Internet Engineering Task Force
Internet Draft
Expires: January 2008

Y. Liu
Huawei
R. White
B. Weis
M. Barnes
Cisco

July 7, 2007

OSPFv3 Automated Group Keying Requirements
draft-liu-ospfv3-automated-keying-req-01.txt

Status of this Memo

This Internet-Draft will expire on January 26, 2008.

Abstract

RFC4552 describes how to provide authentication/confidentiality to OSPFv3 using IPsec. It specifies that same IPsec SA parameters be configured for both inbound and outbound SAs to provide the "one to many" security for multicast OSPFv3 communications over broadcast links (e.g., Ethernet). Manual keying is specified as the mandatory and default group key management solution. However, issues of scalability and security exist with manual keying. It is better to replace manual keying with automated group key management. This document discusses the requirements on OSPFv3 automated group key management, assuming that the centralized group key management architecture introduced in [RFC4046] is used.

Table of Contents

1. Introduction

OSPFv3 [RFC2740] relies on IPSec to provide integrity, authentication, and/or confidentiality. RFC4552 describes how to provide authentication/confidentiality to OSPFv3 using AH/ESP. In Section 7 of RFC4552, it is proved that best scalability and feasibility can be achieved if the same parameters are used for both inbound and outbound AH/ESP SAs to provide the "one to many" communication security while running OSPFv3 over broadcast interfaces. It means that group security model be used to protect OSPFv3 multicast communications over broadcast network.

However, RFC4552 specifies Manual Keying [RFC4301] as the default group key management method, which is neither scalable nor secure. This document discusses replacing manual keying with automated keying using either a "one to many" or "many to many" communication model. It is based on the fact that several GKM protocols have been, or being, standardized by MSEC working group [RFC3547] [RFC3830] [RFC4535] [GKDP], which makes it feasible to provide automated group key management for OSPFv3 using GKM protocols.

Meanwhile, [I-D: draft-ietf-msec-ipsec-extensions] describes multicast extensions to IPsec, which makes it feasible to provide group security to OSPFv3 using standard multicast IPsec.

This document describes the requirements on OSPFv3 automated group key management. It is assumed that the centralized group security architecture [RFC3740] and group key management architecture [RFC4046] would be used, where group SAs are centrally managed on separate key server(s). And one of MSEC GKM protocols will be chosen as the group keying protocols. Use of group key agreement techniques, for example, Cliques [CLIQUES], is out of consideration.

2. Terminology

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.

Group Key Management (GKM) Protocol

> A group key management protocol used by a GCKS to distribute IPsec Security Association policy and keying material. A GKM protocol is used when a group of IPsec devices require the same SAs. For example, when an IPsec SA describes an IP multicast destination, the sender and all receivers MUST have the group SA.

Group Controller Key Server (GCKS)

A Group Key Management (GKM) protocol server that manages IPsec state for a group. A GCKS authenticates and provides the IPsec SA policy and keying material to GKM group members.

GC and KS may also be acted by different entities. GC is responsible defining and distributing group policy and authorization, while KS providing group keying service to a specific group.

Group Member

An IPsec device that belongs to a group. A Group Member is authorized to be a Group Speaker and/or a Group Receiver.

Group Security Association (Group SA/GSA)

A collection of IPsec Security Associations (SAs) and GKM SAs necessary for a Group Member to receive key updates. A GSA describes the working policy for a group. Refer to RFC 4046 [RFC4046] for additional information.

State Synchronization (SS)

A generalization of OSPFv3 communications that occur after two routers discover each other. It includes neighbor discovering, exchanging DDs and LSAs, etc.

3. General Requirements

Requirements discussed in this section are considered general to all keying solutions.

3.1. Authentication and Authorization of Routers

When a router uses a GKM protocol to contact a GCKS, the GCKS authenticates the router and verifies that the router is authorized to participate in the group. Both steps are necessary in order for the Group SA to ensure that the Group SA is kept private to OSPF routers that are allowed to participate in OSPF. Both authentication and authorization MUST be enforced before a Group SA is distributed to a router.

3.2. Secure Distribution of Group SA

Eavesdroppers are not able to access the group keys by intercepting the group SA distribution packets. All existing GKM protocols can meet this requirement.

3.3. Storing Capability of Keys, Authorizations, and Policies

It is expected that a network start up autonomously without a provisioning step after rebooting. To actualize that target, routers SHOULD to be able to store group security context information, such as group policy, authorization, neighbor list, and GSA, etc., in their storages (i.e., flash, hard disk). After rebooting, if this policy is stored a router quickly resumes its group security context to the same state as that of before rebooting and keeps in synchronization with its neighbors. After every router has done that, the OSPF SS process can be securely re-done in the network. Fast routes convergence is assured during this process.

Group security context may change during rebooting. In such cases, synchronization of group security context among routers can not be achieved just by caching and resuming mechanisms. Other measures are needed. Specific requirements on this topic will be discussed case by case in the following sections.

4. GCKS Deployment Example and its Specific Requirements

MSEC GKM protocols, such as GSAKMP and GDOI, are based on a client/server model. This means these protocols rely on reachability between clients and servers for the clients to obtain the group SA from the server. In this case, the GKM is providing protection for OSPF, which is an essential component in providing reachability between the clients and servers. Hence, the client/server model breaks down in this situation.

To overcome this problem, the group SA must be locally available to each group member (each OSPFv3 router). Possible solutions to this circular dependency and their specific requirements are presented in this section.

The expected network where OSPFv3 multicast communications occur and group SA is assumed to be deployed is like the network N1 shown in Fig.1. This network is used in this memo to derive the specific requirements.

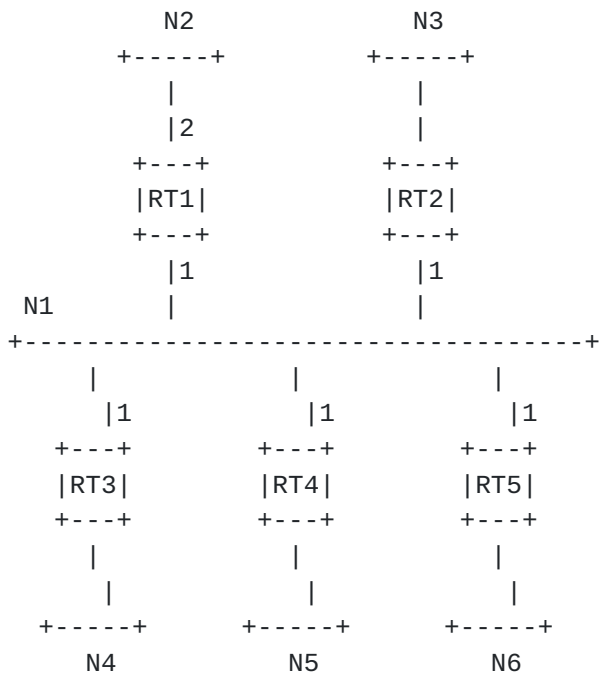

Figure 1 : The scenario used to derive the requirements

N1 is a broadcast network, where GSA is used to protect the OSPFv3 multicast communications among RT1~RT5. Group members attached to N1, including RT1~RT5, get the group SA from GCKS. For convenience of this example, broadcast interfaces attached to N1 are numbered as #1, which each router recognizes is part of a single group.

Types of networks N2~N6 are not fixed. They MAY be either broadcast, P2MP, or NBMA. It is assumed that a GSA plays locally on a multicast network. Whether N2~N6 will share the same GSA with N1 is out the scope of this memo even if they are broadcast type either.

4.1. Decentralized GCKSs

In this deployment example, there will be a GCKS for each multicast network. The GCKS may be either a physical server or a logical one that is acted by an OSPFv3 router, e.g., one router from RT1~RT5 is selected as N1's GCKS. The GCKS and group members are located in the same network, and share a GSA that is used only on that network. (If a group member is attached to multiple networks, then it will install a unique GSA for each network.) The GCKS may deliver either a single SA used by each group member (a "many to many" SA) or deliver an individual SA used by a unique sender (several "one to many" SAs). GSA messages can be distributed from GCKS to group members within one hop. Group members can download the GSA from their GCKS before they establish their routes. It means no routes need to be pre-available

before the OSPFv3 router members run the GKM protocol to download GSA from their GCKS. This solution can solve the contradiction mentioned above. It applies to both small and large ASs.

Requirements introduced by this solution are illustrated as follows.

4.1.1. Single Point of GCKS Failure

This is a requirement common to all client/server protocols/applications. The GCKS MAY be out of service because of attacks, accidents, reboots, etc. Measures MUST be prepared to provide continuous keying service in case of single point of GCKS failure.

4.1.2. GCKS Selecting/Switchover Issue

To deal with the problem of single point of GCKS failure, at least 2 GCKSs need to be deployed in each multicast network. If GCKSs are acted by the routers, switchover mechanisms are necessary to the GKM protocol to provide continuous keying service when the running GCKS becomes out of service and its service needs to be transferred to the backup GCKS.

Another requirement is GCKS selecting mechanism, which is very like OSPF DR/BDR selecting mechanism. It is necessary in the case of a logical GCKS. Routers MUST be able to select their master/backup GCKS when they start up.

4.1.3. Authorization and Authentication of GCKS

When routers begin to select their GCKS and backup GCKS, they need to authenticate the candidates and verify that the selected GCKSs are authorized to act as the GCKSs in the group. Both steps are necessary in order to ensure that attackers can not become the GCKS by cheating others. Both authentication and authorization MUST be enforced during a GCKS selecting process.

4.1.4. New Joiner Issue

New router may join an existing network. It should be able to automatically discover the local GCKS and contact it to get current GSA so it can securely perform OSPF SS process with its neighbors.

However, dynamic changes of group membership may not be pre-known. Thus, the list of authorized routers may not be pre-installed on each GCKS. GCKS must be able to deal with such dynamic changes of authorizations. If a new router registers with it, GCKS MUST be able

to authenticate the new one and verify whether it is authorized to access the GSA even if the new router is not in the GCKS's authorized routers list.

4.2. Decentralized KSs, Centralized GC

In this solution, KS and GC are separate roles. KS is logical and is deployed per network; while GC is centrally deployed. This solution allows for a centralized group management, but allows the KSs to act autonomously. The GC is responsible for defining the group policy and authorization, and pushing them to each KS. Each key server then distributes a GSA conforming to the group policy. As in the Decentralized GCKS case, each key server distributes a GSA that is used only on a single network and SAs may be either have a scope of "many to many" or "one to many".

Requirements in section 4.1.1 and 4.1.2 apply. New requirements introduced by this solution are illustrated as follows.

4.2.1. Bootstrapping Issue

When the network originally starts up, there are no routes, and candidate KS/backup KS routers do not have the authorization information of group membership. After the KS/backup KS are selected, they will not be able to provide the keying service because they don't know the information of authorization and authentication of routers. Measure MUST be provided to handle such issue.

4.2.2. New Joiner Issue

GC is responsible for determining list of authorized routers. Measures must be provided to keep authorization information in synchronization between GC and KS. For example, if a new router is authorized by GC to join the network, KS must be able to authenticate it and verify it is an authorized one even if it does not know the new router and cannot contact with the GC.

4.2.3. Authorization and Authentication of KS

When routers begin to select their master/backup KSs, they must be able to authenticate the candidates and verify that the selected KSs are authorized to play such roles. Both steps of authentication and authorization check are necessary in order to ensure that attackers can not become the KS by cheating others. Both authentication and authorization MUST be enforced during a KS selecting process.

4.3. Decentralized Delegates, Centralized GCKS

In this solution, there is a delegate in each OSPFv3 multicast network. GSA messages are created by the centralized GCKS which may serve multiple OSPFv3 networks at the same time. A delegate is responsible for receiving GSA messages from the GCKS and re-distributing them to its local group members. A delegate can be either physical or logical.

Requirements introduced by this solution are as follows.

4.3.1. Bootstrapping Issue

When neighboring routers initially start, e.g., routers attached to N1, they have no routes to the key server and so, can not download the GSA from the centralized GCKS. Security measures MUST be provided to protect the initial communications among OSPFv3 routers before their adjacencies are established and routes are calculated out.

4.3.2. New Joiner Issue

When a new router joins, it will perform OSPF SS process with its neighbors, and then calculate its routes. Before routes are calculated out, the new coming router cannot contact the remote GCKS to get current GSA. Measures MUST be provided to let the new joining router get current GSA from GCKS so that it can securely communicate with its neighboring routers.

Another problem is that routers may reboot singly. A rebooting router needs to re-join the network after it restarts. However, during the rebooting process, GSA may be updated by the GCKS and distributed to other running routers. Therefore, after a router restarts and resumes the GSA from its storage (e.g., flash), it may find its local GSA is out of synchronization with its neighbors and thus, cannot establish relationship with neighboring routers before it gets the latest GSA. Measures must be provided to let rebooting routers make sure whether its local GSA is out of synchronization with its neighbors and, if yes, be able to get the updated GSA to keep its GSA in synchronization with its neighbors.

4.3.3. Single Point of Delegate Failure

The delegate may crash or reboot. In such cases, the GSA redistribution service will be not available. Measures MUST be prepared to assure continuous GSA relaying service.

4.3.4. Delegate Selecting/Switchover Issue

To deal with the single point of delegate failure, at least 2 delegates must be deployed on every network. In the case of logical delegates, delegate selecting mechanism is required for the GKM protocols so routers can autonomously elect their local master/backup delegates of their local network.

Switchover mechanism is needed to ensure that the backup delegate can immediately supersede the master delegate in case that the later one is out of service.

4.3.5. Authorization and Authentication of Delegate

In the delegates selecting process, routers need to authenticate the candidates and verify that the selected delegates are authorized to play such roles. Both steps of authenticating and authorization checking are necessary in order to ensure that an attacker can not become the delegate by cheating others.

5. Security Considerations

This document lists requirements for automated group keying of OSPFv3. Several possible solutions and their specific requirements are discussed. This section will discuss the security risk of the mentioned solutions.

5.1. Decentralized GCKS

TBD

5.2. Decentralized KS, Centralized GC

TBD

5.3. Decentralized Delegate, Centralized GCKS

TBD

6. IANA Considerations

This document has no IANA considerations.

7. Acknowledgments

Thank Zengjie Kou, Jinliang Feng and Zhenhai Li for technical discussions.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2234] Crocker, D. and Overell, P. (Editors), "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, Internet Mail Consortium and Demon Internet Ltd., November 1997.

[RFC2740] Coltun, R., Ferguson, D. and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.

[RFC3547] Baugher, M., Weis, B., Hardjono, T. and H. Harney, "The Group Domain of Interpretation", RFC3547, July 2003.

[RFC3740] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", RFC3740, March 2004.

[RFC3830] Arkko, J., Carrara, E., Lindholm, F., Naslund, M. and K. Norrman, "MIKEY: Multimedia Internet KEYing", RFC3830, August 2004.

[RFC4046] Baugher, M., Canetti, R., Dondeti, L. and F. Lindholm, "Multicast Security (MSEC) Group Key Management Architecture", RFC4046, April 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC4301, December 2005.

[RFC4535] Harney, H., Meth, U., Colegrove, A. and G. Gross, "GSAKMP: Group Secure Association Key Management Protocol", RFC4535, June 2006.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC4552, June 2006.

[RFC4593] Barbir, A. and Murphy, S. and Y. Yang, "Generic Threats to Routing Protocols", RFC4593, October 2006.

8.2. Informative References

[GKDP] Dondeti, L., Xiang, J. and S. Rowles, "GKDP: Group Key Distribution Protocol", work in progress, October 2006.

[I-D: draft-ietf-msec-ipsec-extensions] Weis, B., Gross, G. and D. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", work in progress, October 2006.

[CLIQUES] Steiner, M., Tsudik, G. and M. Waidner, "CLIQUES: A New Approach to Group key Agreement", IEEE ICDCS'98 , MAY 1998.

Author's Addresses

Ya Liu
Huawei Technologies
Kuike Bld., No.9 Xinxi Rd., Shang-Di Information Industry Base
Hai-Dian District, Beijing 100086
P.R. China

Phone: 8610-82836072
Email: liuya@huawei.com

Russ White
Cisco Systems
7025 Kit Creek Road
RTP, NC 27709
USA

Email: riw@cisco.com

Brian Weis
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134-1706
USA

Phone: +1-408-526-4796
Email: bew@cisco.com

Michael Barnes
Cisco, Inc.
170 W. Tasman Drive
San Jose, CA 95134
USA

Phone: +1-408-525-2785
Email: mjbarnes@cisco.com

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

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