

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: March 18, 2011

H. Chen
Huawei Technology, Inc.
September 14, 2010

Intelligent OSPF Link State Database Exchange
draft-chen-ospf-intelligent-db-exch-01.txt

Abstract

This document presents an intelligent database exchange mechanism for the database exchange procedure in OSPFv2 and OSPFv3. This mechanism is backward compatible. It eliminates the unnecessary database exchanges through using the existing reachability information calculated from the link state database and the un-reachability information about routers recorded. It significantly speeds up the establishment of the full adjacency between two routers in some situations.

Status of this Memo

This Internet-Draft is submitted to IETF 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 March 18, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Eliminating Whole Link State Database Exchange	3
1.2.	Eliminating Part of Link State Database Exchange	4
2.	Terminology	4
3.	Conventions Used in This Document	5
4.	Intelligent Link State Database Exchange	5
5.	Changes to OSPF Protocols	5
5.1.	Changes to Creation of Database Summary List	6
5.2.	Changes to Processing of DD Packet Contents	6
5.3.	New Data Structures and Procedures	7
6.	Some Details about Implementations of New Procedures	7
6.1.	Finding and Recording Failure Time and LSA Change Time	8
6.1.1.	Finding and Recording Failure Time and LSA Change Time	8
6.1.2.	A Simpler Method for Finding Failure Time	9
6.1.3.	Finding and Recording LSA Change Time	10
6.2.	Reusing LSA Age, Finding and Recording Adjust Time	11
7.	Security Considerations	11
8.	IANA Considerations	11
9.	Acknowledgement	12
10.	References	12
10.1.	Normative References	12
10.2.	Informative References	12
	Author's Address	12

1. Introduction

If a full adjacency is to be formed between two OSPF routers, their link state databases will be synchronized through the database exchange procedure described in OSPFv2 [RFC2328] and OSPFv3 [RFC2740]. Each of the routers sends the other the summary of its database through a set of Database Description packets containing the header of every LSA in its database. For every LSA header received in the Database Description packets, the router compares it with the corresponding LSA instance in its database and sends the other router a request for the LSA if the LSA instance in its database is older. The adjacency becomes full when the router finishes sending the summary of its database and processing all the Database Description packets from the other router and gets all the LSAs it requested.

1.1. Eliminating Whole Link State Database Exchange

In some situations, the whole link state database exchange is unnecessary. For example, in the case illustrated in the figure below, we suppose that full adjacencies between routers RT3 and RT4, RT4 and RT5, and RT5 and RT6 have been established, and a new full adjacency between RT3 and RT6 is to be formed over the link directly connecting them. In this case, RT3 and RT6 are reachable each other through RT4 and RT5. They do not need to exchange their link state databases at all since they are reachable each other and already have the same database.




Figure 1: A Full Adjacency between RT3 and RT6 to be Formed

For a large database, eliminating the unnecessary database exchange will significantly speed up the establishment of the full adjacency and save lots of link bandwidth for the transmission of the

3. Conventions Used in This Document

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](#).

4. Intelligent Link State Database Exchange

The intelligent OSPF Link State database exchange mechanism eliminates the unnecessary Database Description packets exchanges and processes through using the existing reachability information calculated from the link state database and the un-reachability information about routers recorded.

When two OSPF routers are going to bring up a full adjacency between them, each of them checks whether it is reachable to the other through using its route table calculated from its link state database. If it can reach the other router, it does not need to send the other any Database Description packets that contain the summary of its database since two routers have the same database. In this case, all the unnecessary Database Description packets exchanges and processes are eliminated. The full adjacency between two routers is formed almost immediately.

If one router can not reach to the other now but it could reach the other at time t and the difference between the current time and t is less than the LSA maximum age MaxAge (3600 seconds), then it does not need to send the other the header of every LSA in its database through Database Description packets. It just needs to send the other the headers of the LSAs that have been changed in its database since time t at which the other router became unreachable. During the intelligent database exchange, when one router detects that the other router was restarted after time t through some way such as comparing the router LSA generated by the other router with its corresponding LSA instance in its link state database, it stops its current intelligent database exchange and starts a new normal database exchange with the other router.

During the intelligent database exchange between two routers, if one router detects that the other router becomes unreachable, it stops its current intelligent database exchange and starts a new normal database exchange with the other router.

5. Changes to OSPF Protocols

The changes to the OSPF protocols include three parts. The first

part is to modify the creation of the neighbor database summary list. The second is to change the processing of a Database Description packet contents. The third is to add some data structures and procedures.

5.1. Changes to Creation of Database Summary List

In the OSPF protocols, when a local router is going to form a full adjacency with a neighboring router, it creates the neighbor database summary list that has the contents of its entire area link state database. The area link state database consists of the router-LSAs, network-LSAs and summary-LSAs contained in the area structure, along with the AS-external-LSAs contained in the global structure. The intelligent database exchange mechanism modifies the creation of the summary list as follows:

If the local router can reach the neighboring router now, then the neighbor database summary list is empty; otherwise (i.e., the router can not reach the neighboring router now), if the router could reach the neighboring router at time t and the difference between the current time and t is less than the LSA maximum age MaxAge, then the neighbor database summary list must have the contents of the LSAs that have been changed in its link state database since time t at which the neighboring router became unreachable; otherwise, the neighbor database summary list has the contents of its entire area link state database.

5.2. Changes to Processing of DD Packet Contents

In the original OSPF protocols [RFC2328, [RFC2740](#)], when the local router accepts a received Database Description Packet as the next in sequence, one part of processing of the packet contents is that the router looks up the LSA contained in the packet 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 in Link State Request Packets. This part of processing of the packet contents is modified as follows:

If the local router can reach the neighboring router now, it does not look up any LSA contained in the packet in its database to see whether it also has an instance of the LSA or if the database copy is less recent. Otherwise, the local router follows the processing of the packet contents as described in OSPF protocols.

5.3. New Data Structures and Procedures

In addition to the above modification, some new data structures and procedures should be added to provide the following functions:

1. Record function, which records the information about:
 - * Every tuple $\langle r, tu \rangle$, where r is the remote router in the entire area that became unreachable from the local router at time tu and the difference between the current time and tu is less than the LSA maximum age MaxAge (3600 seconds).
 - * The time tc for every LSA at which the LSA is changed if there exist some remote routers that became unreachable less than MaxAge ago. It is not necessary to record time tc for any LSA if there is not any remote router that became unreachable within the past MaxAge (3600 seconds).
2. Retrieve function, which provides the information about:
 - * For a given unreachable remote router r , the time tu at which the router r became unreachable.
 - * For a given time tu , all the LSAs in the link state database that have been changed since then. Normally tu is the time when a remote router became unreachable.
3. Delete function, which removes the information about:
 - * Every tuple $\langle r, tu \rangle$ when the remote router r becomes reachable or the difference between the current time and tu (where tu is the time when r became unreachable) is equal to or greater than the LSA maximum age MaxAge (3600 seconds).
 - * The time tc for every LSA recorded if there is not any remote router that are unreachable.

During the full adjacency establishment between two routers, either of them may restart to form the adjacency in the normal way as described in the OSPF protocols [[RFC2328](#)] and [[RFC2740](#)] if it detects that the reachability between them is broken.

6. Some Details about Implementations of New Procedures

This section describes some implementation options for the record function. The implementations of the retrieve and delete functions depend on the implementation of the record function to some extent

and should be trivial after the details of the implementation of the record function are determined.

One implementation finds the failure time and LSA change time and records them accordingly. It finds time t_u for every remote router r at which r became unreachable because some failures happened at time t_u and records the information about $\langle r, t_u \rangle$. It also finds time t_c for every LSA at which the LSA is changed and records the information about t_c for the LSA.

Another implementation reuses some of the existing information in the link state database such as LSA age, and finds some other information such as failure time and records them.

6.1. Finding and Recording Failure Time and LSA Change Time

This subsection specifies two methods for finding the failure time. One is more accurate. The other is simpler. In addition, it describes a method for calculating the time for every LSA at which it is changed.

If a remote router r becomes unreachable from reachable after the shortest path first algorithm is run against the link state database, then tuple $\langle r, t_u \rangle$ is recorded, where t_u is the time at which a failure such as a link down happened that leads to the router r unreachable. Failures can be classified into two classes: link/interface failures and node failures. When an OSPF router detects a failure, it will regenerate and flood LSAs for the failure. Suppose that t_i is the maximum time delay for the OSPF router to detect an interface failure and t_n is the maximum time delay for the OSPF router to detect a node failure. For an LSA, assume that t_p is the time delay for the LSA to reach the local router from its origin.

6.1.1. Finding and Recording Failure Time and LSA Change Time

For a point to point link failure in the network, two router LSAs with the link down are regenerated and flooded. One is by each of two end routers of the link. For a broadcast link or NBMA link failure, one network LSA and one or more router LSAs with the information about link down are generated and flooded. The network LSA is generated by the designated router and the router LSAs are generated by the routers attached to the link.

Among these LSAs, suppose that t_r is the earliest time at which one of the LSAs is received. t_p is less than the age of the LSA, which can be used as an estimated value of t_p . We can also estimate t_p as $(255 - \text{TTL of the LSA update packet})$. We may select the smaller one between these two estimated values for t_p . In all these cases, the

exact t_p is less than its estimated value. The time t_u at which the interface failed can be estimated as $t_u = (t_r - t_i - t_p)$. The estimated value for t_u is earlier than the exact time at which the failure happened. This guarantees that all the LSAs that are changed after the exact t_u will be included in the neighbor database summary list if a full adjacency is to be established with router r .

For n ($n > 1$) link failures, the time at which each of these n link failures happened can be calculated as above. The time t_u is the earliest time among all these n times. If we can identify m ($1 \leq m \leq n$) link failures among these n link failures that results in the remote router r unreachable, we only need to calculate the time at which each of those m link failures happened. In this case, the time t_u is the earliest time among all those m times.

For one node failure in the network, every node that has a full adjacency with the failed node will regenerate and flood the LSA with the link to the failure node down. Among these LSAs, suppose that t_r is the earliest time at which one of these LSAs is received and t_p is the time delay for this LSA to reach the local router from its origin, then the time t_u at which the node failed can be estimated as $t_u = (t_r - t_n - t_p)$. The estimated value for t_u is earlier than the exact time at which the failure happened. This guarantees that all the LSAs that are changed after the exact t_u will be included in the neighbor database summary list if a full adjacency is to be established with router r .

For k ($k > 1$) node failures in the network, there are at most k groups of LSAs will be regenerated and flooded in the network. Every LSA in one of these groups contains the information about the link to the same failed node down. For each group of LSAs, the time at which the node failed can be estimated in the same way as above for one node failure. The time t_u at which the remote router r became unreachable because of k node failures is estimated as the earliest time among all the estimated node failure times.

In the case that there are link and node failures in the network, two times are estimated in the ways described above. One time is the time at which the earliest link failure happened. The other is the time at which the earliest node failure happened. The time t_u at which the remote router r became unreachable because of link and node failures is estimated as the earlier between these two times.

6.1.2. A Simpler Method for Finding Failure Time

One simpler way for finding failure time uses all the changed LSAs received. It derives the failure time from every LSA in the way similar to the ones described above and then selects the earliest

time as t_u . For every changed LSA received at time t_r , the failure time derived from this LSA is $(t_r - \max(t_i, t_n) - t_p)$, where t_i , t_n and t_p are defined above and t_p is calculated in the same way as described above.

This method is simpler than the method described in the above subsection. But the failure time it estimated may not be as accurate as the one the previous method calculated. Thus the LSAs changed between these two times will be included in the neighbor database summary list when a full adjacency is to be created to router r and this failure time is used as the unreachable time for router r . However, it is not necessary to include these LSAs in the summary list.

6.1.3. Finding and Recording LSA Change Time

If there exist some remote routers that became unreachable less than MaxAge ago, for an LSA, we use the time t_r at which the LSA is received as the time t_c at which the LSA is changed for recording the time t_c for the LSA. This time may be a little bit later than the exact time at which the LSA is changed. But it guarantees that the LSA will be included in the neighbor database summary list if a full adjacency is to be established with a router r that became unreachable before the exact t_c .

There are a number of ways for recording the LSA change time. One way is to add a new field similar to the field for LSA age into the data structure for storing LSA and store the estimated change time t_c into this new field for each changed LSA.

Another way for recording the LSA change time is to add a link field into the data structure for storing an LSA, a new array and a variable for the index of the array into the data structure for storing an area. The link field is used to link all the LSAs that have the same change time together. The size of the array is the number of time units in MaxAge (3600 seconds). A time unit can be one second, 1/10 second or other smaller time unit. This depends on the implementation. The array and the variable for the index can be considered as a relative clock. The index variable starts from 0, and goes up by one every time unit. When it reaches the size of the array, it starts from 0 again. If the value of the index variable is k at the current time, the time t_j represented by the element of the array at index j is $t_j = (\text{current time} - \text{time unit} * d)$, where $d = (k - j)$ if $k \geq j$ and $d = (\text{array size} - j + k)$ if $k < j$. The element of the array at index j is a pointer that points to the header of the linked list of all the LSAs that are changed at the time t_c , where $t_c = t_j$ or $(t_c + \text{one time unit}) > t_j$ if $t_c < t_j$ (i.e., rounding up t_c is t_j).

6.2. Reusing LSA Age, Finding and Recording Adjust Time

For every remote router r that became unreachable at time t_u , t_u can be estimated in one of the ways described above. For every changed LSA received at time t_r after time t_u , its age is less than 255, which is the maximum time to live value in the IP packet containing the LSA. Thus we use t_u and reuse the age of an LSA to decide whether the LSA is put into the summary list. When a full adjacency to router r is to be created, every LSA that satisfies the condition " $\text{the age of the LSA} < \text{the current time} - t_u + 255$ " (i.e., " $t_u - 255 < \text{the current time} - \text{the age of the LSA}$ ") must be put into the neighbor database summary list. This will guarantee that all the LSAs that were changed after the router r became unreachable are included in the summary list. However, some LSAs that were changed before time t_u may be included. In order to reduce the number of the LSAs changed before time t_u to be included, we need to find the earliest time t_e in which the changed LSA was regenerated in term of LSA age. Thus we can use t_e in place of $t_u - 255$ to decide whether an LSA must be put into the summary list. When a full adjacency to router r is to be created, every LSA that satisfies the condition " $t_e < \text{the current time} - \text{the age of the LSA}$ " must be put into the neighbor database summary list.

For each changed LSA received at time t_{r-i} after time t_u and before time $(t_u + 255)$, the time at which the LSA was regenerated is estimated as $(t_{r-i} - t_{a-i})$, where t_{a-i} is the age of the LSA. The time t_e is the earliest one among all $(t_{r-i} - t_{a-i})$.

For the case that another remote router r' became unreachable later during the calculation of the time t_e for router r , we stop finding t_e and start to find the earliest time $t_{e'}$ for the remote router r' in the same way as described above and mark that the time t_e relies on the time $t_{e'}$. When a full adjacency to router r is to be created, the time t_e is calculated as follows. It is the earliest time among the t_e partially calculated and the $t_{e'}$ that the t_e depends on.

7. Security Considerations

The mechanism described in this document does not raise any new security issues for the OSPF protocols.

8. IANA Considerations

This document specifies a backward compatible intelligent link state database exchange mechanism for OSPFv2 and OSPFv3, which does not require any new number assignment.

9. Acknowledgement

The author would like to thank Richard Li, Yang Yu, Steve Yao, Quintin Zhao and others for their valuable comments on this draft.

10. References

10.1. 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.
- [RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", [RFC 2740](#), December 1999.

10.2. Informative References

- [RFC5243] Ogier, R., "OSPF Database Exchange Summary List Optimization", [RFC 5243](#), May 2008.

Author's Address

Huaimo Chen
Huawei Technology, Inc.
Boston, MA
US

Email: Huaimochen@huawei.com

