Network Working Group
Internet Draft
Document: draft-ietf-ospf-ospfv3-auth-06.txt
Expires: June 2005

M. Gupta
Nokia
N. Melam
Nokia

December 2004

Authentication/Confidentiality for OSPFv3

Status of this Memo

Abstract

This document describes means/mechanisms to provide authentication/confidentiality to OSPFv3 using an IPv6 AH/ESP Extension Header.

Copyright Notice

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].

Table of Contents

1. Introduction

OSPF (Open Shortest Path First) Version 2 [N1] defines fields AuType and Authentication in its protocol header in order 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]).

2. Transport Mode vs Tunnel Mode

Transport mode Security Association (SA) is the security association between two hosts or routers/gateways when they are acting as hosts. SA must be tunnel mode if either end of the security association is a router/gateway. OSPFv3 packets are exchanged between the routers but as the packets are destined to the routers, the routers act like hosts in this case. So transport mode SA MUST be used in order to provide required security to OSPFv3.

3. Authentication

Implementations conforming to this specification MUST support Authentication for OSPFv3.

In order to provide authentication to OSPFv3, "ESP with NULL encryption" MUST be supported and AH SHOULD be supported by the implementation.

If "ESP with NULL encryption" in transport mode is used, it will provide authentication to only OSPFv3 protocol headers but not to the IPv6 header, extension headers and options.

If AH in transport mode is used, it will provide authentication to OSPFv3 protocol headers, 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 with non-null encryption" 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 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, in order to distinguish OSPFv3 packets from the OSPFv2 packets, OSPF version field in the OSPF header cannot be used. As OSPFv2 is only for IPv4 and OSPFv3 is only for IPv6, version field in 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]

Traffic Selectors
The implementation MUST be able to use interface index, source address, destination address, protocol and direction for choosing the right security action.

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

Encryption and Authentication Algorithms
AES-CBC and HMAC-SHA1 MUST be supported as the encryption and the authentication algorithms respectively. [N6]

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

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. Manual keying MUST be used for this purpose. In manual keying SAs are statically installed on the routers and these static SAs are used to authenticate/encrypt the packets.

The following discussion explains that it is not scalable and practically infeasible to use different security associations for inbound and outbound traffic in order to provide the required "one to many" security. Therefore, the implementations MUST use manually configured keys with same SA for inbound and outbound traffic (as shown in Figure 3).

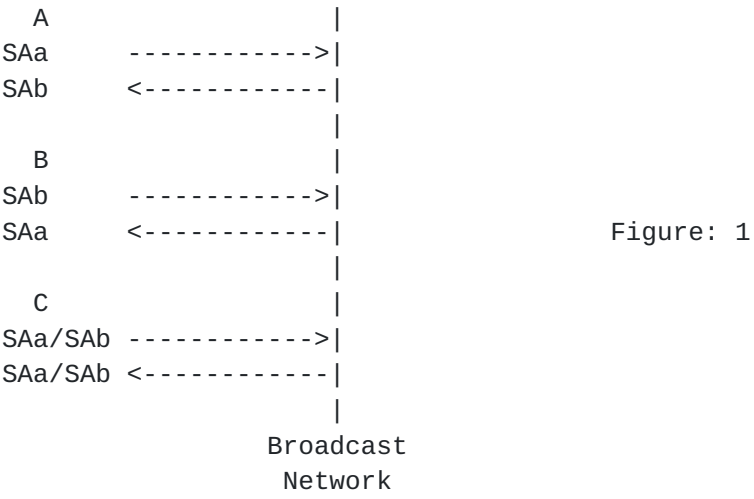

Figure: 1

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 out using SAa then only A will be able to understand it or if C sends the packets out using SAb then only B will be able to understand it. Since the packets are multicast packets and they are going to be processed by both A and B, there is no SA for C to use so that A and B both can understand it.

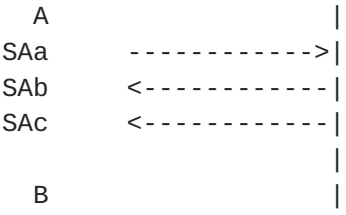

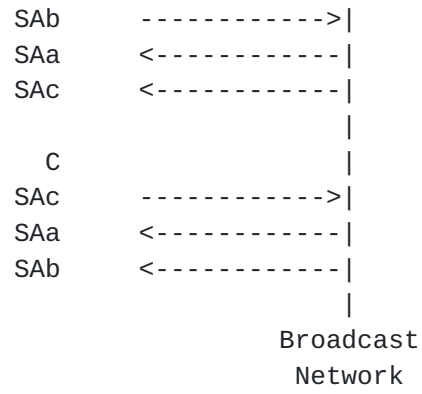

Figure: 2

The problem can be solved by configuring SAs for all the nodes on all the nodes 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 every node will need to be configured with a large number of SAs and addition of a node in the network will cause addition of another SA on all the nodes.

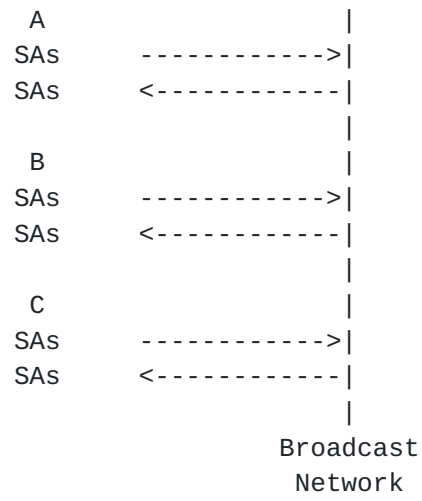

Figure: 3

The problem can also be solved by using the same SA for inbound and outbound traffic as shown in Figure 3.

8. SA Granularity and Selectors

The user SHOULD be given a choice to share the same SA among multiple interfaces or using 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 protocol header to be used as the selectors, it is not possible to use different SAs for different instances of OSPFv3 running over the same interface. Therefore, all the instances of OSPFv3 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

Different SA than the SA of underlying interface MUST be provided for virtual links. Packets sent out on virtual links use unicast non-link local IPv6 addresses as the IPv6 source address and all the other packets use multicast and unicast link local addresses. This difference in the IPv6 source address is used in order to differentiate the packets sent on interfaces and virtual links.

As the end point IP addresses of the virtual links are not known at the time of configuration, the secure channel for these packets needs to be set up dynamically. The end point IP addresses of virtual links are learned during the routing table build up process. The packet exchange over the virtual links starts only after the discovery of end point IP addresses. In order to provide security to these exchanges, the routing module should setup a secure IPsec channel dynamically once it acquires the required information.

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-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 simply suggests using one of the router's own site-local or global IPv6 addresses. In order to install the required security rules for virtual links, the source address also needs to be predictable. So the routers that implement this specification MUST change the way the source and destination addresses are chosen for the packets exchanged over virtual links when the security is enabled on that virtual link.

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 as being 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 as being concatenated and are ordered by ascending Link State ID.

This makes both the source and destination addresses of the packets exchanged over the virtual link, predictable on both the routers for security purposes.

10. Rekeying

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

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 the 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 the 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 will have a single inbound and outbound SA for OSPFv3 on the link 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 involved routers. As this value can vary significantly depending upon the implementation and the deployment, it is left to the administrator to choose the appropriate value.

10.3 Rekeying Interval

This section analyzes the security provided by the 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 is used with the HMAC-MD5 authentication. Any other algorithm combinations will at least be as hard to break as the one mentioned above as shown by the following examples:

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 is not applicable as this specification mandates the authentication when OSPFv3 security is enabled

O DES Encryption and HMAC-MD5 Authentication will be more secure because of the additional security provided by DES

O Other encryption algorithms like 3DES, AES will be more secure than DES

RFC 3562 [I4] analyzes the rekeying requirements for the TCP MD5 signature option. The analysis provided in this RFC is also applicable to OSPFv3 security specification as the analysis is independent of data patterns.

11. IPsec rules

The following set of transport mode rules can be installed in a typical IPsec implementation to provide the authentication/confidentiality to OSPFv3 packets.

Outbound Rules for interface running OSPFv3 security:

No.	source	destination	protocol	action
1	fe80::/10	any	OSPF	apply

Outbound Rules for virtual links running OSPFv3 security:

No.	source	destination	protocol	action
2	src/128	dst/128	OSPF	apply

Inbound Rules for interface running OSPFv3 security:

No.	source	destination	protocol	action
3	fe80::/10	any	ESP/OSPF or AH/OSPF	apply
4	fe80::/10	any	OSPF	drop

Inbound Rules for virtual links running OSPFv3 security:

No.	source	destination	protocol	action
5	src/128	dst/128	ESP/OSPF or AH/OSPF	apply
6	src/128	dst/128	OSPF	drop

For outbound rules, action "apply" means encrypting/calculating ICV and adding ESP or AH header. For inbound rules, action "apply" means decrypting/authenticating the packets and stripping ESP or AH header.

Rules 4 and 6 are to drop the insecure OSPFv3 packets without 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 1, 3 and 4 are meant to secure the unicast and multicast OSPF packets that are not being exchanged over the virtual links. These rules MUST be installed only in the security policy database (SPD) of the interface running OSPFv3 security.

Rules 2, 5 and 6 are meant to secure the packets being exchanged over virtual links. These rules are dynamically installed after learning the end point IP addresses of a virtual link. These rules MUST be installed on at least 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 same source and destination addresses as virtual link end point addresses will be delivered to OSPFv3. Though OSPFv3 drops these packets because they were not received on the right interface, OSPFv3 receives some clear text or malicious packets even when the security is on. Installing these rules on all the interfaces insures that OSPFv3 does not receive these clear text or malicious packets when security is turned on. On the other hand installing these rules on all the interfaces increases the processing overhead on the interfaces where there is no IPsec processing otherwise. The

decision of installing 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. So this decision is left to the implementers.

12. Mandatory Encryption and Authentication Algorithms

The implementation MUST allow the user to choose AES-CBC [N8] as the encryption algorithm and HMAC-SHA1-96 [N9] as the authentication algorithm for securing OSPFv3 packets.

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

13. Entropy of manual keys

The implementations MUST allow the administrator to configure the cryptographic and authentication keys in hexadecimal format instead of restricting it a subset of ASCII characters (letters, numbers etc). Otherwise the entropy of the keys reduces significantly as discussed in [I2].

14. Replay Protection

As it is not possible as per the current standards to provide complete replay protection while using manual keying, the proposed solution will not provide protection against replay attacks. This section analyzes the protection built in OSPF protocol to guard itself against replay attacks.

OSPF protocol has the following types of packets:

Hello: OSPF does not have any protection against replayed hello messages thus OSPF will be vulnerable to the attacks performed by replaying hello messages even after enabling security.

Database Description: These packets are exchanged while building the adjacencies to describe the OSPF database. There are no known threats of replaying these type of packets.

Link State Request, Link State Acknowledgment and Link State Updates: Fields LS age, LS Sequence Number and LS checksum in LSA header are kept intact in OSPFv3. Though these fields do not provide the complete protection, they certainly help against replay attacks of Link State packets.

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

Security Considerations

This memo discusses the use of IPsec AH and ESP headers in order 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 security against all the vulnerabilities identified for scenario 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.

- As the sequence numbers can not be negotiated, replay protection can not be provided. This leaves OSPF insecure against all the attacks that can be performed by replaying OSPF packets.

- Manual keys are usually long lived (changing them very often is a tedious task). This gives an attacker enough time to discover the keys.

- 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 routers on the same link for securing OSPF leaves it insecure against impersonating attacks if one of the routers is compromised, malfunctioning or misconfigured.

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

Normative References

N1. Moy, J., "OSPF version 2", 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 XXXX, date [Note to RFC-Editor: Replace XXXX with the number of the RFC 2401 replacement].

N4. Kent, S., "IP Encapsulating Security Payload (ESP)", RFC XXXY, date [Note to RFC-Editor: Replace XXXY with the number of the RFC 2406 replacement].

N5. Kent, S., "IP Authentication Header (AH)", RFC XXXZ, date [Note to RFC-Editor: Replace XXXZ with the number of the RFC 2402 replacement].

N6. Eastlake, D., "Cryptographic Algorithm Implementation Requirements For ESP And AH", RFC XXYY, date [Note to RFC-Editor: Replace XXYY with the number of the RFC that the draft draft-ietf-ipsec-esp-ah-algorithms-02.txt gets].

N7. Bradner, S., "Key words for use in RFCs to Indicate Requirement Level", BCP 14, RFC 2119, March 1997.

N8. Frankel, S., Glenn, R. and S. Kelly, "The AES-CBC Cipher Algorithm and Its Use with IPsec", RFC 3602, September 2003.

N9. Madson, C. and R. Glenn, "The Use of HMAC-SHA-1-96 within ESP and AH", RFC 2404, November 1998.

Informative References

I1. Kaufman, C., "The Internet Key Exchange (IKEv2) Protocol", RFC XXZZ, date [Note to RFC-Editor: Replace XXZZ with the number of the RFC 2409 replacement].

I2. Jones, E. and O. Moigne, "OSPF Security Vulnerabilities Analysis", draft-ietf-rpsec-ospf-vuln-01.txt, work in progress.

I3. Barbir, A., Murphy, S. and Y. Yang, "Generic Threats to Routing Protocols", draft-ietf-rpsec-routing-threats-07.txt, work in progress.

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

Acknowledgments

Authors would like to extend sincere thanks to Marc Solsona, Janne Peltonen, John Cruz, Dhaval Shah, Abhay Roy and Paul Wells for providing useful information and critiques in order to write this memo.

We would also like to thank IPsec and OSPF WG people to provide valuable review comments.

Authors' Addresses

Mukesh Gupta
Nokia
313 Fairchild Drive
Mountain View, CA 94043
Phone: 650-625-2264
Email: Mukesh.Gupta@nokia.com

Nagavenkata Suresh Melam
Nokia
313 Fairchild Drive
Mountain View, CA 94043
Phone: 650-625-2949
Email: Nagavenkata.Melam@nokia.com

Full copyright statement

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.