Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 26, 2012

J. Dong
X. Zhang
Huawei Technologies
October 24, 2011

Purge LSA for OSPF flushing
draft-dong-ospf-purge-lsa-00

Abstract

In some scenarios current OSPF flushing mechanism may incur problem of delaying the deletion of invalid Link State Advertisement (LSA) and result in desynchronization of link state database. This document proposes a backward compatible solution 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

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 26, 2012.

Copyright Notice

Table of Contents

1. Introduction

The LSA flushing mechanism of OSPF is described in [RFC2328]. In some scenarios such as route oscillation, this flushing mechanism may incur problem of delaying the deletion of invalid LSA on some routers, and result in desynchronization of link state database.

This document proposes a backward compatible solution to solve this problem.

2. Problem with OSPF Flushing Mechanism

As described in section 14.1 of [RFC2328], an LSA can be flushed from the routing domain by setting its LS age to MaxAge, while leaving its LS sequence number alone, and then reflooding the LSA. The MaxAge LSA must be removed immediately from the router's link state database as soon as both a) it is no longer contained on any neighbor Link state retransmission lists and b) none of the router's neighbors are in states Exchange or Loading. And section 12.1.6 of [RFC2328] specifies that "As soon as this flood has been acknowledged by all adjacent neighbors, a new instance can be originated with sequence number of InitialSequenceNumber." Thus after the MaxAge LSA is removed, a new instance with sequence number equal to InitialSequenceNumber would be originated.

Under some scenarios such as route oscillation, such procedure may delay the deletion of invalid LSA in the link state database of some routers thus cause desynchronization of link state database. One example is described as below:

a. Router X and router Y formed OSPF adjacency, X advertised an AS-external-LSA with sequence number set to InitialSequenceNumber to Y. Then route oscillation happens to this external route.

b. X flushes the AS-external-LSA by setting the LS age to MaxAge and sends it to Y. After some time X does not receive the acknowledgment from Y, then X retransmits the flushing LSA to Y.

c. Y receives the first flushing LSA and sends Acknowledgment back to X. Then Y receives the second flushing LSA and sends back the second acknowledgment.

d. X receives the first Acknowledgment and remove the flushing LSA from its link state database. Then X originates this AS-external-LSA with InitialSequenceNumber and sends it to Y. Soon the LSA is flushed again due to route oscillation. After doing this, X receives the second acknowledgment for the previous flushing from Y, then the

flushing LSA is removed from X's link state database.

e. Y receives this new LSA and install it into its link state database. Then in a short time Y receives the flushing LSA. If the time interval between this flushing LSA and previous LSA is shorter than MinLSArrival, this flushing LSA would be discarded by Y.

As a result, router X does not have this LSA in its link state database, but router Y will keep this LSA until the LS age reaches MaxAge. Such desynchronization of link state database may cause traffic blackholing or forwarding loops.

The root cause of this problem is that after LSA flushing the router would set the LS sequence number of the new LSA instance back to InitialSequenceNumber. And in some cases the LSA updates and the acknowledgments may become out of sync.

3. Proposed Flushing Mechanism

This section defines a flushing mechanism to solve the problem described in section 2. This flushing mechanism SHOULD be used when LS sequence number is not to wrap. When it is time for the LS sequence number is to wrap, procedures in Section 14.1 of [RFC2328] MUST be used.

When an LSA needs to be flushed, the router SHOULD keep only the header of the LSA, set its LS age to 0 and increment the LS sequence number by 1. Such an LSA is called Purge LSA. The router SHOULD flood this Purge LSA to neighbors for LSA flushing. The Purge LSA SHOULD be retained in the database until the LS age reaches MaxAge.

When a new instance of the LSA is to be originated, if a corresponding Purge LSA exists in the database, the router SHOULD advance the LSA's LS sequence number one past the LS sequence number of the corresponding Purge LSA. Then the corresponding Purge LSA SHOULD be removed from the router's link state database.

4. Backward Compatibility

The proposed flushing mechanism is backward compatible with legacy OSPF implementations. Legacy OSPF routers would treat the received Purge LSA as a newer instance than its database copy, since the sequence number is larger. And since the Purge LSA contains only LSA header, it will not be used in the routing table calculation. When the LS age of this purge LSA reaches MaxAge, it would be removed from the router's link state database.

5. IANA Considerations

This document makes no request of IANA.

6. Security Considerations

This document does not change the security properties of OSPF.

7. Acknowledgements

The authors would like to thank Yongming Fu and Yong Miao for their helps and suggestions to this document.

8. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

Authors' Addresses

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

Email: jie.dong@huawei.com

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

Email: zhangxudong@huawei.com