

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 22, 2015

M. Dubrovsky
R. Shrivastava
Cisco Systems
D. Cheng
Huawei Technologies
October 19, 2014

**Extensions to OSPF facilitating the deployment of non-backward-compatible changes.
draft-dubrovsky-ospf-non-compatible-02**

Abstract

This document specifies a generic mechanism that facilitates the deployment of non-backward-compatible changes in OSPF protocol. This mechanism allows the OSPF routers to advertise the capability of non-backward-compatible functionality and to make the functionality operational only when supported by all participating routers. Depending on the functionality scope, capability advertisements must be propagated across a link, area or autonomous system (AS). For link and area scope functionality, Router Information Link State Advertisement (LSA) is utilized to propagate the capability information. For the cases when compatibility must be maintained across the whole OSPF autonomous system, new Area Information (AI) LSA is introduced. The AI LSA is a TLV-based analog of Indication-LSA that is used for demand circuit functionality and described in [RFC1793](#).

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 22, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
2.	Method to deploy non-backward-compatible changes	4
3.	Area Information LSA	4
3.1.	OSPFv2 Area Information (AI) Opaque LSA	5
3.2.	OSPFv3 Area Information (AI) LSA	5
3.3.	Area Information LSA TLV format	6
3.4.	Area Information LSA origination	7
3.4.1.	Limiting Area Information LSA origination	7
4.	Capability negotiation before adjacency is fully formed	8
5.	Backward Compatibility	8
6.	IANA Considerations	8
7.	Security Considerations	8
8.	Acknowledgements	8
9.	References	9
9.1.	Normative References	9
9.2.	Informative References	9
	Authors' Addresses	9

1. Introduction

The evolution of OSPF protocol brought up changes that are not backward-compatible. Some of those changes (for example RFC1583Compatibility flag) can cause a routing loop in mixed environments. It therefore requires careful deployment planning, which is difficult to achieve in complex multivendor topologies. Most importantly, the lack of standard extendable mechanism that facilitates the deployment of non-backward-compatible changes obstructs the development of new protocol extensions.

As a solution for the above described problems, this document proposes an extendable mechanism, which guarantees that the non-backward-compatible functionality is turned on only when supported by all participating routers.

The proposed mechanism is not new; the existing demand circuit functionality [[DEMAND](#)] uses the same approach. This document simply makes the solution generic.

2. Method to deploy non-backward-compatible changes

Each participating router advertises the capability of functionality that it supports in the Router Information LSA as described in [RFC 4970](#) [[OSPF-CAP](#)]. Routers only turn on a new functionality when it is supported by every router within the functionality scope. The routers revert back to their original behavior as soon as one incompatible device is detected.

The scope of functionality could be link, area or AS wide. For link and area wide, the router accordingly originates a link or area scope RI LSA. For AS functionality, an area scope RI LSA is used. To propagate compatibility information across area borders, a new LSA type Area Information is introduced.

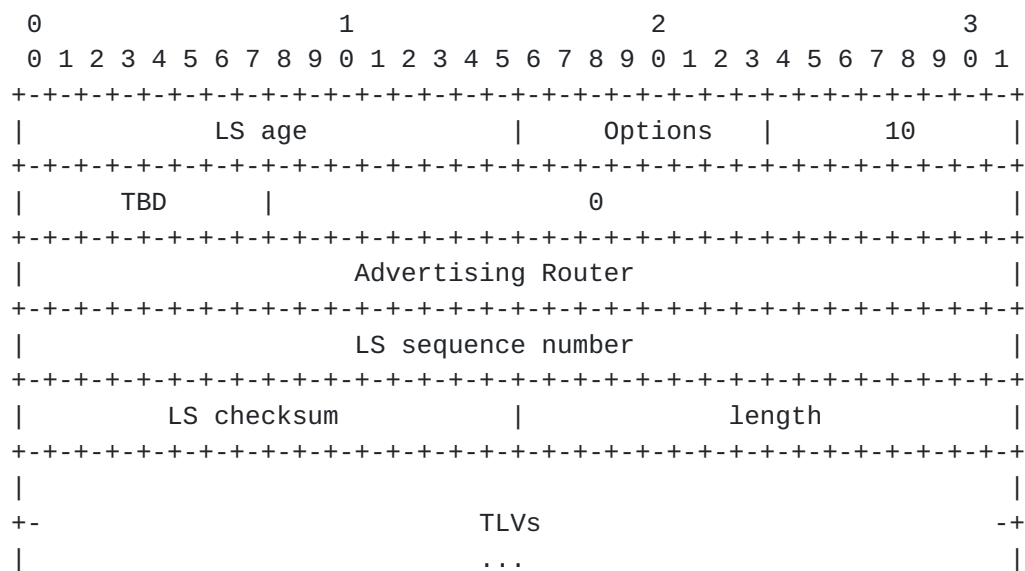
3. Area Information LSA

The Area Border Router inserts a particular capability TLV into an Area Information (AI) LSA to signal that at least one router in the attached areas does not support the functionality. Therefore, the presence of a particular TLV in AI LSA signals the opposite case to the presence of the same TLV in RI LSA. The AI LSA origination algorithm is very similar to the algorithm of Indication-LSA origination [[DEMAND](#)] and outlined below in [Section 3.4](#). The AI LSA format is very similar to RI LSA [[OSPF-CAP](#)]. In OSPFv2, the AI LSA will be implemented with a new opaque LSA type ID. In OSPFv3, the AI

LSA will be implemented with a new LSA type function code. In both protocols, the AI LSA will have an area flooding scope. The exact format of AI LSA is outlined in the sections [3.1](#) and [3.2](#).

[3.1.](#) OSPFv2 Area Information (AI) Opaque LSA

OSPFv2 routers will advertise an area-scoped Opaque-LSA [[OPAQUE](#)]. The OSPFv2 Area Information LSA has a Link-State type of 10 indicating that the flooding scope is area-local, an Opaque type of TBD and Opaque ID of 0.



OSPFv2 Area Information Opaque LSA

The format of the TLVs within the body of an AI LSA is defined in [Section 3.3](#).

[3.2.](#) OSPFv3 Area Information (AI) LSA

The OSPFv3 Area Information LSA has a function code of TBD while the S1/S2 bits are set to 1/0, indicating the area flooding scope for the LSA.

The U bit is set indicating that the OSPFv3 AI 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 AI LSA per flooding scope.

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 a 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.

When new Area 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.

3.4. Area Information LSA origination

Through the origination of AI LSAs, information about the existence of incapable routers propagates from non-backbone areas, to the backbone area and from there to all other areas. The following two cases summarize the requirements for an area border router to originate AI LSAs:

1. Suppose an area border router (Router X) is connected to a non-backbone OSPF area (Area A). Furthermore, assume that Area A has an incapable router i.e. a router LSA without corresponding RI LSA TLV. Then Router X should originate the AI LSAs into all other directly connected areas, including the backbone area, in accordance with the guidelines of [Section 3.4.1](#).

2. Suppose an area border router (Router X) is connected to the backbone OSPF area (Area 0.0.0.0). Furthermore, assume that the backbone has an indication of an existing incapable device via either

- a) the existence of a router LSA without corresponding RI LSA TLV

or

- b) AI LSAs that have been originated by routers other than Router X. Then Router X should originate AI LSAs into all other directly connected non-backbone areas, keeping the guidelines of [Section 3.4.1](#) in mind.

3.4.1. Limiting Area Information LSA origination

The following guidelines should be observed by an area border router (Router X) when originating AI LSAs in order to limit their number. First, AI LSAs with corresponding TLV are not originated into an Area A when A has incapable routers; i.e. router LSAs without corresponding RI LSA TLV. Secondly, if another area border router has originated an AI LSA with corresponding TLV into Area A, and that area border router has a higher OSPF Router ID than Router X (same tie-breaker as for forwarding the address origination; see [Section 12.4.4.1](#) of [OSPF]), then Router X should not originate an AI LSA with corresponding TLV into Area A.

4. Capability negotiation before adjacency is fully formed

For negotiating link scope capability before adjacency is fully formed, link local signaling [[LLS](#)] should be used instead of RI LSA. An example of such a functionality would be a modification to OSPF adjacency formation FSM.

5. Backward Compatibility

The mechanism is backward compatible with the existing OSPF specification. Setting the U bit in OSPFv3 AI LSA allows LSA propagation even if some routers in the area can not decode the LSA content. The Opaque LSA specification [[OPAQUE](#)] also guarantees the propagation of OSPFv2 AI LSA, even if the content is not understood by some of the transit routers.

6. IANA Considerations

The following IANA assignments are to be made from existing registries:

The OSPFv2 opaque LSA option type TBD will need to be reserved for the OSPFv2 AI opaque LSA via IETF Consensus.

OSPFv3 LSA Function Codes TBD will need to be reserved for the OSPFv3 Area Information (AI) LSA via Standards Action.

Both Standards Action and IETF Consensus registration procedures are described in the update of [RFC 2434](#) [I-D.narten-iana-considerations-rfc2434bis].

7. Security Considerations

This document describes a generic mechanism for deployment of non-backward-compatible changes and it introduces Area-Information LSA for AS scope compatibility. The security considerations for those entities are as critical as the topology information currently advertised by the base OSPF protocol. Security considerations for the base OSPF protocol are covered in [[OSPF](#)] and [[OSPFV3](#)].

8. Acknowledgements

The author would like to acknowledge the helpful comments of Cisco OSPF Development team.

This memo is a product of the OSPF Working Group.

9. References

9.1. Normative References

- [LLS] Zinin, A., Roy, A., Nguyen, L., Friedman, B., and D. Yeung, "OSPF Link-Local Signaling", [RFC 5613](#), August 2009.
- [OPAQUE] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), July 2008.
- [OSPF] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [OSPF-CAP] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 4970](#), July 2007.
- [OSPFV3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), July 2008.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.

9.2. Informative References

- [DEMAND] Moy, J., "Extending OSPF to Support Demand Circuits", [RFC 1793](#), April 1995.

Authors' Addresses

Mike Dubrovsky
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: mdubrovs@cisco.com

Rashmi Shrivastava
Cisco Systems
10 West Tasman Drive
San Jose, CA 95134
USA

Email: rashi@cisco.com

Dean Cheng
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: dean.cheng@huawei.com

