

Network Working Group
Request for Comments: 4970
Category: Standards Track

A. Lindem, Ed.
Redback Networks
N. Shen
JP. Vasseur
Cisco Systems
R. Aggarwal
Juniper Networks
S. Shaffer
BridgePort Networks
July 2007

Extensions to OSPF for Advertising Optional Router Capabilities

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 IETF Trust (2007).

Abstract

It is useful for routers in an OSPFv2 or OSPFv3 routing domain to know the capabilities of their neighbors and other routers in the routing domain. This document proposes extensions to OSPFv2 and OSPFv3 for advertising optional router capabilities. A new Router Information (RI) Link State Advertisement (LSA) is proposed for this purpose. In OSPFv2, the RI LSA will be implemented with a new opaque LSA type ID. In OSPFv3, the RI LSA will be implemented with a new LSA type function code. In both protocols, the RI LSA can be advertised at any of the defined flooding scopes (link, area, or autonomous system (AS)).

Table of Contents

1.	Introduction	3
1.1.	Requirements Notation	3
2.	OSPF Router Information (RI) LSA	3
2.1.	OSPFv2 Router Information (RI) Opaque LSA	3
2.2.	OSPFv3 Router Information (RI) Opaque LSA	5
2.3.	OSPF Router Information Capabilities TLV	5
2.4.	Assigned OSPF Router Information Capability Bits	6
2.5.	Flooding Scope of the Router Information LSA	7
3.	Router Information LSA Opaque Usage and Applicability	7
4.	Security Considerations	7
5.	IANA Considerations	8
6.	References	10
6.1.	Normative References	10
6.2.	Informative References	10
	Appendix A. Acknowledgments	11

1. Introduction

It is useful for routers in an OSPFv2 [OSPF] or OSPFv3 [OSPFV3] routing domain to know the capabilities of their neighbors and other routers in the routing domain. This can be useful for both the advertisement and discovery of OSPFv2 and OSPFv3 capabilities. Throughout this document, OSPF will be used when the specification is applicable to both OSPFv2 and OSPFv3. Similarly, OSPFv2 or OSPFv3 will be used when the text is protocol specific.

OSPF uses the options field in LSAs and hello packets to advertise optional router capabilities. In the case of OSPFv2, all the bits in this field have been allocated so new optional capabilities cannot be advertised. This document proposes extensions to OSPF to advertise these optional capabilities via opaque LSAs in OSPFv2 and new LSAs in OSPFv3. For existing OSPF capabilities, backward-compatibility issues dictate that this advertisement is used primarily for informational purposes. For future OSPF features, this advertisement MAY be used as the sole mechanism for advertisement and discovery.

1.1. Requirements Notation

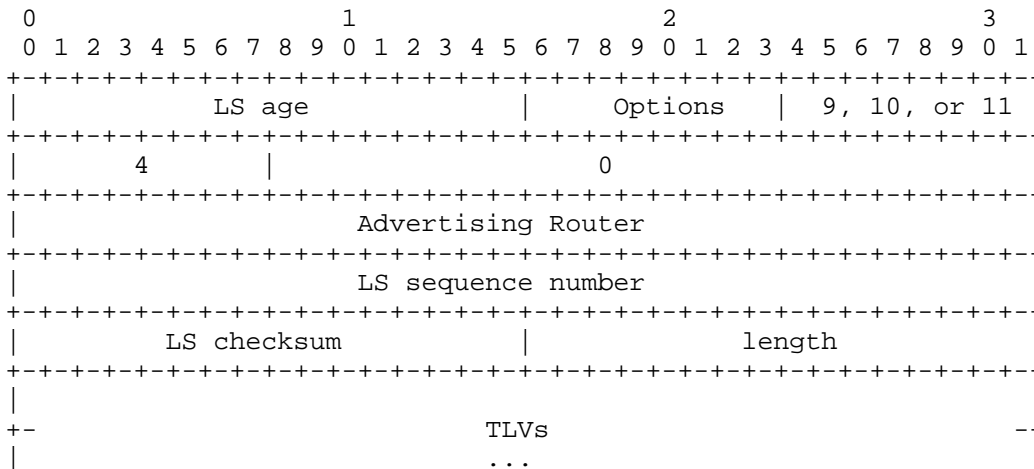
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-KEYWORDS].

2. OSPF Router Information (RI) LSA

OSPF routers MAY optionally advertise their optional capabilities in a link-scoped, area-scoped, or AS-scoped LSA. For existing OSPF capabilities, this advertisement will be used primarily for informational purposes. Future OSPF features could use the RI LSA as the sole mechanism for advertisement and discovery. The RI LSA will be originated initially when an OSPF router instance is created and whenever one of the advertised capabilities is configured or changed.

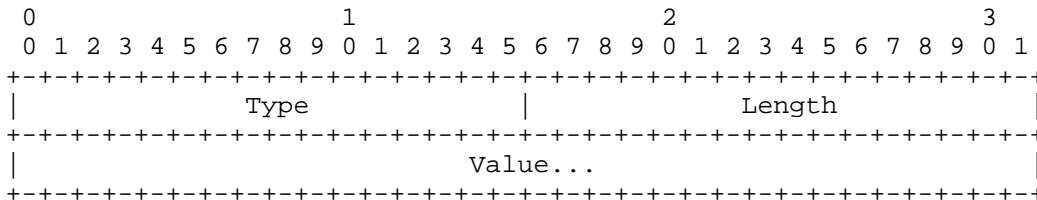
2.1. OSPFv2 Router Information (RI) Opaque LSA

OSPFv2 routers will advertise a link scoped, area-scoped, or AS-scoped Opaque-LSA [OPAQUE]. The OSPFv2 Router Information LSA has an Opaque type of 4 and Opaque ID of 0.



OSPFv2 Router Information Opaque LSA

The format of the TLVs within the body of an RI LSA is the same as the format used by the Traffic Engineering Extensions to OSPF [TE]. The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets. The format of each TLV is:

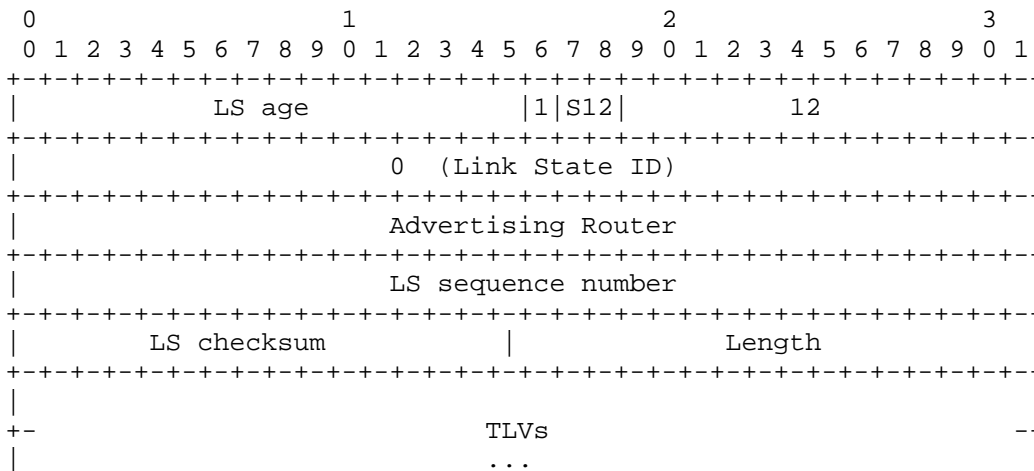


TLV Format

The Length field defines the length of the value portion in octets (thus a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the length field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). Nested TLVs are also 32-bit aligned. For example, a 1-byte value would have the length field set to 1, and 3 octets of padding would be added to the end of the value portion of the TLV. Unrecognized types are ignored.

2.2. OSPFv3 Router Information (RI) Opaque LSA

The OSPFv3 Router Information LSA has a function code of 12 while the S1/S2 bits are dependent on the desired flooding scope for the LSA. The U bit will be set indicating that the OSPFv3 RI LSA should be flooded even if it is not understood. The Link State ID (LSID) value for this LSA is 0. This is unambiguous since an OSPFv3 router will only advertise a single RI LSA per flooding scope.



OSPFv3 Router Information LSA

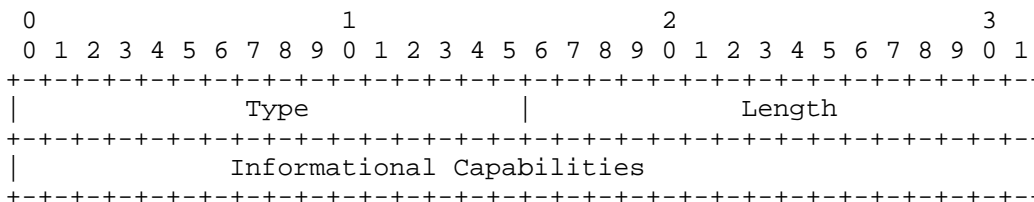
The format of the TLVs within the body of an RI LSA is as defined in Section 2.1

When a new Router Information LSA TLV is defined, the specification MUST explicitly state whether the TLV is applicable to OSPFv2 only, OSPFv3 only, or both OSPFv2 and OSPFv3.

2.3. OSPF Router Informational Capabilities TLV

The first defined TLV in the body of an RI LSA is the Router Informational Capabilities TLV. A router advertising an RI LSA MAY include the Router Informational Capabilities TLV. If included, it MUST be the first TLV in the LSA. Additionally, the TLV MUST accurately reflect the OSPF router's capabilities in the scope advertised. However, the informational capabilities advertised have no impact on the OSPF protocol's operation -- they are advertised purely for informational purposes.

The format of the Router Informational Capabilities TLV is as follows:



- Type A 16-bit field set to 1.
- Length A 16-bit field that indicates the length of the value portion in octets and will be a multiple of 4 octets dependent on the number of capabilities advertised. Initially, the length will be 4, denoting 4 octets of informational capability bits.
- Value A variable length sequence of capability bits rounded to a multiple of 4 octets padded with undefined bits. Initially, there are 4 octets of capability bits. Bits are numbered left-to-right starting with the most significant bit being bit 0.

OSPF Router Informational Capabilities TLV

The Router Informational Capabilities TLV MAY be followed by optional TLVs that further specify a capability.

2.4. Assigned OSPF Router Informational Capability Bits

The following informational capability bits are assigned:

Bit	Capabilities
0	OSPF graceful restart capable [GRACE]
1	OSPF graceful restart helper [GRACE]
2	OSPF Stub Router support [STUB]
3	OSPF Traffic Engineering support [TE]
4	OSPF point-to-point over LAN [P2PLAN]
5	OSPF Experimental TE [EXP-TE]
6-31	Unassigned (Standards Action)

OSPF Router Informational Capabilities Bits

2.5. Flooding Scope of the Router Information LSA

The flooding scope for a Router Information LSA is determined by the LSA type. For OSPFv2, type 9 (link-scoped), type 10 (area-scoped), or a type 11 (AS-scoped) opaque LSA may be flooded. For OSPFv3, the S1 and S2 bits in the LSA type determine the flooding scope. If AS-wide flooding scope is chosen, the originating router should also advertise area-scoped LSA(s) into any attached Not-So-Stubby Area (NSSA) area(s). An OSPF router MAY advertise different capabilities when both NSSA area scoped LSA(s) and an AS-scoped LSA are advertised. This allows functional capabilities to be limited in scope. For example, a router may be an area border router but only support traffic engineering (TE) in a subset of its attached areas.

The choice of flooding scope is made by the advertising router and is a matter of local policy. The originating router MAY advertise multiple RI LSAs as long as the flooding scopes differ. TLV flooding scope rules will be specified on a per-TLV basis and MUST be specified in the accompanying specifications for new Router Information LSA TLVs.

3. Router Information LSA Opaque Usage and Applicability

The purpose of the Router Information (RI) LSA is to advertise information relating to the aggregate OSPF router. Normally, this should be confined to TLVs with a single value or very few values. It is not meant to be a generic container to carry any and all information. The intent is to both limit the size of the RI LSA to the point where an OSPF router will always be able to contain the TLVs in a single LSA and to keep the task of determining what has changed between LSA instances reasonably simple. Hence, discretion and sound engineering judgment will need to be applied when deciding whether newly proposed TLV(s) in support of a new application are advertised in the RI LSA or warrant the creation of an application specific LSA.

4. Security Considerations

This document describes both a generic mechanism for advertising router capabilities and a TLV for advertising informational capability bits. The latter TLV is less critical than the topology information currently advertised by the base OSPF protocol. The security considerations for the generic mechanism are dependent on the future application and, as such, should be described as additional capabilities are proposed for advertisement. Security considerations for the base OSPF protocol are covered in [OSPF] and [OSPFV3].

5. IANA Considerations

The following IANA assignment was made from an existing registry:

The OSPFv2 opaque LSA type 4 has been reserved for the OSPFv2 RI opaque LSA.

The following registries have been defined for the following purposes:

1. Registry for OSPFv3 LSA Function Codes - This new top-level registry will be comprised of the fields Value, LSA function code name, and Document Reference. The OSPFv3 LSA function code is defined in section A.4.2.1 of [OSPFV3]. The OSPFv3 LSA function code 12 has been reserved for the OSPFv3 Router Information (RI) LSA.

Range	Assignment Policy
0	Reserved (not to be assigned)
1-9	Already assigned
10-11	Unassigned (Standards Action)
12	OSPFv3 RI LSA (Assigned herein)
13-255	Unassigned (Standards Action)
256-8175	Reserved (No assignments)
8176-8183	Experimentation (No assignments)
8184-8191	Vendor Private Use (No assignments)

OSPFv3 LSA Function Codes

- * OSPFv3 LSA function codes in the range 256-8175 are not to be assigned at this time. Before any assignments can be made in this range, there MUST be a Standards Track RFC that specifies IANA Considerations that cover the range being assigned.
- * OSPFv3 LSA function codes in the range 8176-8181 are for experimental use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

- * OSPFv3 LSAs with an LSA Function Code in the Vendor Private Use range 8184-8191 MUST include the Vendor Enterprise Code as the first 4 octets following the 20 octets of LSA header.
 - * If a new LSA Function Code is documented, the documentation MUST include the valid combinations of the U, S2, and S1 bits for the LSA. It SHOULD also describe how the Link State ID is to be assigned.
2. Registry for OSPF RI TLVs - This top-level registry will be comprised of the fields Value, TLV Name, and Document Reference. The value of 1 for the capabilities TLV is defined herein.

Range	Assignment Policy
0	Reserved (not to be assigned)
1	Already assigned
2-32767	Unassigned (Standards Action)
32768-32777	Experimentation (No assignments)
32778-65535	Reserved (Not to be assigned)

OSPF RI TLVs

- * Types in the range 32768-32777 are for experimental use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.
 - * Types in the range 32778-65535 are reserved and are not to be assigned at this time. Before any assignments can be made in this range, there MUST be a Standards Track RFC that specifies IANA Considerations that covers the range being assigned.
3. Registry for OSPF Router Informational Capability Bits - This sub-registry of the OSPF RI TLV registry will be comprised of the fields Bit Number, Capability Name, and Document Reference. The values are defined in Section 2.4. All Router Informational Capability TLV additions are to be assigned through standards action.

6. References

6.1. Normative References

- [OPAQUE] Coltun, R., "The OSPF Opaque LSA Option", RFC 2370, July 1998.
- [OSPF] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [OSPFV3] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.
- [RFC-KEYWORDS] Bradner, S., "Key words for use in RFC's to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering Extensions to OSPF", RFC 3630, September 2003.

6.2. Informative References

- [EXP-TE] Srisuresh, P. and P. Joseph, "OSPF-xTE: Experimental Extension to OSPF for Traffic Engineering", RFC 4973, July 2007.
- [GRACE] Moy, J., Pillay-Esnault, P., and A. Lindem, "Graceful OSPF Restart", RFC 3623, November 2003.
- [P2PLAN] Shen, N. and A. Zinin, "Point-to-point operation over LAN in link-state routing protocols", Work in Progress, April 2006.
- [STUB] Retana, A., Nguyen, L., White, R., Zinin, A., and D. McPherson, "OSPF Stub Router Advertisement", RFC 3137, June 2001.

Appendix A. Acknowledgments

The idea for this work grew out of a conversation with Andrew Partan and we would like to thank him for his contribution. The authors would like to thanks Peter Psenak for his review and helpful comments on early versions of the document.

Comments from Abhay Roy, Vishwas Manral, Vivek Dubey, and Adrian Farrel have been incorporated into later versions.

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

Authors' Addresses

Acee Lindem (editor)
Redback Networks
102 Carric Bend Court
Cary, NC 27519
USA

EEmail: acee@redback.com

Naiming Shen
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA

EEmail: naiming@cisco.com

Jean-Philippe Vasseur
Cisco Systems
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

EEmail: jpv@cisco.com

Rahul Aggarwal
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

EEmail: rahul@juniper.net

Scott Shaffer
BridgePort Networks
One Main Street, 7th Floor
Cambridge, MA 02142
USA

EEmail: sshaffer@bridgeport-networks.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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

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

Intellectual Property

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

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

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

Acknowledgement

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