

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 14, 2013

A. Retana
Hewlett-Packard Co.
A. Lindem
Ericsson
July 13, 2012

OSPF Incremental Link State Database Synchronization
draft-retana-ospf-ils-01

Abstract

The ability of OSPF to transport non-routing information to be used by other applications was defined by the Opaque LSA Option. In order to not impact the convergence of routing information, this document describes a simple process to incrementally synchronize the routing and non-routing information residing in an OSPF link-state database.

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 January 14, 2013.

Copyright Notice

Copyright (c) 2012 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	3
2.	Requirements Language	3
3.	Incremental LSDB Synchronization Process	3
3.1.	Graceful Restart	4
3.2.	Flooding and Database Synchronization	5
4.	Backward Compatibility	5
5.	IANA Considerations	6
6.	Security Considerations	6
7.	Acknowledgements	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	6
Appendix A.	Changes between the -00 and -01 versions.	7
	Authors' Addresses	7

1. Introduction

Opaque LSAs [[RFC5250](#)] provide the ability for OSPFv2 [[RFC2328](#)] to transport non-routing information to be used by other applications. A similar capability exists in OSPFv3 [[RFC5340](#)] through the use of the U-bit and an appropriate LSA Function Code. Throughout this document Opaque LSAs and ones with unrecognized link-state types will be referred to simply as "opaque".

The presence of opaque information in the OSPF Link-State Database (LSDB) may result in longer database exchange times, especially in cases where the amount of data is significantly larger than the routing-specific information. In order to not impact the convergence of routing information, this document describes a simple process to incrementally synchronize the information residing in an OSPF LSDB. The process uses existing mechanisms.

2. 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 [[RFC2119](#)].

3. Incremental LSDB Synchronization Process

The Incremental LSDB Synchronization (ILS) process consists of the following steps:

LSA Prioritization

The contents of the local LSDB are classified to determine which LSAs require prioritized synchronization.

In general, LSAs containing routing-specific information MUST be classified as requiring prioritized synchronization, while

other LSAs MAY be classified as not requiring it.

All routers in the OSPF routing domain should use the same criteria for LSA prioritization. While this is not required for correct OSPF operation, it is required for deterministic operation of the applications using the opaque LSAs.

Prioritized LSDB Synchronization

This step corresponds to the adjacency establishment process as described in [[RFC2328](#)].

LSAs classified as not requiring prioritized synchronization MUST NOT be included in Database Description (DBD) Packets during the Database Exchange Process. The OSPF routing table structure SHOULD be calculated before moving on to the next step.

Final LSDB Synchronization

In this step, any remaining LSAs in the LSDB SHOULD be synchronized. The routers MUST use the Out-of-Band LSDB Resynchronization [[RFC4811](#)] (OOB Resync) mechanism, which provides a way to resynchronize the LSDB without affecting the advertised neighbor state.

The process is described in terms of LSAs containing (or not) routing-specific information, but it may be generalized to include any other criteria considered significant in the local network and protocol instance.

The last step in the process MAY be used recursively to achieve an incremental LSDB synchronization through different types of data, making it also applicable to environments where only non-routing information exists.

[3.1](#). Graceful Restart

The restart of the OSPF software in a router also presents an opportunity for LSDB synchronization. Because the restarting router is still in the forwarding path, it is important for the routing information in the LSDB to be synchronized as fast as possible. ILS

can be used, with minor modifications, to reduce the synchronization time and the probability of network topology changes.

Graceful OSPF Restart

Graceful OSPF Restart for OSPFv2 [[RFC3623](#)] and OSPFv3 [[RFC5187](#)] don't specify any changes to the adjacency establishment process.

ILS can be used by the Helper Neighbor during the Grace Period; if so, then the Helper Node MUST include any Grace-LSAs in the DBD Packets during the Prioritized LSDB Synchronization step.

OSPF Restart Signaling

OSPF Restart Signaling [[RFC4812](#)] defines a mechanism to inform neighbors about a local restart, in which the LSDB synchronization is achieved using OOB Resync. In other words, the Prioritized LSDB Synchronization step would use OOB Resync if the non-restarting router uses ILS. No other

changes to the process are needed.

[3.2.](#) Flooding and Database Synchronization

In order to maintain OSPF reliable flooding, separate neighbor flooding state must be maintained for the lower priority LSAs and used when determining whether a lower priority LSA has been flooded or put on a neighbor's retransmission list.

The rules for flooding lower-priority LSAs and deleting lower priority MaxAge LSAs are modified for ILS synchronization. Neighbor state MUST be maintained as to whether low-priority LSAs have been synchronized with a given neighbor. In this document, this lower priority state will be referred to as ILS-state. This ILS-state will not advance until a neighbor reaches Full state and will return to Down ILS-state should the neighbor fall from Full state.

If a given neighbor's ILS-state is Exchange or higher, then lower priority LSAs should be flooded to that neighbor. This similar to the general LSA flooding rules in [Section 13.1 of \[RFC2328\]](#).

Consistent with [Section 14.1 in \[RFC2328\]](#), a lower priority MaxAge

LSA cannot be removed the Link State Database if any the router is in Exchange or Loading ILS-state.

If multiple ILS synchronizations are used for synchronization different priorities of LSAs, separate ILS-state MUST be maintained for each priority and the previously described rules MUST be applied at each priority.

[4.](#) Backward Compatibility

The operation of ILS depends on the support of OOB Resync during synchronization; no backwards compatibility issues exist there [[RFC4811](#)]. If OOB Resync is not supported by one of the routers, then the LSDB synchronization would fall back to the adjacency establishment process as described in [[RFC2328](#)].

If OOB Resync is supported, but ILS has not been implemented by all the routers involved, the operation is still backwards compatible. Note that the process ([Section 3](#)) depends on the database description by the local router. In other words, a router may decide to not fully describe the contents of its LSDB to its neighbor during the adjacency establishment process, and later use OOB Resync to incrementally describe the difference; the receiver doesn't need to be aware of ILS. The benefits of ILS may only be partially realized if not supported by all the routers involved in synchronization.

[5.](#) IANA Considerations

This memo includes no request to IANA.

[6.](#) Security Considerations

The process described in this document does not introduce any new security issues into the OSPF protocol.

[7.](#) Acknowledgements

The authors would like to thank Abhay Roy and Liem Nguyen for their comments, and Dimitri Papadimitriou for his comments and for

providing the motivation for this document.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4811] Nguyen, L., Roy, A., and A. Zinin, "OSPF Out-of-Band Link State Database (LSDB) Resynchronization", [RFC 4811](#), March 2007.

8.2. Informative References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [RFC3623] Moy, J., Pillay-Esnault, P., and A. Lindem, "Graceful OSPF Restart", [RFC 3623](#), November 2003.
- [RFC4812] Nguyen, L., Roy, A., and A. Zinin, "OSPF Restart Signaling", [RFC 4812](#), March 2007.
- [RFC5187] Pillay-Esnault, P. and A. Lindem, "OSPFv3 Graceful Restart", [RFC 5187](#), June 2008.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), July 2008.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), July 2008.

Retana & Lindem	Expires January 14, 2013	[Page 6]
-----------------	--------------------------	----------

Internet-Draft	OSPF Incremental LSDB Sync	July 2012
----------------	----------------------------	-----------

Appendix A. Changes between the -00 and -01 versions.

- o Added a new section explaining how to maintain reliable flooding for the low priority LSAs.
- o Updated authors and contact information.

Authors' Addresses

Alvaro Retana
Hewlett-Packard Co.
2610 Wycliff Road
Raleigh, NC 27607
USA

Email: alvaro.retana@hp.com

Acee Lindem
Ericsson
102 Carric Bend Court
Cary, NC 27519
USA

Email: acee.lindem@ericsson.com