

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 31, 2008

A. Zinin
Alcatel
A. Roy
L. Nguyen
Cisco Systems
B. Friedman
Redback Networks
D. Yeung
August 2007

OSPF Link-local Signaling
draft-ietf-ospf-lls-03.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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 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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on January 31, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

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 memo describes a backward-compatible technique to perform link-local signaling, i.e., exchange arbitrary data on a link.

Table of Contents

1.	Introduction	3
1.1.	Conventions Used In This Document	3
2.	Proposed solution	4
2.1.	Options Field	5
2.2.	LLS Data Block	5
2.3.	LLS TLVs	6
2.4.	Extended Options TLV	6
2.5.	Cryptographic Authentication TLV (OSPFv2 ONLY)	7
3.	IANA Considerations	9
4.	Compatibility Issues	10
5.	Security Considerations	11
6.	References	12
6.1.	Normative References	12
6.2.	Informative References	12
Appendix A.	Acknowledgements	13
	Authors' Addresses	14
	Intellectual Property and Copyright Statements	15

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. 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](#) [[KEY](#)].

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. 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.




Figure 1: LLS Data Block in OSPFv2 and OSPFv3

The LLS data block MAY be attached to OSPF hello and DD packets. The data included in 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 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 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 OSPF Options field (see Figure 2a/2b). Routers set L bit in Hello and DBD packets to indicate that the packet contains LLS data block. In other words, LLS data block is only examined if L bit is set.

```

+---+---+---+---+---+---+---+---+
| * | 0 | 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 is only set in Hello and DBD packets.

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 3: Format of LLS Data Block

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

The 16-bit LLS Data Length field contains the length (in 32-bit words) of the LLS block including the header and payload. Implementations MUST 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 2.5](#)).

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 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 TLV Length field (though it is included in the LLS Data Length field of the LLS block header).

2.4. 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 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 E0-TLV is 4 bytes.

The value of the type field in E0-TLV is 1.

E0-TLV MUST only appear once in the LLS data block.




Figure 5: Format of E0 TLV

Currently, [[OOB](#)] and [[RESTART](#)] use bits in the Extended Options field of the E0-TLV.

The Extended Options 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 be included in the LLS block when the cryptographic (MD5) 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 that used for the main OSPFv2 packet. The cryptographic sequence number is included in the TLV and MUST be the same as the one in the main OSPFv2 packet for the LLS block to be considered authentic.

The TLV is constructed as shown Figure 6.

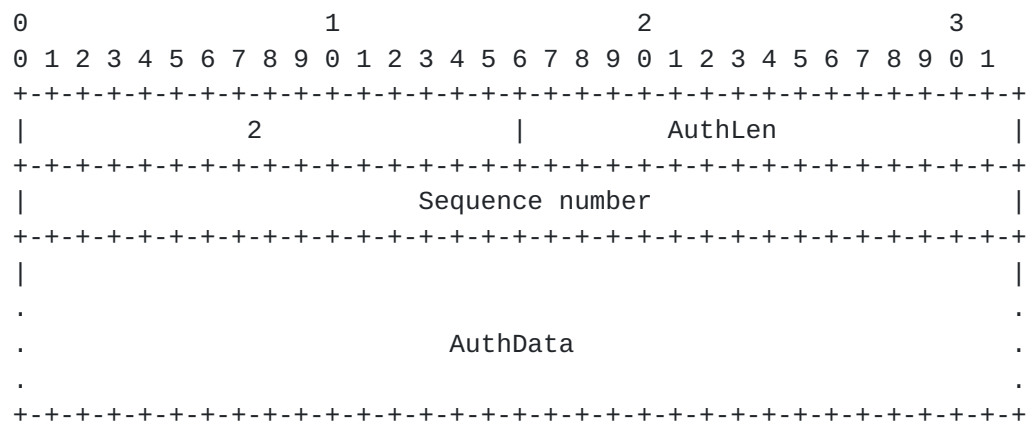


Figure 6: Format of Cryptographic Authentication TLV

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 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 OSPFv2 packet.

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

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

3. IANA Considerations

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.

Following the policies outlined in [[IANA](#)], 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 the following LLS TLV types in OSPFv2/OSPFv3.

TLV Type	Name	Reference
0	Reserved	
1	Extended Options	[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

This document also assigns the following bits for the Extended Options bits field in the EO-TLV outlined in [Section 2.5](#):

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

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

4. Compatibility Issues

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 which do not support the LLS extension.

5. Security Considerations

The described technique provides the same level of security as OSPF protocol by allowing LLS data to be authenticated (see [Section 2.5](#) for more details).

OSPFv3 has IPSec authentication built in. There are AH/ESP techniques which operate on the whole OSPFv3 payload. So we do not need a separate cryptographic TLV for OSPFv3.

6. References

6.1. Normative References

- [IANA] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 2334](#), October 1998.
- [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., and J. Moy, "OSPF for IPv6", [RFC 2740](#), December 1999.

6.2. Informative References

- [OOB] Zinin, A., Roy, A., and L. Nguyen, "OSPF Out-of-band LSDB resynchronization", [draft-nguyen-ospf-oob-resync-06.txt](#) (work in progress), October 2006.
- [RESTART] Zinin, A., Roy, A., and L. Nguyen, "OSPF Restart Signaling", [draft-nguyen-ospf-restart-06.txt](#) (work in progress), October 2006.

[Appendix A](#). Acknowledgements

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

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
300 Holger Way
San Jose, CA 95134
USA

Email: friedman@redback.com

Derek Young
Email: derekmyeung@yahoo.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).