
| 3.2.2 | Receiving protocol packets | 17 |
| 3.2.2.1 | Receiving Hello Packets | 19 |
| 3.3 | The Routing table Structure | 19 |
| 3.3.1 | Routing table lookup | 20 |
| 3.4 | Link State Advertisements | 20 |
| 3.4.1 | The LSA Header | 21 |
| 3.4.2 | The link-state database | 22 |
| 3.4.3 | Originating LSAs | 22 |
| 3.4.3.1 | Router-LSAs | 25 |
| 3.4.3.2 | Network-LSAs | 27 |
| 3.4.3.3 | Inter-Area-Prefix-LSAs | 28 |
| 3.4.3.4 | Inter-Area-Router-LSAs | 29 |
| 3.4.3.5 | AS-external-LSAs | 29 |
| 3.4.3.6 | Link-LSAs | 31 |
| 3.4.3.7 | Intra-Area-Prefix-LSAs | 32 |
| 3.5 | Flooding | 35 |
| 3.5.1 | Receiving Link State Update packets | 36 |
| 3.5.2 | Sending Link State Update packets | 36 |
| 3.5.3 | Installing LSAs in the database | 38 |

| | | |
|---------|--|----|
| 3.6 | Definition of self-originated LSAs | 39 |
| 3.7 | Virtual links | 39 |
| 3.8 | Routing table calculation | 39 |
| 3.8.1 | Calculating the shortest path tree for an area | 40 |
| 3.8.1.1 | The next hop calculation | 41 |
| 3.8.2 | Calculating the inter-area routes | 42 |
| 3.8.3 | Examining transit areas' summary-LSAs | 42 |
| 3.8.4 | Calculating AS external routes | 42 |
| 3.9 | Multiple interfaces to a single link | 43 |
| | References | 44 |
| A | OSPF data formats | 46 |
| A.1 | Encapsulation of OSPF packets | 46 |
| A.2 | The Options field | 47 |
| A.3 | OSPF Packet Formats | 48 |
| A.3.1 | The OSPF packet header | 49 |
| A.3.2 | The Hello packet | 50 |
| A.3.3 | The Database Description packet | 52 |
| A.3.4 | The Link State Request packet | 54 |
| A.3.5 | The Link State Update packet | 55 |
| A.3.6 | The Link State Acknowledgment packet | 56 |
| A.4 | LSA formats | 57 |
| A.4.1 | IPv6 Prefix Representation | 58 |
| A.4.1.1 | Prefix Options | 58 |
| A.4.2 | The LSA header | 59 |
| A.4.2.1 | LS type | 60 |
| A.4.3 | Router-LSAs | 61 |
| A.4.4 | Network-LSAs | 64 |
| A.4.5 | Inter-Area-Prefix-LSAs | 65 |
| A.4.6 | Inter-Area-Router-LSAs | 66 |
| A.4.7 | AS-external-LSAs | 67 |
| A.4.8 | Link-LSAs | 69 |
| A.4.9 | Intra-Area-Prefix-LSAs | 71 |
| B | Architectural Constants | 73 |
| C | Configurable Constants | 73 |
| C.1 | Global parameters | 73 |
| C.2 | Area parameters | 74 |
| C.3 | Router interface parameters | 75 |
| C.4 | Virtual link parameters | 77 |
| C.5 | NBMA network parameters | 77 |
| C.6 | Point-to-MultiPoint network parameters | 78 |
| C.7 | Host route parameters | 78 |
| | Security Considerations | 79 |
| | Authors' Addresses | 79 |
| | Full Copyright Statement | 80 |

1. Introduction

This document describes the modifications to OSPF to support version 6 of the Internet Protocol (IPv6). The fundamental mechanisms of OSPF (flooding, DR election, area support, SPF calculations, etc.) remain unchanged. However, some changes have been necessary, either due to changes in protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6.

This document is organized as follows. Section 2 describes the differences between OSPF for IPv4 and OSPF for IPv6 in detail. Section 3 provides implementation details for the changes. Appendix A gives the OSPF for IPv6 packet and LSA formats. Appendix B lists the OSPF architectural constants. Appendix C describes configuration parameters.

1.1. Terminology

This document attempts to use terms from both the OSPF for IPv4 specification ([Ref1]) and the IPv6 protocol specifications ([Ref14]). This has produced a mixed result. Most of the terms used both by OSPF and IPv6 have roughly the same meaning (e.g., interfaces). However, there are a few conflicts. IPv6 uses "link" similarly to IPv4 OSPF's "subnet" or "network". In this case, we have chosen to use IPv6's "link" terminology. "Link" replaces OSPF's "subnet" and "network" in most places in this document, although OSPF's Network-LSA remains unchanged (and possibly unfortunately, a new Link-LSA has also been created).

The names of some of the OSPF LSAs have also changed. See Section 2.8 for details.

2. Differences from OSPF for IPv4

Most of the algorithms from OSPF for IPv4 [Ref1] have preserved in OSPF for IPv6. However, some changes have been necessary, either due to changes in protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6.

The following subsections describe the differences between this document and [Ref1].

2.1. Protocol processing per-link, not per-subnet

IPv6 uses the term "link" to indicate "a communication facility or medium over which nodes can communicate at the link layer" ([Ref14]). "Interfaces" connect to links. Multiple IP subnets can be assigned to a single link, and two nodes can talk directly over a single link, even if they do not share a common IP subnet (IPv6 prefix).

For this reason, OSPF for IPv6 runs per-link instead of the IPv4 behavior of per-IP-subnet. The terms "network" and "subnet" used in the IPv4 OSPF specification ([Ref1]) should generally be relaced by link. Likewise, an OSPF interface now connects to a link instead of an IP subnet, etc.

This change affects the receiving of OSPF protocol packets, and the contents of Hello Packets and Network-LSAs.

2.2. Removal of addressing semantics

In OSPF for IPv6, addressing semantics have been removed from the OSPF protocol packets and the main LSA types, leaving a network-protocol-independent core. In particular:

- o IPv6 Addresses are not present in OSPF packets, except in LSA payloads carried by the Link State Update Packets. See Section 2.7 for details.
- o Router-LSAs and Network-LSAs no longer contain network addresses, but simply express topology information. See Section 2.8 for details.
- o OSPF Router IDs, Area IDs and LSA Link State IDs remain at the IPv4 size of 32-bits. They can no longer be assigned as (IPv6) addresses.
- o Neighboring routers are now always identified by Router ID, where previously they had been identified by IP address on broadcast and NBMA "networks".

2.3. Addition of Flooding scope

Flooding scope for LSAs has been generalized and is now explicitly coded in the LSA's LS type field. There are now three separate flooding scopes for LSAs:

- o Link-local scope. LSA is flooded only on the local link, and no further. Used for the new Link-LSA (see Section A.4.8).
- o Area scope. LSA is flooded throughout a single OSPF area only. Used for Router-LSAs, Network-LSAs, Inter-Area-Prefix-LSAs, Inter-Area-Router-LSAs and Intra-Area-Prefix-LSAs.
- o AS scope. LSA is flooded throughout the routing domain. Used for AS-external-LSAs.

2.4. Explicit support for multiple instances per link

OSPF now supports the ability to run multiple OSPF protocol instances on a single link. For example, this may be required on a NAP segment shared between several providers -- providers may be running separate OSPF routing domains that want to remain separate even though they have one or more physical network segments (i.e., links) in common. In OSPF for IPv4 this was supported in a haphazard fashion using the authentication fields in the OSPF for IPv4 header.

Another use for running multiple OSPF instances is if you want, for one reason or another, to have a single link belong to two or more OSPF areas.

Support for multiple protocol instances on a link is accomplished via an "Instance ID" contained in the OSPF packet header and OSPF interface structures. Instance ID solely affects the reception of OSPF packets.

2.5. Use of link-local addresses

IPv6 link-local addresses are for use on a single link, for purposes of neighbor discovery, auto-configuration, etc. IPv6 routers do not forward IPv6 datagrams having link-local source addresses [Ref15]. Link-local unicast addresses are assigned from the IPv6 address range FF80/10.

OSPF for IPv6 assumes that each router has been assigned link-local unicast addresses on each of the router's attached physical segments. On all OSPF interfaces except virtual links, OSPF packets are sent using the interface's associated link-local unicast address as source. A router learns the link-local addresses of all other routers attached to its links, and uses these addresses as next hop information during packet forwarding.

On virtual links, global scope or site-local IP addresses must be used as the source for OSPF protocol packets.

Link-local addresses appear in OSPF Link-LSAs (see Section 3.4.3.6). However, link-local addresses are not allowed in other OSPF LSA types. In particular, link-local addresses must not be advertised in inter-area-prefix-LSAs (Section 3.4.3.3), AS-external-LSAs (Section 3.4.3.5) or intra-area-prefix-LSAs (Section 3.4.3.7).

2.6. Authentication changes

In OSPF for IPv6, authentication has been removed from OSPF itself. The "AuType" and "Authentication" fields have been removed from the OSPF packet header, and all authentication related fields have been removed from the OSPF area and interface structures.

When running over IPv6, OSPF relies on the IP Authentication Header (see [Ref19]) and the IP Encapsulating Security Payload (see [Ref20]) to ensure integrity and authentication/confidentiality of routing exchanges.

Protection of OSPF packet exchanges against accidental data corruption is provided by the standard IPv6 16-bit one's complement checksum, covering the entire OSPF packet and prepended IPv6 pseudo-header (see Section A.3.1).

2.7. Packet format changes

OSPF for IPv6 runs directly over IPv6. Aside from this, all addressing semantics have been removed from the OSPF packet headers, making it essentially "network-protocol-independent". All addressing information is now contained in the various LSA types only.

In detail, changes in OSPF packet format consist of the following:

- o The OSPF version number has been increased from 2 to 3.
- o The Options field in Hello Packets and Database description Packet has been expanded to 24-bits.
- o The Authentication and AuType fields have been removed from the OSPF packet header (see Section 2.6).
- o The Hello packet now contains no address information at all, and includes an Interface ID which the originating router has assigned to uniquely identify (among its own interfaces) its interface to the link. This Interface ID becomes the Netowrk-LSA's Link State ID, should the router become Designated-Router on the link.

- o Two option bits, the "R-bit" and the "V6-bit", have been added to the Options field for processing Router-LSAs during the SPF calculation (see Section A.2). If the "R-bit" is clear an OSPF speaker can participate in OSPF topology distribution without being used to forward transit traffic; this can be used in multi-homed hosts that want to participate in the routing protocol. The V6-bit specializes the R-bit; if the V6-bit is clear an OSPF speaker can participate in OSPF topology distribution without being used to forward IPv6 datagrams. If the R-bit is set and the V6-bit is clear, IPv6 datagrams are not forwarded but diagrams belonging to another protocol family may be forwarded.
- o The OSPF packet header now includes an "Instance ID" which allows multiple OSPF protocol instances to be run on a single link (see Section 2.4).

2.8. LSA format changes

All addressing semantics have been removed from the LSA header, and from Router-LSAs and Network-LSAs. These two LSAs now describe the routing domain's topology in a network-protocol-independent manner. New LSAs have been added to distribute IPv6 address information, and data required for next hop resolution. The names of some of IPv4's LSAs have been changed to be more consistent with each other.

In detail, changes in LSA format consist of the following:

- o The Options field has been removed from the LSA header, expanded to 24 bits, and moved into the body of Router-LSAs, Network-LSAs, Inter-Area-Router-LSAs and Link-LSAs. See Section A.2 for details.
- o The LSA Type field has been expanded (into the former Options space) to 16 bits, with the upper three bits encoding flooding scope and the handling of unknown LSA types (see Section 2.9).
- o Addresses in LSAs are now expressed as [prefix, prefix length] instead of [address, mask] (see Section A.4.1). The default route is expressed as a prefix with length 0.
- o The Router and Network LSAs now have no address information, and are network-protocol-independent.
- o Router interface information may be spread across multiple Router LSAs. Receivers must concatenate all the Router-LSAs originated by a given router when running the SPF calculation.

- o A new LSA called the Link-LSA has been introduced. The LSAs have local-link flooding scope; they are never flooded beyond the link that they are associated with. Link-LSAs have three purposes: 1) they provide the router's link-local address to all other routers attached to the link, 2) they inform other routers attached to the link of a list of IPv6 prefixes to associate with the link and 3) they allow the router to assert a collection of Options bits to associate with the Network-LSA that will be originated for the link. See Section A.4.8 for details.

In IPv4, the router-LSA carries a router's IPv4 interface addresses, the IPv4 equivalent of link-local addresses. These are only used when calculating next hops during the OSPF routing calculation (see Section 16.1.1 of [Ref1]), so they do not need to be flooded past the local link; hence using link-LSAs to distribute these addresses is more efficient. Note that link-local addresses cannot be learned through the reception of Hellos in all cases: on NBMA links next hop routers do not necessarily exchange hellos, but rather learn of each other's existence by way of the Designated Router.

- o The Options field in the Network LSA is set to the logical OR of the Options that each router on the link advertises in its Link-LSA.
- o Type-3 summary-LSAs have been renamed "Inter-Area-Prefix-LSAs". Type-4 summary LSAs have been renamed "Inter-Area-Router-LSAs".
- o The Link State ID in Inter-Area-Prefix-LSAs, Inter-Area-Router-LSAs and AS-external-LSAs has lost its addressing semantics, and now serves solely to identify individual pieces of the Link State Database. All addresses or Router IDs that were formerly expressed by the Link State ID are now carried in the LSA bodies.
- o Network-LSAs and Link-LSAs are the only LSAs whose Link State ID carries additional meaning. For these LSAs, the Link State ID is always the Interface ID of the originating router on the link being described. For this reason, Network-LSAs and Link-LSAs are now the only LSAs whose size cannot be limited: a Network-LSA must list all routers connected to the link, and a Link-LSA must list all of a router's addresses on the link.
- o A new LSA called the Intra-Area-Prefix-LSA has been introduced. This LSA carries all IPv6 prefix information that in IPv4 is included in Router-LSAs and Network-LSAs. See Section A.4.9 for details.

- o Inclusion of a forwarding address in AS-external-LSAs is now optional, as is the inclusion of an external route tag (see [Ref5]). In addition, AS-external-LSAs can now reference another LSA, for inclusion of additional route attributes that are outside the scope of the OSPF protocol itself. For example, this can be used to attach BGP path attributes to external routes as proposed in [Ref10].

2.9. Handling unknown LSA types

Handling of unknown LSA types has been made more flexible so that, based on LS type, unknown LSA types are either treated as having link-local flooding scope, or are stored and flooded as if they were understood (desirable for things like the proposed External-Attributes-LSA in [Ref10]). This behavior is explicitly coded in the LSA Handling bit of the link state header's LS type field (see Section A.4.2.1).

The IPv4 OSPF behavior of simply discarding unknown types is unsupported due to the desire to mix router capabilities on a single link. Discarding unknown types causes problems when the Designated Router supports fewer options than the other routers on the link.

2.10. Stub area support

In OSPF for IPv4, stub areas were designed to minimize link-state database and routing table sizes for the areas' internal routers. This allows routers with minimal resources to participate in even very large OSPF routing domains.

In OSPF for IPv6, the concept of stub areas is retained. In IPv6, of the mandatory LSA types, stub areas carry only router-LSAs, network-LSAs, Inter-Area-Prefix-LSAs, Link-LSAs, and Intra-Area-Prefix-LSAs. This is the IPv6 equivalent of the LSA types carried in IPv4 stub areas: router-LSAs, network-LSAs and type 3 summary-LSAs.

However, unlike in IPv4, IPv6 allows LSAs with unrecognized LS types to be labeled "Store and flood the LSA, as if type understood" (see the U-bit in Section A.4.2.1). Uncontrolled introduction of such LSAs could cause a stub area's link-state database to grow larger than its component routers' capacities.

To guard against this, the following rule regarding stub areas has been established: an LSA whose LS type is unrecognized may only be flooded into/throughout a stub area if both a) the LSA has area or link-local flooding scope and b) the LSA has U-bit set to 0. See Section 3.5 for details.

2.11. Identifying neighbors by Router ID

In OSPF for IPv6, neighboring routers on a given link are always identified by their OSPF Router ID. This contrasts with the IPv4 behavior where neighbors on point-to-point networks and virtual links are identified by their Router IDs, and neighbors on broadcast, NBMA and Point-to-MultiPoint links are identified by their IPv4 interface addresses.

This change affects the reception of OSPF packets (see Section 8.2 of [Ref1]), the lookup of neighbors (Section 10 of [Ref1]) and the reception of Hello Packets (Section 10.5 of [Ref1]).

The Router ID of 0.0.0.0 is reserved, and should not be used.

3. Implementation details

When going from IPv4 to IPv6, the basic OSPF mechanisms remain unchanged from those documented in [Ref1]. These mechanisms are briefly outlined in Section 4 of [Ref1]. Both IPv6 and IPv4 have a link-state database composed of LSAs and synchronized between adjacent routers. Initial synchronization is performed through the Database Exchange process, through the exchange of Database Description, Link State Request and Link State Update packets. Thereafter database synchronization is maintained via flooding, utilizing Link State Update and Link State Acknowledgment packets. Both IPv6 and IPv4 use OSPF Hello Packets to discover and maintain neighbor relationships, and to elect Designated Routers and Backup Designated Routers on broadcast and NBMA links. The decision as to which neighbor relationships become adjacencies, along with the basic ideas behind inter-area routing, importing external information in AS-external-LSAs and the various routing calculations are also the same.

In particular, the following IPv4 OSPF functionality described in [Ref1] remains completely unchanged for IPv6:

- o Both IPv4 and IPv6 use OSPF packet types described in Section 4.3 of [Ref1], namely: Hello, Database Description, Link State Request, Link State Update and Link State Acknowledgment packets. While in some cases (e.g., Hello packets) their format has changed somewhat, the functions of the various packet types remains the same.
- o The system requirements for an OSPF implementation remain unchanged, although OSPF for IPv6 requires an IPv6 protocol stack (from the network layer on down) since it runs directly over the IPv6 network layer.

- o The discovery and maintenance of neighbor relationships, and the selection and establishment of adjacencies remain the same. This includes election of the Designated Router and Backup Designated Router on broadcast and NBMA links. These mechanisms are described in Sections 7, 7.1, 7.2, 7.3, 7.4 and 7.5 of [Ref1].
- o The link types (or equivalently, interface types) supported by OSPF remain unchanged, namely: point-to-point, broadcast, NBMA, Point-to-MultiPoint and virtual links.
- o The interface state machine, including the list of OSPF interface states and events, and the Designated Router and Backup Designated Router election algorithm, remain unchanged. These are described in Sections 9.1, 9.2, 9.3 and 9.4 of [Ref1].
- o The neighbor state machine, including the list of OSPF neighbor states and events, remain unchanged. These are described in Sections 10.1, 10.2, 10.3 and 10.4 of [Ref1].
- o Aging of the link-state database, as well as flushing LSAs from the routing domain through the premature aging process, remains unchanged from the description in Sections 14 and 14.1 of [Ref1].

However, some OSPF protocol mechanisms have changed, as outlined in Section 2 above. These changes are explained in detail in the following subsections, making references to the appropriate sections of [Ref1].

The following subsections provide a recipe for turning an IPv4 OSPF implementation into an IPv6 OSPF implementation.

3.1. Protocol data structures

The major OSPF data structures are the same for both IPv4 and IPv6: areas, interfaces, neighbors, the link-state database and the routing table. The top-level data structures for IPv6 remain those listed in Section 5 of [Ref1], with the following modifications:

- o All LSAs with known LS type and AS flooding scope appear in the top-level data structure, instead of belonging to a specific area or link. AS-external-LSAs are the only LSAs defined by this specification which have AS flooding scope. LSAs with unknown LS type, U-bit set to 1 (flood even when unrecognized) and AS flooding scope also appear in the top-level data structure.

3.1.1. The Area Data structure

The IPv6 area data structure contains all elements defined for IPv4 areas in Section 6 of [Ref1]. In addition, all LSAs of known type which have area flooding scope are contained in the IPv6 area data structure. This always includes the following LSA types: router-LSAs, network-LSAs, inter-area-prefix-LSAs, inter-area-router-LSAs and intra-area-prefix-LSAs. LSAs with unknown LS type, U-bit set to 1 (flood even when unrecognized) and area scope also appear in the area data structure. IPv6 routers implementing MOSPF add group-membership-LSAs to the area data structure. Type-7-LSAs belong to an NSSA area's data structure.

3.1.2. The Interface Data structure

In OSPF for IPv6, an interface connects a router to a link. The IPv6 interface structure modifies the IPv4 interface structure (as defined in Section 9 of [Ref1]) as follows:

Interface ID

Every interface is assigned an Interface ID, which uniquely identifies the interface with the router. For example, some implementations may be able to use the MIB-II IfIndex ([Ref3]) as Interface ID. The Interface ID appears in Hello packets sent out the interface, the link-local-LSA originated by router for the attached link, and the router-LSA originated by the router-LSA for the associated area. It will also serve as the Link State ID for the network-LSA that the router will originate for the link if the router is elected Designated Router.

Instance ID

Every interface is assigned an Instance ID. This should default to 0, and is only necessary to assign differently on those links that will contain multiple separate communities of OSPF Routers. For example, suppose that there are two communities of routers on a given ethernet segment that you wish to keep separate.

The first community is given an Instance ID of 0, by assigning 0 as the Instance ID of all its routers' interfaces to the ethernet. An Instance ID of 1 is assigned to the other routers' interfaces to the ethernet. The OSPF transmit and receive processing (see Section 3.2) will then keep the two communities separate.

List of LSAs with link-local scope

All LSAs with link-local scope and which were originated/flooded on the link belong to the interface structure which connects to the link. This includes the collection of the link's link-LSAs.

List of LSAs with unknown LS type

All LSAs with unknown LS type and U-bit set to 0 (if unrecognized, treat the LSA as if it had link-local flooding scope) are kept in the data structure for the interface that received the LSA.

IP interface address

For IPv6, the IPv6 address appearing in the source of OSPF packets sent out the interface is almost always a link-local address. The one exception is for virtual links, which must use one of the router's own site-local or global IPv6 addresses as IP interface address.

List of link prefixes

A list of IPv6 prefixes can be configured for the attached link. These will be advertised by the router in link-LSAs, so that they can be advertised by the link's Designated Router in intra-area-prefix-LSAs.

In OSPF for IPv6, each router interface has a single metric, representing the cost of sending packets out the interface. In addition, OSPF for IPv6 relies on the IP Authentication Header (see [Ref19]) and the IP Encapsulating Security Payload (see [Ref20]) to ensure integrity and authentication/confidentiality of routing exchanges. For that reason, AuType and Authentication key are not associated with IPv6 OSPF interfaces.

Interface states, events, and the interface state machine remain unchanged from IPv4, and are documented in Sections 9.1, 9.2 and 9.3 of [Ref1] respectively. The Designated Router and Backup Designated Router election algorithm also remains unchanged from the IPv4 election in Section 9.4 of [Ref1].

3.1.3. The Neighbor Data Structure

The neighbor structure performs the same function in both IPv6 and IPv4. Namely, it collects all information required to form an adjacency between two routers, if an adjacency becomes necessary. Each neighbor structure is bound to a single OSPF interface. The differences between the IPv6 neighbor structure and the neighbor structure defined for IPv4 in Section 10 of [Ref1] are:

Neighbor's Interface ID

The Interface ID that the neighbor advertises in its Hello Packets must be recorded in the neighbor structure. The router will include the neighbor's Interface ID in the router's router-LSA when either a) advertising a point-to-point link to the neighbor or b) advertising a link to a network where the neighbor has become Designated Router.

Neighbor IP address

Except on virtual links, the neighbor's IP address will be an IPv6 link-local address.

Neighbor's Designated Router

The neighbor's choice of Designated Router is now encoded as a Router ID, instead of as an IP address.

Neighbor's Backup Designated Router

The neighbor's choice of Designated Router is now encoded as a Router ID, instead of as an IP address.

Neighbor states, events, and the neighbor state machine remain unchanged from IPv4, and are documented in Sections 10.1, 10.2 and 10.3 of [Ref1] respectively. The decision as to which adjacencies to form also remains unchanged from the IPv4 logic documented in Section 10.4 of [Ref1].

3.2. Protocol Packet Processing

OSPF for IPv6 runs directly over IPv6's network layer. As such, it is encapsulated in one or more IPv6 headers, with the Next Header field of the immediately encapsulating IPv6 header set to the value 89.

As for IPv4, in IPv6 OSPF routing protocol packets are sent along adjacencies only (with the exception of Hello packets, which are used to discover the adjacencies). OSPF packet types and functions are the same in both IPv4 and IPv6, encoded by the

Type field of the standard OSPF packet header.

3.2.1. Sending protocol packets

When an IPv6 router sends an OSPF routing protocol packet, it fills in the fields of the standard OSPF for IPv6 packet header (see Section A.3.1) as follows:

Version

Set to 3, the version number of the protocol as documented in this specification.

Type

The type of OSPF packet, such as Link state Update or Hello Packet.

Packet length

The length of the entire OSPF packet in bytes, including the standard OSPF packet header.

Router ID

The identity of the router itself (who is originating the packet).

Area ID

The OSPF area that the packet is being sent into.

Instance ID

The OSPF Instance ID associated with the interface that the packet is being sent out of.

Checksum

The standard IPv6 16-bit one's complement checksum, covering the entire OSPF packet and prepended IPv6 pseudo-header (see Section A.3.1).

Selection of OSPF routing protocol packets' IPv6 source and destination addresses is performed identically to the IPv4 logic in Section 8.1 of [Ref1]. The IPv6 destination address is chosen from among the addresses AllSPFRouters, AllDRouters and the Neighbor IP address associated with the other end of the adjacency (which in IPv6, for all links except virtual links, is an IPv6 link-local address).

The sending of Link State Request Packets and Link State Acknowledgment Packets remains unchanged from the IPv4 procedures documented in Sections 10.9 and 13.5 of [Ref1] respectively. Sending Hello Packets is documented in Section 3.2.1.1, and the sending of Database Description Packets in Section 3.2.1.2. The sending of Link State Update Packets is documented in Section 3.5.2.

3.2.1.1. Sending Hello packets

IPv6 changes the way OSPF Hello packets are sent in the following ways (compare to Section 9.5 of [Ref1]):

- o Before the Hello Packet is sent out an interface, the interface's Interface ID must be copied into the Hello Packet.
- o The Hello Packet no longer contains an IP network mask, as OSPF for IPv6 runs per-link instead of per-subnet.
- o The choice of Designated Router and Backup Designated Router are now indicated within Hellos by their Router IDs, instead of by their IP interface addresses. Advertising the Designated Router (or Backup Designated Router) as 0.0.0.0 indicates that the Designated Router (or Backup Designated Router) has not yet been chosen.

- o The Options field within Hello packets has moved around, getting larger in the process. More options bits are now possible. Those that must be set correctly in Hello packets are: The E-bit is set if and only if the interface attaches to a non-stub area, the N-bit is set if and only if the interface attaches to an NSSA area (see [Ref9]), and the DC-bit is set if and only if the router wishes to suppress the sending of future Hellos over the interface (see [Ref11]). Unrecognized bits in the Hello Packet's Options field should be cleared.

Sending Hello packets on NBMA networks proceeds for IPv6 in exactly the same way as for IPv4, as documented in Section 9.5.1 of [Ref1].

3.2.1.2. Sending Database Description Packets

The sending of Database Description packets differs from Section 10.8 of [Ref1] in the following ways:

- o The Options field within Database Description packets has moved around, getting larger in the process. More options bits are now possible. Those that must be set correctly in Database Description packets are: The MC-bit is set if and only if the router is forwarding multicast datagrams according to the MOSPF specification in [Ref7], and the DC-bit is set if and only if the router wishes to suppress the sending of Hellos over the interface (see [Ref11]). Unrecognized bits in the Database Description Packet's Options field should be cleared.

3.2.2. Receiving protocol packets

Whenever an OSPF protocol packet is received by the router it is marked with the interface it was received on. For routers that have virtual links configured, it may not be immediately obvious which interface to associate the packet with. For example, consider the Router RT11 depicted in Figure 6 of [Ref1]. If RT11 receives an OSPF protocol packet on its interface to Network N8, it may want to associate the packet with the interface to Area 2, or with the virtual link to Router RT10 (which is part of the backbone). In the following, we assume that the packet is initially associated with the non-virtual link.

In order for the packet to be passed to OSPF for processing, the following tests must be performed on the encapsulating IPv6 headers:

- o The packet's IP destination address must be one of the IPv6 unicast addresses associated with the receiving interface (this includes link-local addresses), or one of the IP multicast addresses AllSPFRouters or AllDRouters.

- o The Next Header field of the immediately encapsulating IPv6 header must specify the OSPF protocol (89).
- o Any encapsulating IP Authentication Headers (see [Ref19]) and the IP Encapsulating Security Payloads (see [Ref20]) must be processed and/or verified to ensure integrity and authentication/confidentiality of OSPF routing exchanges.
- o Locally originated packets should not be passed on to OSPF. That is, the source IPv6 address should be examined to make sure this is not a multicast packet that the router itself generated.

After processing the encapsulating IPv6 headers, the OSPF packet header is processed. The fields specified in the header must match those configured for the receiving interface. If they do not, the packet should be discarded:

- o The version number field must specify protocol version 3.
- o The standard IPv6 16-bit one's complement checksum, covering the entire OSPF packet and prepended IPv6 pseudo-header, must be verified (see Section A.3.1).
- o The Area ID found in the OSPF header must be verified. If both of the following cases fail, the packet should be discarded. The Area ID specified in the header must either:
 - (1) Match the Area ID of the receiving interface. In this case, unlike for IPv4, the IPv6 source address is not restricted to lie on the same IP subnet as the receiving interface. IPv6 OSPF runs per-link, instead of per-IP-subnet.
 - (2) Indicate the backbone. In this case, the packet has been sent over a virtual link. The receiving router must be an area border router, and the Router ID specified in the packet (the source router) must be the other end of a configured virtual link. The receiving interface must also attach to the virtual link's configured Transit area. If all of these checks succeed, the packet is accepted and is from now on associated with the virtual link (and the backbone area).
- o The Instance ID specified in the OSPF header must match the receiving interface's Instance ID.

- o Packets whose IP destination is AllDRouters should only be accepted if the state of the receiving interface is DR or Backup (see Section 9.1).

After header processing, the packet is further processed according to its OSPF packet type. OSPF packet types and functions are the same for both IPv4 and IPv6.

If the packet type is Hello, it should then be further processed by the Hello Protocol. All other packet types are sent/received only on adjacencies. This means that the packet must have been sent by one of the router's active neighbors. The neighbor is identified by the Router ID appearing in the received packet's OSPF header. Packets not matching any active neighbor are discarded.

The receive processing of Database Description Packets, Link State Request Packets and Link State Acknowledgment Packets remains unchanged from the IPv4 procedures documented in Sections 10.6, 10.7 and 13.7 of [Ref1] respectively. The receiving of Hello Packets is documented in Section 3.2.2.1, and the receiving of Link State Update Packets is documented in Section 3.5.1.

3.2.2.1. Receiving Hello Packets

The receive processing of Hello Packets differs from Section 10.5 of [Ref1] in the following ways:

- o On all link types (e.g., broadcast, NBMA, point-to-point, etc), neighbors are identified solely by their OSPF Router ID. For all link types except virtual links, the Neighbor IP address is set to the IPv6 source address in the IPv6 header of the received OSPF Hello packet.
- o There is no longer a Network Mask field in the Hello Packet.
- o The neighbor's choice of Designated Router and Backup Designated Router is now encoded as an OSPF Router ID instead of an IP interface address.

3.3. The Routing table Structure

The routing table used by OSPF for IPv4 is defined in Section 11 of [Ref1]. For IPv6 there are analogous routing table entries: there are routing table entries for IPv6 address prefixes, and also for AS boundary routers. The latter routing table entries are only used to hold intermediate results during the routing table build process (see Section 3.8).

Also, to hold the intermediate results during the shortest-path calculation for each area, there is a separate routing table for each area holding the following entries:

- o An entry for each router in the area. Routers are identified by their OSPF router ID. These routing table entries hold the set of shortest paths through a given area to a given router, which in turn allows calculation of paths to the IPv6 prefixes advertised by that router in Intra-area-prefix-LSAs. If the router is also an area-border router, these entries are also used to calculate paths for inter-area address prefixes. If in addition the router is the other endpoint of a virtual link, the routing table entry describes the cost and viability of the virtual link.
- o An entry for each transit link in the area. Transit links have associated network-LSAs. Both the transit link and the network-LSA are identified by a combination of the Designated Router's Interface ID on the link and the Designated Router's OSPF Router ID. These routing table entries allow later calculation of paths to IP prefixes advertised for the transit link in intra-area-prefix-LSAs.

The fields in the IPv4 OSPF routing table (see Section 11 of [Ref1]) remain valid for IPv6: Optional capabilities (routers only), path type, cost, type 2 cost, link state origin, and for each of the equal cost paths to the destination, the next hop and advertising router.

For IPv6, the link-state origin field in the routing table entry is the router-LSA or network-LSA that has directly or indirectly produced the routing table entry. For example, if the routing table entry describes a route to an IPv6 prefix, the link state origin is the router-LSA or network-LSA that is listed in the body of the intra-area-prefix-LSA that has produced the route (see Section A.4.9).

3.3.1. Routing table lookup

Routing table lookup (i.e., determining the best matching routing table entry during IP forwarding) is the same for IPv6 as for IPv4.

3.4. Link State Advertisements

For IPv6, the OSPF LSA header has changed slightly, with the LS type field expanding and the Options field being moved into the body of appropriate LSAs. Also, the formats of some LSAs have changed somewhat (namely router-LSAs, network-LSAs and AS-external-LSAs), while the names of other LSAs have been changed (type 3 and 4 summary-LSAs are now inter-area-prefix-LSAs and inter-area-router-

LSAs respectively) and additional LSAs have been added (Link-LSAs and Intra-Area-Prefix-LSAs). Type of Service (TOS) has been removed from the OSPFv2 specification [Ref1], and is not encoded within OSPF for IPv6's LSAs.

These changes will be described in detail in the following subsections.

3.4.1. The LSA Header

In both IPv4 and IPv6, all OSPF LSAs begin with a standard 20 byte LSA header. However, the contents of this 20 byte header have changed in IPv6. The LS age, Advertising Router, LS Sequence Number, LS checksum and length fields within the LSA header remain unchanged, as documented in Sections 12.1.1, 12.1.5, 12.1.6, 12.1.7 and A.4.1 of [Ref1] respectively. However, the following fields have changed for IPv6:

Options

The Options field has been removed from the standard 20 byte LSA header, and into the body of router-LSAs, network-LSAs, inter-area-router-LSAs and link-LSAs. The size of the Options field has increased from 8 to 24 bits, and some of the bit definitions have changed (see Section A.2). In addition a separate PrefixOptions field, 8 bits in length, is attached to each prefix advertised within the body of an LSA.

LS type

The size of the LS type field has increased from 8 to 16 bits, with the top two bits encoding flooding scope and the next bit encoding the handling of unknown LS types. See Section A.4.2.1 for the current coding of the LS type field.

Link State ID

Link State ID remains at 32 bits in length, but except for network-LSAs and link-LSAs, Link State ID has shed any addressing semantics. For example, an IPv6 router originating multiple AS-external-LSAs could start by assigning the first a Link State ID of 0.0.0.1, the second a Link State ID of 0.0.0.2, and so on. Instead of the IPv4 behavior of encoding the network number within the AS-external-LSA's Link State ID, the IPv6 Link State ID simply serves as a way to differentiate multiple LSAs originated by the same router.

For network-LSAs, the Link State ID is set to the Designated Router's Interface ID on the link. When a router originates a Link-LSA for a given link, its Link State ID is set equal to the router's Interface ID on the link.

3.4.2. The link-state database

In IPv6, as in IPv4, individual LSAs are identified by a combination of their LS type, Link State ID and Advertising Router fields. Given two instances of an LSA, the most recent instance is determined by examining the LSAs' LS Sequence Number, using LS checksum and LS age as tiebreakers (see Section 13.1 of [Ref1]).

In IPv6, the link-state database is split across three separate data structures. LSAs with AS flooding scope are contained within the top-level OSPF data structure (see Section 3.1) as long as either their LS type is known or their U-bit is 1 (flood even when unrecognized); this includes the AS-external-LSAs. LSAs with area flooding scope are contained within the appropriate area structure (see Section 3.1.1) as long as either their LS type is known or their U-bit is 1 (flood even when unrecognized); this includes router-LSAs, network-LSAs, inter-area-prefix-LSAs, inter-area-router-LSAs, and intra-area-prefix-LSAs. LSAs with unknown LS type and U-bit set to 0 and/or link-local flooding scope are contained within the appropriate interface structure (see Section 3.1.2); this includes link-LSAs.

To lookup or install an LSA in the database, you first examine the LS type and the LSA's context (i.e., to which area or link does the LSA belong). This information allows you to find the correct list of LSAs, all of the same LS type, where you then search based on the LSA's Link State ID and Advertising Router.

3.4.3. Originating LSAs

The process of reoriginating an LSA in IPv6 is the same as in IPv4: the LSA's LS sequence number is incremented, its LS age is set to 0, its LS checksum is calculated, and the LSA is added to the link state database and flooded out the appropriate interfaces.

To the list of events causing LSAs to be reoriginated, which for IPv4 is given in Section 12.4 of [Ref1], the following events and/or actions are added for IPv6:

- o The state of one of the router's interfaces changes. The router may need to (re)originate or flush its Link-LSA and one or more router-LSAs and/or intra-area-prefix-LSAs.
- o The identity of a link's Designated Router changes. The router may need to (re)originate or flush the link's network-LSA and one or more router-LSAs and/or intra-area-prefix-LSAs.

- o A neighbor transitions to/from "Full" state. The router may need to (re)originate or flush the link's network-LSA and one or more router-LSAs and/or intra-area-prefix-LSAs.
- o The Interface ID of a neighbor changes. This may cause a new instance of a router-LSA to be originated for the associated area, and the reorigination of one or more intra-area-prefix-LSAs.
- o A new prefix is added to an attached link, or a prefix is deleted (both through configuration). This causes the router to reoriginate its link-LSA for the link, or, if it is the only router attached to the link, causes the router to reoriginate an intra-area-prefix-LSA.
- o A new link-LSA is received, causing the link's collection of prefixes to change. If the router is Designated Router for the link, it originates a new intra-area-prefix-LSA.

Detailed construction of the seven required IPv6 LSA types is supplied by the following subsections. In order to display example LSAs, the network map in Figure 15 of [Ref1] has been reworked to show IPv6 addressing, resulting in Figure 1. The OSPF cost of each interface is has been displayed in Figure 1. The assignment of IPv6 prefixes to network links is shown in Table 1. A single area address range has been configured for Area 1, so that outside of Area 1 all of its prefixes are covered by a single route to 5f00:0000:c001::/48. The OSPF interface IDs and the link-local addresses for the router interfaces in Figure 1 are given in Table 2.

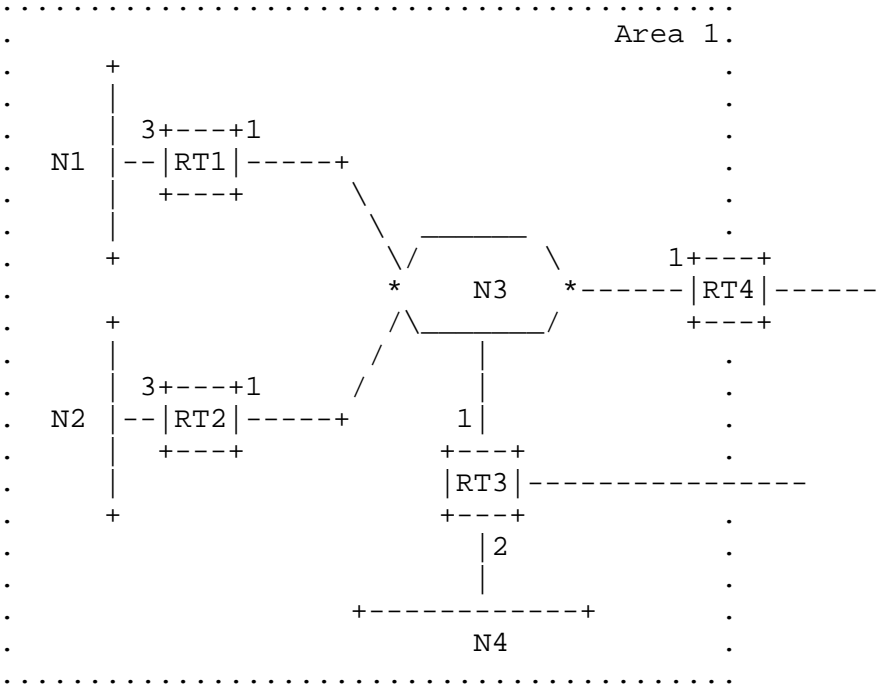


Figure 1: Area 1 with IP addresses shown

| Network | IPv6 prefix |
|---------|--------------------------|
| N1 | 5f00:0000:c001:0200::/56 |
| N2 | 5f00:0000:c001:0300::/56 |
| N3 | 5f00:0000:c001:0100::/56 |
| N4 | 5f00:0000:c001:0400::/56 |

Table 1: IPv6 link prefixes for sample network

| Router | interface | Interface ID | link-local address |
|--------|-----------|--------------|--------------------|
| RT1 | to N1 | 1 | fe80:0001::RT1 |
| | to N3 | 2 | fe80:0002::RT1 |
| RT2 | to N2 | 1 | fe80:0001::RT2 |
| | to N3 | 2 | fe80:0002::RT2 |
| RT3 | to N3 | 1 | fe80:0001::RT3 |
| | to N4 | 2 | fe80:0002::RT3 |
| RT4 | to N3 | 1 | fe80:0001::RT4 |

Table 2: OSPF Interface IDs and link-local addresses

3.4.3.1. Router-LSAs

The LS type of a router-LSA is set to the value 0x2001. Router-LSAs have area flooding scope. A router may originate one or more router-LSAs for a given area. Each router-LSA contains an integral number of interface descriptions; taken together, the collection of router-LSAs originated by the router for an area describes the collected states of all the router's interfaces to the area. When multiple router-LSAs are used, they are distinguished by their Link State ID fields.

The Options field in the router-LSA should be coded as follows. The V6-bit should be set. The E-bit should be clear if and only if the attached area is an OSPF stub area. The MC-bit should be set if and only if the router is running MOSPF (see [Ref8]). The N-bit should be set if and only if the attached area is an OSPF NSSA area. The R-bit should be set. The DC-bit should be set if and only if the router can correctly process the DoNotAge bit when it appears in the LS age field of LSAs (see [Ref11]). All unrecognized bits in the Options field should be cleared

To the left of the Options field, the router capability bits V, E and B should be coded according to Section 12.4.1 of [Ref1]. Bit W should be coded according to [Ref8].

Each of the router's interfaces to the area are then described by appending "link descriptions" to the router-LSA. Each link description is 16 bytes long, consisting of 5 fields: (link) Type, Metric, Interface ID, Neighbor Interface ID and Neighbor Router ID (see Section A.4.3). Interfaces in state "Down" or "Loopback" are not described (although looped back interfaces can contribute prefixes to Intra-Area-Prefix-LSAs). Nor are interfaces without any full adjacencies described. All other interfaces to the area add zero, one or more link descriptions, the number and content of which depend on the interface type. Within each link description, the Metric field is always set the interface's output cost and the Interface ID field is set to the interface's OSPF Interface ID.

Point-to-point interfaces

If the neighboring router is fully adjacent, add a Type 1 link description (point-to-point). The Neighbor Interface ID field is set to the Interface ID advertised by the neighbor in its Hello packets, and the Neighbor Router ID field is set to the neighbor's Router ID.

Broadcast and NBMA interfaces

If the router is fully adjacent to the link's Designated Router, or if the router itself is Designated Router and is fully adjacent to at least one other router, add a single Type 2 link description (transit network). The Neighbor Interface ID field is set to the Interface ID advertised by the Designated Router in its Hello packets, and the Neighbor Router ID field is set to the Designated Router's Router ID.

Virtual links

If the neighboring router is fully adjacent, add a Type 4 link description (virtual). The Neighbor Interface ID field is set to the Interface ID advertised by the neighbor in its Hello packets, and the Neighbor Router ID field is set to the neighbor's Router ID. Note that the output cost of a virtual link is calculated during the routing table calculation (see Section 3.7).

Point-to-MultiPoint interfaces

For each fully adjacent neighbor associated with the interface, add a separate Type 1 link description (point-to-point) with Neighbor Interface ID field set to the Interface ID advertised by the neighbor in its Hello packets, and Neighbor Router ID field set to the neighbor's Router ID.

As an example, consider the router-LSA that router RT3 would originate for Area 1 in Figure 1. Only a single interface must be described, namely that which connects to the transit network N3. It assumes that RT4 has been elected Designated Router of Network N3.

; RT3's router-LSA for Area 1

```

LS age = 0                                ;newly (re)originated
LS type = 0x2001                          ;router-LSA
Link State ID = 0                         ;first fragment
Advertising Router = 192.1.1.3           ;RT3's Router ID
bit E = 0                                ;not an AS boundary router
bit B = 1                                ;area border router
Options = (V6-bit|E-bit|R-bit)
  Type = 2                                ;connects to N3
  Metric = 1                              ;cost to N3
  Interface ID = 1                        ;RT3's Interface ID on N3
  Neighbor Interface ID = 1               ;RT4's Interface ID on N3
  Neighbor Router ID = 192.1.1.4 ; RT4's Router ID

```

If for example another router was added to Network N4, RT3 would have to advertise a second link description for its connection to (the now transit) network N4. This could be accomplished by reoriginating the above router-LSA, this time with two link descriptions. Or, a

separate router-LSA could be originated with a separate Link State ID (e.g., using a Link State ID of 1) to describe the connection to N4.

Host routes no longer appear in the router-LSA, but are instead included in intra-area-prefix-LSAs.

3.4.3.2. Network-LSAs

The LS type of a network-LSA is set to the value 0x2002. Network-LSAs have area flooding scope. A network-LSA is originated for every broadcast or NBMA link having two or more attached routers, by the link's Designated Router. The network-LSA lists all routers attached to the link.

The procedure for originating network-LSAs in IPv6 is the same as the IPv4 procedure documented in Section 12.4.2 of [Ref1], with the following exceptions:

- o An IPv6 network-LSA's Link State ID is set to the Interface ID of the Designated Router on the link.
- o IPv6 network-LSAs do not contain a Network Mask. All addressing information formerly contained in the IPv4 network-LSA has now been consigned to intra-Area-Prefix-LSAs.
- o The Options field in the network-LSA is set to the logical OR of the Options fields contained within the link's associated link-LSAs. In this way, the network link exhibits a capability when at least one of the link's routers requests that the capability be asserted.

As an example, assuming that Router RT4 has been elected Designated Router of Network N3 in Figure 1, the following network-LSA is originated:

; Network-LSA for Network N3

```

LS age = 0                      ;newly (re)originated
LS type = 0x2002                 ;network-LSA
Link State ID = 1                ;RT4's Interface ID on N3
Advertising Router = 192.1.1.4  ;RT4's Router ID
Options = (V6-bit|E-bit|R-bit)
    Attached Router = 192.1.1.4    ;Router ID
    Attached Router = 192.1.1.1    ;Router ID
    Attached Router = 192.1.1.2    ;Router ID
    Attached Router = 192.1.1.3    ;Router ID

```

3.4.3.3. Inter-Area-Prefix-LSAs

The LS type of an inter-area-prefix-LSA is set to the value 0x2003. Inter-area-prefix-LSAs have area flooding scope. In IPv4, inter-area-prefix-LSAs were called type 3 summary-LSAs. Each inter-area-prefix-LSA describes a prefix external to the area, yet internal to the Autonomous System.

The procedure for originating inter-area-prefix-LSAs in IPv6 is the same as the IPv4 procedure documented in Sections 12.4.3 and 12.4.3.1 of [Ref1], with the following exceptions:

- o The Link State ID of an inter-area-prefix-LSA has lost all of its addressing semantics, and instead simply serves to distinguish multiple inter-area-prefix-LSAs that are originated by the same router.
- o The prefix is described by the PrefixLength, PrefixOptions and Address Prefix fields embedded within the LSA body. Network Mask is no longer specified.
- o The NU-bit in the PrefixOptions field should be clear. The coding of the MC-bit depends upon whether, and if so how, MOSPF is operating in the routing domain (see [Ref8]).
- o Link-local addresses must never be advertised in inter-area-prefix-LSAs.

As an example, the following shows the inter-area-prefix-LSA that Router RT4 originates into the OSPF backbone area, condensing all of Area 1's prefixes into the single prefix 5f00:0000:c001::/48. The cost is set to 4, which is the maximum cost to all of the prefix' individual components. The prefix is padded out to an even number of 32-bit words, so that it consumes 64-bits of space instead of 48 bits.

```

; Inter-area-prefix-LSA for Area 1 addresses
; originated by Router RT4 into the backbone

LS age = 0                ;newly (re)originated
LS type = 0x2003          ;inter-area-prefix-LSA
Advertising Router = 192.1.1.4 ;RT4's ID
Metric = 4                ;maximum to components
PrefixLength = 48
PrefixOptions = 0
Address Prefix = 5f00:0000:c001 ;padded to 64-bits
```

3.4.3.4. Inter-Area-Router-LSAs

The LS type of an inter-area-router-LSA is set to the value 0x2004. Inter-area-router-LSAs have area flooding scope. In IPv4, inter-area-router-LSAs were called type 4 summary-LSAs. Each inter-area-router-LSA describes a path to a destination OSPF router (an ASBR) that is external to the area, yet internal to the Autonomous System.

The procedure for originating inter-area-router-LSAs in IPv6 is the same as the IPv4 procedure documented in Section 12.4.3 of [Ref1], with the following exceptions:

- o The Link State ID of an inter-area-router-LSA is no longer the destination router's OSPF Router ID, but instead simply serves to distinguish multiple inter-area-router-LSAs that are originated by the same router. The destination router's Router ID is now found in the body of the LSA.
- o The Options field in an inter-area-router-LSA should be set equal to the Options field contained in the destination router's own router-LSA. The Options field thus describes the capabilities supported by the destination router.

As an example, consider the OSPF Autonomous System depicted in Figure 6 of [Ref1]. Router RT4 would originate into Area 1 the following inter-area-router-LSA for destination router RT7.

```
; inter-area-router-LSA for AS boundary router RT7
; originated by Router RT4 into Area 1

LS age = 0                ;newly (re)originated
LS type = 0x2004           ;inter-area-router-LSA
Advertising Router = 192.1.1.4 ;RT4's ID
Options = (V6-bit|E-bit|R-bit) ;RT7's capabilities
Metric = 14                ;cost to RT7
Destination Router ID = Router RT7's ID
```

3.4.3.5. AS-external-LSAs

The LS type of an AS-external-LSA is set to the value 0x4005. AS-external-LSAs have AS flooding scope. Each AS-external-LSA describes a path to a prefix external to the Autonomous System.

The procedure for originating AS-external-LSAs in IPv6 is the same as the IPv4 procedure documented in Section 12.4.4 of [Ref1], with the following exceptions:

- o The Link State ID of an AS-external-LSA has lost all of its addressing semantics, and instead simply serves to distinguish multiple AS-external-LSAs that are originated by the same router.
- o The prefix is described by the PrefixLength, PrefixOptions and Address Prefix fields embedded within the LSA body. Network Mask is no longer specified.
- o The NU-bit in the PrefixOptions field should be clear. The coding of the MC-bit depends upon whether, and if so how, MOSPF is operating in the routing domain (see [Ref8]).
- o Link-local addresses can never be advertised in AS-external-LSAs.
- o The forwarding address is present in the AS-external-LSA if and only if the AS-external-LSA's bit F is set.
- o The external route tag is present in the AS-external-LSA if and only if the AS-external-LSA's bit T is set.
- o The capability for an AS-external-LSA to reference another LSA has been included, by inclusion of the Referenced LS Type field and the optional Referenced Link State ID field (the latter present if and only if Referenced LS Type is non-zero). This capability is for future use; for now Referenced LS Type should be set to 0 and received non-zero values for this field should be ignored.

As an example, consider the OSPF Autonomous System depicted in Figure 6 of [Ref1]. Assume that RT7 has learned its route to N12 via BGP, and that it wishes to advertise a Type 2 metric into the AS. Further assume the the IPv6 prefix for N12 is the value 5f00:0000:0a00::/40. RT7 would then originate the following AS-external-LSA for the external network N12. Note that within the AS-external-LSA, N12's prefix occupies 64 bits of space, to maintain 32-bit alignment.

```
; AS-external-LSA for Network N12,
; originated by Router RT7
```

```
LS age = 0                ;newly (re)originated
LS type = 0x4005          ;AS-external-LSA
Link State ID = 123       ;or something else
Advertising Router = Router RT7's ID
bit E = 1                ;Type 2 metric
bit F = 0                ;no forwarding address
bit T = 1                ;external route tag included
Metric = 2
PrefixLength = 40
PrefixOptions = 0
```

Referenced LS Type = 0 ;no Referenced Link State ID
Address Prefix = 5f00:0000:0a00 ;padded to 64-bits
External Route Tag = as per BGP/OSPF interaction

3.4.3.6. Link-LSAs

The LS type of a Link-LSA is set to the value 0x0008. Link-LSAs have link-local flooding scope. A router originates a separate Link-LSA for each attached link that supports 2 or more (including the originating router itself) routers.

Link-LSAs have three purposes: 1) they provide the router's link-local address to all other routers attached to the link and 2) they inform other routers attached to the link of a list of IPv6 prefixes to associate with the link and 3) they allow the router to assert a collection of Options bits in the Network-LSA that will be originated for the link.

A Link-LSA for a given Link L is built in the following fashion:

- o The Link State ID is set to the router's Interface ID on Link L.
- o The Router Priority of the router's interface to Link L is inserted into the Link-LSA.
- o The Link-LSA's Options field is set to those bits that the router wishes set in Link L's Network LSA.
- o The router inserts its link-local address on Link L into the Link-LSA. This information will be used when the other routers on Link L do their next hop calculations (see Section 3.8.1.1).
- o Each IPv6 address prefix that has been configured into the router for Link L is added to the Link-LSA, by specifying values for PrefixLength, PrefixOptions, and Address Prefix fields.

After building a Link-LSA for a given link, the router installs the link-LSA into the associate interface data structure and floods the Link-LSA onto the link. All other routers on the link will receive the Link-LSA, but it will go no further.

As an example, consider the Link-LSA that RT3 will build for N3 in Figure 1. Suppose that the prefix 5f00:0000:c001:0100::/56 has been configured within RT3 for N3. This will give rise to the following Link-LSA, which RT3 will flood onto N3, but nowhere else. Note that not all routers on N3 need be configured with the prefix; those not configured will learn the prefix when receiving RT3's Link-LSA.

```

; RT3's Link-LSA for N3

LS age = 0                      ;newly (re)originated
LS type = 0x0008                ;Link-LSA
Link State ID = 1               ;RT3's Interface ID on N3
Advertising Router = 192.1.1.3 ;RT3's Router ID
Rtr Pri = 1                     ;RT3's N3 Router Priority
Options = (V6-bit|E-bit|R-bit)
Link-local Interface Address = fe80:0001::RT3
# prefixes = 1
PrefixLength = 56
PrefixOptions = 0
Address Prefix = 5f00:0000:c001:0100 ;pad to 64-bits

```

3.4.3.7. Intra-Area-Prefix-LSAs

The LS type of an intra-area-prefix-LSA is set to the value 0x2009. Intra-area-prefix-LSAs have area flooding scope. An intra-area-prefix-LSA has one of two functions. It associates a list of IPv6 address prefixes with a transit network link by referencing a network-LSA, or associates a list of IPv6 address prefixes with a router by referencing a router-LSA. A stub link's prefixes are associated with its attached router.

A router may originate multiple intra-area-prefix-LSAs for a given area, distinguished by their Link State ID fields. Each intra-area-prefix-LSA contains an integral number of prefix descriptions.

A link's Designated Router originates one or more intra-area-prefix-LSAs to advertise the link's prefixes throughout the area. For a link L, L's Designated Router builds an intra-area-prefix-LSA in the following fashion:

- o In order to indicate that the prefixes are to be associated with the Link L, the fields Referenced LS type, Referenced Link State ID, and Referenced

Advertising Router are set to the corresponding fields in Link L's network-LSA (namely LS type, Link State ID, and Advertising Router respectively). This means that Referenced LS Type is set to 0x2002, Referenced Link State ID is set to the Designated Router's Interface ID on Link L, and Referenced Advertising Router is set to the Designated Router's Router ID.

- o Each Link-LSA associated with Link L is examined (these are in the Designated Router's interface structure for Link L). If the Link-LSA's Advertising Router is fully adjacent to the Designated Router, the list of prefixes in the Link-LSA is copied into the

intra-area-prefix-LSA that is being built. Prefixes having the NU-bit and/or LA-bit set in their Options field should not be copied, nor should link-local addresses be copied. Each prefix is described by the PrefixLength, PrefixOptions, and Address Prefix fields. Multiple prefixes having the same PrefixLength and Address Prefix are considered to be duplicates; in this case their Prefix Options fields should be merged by logically OR'ing the fields together, and a single resulting prefix should be copied into the intra-area-prefix-LSA. The Metric field for all prefixes is set to 0.

- o The "# prefixes" field is set to the number of prefixes that the router has copied into the LSA. If necessary, the list of prefixes can be spread across multiple intra-area-prefix-LSAs in order to keep the LSA size small.

A router builds an intra-area-prefix-LSA to advertise its own prefixes, and those of its attached stub links. A Router RTX would build its intra-area-prefix-LSA in the following fashion:

- o In order to indicate that the prefixes are to be associated with the Router RTX itself, RTX sets Referenced LS type to 0x2001, Referenced Link State ID to 0, and Referenced Advertising Router to RTX's own Router ID.
- o Router RTX examines its list of interfaces to the area. If the interface is in state Down, its prefixes are not included. If the interface has been reported in RTX's router-LSA as a Type 2 link description (link to transit network), its prefixes are not included (they will be included in the intra-area-prefix-LSA for the link instead). If the interface type is Point-to-MultiPoint, or the interface is in state Loopback, or the interface connects to a point-to-point link which has not been assigned a prefix, then the site-local and global scope IPv6 addresses associated with the interface (if any) are copied into the intra-area-prefix-LSA, setting the LA-bit in the PrefixOptions field, and setting the PrefixLength to 128 and the Metric to 0. Otherwise, the list of site-local and global prefixes configured in RTX for the link are copied into the intra-area-prefix-LSA by specifying the PrefixLength, PrefixOptions, and Address Prefix fields. The Metric field for each of these prefixes is set to the interface's output cost.
- o RTX adds the IPv6 prefixes for any directly attached hosts belonging to the area (see Section C.7) to the intra-area-prefix-LSA.

- o If RTX has one or more virtual links configured through the area, it includes one of its site-local or global scope IPv6 interface addresses in the LSA (if it hasn't already), setting the LA-bit in the PrefixOptions field, and setting the PrefixLength to 128 and the Metric to 0. This information will be used later in the routing calculation so that the two ends of the virtual link can discover each other's IPv6 addresses.
- o The "# prefixes" field is set to the number of prefixes that the router has copied into the LSA. If necessary, the list of prefixes can be spread across multiple intra-area-prefix-LSAs in order to keep the LSA size small.

For example, the intra-area-prefix-LSA originated by RT4 for Network N3 (assuming that RT4 is N3's Designated Router), and the intra-area-prefix-LSA originated into Area 1 by Router RT3 for its own prefixes, are pictured below.

```

; Intra-area-prefix-LSA
; for network link N3

LS age = 0                      ;newly (re)originated
LS type = 0x2009                 ;Intra-area-prefix-LSA
Link State ID = 5                ;or something
Advertising Router = 192.1.1.4 ;RT4's Router ID
# prefixes = 1
Referenced LS type = 0x2002 ;network-LSA reference
Referenced Link State ID = 1
Referenced Advertising Router = 192.1.1.4
PrefixLength = 56                ;N3's prefix
PrefixOptions = 0
Metric = 0
Address Prefix = 5f00:0000:c001:0100 ;pad

; RT3's Intra-area-prefix-LSA
; for its own prefixes

LS age = 0                      ;newly (re)originated
LS type = 0x2009                 ;Intra-area-prefix-LSA
Link State ID = 177              ;or something
Advertising Router = 192.1.1.3 ;RT3's Router ID
# prefixes = 1
Referenced LS type = 0x2001 ;router-LSA reference
Referenced Link State ID = 0
Referenced Advertising Router = 192.1.1.3
PrefixLength = 56                ;N4's prefix

```

```
PrefixOptions = 0
Metric = 2 ;N4 interface cost
Address Prefix = 5f00:0000:c001:0400 ;pad
```

When network conditions change, it may be necessary for a router to move prefixes from one intra-area-prefix-LSA to another. For example, if the router is Designated Router for a link but the link has no other attached routers, the link's prefixes are advertised in an intra-area-prefix-LSA referring to the Designated Router's router-LSA. When additional routers appear on the link, a network-LSA is originated for the link and the link's prefixes are moved to an intra-area-prefix-LSA referring to the network-LSA.

Note that in the intra-area-prefix-LSA, the "Referenced Advertising Router" is always equal to the router that is originating the intra-area-prefix-LSA (i.e., the LSA's Advertising Router). The reason that the Referenced Advertising Router field appears is that, even though it is currently redundant, it may not be in the future. We may sometime want to use the same LSA format to advertise address prefixes for other protocol suites. In that event, the Designated Router may not be running the other protocol suite, and so another of the link's routers may need to send out the prefix-LSA. In that case, "Referenced Advertising Router" and "Advertising Router" would be different.

3.5. Flooding

Most of the flooding algorithm remains unchanged from the IPv4 flooding mechanisms described in Section 13 of [Ref1]. In particular, the processes for determining which LSA instance is newer (Section 13.1 of [Ref1]), responding to updates of self-originated LSAs (Section 13.4 of [Ref1]), sending Link State Acknowledgment packets (Section 13.5 of [Ref1]), retransmitting LSAs (Section 13.6 of [Ref1]) and receiving Link State Acknowledgment packets (Section 13.7 of [Ref1]) are exactly the same for IPv6 and IPv4.

However, the addition of flooding scope and handling options for unrecognized LSA types (see Section A.4.2.1) has caused some changes in the OSPF flooding algorithm: the reception of Link State Updates (Section 13 in [Ref1]) and the sending of Link State Updates (Section 13.3 of [Ref1]) must take into account the LSA's scope and U-bit setting. Also, installation of LSAs into the OSPF database (Section 13.2 of [Ref1]) causes different events in IPv6, due to the reorganization of LSA types and contents in IPv6. These changes are described in detail below.

3.5.1. Receiving Link State Update packets

The encoding of flooding scope in the LS type and the need to process unknown LS types causes modifications to the processing of received Link State Update packets. As in IPv4, each LSA in a received Link State Update packet is examined. In IPv4, eight steps are executed for each LSA, as described in Section 13 of [Ref1]. For IPv6, all the steps are the same, except that Steps 2 and 3 are modified as follows:

- (2) Examine the LSA's LS type. If the LS type is unknown, the area has been configured as a stub area, and either the LSA's flooding scope is set to "AS flooding scope" or the U-bit of the LS type is set to 1 (flood even when unrecognized), then discard the LSA and get the next one from the Link State Update Packet. This generalizes the IPv4 behavior where AS-external-LSAs are not flooded into/throughout stub areas.
- (3) Else if the flooding scope of the LSA is set to "reserved", discard the LSA and get the next one from the Link State Update Packet.

Steps 5b (sending Link State Update packets) and 5d (installing LSAs in the link state database) in Section 13 of [Ref1] are also somewhat different for IPv6, as described in Sections 3.5.2 and 3.5.3 below.

3.5.2. Sending Link State Update packets

The sending of Link State Update packets is described in Section 13.3 of [Ref1]. For IPv4 and IPv6, the steps for sending a Link State Update packet are the same (steps 1 through 5 of Section 13.3 in [Ref1]). However, the list of eligible interfaces out which to flood the LSA is different. For IPv6, the eligible interfaces are selected based on the following factors:

- o The LSA's flooding scope.
- o For LSAs with area or link-local flooding scoping, the particular area or interface that the LSA is associated with.
- o Whether the LSA has a recognized LS type.
- o The setting of the U-bit in the LS type. If the U-bit is set to 0, unrecognized LS types are treated as having link-local scope. If set to 1, unrecognized LS types are stored and flooded as if they were recognized.

Choosing the set of eligible interfaces then breaks into the following cases:

Case 1

The LSA's LS type is recognized. In this case, the set of eligible interfaces is set depending on the flooding scope encoded in the LS type. If the flooding scope is "AS flooding scope", the eligible interfaces are all router interfaces excepting virtual links. In addition, AS-external-LSAs are not flooded out interfaces connecting to stub areas. If the flooding scope is "area flooding scope", the set of eligible interfaces are those interfaces connecting to the LSA's associated area. If the flooding scope is "link-local flooding scope", then there is a single eligible interface, the one connecting to the LSA's associated link (which, when the LSA is received in a Link State Update packet, is also the interface the LSA was received on).

Case 2

The LS type is unrecognized, and the U-bit in the LS Type is set to 0 (treat the LSA as if it had link-local flooding scope). In this case there is a single eligible interface, namely, the interface on which the LSA was received.

Case 3

The LS type is unrecognized, and the U-bit in the LS Type is set to 1 (store and flood the LSA, as if type understood). In this case, select the eligible interfaces based on the encoded flooding scope as in Case 1 above. However, in this case interfaces attached to stub areas are always excluded.

A further decision must sometimes be made before adding an LSA to a given neighbor's link-state retransmission list (Step 1d in Section 13.3 of [Ref1]). If the LS type is recognized by the router, but not by the neighbor (as can be determined by examining the Options field that the neighbor advertised in its Database Description packet) and the LSA's U-bit is set to 0, then the LSA should be added to the neighbor's link-state retransmission list if and only if that neighbor is the Designated Router or Backup Designated Router for the attached link. The LS types described in detail by this memo, namely router-LSAs (LS type 0x2001), network-LSAs (0x2002), Inter-Area-Prefix-LSAs (0x2003), Inter-Area-Router-LSAs (0x2004), AS-External-LSAs (0x4005), Link-LSAs (0x0008) and Intra-Area-Prefix-LSAs (0x2009) are assumed to be understood by all routers. However, as an example the group-membership-LSA (0x2006) is understood only by MOSPF routers and since it has its U-bit set to 0, it should only be forwarded to a non-MOSPF neighbor (determined by examining the MC-bit in the neighbor's Database Description packets' Options field) when the neighbor is Designated Router or Backup Designated Router for the

attached link.

The previous paragraph solves a problem in IPv4 OSPF extensions such as MOSPF, which require that the Designated Router support the extension in order to have the new LSA types flooded across broadcast and NBMA networks (see Section 10.2 of [Ref8]).

3.5.3. Installing LSAs in the database

There are three separate places to store LSAs, depending on their flooding scope. LSAs with AS flooding scope are stored in the global OSPF data structure (see Section 3.1) as long as their LS type is known or their U-bit is 1. LSAs with area flooding scope are stored in the appropriate area data structure (see Section 3.1.1) as long as their LS type is known or their U-bit is 1. LSAs with link-local flooding scope, and those LSAs with unknown LS type and U-bit set to 0 (treat the LSA as if it had link-local flooding scope) are stored in the appropriate interface structure.

When storing the LSA into the link-state database, a check must be made to see whether the LSA's contents have changed. Changes in contents are indicated exactly as in Section 13.2 of [Ref1]. When an LSA's contents have been changed, the following parts of the routing table must be recalculated, based on the LSA's LS type:

Router-LSAs, Network-LSAs, Intra-Area-Prefix-LSAs and Link-LSAs

The entire routing table is recalculated, starting with the shortest path calculation for each area (see Section 3.8).

Inter-Area-Prefix-LSAs and Inter-Area-Router-LSAs

The best route to the destination described by the LSA must be recalculated (see Section 16.5 in [Ref1]). If this destination is an AS boundary router, it may also be necessary to re-examine all the AS-external-LSAs.

AS-external-LSAs

The best route to the destination described by the AS-external-LSA must be recalculated (see Section 16.6 in [Ref1]).

As in IPv4, any old instance of the LSA must be removed from the database when the new LSA is installed. This old instance must also be removed from all neighbors' Link state retransmission lists.

3.6. Definition of self-originated LSAs

In IPv6 the definition of a self-originated LSA has been simplified from the IPv4 definition appearing in Sections 13.4 and 14.1 of [Ref1]. For IPv6, self-originated LSAs are those LSAs whose Advertising Router is equal to the router's own Router ID.

3.7. Virtual links

OSPF virtual links for IPv4 are described in Section 15 of [Ref1]. Virtual links are the same in IPv6, with the following exceptions:

- o LSAs having AS flooding scope are never flooded over virtual adjacencies, nor are LSAs with AS flooding scope summarized over virtual adjacencies during the Database Exchange process. This is a generalization of the IPv4 treatment of AS-external-LSAs.
- o The IPv6 interface address of a virtual link must be an IPv6 address having site-local or global scope, instead of the link-local addresses used by other interface types. This address is used as the IPv6 source for OSPF protocol packets sent over the virtual link.
- o Likewise, the virtual neighbor's IPv6 address is an IPv6 address with site-local or global scope. To enable the discovery of a virtual neighbor's IPv6 address during the routing calculation, the neighbor advertises its virtual link's IPv6 interface address in an Intra-Area-Prefix-LSA originated for the virtual link's transit area (see Sections 3.4.3.7 and 3.8.1).
- o Like all other IPv6 OSPF interfaces, virtual links are assigned unique (within the router) Interface IDs. These are advertised in Hellos sent over the virtual link, and in the router's router-LSAs.

3.8. Routing table calculation

The IPv6 OSPF routing calculation proceeds along the same lines as the IPv4 OSPF routing calculation, following the five steps specified by Section 16 of [Ref1]. High level differences between the IPv6 and IPv4 calculations include:

- o Prefix information has been removed from router-LSAs, and now is advertised in intra-area-prefix-LSAs. Whenever [Ref1] specifies that stub networks within router-LSAs be examined, IPv6 will instead examine prefixes within intra-area-prefix-LSAs.

- o Type 3 and 4 summary-LSAs have been renamed inter-area-prefix-LSAs and inter-area-router-LSAs (respectively).
- o Addressing information is no longer encoded in Link State IDs, and must instead be found within the body of LSAs.
- o In IPv6, a router can originate multiple router-LSAs within a single area, distinguished by Link State ID. These router-LSAs must be treated as a single aggregate by the area's shortest path calculation (see Section 3.8.1).

For each area, routing table entries have been created for the area's routers and transit links, in order to store the results of the area's shortest-path tree calculation (see Section 3.8.1). These entries are then used when processing intra-area-prefix-LSAs, inter-area-prefix-LSAs and inter-area-router-LSAs, as described in Section 3.8.2.

Events generated as a result of routing table changes (Section 16.7 of [Ref1]), and the equal-cost multipath logic (Section 16.8 of [Ref1]) are identical for both IPv4 and IPv6.

3.8.1. Calculating the shortest path tree for an area

The IPv4 shortest path calculation is contained in Section 16.1 of [Ref1]. The graph used by the shortest-path tree calculation is identical for both IPv4 and IPv6. The graph's vertices are routers and transit links, represented by router-LSAs and network-LSAs respectively. A router is identified by its OSPF Router ID, while a transit link is identified by its Designated Router's Interface ID and OSPF Router ID. Both routers and transit links have associated routing table entries within the area (see Section 3.3).

Section 16.1 of [Ref1] splits up the shortest path calculations into two stages. First the Dijkstra calculation is performed, and then the stub links are added onto the tree as leaves. The IPv6 calculation maintains this split.

The Dijkstra calculation for IPv6 is identical to that specified for IPv4, with the following exceptions (referencing the steps from the Dijkstra calculation as described in Section 16.1 of [Ref1]):

- o The Vertex ID for a router is the OSPF Router ID. The Vertex ID for a transit network is a combination of the Interface ID and OSPF Router ID of the network's Designated Router.

- o In Step 2, when a router Vertex V has just been added to the shortest path tree, there may be multiple LSAs associated with the router. All Router-LSAs with Advertising Router set to V's OSPF Router ID must be processed as an aggregate, treating them as fragments of a single large router-LSA. The Options field and the router type bits (bits W, V, E and B) should always be taken from "fragment" with the smallest Link State ID.
- o Step 2a is not needed in IPv6, as there are no longer stub network links in router-LSAs.
- o In Step 2b, if W is a router, there may again be multiple LSAs associated with the router. All Router-LSAs with Advertising Router set to W's OSPF Router ID must be processed as an aggregate, treating them as fragments of a single large router-LSA.
- o In Step 4, there are now per-area routing table entries for each of an area's routers, instead of just the area border routers. These entries subsume all the functionality of IPv4's area border router routing table entries, including the maintenance of virtual links. When the router added to the area routing table in this step is the other end of a virtual link, the virtual neighbor's IP address is set as follows: The collection of intra-area-prefix-LSAs originated by the virtual neighbor is examined, with the virtual neighbor's IP address being set to the first prefix encountered having the "LA-bit" set.
- o Routing table entries for transit networks, which are no longer associated with IP networks, are also modified in Step 4.

The next stage of the shortest path calculation proceeds similarly to the two steps of the second stage of Section 16.1 in [Ref1]. However, instead of examining the stub links within router-LSAs, the list of the area's intra-area-prefix-LSAs is examined. A prefix advertisement whose "NU-bit" is set should not be included in the routing calculation. The cost of any advertised prefix is the sum of the prefix' advertised metric plus the cost to the transit vertex (either router or transit network) identified by intra-area-prefix-LSA's Referenced LS type, Referenced Link State ID and Referenced Advertising Router fields. This latter cost is stored in the transit vertex' routing table entry for the area.

3.8.1.1. The next hop calculation

In IPv6, the calculation of the next hop's IPv6 address (which will be a link-local address) proceeds along the same lines as the IPv4 next hop calculation (see Section 16.1.1 of [Ref1]). The only difference is in calculating the next hop IPv6 address for a router

(call it Router X) which shares a link with the calculating router. In this case the calculating router assigns the next hop IPv6 address to be the link-local interface address contained in Router X's Link-LSA (see Section A.4.8) for the link. This procedure is necessary since on some links, such as NBMA links, the two routers need not be neighbors, and therefore might not be exchanging OSPF Hellos.

3.8.2. Calculating the inter-area routes

Calculation of inter-area routes for IPv6 proceeds along the same lines as the IPv4 calculation in Section 16.2 of [Ref1], with the following modifications:

- o The names of the Type 3 summary-LSAs and Type 4 summary-LSAs have been changed to inter-area-prefix-LSAs and inter-area-router-LSAs respectively.
- o The Link State ID of the above LSA types no longer encodes the network or router described by the LSA. Instead, an address prefix is contained in the body of an inter-area-prefix-LSA, and a described router's OSPF Router ID is carried in the body of an inter-area-router-LSA.
- o Prefixes having the "NU-bit" set in their Prefix Options field should be ignored by the inter-area route calculation.

When a single inter-area-prefix-LSA or inter-area-router-LSA has changed, the incremental calculations outlined in Section 16.5 of [Ref1] can be performed instead of recalculating the entire routing table.

3.8.3. Examining transit areas' summary-LSAs

Examination of transit areas' summary-LSAs in IPv6 proceeds along the same lines as the IPv4 calculation in Section 16.3 of [Ref1], modified in the same way as the IPv6 inter-area route calculation in Section 3.8.2.

3.8.4. Calculating AS external routes

The IPv6 AS external route calculation proceeds along the same lines as the IPv4 calculation in Section 16.4 of [Ref1], with the following exceptions:

- o The Link State ID of the AS-external-LSA types no longer encodes the network described by the LSA. Instead, an address prefix is contained in the body of an AS-external-LSA.

- o The default route is described by AS-external-LSAs which advertise zero length prefixes.
- o Instead of comparing the AS-external-LSA's Forwarding address field to 0.0.0.0 to see whether a forwarding address has been used, bit F of the external-LSA is examined. A forwarding address is in use if and only if bit F is set.
- o Prefixes having the "NU-bit" set in their Prefix Options field should be ignored by the inter-area route calculation.

When a single AS-external-LSA has changed, the incremental calculations outlined in Section 16.6 of [Ref1] can be performed instead of recalculating the entire routing table.

3.9. Multiple interfaces to a single link

In OSPF for IPv6, a router may have multiple interfaces to a single link. All interfaces are involved in the reception and transmission of data traffic, however only a single interface sends and receives OSPF control traffic. In more detail:

- o Each of the multiple interfaces are assigned different Interface IDs. In this way the router can automatically detect when multiple interfaces attach to the same link, when receiving Hellos from its own Router ID but with an Interface ID other than the receiving interface's.
- o The router turns off the sending and receiving of OSPF packets (that is, control traffic) on all but one of the interfaces to the link. The choice of interface to send and receive control traffic is implementation dependent; as one example, the interface with the highest Interface ID could be chosen. If the router is elected DR, it will be this interface's Interface ID that will be used as the network-LSA's Link State ID.
- o All the multiple interfaces to the link will however appear in the router-LSA. In addition, a Link-LSA will be generated for each of the multiple interfaces. In this way, all interfaces will be included in OSPF's routing calculations.
- o If the interface which is responsible for sending and receiving control traffic fails, another will have to take over, reforming all neighbor adjacencies from scratch. This failure can be detected by the router itself, when the other interfaces to the same link cease to hear the router's own Hellos.

References

- [Ref1] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [Ref2] McKenzie, A., "ISO Transport Protocol specification ISO DP 8073", RFC 905, April 1984.
- [Ref3] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB using SMIV2", RFC 2233, November 1997.
- [Ref4] Fuller, V., Li, T, Yu, J. and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.
- [Ref5] Varadhan, K., Hares, S. and Y. Rekhter, "BGP4/IDRP for IP---OSPF Interaction", RFC 1745, December 1994
- [Ref6] Reynolds, J. and J. Postel, "Assigned Numbers", STD 2, RFC 1700, October 1994.
- [Ref7] deSouza, O. and M. Rodrigues, "Guidelines for Running OSPF Over Frame Relay Networks", RFC 1586, March 1994.
- [Ref8] Moy, J., "Multicast Extensions to OSPF", RFC 1584, March 1994.
- [Ref9] Coltun, R. and V. Fuller, "The OSPF NSSA Option", RFC 1587, March 1994.
- [Ref10] Ferguson, D., "The OSPF External Attributes LSA", unpublished.
- [Ref11] Moy, J., "Extending OSPF to Support Demand Circuits", RFC 1793, April 1995.
- [Ref12] Mogul, J. and S. Deering, "Path MTU Discovery", RFC 1191, November 1990.
- [Ref13] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [Ref14] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [Ref15] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.

- [Ref16] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification" RFC 2463, December 1998.
- [Ref17] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [Ref18] McCann, J., Deering, S. and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.
- [Ref19] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [Ref20] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

A. OSPF data formats

This appendix describes the format of OSPF protocol packets and OSPF LSAs. The OSPF protocol runs directly over the IPv6 network layer. Before any data formats are described, the details of the OSPF encapsulation are explained.

Next the OSPF Options field is described. This field describes various capabilities that may or may not be supported by pieces of the OSPF routing domain. The OSPF Options field is contained in OSPF Hello packets, Database Description packets and in OSPF LSAs.

OSPF packet formats are detailed in Section A.3.

A description of OSPF LSAs appears in Section A.4. This section describes how IPv6 address prefixes are represented within LSAs, details the standard LSA header, and then provides formats for each of the specific LSA types.

A.1 Encapsulation of OSPF packets

OSPF runs directly over the IPv6's network layer. OSPF packets are therefore encapsulated solely by IPv6 and local data-link headers.

OSPF does not define a way to fragment its protocol packets, and depends on IPv6 fragmentation when transmitting packets larger than the link MTU. If necessary, the length of OSPF packets can be up to 65,535 bytes. The OSPF packet types that are likely to be large (Database Description Packets, Link State Request, Link State Update, and Link State Acknowledgment packets) can usually be split into several separate protocol packets, without loss of functionality. This is recommended; IPv6 fragmentation should be avoided whenever possible. Using this reasoning, an attempt should be made to limit the sizes of OSPF packets sent over virtual links to 1280 bytes unless Path MTU Discovery is being performed [Ref14].

The other important features of OSPF's IPv6 encapsulation are:

- o Use of IPv6 multicast. Some OSPF messages are multicast, when sent over broadcast networks. Two distinct IP multicast addresses are used. Packets sent to these multicast addresses should never be forwarded; they are meant to travel a single hop only. As such, the multicast addresses have been chosen with link-local scope, and packets sent to these addresses should have their IPv6 Hop Limit set to 1.

AllSPFRouters

This multicast address has been assigned the value FF02::5. All routers running OSPF should be prepared to receive packets sent to this address. Hello packets are always sent to this destination. Also, certain OSPF protocol packets are sent to this address during the flooding procedure.

AllDRouters

This multicast address has been assigned the value FF02::6. Both the Designated Router and Backup Designated Router must be prepared to receive packets destined to this address. Certain OSPF protocol packets are sent to this address during the flooding procedure.


- o OSPF is IP protocol 89. This number should be inserted in the Next Header field of the encapsulating IPv6 header.

A.2 The Options field

The 24-bit OSPF Options field is present in OSPF Hello packets, Database Description packets and certain LSAs (router-LSAs, network-LSAs, inter-area-router-LSAs and link-LSAs). The Options field enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers. Through this mechanism routers of differing capabilities can be mixed within an OSPF routing domain.

An option mismatch between routers can cause a variety of behaviors, depending on the particular option. Some option mismatches prevent neighbor relationships from forming (e.g., the E-bit below); these mismatches are discovered through the sending and receiving of Hello packets. Some option mismatches prevent particular LSA types from being flooded across adjacencies (e.g., the MC-bit below); these are discovered through the sending and receiving of Database Description packets. Some option mismatches prevent routers from being included in one or more of the various routing calculations because of their reduced functionality (again the MC-bit is an example); these mismatches are discovered by examining LSAs.

Six bits of the OSPF Options field have been assigned. Each bit is described briefly below. Routers should reset (i.e. clear) unrecognized bits in the Options field when sending Hello packets or Database Description packets and when originating LSAs. Conversely, routers encountering unrecognized Option bits in received Hello Packets, Database Description packets or LSAs should ignore the capability and process the packet/LSA normally.



The Options field

v6-bit

If this bit is clear, the router/link should be excluded from IPv6 routing calculations. See Section 3.8 of this memo.

E-bit

This bit describes the way AS-external-LSAs are flooded, as described in Sections 3.6, 9.5, 10.8 and 12.1.2 of [Ref1].

MC-bit

This bit describes whether IP multicast datagrams are forwarded according to the specifications in [Ref7].

N-bit

This bit describes the handling of Type-7 LSAs, as specified in [Ref8].

R-bit

This bit (the 'Router' bit) indicates whether the originator is an active router. If the router bit is clear routes which transit the advertising node cannot be computed. Clearing the router bit would be appropriate for a multi-homed host that wants to participate in routing, but does not want to forward non-locally addressed packets.

DC-bit

This bit describes the router's handling of demand circuits, as specified in [Ref10].

A.3 OSPF Packet Formats

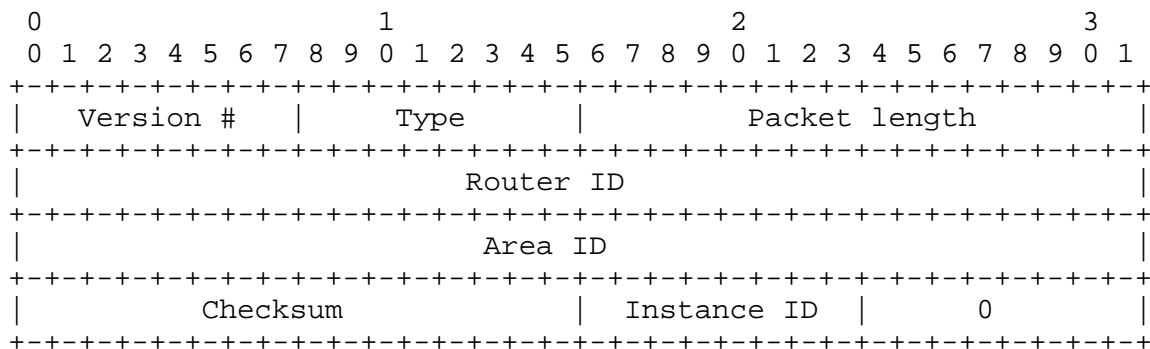
There are five distinct OSPF packet types. All OSPF packet types begin with a standard 16 byte header. This header is described first. Each packet type is then described in a succeeding section. In these sections each packet's division into fields is displayed, and then the field definitions are enumerated.

All OSPF packet types (other than the OSPF Hello packets) deal with lists of LSAs. For example, Link State Update packets implement the flooding of LSAs throughout the OSPF routing domain. The format of LSAs is described in Section A.4.

The receive processing of OSPF packets is detailed in Section 3.2.2. The sending of OSPF packets is explained in Section 3.2.1.

A.3.1 The OSPF packet header

Every OSPF packet starts with a standard 16 byte header. Together with the encapsulating IPv6 headers, the OSPF header contains all the information necessary to determine whether the packet should be accepted for further processing. This determination is described in Section 3.2.2 of this memo.



Version

The OSPF version number. This specification documents version 3 of the OSPF protocol.

Type

The OSPF packet types are as follows. See Sections A.3.2 through A.3.6 for details.

| Type | Description |
|-------|---------------------------|
| ----- | |
| 1 | Hello |
| 2 | Database Description |
| 3 | Link State Request |
| 4 | Link State Update |
| 5 | Link State Acknowledgment |

Packet length

The length of the OSPF protocol packet in bytes. This length includes the standard OSPF header.

Router ID

The Router ID of the packet's source.

Area ID

A 32 bit number identifying the area that this packet belongs to. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.

Checksum

OSPF uses the standard checksum calculation for IPv6 applications: The 16-bit one's complement of the one's complement sum of the entire contents of the packet, starting with the OSPF packet header, and prepending a "pseudo-header" of IPv6 header fields, as specified in [Ref14, section 8.1]. The "Upper-Layer Packet Length" in the pseudo-header is set to value of the OSPF packet header's length field. The Next Header value used in the pseudo-header is 89. If the packet's length is not an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. Before computing the checksum, the checksum field in the OSPF packet header is set to 0.

Instance ID


Enables multiple instances of OSPF to be run over a single link. Each protocol instance would be assigned a separate Instance ID; the Instance ID has local link significance only. Received packets whose Instance ID is not equal to the receiving interface's Instance ID are discarded.

0 These fields are reserved. They must be 0.

A.3.2 The Hello packet

Hello packets are OSPF packet type 1. These packets are sent periodically on all interfaces (including virtual links) in order to establish and maintain neighbor relationships. In addition, Hello Packets are multicast on those links having a multicast or broadcast capability, enabling dynamic discovery of neighboring routers.

All routers connected to a common link must agree on certain parameters (HelloInterval and RouterDeadInterval). These parameters are included in Hello packets, so that differences can inhibit the forming of neighbor relationships. The Hello packet also contains fields used in Designated Router election (Designated Router ID and Backup Designated Router ID), and fields used to detect bi-directionality (the Router IDs of all neighbors whose Hellos have been recently received).

**Interface ID**

32-bit number uniquely identifying this interface among the collection of this router's interfaces. For example, in some implementations it may be possible to use the MIB-II IfIndex ([Ref3]).

Rtr Pri

This router's Router Priority. Used in (Backup) Designated Router election. If set to 0, the router will be ineligible to become (Backup) Designated Router.

Options

The optional capabilities supported by the router, as documented in Section A.2.

HelloInterval

The number of seconds between this router's Hello packets.

RouterDeadInterval

The number of seconds before declaring a silent router down.

Designated Router ID

The identity of the Designated Router for this network, in the view of the sending router. The Designated Router is identified by its Router ID. Set to 0.0.0.0 if there is no Designated Router.

Backup Designated Router ID


The identity of the Backup Designated Router for this network, in the view of the sending router. The Backup Designated Router is identified by its IP Router ID. Set to 0.0.0.0 if there is no Backup Designated Router.

Neighbor ID

The Router IDs of each router from whom valid Hello packets have been seen recently on the network. Recently means in the last RouterDeadInterval seconds.

A.3.3 The Database Description packet

Database Description packets are OSPF packet type 2. These packets are exchanged when an adjacency is being initialized. They describe the contents of the link-state database. Multiple packets may be used to describe the database. For this purpose a poll-response procedure is used. One of the routers is designated to be the master, the other the slave. The master sends Database Description packets (polls) which are acknowledged by Database Description packets sent by the slave (responses). The responses are linked to the polls via the packets' DD sequence numbers.



The format of the Database Description packet is very similar to both the Link State Request and Link State Acknowledgment packets. The main part of all three is a list of items, each item describing a piece of the link-state database. The sending of Database Description Packets is documented in Section 10.8 of [Ref1]. The reception of Database Description packets is documented in Section 10.6 of [Ref1].

Options

The optional capabilities supported by the router, as documented in Section A.2.

Interface MTU

The size in bytes of the largest IPv6 datagram that can be sent out the associated interface, without fragmentation. The MTUs of common Internet link types can be found in Table 7-1 of [Ref12]. Interface MTU should be set to 0 in Database Description packets sent over virtual links.

I-bit

The Init bit. When set to 1, this packet is the first in the sequence of Database Description Packets.

M-bit

The More bit. When set to 1, it indicates that more Database Description Packets are to follow.

MS-bit

The Master/Slave bit. When set to 1, it indicates that the router is the master during the Database Exchange process. Otherwise, the router is the slave.

DD sequence number

Used to sequence the collection of Database Description Packets. The initial value (indicated by the Init bit being set) should be unique. The DD sequence number then increments until the complete database description has been sent.


The rest of the packet consists of a (possibly partial) list of the link-state database's pieces. Each LSA in the database is described by its LSA header. The LSA header is documented in Section A.4.1. It contains all the information required to uniquely identify both the LSA and the LSA's current instance.

A.3.4 The Link State Request packet

Link State Request packets are OSPF packet type 3. After exchanging Database Description packets with a neighboring router, a router may find that parts of its link-state database are out-of-date. The Link State Request packet is used to request the pieces of the neighbor's database that are more up-to-date. Multiple Link State Request packets may need to be used.

A router that sends a Link State Request packet has in mind the precise instance of the database pieces it is requesting. Each instance is defined by its LS sequence number, LS checksum, and LS age, although these fields are not specified in the Link State Request Packet itself. The router may receive even more recent instances in response.

The sending of Link State Request packets is documented in Section 10.9 of [Ref1]. The reception of Link State Request packets is documented in Section 10.7 of [Ref1].




Each LSA requested is specified by its LS type, Link State ID, and Advertising Router. This uniquely identifies the LSA, but not its instance. Link State Request packets are understood to be requests for the most recent instance (whatever that might be).

A.3.5 The Link State Update packet

Link State Update packets are OSPF packet type 4. These packets implement the flooding of LSAs. Each Link State Update packet carries a collection of LSAs one hop further from their origin. Several LSAs may be included in a single packet.

Link State Update packets are multicast on those physical networks that support multicast/broadcast. In order to make the flooding procedure reliable, flooded LSAs are acknowledged in Link State Acknowledgment packets. If retransmission of certain LSAs is necessary, the retransmitted LSAs are always carried by unicast Link State Update packets. For more information on the reliable flooding of LSAs, consult Section 3.5.



LSAs

The number of LSAs included in this update.


The body of the Link State Update packet consists of a list of LSAs. Each LSA begins with a common 20 byte header, described in Section A.4.2. Detailed formats of the different types of LSAs are described in Section A.4.

A.3.6 The Link State Acknowledgment packet

Link State Acknowledgment Packets are OSPF packet type 5. To make the flooding of LSAs reliable, flooded LSAs are explicitly acknowledged. This acknowledgment is accomplished through the sending and receiving of Link State Acknowledgment packets. The sending of Link State Acknowledgement packets is documented in Section 13.5 of [Ref1]. The reception of Link State Acknowledgement packets is documented in Section 13.7 of [Ref1].

Multiple LSAs can be acknowledged in a single Link State Acknowledgment packet. Depending on the state of the sending interface and the sender of the corresponding Link State Update packet, a Link State Acknowledgment packet is sent either to the multicast address AllSPFRouters, to the multicast address AllDRouters, or as a unicast (see Section 13.5 of [Ref1] for details).

The format of this packet is similar to that of the Data Description packet. The body of both packets is simply a list of LSA headers.



Each acknowledged LSA is described by its LSA header. The LSA header is documented in Section A.4.2. It contains all the information required to uniquely identify both the LSA and the LSA's current instance.

A.4 LSA formats

This memo defines seven distinct types of LSAs. Each LSA begins with a standard 20 byte LSA header. This header is explained in Section A.4.2. Succeeding sections then diagram the separate LSA types.

Each LSA describes a piece of the OSPF routing domain. Every router originates a router-LSA. A network-LSA is advertised for each link by its Designated Router. A router's link-local addresses are advertised to its neighbors in link-LSAs. IPv6 prefixes are advertised in intra-area-prefix-LSAs, inter-area-prefix-LSAs and AS-external-LSAs. Location of specific routers can be advertised across area boundaries in inter-area-router-LSAs. All LSAs are then flooded throughout the OSPF routing domain. The flooding algorithm is reliable, ensuring that all routers have the same collection of LSAs. (See Section 3.5 for more information concerning the flooding algorithm). This collection of LSAs is called the link-state database.

From the link state database, each router constructs a shortest path tree with itself as root. This yields a routing table (see Section 11 of [Ref1]). For the details of the routing table build process, see Section 3.8.

A.4.1 IPv6 Prefix Representation

IPv6 addresses are bit strings of length 128. IPv6 routing algorithms, and OSPF for IPv6 in particular, advertise IPv6 address prefixes. IPv6 address prefixes are bit strings whose length ranges between 0 and 128 bits (inclusive).

Within OSPF, IPv6 address prefixes are always represented by a combination of three fields: PrefixLength, PrefixOptions, and Address Prefix. PrefixLength is the length in bits of the prefix. PrefixOptions is an 8-bit field describing various capabilities associated with the prefix (see Section A.4.2). Address Prefix is an encoding of the prefix itself as an even multiple of 32-bit words, padding with zero bits as necessary; this encoding consumes $(\text{PrefixLength} + 31) / 32$ 32-bit words.

The default route is represented by a prefix of length 0.

Examples of IPv6 Prefix representation in OSPF can be found in Sections A.4.5, A.4.7, A.4.8 and A.4.9.

A.4.1.1 Prefix Options

Each prefix is advertised along with an 8-bit field of capabilities. These serve as input to the various routing calculations, allowing, for example, certain prefixes to be ignored in some cases, or to be marked as not readvertisable in others.

```

      0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+
|   |   |   |   |   | P|MC|LA|NU|
+---+---+---+---+---+---+---+
```

The Prefix Options field

NU-bit

The "no unicast" capability bit. If set, the prefix should be excluded from IPv6 unicast calculations, otherwise it should be included.

LA-bit

The "local address" capability bit. If set, the prefix is actually an IPv6 interface address of the advertising router.

MC-bit


The "multicast" capability bit. If set, the prefix should be included in IPv6 multicast routing calculations, otherwise it should be excluded.

P-bit

The "propagate" bit. Set on NSSA area prefixes that should be readvertised at the NSSA area border.

A.4.2 The LSA header

All LSAs begin with a common 20 byte header. This header contains enough information to uniquely identify the LSA (LS type, Link State ID, and Advertising Router). Multiple instances of the LSA may exist in the routing domain at the same time. It is then necessary to determine which instance is more recent. This is accomplished by examining the LS age, LS sequence number and LS checksum fields that are also contained in the LSA header.

**LS age**

The time in seconds since the LSA was originated.

LS type

The LS type field indicates the function performed by the LSA. The high-order three bits of LS type encode generic properties of the LSA, while the remainder (called LSA function code) indicate the LSA's specific functionality. See Section A.4.2.1 for a detailed description of LS type.

Link State ID

Together with LS type and Advertising Router, uniquely identifies the LSA in the link-state database.

Advertising Router

The Router ID of the router that originated the LSA. For example, in network-LSAs this field is equal to the Router ID of the network's Designated Router.

LS sequence number

Detects old or duplicate LSAs. Successive instances of an LSA are given successive LS sequence numbers. See Section 12.1.6 in [Ref1] for more details.

LS checksum

The Fletcher checksum of the complete contents of the LSA, including the LSA header but excluding the LS age field. See Section 12.1.7 in [Ref1] for more details.

length

The length in bytes of the LSA. This includes the 20 byte LSA header.

A.4.2.1 LS type

The LS type field indicates the function performed by the LSA. The high-order three bits of LS type encode generic properties of the LSA, while the remainder (called LSA function code) indicate the LSA's specific functionality. The format of the LS type is as follows:

```

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

The U bit indicates how the LSA should be handled by a router which does not recognize the LSA's function code. Its values are:

| U-bit | LSA Handling |
|-------|--|
| 0 | Treat the LSA as if it had link-local flooding scope |
| 1 | Store and flood the LSA, as if type understood |

The S1 and S2 bits indicate the flooding scope of the LSA. The values are:


| S2 | S1 | Flooding Scope |
|----|----|---|
| 0 | 0 | Link-Local Scoping. Flooded only on link it is originated on. |
| 0 | 1 | Area Scoping. Flooded to all routers in the originating area |
| 1 | 0 | AS Scoping. Flooded to all routers in the AS |
| 1 | 1 | Reserved |

The LSA function codes are defined as follows. The origination and processing of these LSA function codes are defined elsewhere in this memo, except for the group-membership-LSA (see [Ref7]) and the Type-7-LSA (see [Ref8]). Each LSA function code also implies a specific setting for the U, S1 and S2 bits, as shown below.

| LSA function code | LS Type | Description |
|-------------------|---------|-----------------------|
| 1 | 0x2001 | Router-LSA |
| 2 | 0x2002 | Network-LSA |
| 3 | 0x2003 | Inter-Area-Prefix-LSA |
| 4 | 0x2004 | Inter-Area-Router-LSA |
| 5 | 0x4005 | AS-External-LSA |
| 6 | 0x2006 | Group-membership-LSA |
| 7 | 0x2007 | Type-7-LSA |
| 8 | 0x0008 | Link-LSA |
| 9 | 0x2009 | Intra-Area-Prefix-LSA |

A.4.3 Router-LSAs

Router-LSAs have LS type equal to 0x2001. Each router in an area originates one or more router-LSAs. The complete collection of router-LSAs originated by the router describe the state and cost of the router's interfaces to the area. For details concerning the construction of router-LSAs, see Section 3.4.3.1. Router-LSAs are flooded throughout a single area only.



A single router may originate one or more Router LSAs, distinguished by their Link-State IDs (which are chosen arbitrarily by the originating router). The Options field and V, E and B bits should be the same in all Router LSAs from a single originator. However, in the case of a mismatch the values in the LSA with the lowest Link State ID take precedence. When more than one Router LSA is received from a single router, the links are processed as if concatenated into a single LSA.

bit V

When set, the router is an endpoint of one or more fully adjacent virtual links having the described area as Transit area (V is for virtual link endpoint).

bit E

When set, the router is an AS boundary router (E is for external).

bit B

When set, the router is an area border router (B is for border).

bit W

When set, the router is a wild-card multicast receiver. When running MOSPF, these routers receive all multicast datagrams, regardless of destination. See Sections 3, 4 and A.2 of [Ref8] for details.

Options

The optional capabilities supported by the router, as documented in Section A.2.

The following fields are used to describe each router interface. The Type field indicates the kind of interface being described. It may be an interface to a transit network, a point-to-point connection to another router or a virtual link. The values of all the other fields describing a router interface depend on the interface's Type field.

Type

The kind of interface being described. One of the following:

| Type | Description |
|-------|---|
| ----- | |
| 1 | Point-to-point connection to another router |
| 2 | Connection to a transit network |
| 3 | Reserved |
| 4 | Virtual link |

Metric

The cost of using this router interface, for outbound traffic.

Interface ID

The Interface ID assigned to the interface being described. See Sections 3.1.2 and C.3.

Neighbor Interface ID

The Interface ID the neighbor router (or the attached link's Designated Router, for Type 2 interfaces) has been advertising in hello packets sent on the attached link.

Neighbor Router ID

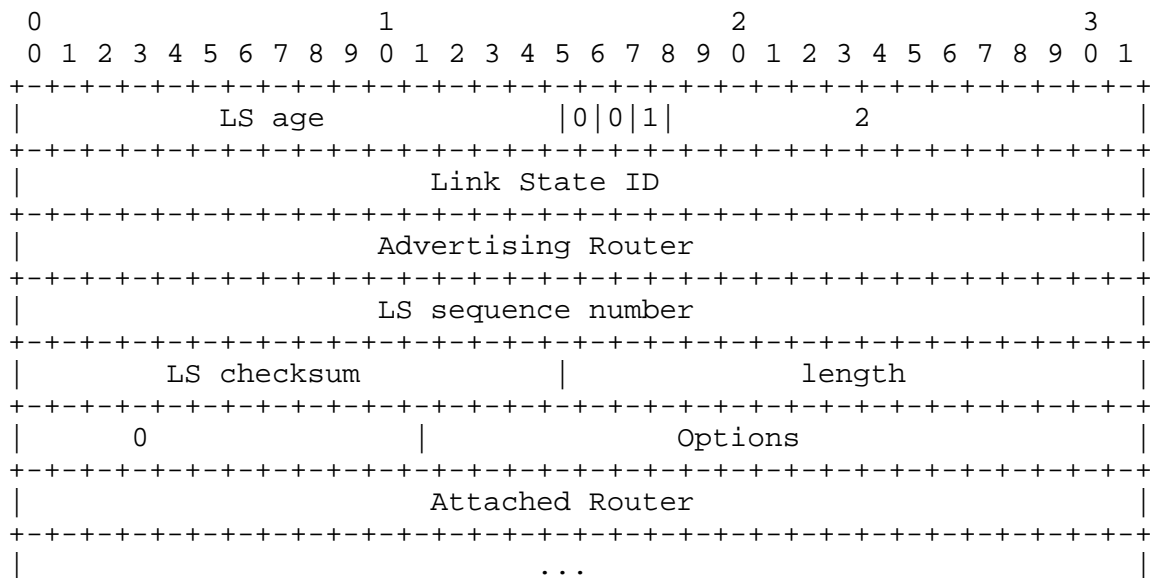
The Router ID the neighbor router (or the attached link's Designated Router, for Type 2 interfaces).

For Type 2 links, the combination of Neighbor Interface ID and Neighbor Router ID allows the network-LSA for the attached link to be found in the link-state database.

A.4.4 Network-LSAs

Network-LSAs have LS type equal to 0x2002. A network-LSA is originated for each broadcast and NBMA link in the area which supports two or more routers. The network-LSA is originated by the link's Designated Router. The LSA describes all routers attached to the link, including the Designated Router itself. The LSA's Link State ID field is set to the Interface ID that the Designated Router has been advertising in Hello packets on the link.

The distance from the network to all attached routers is zero. This is why the metric fields need not be specified in the network-LSA. For details concerning the construction of network-LSAs, see Section 3.4.3.2.



Attached Router

The Router IDs of each of the routers attached to the link. Actually, only those routers that are fully adjacent to the Designated Router are listed. The Designated Router includes itself in this list. The number of routers included can be deduced from the LSA header's length field.

A.4.5 Inter-Area-Prefix-LSAs

Inter-Area-Prefix-LSAs have LS type equal to 0x2003. These LSAs are the IPv6 equivalent of OSPF for IPv4's type 3 summary-LSAs (see Section 12.4.3 of [Ref1]). Originated by area border routers, they describe routes to IPv6 address prefixes that belong to other areas. A separate Inter-Area-Prefix-LSA is originated for each IPv6 address prefix. For details concerning the construction of Inter-Area-Prefix-LSAs, see Section 3.4.3.3.

For stub areas, Inter-Area-Prefix-LSAs can also be used to describe a (per-area) default route. Default summary routes are used in stub areas instead of flooding a complete set of external routes. When describing a default summary route, the Inter-Area-Prefix-LSA's PrefixLength is set to 0.

| 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 |
| LS age | | | | | | | | | | 0 0 1 | | | | | | | | | | 3 | | | | | | | | | | | | | | | | | | | |
| Link State ID | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Advertising Router | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LS sequence number | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LS checksum | | | | | | | | | | | | | | | | | | | | length | | | | | | | | | | | | | | | | | | | |
| 0 | | | | | | | | | | | | | | | | | | | | Metric | | | | | | | | | | | | | | | | | | | |
| PrefixLength | | | | | | | | | | PrefixOptions | | | | | | | | | | (0) | | | | | | | | | | | | | | | | | | | |
| Address Prefix | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| ... | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Metric

The cost of this route. Expressed in the same units as the interface costs in the router-LSAs. When the Inter-Area-Prefix-LSA is describing a route to a range of addresses (see Section C.2) the cost is set to the maximum cost to any reachable component of the address range.

PrefixLength, PrefixOptions and Address Prefix

Representation of the IPv6 address prefix, as described in Section A.4.1.

A.4.6 Inter-Area-Router-LSAs

Inter-Area-Router-LSAs have LS type equal to 0x2004. These LSAs are the IPv6 equivalent of OSPF for IPv4's type 4 summary-LSAs (see Section 12.4.3 of [Ref1]). Originated by area border routers, they describe routes to routers in other areas. (To see why it is necessary to advertise the location of each ASBR, consult Section 16.4 in [Ref1].) Each LSA describes a route to a single router. For details concerning the construction of Inter-Area-Router-LSAs, see Section 3.4.3.4.

| 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 |
| LS age | | | | | | | | | | 0 0 1 | | | | | | | | | | 4 | | | | | | | | | | | | | | | | | | | |
| Link State ID | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Advertising Router | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LS sequence number | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LS checksum | | | | | | | | | | | | | | | | | | | | length | | | | | | | | | | | | | | | | | | | |
| 0 | | | | | | | | | | | | | | | | | | | | Options | | | | | | | | | | | | | | | | | | | |
| 0 | | | | | | | | | | | | | | | | | | | | Metric | | | | | | | | | | | | | | | | | | | |
| Destination Router ID | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Options

The optional capabilities supported by the router, as documented in Section A.2.

Metric

The cost of this route. Expressed in the same units as the interface costs in the router-LSAs.

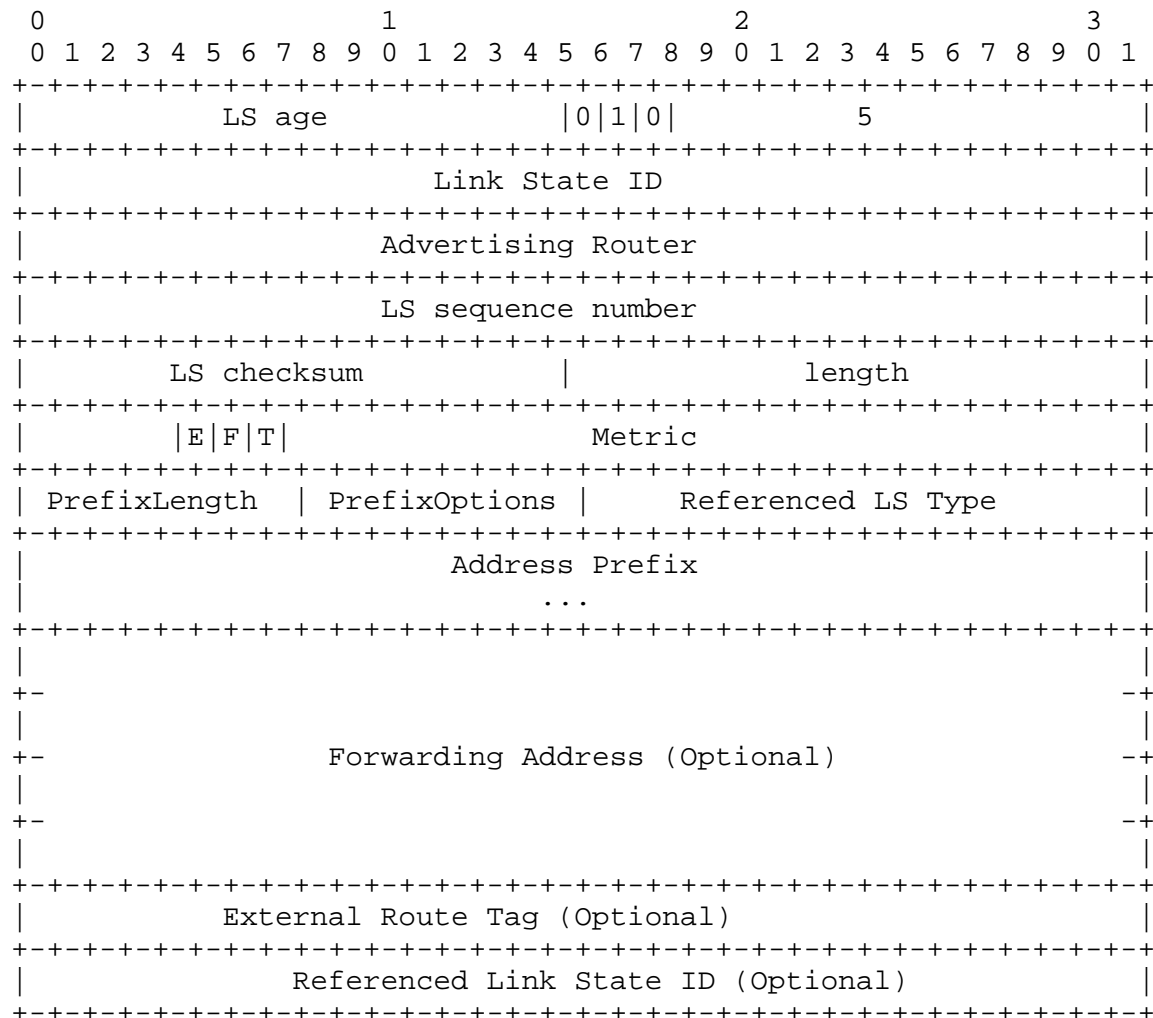
Destination Router ID

The Router ID of the router being described by the LSA.

A.4.7 AS-external-LSAs

AS-external-LSAs have LS type equal to 0x4005. These LSAs are originated by AS boundary routers, and describe destinations external to the AS. Each LSA describes a route to a single IPv6 address prefix. For details concerning the construction of AS-external-LSAs, see Section 3.4.3.5.

AS-external-LSAs can be used to describe a default route. Default routes are used when no specific route exists to the destination. When describing a default route, the AS-external-LSA's PrefixLength is set to 0.



bit E

The type of external metric. If bit E is set, the metric specified is a Type 2 external metric. This means the metric is considered larger than any intra-AS path. If bit E is zero, the specified metric is a Type 1 external metric. This means that it is expressed in the same units as the link state metric (i.e., the same units as interface cost).

bit F

If set, a Forwarding Address has been included in the LSA.

bit T

If set, an External Route Tag has been included in the LSA.

Metric

The cost of this route. Interpretation depends on the external type indication (bit E above).

PrefixLength, PrefixOptions and Address Prefix

Representation of the IPv6 address prefix, as described in Section A.4.1.

Referenced LS type

If non-zero, an LSA with this LS type is to be associated with this LSA (see Referenced Link State ID below).

Forwarding address

A fully qualified IPv6 address (128 bits). Included in the LSA if and only if bit F has been set. If included, Data traffic for the advertised destination will be forwarded to this address. Must not be set to the IPv6 Unspecified Address (0:0:0:0:0:0:0:0).

External Route Tag

A 32-bit field which may be used to communicate additional information between AS boundary routers; see [Ref5] for example usage. Included in the LSA if and only if bit T has been set.

Referenced Link State ID Included if and only if Reference LS Type is non-zero. If included, additional information concerning the advertised external route can be found in the LSA having LS type equal to "Referenced LS Type", Link State ID equal to "Referenced Link State ID" and Advertising Router the same as that specified in the AS-external-LSA's link state header. This additional information is not used by the OSPF protocol itself. It may be used to communicate information between AS boundary routers; the precise nature of such information is outside the scope of this specification.

All, none or some of the fields labeled Forwarding address, External Route Tag and Referenced Link State ID may be present in the AS-external-LSA (as indicated by the setting of bit F, bit T and Referenced LS type respectively). However, when present Forwarding Address always comes first, with External Route Tag always preceding Referenced Link State ID.

A.4.8 Link-LSAs

Link-LSAs have LS type equal to 0x0008. A router originates a separate Link-LSA for each link it is attached to. These LSAs have local-link flooding scope; they are never flooded beyond the link that they are associated with. Link-LSAs have three purposes: 1) they provide the router's link-local address to all other routers attached to the link and 2) they inform other routers attached to the link of a list of IPv6 prefixes to associate with the link and 3) they allow the router to assert a collection of Options bits to associate with the Network-LSA that will be originated for the link.

A link-LSA's Link State ID is set equal to the originating router's Interface ID on the link.

**Rtr Pri**

The Router Priority of the interface attaching the originating router to the link.

Options

The set of Options bits that the router would like set in the Network-LSA that will be originated for the link.

Link-local Interface Address

The originating router's link-local interface address on the link.

prefixes

The number of IPv6 address prefixes contained in the LSA.

The rest of the link-LSA contains a list of IPv6 prefixes to be associated with the link.

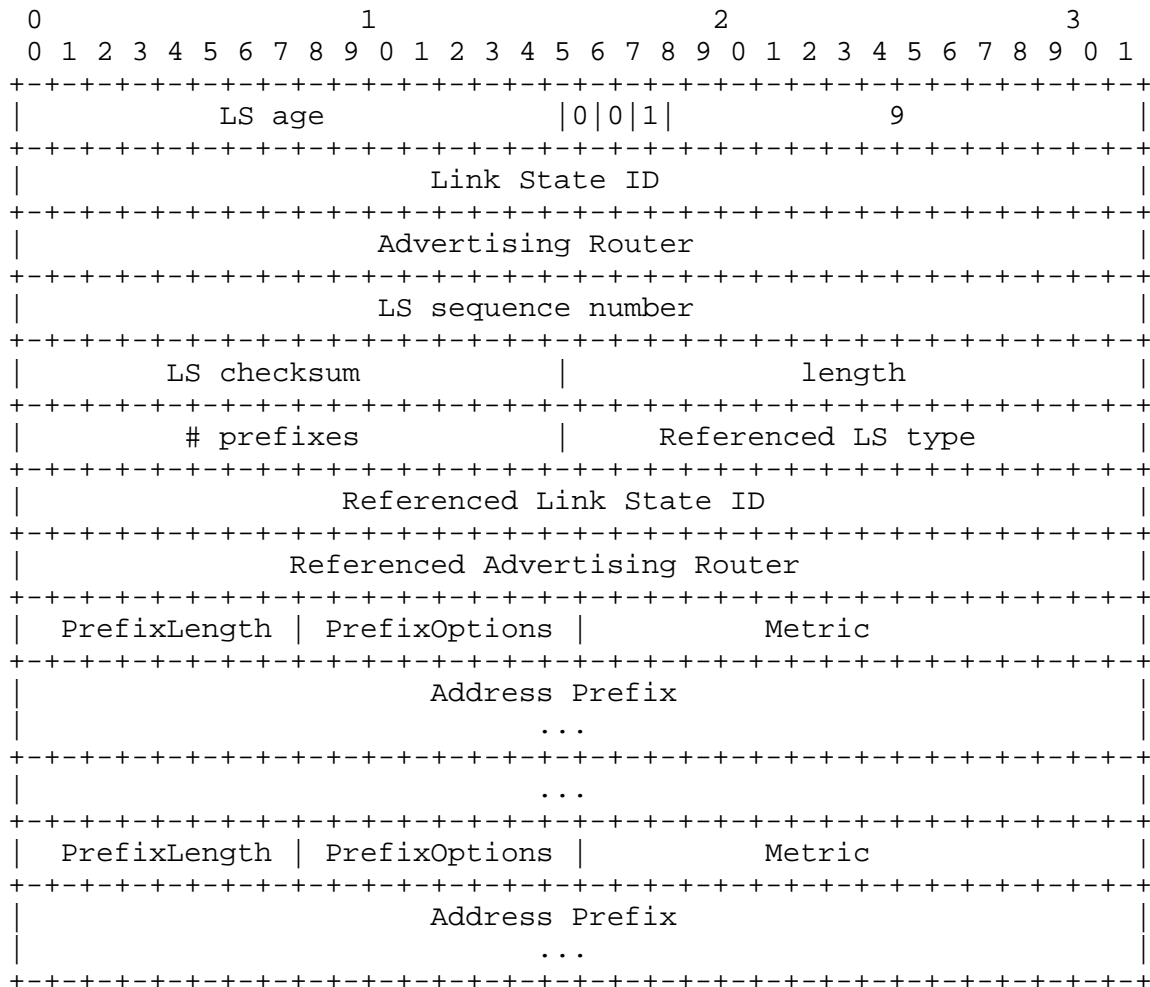
PrefixLength, PrefixOptions and Address Prefix

Representation of an IPv6 address prefix, as described in Section A.4.1.

A.4.9 Intra-Area-Prefix-LSAs

Intra-Area-Prefix-LSAs have LS type equal to 0x2009. A router uses Intra-Area-Prefix-LSAs to advertise one or more IPv6 address prefixes that are associated with a) the router itself, b) an attached stub network segment or c) an attached transit network segment. In IPv4, a) and b) were accomplished via the router's router-LSA, and c) via a network-LSA. However, in OSPF for IPv6 all addressing information has been removed from router-LSAs and network-LSAs, leading to the introduction of the Intra-Area-Prefix-LSA.

A router can originate multiple Intra-Area-Prefix-LSAs for each router or transit network; each such LSA is distinguished by its Link State ID.



prefixes

The number of IPv6 address prefixes contained in the LSA.

Router

Referenced LS type, Referenced Link State ID and Referenced Advertising

Identifies the router-LSA or network-LSA with which the IPv6 address prefixes should be associated. If Referenced LS type is 1, the prefixes are associated with a router-LSA, Referenced Link State ID should be 0 and Referenced Advertising Router should be the originating router's Router ID. If Referenced LS type is 2, the prefixes are associated with a network-LSA, Referenced Link State ID should be the Interface ID of the link's Designated Router and Referenced Advertising Router should be the Designated Router's Router ID.

The rest of the Intra-Area-Prefix-LSA contains a list of IPv6 prefixes to be associated with the router or transit link, together with the cost of each prefix.

PrefixLength, PrefixOptions and Address Prefix

Representation of an IPv6 address prefix, as described in Section A.4.1.

Metric

The cost of this prefix. Expressed in the same units as the interface costs in the router-LSAs.

B. Architectural Constants

Architectural constants for the OSPF protocol are defined in Appendix B of [Ref1]. The only difference for OSPF for IPv6 is that DefaultDestination is encoded as a prefix of length 0 (see Section A.4.1).

C. Configurable Constants

The OSPF protocol has quite a few configurable parameters. These parameters are listed below. They are grouped into general functional categories (area parameters, interface parameters, etc.). Sample values are given for some of the parameters.

Some parameter settings need to be consistent among groups of routers. For example, all routers in an area must agree on that area's parameters, and all routers attached to a network must agree on that network's HelloInterval and RouterDeadInterval.

Some parameters may be determined by router algorithms outside of this specification (e.g., the address of a host connected to the router via a SLIP line). From OSPF's point of view, these items are still configurable.

C.1 Global parameters

In general, a separate copy of the OSPF protocol is run for each area. Because of this, most configuration parameters are defined on a per-area basis. The few global configuration parameters are listed below.

Router ID

This is a 32-bit number that uniquely identifies the router in the Autonomous System. If a router's OSPF Router ID is changed, the router's OSPF software should be restarted before the new Router ID takes effect. Before restarting in order to change its Router

ID, the router should flush its self-originated LSAs from the routing domain (see Section 14.1 of [Ref1]), or they will persist for up to MaxAge minutes.

Because the size of the Router ID is smaller than an IPv6 address, it cannot be set to one of the router's IPv6 addresses (as is commonly done for IPv4). Possible Router ID assignment procedures for IPv6 include: a) assign the IPv6 Router ID as one of the router's IPv4 addresses or b) assign IPv6 Router IDs through some local administrative procedure (similar to procedures used by manufacturers to assign product serial numbers).

The Router ID of 0.0.0.0 is reserved, and should not be used.

C.2 Area parameters

All routers belonging to an area must agree on that area's configuration. Disagreements between two routers will lead to an inability for adjacencies to form between them, with a resulting hindrance to the flow of routing protocol and data traffic. The following items must be configured for an area:

Area ID

This is a 32-bit number that identifies the area. The Area ID of 0 is reserved for the backbone.

List of address ranges

Address ranges control the advertisement of routes across area boundaries. Each address range consists of the following items:

[IPv6 prefix, prefix length]

Describes the collection of IPv6 addresses contained in the address range.

Status Set to either Advertise or DoNotAdvertise. Routing information is condensed at area boundaries. External to the area, at most a single route is advertised (via a inter-area-prefix-LSA) for each address range. The route is advertised if and only if the address range's Status is set to Advertise. Unadvertised ranges allow the existence of certain networks to be intentionally hidden from other areas. Status is set to Advertise by default.

ExternalRoutingCapability

Whether AS-external-LSAs will be flooded into/throughout the area. If AS-external-LSAs are excluded from the area, the area is called a "stub". Internal to stub areas, routing to external destinations will be based solely on a default inter-area route. The backbone cannot be configured as a stub area. Also, virtual links cannot be configured through stub areas. For more information, see Section 3.6 of [Ref1].

StubDefaultCost

If the area has been configured as a stub area, and the router itself is an area border router, then the StubDefaultCost indicates the cost of the default inter-area-prefix-LSA that the router should advertise into the area. See Section 12.4.3.1 of [Ref1] for more information.

C.3 Router interface parameters

Some of the configurable router interface parameters (such as Area ID, HelloInterval and RouterDeadInterval) actually imply properties of the attached links, and therefore must be consistent across all the routers attached to that link. The parameters that must be configured for a router interface are:

IPv6 link-local address

The IPv6 link-local address associated with this interface. May be learned through auto-configuration.

Area ID

The OSPF area to which the attached link belongs.

Instance ID

The OSPF protocol instance associated with this OSPF interface. Defaults to 0.

Interface ID

32-bit number uniquely identifying this interface among the collection of this router's interfaces. For example, in some implementations it may be possible to use the MIB-II IfIndex ([Ref3]).

IPv6 prefixes

The list of IPv6 prefixes to associate with the link. These will be advertised in intra-area-prefix-LSAs.

Interface output cost(s)

The cost of sending a packet on the interface, expressed in the link state metric. This is advertised as the link cost for this interface in the router's router-LSA. The interface output cost must always be greater than 0.

RxmtInterval

The number of seconds between LSA retransmissions, for adjacencies belonging to this interface. Also used when retransmitting Database Description and Link State Request Packets. This should be well over the expected round-trip delay between any two routers on the attached link. The setting of this value should be conservative or needless retransmissions will result. Sample value for a local area network: 5 seconds.

InfTransDelay

The estimated number of seconds it takes to transmit a Link State Update Packet over this interface. LSAs contained in the update packet must have their age incremented by this amount before transmission. This value should take into account the transmission and propagation delays of the interface. It must be greater than 0. Sample value for a local area network: 1 second.

Router Priority

An 8-bit unsigned integer. When two routers attached to a network both attempt to become Designated Router, the one with the highest Router Priority takes precedence. If there is still a tie, the router with the highest Router ID takes precedence. A router whose Router Priority is set to 0 is ineligible to become Designated Router on the attached link. Router Priority is only configured for interfaces to broadcast and NBMA networks.

HelloInterval

The length of time, in seconds, between the Hello Packets that the router sends on the interface. This value is advertised in the router's Hello Packets. It must be the same for all routers attached to a common link. The smaller the HelloInterval, the faster topological changes will be detected; however, more OSPF routing protocol traffic will ensue. Sample value for a X.25 PDN: 30 seconds. Sample value for a local area network (LAN): 10 seconds.

RouterDeadInterval

After ceasing to hear a router's Hello Packets, the number of seconds before its neighbors declare the router down. This is also advertised in the router's Hello Packets in their

RouterDeadInterval field. This should be some multiple of the HelloInterval (say 4). This value again must be the same for all routers attached to a common link.

C.4 Virtual link parameters

Virtual links are used to restore/increase connectivity of the backbone. Virtual links may be configured between any pair of area border routers having interfaces to a common (non-backbone) area. The virtual link appears as an unnumbered point-to-point link in the graph for the backbone. The virtual link must be configured in both of the area border routers.

A virtual link appears in router-LSAs (for the backbone) as if it were a separate router interface to the backbone. As such, it has most of the parameters associated with a router interface (see Section C.3). Virtual links do not have link-local addresses, but instead use one of the router's global-scope or site-local IPv6 addresses as the IP source in OSPF protocol packets it sends along the virtual link. Router Priority is not used on virtual links. Interface output cost is not configured on virtual links, but is dynamically set to be the cost of the intra-area path between the two endpoint routers. The parameter RxmtInterval must be configured, and should be well over the expected round-trip delay between the two routers. This may be hard to estimate for a virtual link; it is better to err on the side of making it too large.

A virtual link is defined by the following two configurable parameters: the Router ID of the virtual link's other endpoint, and the (non-backbone) area through which the virtual link runs (referred to as the virtual link's Transit area). Virtual links cannot be configured through stub areas.

C.5 NBMA network parameters

OSPF treats an NBMA network much like it treats a broadcast network. Since there may be many routers attached to the network, a Designated Router is selected for the network. This Designated Router then originates a network-LSA, which lists all routers attached to the NBMA network.

However, due to the lack of broadcast capabilities, it may be necessary to use configuration parameters in the Designated Router selection. These parameters will only need to be configured in those routers that are themselves eligible to become Designated Router (i.e., those router's whose Router Priority for the network is non-zero), and then only if no automatic procedure for discovering neighbors exists:

List of all other attached routers

The list of all other routers attached to the NBMA network. Each router is configured with its Router ID and IPv6 link-local address on the network. Also, for each router listed, that router's eligibility to become Designated Router must be defined. When an interface to a NBMA network comes up, the router sends Hello Packets only to those neighbors eligible to become Designated Router, until the identity of the Designated Router is discovered.

PollInterval If a neighboring router has become inactive (Hello Packets have not been seen for RouterDeadInterval seconds), it may still be necessary to send Hello Packets to the dead neighbor. These Hello Packets will be sent at the reduced rate PollInterval, which should be much larger than HelloInterval. Sample value for a PDN X.25 network: 2 minutes.

C.6 Point-to-MultiPoint network parameters

On Point-to-MultiPoint networks, it may be necessary to configure the set of neighbors that are directly reachable over the Point-to-MultiPoint network. Each neighbor is configured with its Router ID and IPv6 link-local address on the network. Designated Routers are not elected on Point-to-MultiPoint networks, so the Designated Router eligibility of configured neighbors is undefined.

C.7 Host route parameters

Host prefixes are advertised in intra-area-prefix-LSAs. They indicate either internal router addresses, router interfaces to point-to-point networks, looped router interfaces, or IPv6 hosts that are directly connected to the router (e.g., via a PPP connection). For each host directly connected to the router, the following items must be configured:

Host IPv6 prefix

The IPv6 prefix belonging to the host.

Cost of link to host

The cost of sending a packet to the host, in terms of the link state metric. However, since the host probably has only a single connection to the internet, the actual configured cost(s) in many cases is unimportant (i.e., will have no effect on routing).

Area ID

The OSPF area to which the host's prefix belongs.

Security Considerations

When running over IPv6, OSPF relies on the IP Authentication Header (see [Ref19]) and the IP Encapsulating Security Payload (see [Ref20]) to ensure integrity and authentication/confidentiality of routing exchanges.

Most OSPF implementations will be running on systems that support multiple protocols, many of them having independent security assumptions and domains. When IPSEC is used to protect OSPF packets, it is important for the implementation to check the IPSEC SA, and local SA database to make sure that the packet originates from a source THAT IS TRUSTED FOR OSPF PURPOSES.

Authors' Addresses

Rob Coltun
Siara Systems
300 Ferguson Drive
Mountain View, CA 94043

Phone: (650) 390-9030
EMail: rcoltun@siara.com

Dennis Ferguson
Juniper Networks
385 Ravendale Drive
Mountain View, CA 94043

Phone: +1 650 526 8004
EMail: dennis@juniper.com

John Moy
Sycamore Networks, Inc.
10 Elizabeth Drive
Chelmsford, MA 01824

Phone: (978) 367-2161
Fax: (978) 250-3350
EMail: jmoy@sycamorenet.com

Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

