OSPF Out-of-band LSDB resynchronization draft-ietf-ospf-oob-resync-01.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of <u>Section 10 of RFC2026</u>.

Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/lid-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Abstract

OSPF is a link-state intra-domain routing protocol used in IP networks. LSDB synchronization in OSPF is achieved via two methods-initial LSDB synchronization when an OSPF router has just been connected to the network and asynchronous flooding that ensures continuous LSDB synchronization in the presence of topology changes after the initial procedure was completed. It may sometime be necessary for OSPF routers to resynchronize their LSDBs. OSPF standard, however, does not allow routers to do so without actually changing the topology view of the network. This memo describes a mechanism to perform such form of out-of-band LSDB synchronization.

1 Motivation

According to the OSPF standard [RFC2328], after two OSPF routers have established an adjacency (the neighbor FSMs have reached Full state), routers announce the adjacency states in their router-LSAs. Asynchronous flooding algorithm ensures routers' LSDBs stay in sync in the presence of topology changes. However, if routers need (for some reason) to resynchronize their LSDBs, they cannot do that without actually putting the neighbor FSMs into the ExStart state. This effectively causes the adjacencies to be removed from the router-LSAs, which may not be acceptable in some cases. In this document, we provide the means for so-called out-of-band (OOB) LSDB resynchronization.

The described mechanism can be used in a number of situations including those where the routers are picking the adjacencies up after a reload. The process of adjacency preemption is outside the scope of this document. Only the details related to LSDB resynchronization are mentioned herein.

2 Proposed solution

The format of the OSPF Database Description packet is changed to include a new R-bit indicating OOB LSDB resynchronization. All DBD packets sent during the OOB resynchronization procedure are sent with the R-bit set.

Also, two new fields are added to the neighbor data structure. The first field indicates neighbor's OOB resynchronization capability. The second indicates that OOB LSDB resynchronization is in process. The latter field allows OSPF implementations to utilize the existing neighbor FSM code.

A bit is occupied in the Extended Options TLV (see [LLS]). Routers set this bit to indicate their capability to support the described technique.

2.1 The LR bit

A new bit, called LR (LR stands for LSDB Resynchronization) is introduced to the LLS Extended Options TLV (see [LLS]). The value of the bit is TBD (0x00000001 is the temporarily used value, see Figure 1). Routers set LR bit to announce OOB LSDB resynchronization capability.

+		-+-		-+-		· + ·		+ -		+ -		+ -		+ •	-	-+-		+ -		- + -		· + ·		+ -		· + ·		+ •		+-	+
I	*	Ι	*	Ι	*	Ι	*	Ι	*	Ι	*	Ι	*			.	*	Ι	*	Ι	*	Ι	*	Ι	*	Ι	*		*	L	LR
+		-+-		-+-		+ •		+ -		+ -		- + -		+ •	-	-+-		+ •		- + -		- + -		+ -		+ •		+ •		+ -	+

Figure 1. The Options field

[Page 2]

OSPF OOB LSDB Resync

Routers supporting the OOB LSDB resynchronization technique set the LR bit in the EO-TLV in the LLS block attached to both Hello and DBD packets. Note that no bit is set in the standard OSPF Options field, neither in OSPF packets, nor in LSAs.

2.2 OSPF Neighbor Data Structure

A field is introduced into OSPF neighbor data structure, as described below. The name of the field is OOBResync and it is a flag indicating that the router is currently performing OOB LSDB resynchronization with the neighbor.

OOBResync flag is set when the router is initiating the OOB LSDB resynchronization (see <u>Section 2.7</u> for more details).

Routers clear OOBResync flag on the following conditions.

- o The neighbor data structure is first created
- The neighbor FSM transitions to any state lower than ExStart
- o The neighbor FSM transitions to ExStart state because a DBD packet with R-bit clear has been received.
- o The neighbor FSM reaches state Full

Note that OOBResync flag may have TRUE value only if the neighbor FSM is in states ExStart, Exchange, or Loading. As indicated above, if the FSM transitions to any other state, the OOBResync flag should be cleared.

It is important to mention that operation of OSPF neighbor FSM is not changed by this document. However, depending on the state of the OOBResync flag, the router sends either normal DBD packets or DBD packets with the R-bit set.

2.3 Hello Packets

Routers capable of performing OOB LSDB resynchronization should always set the LR bit in their Hello packets.

2.4 DBD Packets

Routers supporting the described technique should always set the LR bit in the DBD packets. Since the Options field of the initial DBD packet is stored in corresponding neighbor data structure, the LR bit may be used later to check if a neighbor is capable of performing OOB

[Page 3]

LSB resynchronization.

The format of type-2 (DBD) OSPF packets is changed to include a flag indicating OOB LSDB resynchronization procedure. Figure 2 illustrates the new packet format.

0 2 3 1 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 Version # 2 Packet length 1 Router ID Area ID Checksum | AuType Authentication Authentication Interface MTU | Options |0|0|0|0|R|I|M|MS DD sequence number + --+ An LSA Header + -- + L +--+ + --+ . . .

The R-bit in OSPF type-2 packets is set when the OOBResync flag for the specific neighbor is set to TRUE. If a DBD packets with R-bit clear is received for a neighbor with active OOBResync flag, the OOB LSDB resynchronization process is cancelled and normal LSDB synchronization procedure is initiated.

When a DBD packet is received with R-bit set and the sender is known to be OOB-incapable, the packet should be dropped and a SeqNumber-Mismatch event should be generated for the neighbor.

Processing of DBD packets is modified as follows.

[Page 4]

- 1) If the R-bit is set, do the following
 - o If bits I, M, and MS are set and the state of the neighbor FSM is Full and OOBResync flag is not set, the packet is accepted, the OOBResync flag is set and the FSM is put into ExStart state.
 - o Otherwise, if OOBResync flag is set and the state of the neighbor FSM is ExStart, Exchange, or Loading, the packet is processed just as described in [<u>RFC2328</u>].
 - Otherwise, if neighbor state is Full and the receiving router was the Slave in the LSDB exchange process, it must be ready to identify duplicate DBDs with R-bit set from the master and resend the acknowledging packet.
 - o Otherwise (the OOBResync flag is off, or the state is not Full, or the packet is not a duplicate), a SeqNumberMismatch is generated for the neighbor FSM that causes transition to state ExStart.
- 2) Otherwise (the R-bit is not set) do the following
 - If OOBResync flag for the neighbor is set, OOBResync flag is cleared and a SeqNumberMismatch event is generated for the neighbor FSM.
 - o Otherwise, process the DBD packet as described in [<u>RFC2328</u>].

It is also necessary to limit the time an adjacency can spend in ExStart, Exchange, and Loading states with OOBResync flag set to a finite period of time (e.g., by limitting the number of times DBD and link state request packets can be retransmitted). If the adjacency does not proceed to Full state before the timeout, the neighboring routers experience problems in LSDB resynchronization. The requesting router may decide to stop trying to resynchronize the LSDB over this adjacency (if, for example, it can be resynchronized via another neighbor on the same segment) or to resynchronize using the legacy method by clearing the OOBResync flag and leaving the FSM in ExStart state. The neighboring router may decide to cancel the OOB procedure for the neighbor.

[Page 5]

2.5 Neighbor State Treatment

OSPF implementation supporting the described technique should modify the logic consulting the state of a neighbor FSM as described below.

- FSM state transitioning from and to the Full state with
 OOBResync flag set should not cause origination of a new version
 of router-LSA or network-LSA.
- o Any explicit checks for the Full state of a neighbor FSM for the purposes other than LSDB synchronization and flooding should treat states ExStart, Exchange, and Loading as state Full, provided that OOBResync flag is set for the neighbor. (Flooding and MaxAge-LSA-specific procedures should not check the state of OOBResync flag, but should continue consulting only the FSM state.)

2.6 Initiating OOB LSDB Resynchronization

To initiate out-of-band LSDB resynchronization, the router must first make sure that the corresponding neighbor supports this technology (by checking the LR bit in Options field of the neighbor data structure). If the neighboring router is capable, the OOBResync flag for the neighbor should be set to TRUE and the FSM state should be forced to ExStart.

<u>3</u> Compatibility Issues

Because OOB-capable routers explicitly indicate their capability by setting the corresponding bit in the Options field, no DBD packets with R-bit set are sent to OOB-incapable routers.

The LR bit itself is transparent for OSPF routers and does not affect communication between routers.

<u>4</u> Security Considerations

The described technique does not introduce any new security issues into OSPF protocol.

5 Acknowledgements

The authors would like to thank Acee Lindem, Russ White, Don Slice, and Alvaro Retana for their valuable comments.

[Page 6]

<u>6</u> References

[RFC2328]

J. Moy. OSPF version 2. Technical Report <u>RFC 2328</u>, Internet Engineering Task Force, 1998. <u>ftp://ftp.isi.edu/in-</u> notes/rfc2328.txt.

[LLS] Zinin, Friedman, Roy, Nguyen, Yeung, "OSPF Link-local Signaling", draft-ietf-ospf-lls-00.txt, Work in progress.

7 Authors' addresses

Alex Zinin	Abhay Roy							
Cisco Systems	Cisco Systems							
150 W. Tasman Dr.	170 W. Tasman Dr.							
San Jose,CA 95134	San Jose,CA 95134							
USA	USA							
E-mail: azinin@cisco.com	E-mail: akr@cisco.com							

Liem Nguyen 7025 Kit Creek Rd. Research Triangle Park, NC 27709 USA e-mail: lhnguyen@cisco.com

[Page 7]