

The OSPF Opaque LSA Option

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 (1998). All Rights Reserved.

Table Of Contents

1.0 Abstract	1
2.0 Overview	2
2.1 Organization Of This Document	2
2.2 Acknowledgments	3
3.0 The Opaque LSA	3
3.1 Flooding Opaque LSAs	4
3.2 Modifications To The Neighbor State Machine	5
4.0 Protocol Data Structures	6
4.1 Additions To The OSPF Neighbor Structure	6
5.0 Management Considerations	7
6.0 Security Considerations	9
7.0 IANA Considerations	10
8.0 References	10
9.0 Author's Information	11
Appendix A: OSPF Data Formats	12
A.1 The Options Field	12
A.2 The Opaque LSA	13
Appendix B: Full Copyright Statment	15

1.0 Abstract

This memo defines enhancements to the OSPF protocol to support a new class of link-state advertisements (LSA) called Opaque LSAs. Opaque LSAs provide a generalized mechanism to allow for the future extensibility of OSPF. Opaque LSAs consist of a standard LSA header followed by application-specific information. The information field

may be used directly by OSPF or by other applications. Standard OSPF link-state database flooding mechanisms are used to distribute Opaque LSAs to all or some limited portion of the OSPF topology.

2.0 Overview

Over the last several 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 defines enhancements to the OSPF protocol to support a new class of link-state advertisements (LSA) called Opaque LSAs. Opaque LSAs provide a generalized mechanism to allow for the future extensibility of OSPF. The information contained in Opaque LSAs may be used directly by OSPF or indirectly by some application wishing to distribute information throughout the OSPF domain. For example, the OSPF LSA may be used by routers to distribute IP to link-layer address resolution information (see [ARA] for more information). The exact use of Opaque LSAs is beyond the scope of this memo.

Opaque LSAs consist of a standard LSA header followed by a 32-bit qaligned application-specific information field. 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 field of the Opaque LSA identifies the LSA's range of topological distribution. This range is referred to as the Flooding Scope.

It is envisioned that an implementation of the Opaque option provides an application interface for 1) encapsulating application-specific information in a specific Opaque type, 2) sending and receiving application-specific information, and 3) if required, informing the application of the change in validity of previously received information when topological changes are detected.

2.1 Organization Of This Document

This document first defines the three types of Opaque LSAs followed by a description of OSPF packet processing. The packet processing sections include modifications to the flooding procedure and to the neighbor state machine. Appendix A then gives the packet formats.

2.2 Acknowledgments

The author would like to thank Dennis Ferguson, Acee Lindem, John Moy, Sandra Murphy, Man-Kit Yeung, 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. Opaque LSAs consist of a standard LSA header followed by a 32-bit aligned application-specific information field. Standard link-state database flooding mechanisms are used for distribution of Opaque LSAs. The range of topological distribution (i.e., the flooding scope) of an Opaque LSA is identified by its link-state type. This section documents the flooding of Opaque LSAs.

The flooding scope associated with each Opaque link-state type is defined as follows.

- o Link-state type 9 denotes a link-local scope. Type-9 Opaque LSAs are not flooded beyond the local (sub)network.
- o Link-state type 10 denotes an area-local scope. Type-10 Opaque LSAs are not flooded beyond the borders of their associated area.
- o Link-state type 11 denotes that the LSA is flooded throughout the Autonomous System (AS). The flooding scope of type-11 LSAs are equivalent to the flooding scope of AS-external (type-5) LSAs. 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 by routers into their connected 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.

The link-state ID of the Opaque LSA is divided into an Opaque type field (the first 8 bits) and a type-specific ID (the remaining 24 bits). The packet format of the Opaque LSA is given in Appendix A. Section 7.0 describes Opaque type allocation and assignment.

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 (e.g., a type-11 (domain-wide) Opaque LSA is not accepted in a stub area). The flooding scope effects both the

synchronization of the link-state database and the flooding procedure.

The following describes the modifications to these procedures that are necessary to insure conformance to 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. Section 13 of [OSPF] describes the OSPF flooding procedure. The following describes the Opaque LSA's type-specific flooding restrictions.

- 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 target 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 (upon reception) is not the same as the area associated with the target 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 LSA is flooded throughout the AS) and the target interface is associated with a stub area the Opaque LSA must not be flooded out the interface. A type-11 Opaque LSA that is received on an interface associated with a stub area must be discarded and not acknowledged (the neighboring router has flooded the LSA in error).

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; the O-bit is not set in packets other than Database Description

packets. Then, in the next step of the Database Exchange process, Opaque LSAs are included in the Database summary list that is 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. However, 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 this 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 area (see Section 3.6 of [OSPF]).

Type-9 Opaque LSAs are omitted from the Database summary list if the interface associated with the neighbor is not the interface associated with the Opaque LSA (as noted upon reception).

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 10.6 and 10.8 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 implementation. In addition to the data structures listed below, this specification references 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 Management Considerations

This section identifies the current OSPF MIB [OSPFMIB] capabilities that are applicable to the Opaque option and lists the additional management information which is required for its support.

Opaque LSAs are types 9, 10 and 11 link-state advertisements. The link-state ID of the Opaque LSA is divided into an Opaque type field (the first 8 bits) and a type-specific ID (the remaining 24 bits). The packet format of the Opaque LSA is given in Appendix A. The range of topological distribution (i.e., the flooding scope) of an Opaque LSA is identified by its link-state type.

- o Link-State type 9 Opaque LSAs have a link-local scope. Type-9 Opaque LSAs are flooded on a single local (sub)network but are not flooded beyond the local (sub)network.
- o Link-state type 10 Opaque LSAs have an area-local scope. Type-10 Opaque LSAs are flooded throughout a single area but are not flooded beyond the borders of the associated area.
- o Link-state type 11 Opaque LSAs have an Autonomous-System-wide scope. The flooding scope of type-11 LSAs are equivalent to the flooding scope of AS-external (type-5) LSAs.

The OSPF MIB provides a number of objects that can be used to manage and monitor an OSPF router's Link-State Database. The ones that are relevant to the Opaque option are as follows.

The `ospfGeneralGroup` defines two objects for keeping track of newly originated and newly received LSAs (`ospfOriginateNewLsas` and `ospfRxNewLsas` respectively).

The OSPF MIB defines a set of optional traps. The `ospfOriginateLsa` trap signifies that a new LSA has been originated by a router and the `ospfMaxAgeLsa` trap signifies that one of the LSAs in the router's link-state database has aged to `MaxAge`.

The `ospfAreaTable` describes the configured parameters and cumulative statistics of the router's attached areas. This table includes a count of the number of LSAs contained in the area's link-state database (`ospfAreaLsaCount`), and a sum of the LSA's LS checksums contained in this area (`ospfAreaLsaCksumSum`). This sum can be used to determine if there has been a change in a router's link-state database, and to compare the link-state database of two routers.

The `ospfLsdbTable` describes the OSPF Process's link-state database (excluding AS-external LSAs). Entries in this table are indexed with an Area ID, a link-state type, a link-state ID and the originating router's Router ID.

The management objects that are needed to support the Opaque option are as follows.

An Opaque-option-enabled object is needed to indicate if the Opaque option is enabled on the router.

The origination and reception of new Opaque LSAs should be reflected in the counters `ospfOriginateNewLsas` and `ospfRxNewLsas` (inclusive for types 9, 10 and 11 Opaque LSAs).

If the OSPF trap option is supported, the origination of new Opaque LSAs and purging of MaxAge Opaque LSAs should be reflected in the `ospfOriginateLsa` and `ospfMaxAgeLsa` traps (inclusive for types 9, 10 and 11 Opaque LSAs).

The number of type-10 Opaque LSAs should be reflected in `ospfAreaLsaCount`; the checksums of type-10 Opaque LSAs should be included in `ospfAreaLsaChksumSum`.

Type-10 Opaque LSAs should be included in the `ospfLsdbTable`. Note that this table does not include a method of examining the Opaque type field (in the Opaque option this is a sub-field of the link-state ID).

Up until now, LSAs have not had a link-local scope so there is no method of requesting the number of, or examining the LSAs that are associated with a specific OSPF interface. A new group of management objects are required to support type-9 Opaque LSAs. These objects should include a count of type-9 Opaque LSAs, a checksum sum and a table for displaying the link-state database for type-9 Opaque LSAs on a per-interface basis. Entries in this table should be indexed with an Area ID, interface's IP address, Opaque type, link-state ID and the originating router's Router ID.

Prior to the introduction of type-11 Opaque LSAs, AS-External (type-5) LSAs have been the only link-state types which have an Autonomous-System-wide scope. A new group of objects are required to support type-11 Opaque LSAs. These objects should include a count of type-11 Opaque LSAs, a type-11 checksum sum and a table for displaying the type-11 link-state database. Entries in this table should be indexed with the Opaque type, link-state ID and the

originating router's Router ID. The type-11 link-state database table will allow type-11 LSAs to be displayed once for the router rather than once in each non-stub area.

6.0 Security Considerations

There are two types of issues that need be addressed when looking at protecting routing protocols from misconfigurations and malicious attacks. The first is authentication and certification of routing protocol information. The second is denial of service attacks resulting from repetitive origination of the same router advertisement or origination a large number of distinct advertisements resulting in database overflow. Note that both of these concerns exist independently of a router's support for the Opaque option.

To address the authentication concerns, OSPF protocol exchanges are authenticated. OSPF supports multiple types of authentication; the type of authentication in use can be configured on a per network segment basis. One of OSPF's authentication types, namely the Cryptographic authentication option, is believed to be secure against passive attacks and provide significant protection against active attacks. When using the Cryptographic authentication option, each router appends a "message digest" to its transmitted OSPF packets. Receivers then use the shared secret key and received digest to verify that each received OSPF packet is authentic.

The quality of the security provided by the Cryptographic authentication option depends completely on the strength of the message digest algorithm (MD5 is currently the only message digest algorithm specified), the strength of the key being used, and the correct implementation of the security mechanism in all communicating OSPF implementations. It also requires that all parties maintain the secrecy of the shared secret key. None of the standard OSPF authentication types provide confidentiality. Nor do they protect against traffic analysis. For more information on the standard OSPF security mechanisms, see Sections 8.1, 8.2, and Appendix D of [OSPF].

[DIGI] describes the extensions to OSPF required to add digital signature authentication to Link State data and to provide a certification mechanism for router data. [DIGI] also describes the added LSA processing and key management as well as a method for migration from, or co-existence with, standard OSPF V2.

Repetitive origination of advertisements are addressed by OSPF by mandating a limit on the frequency that new instances of any particular LSA can be originated and accepted during the flooding procedure. The frequency at which new LSA instances may be

originated is set equal to once every MinLSInterval seconds, whose value is 5 seconds (see Section 12.4 of [OSPF]). The frequency at which new LSA instances are accepted during flooding is once every MinLSArrival seconds, whose value is set to 1 (see Section 13, Appendix B and G.5 of [OSPF]).

Proper operation of the OSPF protocol requires that all OSPF routers maintain an identical copy of the OSPF link-state database. However, when the size of the link-state database becomes very large, some routers may be unable to keep the entire database due to resource shortages; we term this "database overflow". When database overflow is anticipated, the routers with limited resources can be accommodated by configuring OSPF stub areas and NSSAs. [OVERFLOW] details a way of gracefully handling unanticipated database overflows.

7.0 IANA Considerations

Opaque types are maintained by the IANA. Extensions to OSPF which require a new Opaque type must be reviewed by the OSPF working group. In the event that the OSPF working group has disbanded the review shall be performed by a recommended Designated Expert.

Following the policies outlined in [IANA], Opaque type values in the range of 0-127 are allocated through an IETF Consensus action and Opaque type values in the range of 128-255 are reserved for private and experimental use.

8.0 References

- [ARA] Coltun, R., and J. Heinanen, "The OSPF Address Resolution Advertisement Option", Work in Progress.
- [DEMD] Moy, J., "Extending OSPF to Support Demand Circuits", RFC 1793, April 1995.
- [DIGI] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.
- [IANA] Narten, T., and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", Work in Progress.
- [MOSPF] Moy, J., "Multicast Extensions to OSPF", RFC 1584, March 1994.

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

[OSPF] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[OSPFMIB] Baker, F., and R. Coltun, "OSPF Version 2 Management Information Base", RFC 1850, November 1995.

[OVERFLOW] Moy, J., "OSPF Database Overflow", RFC 1765, March 1995.

9.0 Author's Information

Rob Coltun
FORE Systems

Phone: (703) 245-4543
EMail: rcoltun@fore.com

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.

Six 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.

```

+-----+
| * | O | DC | EA | N/P | MC | E | * |
+-----+

```

The Options Field

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].

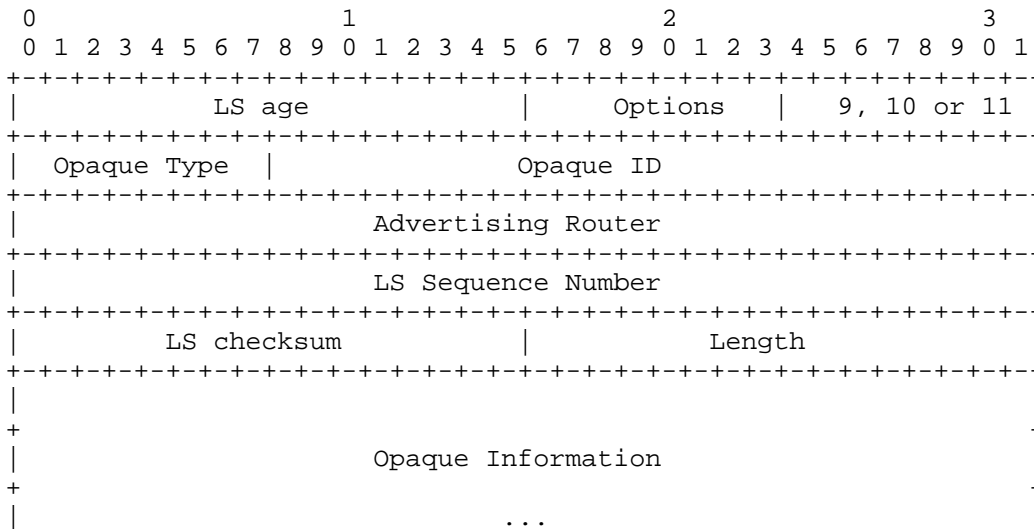
O-bit

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

A.2 The Opaque LSA

Opaque LSAs are Type 9, 10 and 11 link-state advertisements. These advertisements may be used directly by OSPF or indirectly by some application wishing to distribute information throughout the OSPF domain. The function of the Opaque LSA option is to provide for future extensibility of OSPF.

Opaque LSAs contain some number of octets (of application-specific data) 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). Section 3 of this document describes the flooding procedures for the Opaque LSA.



Link-State Type

The link-state type of the Opaque LSA identifies the LSA's range of topological distribution. This range is referred to as the Flooding Scope. 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). See section 7.0 of this document for a description of Opaque type allocation and assignment.

Appendix B. Full Copyright Statement

Copyright (C) The Internet Society (1998). 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.