Network Working Group
Request for Comments: 4813
Category: Experimental

B. Friedman
L. Nguyen
A. Roy
D. Yeung
Cisco Systems
A. Zinin
Alcatel

February 2007

OSPF Link-Local Signaling

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Abstract

OSPF is a link-state intra-domain routing protocol used in IP networks. OSPF routers exchange information on a link using packets that follow a well-defined format. The format of OSPF packets is not flexible enough to enable applications to exchange arbitrary data, which may be necessary in certain situations. This memo describes a vendor-specific, backward-compatible technique to perform link-local signaling, i.e., exchange arbitrary data on a link.

Table of Contents

1. Introduction

Formats of OSPF [RFC2328] packets are not very flexible to provide an acceptable mechanism for opaque data transfer. However, this appears to be very useful to allow OSPF routers to do so. An example where such a technique could be used is exchanging some capabilities on a link (standard OSPF utilizes the Options field in Hello and Exchange packets, but there are not so many bits left in it).

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, so this document describes how to exchange data using existing, standard OSPF packet types.

2. Proposed Solution

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

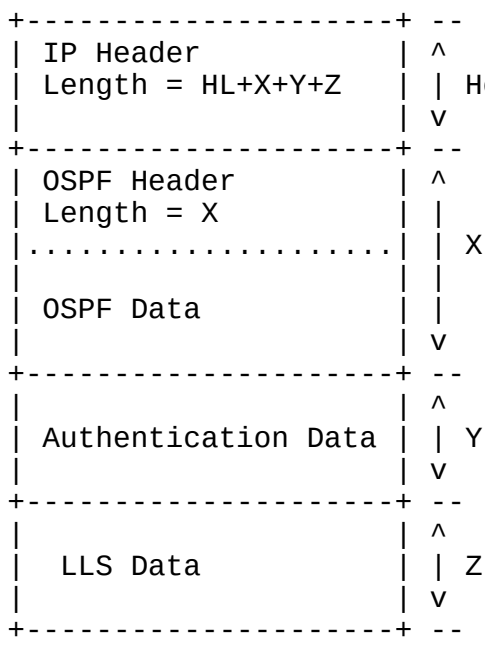

Figure 1: Attaching LLS Data Block

The LLS data block may be attached to OSPF packets of two types -- type 1 (OSPF Hello), and type 2 (OSPF DBD). 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 moment in time. However, delivery of LLS data in Hello packets is not guaranteed. The data sent with Database Description (DBD) packets is guaranteed to be delivered as part of the adjacency forming process.

This memo does not specify how the data transmitted by the LLS mechanism should be interpreted by OSPF routers. The interface between the OSPF LLS component and its clients is implementation-specific.

2.1. Options Field

A new bit, called L (L stands for LLS), is introduced to the OSPF Options field (see Figure 2). The value of the bit is 0x10. Routers set the L-bit in Hello and DBD packets to indicate that the packet contains the LLS data block.

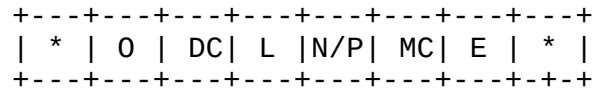

Figure 2: The Options Field

L-bit

This bit is set only in Hello and DBD packets. It is not set in OSPF Link State Advertisements (LSAs) and may be used in them for different purposes.

2.2. LLS Data Block

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

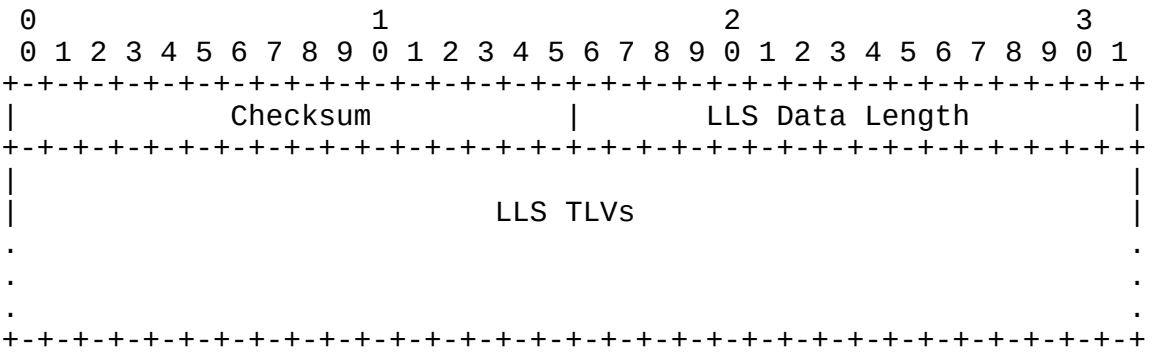

Figure 3: Format of the LLS Data Block

Checksum

The Checksum field contains the standard IP checksum of the entire contents of the LLS block.

LLS Length

The 16-bit LLS Data Length field contains the length (in 32-bit words) of the LLS block including the header and payload. Implementations should not use the Length field in the IP packet header to determine the length of the LLS data block.

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 and the LLS block will contain a cryptographic authentication TLV (see Section 2.4.2).

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

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

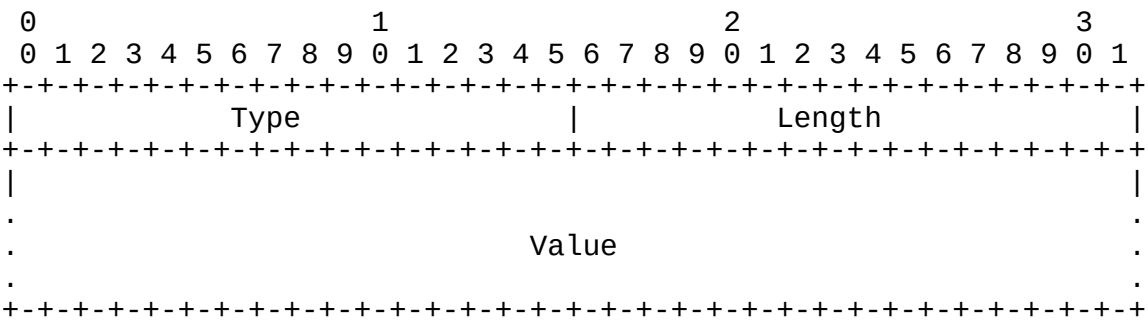

Figure 4: Format of LLS TLVs

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

2.4. Predefined TLV

2.4.1. Extended Options TLV

This subsection describes a TLV called Extended Options (EO) TLV. The format of EO-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. This field may be used to announce some OSPF capabilities that are link-specific. Also, other OSPF extensions may allocate bits in the bit vector to perform boolean link-local signaling.

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

The value of the Type field in EO-TLV is 1.

EO-TLV should only appear once in the LLS data block.

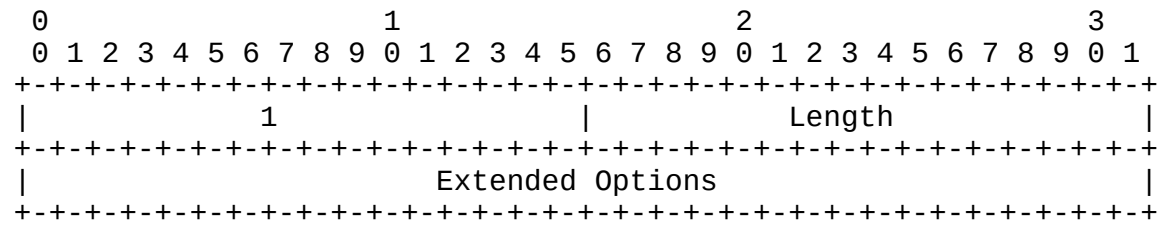

Figure 5: Format of EO-TLV

Currently, [RFC4811] and [RFC4812] use bits in the Extended Options field of the EO-TLV. The Extended Options bits are also defined in Section 5.

2.4.2. Cryptographic Authentication TLV

This document defines a special TLV that is used for cryptographic authentication (CA-TLV) of the LLS data block. This TLV should be included in the LLS block when the cryptographic (MD5) authentication is enabled on the corresponding interface. The message digest of the LLS block should be calculated using the same key as that used for the main OSPF packet. The cryptographic sequence number is included in the TLV and must be the same as the one in the main OSPF packet for the LLS block to be considered authentic.

The TLV is constructed as shown Figure 6.

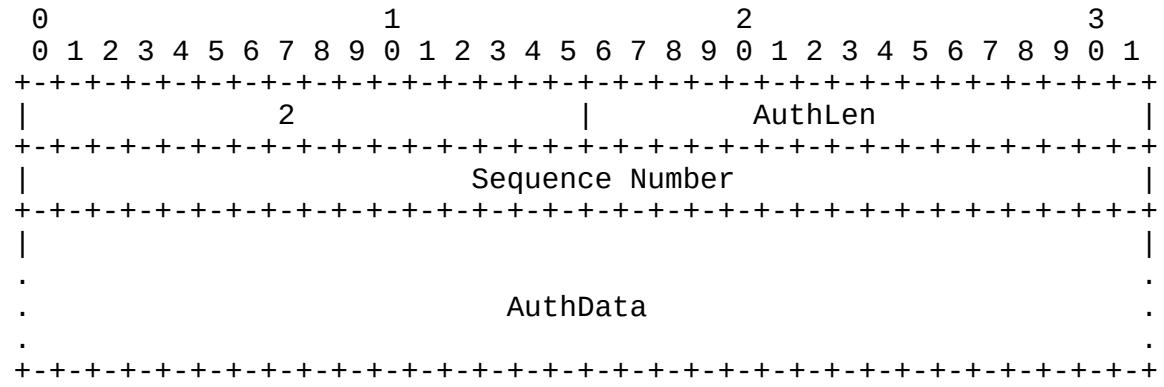

Figure 6: Format of Cryptographic Authentication TLV

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

The Length field in the header contains the length of the data portion of the TLV that includes 4 bytes for the sequence number and the length of the message digest (MD5) block for the whole LLS block

in bytes (this will always be 16 bytes for MD5). So the AuthLen field will have value of 20.

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 OSPF packet.

The AuthData field contains the message digest calculated for the LLS data block.

The CA-TLV may appear in the LLS block only once. Also, when present, this TLV should be the last in the LLS block.

3. Backward Compatibility

The modifications to OSPF packet formats are compatible with standard OSPF because LLS-incapable routers will not consider the extra data after the packet; i.e., the LLS data block will be ignored by routers that do not support the LLS extension.

4. Security Considerations

The function described in this document does not create any new security issues for the OSPF protocol. The described technique provides the same level of security as the OSPF protocol by allowing LLS data to be authenticated (see Section 2.4.2 for more details).

5. IANA Considerations

LLS TLV types are maintained by the IANA. Extensions to OSPF that require a new LLS TLV type must be reviewed by a designated expert from the routing area.

Following the policies outlined in [RFC2434], LLS type values in the range of 0-32767 are allocated through an IETF consensus action, and LLS type values in the range of 32768-65536 are reserved for private and experimental use.

This document assigns LLS types 1 and 2, as follows:

LLS Type	Name	Reference
0	Reserved	
1	Extended Options	[RFC4813]
2	Cryptographic Authentication	[RFC4813]
3-32767	Reserved for assignment by the IANA	
32768-65535	Private Use	

This document also assigns the following bits for the Extended Options bits field in the EO-TLV outlined in Section 2.4.1:

Extended Options Bit	Name	Reference
0x00000001	LSDB Resynchronization (LR)	[RFC4811]
0x00000002	Restart Signal (RS-bit)	[RFC4812]

Other Extended Options bits will be allocated through an IETF consensus action.

6. References

6.1. Normative References

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

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

6.2. Informative References

[RFC4811] Nguyen, L., Roy, A., and A. Zinin, "OSPF Out-of-Band Link State Database (LSDB) Resynchronization", RFC 4811, February 2007.

[RFC4812] Nguyen, L., Roy, A., and A. Zinin, "OSPF Restart Signaling", RFC 4812, February 2007.

Appendix A. Acknowledgments

The authors would like to acknowledge Russ White for his review of this document.

Authors' Addresses

Barry Friedman
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: friedman@cisco.com

Liem Nguyen
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: lhnguyen@cisco.com

Abhay Roy
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: akr@cisco.com

Derek Yeung
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: myeung@cisco.com

Alex Zinin
Alcatel
Sunnyvale, CA
USA
EMail: zinin@psg.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.