

OSPF Database Exchange Summary List Optimization

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.

Abstract

This document describes a backward-compatible optimization for the Database Exchange process in OSPFv2 and OSPFv3. In this optimization, a router does not list a Link State Advertisement (LSA) in Database Description packets sent to a neighbor, if the same or a more recent instance of the LSA was listed in a Database Description packet already received from the neighbor. This optimization reduces Database Description overhead by about 50% in large networks. This optimization does not affect synchronization, since it only omits unnecessary information from Database Description packets.

1. Introduction

In OSPFv2 [RFC2328] and OSPFv3 [RFC2740], when two neighboring routers become adjacent, they synchronize their link-state databases via the Database Exchange process. Each router sends the other router a set of Database Description (DD) packets that describes the router's link-state database. This is done by listing each LSA (i.e., including the header of each LSA) in one of the sent DD packets. This procedure allows each router to determine whether the other router has newer LSA instances that should be requested via Link State Request packets.

The optimization simply observes that it is not necessary for a router (master or slave) to list an LSA in a DD packet if it knows the neighbor already has an instance of the LSA that is the same or more recent (and therefore will not request the LSA). To avoid listing such LSAs in DD packets, when an LSA is listed in a DD packet received from the neighbor, and the Database summary list for the neighbor has an instance of the LSA that is the same as or less recent than the one received, the LSA is removed from the summary list.

The optimization, called the Database Exchange summary list optimization, does not affect synchronization, since the LSAs that are omitted from DD packets are unnecessary. The optimization is fully backward compatible with OSPF. The optimization reduces Database Description overhead by about 50% in large networks in which routers are usually already nearly synchronized when they become adjacent, since it reduces the total number of LSA headers exchanged by about one-half in such networks. The optimization is especially beneficial in large networks with limited bandwidth, such as large mobile ad hoc networks.

2. Specification of Optimization

The Database Exchange summary list optimization is defined by modifying Section 10.6 "Receiving Database Description Packets" of RFC 2328 as follows. The second-to-last paragraph of Section 10.6 is replaced with the following augmented paragraph:

When the router accepts a received Database Description Packet as the next in sequence, the packet contents are processed as follows. For each LSA listed, the LSA's LS type is checked for validity. If the LS type is unknown (e.g., not one of the LS types 1-5 defined by this specification), or if this is an AS-external-LSA (LS type = 5) and the neighbor is associated with a stub area, generate the neighbor event SeqNumberMismatch and stop processing the packet. Otherwise, the router looks up the LSA in its database to see whether it also has an instance of the LSA. If it does not, or if the database copy is less recent, the LSA is put on the Link state request list so that it can be requested (immediately or at some later time) in Link State Request Packets. In addition, if the Database summary list contains an instance of the LSA that is the same as or less recent than the listed LSA, the LSA is removed from the Database summary list.

The above additional step (which updates the Database summary list) may be performed either before or after the router looks up the listed LSA in its database and possibly adds the LSA to the Link state request list. However, to implement the optimization, the additional step must be performed for each LSA listed in the received DD packet (to fully update the Database summary list) before the next DD packet is sent in response.

Although the optimization does not require that LSAs be listed in DD packets in any particular order, faster lookup of LSAs in the Database summary list may be possible if LSAs are listed in the same order by all routers. If such an ordering is used, then to encourage different implementations to use the same ordering, this document recommends that LSAs be listed in lexicographically increasing order of (LS type, Link State ID, Advertising Router) for OSPFv2 and (LS type, Advertising Router, Link State ID) for OSPFv3.

3. Example

This section describes an example to illustrate the Database Exchange summary list optimization. Assume that routers RT1 and RT2 already have identical databases when they start Database Exchange. Also assume that the list of LSA headers for the database fits into two DD packets. Then, the standard Database Exchange is as follows when RT1 is the first to change the neighbor state to ExStart. Note that each router sends two full DD packets.

RT1 (slave)		RT2 (master)
ExStart	Empty DD (Seq=x,I,M,Master) ----->	
	Empty DD (Seq=y,I,M,Master) <-----	ExStart
Exchange	Full DD (Seq=y,M,Slave) ----->	
	Full DD (Seq=y+1,M,Master) <-----	Exchange
	Full DD (Seq=y+1,Slave) ----->	
	Full DD (Seq=y+2, Master) <-----	
Full	Empty DD (Seq=y+2, Slave) ----->	
		Full

If the optimization is used, when RT2 receives the first full DD packet from RT1, it removes from its summary list all LSAs that are listed in the DD packet. Then RT2 sends a DD packet that lists the remaining LSAs (since all of the LSA headers fit into two DD packets). When RT1 receives this DD packet, it removes these remaining LSAs from its summary list (causing it to be empty) and sends an empty DD packet to RT2.

With the optimization, each router sends only one full DD packet instead of two, as shown below.

RT1 (slave)		RT2 (master)
ExStart	Empty DD (Seq=x,I,M,Master) ----->	
	Empty DD (Seq=y,I,M,Master) <-----	ExStart
Exchange	Full DD (Seq=y,M,Slave) ----->	
	Full DD (Seq=y+1,Master) <-----	Exchange
Full	Empty DD (Seq=y+1, Slave) ----->	
		Full

4. Security Considerations

This document does not raise any new security concerns.

5. IANA Considerations

This document specifies a simple backward-compatible optimization for OSPFv2 and OSPFv3 that does not require any new number assignment.

6. Normative References

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

[RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.

7. Acknowledgments

The recommendation to list LSAs in lexicographical order was suggested by Acee Lindem.

Author's Address

Richard G. Ogier
 EMail: rich.ogier@earthlink.net

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

