

Network Working Group  
Request for Comments: 4552  
Category: Standards Track

M. Gupta  
Tropos Networks  
N. Melam  
Juniper Networks  
June 2006

## Authentication/Confidentiality for OSPFv3

### Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2006).

### Abstract

This document describes means and mechanisms to provide authentication/confidentiality to OSPFv3 using an IPv6 Authentication Header/Encapsulating Security Payload (AH/ESP) extension header.

## Table of Contents

1. Introduction .....	2
1.1. Conventions Used in This Document .....	2
2. Transport Mode vs. Tunnel Mode .....	3
3. Authentication .....	3
4. Confidentiality .....	3
5. Distinguishing OSPFv3 from OSPFv2 .....	4
6. IPsec Requirements .....	4
7. Key Management .....	5
8. SA Granularity and Selectors .....	7
9. Virtual Links .....	8
10. Rekeying .....	9
10.1. Rekeying Procedure .....	9
10.2. KeyRolloverInterval .....	9
10.3. Rekeying Interval .....	10
11. IPsec Protection Barrier and SPD .....	10
12. Entropy of Manual Keys .....	12
13. Replay Protection .....	12
14. Security Considerations .....	12
15. References .....	13
15.1. Normative References .....	13
15.2. Informative References .....	13

## 1. Introduction

OSPF (Open Shortest Path First) Version 2 [N1] defines the fields AuType and Authentication in its protocol header to provide security. In OSPF for IPv6 (OSPFv3) [N2], both of the authentication fields were removed from OSPF headers. OSPFv3 relies on the IPv6 Authentication Header (AH) and IPv6 Encapsulating Security Payload (ESP) to provide integrity, authentication, and/or confidentiality.

This document describes how IPv6 AH/ESP extension headers can be used to provide authentication/confidentiality to OSPFv3.

It is assumed that the reader is familiar with OSPFv3 [N2], AH [N5], ESP [N4], the concept of security associations, tunnel and transport mode of IPsec, and the key management options available for AH and ESP (manual keying [N3] and Internet Key Exchange (IKE)[I1]).

## 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 [N7].

## 2. Transport Mode vs. Tunnel Mode

The transport mode Security Association (SA) is generally used between two hosts or routers/gateways when they are acting as hosts. The SA must be a tunnel mode SA if either end of the security association is a router/gateway. Two hosts MAY establish a tunnel mode SA between themselves. OSPFv3 packets are exchanged between routers. However, since the packets are locally delivered, the routers assume the role of hosts in the context of tunnel mode SA. All implementations conforming to this specification MUST support transport mode SA to provide required IPsec security to OSPFv3 packets. They MAY also support tunnel mode SA to provide required IPsec security to OSPFv3 packets.

## 3. Authentication

Implementations conforming to this specification MUST support authentication for OSPFv3.

In order to provide authentication to OSPFv3, implementations MUST support ESP and MAY support AH.

If ESP in transport mode is used, it will only provide authentication to OSPFv3 protocol packets excluding the IPv6 header, extension headers, and options.

If AH in transport mode is used, it will provide authentication to OSPFv3 protocol packets, selected portions of IPv6 header, selected portions of extension headers, and selected options.

When OSPFv3 authentication is enabled,

- o OSPFv3 packets that are not protected with AH or ESP MUST be silently discarded.
- o OSPFv3 packets that fail the authentication checks MUST be silently discarded.

## 4. Confidentiality

Implementations conforming to this specification SHOULD support confidentiality for OSPFv3.

If confidentiality is provided, ESP MUST be used.

When OSPFv3 confidentiality is enabled,

- o OSPFv3 packets that are not protected with ESP MUST be silently discarded.
- o OSPFv3 packets that fail the confidentiality checks MUST be silently discarded.

#### 5. Distinguishing OSPFv3 from OSPFv2

The IP/IPv6 Protocol Type for OSPFv2 and OSPFv3 is the same (89), and OSPF distinguishes them based on the OSPF header version number. However, current IPsec standards do not allow using arbitrary protocol-specific header fields as the selectors. Therefore, the OSPF version field in the OSPF header cannot be used to distinguish OSPFv3 packets from OSPFv2 packets. As OSPFv2 is only for IPv4 and OSPFv3 is only for IPv6, the version field in the IP header can be used to distinguish OSPFv3 packets from OSPFv2 packets.

#### 6. IPsec Requirements

In order to implement this specification, the following IPsec capabilities are required.

##### Transport mode

IPsec in transport mode MUST be supported. [N3]

##### Multiple Security Policy Databases (SPDs)

The implementation MUST support multiple SPDs with an SPD selection function that provides an ability to choose a specific SPD based on interface. [N3]

##### Selectors

The implementation MUST be able to use source address, destination address, protocol, and direction as selectors in the SPD.

##### Interface ID tagging

The implementation MUST be able to tag the inbound packets with the ID of the interface (physical or virtual) via which it arrived. [N3]

##### Manual key support

Manually configured keys MUST be able to secure the specified traffic. [N3]

#### Encryption and authentication algorithms

The implementation MUST NOT allow the user to choose stream ciphers as the encryption algorithm for securing OSPFv3 packets since the stream ciphers are not suitable for manual keys.

Except when in conflict with the above statement, the key words "MUST", "MUST NOT", "REQUIRED", "SHOULD", and "SHOULD NOT" that appear in the [N6] document for algorithms to be supported are to be interpreted as described in [N7] for OSPFv3 support as well.

#### Dynamic IPsec rule configuration

The routing module SHOULD be able to configure, modify, and delete IPsec rules on the fly. This is needed mainly for securing virtual links.

#### Encapsulation of ESP packet

IP encapsulation of ESP packets MUST be supported. For simplicity, UDP encapsulation of ESP packets SHOULD NOT be used.

#### Different SAs for different Differentiated Services Code Points (DSCPs)

As per [N3], the IPsec implementation MUST support the establishment and maintenance of multiple SAs with the same selectors between a given sender and receiver. This allows the implementation to associate different classes of traffic with the same selector values in support of Quality of Service (QoS).

### 7. Key Management

OSPFv3 exchanges both multicast and unicast packets. While running OSPFv3 over a broadcast interface, the authentication/confidentiality required is "one to many". Since IKE is based on the Diffie-Hellman key agreement protocol and works only for two communicating parties, it is not possible to use IKE for providing the required "one to many" authentication/confidentiality. This specification mandates the usage of Manual Keying with current IPsec implementations. Future specifications can explore the usage of protocols like Kerberized Internet Negotiation of Keys/Group Secure Association Key Management Protocol (KINK/GSAKMP) when they are widely available. In manual keying, SAs are statically installed on the routers and these static SAs are used to authenticate/encrypt packets.

The following discussion explains that it is not scalable and is practically infeasible to use different security associations for inbound and outbound traffic to provide the required "one to many" security. Therefore, the implementations MUST use manually

configured keys with the same SA parameters (Security Parameter Index (SPI), keys, etc.) for both inbound and outbound SAs (as shown in Figure 3).




Figure 1

If we consider communication between A and B in Figure 1, everything seems to be fine. A uses security association SAa for outbound packets and B uses the same for inbound packets and vice versa. Now if we include C in the group and C sends a packet using SAa, then only A will be able to understand it. Similarly, if C sends a packet using SAb, then only B will be able to understand it. Since the packets are multicast and they are going to be processed by both A and B, there is no SA for C to use so that both A and B can understand them.




Figure 2

The problem can be solved by configuring SAs for all the nodes on every other node as shown in Figure 2. So A, B, and C will use SAa, SAb, and SAc, respectively, for outbound traffic. Each node will lookup the SA to be used based on the source (A will use SAb and SAc for packets received from B and C, respectively). This solution is not scalable and practically infeasible because a large number of SAs will need to be configured on each node. Also, the addition of a node in the broadcast network will require the addition of another SA on every other node.




Figure 3

The problem can be solved by using the same SA parameters (SPI, keys, etc.) for both inbound (SAi) and outbound (SAo) SAs as shown in Figure 3.

### 8. SA Granularity and Selectors

The user SHOULD be given the choice of sharing the same SA among multiple interfaces or using a unique SA per interface.

OSPFv3 supports running multiple instances over one interface using the "Instance Id" field contained in the OSPFv3 header. As IPsec does not support arbitrary fields in the protocol header to be used as the selectors, it is not possible to use different SAs for different OSPFv3 instances running over the same interface. Therefore, all OSPFv3 instances running over the same interface will have to use the same SA. In OSPFv3 RFC terminology, SAs are per-link and not per-interface.

## 9. Virtual Links

A different SA than the SA of the underlying interface MUST be provided for virtual links. Packets sent on virtual links use unicast non-link local IPv6 addresses as the IPv6 source address, while packets sent on other interfaces use multicast and unicast link local addresses. This difference in the IPv6 source address differentiates the packets sent on virtual links from other OSPFv3 interface types.

As the virtual link end point IPv6 addresses are not known, it is not possible to install SPD/Security Association Database (SAD) entries at the time of configuration. The virtual link end point IPv6 addresses are learned during the routing table computation process. The packet exchange over the virtual links starts only after the discovery of the end point IPv6 addresses. In order to protect these exchanges, the routing module must install the corresponding SPD/SAD entries before starting these exchanges. Note that manual SA parameters are preconfigured but not installed in the SAD until the end point addresses are learned.

According to the OSPFv3 RFC [N2], the virtual neighbor's IP address is set to the first prefix with the "LA-bit" set from the list of prefixes in intra-area-prefix-Link State Advertisements (LSAs) originated by the virtual neighbor. But when it comes to choosing the source address for the packets that are sent over the virtual link, the RFC [N2] simply suggests using one of the router's own global IPv6 addresses. In order to install the required security rules for virtual links, the source address also needs to be predictable. Hence, routers that implement this specification MUST change the way the source and destination addresses are chosen for packets exchanged over virtual links when IPsec is enabled.

The first IPv6 address with the "LA-bit" set in the list of prefixes advertised in intra-area-prefix-LSAs in the transit area MUST be used as the source address for packets exchanged over the virtual link. When multiple intra-area-prefix-LSAs are originated, they are considered concatenated and are ordered by ascending Link State ID.

The first IPv6 address with the "LA-bit" set in the list of prefixes received in intra-area-prefix-LSAs from the virtual neighbor in the transit area MUST be used as the destination address for packets exchanged over the virtual link. When multiple intra-area-prefix-LSAs are received, they are considered concatenated and are ordered by ascending Link State ID.

This makes both the source and destination addresses of packets exchanged over the virtual link predictable when IPsec is enabled.



## 10. Rekeying

To maintain the security of a link, the authentication and encryption key values SHOULD be changed periodically.

### 10.1. Rekeying Procedure

The following three-step procedure SHOULD be provided to rekey the routers on a link without dropping OSPFv3 protocol packets or disrupting the adjacency.

- (1) For every router on the link, create an additional inbound SA for the interface being rekeyed using a new SPI and the new key.
- (2) For every router on the link, replace the original outbound SA with one using the new SPI and key values. The SA replacement operation should be atomic with respect to sending OSPFv3 packets on the link so that no OSPFv3 packets are sent without authentication/encryption.
- (3) For every router on the link, remove the original inbound SA.

Note that all routers on the link must complete step 1 before any begin step 2. Likewise, all the routers on the link must complete step 2 before any begin step 3.

One way to control the progression from one step to the next is for each router to have a configurable time constant `KeyRolloverInterval`. After the router begins step 1 on a given link, it waits for this interval and then moves to step 2. Likewise, after moving to step 2, it waits for this interval and then moves to step 3.

In order to achieve smooth key transition, all routers on a link should use the same value for `KeyRolloverInterval` and should initiate the key rollover process within this time period.

At the end of this procedure, all the routers on the link will have a single inbound and outbound SA for OSPFv3 with the new SPI and key values.

### 10.2. `KeyRolloverInterval`

The configured value of `KeyRolloverInterval` should be long enough to allow the administrator to change keys on all the OSPFv3 routers. As this value can vary significantly depending upon the implementation and the deployment, it is left to the administrator to choose an appropriate value.

### 10.3. Rekeying Interval

This section analyzes the security provided by manual keying and recommends that the encryption and authentication keys SHOULD be changed at least every 90 days.

The weakest security provided by the security mechanisms discussed in this specification is when NULL encryption (for ESP) or no encryption (for AH) is used with the HMAC-MD5 authentication. Any other algorithm combinations will at least be as hard to break as the ones mentioned above. This is shown by the following reasonable assumptions:

- o NULL Encryption and HMAC-SHA-1 Authentication will be more secure as HMAC-SHA-1 is considered to be more secure than HMAC-MD5.
- o NON-NULL Encryption and NULL Authentication combination is not applicable as this specification mandates authentication when OSPFv3 security is enabled.
- o Data Encryption Security (DES) Encryption and HMAC-MD5 Authentication will be more secure because of the additional security provided by DES.
- o Other encryption algorithms like 3DES and the Advanced Encryption Standard (AES) will be more secure than DES.

RFC 3562 [I4] analyzes the rekeying requirements for the TCP MD5 signature option. The analysis provided in RFC 3562 is also applicable to this specification as the analysis is independent of data patterns.

## 11. IPsec Protection Barrier and SPD

The IPsec protection barrier MUST be around the OSPF protocol. Therefore, all the inbound and outbound OSPF traffic goes through IPsec processing.

The SPD selection function MUST return an SPD with the following rule for all the interfaces that have OSPFv3 authentication/confidentiality disabled.

No.	source	destination	protocol	action
1	any	any	OSPF	bypass

The SPD selection function MUST return an SPD with the following rules for all the interfaces that have OSPFv3 authentication/confidentiality enabled.

No.	source	destination	protocol	action
2	fe80::/10	any	OSPF	protect
3	fe80::/10	any	ESP/OSPF or AH/OSPF	protect
4	src/128	dst/128	OSPF	protect
5	src/128	dst/128	ESP/OSPF or AH/OSPF	protect

For rules 2 and 4, action "protect" means encrypting/calculating Integrity Check Value (ICV) and adding an ESP or AH header. For rules 3 and 5, action "protect" means decrypting/authenticating the packets and stripping the ESP or AH header.

Rule 1 will bypass the OSPFv3 packets without any IPsec processing on the interfaces that have OSPFv3 authentication/confidentiality disabled.

Rules 2 and 4 will drop the inbound OSPFv3 packets that have not been secured with ESP/AH headers.

ESP/OSPF or AH/OSPF in rules 3 and 5 mean that it is an OSPF packet secured with ESP or AH.

Rules 2 and 3 are meant to secure the unicast and multicast OSPF packets that are not being exchanged over the virtual links.

Rules 4 and 5 are meant to secure the packets being exchanged over virtual links. These rules are installed after learning the virtual link end point IPv6 addresses. These rules MUST be installed in the SPD for the interfaces that are connected to the transit area for the virtual link. These rules MAY alternatively be installed on all the interfaces. If these rules are not installed on all the interfaces, clear text or malicious OSPFv3 packets with the same source and destination addresses as the virtual link end point IPv6 addresses will be delivered to OSPFv3. Though OSPFv3 drops these packets as they were not received on the right interface, OSPFv3 receives some clear text or malicious packets even when the security is enabled. Installing these rules on all the interfaces ensures that OSPFv3 does not receive these clear text or malicious packets when security is enabled. On the other hand, installing these rules on all the interfaces increases the processing overhead on the interfaces where there is no other IPsec processing. The decision of whether to install these rules on all the interfaces or on just the interfaces that are connected to the transit area is a private decision and doesn't affect the interoperability in any way. Hence it is an implementation choice.

## 12. Entropy of Manual Keys

The implementations MUST allow the administrator to configure the cryptographic and authentication keys in hexadecimal format rather than restricting it to a subset of ASCII characters (letters, numbers, etc.). A restricted character set will reduce key entropy significantly as discussed in [I2].

## 13. Replay Protection

Since it is not possible using the current standards to provide complete replay protection while using manual keying, the proposed solution will not provide protection against replay attacks.

Detailed analysis of various vulnerabilities of the routing protocols and OSPF in particular is discussed in [I3] and [I2]. The conclusion is that replay of OSPF packets can cause adjacencies to be disrupted, which can lead to a DoS attack on the network. It can also cause database exchange process to occur continuously thus causing CPU overload as well as micro loops in the network.

## 14. Security Considerations

This memo discusses the use of IPsec AH and ESP headers to provide security to OSPFv3 for IPv6. Hence, security permeates throughout this document.

OSPF Security Vulnerabilities Analysis [I2] identifies OSPF vulnerabilities in two scenarios -- one with no authentication or simple password authentication and the other with cryptographic authentication. The solution described in this specification provides protection against all the vulnerabilities identified for scenarios with cryptographic authentication with the following exceptions:

Limitations of manual key:

This specification mandates the usage of manual keys. The following are the known limitations of the usage of manual keys.

- o As the sequence numbers cannot be negotiated, replay protection cannot be provided. This leaves OSPF insecure against all the attacks that can be performed by replaying OSPF packets.
- o Manual keys are usually long lived (changing them often is a tedious task). This gives an attacker enough time to discover the keys.

- o As the administrator is manually configuring the keys, there is a chance that the configured keys are weak (there are known weak keys for DES/3DES at least).

#### Impersonating attacks:

The usage of the same key on all the OSPF routers connected to a link leaves them all insecure against impersonating attacks if any one of the OSPF routers is compromised, malfunctioning, or misconfigured.

Detailed analysis of various vulnerabilities of routing protocols is discussed in [I3].

## 15. References

### 15.1. Normative References

- [N1] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [N2] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.
- [N3] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [N4] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [N5] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [N6] Eastlake 3rd, D., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4305, December 2005.
- [N7] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

### 15.2. Informative References

- [I1] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.
- [I2] Jones, E. and O. Moigne, "OSPF Security Vulnerabilities Analysis", Work in Progress.
- [I3] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", Work in Progress.

[I4] Leech, M., "Key Management Considerations for the TCP MD5 Signature Option", RFC 3562, July 2003.

#### Acknowledgements

The authors would like to extend sincere thanks to Marc Solsona, Janne Peltonen, John Cruz, Dhaval Shah, Abhay Roy, Paul Wells, Vishwas Manral, and Sam Hartman for providing useful information and critiques on this memo. The authors would like to extend special thanks to Acee Lindem for many editorial changes.

We would also like to thank members of the IPsec and OSPF WG for providing valuable review comments.

#### Authors' Addresses

Mukesh Gupta  
Tropos Networks  
555 Del Rey Ave  
Sunnyvale, CA 94085

Phone: 408-331-6889  
EMail: mukesh.gupta@tropos.com

Nagavenkata Suresh Melam  
Juniper Networks  
1194 N. Mathilda Ave  
Sunnyvale, CA 94089

Phone: 408-505-4392  
EMail: nmelam@juniper.net

## Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 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](mailto:ietf-ipr@ietf.org).

## Acknowledgement

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