Expiration Date: November 1997 FORE Systems
File name: draft-ietf-ospf-opaque-01.txt

# The OSPF Opaque LSA Option

Rob Coltun
FORE Systems
(301) 571-2521
rcoltun@fore.com

## Status Of This Memo

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet- Drafts as reference material or to cite them other than as "work in progress".

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet- Drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

Table Of Contents

## 1.0 Abstract

This memo documents enhancements to the OSPF protocol to support a new class of link-state advertisements (LSA) called Opaque LSAs. The Opaque LSA option defines a general mechanism to allow for future extensibility of OSPF. The information contained in Opaque LSAs may be used directly by OSPF or by other protocols. Opaque LSAs contain some number of octets padded to 32-bit alignment. The standard OSPF link-state database flooding mechanisms are use for distribution of Opaque LSAs. Opaque LSAs are flooded throughout all or some limited portion of the OSPF topology.

## 2.0 Overview

Over the last few years the OSPF routing protocol [OSPF] has been widely deployed throughout the Internet. As a result of this deployment and the evolution of networking technology, OSPF has been extended to support many options; this evolution will obviously continue.

This memo documents enhancements to the OSPF protocol to support a new class of link-state advertisements (LSA) called Opaque LSAs which defines an optional generalized mechanism to allow for future extensibility of OSPF. The information contained in Opaque LSAs may be used directly by OSPF or by other protocols. For example, the OSPF LSA may be used to distribute BGP AS Path information (as documented in The OSPF External Attributes LSA [EAL]) which is then used by BGP route-leaking mechanisms. The option may also be used to distribute IP QoS information which may be used directly by an OSPF path computation. The exact use of Opaque LSAs is beyond the scope of this draft.

The data contained in an Opaque LSA is some number of 32-bit aligned octets. Like any other LSA, the Opaque LSA uses the link-state database distribution mechanism for flooding this information throughout the topology. The link-state type of the Opaque LSA identifies the Flooding Scope or range of the topology to which this LSA may be distributed to.

### 2.1 Organization Of This Document

This document first defines the three types of Opaque LSAs followed by a description of OSPF packet processing which includes modifications to the flooding procedure and the neighbor state machine needed to support the Opaque LSA. Appendix A then gives the packet formats.

## 2.2 Acknowledgments

The author would like to thank Dennis Ferguson, Acee Lindem, John Moy, Sandra Murphy, Zhaohui "Jeffrey" Zhang and the rest of the OSPF Working Group for the ideas and support they have given to this project.

# 3.0 The Opaque LSA

Opaque LSAs are types 9, 10 and 11 link-state advertisements. Each type has a unique flooding scope and may be originated by any router. The data contained in the Opaque LSA consists of some number of octets aligned to a 32-bit boundary. Like any other LSA, Opaque LSAs use the link-state database distribution mechanism for flooding this information throughout the topology. The Opaque LSA's link-state type identifies the range of the topology to which this LSA may be distributed to. This section documents the flooding of Opaque LSAs.

The following are possible values of the link-state type filed with their related flooding scope.

- Link-state type 9 denotes a link-local scope. Type 9 Opaque LSAs are not flooded beyond the local (sub)network.

- Link-state type 10 denotes an area-local scope. Type 10 Opaque LSAs are not flooded beyond the area that they are originated into.

- Link-state type 11 denotes that the LSA is flooded throughout the Autonomous System (AS). Type 11 LSAs maintain the flooding scope of existing OSPF areas. Specifically type 11 Opaque LSAs are 1) flooded throughout all transit areas, 2) not flooded into stub areas from the backbone and 3) not originated within stub areas. As with type 5 LSAs, if a type 11 Opaque LSA is received in a stub area from a neighboring router within the stub area the LSA is rejected.

Origination of Opaque LSAs are unique to the application using it. The link-state ID of the Opaque LSA is divided into a type field (the first 8 bits) a type-specific ID (the remaining 24 bits). The packet format of the Opaque LSA is given in Appendix A.

The responsibility for proper handling of the Opaque LSA's flooding scope is placed on both the sender and receiver of the LSA. The receiver must always store a valid received Opaque LSA in its link-state database. The receiver must not accept Opaque LSAs that violate the flooding scope (i.e., a type 11 (domain-wide) Opaque LSA is not

accepted in a stub area). Flooding scope effects both the building of the Database summary list during the initial synchronization of the link-state database and the flooding procedure.

In order to make the use of the Opaque LSAs predictable, it is recommended that all routers within the scope of use have the same Opaque LSA capabilities. For example, if the Opaque LSA is to be used for flooding Opaque information throughout a single area, all routers within the area should support the Opaque option.

The following describes the modifications to these procedures that are necessary to insure proper use of the Opaque LSA's Scoping Rules.

### 3.1 Flooding Opaque LSAs

The flooding of Opaque LSAs must follow the rules of Flooding Scope as specified in this section. The flooding mechanisms must suppress the flooding of Opaque LSAs as described in the following.

- o If the Opaque LSA is type 9 (the flooding scope is link-local) and the interface that the LSA was received on is not the same as the target interface (e.g., the interface associated with a particular neighbor), the Opaque LSA must not be flooded out that interface (or to that neighbor). An implementation should keep track of the IP interface associated with each Opaque LSA having a link-local flooding scope.

- o If the Opaque LSA is type 10 (the flooding scope is area-local) and the area associated with Opaque LSA is not the area associated with a particular interface, the Opaque LSA must not be flooded out the interface. An implementation should keep track of the OSPF area associated with each Opaque LSA having an area-local flooding scope.

- o If the Opaque LSA is type 11 (the flooding scope is the entire AS) and 1) the area associated with Opaque LSA is not the area associated with a particular interface (i.e., the target area is not the same) and 2) the target area is a stub area, the Opaque LSA must not be flooded out the interface.

- o If an Opaque LSA is received on an interface associated with a stub area the LSA, the LSA is to be discarded and not acknowledged since neighboring router has flooded the LSA in error (see Section 13 of [OSPF], receiving LSAs having unknown LS types).

When opaque-capable routers and non-opaque-capable OSPF routers are mixed together in a routing domain, the Opaque LSAs are not flooded to the non-opaque-capable routers. As a general design principle, optional OSPF advertisements are only flooded to those routers that understand them.

An opaque-capable router learns of its neighbor's opaque capability at the beginning of the "Database Exchange Process" (see Section 10.6 of [OSPF], receiving Database Description packets from a neighbor in state ExStart). A neighbor is opaque-capable if and only if it sets the O-bit in the Options field of its Database Description packets. Then, in the next step of the Database Exchange process, Opaque LSAs are included in the Database summary list sent to the neighbor (see Sections 3.2 below and 10.3 of [OSPF]) if and only if the neighbor is opaque capable.

When flooding Opaque-LSAs to adjacent neighbors, a opaque-capable router looks at the neighbor's opaque capability. Opaque LSAs are only flooded to opaque capable neighbors. To be more precise, in Section 13.3 of [OSPF], Opaque LSAs are only placed on the link-state retransmission lists of opaque-capable neighbors. Note however that when sending Link State Update packets as multicasts, a non-opaque-capable neighbor may (inadvertently) receive Opaque LSAs. The non-opaque-capable router will then simply discard the LSA (see Section 13 of [OSPF], receiving LSAs having unknown LS types).

### 3.2 Modifications To The Neighbor State Machine

The state machine as it exists in section 10.3 of [OSPF] remains unchanged except for the action associated with State: ExStart, Event: NegotiationDone which is where the Database summary list is built. To incorporate the Opaque LSA in OSPF the action is changed to the following.

| | |
|---|---|
| State(s): | ExStart |
| Event: | NegotiationDone |
| New state: | Exchange |
| Action: | The router must list the contents of its entire area link-state database in the neighbor Database summary list. The area link-state database consists of the Router LSAs, Network LSAs, Summary LSAs and types 9 and 10 Opaque LSAs contained in the area structure, along with AS External and type 11 Opaque LSAs contained in the global structure. AS External and |

type 11 Opaque LSAs are omitted from a virtual neighbor's Database summary list. AS External LSAs and type 11 Opaque LSAs are omitted from the Database summary list if the area has been configured as a stub (see Section 3.6 of [OSPF]).

Opaque LSAs are omitted from the Database summary list if the following conditions are met:
1) the LSA type is type 9 (the flooding scope is link-local) and the interface associated with the Opaque LSA (upon reception) does not equal the interface associated with the neighbor;
2) the LSA type is 10 (the flooding scope is area-local) and the area associated with Opaque LSA is not the area associated with the neighbor's interface.

Any advertisement whose age is equal to MaxAge is omitted from the Database summary list. It is instead added to the neighbor's link-state retransmission list. A summary of the Database summary list will be sent to the neighbor in Database Description packets. Each Database Description Packet has a DD sequence number, and is explicitly acknowledged. Only one Database Description Packet is allowed to be outstanding at any one time. For more detail on the sending and receiving of Database Description packets, see Sections 10.8 and 10.6 of [OSPF].

## 4.0 Protocol data structures

The Opaque option is described herein in terms of its operation on various protocol data structures. These data structures are included for explanatory uses only, and are not intended to constrain an OSPF implementation. Besides the data structures listed below, this specification will also reference the various data structures (e.g., OSPF neighbors) defined in [OSPF].

In an OSPF router, the following item is added to the list of global OSPF data structures described in Section 5 of [OSPF]:

- o Opaque capability. Indicates whether the router is running the Opaque option (i.e., capable of storing Opaque LSAs). Such a router will continue to inter-operate with non-opaque-capable OSPF routers.

## 4.1 Additions To The OSPF Neighbor Structure

The OSPF neighbor structure is defined in Section 10 of [OSPF]. In an opaque-capable router, the following items are added to the OSPF neighbor structure:

- o Neighbor Options. This field was already defined in the OSPF specification. However, in opaque-capable routers there is a new option which indicates the neighbor's Opaque capability. This new option is learned in the Database Exchange process through reception of the neighbor's Database Description packets, and determines whether Opaque LSAs are flooded to the neighbor. For a more detailed explanation of the flooding of the Opaque LSA see section 3 of this document.

## 5.0 References

[OSPF] Moy, J., "OSPF Version 2", IETF Internet Draft, Cascade, April 1997.

[MOSPF] Moy, J., "Multicast Extensions to OSPF", RFC 1584, Proteon, Inc., March 1994.

[NSSA] Coltun, R. and V. Fuller, "The OSPF NSSA Option", RFC 1587, RainbowBridge Communications, Stanford University, March 1994.

[DEMD] Moy, J., "Extending OSPF to Support Demand Circuits", RFC 1793, Cascade, April 1995.

[EAL] Ferguson, D., "The OSPF External Attributes LSA", work in progress.

## Appendix A: OSPF Data formats

This appendix describes the format of the Options Field followed by the packet format of the Opaque LSA.

### A.1 The Options Field

The OSPF Options field is present in OSPF Hello packets, Database Description packets and all link-state advertisements. 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.

When used in Hello packets, the Options field allows a router to reject a neighbor because of a capability mismatch. Alternatively, when capabilities are exchanged in Database Description packets a router can choose not to forward certain link-state advertisements to a neighbor because of its reduced functionality. Lastly, listing capabilities in link-state advertisements allows routers to forward traffic around reduced functionality routers, by excluding them from parts of the routing table calculation.

Seven bits of the OSPF Options field have been assigned, although only the O-bit is described completely by this memo. 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 link-state advertisements. Conversely, routers encountering unrecognized Option bits in received Hello Packets, Database Description packets or link-state advertisements should ignore the capability and process the packet/advertisement normally.


The Options Field

T-bit
: This bit describes the router's TOS-based routing capability, as specified in Sections 9.5, 10.8, 12.1.2 and 16.9 of [OSPF].

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 [OSPF].

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

N/P-bit
This bit describes the handling of Type-7 LSAs, as specified in [NSSA].

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

EA-bit
This bit describes the router's willingness to receive and forward External-Attributes-LSAs, as specified in [EAL].

O-bit
This bit describes the router's willingness to receive and forward Opaque-LSAs as specified in this document.

A.2 Opaque LSA

Opaque LSAs are the Type 9, 10 and 11 link-state advertisements. These advertisements are originated by any router and may be used directly by OSPF or indirectly by other protocols such as BGP wishing to distribute information throughout the OSPF domain. The primary function of the Opaque LSA is to provide for future extensibility to OSPF.

The data contained in the Opaque LSA consists of some number of octets padded to 32-bit alignment. Like any other LSA, the Opaque LSA uses the link-state database distribution mechanism for flooding this information throughout the topology. However, the Opaque LSA has a flooding scope associated with it so that the scope of flooding may be link-local (type 9), area-local (type 10) or the entire OSPF routing domain (type 11).

Origination of Opaque LSAs are unique to the application using it. Section 3 of this document describes the flooding procedures for the Opaque LSA.

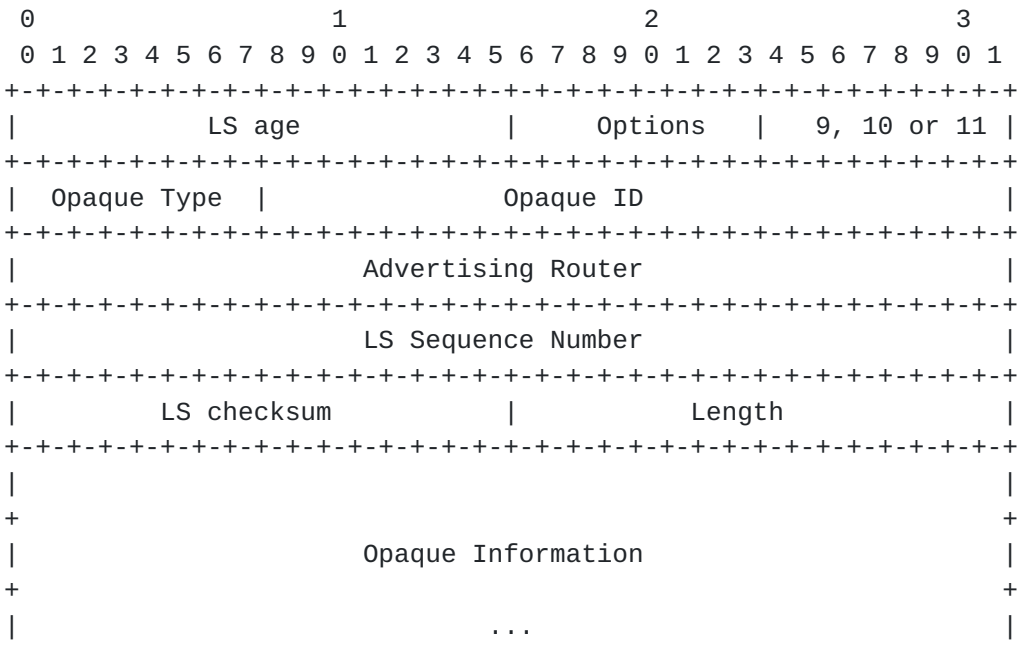


Link-State Type

The link-state type identifies the flooding scope (or range) of the topology to which this LSA may be distributed to. The following explains the flooding scope of each of the link-state types.

o A value of 9 denotes a link-local scope. Opaque LSAs with a

link-local scope are not flooded beyond the local (sub)network.

o A value of 10 denotes an area-local scope. Opaque LSAs with a area-local scope are not flooded beyond the area that they are originated into.

o A value of 11 denotes that the LSA is flooded throughout the Autonomous System (e.g., has the same scope as type-5 LSAs). Opaque LSAs with AS-wide scope are not flooded into stub areas.

### Syntax Of The Opaque LSA's Link-State ID

The link-state ID of the Opaque LSA is divided into an Opaque Type field (the first 8 bits) and an Opaque ID (the remaining 24 bits). Opaque type values in the range of 128-255 are reserved for private and experimental use.