

Network Working Group
Request for Comments: 4812
Category: Informational

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

OSPF Restart Signaling

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. Routers find new and detect unreachable neighbors via the Hello subprotocol. Hello OSPF packets are also used to ensure two-way connectivity within time. When a router restarts its OSPF software, it may not know its neighbors. If such a router sends a Hello packet on an interface, its neighbors are going to reset the adjacency, which may not be desirable in certain conditions.

This memo describes a vendor-specific mechanism that allows OSPF routers to inform their neighbors about the restart process. Note that this mechanism requires support from neighboring routers. 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. Introduction2
- 2. Proposed Solution2
 - 2.1. Sending Hello Packets with the RS-bit Set3
 - 2.2. Receiving Hello Packets with the RS-Bit Set3
 - 2.3. Ensuring Topology Stability4
- 3. Backward Compatibility4
- 4. Security Considerations4
- 5. IANA Considerations4
- 6. References5
 - 6.1. Normative References5
 - 6.2. Informative References5
- Appendix A. Acknowledgements6

1. Introduction

While performing a graceful restart of OSPF software [RFC3623], routers need to prevent their neighbors from resetting their adjacencies. However, after a reload, routers may not be aware of the neighbors they had adjacencies with in their previous incarnations. If such a router sends a Hello packet on an interface and this packet does not list some neighbors, those neighbors will reset the adjacency with the restarting router.

This document describes a technique that allows restarting routers to inform their neighbors that they may not know about some neighbors yet and the absence of some router IDs in the Hello packets should be ignored.

2. Proposed Solution

With this Restart Signaling Solution, a new bit, called RS (restart signal), is introduced into the Extended Options (EO) TLV in the Link-Local Signaling (LLS) block (see [RFC4813]). The value of this bit is 0x00000002; see Figure 1 below.




Figure 1. Bits in Extended Options TLV

For a definition of the LR-bit, see [RFC4811].

2.1. Sending Hello Packets with the RS-bit Set

OSPF routers should set the RS-bit in the EO-TLV attached to a Hello packet when it is not known that all neighbors are listed in this packet, but the restarting router wants them to preserve their adjacencies. The RS-bit must not be set in Hello packets longer than RouterDeadInterval seconds.

2.2. Receiving Hello Packets with the RS-Bit Set

When an OSPF router receives a Hello packet containing the LLS block with the EO-TLV that has the RS-bit set, the router should skip the two-way connectivity check with the announcing neighbor (i.e., the router should not generate a 1-WayReceived event for the neighbor if it does not find its own router ID in the list of neighbors as described in Section 10.5 of [RFC2328]), provided that the neighbor Finite State Machine (FSM) for this neighbor is in the Full state.

The router should also send a unicast Hello back to the sender in reply to a Hello packet with RS-bit set. This is to speed up learning of previously known neighbors. When sending such a reply packet, care must be taken to ensure that the RS-bit is clear in it.

Two additional fields are introduced in the neighbor data structure: RestartState flag and ResyncTimeout timer. RestartState flag indicates that a Hello packet with the RS-bit set has been received and the local router expects its neighbor to go through the Link State Database (LSDB) resynchronization procedure using [RFC4811]. ResyncTimeout is a single-shot timer limiting the delay between the first seen Hello packet with the RS-bit set and initialization of the LSDB resynchronization procedure. The length of ResyncTimeout timer is RouterDeadInterval seconds.

When a Hello packet with the RS-bit set is received and RestartState flag is not set for the neighbor, the router sets RestartState flag and starts ResyncTimeout timer. If ResyncTimeout expires, RestartState flag is cleared and a 1-WayReceived event is generated for the neighbor. If, while ResyncTimeout timer is running, the neighbor starts LSDB resynchronization procedure using [RFC4811], ResyncTimeout timer is canceled. The router also clears RestartState flag on completion of the LSDB resynchronization process.

Two or more routers on the same segment cannot have Hello packets with the RS-bit set at the same time, as can be the case when two or more routers restart at about the same time. In such a scenario, the routers should clear the RestartState flag, cancel the ResyncTimeout timer, and generate a 1-WayReceived event.

2.3. Ensuring Topology Stability

Under certain circumstances, it might be desirable to stop announcing the restarting router as fully adjacent if this may lead to possible routing loops. In order to provide this functionality, a configurable option is provided on the neighboring routers that instructs the OSPF process to follow the logics described below.

When an OSPF router schedules a routing table calculation due to a change in the contents of its LSDB, it should also reset all adjacencies with restarting routers (those with RestartState set to TRUE) by clearing the RestartState neighbor flags, canceling ResyncTimeout timers (if running), and generating the 1-WayReceived events for the neighbor FSMs.

3. Backward Compatibility

The described technique requires cooperation from neighboring routers. However, if neighbors do not support this technique, they will just reset the adjacency.

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.

[RFC4811] Nguyen, L., Roy, A., and A. Zinin, "OSPF Out-of-Band Link State Database (LSDB) Resynchronization", RFC 4811, March 2007.

Appendix A. Acknowledgments

The authors would like to thank John Moy, 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.