Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 4, 2017

J. Dong
X. Zhang
Huawei Technologies
Z. Li
China Mobile

October 31, 2016

OSPF LSA Flushing Problem Statement
draft-dong-ospf-maxage-flush-problem-statement-01

Abstract

In OSPF protocol, Link State Advertisements (LSAs) are exchanged in Link State Update (LSU) packets to achieve link state database (LSDB) synchronization and consistent route calculation. OSPF protocol specifies several scenarios in which an LSA is flushed with the LS age field set to MaxAge. In some cases, the flushing of MaxAge LSAs may cause flooding storm of OSPF packets and severely impact the services provided by the network.

This document describes the problem of OSPF LSA flushing, and ask for solutions to solve this problem.

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

Table of Contents

1. Introduction

In OSPF protocol [RFC2328], Link State Updates (LSAs) are exchanged in Link State Update (LSU) packets to achieve link state database (LSDB) synchronization and consistent route calculation. OSPF specifies several scenarios in which an LSA is flushed with the LS age field set to MaxAge. In some cases, the flushing of MaxAge LSAs may cause flooding storm of OSPF packets and severely impact the services in the network. Since the MaxAge LSA may be flushed by any OSPF router, usually it would take a long time for troubleshooting and could cause huge damage to both the network provider and its customers.

2. Typical Scenarios of LSA Flushing

[RFC2328] specifies several scenarios in which an LSA should be flushed with the LS age field set to MaxAge. Under normal circumstances, the LSA flushing happens when the LS age of an LSA naturally reaches MaxAge, this can be done by any OSPF router. Since OSPF router would generate a new instance of the self-originated LSA when its LS age reaches LSRefreshTime, which is usually the half of the value of MaxAge, the naturally aging to MaxAge case would only happen when the originator of the LSA is not reachable in the network and cannot refresh the LSA.

Another case of LSA flushing is "Premature aging", which is to set the LS age of a self-originated LSA to MaxAge and then flood the LSA. Premature aging is used when the self-originated LSA's sequence number field is about to wrap, or all the external routes previously advertised by the LSA are no longer reachable. Premature aging and flushing of LSA can also happen when a router is changed from the Designated Router (DR) to a non-DR, or in some rare cases the router's Router ID is changed.

Field experience has shown several circumstances where MaxAge LSA flushing may be generated by the misbehaved router in the network. For example, the LS age may be corrupted to reach the MaxAge much earlier than normally expected. This is difficult to detect with the existing OSPF checksum mechanism, as the LS age field is excluded from the checksum calculation of LSA. Besides, OSPF cryptographic authentication can not detect the corruption of the LS age field if it happens before the LSA is assembled to LSU packet.

3. Consequence of LSA Flushing

While MaxAge LSA flushing is important for fast convergence and the consistency of the Link-State DataBase (LSDB) of all OSPF routers, as shown in several accidents happened in the production network, improper LSA flushing can have severe impact to the network and the services provided by the network. This section evaluates the impacts of MaxAge LSA flushing.

According to section 14 of [RFC2328], the MaxAge LSA can be flushed by any router, no matter whether this LSA is self-originated or not. According to the flooding scope of the LSA, this MaxAge LSA would be flooded either in the whole routing domain or in the specific area. On all the routers receiving this MaxAge LSA, this would cause the old LSA instance being replaced, and consequently triggers route calculation and installation. When the MaxAge LSA is received by the originating router of this LSA, the originating router would increase the LSA's LS sequence number one past the received LS sequence

number, and originate a new instance of the LSA. If the LSA flushing is due to systematic problem and cannot recover automatically, this flooding and processing would last forever, which severely impacts network reachability and stability. Since OSPF is the fundamental protocol to build the infrastructure for other protocols such as BGP, LDP, etc., and various services provided by the network, it will cause huge damage to both the network provider and its customers.

As the MaxAge LSA may be flushed by any OSPF router, usually it would take a long time for troubleshooting to locate the misbehaved router in the network, and during this time the LSA flushing could have caused huge damage to both the network provider and its customers.

4. Requirements on Potential Solutions

Considering the importance of OSPF protocol to the networks and the services carried in the networks, and the potential severe impact of MaxAge LSA flooding, this document calls for solutions to protect against or mitigate the impact of improper MaxAge LSA flushing.

The potential solutions can be classified into two categories, and the requirements are provided in following sections respectively.

4.1. Solution for Problem Localization

Since OSPF allows the flushing of non-self originated LSAs, for troubleshooting and problem localization, some mechanism to identify the misbehaved router quickly is needed. If the improper MaxAge LSA flushing is caused by systematic problem, operators would need to locate the misbehaved router and shut it down to stop the flooding storm.

[RFC6232] proposes to add the Purge Originator Identification (POI) TLV into IS-IS Purge LSPs to identify the originator of IS-IS Purges. Although a similar TLV may be added into the OSPF extended LSAs as defined in [RFC7684] and [I-D.ietf-ospf-ospfv3-lsa-extend], the structure of the legacy OSPF LSAs as defined in [RFC2328] is not TLV-based and such mechanism does not apply. Some problem localization solution which is backward compatible and applicable to all the OSPF LSAs would be preferred.

4.2. Solution for Impact Mitigation

Since the flooding storm caused by improper LSA flushing can have severe impact to network stability and the services provided by the network, it is important to alleviate such impact even before the root cause or the misbehaved router can be identified. In addition, some problem localization mechanisms may rely on the availability of

the network, which means the impact mitigation mechanism is necessary to ensure that the problem localization mechanisms do work when severe flooding storm caused by LSA flushing happens in the network.

It is important that the impact mitigation solution is backward compatible and can support incremental deployment. Preferably, the mitigation solution should not delay the route convergence triggered by normal LSA flushing.

5. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

6. Security Considerations

This document describes the problem of MaxAge LSA flushing, which in some cases is due to the lack of integrity protection of the LS age field. The LS age field may be altered as a result of software or hardware problem, such modification cannot be detected by LSA checksum nor OSPF packet cryptographic authentication. LSA flushing could have severe impact on network stability and the services provided by the network. This may be considered as a security vulnerability.

7. Acknowledgements

The authors would like to thank Bruno Decraene, Acee Lindom and Les Ginsberg for the discussion on this topic.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <http://www.rfc-editor.org/info/rfc2328>.

8.2. Informative References

[I-D.ietf-ospf-ospfv3-lsa-extend]
Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", draft-ietf-ospf-ospfv3-lsa-extend-13 (work in progress), October 2016.

[RFC6232] Wei, F., Qin, Y., Li, Z., Li, T., and J. Dong, "Purge Originator Identification TLV for IS-IS", RFC 6232, DOI 10.17487/RFC6232, May 2011, <http://www.rfc-editor.org/info/rfc6232>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <http://www.rfc-editor.org/info/rfc7684>.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No.156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

Xudong Zhang
Huawei Technologies
Huawei Campus, No.156 Beiqing Rd.
Beijing 100095
China

Email: zhangxudong@huawei.com

Zhenqiang Li
China Mobile
No.32 Xuanwumenxi Ave., Xicheng District
Beijing 100032
China

Email: li_zhenqiang@hotmail.com