Network Working Group
Internet-Draft
Obsoletes: 4813 (if approved)
Intended status: Standards Track
Expires: December 10, 2009

A. Zinin
Alcatel
A. Roy
L. Nguyen
Cisco Systems
B. Friedman
Redback Networks
D. Yeung
Cisco Systems
June 8, 2009

OSPF Link-local Signaling
draft-ietf-ospf-lls-08.txt

Status of this Memo

Copyright Notice

Abstract

OSPF is a link-state intra-domain routing protocol. OSPF routers exchange information on a link using packets that follow a well-defined fixed format. The format is not flexible enough to enable new features which need to exchange arbitrary data. This document describes a backward-compatible technique to perform link-local signaling, i.e., exchange arbitrary data on a link. This document replaces the experimental specification published in RFC4813 to bring it on the Standards Track.

Table of Contents

1. Introduction

This document describes an extension to OSPFv2 [OSPFV2] and OSPFv3 [OSPFV3] allowing additional information to be exchanged between routers on the same link. OSPFv2 and OSPFv3 packet formats are fixed and do not allow for extension. This document proposes appending an optional data block composed of Type/Length/Value (TLV) triplets to existing OSPFv2 and OSPFv3 packets to carry this additional information. Throughout this document, OSPF will be used when the specification is applicable to both OSPFv2 and OSPFv3. Similarly, OSPFv2 or OSPFv3 will be used when the text is protocol specific.

One potential way of solving this task could be introducing a new packet type. However, that would mean introducing extra packets on the network which may not be desirable and may cause backward compatibility issues. This document describes how to exchange data using standard OSPF packet types.

1.1. Requirements notation

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 [KEY].

2. Proposed solution

To perform link-local signaling (LLS), OSPF routers add a special data block to the end of OSPF packets or right after the authentication data block when cryptographic authentication is used. The length of the LLS block is not included into the length of OSPF packet, but is included in the IPv4/IPv6 packet length. Figure 1 illustrates how the LLS data block is attached.

```
   +---------------------+ --             --  +---------------------+
   | IP Header           | ^              ^   | IPv6 Header         |
   | Length = HL+X+Y+Z   | | Header Length |   | Length = HL+X+Y     |
   |                     | v              v   |                     |
   +---------------------+ --             --  +---------------------+
   | OSPF Header         | ^              ^   | OSPFv3 Header       |
   | Length = X          | |              |   | Length = X          |
   |.....................| | X            | X |.....................|
   |                     | |              |   |                     |
   | OSPFv2 Data         | |              |   | OSPFv3 Data         |
   |                     | v              v   |                     |
   +---------------------+ --             --  +---------------------+
   |                     | ^              ^   |                     |
   | Authentication Data | | Y            | Y |  LLS Data           |
   |                     | v              v   |                     |
   +---------------------+ --             --  +---------------------+
   |                     | ^
   |  LLS Data           | | Z
   |                     | v
   +---------------------+ --
```

Figure 1: LLS Data Block in OSPFv2 and OSPFv3

The LLS block MAY be attached to OSPF Hello and DD packets. LLS block MUST NOT be attached to any other OSPF packet types on generation and MUST be ignored on reception.

The data included in the LLS block attached to a Hello packet MAY be used for dynamic signaling since Hello packets may be sent at any time. However, delivery of LLS data in Hello packets is not guaranteed. The data sent with DD packets is guaranteed to be delivered as part of the adjacency forming process.

This document does not specify how the data transmitted by the LLS mechanism should be interpreted by OSPF routers. As routers that do not understand LLS may receive these packets, changes made due to LLS block TLV's do not affect the basic routing when interacting with non-LLS routers.

2.1. L-bit in Options Field

A new L bit (L stands for LLS) is introduced into the OSPF Options field (see Figure 2a/2b). Routers set the L bit in Hello and DD packets to indicate that the packet contains an LLS data block. In other words, the LLS data block is only examined if the L bit is set.

```
              +---+---+---+---+---+---+---+---+
              | * | O | DC| L |N/P| MC| E | * |
              +---+---+---+---+---+---+---+-+-+
```

Figure 2a: OSPFv2 Options field

```
    0                   1                       2
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4  5 6 7  8  9  0  1  2  3
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+--+-+-+--+--+--+--+--+--+
   | | | | | | | | | | | | | | |L|AF|*|*|DC| R| N|MC| E|V6|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+--+-+-+--+--+--+--+--+--+
```

Figure 2b: OSPFv3 Options field

The L bit MUST NOT be set except in Hello and DD packets that contain LLS block.

2.2. LLS Data Block

The data block used for link-local signaling is formatted as described below (see Figure 3 for illustration).

```
    0                   1                   2                   3
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |            Checksum           |       LLS Data Length         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   |                           LLS TLVs                            |
   .                                                               .
   .                                                               .
   .                                                               .
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Format of LLS Data Block

The Checksum field contains the standard IP checksum for the entire contents of the LLS block. Before computing the checksum, the

checksum field is set to 0. If the checksum is incorrect, the OSPF packet MUST be processed, but the LLS block MUST be discarded.

The 16-bit LLS Data Length field contains the length (in 32-bit words) of the LLS block including the header and payload.

Note that if the OSPF packet is cryptographically authenticated, the LLS data block MUST also be cryptographically authenticated. In this case, the regular LLS checksum is not calculated, but is instead set to 0.

The rest of the block contains a set of Type/Length/Value (TLV) triplets as described in Section 2.3. All TLVs MUST be 32-bit aligned (with padding if necessary).

2.3. LLS TLVs

The contents of LLS data block is constructed using TLVs. See Figure 4 for the TLV format.

The type field contains the TLV ID which is unique for each type of TLV. The Length field contains the length of the Value field (in bytes). The value field is variable and contains arbitrary data.

```
 0                   1                   2                   3
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |            Length             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                             Value                             .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Format of LLS TLVs

Note that TLVs are always padded to 32-bit boundary, but padding bytes are not included in the TLV Length field (though they are included in the LLS Data Length field in the LLS block header). Unrecognized TLV types are ignored.

2.4. Extended Options and Flags TLV

This subsection describes a TLV called the Extended Options and Flags (EOF) TLV. The format of EOF-TLV is shown in Figure 5.

Bits in the Value field do not have any semantics from the point of

view of the LLS mechanism. Bits MAY be allocated to announce OSPF link-local capabilities. Bits MAY also be allocated to perform boolean link-local signaling.

The length of the Value field in the EOF-TLV is 4 bytes.

The value of the type field in the EOF-TLV is 1.

The EOF-TLV MUST only appear once in the LLS data block.

```
    0                   1                   2                   3
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |             1                 |            4                  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                  Extended Options and Flags                   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Format of EOF-TLV

Currently, [OOB] and [RESTART] use bits in the Extended Options field of the EOF-TLV.

The Extended Options and Flags bits are defined in Section 3.

2.5. Cryptographic Authentication TLV (OSPFv2 ONLY)

This document defines a special TLV that is used for cryptographic authentication (CA-TLV) of the LLS data block. This TLV MUST only be included in the LLS block when cryptographic authentication is enabled on the corresponding interface. The message digest of the LLS block MUST be calculated using the same key and authentication algorithm as used for the OSPFv2 packet. The cryptographic sequence number is included in the TLV and MUST be the same as the one in the OSPFv2 authentication data for the LLS block to be considered authentic.

The TLV is constructed as shown in Figure 6.

```
    0                   1                   2                   3
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |              2                |         AuthLen               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                         Sequence number                       |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   .                                                               .
   .                           AuthData                            .
   .                                                               .
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: Format of Cryptographic Authentication TLV

The value of the Type field for the CA-TLV is 2.

The Length field in the header contains the length of the data portion of the TLV including 4 bytes for Sequence Number and the length of the message digest block for the whole LLS block in bytes.

The Sequence Number field contains the cryptographic sequence number that is used to prevent simple replay attacks. For the LLS block to be considered authentic, the Sequence Number in the CA-TLV MUST match the Sequence Number in the OSPFv2 packet header Authentication field (which MUST be present). In the event of Sequence Number mismatch or Authentication failure, the whole LLS block MUST be ignored.

The CA-TLV MUST NOT appear more than once in the LLS block. Also, when present, this TLV MUST be the last TLV in the LLS block. If it appears more than once, only the first occurrence is processed and any others MUST be ignored.

The AuthData contains the message digest calculated for the LLS data block up to CA-TLV (i.e. exludes CA-TLV data).

The CA-TLV is not applicable to OSPFv3 and it MUST NOT be added to any OSPFv3 packet. If found on reception, this TLV MUST be ignored.

2.6. Private TLVs

LLS type values in the range of 32768-65536 are reserved for private use. The first four octets of the Value field MUST be the private enterprise code [ENTNUM]. This allows multiple vendor private extensions to coexist in a network.

3. IANA Considerations

This document uses the registry that was originally created in [RFC4813]. IANA is requested to update the following registry to point to this document instead:

- "Open Shortest Path First (OSPF) Link Local Signalling (LLS) - Type/Length/Value Identifiers (TLV)"

IANA is requested to allocate L-bit in "OSPFv2 Options Registry" and "OSPFv3 Options Registry" as per Section 2.1.

LLS TLV types are maintained by the IANA. Extensions to OSPF which require a new LLS TLV type MUST be reviewed by an Designated Expert from the routing area.

The criteria for allocating LLS TLVs are:

- LLS should not be used for information that would be better suited to be advertised in a link-local LSA.

- LLS should be confined to signaling between direct neighbors.

- Discretion should be used in the volume of information signaled using LLS due to the obvious MTU and performance implications.

Following the policies outlined in [IANA], LLS type values in the range of 0-32767 are allocated through an IETF Review and LLS type values in the range of 32768-65536 are reserved for private use.

This document assigns the following LLS TLV types in OSPFv2/OSPFv3.

TLV Type	Name	Reference
0	Reserved	
1	Extended Options and Flags	[RFCNNNN]*
2	Cryptographic Authentication+	[RFCNNNN]*
3-32767	Reserved for assignment by the IANA	
32768-65535	Private Use	

*[RFCNNNN] refers to the RFC number-to-be for this document.
+ Cryptographic Authentication TLV is only defined for OSPFv2

IANA is requested to rename the sub-registry "LLS Type 1 Extended Options" to "LLS Type 1 Extended Options and Flags".

This document also assigns the following bits in the EOF-TLV outlined in Section 2.5:

Bit	Name	Reference
0x00000001	LSDB Resynchronization (LR)	[OOB]
0x00000002	Restart Signal (RS-bit)	[RESTART]

Future allocation of Extended Options and Flags bits MUST be reviewed by an Designated Expert from the routing area.

4. Compatibility Issues

The modifications to OSPF packet formats are compatible with standard OSPF since OSPF routers not supporting LLS will ignore the LLS data block after the OSPF packet or cryptographic message digest. As of this writing, there are implementations deployed with [RFC4813] compliant software. Routers not implementing [RFC4813] ignore the LLS data at the end of OSPF packet.

Careful consideration should be given to carrying additional LLS data, as it may affect the OSPF adjacency bring-up time due to additional propagation delay and/or processing time.

5. Security Considerations

Security Considerations inherited from OSPFv2 are described in [OSPFV2]. This technique provides the same level of security as the basic OSPFv2 protocol by allowing LLS data to be authenticated using the same cryptographic authentication that OSPFv2 uses (see Section 2.5 for more details).

Security considerations inherited from OSPFv3 are described in [OSPFV3] and [OSPFV3AUTH]. OSPFv3 utilizes IPSec for authentication and encryption. With IPsec, the AH (Authentication Header), ESP (Encapsulating Security Payload), or both are applied to the entire OSPFv3 payload including the LLS block.

6. References

6.1. Normative References

[IANA]   Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, May 2008.

[KEY]   Bradner, S., "Key words for use in RFC's to Indicate Requirement Levels", RFC 2119, March 1997.

[OSPFV2]   Moy, J., "OSPF Version 2", RFC 2328, April 1998.

[OSPFV3]   Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[OSPFV3AUTH] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, June 2006.

6.2. Informative References

[ENTNUM]   IANA, "http://www.iana.org/assignments/enterprise-numbers".

[OOB]   Zinin, A., Roy, A., and L. Nguyen, "OSPF Out-of-band LSDB resynchronization", RFC 4811, March 2007.

[RESTART]   Zinin, A., Roy, A., and L. Nguyen, "OSPF Restart Signaling", RFC 4812, March 2007.

[RFC4813]   Friedman, B., Nguyen, L., Roy, A., Yeung, D., and A. Zinin, "OSPF Link-Local Signaling", RFC 4813, March 2007.

Appendix A. Acknowledgements

The authors would like to acknowledge Russ White, Acee Lindem and Manral Vishwas for their review of this document.

Appendix B. Changes from RFC 4813

This section describes the substantive change from [RFC4813].

- Added OSPFv3 support

- Private TLVs MUST use private enterprise code

- Clarified requirement levels at several places

- Changed from Experimental to Standards Track

Authors' Addresses

Alex Zinin
Alcatel
Sunnyvale
USA

Email: zinin@psg.com

Abhay Roy
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Email: akr@cisco.com

Liem Nguyen
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Email: lhnguyen@cisco.com

Barry Friedman
Redback Networks
100 Headquarters Drive
San Jose, CA 95134
USA

Email: friedman@redback.com

Derek Yeung
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Email: myeung@cisco.com