

Network Working Group
Request for Comments: 4811
Category: Informational

L. Nguyen
A. Roy
Cisco Systems
A. Zinin
Alcatel-Lucent
March 2007

OSPF Out-of-Band Link State Database (LSDB) Resynchronization

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

OSPF is a link-state intra-domain routing protocol used in IP networks. Link State Database (LSDB) synchronization in OSPF is achieved via two methods -- initial LSDB synchronization when an OSPF router has just been connected to the network and asynchronous flooding that ensures continuous LSDB synchronization in the presence of topology changes after the initial procedure was completed. It may sometime be necessary for OSPF routers to resynchronize their LSDBs. The OSPF standard, however, does not allow routers to do so without actually changing the topology view of the network.

This memo describes a vendor-specific mechanism to perform such a form of out-of-band LSDB synchronization. The mechanism described in this document was proposed before Graceful OSPF Restart, as described in RFC 3623, came into existence. It is implemented/supported by at least one major vendor and is currently deployed in the field. The purpose of this document is to capture the details of this mechanism for public use. This mechanism is not an IETF standard.

Table of Contents

1. Introduction	2
2. Proposed Solution	2
2.1. The LR-Bit	3
2.2. OSPF Neighbor Data Structure	3
2.3. Hello Packets	4
2.4. DBD Packets	4
2.5. Neighbor State Treatment	7
2.6. Initiating OOB LSDB Resynchronization	7
3. Backward Compatibility	7
4. Security Considerations	7
5. IANA Considerations	7
6. References	8
6.1. Normative References	8
6.2. Informative References	8
Appendix A. Acknowledgements	9

1. Introduction

According to the OSPF standard [RFC2328], after two OSPF routers have established an adjacency (the neighbor Finite State Machines (FSMs) have reached Full state), routers announce the adjacency states in their router-Link State Advertisements (LSAs). Asynchronous flooding algorithm ensures that routers' LSDBs stay in sync in the presence of topology changes. However, if routers need (for some reason) to resynchronize their LSDBs, they cannot do that without actually putting the neighbor FSMs into the ExStart state. This effectively causes the adjacencies to be removed from the router-LSAs, which may not be acceptable if the desire is to prevent routing table flaps during database resynchronization. In this document, we provide the means for so-called out-of-band (OOB) LSDB resynchronization.

The described mechanism can be used in a number of situations including those where the routers are picking up the adjacencies after a reload. The process of adjacency preemption is outside the scope of this document. Only the details related to LSDB resynchronization are mentioned herein.

2. Proposed Solution

With this Out-of-Band Resynchronization Solution, the format of the OSPF Database Description (DBD) packet is changed to include a new R-bit indicating OOB LSDB resynchronization. All DBD packets sent during the OOB resynchronization procedure are sent with the R-bit set.

Also, two new fields are added to the neighbor data structure. The first field indicates a neighbor's OOB resynchronization capability. The second indicates that OOB LSDB resynchronization is in process. The latter field allows OSPF implementations to utilize the existing neighbor FSM code.

A bit is occupied in the Extended Options (EO) TLV (see [RFC4813]). Routers set this bit to indicate their capability to support the described technique.

2.1. The LR-Bit

A new bit, called LR (LR stands for LSDB Resynchronization), is introduced to the LLS Extended Options TLV (see [RFC4813]). The value of the bit is 0x00000001; see Figure 1. See the "IANA Considerations" section of [RFC4813] for more information on the Extended Options bit definitions. Routers set the LR-bit to announce OOB LSDB resynchronization capability.




Figure 1. The Options Field

Routers supporting the OOB LSDB resynchronization technique set the LR-bit in the EO-TLV in the LLS block attached to both Hello and DBD packets. Note that no bit is set in the standard OSPF Options field, neither in OSPF packets nor in LSAs.

2.2. OSPF Neighbor Data Structure

A field is introduced into OSPF neighbor data structure, as described below. The name of the field is OOBResync, and it is a flag indicating that the router is currently performing OOB LSDB resynchronization with the neighbor.

The OOBResync flag is set when the router is initiating OOB LSDB resynchronization (see Section 2.6 for more details).

Routers clear the OOBResync flag on the following conditions:

- o The neighbor data structure is first created.
- o The neighbor FSM transitions to any state lower than ExStart.
- o The neighbor FSM transitions to the ExStart state because a DBD packet with the R-bit clear has been received.

- o The neighbor FSM reaches the state Full.

Note that the OOBResync flag may have a TRUE value only if the neighbor FSM is in states ExStart, Exchange, or Loading. As indicated above, if the FSM transitions to any other state, the OOBResync flag should be cleared.

It is important to mention that operation of the OSPF neighbor FSM is not changed by this document. However, depending on the state of the OOBResync flag, the router sends either normal DBD packets or DBD packets with the R-bit set.

2.3. Hello Packets

Routers capable of performing OOB LSDB resynchronization should always set the LR-bit in their Hello packets.

2.4. DBD Packets

Routers supporting the described technique should always set the LR-bit in the DBD packets. Since the Options field of the initial DBD packet is stored in corresponding neighbor data structure, the LR-bit may be used later to check if a neighbor is capable of performing OOB LSDB resynchronization.

The format of type 2 (DBD) OSPF packets is changed to include a flag indicating the OOB LSDB resynchronization procedure. Figure 2 illustrates the new packet format.

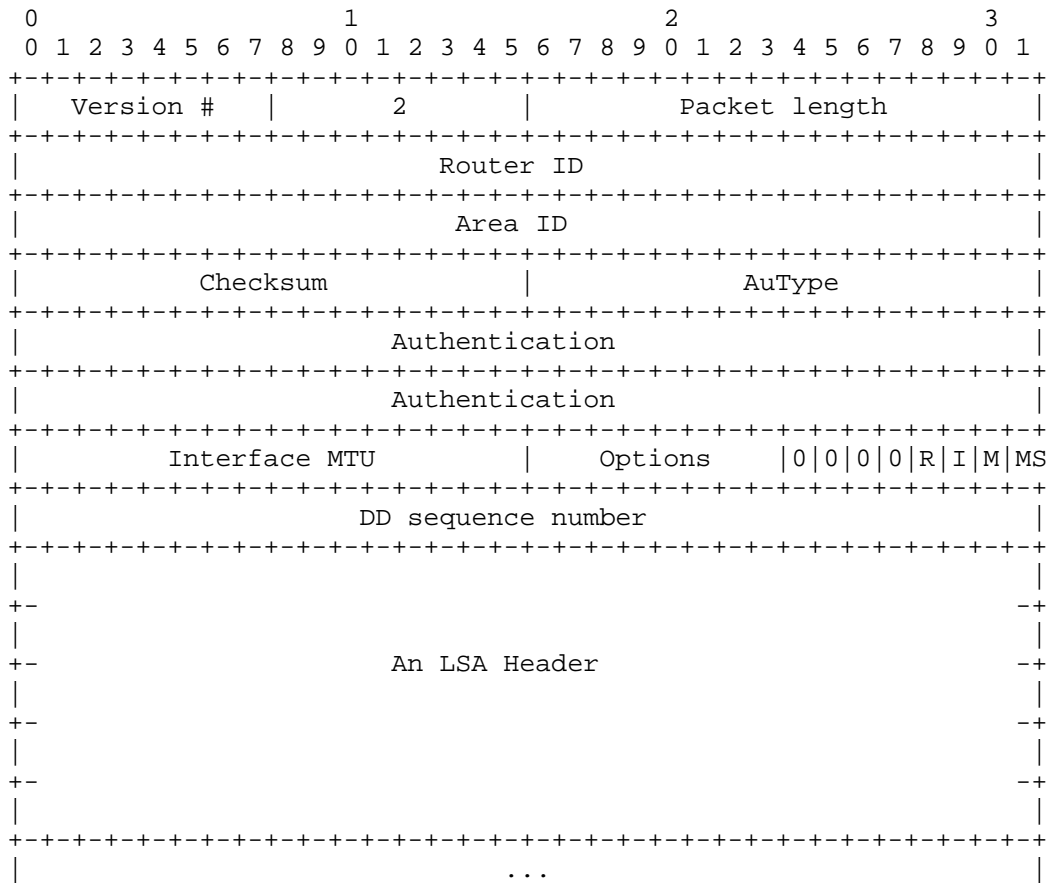


Figure 2. Modified DBD Packet

The R-bit in OSPF type 2 packets is set when the OOBResync flag for the specific neighbor is set to TRUE. If a DBD packets with the R-bit clear is received for a neighbor with active OOBResync flag, the OOB LSDB resynchronization process is canceled and normal LSDB synchronization procedure is initiated.

When a DBD packet is received with the R-bit set and the sender is known to be OOB-incapable, the packet should be dropped and a SeqNumber-Mismatch event should be generated for the neighbor.

Processing of DBD packets is modified as follows:

1. If the OOBResync flag for the neighbor is set (the LSDB resynchronization process has been started) and the received DBD packet does not have the R-bit set, ignore the packet and generate a SeqNumberMismatch event for the neighbor FSM.
2. Otherwise, if the OOBResync flag for the neighbor is clear and the received DBD packet has the R-bit set, perform the following steps:
 - * If the neighbor FSM is in state Full and bits I, M, and MS are set in the DBD packet, set the OOBResync flag for the neighbor, put the FSM in ExStart state, and continue processing the DBD packet as described in [RFC2328].
 - * Otherwise, ignore received DBD packet (no OOB DBD packets are allowed with OOBResync flag clear and FSM in state other than Full). Also, if the state of the FSM is Exchange or higher, generate a SeqNumberMismatch event for the neighbor FSM.
3. Otherwise, process the DBD packet as described in [RFC2328].

During normal processing of the initial OOB DBD packet (with bits R, I, M, and MS set), if the receiving router is selected to be the Master, it may speed up the resynchronization process by immediately replying to the received packet.

It is also necessary to limit the time an adjacency can spend in ExStart, Exchange, and Loading states with OOBResync flag set to a finite period of time (e.g., by limiting the number of times DBD and link state request packets can be retransmitted). If the adjacency does not proceed to Full state before the timeout, it is indicative that the neighboring router cannot resynchronize its LSDB with the local router. The requesting router may decide to stop trying to resynchronize the LSDB over this adjacency (if, for example, it can be resynchronized via another neighbor on the same segment) or to resynchronize using the legacy method by clearing the OOBResync flag and leaving the FSM in ExStart state. The neighboring router may decide to cancel the OOB procedure for the neighbor.

2.5. Neighbor State Treatment

An OSPF implementation supporting the described technique should modify the logic consulting the state of a neighbor FSM as described below.

- o FSM state transitioning from and to the Full state with the OOBResync flag set should not cause origination of a new version of router-LSA or network-LSA.
- o Any explicit checks for the Full state of a neighbor FSM for the purposes other than LSDB synchronization and flooding should treat states ExStart, Exchange, and Loading as state Full, provided that OOBResync flag is set for the neighbor. (Flooding and MaxAge-LSA-specific procedures should not check the state of the OOBResync flag, but should continue consulting only the FSM state.)

2.6. Initiating OOB LSDB Resynchronization

To initiate out-of-band LSDB resynchronization, the router must first make sure that the corresponding neighbor supports this technology (by checking the LR-bit in the Options field of the neighbor data structure). If the neighboring router is capable, the OOBResync flag for the neighbor should be set to TRUE and the FSM state should be forced to ExStart.

3. Backward Compatibility

Because OOB-capable routers explicitly indicate their capability by setting the corresponding bit in the Options field, no DBD packets with the R-bit set are sent to OOB-incapable routers.

The LR-bit itself is transparent for OSPF implementations and does not affect communication between routers.

4. Security Considerations

The described technique does not introduce any new security issues into the OSPF protocol.

5. IANA Considerations

Please refer to the "IANA Considerations" section of [RFC4813] for more information on the Extended Options bit definitions.

6. References

6.1. Normative 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.

6.2. Informative References

- [RFC4813] Friedman, B., Nguyen, L., Roy, A., Yeung, D., and A. Zinin, "OSPF Link-Local Signaling", RFC 4813, March 2007.

Appendix A. Acknowledgments

The authors would like to thank Acee Lindem, Russ White, Don Slice, and Alvaro Retana for their valuable comments.

Authors' Addresses

Liem Nguyen
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: lhnguyen@cisco.com

Abhay Roy
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: akr@cisco.com

Alex Zinin
Alcatel-Lucent
Mountain View, CA
USA
EMail: alex.zinin@alcatel-lucent.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.