

Network Working Group
Request for Comments: 4046
Category: Informational

M. Baugher
Cisco
R. Canetti
IBM
L. Dondeti
Qualcomm
F. Lindholm
Ericsson
April 2005

Multicast Security (MSEC) Group Key Management Architecture

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document defines the common architecture for Multicast Security (MSEC) key management protocols to support a variety of application, transport, and network layer security protocols. It also defines the group security association (GSA), and describes the key management protocols that help establish a GSA. The framework and guidelines described in this document permit a modular and flexible design of group key management protocols for a variety of different settings that are specialized to applications needs. MSEC key management protocols may be used to facilitate secure one-to-many, many-to-many, or one-to-one communication.

Table of Contents

1. Introduction: Purpose of this Document	2
2. Requirements of a Group Key Management Protocol	4
3. Overall Design of Group Key Management Architecture	6
3.1. Overview	6
3.2. Detailed Description of the GKM Architecture	8
3.3. Properties of the Design	11
3.4. Group Key Management Block Diagram	11
4. Registration Protocol	13
4.1. Registration Protocol via Piggybacking or Protocol Reuse ..	13
4.2. Properties of Alternative Registration Exchange Types	14

- 4.3. Infrastructure for Alternative Registration Exchange Types15
- 4.4. De-registration Exchange16
- 5. Rekey Protocol16
 - 5.1. Goals of the Rekey Protocol17
 - 5.2. Rekey Message Transport and Protection17
 - 5.3. Reliable Transport of Rekey Messages18
 - 5.4. State-of-the-art on Reliable Multicast Infrastructure20
 - 5.5. Implosion21
 - 5.6. Incorporating Group Key Management Algorithms22
 - 5.7. Stateless, Stateful, and Self-healing Rekeying Algorithms22
 - 5.8. Interoperability of a GKMA23
- 6. Group Security Association24
 - 6.1. Group Policy24
 - 6.2. Contents of the Rekey SA25
 - 6.2.1. Rekey SA Policy26
 - 6.2.2. Group Identity27
 - 6.2.3. KEKs27
 - 6.2.4. Authentication Key27
 - 6.2.5. Replay Protection27
 - 6.2.6. Security Parameter Index (SPI)27
 - 6.3. Contents of the Data SA27
 - 6.3.1. Group Identity28
 - 6.3.2. Source Identity28
 - 6.3.3. Traffic Protection Keys28
 - 6.3.4. Data Authentication Keys28
 - 6.3.5. Sequence Numbers28
 - 6.3.6. Security Parameter Index (SPI)28
 - 6.3.7. Data SA Policy28
- 7. Scalability Considerations29
- 8. Security Considerations31
- 9. Acknowledgments32
- 10. Informative References33

1. Introduction: Purpose of this Document

This document defines a common architecture for Multicast Security (MSEC) key management protocols to support a variety of application-, transport-, and network-layer security protocols. It also defines the group security association (GSA) and describes the key management protocols that help establish a GSA. The framework and guidelines described in this document permit a modular and flexible design of group key management protocols for a variety of different settings that are specialized to applications needs. MSEC key management protocols may be used to facilitate secure one-to-many, many-to-many, or one-to-one communication.

Group and multicast applications in IP networks have diverse security requirements [TAXONOMY]. Their key management requirements, briefly reviewed in Section 2.0, include support for internetwork-, transport- and application-layer security protocols. Some applications achieve simpler operation by running key management messaging over a pre-established secure channel (e.g., TLS or IPsec). Other security protocols benefit from a key management protocol that can run over an already-deployed session initiation or management protocol (e.g., SIP or RTSP). Finally, some benefit from a lightweight key management protocol that requires few round trips. For all these reasons, application-, transport-, and IP-layer data security protocols (e.g., SRTP [RFC3711] and IPsec [RFC2401]) benefit from different group key management systems. This document defines a common architecture and design for all group key management (GKM) protocols.

This common architecture for group key management is called the MSEC group key management architecture. It is based on the group control or key server model developed in GKMP [RFC2094] and assumed by group key management algorithms such as LKH [RFC2627], OFT [OFT], and MARKS [MARKS]. There are other approaches that are not considered in this architecture, such as the highly distributed Cliques group key management protocol [CLIQUES] or broadcast key management schemes [FN93,Wool]. MSEC key management may in fact be complementary to other group key management designs, but the integration of MSEC group key management with Cliques, broadcast key management, or other group key systems is not considered in this document.

Key management protocols are difficult to design and validate. The common architecture described in this document eases this burden by defining common abstractions and an overall design that can be specialized for different uses.

This document builds on and extends the Group Key Management Building Block document of the IRTF SMuG research group [GKMBB] and is part of the MSEC document roadmap. The MSEC architecture [MSEC-Arch] defines a complete multicast or group security architecture, of which key management is a component.

The rest of this document is organized as follows. Section 2 discusses the security, performance and architectural requirements for a group key management protocol. Section 3 presents the overall architectural design principles. Section 4 describes the registration protocol in detail, and Section 5 does the same for rekey protocol. Section 6 considers the interface to the Group Security Association (GSA). Section 7 reviews the scalability issues for group key management protocols and Section 8 discusses security considerations.

2. Requirements of a Group Key Management Protocol

A group key management (GKM) protocol supports protected communication between members of a secure group. A secure group is a collection of principals, called members, who may be senders, receivers, or both receivers and senders to other members of the group. Group membership may vary over time. A group key management protocol helps to ensure that only members of a secure group can gain access to group data (by gaining access to group keys) and can authenticate group data. The goal of a group key management protocol is to provide legitimate group members with the up-to-date cryptographic state they need for secrecy and authentication.

Multicast applications, such as video broadcast and multicast file transfer, typically have the following key management requirements (see also [TAXONOMY]). Note that the list is neither applicable to all applications nor exhaustive.

1. Group members receive security associations that include encryption keys, authentication/integrity keys, cryptographic policy that describes the keys, and attributes such as an index for referencing the security association (SA) or particular objects contained in the SA.
2. In addition to the policy associated with group keys, the group owner or the Group Controller and Key Server (GCKS) may define and enforce group membership, key management, data security, and other policies that may or may not be communicated to the entire membership.
3. Keys will have a pre-determined lifetime and may be periodically refreshed.
4. Key material should be delivered securely to members of the group so that they are secret, integrity-protected and verifiably obtained from an authorized source.
5. The key management protocol should be secure against replay attacks and Denial of Service(DoS) attacks (see the Security Considerations section of this memo).
6. The protocol should facilitate addition and removal of group members. Members who are added may optionally be denied access to the key material used before they joined the group, and removed members should lose access to the key material following their departure.

7. The protocol should support a scalable group rekey operation without unicast exchanges between members and a Group Controller and Key Server (GCKS), to avoid overwhelming a GCKS managing a large group.
8. The protocol should be compatible with the infrastructure and performance needs of the data security application, such as the IPsec security protocols AH and ESP, and/or application layer security protocols such as SRTP [RFC3711].
9. The key management protocol should offer a framework for replacing or renewing transforms, authorization infrastructure, and authentication systems.
10. The key management protocol should be secure against collusion among excluded members and non-members. Specifically, collusion must not result in attackers gaining any additional group secrets than each of them individually are privy to. In other words, combining the knowledge of the colluding entities must not result in revealing additional group secrets.
11. The key management protocol should provide a mechanism to securely recover from a compromise of some or all of the key material.
12. The key management protocol may need to address real-world deployment issues such as NAT-traversal and interfacing with legacy authentication mechanisms.

In contrast to typical unicast key and SA negotiation protocols such as TLS and IKE, multicast group key management protocols provide SA and key download capability. This feature may be useful for point-to-point as well as multicast communication, so that a group key management protocol may be useful for unicast applications. Group key management protocols may be used for protecting multicast or unicast communications between members of a secure group. Secure sub-group communication is also plausible using the group SA.

There are other requirements for small group operation with many all members as potential senders. In this case, the group setup time may need to be optimized to support a small, highly interactive group environment [RFC2627].

The current key management architecture covers secure communication in large single-sender groups, such as source-specific multicast groups. Scalable operation to a range of group sizes is also a desirable feature, and a better group key management protocol will support large, single-sender groups as well as groups that have many

senders. It may be that no single key management protocol can satisfy the scalability requirements of all group-security applications.

It is useful to emphasize two non-requirements: technical protection measures (TPM) [TPM] and broadcast key management. TPM are used for such things as copy protection by preventing the device user from getting easy access to the group keys. There is no reason why a group key management protocol cannot be used in an environment where the keys are kept in a tamper-resistant store, using various types of hardware or software to implement TPM. For simplicity, however, the MSEC key management architecture described in this document does not consider design for technical protection.

The second non-requirement is broadcast key management when there is no back channel [FN93,JKKV94] or for a non-networked device such as a digital videodisc player. We assume IP network operation with two-way communication, however asymmetric, and authenticated key-exchange procedures that can be used for member registration. Broadcast applications may use a one-way Internet group key management protocol message and a one-way rekey message, as described below.

3. Overall Design of Group Key Management Architecture

The overall group key management architecture is based upon a group controller model [RFC2093,RFC2094,RFC2627,OFT,GSAKMP,RFC3547] with a single group owner as the root-of-trust. The group owner designates a group controller for member registration and GSA rekeying.

3.1. Overview

The main goal of a group key management protocol is to securely provide group members with an up-to-date security association (SA), which contains the needed information for securing group communication (i.e., the group data). We call this SA the Data SA. In order to obtain this goal, the group key management architecture defines the following protocols.

(1) Registration Protocol

This is a unicast protocol between the Group Controller and Key Server (GCKS) and a joining group member. In this protocol, the GCKS and joining member mutually authenticate each other. If the authentication succeeds and the GCKS finds that the joining member is authorized, then the GCKS supplies the joining member with the following information:

- (a) Sufficient information to initialize the Data SA within the joining member. This information is given only if the group security policy calls for initializing the Data SA at registration, instead of, or in addition to, as part of the rekey protocol.
- (b) Sufficient information to initialize a Rekey SA within the joining member (see more details about this SA below). This information is given if the group security policy calls for a rekey protocol.

The registration protocol must ensure that the transfer of information from GCKS to member is done in an authenticated and confidential manner over a security association. We call this SA the Registration SA. A complementary de-registration protocol serves to explicitly remove Registration SA state. Members may choose to delete Registration SA state.

(2) Rekey Protocol

A GCKS may periodically update or change the Data SA, by sending rekey information to the group members. Rekey messages may result from group membership changes, from changes in group security policy, from the creation of new traffic-protection keys (TPKs, see next section) for the particular group, or from key expiration. Rekey messages are protected by the Rekey SA, which is initialized in the registration protocol. They contain information for updating the Rekey SA and/or the Data SA and can be sent via multicast to group members or via unicast from the GCKS to a particular group member.

Note that there are other means for managing (e.g., expiring or refreshing) the Data SA without interaction between the GCKS and the members. For example in MARKS [MARKS], the GCKS pre-determines TPKs for different periods in the lifetime of the secure group and distributes keys to members based on their membership periods. Alternative schemes such as the GCKS disbanding the secure group and starting a new group with a new Data SA are also possible, although this is typically limited to small groups.

Rekey messages are authenticated using one of the two following options:

- (1) Using source authentication [TAXONOMY], that is, enabling each group member to verify that a rekey message originates with the GCKS and none other.

- (2) Using only group-based authentication with a symmetric key. Members can only be assured that the rekey messages originated within the group. Therefore, this is applicable only when all members of the group are trusted not to impersonate the GCKS. Group authentication for rekey messages is typically used when public-key cryptography is not suitable for the particular group.

The rekey protocol ensures that all members receive the rekey information in a timely manner. In addition, the rekey protocol specifies mechanisms for the parties to contact the GCKS and re-synch if their keys expired and an updated key has not been received. The rekey protocol for large-scale groups offers mechanisms to avoid implosion problems and to ensure reliability in its delivery of keying material.

Although the Rekey SA is established by the registration protocol, it is updated using a rekey protocol. When a member leaves the group, it destroys its local copy of the GSA. Using a de-registration message may be an efficient way for a member to inform the GCKS that it has destroyed, or is about to destroy, the SAs. Such a message may prompt the GCKS to cryptographically remove the member from the group (i.e., to prevent the member from having access to future group communication). In large-scale multicast applications, however, de-registration can potentially cause implosion at the GCKS.

3.2. Detailed Description of the GKM Architecture

Figure 1 depicts the overall design of a GKM protocol. Each group member, sender or receiver, uses the registration protocol to get authorized and authenticated access to a particular Group, its policies, and its keys. The two types of group keys are the key encryption keys (KEKs) and the traffic encryption keys (TEKs). For group authentication of rekey messages or data, key integrity or traffic integrity keys may be used, as well. We use the term protection keys to refer to both integrity and encryption keys. For example, the term traffic protection key (TPK) is used to denote the combination of a TEK and a traffic integrity key, or the key material used to generate them.

The KEK may be a single key that protects the rekey message, typically containing a new Rekey SA (containing a KEK) and/or Data SA (containing a TPK/TEK). A Rekey SA may also contain a vector of keys that are part of a group key membership algorithm [RFC2627,OFT,TAXONOMY,SD1,SD2]. The data security protocol uses TPKs to protect streams, files, or other data sent and received by

the data security protocol. Thus the registration protocol and/or the rekey protocol establish the KEK(s) and/or the TPKs.

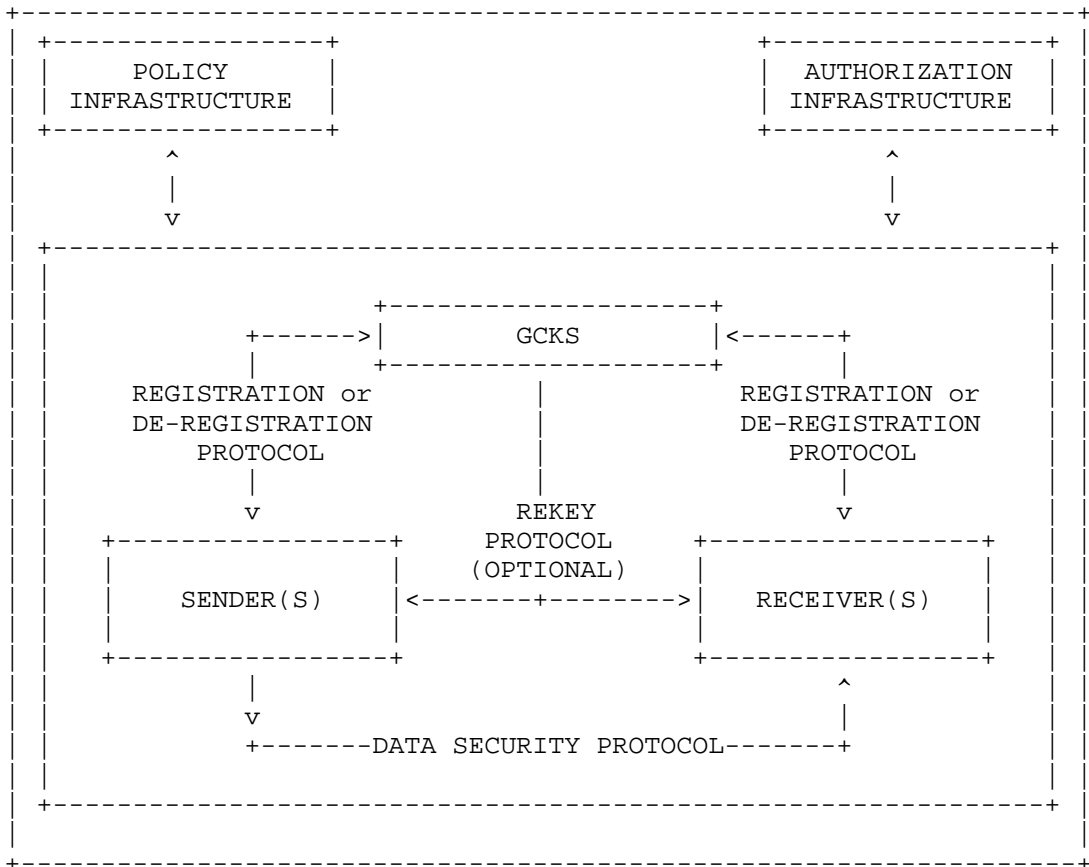


Figure 1: Group Security Association Model

There are a few distinct outcomes to a successful registration Protocol exchange.

- o If the GCKS uses rekey messages, then the admitted member receives the Rekey SA. The Rekey SA contains the group's rekey policy (note that not all of the policy need to be revealed to members), and at least a group KEK. In addition, the GCKS sends a group key integrity key for integrity protection of rekey messages. If a group key management algorithm is used for efficient rekeying, the GCKS also sends one or more KEKs as specified by the key distribution policy of the group key management algorithm.

- o If rekey messages are not used for the Group, then the admitted member receives TPKs (as part of the Data Security SAs) that are passed to the member's Data Security Protocol (as IKE does for IPsec).
- o The GCKS may pass one or more TPKs to the member even if rekey messages are used, for efficiency reasons and according to group policy.

The GCKS creates the KEK and TPKs and downloads them to each member, as the KEK and TPKs are common to the entire group. The GCKS is a separate logical entity that performs member authentication and authorization according to the group policy that is set by the group owner. The GCKS may present a credential signed by the group owner to the group member, so that member can check the GCKS's authorization. The GCKS, which may be co-located with a member or be physically separate, runs the rekey protocol to push rekey messages containing refreshed KEKs, new TPKs, and/or refreshed TPKs to members. Note that some group key management algorithms refresh any of the KEKs (potentially), whereas others only refresh the group KEK.

Alternatively, the sender may forward rekey messages on behalf of the GCKS when it uses a credential mechanism that supports delegation. Thus, it is possible for the sender, or other members, to source keying material (TPKs encrypted in the Group KEK) as it sources multicast or unicast data. As mentioned above, the rekey message can be sent using unicast or multicast delivery. Upon receipt of a TPK (as part of a Data SA) via a rekey message or a registration protocol exchange, the member's group key management functional block will provide the new or updated security association (SA) to the data security protocol. This protects the data sent from sender to receiver.

The Data SA protects the data sent on the arc labeled DATA SECURITY PROTOCOL shown in Figure 1. A second SA, the Rekey SA, is optionally established by the key management protocol for rekey messages as shown in Figure 1 by the arc labeled REKEY PROTOCOL. The rekey message is optional because all keys, KEKs and TPKs, can be delivered by the registration protocol exchanges shown in Figure 1, and those keys may not need to be updated. The registration protocol is protected by a third, unicast, SA between the GCKS and each member. This is called the Registration SA. There may be no need for the Registration SA to remain in place after the completion of the registration protocol exchanges. The de-registration protocol may be used when explicit teardown of the SA is desirable (such as when a phone call or conference terminates). The three SAs compose the GSA. The only optional SA is the Rekey SA.

Figure 1 shows two blocks that are external to the group key management protocol: The policy and authorization infrastructures are discussed in Section 6.1. The Multicast Security Architecture document further clarifies the SAs and their use as part of the complete architecture of a multicast security solution [MSEC-Arch].

3.3. Properties of the Design

The design of Section 3.2 achieves scalable operation by (1) allowing the de-coupling of authenticated key exchange in a registration protocol from a rekey protocol, (2) allowing the rekey protocol to use unicast push or multicast distribution of group and data keys as an option, (3) allowing all keys to be obtained by the unicast registration protocol, and (4) delegating the functionality of the GCKS among multiple entities, i.e., to permit distributed operation of the GCKS.

High-capacity operation is obtained by (1) amortizing computationally-expensive asymmetric cryptography over multiple data keys used by data security protocols, (2) supporting multicast distribution of symmetric group and data keys, and (3) supporting key revocation algorithms such as LKH [RFC2627,OFT,SD1,SD2] that allow members to be added or removed at logarithmic rather than linear space/time complexity. The registration protocol may use asymmetric cryptography to authenticate joining members and optionally establish the group KEK. Asymmetric cryptography such as Diffie-Hellman key agreement and/or digital signatures are amortized over the life of the group KEK. A Data SA can be established without the use of asymmetric cryptography; the TPKs are simply encrypted in the symmetric KEK and sent unicast or multicast in the rekey protocol.

The design of the registration and rekey protocols is flexible. The registration protocol establishes a Rekey SA or one or more Data SAs or both types of SAs. At least one of the SAs is present (otherwise, there is no purpose to the Registration SA). The Rekey SA may update the Rekey SA, or establish or update one or more Data SAs. Individual protocols or configurations may use this flexibility to obtain efficient operation.

3.4. Group Key Management Block Diagram

In the block diagram of Figure 2, group key management protocols run between a GCKS and member principal to establish a Group Security Association (GSA). The GSA consists of a Data SA, an optional Rekey SA, and a Registration SA. The GCKS may use a delegated principal, such as the sender, which has a delegation credential signed by the GCKS. The Member of Figure 2 may be a sender or receiver of multicast or unicast data. There are two functional blocks in Figure

2 labeled GKM, and there are two arcs between them depicting the group key-management registration (reg) and rekey (rek) protocols. The message exchanges are in the GSA establishment protocols, which are the registration protocol and the rekey protocol described above.

Figure 2 shows that a complete group-key management functional specification includes much more than the message exchange. Some of these functional blocks and the arcs between them are peculiar to an operating system (OS) or vendor product, such as vendor specifications for products that support updates to the IPsec

Security Association Database (SAD) and Security Policy Database (SPD) [RFC2367]. Various vendors also define the functions and interface of credential stores, CRED in Figure 2.

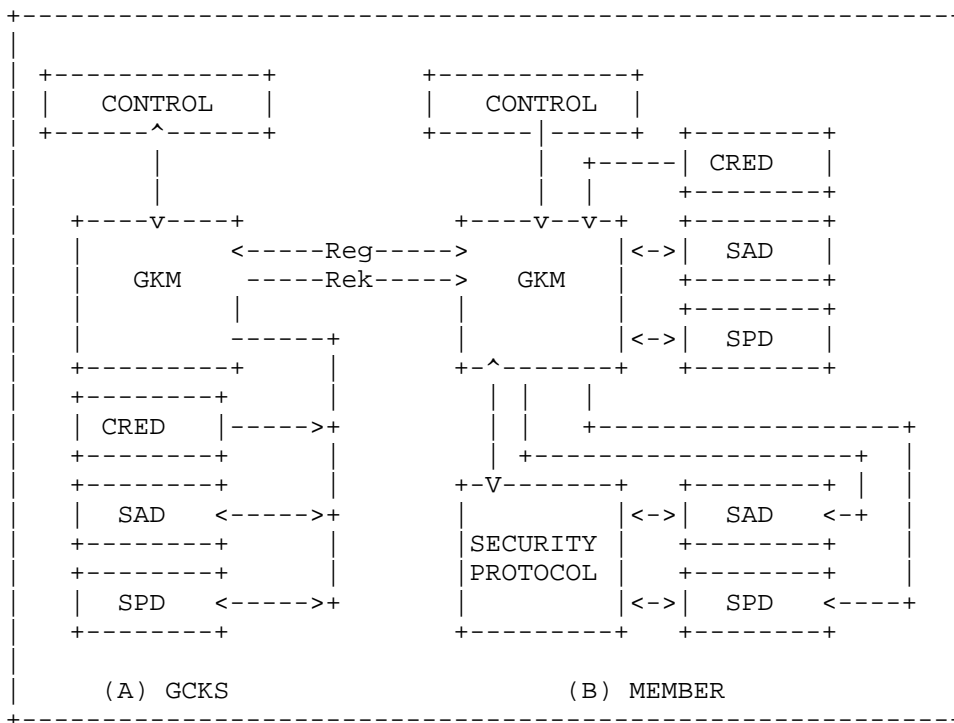


Figure 2: Group Key Management Block in a Host

The CONTROL function directs the GCKS to establish a group, admit a member, or remove a member, or it directs a member to join or leave a group. CONTROL includes authorization that is subject to group policy [GSPT] but its implementation is specific to the GCKS. For large scale multicast sessions, CONTROL could perform session

announcement functions to inform a potential group member that it may join a group or receive group data (e.g., a stream of file transfer protected by a data security protocol). Announcements notify group members to establish multicast SAs in advance of secure multicast data transmission. Session Description Protocol (SDP) is one form that the announcements might take [RFC2327]. The announcement function may be implemented in a session directory tool, an electronic program guide (EPG), or by other means. The Data Security or the announcement function directs group key management using an application programming interface (API), which is peculiar to the host OS in its specifics. A generic API for group key management is for further study, but this function is necessary to allow Group (KEK) and Data (TPKs) key establishment to be scalable to the particular application. A GCKS application program will use the API to initiate the procedures for establishing SAs on behalf of a Security Protocol in which members join secure groups and receive keys for streams, files, or other data.

The goal of the exchanges is to establish a GSA through updates to the SAD of a key management implementation and particular Security Protocol. The Data Security Protocol ("SECURITY PROTOCOL") of Figure 2 may span internetwork and application layers or operate at the internetwork layer, such as AH and ESP.

4. Registration Protocol

The design of the registration protocol is flexible and can support different application scenarios. The chosen registration protocol solution reflects the specific requirements of specific scenarios. In principle, it is possible to base a registration protocol on any secure-channel protocol, such as IPsec and TLS, which is the case in tunneled GSAKMP [tGSAKMP]. GDOI [RFC3547] reuses IKE Phase 1 as the secure channel to download Rekey and/or Data SAs. Other protocols, such as MIKEY and GSAKMP, use authenticated Diffie-Hellman exchanges similar to IKE Phase 1, but they are specifically tailored for key download to achieve efficient operation. We discuss the design of a registration protocol in detail in the rest of this section.

4.1. Registration Protocol via Piggybacking or Protocol Reuse

Some registration protocols need to tunnel through a data-signaling protocol to take advantage of already existing security functionality, and/or to optimize the total session setup time. For example, a telephone call has strict bounds for delay in setup time. It is not feasible to run security exchanges in parallel with call setup, since the latter often resolves the address. Call setup must complete before the caller knows the callee's address. In this case, it may be advantageous to tunnel the key exchange procedures inside

call establishment [H.235,MIKEY], so that both can complete (or fail, see below) at the same time.

The registration protocol has different requirements depending on the particular integration/tunneling approach. These requirements are not necessarily security requirements, but will have an impact on the chosen security solution. For example, the security association will certainly fail if the call setup fails in the case of IP telephony.

Conversely, the registration protocol imposes requirements on the protocol that tunnels it. In the case of IP telephony, the call setup usually will fail when the security association is not successfully established. In the case of video-on-demand, protocols such as RTSP that convey key management data will fail when a needed security association cannot be established.

Both GDOI and MIKEY use this approach, but in different ways. MIKEY can be tunneled in SIP and RTSP. It takes advantage of the session information contained in these protocols and the possibility to optimize the setup time for the registration procedure. SIP requires that a tunneled protocol must use at most one roundtrip (i.e., two messages). This is also a desirable requirement from RTSP.

The GDOI approach takes advantage of the already defined ISAKMP phase 1 exchange [RFC2409], and extends the phase 2 exchange for the registration. The advantage here is the reuse of a successfully deployed protocol and the code base, where the defined phase 2 exchange is protected by the SA created by phase 1. GDOI also inherits other functionality of the ISAKMP, and thus it is readily suitable for running IPsec protocols over IP multicast services.

4.2. Properties of Alternative Registration Exchange Types

The required design properties of a registration protocol have different trade-offs. A protocol that provides perfect forward secrecy and identity protection trades performance or efficiency for better security, while a protocol that completes in one or two messages may trade security functionality (e.g., identity protection) for efficiency.

Replay protection generally uses either a timestamp or a sequence number. The first requires synchronized clocks, while the latter requires retention of state. In a timestamp-based protocol, a replay cache is needed to store the authenticated messages (or the hashes of the messages) received within the allowable clock skew. The size of the replay cache depends on the number of authenticated messages received during the allowable clock skew. During a DoS attack, the replay cache might become overloaded. One solution is to over-

provision the replay cache, but this may lead to a large replay cache. Another solution is to let the allowable clock skew be changed dynamically during runtime. During a suspected DoS attack, the allowable clock skew is decreased so that the replay cache becomes manageable.

A challenge-response mechanism (using Nonces) obviates the need for synchronized clocks for replay protection when the exchange uses three or more messages [MVV].

Additional security functions become possible as the number of allowable messages in the registration protocol increase. ISAKMP offers identity protection, for example, as part of a six-message exchange. With additional security features, however, comes added complexity: Identity protection, for example, not only requires additional messages, but may result in DoS vulnerabilities since authentication is performed in a late stage of the exchange after resources already have been devoted.

In all cases, there are tradeoffs with the number of message exchanged, the desired security services, and the amount of infrastructure that is needed to support the group key management service. Whereas protocols that use two or even one-message setup have low latency and computation requirements, they may require more infrastructure such as secure time or offer less security such as the absence of identity protection. What tradeoffs are acceptable and what are not is very much dictated by the application and application environment.

4.3. Infrastructure for Alternative Registration Exchange Types

The registration protocol may need external infrastructures to handle authentication and authorization, replay protection, protocol-run integrity, and possibly other security services such as secure synchronized clocks. For example, authentication and authorization may need a PKI deployment (with either authorization-based certificates or a separate management) or may be handled using AAA infrastructure. Replay protection using timestamps requires an external infrastructure or protocol for clock synchronization.

However, external infrastructures may not always be needed; for example pre-shared keys are used for authentication and authorization. This may be the case if the subscription base is relatively small. In a conversational multimedia scenario (e.g., a VoIP call between two or more people), it may be the end user who handles the authorization by manually accepting/rejecting the incoming calls. In that case, infrastructure support may not be required.

4.4. De-registration Exchange

The session-establishment protocol (e.g., SIP, RTSP) that conveys a registration exchange often has a session-disestablishment protocol such as RTSP TEARDOWN [RFC2326] or SIP BYE [RFC3261]. The session-disestablishment exchange between endpoints offers an opportunity to signal the end of the GSA state at the endpoints. This exchange need only be a unidirectional notification by one side that the GSA is to be destroyed. For authentication of this notification, we may use a proof-of-possession of the group key(s) by one side to the other. Some applications benefit from acknowledgement in a mutual, two-message exchange signaling disestablishment of the GSA concomitant with disestablishment of the session, e.g., RTSP or SIP session. In this case, a two-way proof-of-possession might serve for mutual acknowledgement of the GSA disestablishment.

5. Rekey Protocol

The group rekey protocol is for transport of keys and SAs between a GCKS and the members of a secure communications group. The GCKS sends rekey messages to update a Rekey SA, or initialize/update a Data SA or both. Rekey messages are protected by a Rekey SA. The GCKS may update the Rekey SA when group membership changes or when KEKs or TPKs expire. Recall that KEKs correspond to a Rekey SA and TPKs correspond to a Data SA.

The following are some desirable properties of the rekey protocol.

- o The rekey protocol ensures that all members receive the rekey information in a timely manner.
- o The rekey protocol specifies mechanisms allowing the parties to contact the GCKS and re-sync when their keys expire and no updates have been received.
- o The rekey protocol avoids implosion problems and ensures reliability in delivering Rekey information.

We further note that the rekey protocol is primarily responsible for scalability of the group key management architecture. Hence, it is imperative that we provide the above listed properties in a scalable manner. Note that solutions exist in the literature (both IETF standards and research articles) for parts of the problem. For instance, the rekey protocol may use a scalable group key management algorithm (GKMA) to reduce the number of keys sent in a rekey message. Examples of a GKMA include LKH, OFT, Subset difference based schemes etc.

5.1. Goals of the Rekey Protocol

The goals of the rekey protocol are:

- o to synchronize a GSA,
- o to provide privacy and (symmetric or asymmetric) authentication, replay protection and DoS protection,
- o efficient rekeying after changes in group membership or when keys (KEKs) expire,
- o reliable delivery of rekey messages,
- o member recovery from an out-of-sync GSA,
- o high throughput and low latency, and
- o support IP Multicast or multi-unicast.

We identify several major issues in the design of a rekey protocol:

1. rekey message format,
2. reliable transport of rekey messages,
3. implosion,
4. recovery from out-of-sync GSA,
5. incorporating GKMA in rekey messages, and
6. interoperability of GKMA.

Note that interoperation of rekey protocol implementations is insufficient for a GCKS to successfully rekey a group. The GKMA must also interoperate, i.e., standard versions of the group key management algorithms such as LKH, OFT, or Subset Difference must be used.

The rest of this section discusses these topics in detail.

5.2. Rekey Message Transport and Protection

Rekey messages contain Rekey and/or Data SAs along with KEKs and TPKs. These messages need to be confidential, authenticated, and protected against replay and DoS attacks. They are sent via multicast or multi-unicast from the GCKS to the members.

Rekey messages are encrypted with the Group KEK for confidentiality. When used in conjunction with a GKMA, portions of the rekey message are first encrypted with the appropriate KEKs as specified by the GKMA. The GCKS authenticates rekey messages using either a MAC, computed using the group Authentication key, or a digital signature. In both cases, a sequence number is included in computation of the MAC or the signature to protect against replay attacks.

When group authentication is provided with a symmetric key, rekey messages are vulnerable to attacks by other members of the group. Rekey messages are digitally signed when group members do not trust each other. When asymmetric authentication is used, members receiving rekey messages are vulnerable to DoS attacks. An external adversary may send a bogus rekey message, which a member cannot identify until after it performs an expensive digital signature operation. To protect against such an attack, a MAC may be sent as part of the rekey message. Members verify the signature only upon successful verification of the MAC.

Rekey messages contain group key updates corresponding to a single [RFC2627,OFT] or multiple membership changes [SD1,SD2,BatchRekey] and may contain group key initialization messages [OFT].

5.3. Reliable Transport of Rekey Messages

The GCKS must ensure that all members have the current Data Security and Rekey SAs. Otherwise, authorized members may be inadvertently excluded from receiving group communications. Thus, the GCKS needs to use a rekey algorithm that is inherently reliable or employ a reliable transport mechanism to send rekey messages.

There are two dimensions to the problem. Messages that update group keys may be lost in transit or may be missed by a host when it is offline. LKH and OFT group key management algorithms rely on past history of updates being received by the host. If the host goes offline, it will need to resynchronize its group-key state when it comes online; this may require a unicast exchange with the GCKS. The Subset Difference algorithm, however, conveys all the necessary state in its rekey messages and does not need members to be always online or keeping state. The Subset Difference algorithm does not require a back channel and can operate on a broadcast network. If a rekey message is lost in transmission, the Subset Difference algorithm cannot decrypt messages encrypted with the TPK sent via the lost rekey message. There are self-healing GKMA's proposed in the literature that allow a member to recover lost rekey messages, as long as rekey messages before and after the lost rekey message are received.

Rekey messages are typically short (for single membership change as well as for small groups), which makes it easy to design a reliable delivery protocol. On the other hand, the security requirements may add an additional dimension to address. There are some special cases in which membership changes are processed as a batch, reducing the frequency of rekey messages but increasing their size. Furthermore, among all the KEKs sent in a rekey message, as many as half the members need only a single KEK. We may take advantage of these properties in designing a rekey message(s) and a protocol for their reliable delivery.

Three categories of solutions have been proposed:

1. Repeatedly transmit the rekey message. In many cases rekey messages translate to only one or two IP packets.
2. Use an existing reliable multicast protocol/infrastructure.
3. Use FEC for encoding rekey packets (with NACKs as feedback) [BatchRekey].

Note that for small messages, category 3 is essentially the same as category 1.

The group member might be out of synchrony with the GCKS if it receives a rekey message having a sequence number that is more than one greater than the last sequence number processed. This is one means by which the GCKS member detects that it has missed a rekey message. Alternatively, the data-security application, upon detecting that it is using an out-of-date key, may notify the group key management module. The action taken by the GCKS member is a matter of group policy. The GCKS member should log the condition and may contact the GCKS to rerun the re-registration protocol to obtain a fresh group key. The group policy needs to take into account boundary conditions, such as reordered rekey messages when rekeying is so frequent that two messages might get reordered in an IP network. The group key policy also needs to take into account the potential for denial of service attacks where an attacker delays or deletes a rekey message in order to force a subnetwork or subset of the members to simultaneously contact the GCKS.

If a group member becomes out-of-synch with the GSA then it should re-register with the GCKS. However, in many cases there are other, simpler methods for re-synching with the group:

- o The member can open a simple unprotected connection (e.g., TCP) with the GCKS and obtain the current (or several recent) rekey messages. Note that there is no need for authentication or

encryption here, since the rekey message is already signed and is multicast in the clear. One may think that this opens the GCKS to DoS attacks by many bogus such requests. This, however, does not seem to worsen the situation; in fact, bombarding the GCKS with bogus resynch requests would be much more problematic.

- o The GCKS can post the rekey messages on some public site (e.g., a web site) and the out-of-synch member can obtain the rekey messages from that site.

The GCKS may always provide all three ways of resynching (i.e., re-registration, simple TCP, and public posting). This way, the member may choose how to resynch; it also avoids adding yet another field to the policy token [GSPT]. Alternatively, a policy token may contain a field specifying one or more methods supported for resynchronization of a GSA.

5.4. State-of-the-art on Reliable Multicast Infrastructure

The rekey message may be sent using reliable multicast. There are several types of reliable multicast protocols with different properties. However, there are no standards track reliable multicast protocols published at this time, although IETF consensus has been reached on two protocols that are intended to go into the standards track [NORM,RFC3450]. Thus, this document does not recommend a particular reliable multicast protocol or set of protocols for the purpose of reliable group rekeying. The suitability of NAK-based, ACK-based or other reliable multicast methods is determined by the application needs and operational environment. In the future, group key management protocols may choose to use particular standards-based approaches that meet the needs of the particular application. A secure announcement facility may be needed to signal the use of a reliable multicast protocol, which could be specified as part of group policy. The reliable multicast announcement and policy specification, however, can only follow the establishment of reliable multicast standards and are not considered further in this document.

Today, the several MSEC group key management protocols support sequencing of the rekey messages through a sequence number, which is authenticated along with the rekey message. A sender of rekey messages may re-transmit multiple copies of the message provided that they have the same sequence number. Thus, re-sending the message is a rudimentary means of overcoming loss along the network path. A member who receives the rekey message will check the sequence number to detect duplicate and missing rekey messages. The member receiver will discard duplicate messages that it receives. Large rekey messages, such as those that contain LKH or OFT tree structures,

might benefit from transport-layer FEC in the future, when standards-based methods become available. It is unlikely that forward error correction (FEC) methods will benefit short rekey messages that fit within a single message. In this case, FEC degenerates to simple retransmission of the message.

5.5. Implosion

Implosion may occur due to one of two reasons. First, recall that one of the goals of the rekey protocol is to synchronize a GSA. When a rekey or Data SA expires, members may contact the GCKS for an update. If all, or even many, members contact the GCKS at about the same time, the GCKS might not be able to handle all those messages. We refer to this as an out-of-sync implosion.

The second case is in the reliable delivery of rekey messages. Reliable multicast protocols use feedback (NACK or ACK) to determine which packets must be retransmitted. Packet losses may result in many members sending NACKs to the GCKS. We refer to this as feedback implosion.

The implosion problem has been studied extensively in the context of reliable multicasting. The proposed feedback suppression and aggregation solutions might be useful in the GKM context as well. Members may wait a random time before sending an out-of-sync or feedback message. Meanwhile, members might receive the necessary key updates and therefore not send a feedback message. An alternative solution is to have the members contact one of several registration servers when they are out-of-sync. This requires GSA synchronization between the multiple registration servers.

Feedback aggregation and local recovery employed by some reliable multicast protocols are not easily adaptable to transport of rekey messages. Aggregation raises authentication issues. Local recovery is more complex because members need to establish SAs with the local repair server. Any member of the group or a subordinate GCKS may serve as a repair server, which can be responsible for resending rekey messages.

Members may use the group SA, more specifically the Rekey SA, to authenticate requests sent to the repair server. However, replay protection requires maintaining state at members as well as repair servers. Authentication of repair requests is meant to protect against DoS attacks. Note also that an out-of-sync member may use an expired Rekey SA to authenticate repair requests, which requires repair servers to accept messages protected by old SAs.

Alternatively, a simple mechanism may be employed to achieve local repair efficiently. Each member receives a set of local repair server addresses as part of group operation policy information. When a member does not receive a rekey message, it can send a "Retransmit replay message(s) with sequence number n and higher" message to one of the local repair servers. The repair server can either ignore the request if it is busy or retransmit the requested rekey messages as received from the GCKS. The repair server, which is also another member may choose to serve only m requests in a given time period (i.e., rate limits responses) or per a given rekey message. Rate limiting the requests and responses protects the repair servers as well as other members of the group from DoS attacks.

5.6. Incorporating Group Key Management Algorithms

Group key management algorithms make rekeying scalable. Large group rekeying without employing GKMA is prohibitively expensive.

Following are some considerations in selecting a GKMA:

- o Protection against collusion.

Members (or non-members) should not be able to collaborate to deduce keys for which they are not privileged (following the GKMA key distribution rules).

- o Forward access control

The GKMA should ensure that departing members cannot get access to future group data.

- o Backward access control

The GKMA should ensure that joining members cannot decrypt past data.

5.7. Stateless, Stateful, and Self-healing Rekeying Algorithms

We classify group key management algorithms into three categories: stateful, stateless, and self-healing.

Stateful algorithms [RFC2627,OFT] use KEKs from past rekeying instances to encrypt (protect) KEKs corresponding to the current and future rekeying instances. The main disadvantage in these schemes is that if a member were offline or otherwise failed to receive KEKs from a past rekeying instance, it may no longer be able to synchronize its GSA even though it can receive KEKs from all future rekeying instances. The only solution is to contact the GCKS

explicitly for resynchronization. Note that the KEKs for the first rekeying instance are protected by the Registration SA. Recall that communication in that phase is one to one, and therefore it is easy to ensure reliable delivery.

Stateless GKMA [SD1,SD2] encrypt rekey messages with KEKs sent during the registration protocol. Since rekey messages are independent of any past rekey messages (i.e., that are not protected by KEKs therein), a member may go offline but continue to decipher future communications. However, stateless GKMA offer no mechanisms to recover past rekeying messages. Stateless rekeying may be relatively inefficient, particularly for immediate (not batch) rekeying in highly dynamic groups.

In self-healing schemes [Self-Healing], a member can reconstruct a lost rekey message as long as it receives some past and some future rekey messages.

5.8. Interoperability of a GKMA

Most GKMA specifications do not specify packet formats, although many group key management algorithms need format specification for interoperability. There are several alternative ways to manage key trees and to number nodes within key trees. The following information is needed during initialization of a Rekey SA or included with each GKMA packet.

- o GKMA name (e.g., LKH, OFT, Subset Difference)
- o GKMA version number (implementation specific). Version may imply several things such as the degree of a key tree, proprietary enhancements, and qualify another field such as a key ID.
- o Number of keys or largest ID
- o Version-specific data
- o Per-key information:
 - key ID,
 - key lifetime (creation/expiration data) ,
 - encrypted key, and
 - encryption key's ID (optional).

Key IDs may change in some implementations in which case one needs to send:

- o List of <old id, new id> pairs.

6. Group Security Association

The GKM architecture defines the interfaces between the registration, rekey, and data security protocols in terms of the Security Associations (SAs) of those protocols. By isolating these protocols behind a uniform interface, the architecture allows implementations to use protocols best suited to their needs. For example, a rekey protocol for a small group could use multiple unicast transmissions with symmetric authentication, while a rekey protocol for a large group could use IP Multicast with packet-level Forward Error Correction and source authentication.

The group key management architecture provides an interface between the security protocols and the group SA (GSA). The GSA consists of three SAs: Registration SA, Rekey SA, and Data SA. The Rekey SA is optional. There are two cases in defining the relationships between the three SAs. In both cases, the Registration SA protects the registration protocol.

Case 1: Group key management is done WITHOUT using a Rekey SA. The registration protocol initializes and updates one or more Data SAs (having TPKs to protect files or streams). Each Data SA corresponds to a single group, which may have more than one Data SA.

Case 2: Group key management is done WITH a Rekey SA to protect the rekey protocol. The registration protocol initializes the one or more Rekey SAs as well as zero or more Data SAs, upon successful completion. When a Data SA is not initialized in the registration protocol, initialization is done in the rekey protocol. The rekey protocol updates Rekey SA(s) AND establishes Data SA(s).

6.1. Group Policy

Group policy is described in detail in the Group Security Policy Token document [GSPT]. Group policy can be distributed through group announcements, key management protocols, and other out-of-band means (e.g., via a web page). The group key management protocol carries cryptographic policies of the SAs and the keys it establishes, as well as additional policies for the secure operation of the group.

The acceptable cryptographic policies for the registration protocol, which may run over TLS [TLS], IPsec, or IKE, are not conveyed in the group key management protocol since they precede any of the key management exchanges. Thus, a security policy repository having some access protocol may need to be queried prior to establishing the key-management session, to determine the initial cryptographic policies for that establishment. This document assumes the existence of such a repository and protocol for GCKS and member policy queries. Thus group security policy will be represented in a policy repository and accessible using a policy protocol. Policy distribution may be a push or a pull operation.

The group key management architecture assumes that the following group policy information may be externally managed, e.g., by the content owner, group conference administrator or group owner:

- o the identity of the Group owner, the authentication method, and the delegation method for identifying a GCKS for the group;
- o the group GCKS, authentication method, and delegation method for any subordinate GCKSs for the group;
- o the group membership rules or list and authentication method.

There are two additional policy-related requirements external to group key management.

- o There is an authentication and authorization infrastructure such as X.509 [RFC3280], SPKI [RFC2693], or a pre-shared key scheme, in accordance with the group policy for a particular group.
- o There is an announcement mechanism for secure groups and events, which operates according to group policy for a particular group.

Group policy determines how the registration and rekey protocols initialize or update Rekey and Data SAs. The following sections describe potential information sent by the GCKS for the Rekey and Data SAs. A member needs the information specified in the next sections to establish Rekey and Data SAs.

6.2. Contents of the Rekey SA

The Rekey SA protects the rekey protocol. It contains cryptographic policy, Group Identity, and Security Parameter Index (SPI) [RFC2401] to uniquely identify an SA, replay protection information, and key protection keys.

6.2.1. Rekey SA Policy

- o GROUP KEY MANAGEMENT ALGORITHM

This represents the group key revocation algorithm that enforces forward and backward access control. Examples of key revocation algorithms include LKH, LKH+, OFT, OFC, and Subset Difference [RFC2627,OFT,TAXONOMY,SD1,SD2]. If the key revocation algorithm is NULL, the Rekey SA contains only one KEK, which serves as the group KEK. The rekey messages initialize or update Data SAs as usual. However, the Rekey SA itself can be updated (the group KEK can be rekeyed) when members join or the KEK is about to expire. Leave rekeying is done by re-initializing the Rekey SA through the rekey protocol.

- o KEK ENCRYPTION ALGORITHM

This specifies a standard encryption algorithm such as 3DES or AES, and also the KEK KEY LENGTH.

- o AUTHENTICATION ALGORITHM

This algorithm uses digital signatures for GCKS authentication (since all shared secrets are known to some or all members of the group), or some symmetric secret in computing MACs for group authentication. Symmetric authentication provides weaker authentication in that any group member can impersonate a particular source. The AUTHENTICATION KEY LENGTH is also to be specified.

- o CONTROL GROUP ADDRESS

This address is used for multicast transmission of rekey messages. This information is sent over the control channel such as in an ANNOUNCEMENT protocol or call setup message. The degree to which the control group address is protected is a matter of group policy.

- o REKEY SERVER ADDRESS

This address allows the registration server to be a different entity from the server used for rekeying, such as for future invocations of the registration and rekey protocols. If the registration server and the rekey server are two different entities, the registration server sends the rekey server's address as part of the Rekey SA.

6.2.2. Group Identity

The group identity accompanies the SA (payload) information as an identifier if the specific group key management protocol allows multiple groups to be initialized in a single invocation of the registration protocol, or multiple groups to be updated in a single rekey message. It is often simpler to restrict each registration invocation to a single group, but such a restriction is unnecessary. It is always necessary to identify the group when establishing a Rekey SA, either implicitly through an SPI or explicitly as an SA parameter.

6.2.3. KEKs

Corresponding to the key management algorithm, the Rekey SA contains one or more KEKs. The GCKS holds the key encrypting keys of the group, while the members receive keys following the specification of the key management algorithm. When there are multiple KEKs for a group (as in an LKH tree), each KEK needs to be associated with a Key ID, which is used to identify the key needed to decrypt it. Each KEK has a LIFETIME associated with it, after which the KEK expires.

6.2.4. Authentication Key

The GCKS provides a symmetric or public key for authentication of its rekey messages. Symmetric key authentication is appropriate only when all group members can be trusted not to impersonate the GCKS. The architecture does not rule out methods for deriving symmetric authentication keys at the member [RFC2409] rather than pushing them from the GCKS.

6.2.5. Replay Protection

Rekey messages need to be protected from replay/reflection attacks. Sequence numbers are used for this purpose, and the Rekey SA (or protocol) contains this information.

6.2.6. Security Parameter Index (SPI)

The tuple <Group identity, SPI> uniquely identifies a Rekey SA. The SPI changes each time the KEKs change.

6.3. Contents of the Data SA

The GCKS specifies the data security protocol used for secure transmission of data from sender(s) to receiving members. Examples of data security protocols include IPsec ESP [RFC2401] and SRTP [RFC3711]. While the contents of each of these protocols are out of

the scope of this document, we list the information sent by the registration protocol (or the rekey protocol) to initialize or update the Data SA.

6.3.1. Group Identity

The Group identity accompanies SA information when Data SAs are initialized or rekeyed for multiple groups in a single invocation of the registration protocol or in a single Rekey message.

6.3.2. Source Identity

The SA includes source identity information when the group owner chooses to reveal source identity to authorized members only. A public channel such as the announcement protocol is only appropriate when there is no need to protect source or group identities.

6.3.3. Traffic Protection Keys

Regardless of the data security protocol used, the GCKS supplies the TPKs, or information to derive TPKs for traffic protection.

6.3.4. Data Authentication Keys

Depending on the data authentication method used by the data security protocol, group key management may pass one or more keys, functions (e.g., TESLA [TESLA-INFO, TESLA-SPEC]), or other parameters used for authenticating streams or files.

6.3.5. Sequence Numbers

The GCKS passes sequence numbers when needed by the data security protocol, for SA synchronization and replay protection.

6.3.6. Security Parameter Index (SPI)

The GCKS may provide an identifier as part of the Data SA contents for data security protocols that use an SPI or similar mechanism to identify an SA or keys within an SA.

6.3.7. Data SA policy

The Data SA parameters are specific to the data security protocol but generally include encryption algorithm and parameters, the source authentication algorithm and parameters, the group authentication algorithm and parameters, and/or replay protection information.

7. Scalability Considerations

The area of group communications is quite diverse. In teleconferencing, a multipoint control unit (MCU) may be used to aggregate a number of teleconferencing members into a single session; MCUs may be hierarchically organized as well. A loosely coupled teleconferencing session [RFC3550] has no central controller but is fully distributed and end-to-end. Teleconferencing sessions tend to have at most dozens of participants. However, video broadcast that uses multicast communications and media-on-demand that uses unicast are large-scale groups numbering hundreds to millions of participants.

As described in the Requirements section, Section 2, the group key management architecture supports multicast applications with a single sender. The architecture described in this paper supports large-scale operation through the following features.

1. There is no need for a unicast exchange to provide data keys to a security protocol for members who have previously registered in the particular group; data keys can be pushed in the rekey protocol.
2. The registration and rekey protocols are separable to allow flexibility in how members receive group secrets. A group may use a smart-card based system in place of the registration protocol, for example, to allow the rekey protocol to be used with no back channel for broadcast applications such as television conditional access systems.
3. The registration and rekey protocols support new keys, algorithms, authentication mechanisms and authorization infrastructures in the architecture. When the authorization infrastructure supports delegation, as in X.509 and SPKI, the GCKS function can be distributed as shown in Figure 3 below.

The first feature in the list allows fast keying of data security protocols when the member already belongs to the group. While this is realistic for subscriber groups and customers of service providers who offer content events, it may be too restrictive for applications that allow member enrollment at the time of the event. The MSEC group key management architecture suggests hierarchically organized key distribution to handle potential mass simultaneous registration requests. The Figure 3 configuration may be needed when conventional clustering and load balancing solutions of a central GCKS site cannot meet customer requirements. Unlike conventional caching and content

distribution networks, however, the configuration shown in Figure 3 has additional security ramifications for physical security of a GCKS.

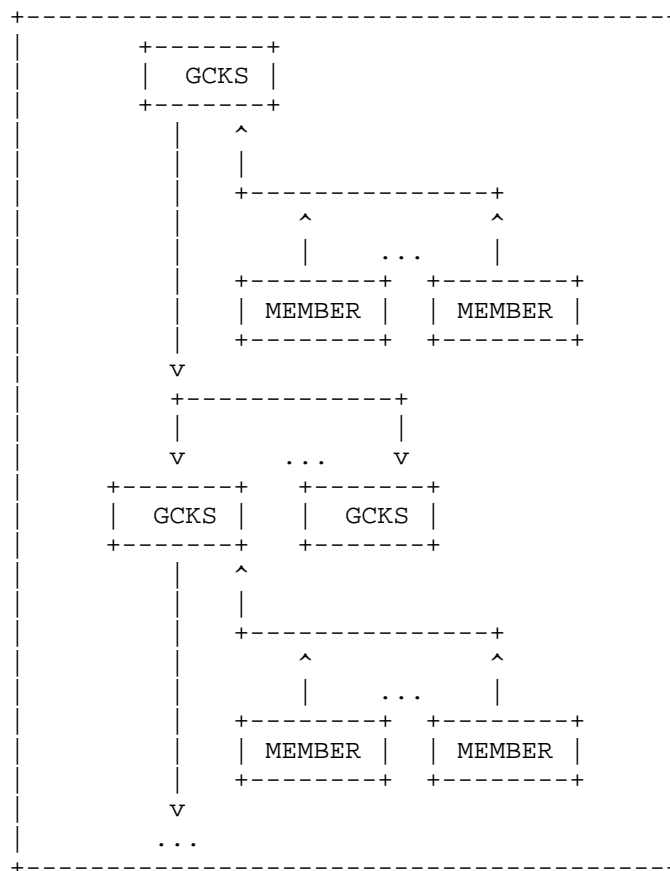


Figure 3: Hierarchically Organized Key Distribution

More analysis and work is needed on the protocol instantiations of the group key management architecture, to determine how effectively and securely the architecture can support large-scale multicast applications. In addition to being as secure as pairwise key management against man-in-the-middle, replay, and reflection attacks, group key management protocols have additional security needs. Unlike pairwise key management, group key management needs to be secure against attacks by group members who attempt to impersonate a GCKS or disrupt the operation of a GCKS, as well as by non-members.

Thus, secure groups need to converge to a common group key when members are attacking the group, joining and leaving the group, or being evicted from the group. Group key management protocols also need to be robust when DoS attacks or network partition leads to large numbers of synchronized requests. An instantiation of group key management, therefore, needs to consider how GCKS operation might be distributed across multiple GCKSs designated by the group owner to serve keys on behalf of a designated GCKS. GSAKMP [GSAKMP] protocol uses the policy token and allows designating some of the members as subordinate GCKSs to address this scalability issue.

8. Security Considerations

This memo describes MSEC key management architecture. This architecture will be instantiated in one or more group key management protocols, which must be protected against man-in-the-middle, connection hijacking, replay, or reflection of past messages, and denial of service attacks.

Authenticated key exchange [STS,SKEME,RFC2408,RFC2412,RFC2409] techniques limit the effects of man-in-the-middle and connection hijacking attacks. Sequence numbers and low-computation message authentication techniques can be effective against replay and reflection attacks. Cookies [RFC2522], when properly implemented, provide an efficient means to reduce the effects of denial of service attacks.

This memo does not address attacks against key management or security protocol implementations such as so-called type attacks that aim to disrupt an implementation by such means as buffer overflow. The focus of this memo is on securing the protocol, not on implementing the protocol.

While classical techniques of authenticated key exchange can be applied to group key management, new problems arise with the sharing of secrets among a group of members: group secrets may be disclosed by a member of the group, and group senders may be impersonated by other members of the group. Key management messages from the GCKS should not be authenticated using shared symmetric secrets unless all members of the group can be trusted not to impersonate the GCKS or each other. Similarly, members who disclose group secrets undermine the security of the entire group. Group owners and GCKS administrators must be aware of these inherent limitations of group key management.

Another limitation of group key management is policy complexity. While peer-to-peer security policy is an intersection of the policy of the individual peers, a group owner sets group security policy

externally in secure groups. This document assumes there is no negotiation of cryptographic or other security parameters in group key management. Group security policy, therefore, poses new risks to members who send and receive data from secure groups. Security administrators, GCKS operators, and users need to determine minimal acceptable levels of security (e.g., authentication and admission policy of the group, key lengths, cryptographic algorithms and protocols used) when joining secure groups.

Given the limitations and risks of group security, the security of the group key management registration protocol should be as good as the base protocols on which it is developed, such as IKE, IPsec, TLS, or SSL. The particular instantiations of this group key management architecture must ensure that the high standards for authenticated key exchange are preserved in their protocol specifications, which will be Internet standards-track documents that are subject to review, analysis, and testing.

The second protocol, the group key management rekey protocol, is new and has unknown risks. The source-authentication risks described above are obviated by the use of public-key cryptography. The use of multicast delivery may raise additional security issues such as reliability, implosion, and denial-of-service attacks based upon the use of multicast. The rekey protocol specification needs to offer secure solutions to these problems. Each instantiation of the rekey protocol, such as the GSAKMP Rekey or the GDOI Groupkey-push operations, need to validate the security of their rekey specifications.

Novelty and complexity are the biggest risks to group key management protocols. Much more analysis and experience are needed to ensure that the architecture described in this document can provide a well-articulated standard for security and risks of group key management.

9. Acknowledgments

The GKM Building Block [GKMBB] I-D by SMuG was a precursor to this document; thanks to Thomas Hardjono and Hugh Harney for their efforts. During the course of preparing this document, Andrea Colegrove, Brian Weis, George Gross, and several others in the MSEC WG and GSEC and SMuG research groups provided valuable comments that helped improve this document. The authors appreciate their contributions to this document.

10. Informative References

- [BatchRekey] Yang, Y. R., et al., "Reliable Group Rekeying: Design and Performance Analysis", Proc. ACM SIGCOMM, San Diego, CA, August 2001.
- [CLIQUES] Steiner, M., Tsudik, G., and M. Waidner, "CLIQUES: A New Approach to Group Key Agreement", IEEE ICDCS 97, May 1997
- [FN93] Fiat, A. and M. Naor, "Broadcast Encryption, Advances in Cryptology", CRYPTO 93 Proceedings, Lecture Notes in Computer Science, Vol. 773, pp. 480-491, 1994.
- [GKMBB] Harney, H., M. Baugher, and T. Hardjono, "GKM Building Block: Group Security Association (GSA) Definition," Work in Progress, September 2000.
- [GSAKMP] Harney, H., Colegrove, A., Harder, E., Meth, U., and R. Fleischer, "Group Secure Association Key Management Protocol", Work in Progress, February 2003.
- [GSPT] Hardjono, T., Harney, H., McDaniel, P., Colegrove, A., and P. Dinsmore, "The MSEC Group Security Policy Token", Work in Progress, August 2003.
- [H.235] International Telecommunications Union, "Security and Encryption for H-Series (H.323 and other H.245-based) Multimedia Terminals", ITU-T Recommendation H.235 Version 3, Work in progress, 2001.
- [JKKV94] Just, M., Kranakis, E., Krizanc, D., and P. van Oorschot, "On Key Distribution via True Broadcasting", Proc. 2nd ACM Conference on Computer and Communications Security, pp. 81-88, November 1994.
- [MARKS] Briscoe, B., "MARKS: Zero Side Effect Multicast Key Management Using Arbitrarily Revealed Key Sequences", Proc. First International Workshop on Networked Group Communication (NGC), Pisa, Italy, November 1999.
- [MIKEY] Arkko, J., Carrara, E., Lindholm, F., Naslund, M., and K. Norrman, "MIKEY: Multimedia Internet KEYing", RFC 3830, August 2004.

- [MSEC-Arch] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", RFC 3740, March 2004.
- [MVV] Menzes, A.J., van Oorschot, P.C., and S.A. Vanstone, "Handbook of Applied Cryptography", CRC Press, 1996.
- [NORM] Adamon, B., Bormann, C., Handley, M., and J. Macker, "Negative-acknowledgment (NACK)-Oriented Reliable Multicast (NORM) Protocol", RFC 3940, November 2004.
- [OFT] Balenson, D., McGrew, P.C., and A. Sherman, "Key Management for Large Dynamic Groups: One-Way Function Trees and Amortized Initialization", IRTF Work in Progress, August 2000.
- [RFC2093] Harney, H. and C. Muckenhirn, "Group Key Management Protocol (GKMP) Specification", RFC 2093, July 1997.
- [RFC2094] Harney, H., and C. Muckenhirn, "Group Key Management Protocol (GKMP) Architecture" RFC 2094, July 1997.
- [RFC2326] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", RFC 2326, April 1998.
- [RFC2327] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.
- [RFC2367] McDonald, D., Metz, C., and B. Phan, "PF_KEY Key Management API, Version 2", RFC 2367, July 1998.
- [RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [RFC2408] Maughan, D., Schertler, M., Schneider, M., and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", RFC 2408, November 1998.
- [RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.
- [RFC2412] Orman, H., "The OAKLEY Key Determination Protocol", RFC 2412, November 1998.
- [RFC2522] Karn, P. and W. Simpson, "Photuris: Session-Key Management Protocol", RFC 2522, March 1999.

- [RFC2693] Ellison, C., Frantz, B., Lampson, B., Rivest, R., Thomas, B., and T. Ylonen, "SPKI Certificate Theory", RFC 2693, September 1999.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3280] Housley, R., Polk, W., Ford, W., and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3280, April 2002.
- [RFC2627] Wallner, D., Harder, E., and R. Agee, "Key Management for Multicast: Issues and Architectures", RFC 2627, June 1999.
- [RFC3450] Luby, M., Gemmell, J., Vicisano, L., Rizzo, L., and J. Crowcroft, "Asynchronous Layered Coding (ALC) Protocol Instantiation", RFC 3450, December 2002.
- [RFC3547] Baugher, M., Weis, B., Hardjono, T., and H. Harney, "The Group Domain of Interpretation", RFC 3547, July 2003.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [SD1] Naor, D., Naor, M., and J. Lotspiech, "Revocation and Tracing Schemes for Stateless Receiver", Advances in Cryptology - CRYPTO, Santa Barbara, CA: Springer-Verlag Inc., LNCS 2139, August 2001.
- [SD2] Naor, M. and B. Pinkas, "Efficient Trace and Revoke Schemes", Proceedings of Financial Cryptography 2000, Anguilla, British West Indies, February 2000.
- [Self-Healing] Staddon, J., et. al., "Self-healing Key Distribution with Revocation", Proc. 2002 IEEE Symposium on Security and Privacy, Oakland, CA, May 2002.

- [SKEME] H. Krawczyk, "SKEME: A Versatile Secure Key Exchange Mechanism for Internet", ISOC Secure Networks and Distributed Systems Symposium, San Diego, 1996.
- [STS] Diffie, P. van Oorschot, M., and J. Wiener, "Authentication and Authenticated Key Exchanges", Designs, Codes and Cryptography, 2, 107-125 (1992), Kluwer Academic Publishers.
- [TAXONOMY] Canetti, R., et. al., "Multicast Security: A Taxonomy and some Efficient Constructions", IEEE INFOCOM, 1999.
- [TESLA-INFO] Perrig, A., Canetti, R., Song, D., Tygar, D., and B. Briscoe, "TESLA: Multicast Source Authentication Transform Introduction", Work in Progress, December 2004.
- [TESLA-SPEC] Perrig, A., R. Canetti, and Whillock, "TESLA: Multicast Source Authentication Transform Specification", Work in Progress, April 2002.
- [tGSAKMP] Harney, H., et. al., "Tunneled Group Secure Association Key Management Protocol", Work in Progress, May 2003.
- [TLS] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0," RFC 2246, January 1999.
- [TPM] Marks, D. and B. Turnbull, "Technical protection measures: The Intersection of Technology, Law, and Commercial Licenses", Workshop on Implementation Issues of the WIPO Copyright Treaty (WCT) and the WIPO Performances and Phonograms Treaty (WPPT), World Intellectual Property Organization, Geneva, December 6 and 7, 1999.
- [Wool] Wool, A., "Key Management for Encrypted broadcast", 5th ACM Conference on Computer and Communications Security, San Francisco, CA, Nov. 1998.

Authors' Addresses

Mark Baugher
Cisco Systems
5510 SW Orchid St.
Portland, OR 97219, USA

Phone: +1 408-853-4418
EMail: mbaugher@cisco.com

Ran Canetti
IBM Research
30 Saw Mill River Road
Hawthorne, NY 10532, USA

Phone: +1 914-784-7076
EMail: canetti@watson.ibm.com

Lakshminath R. Dondeti
Qualcomm
5775 Morehouse Drive
San Diego, CA 92121

Phone: +1 858 845 1267
EMail: ldondeti@qualcomm.com

Fredrik Lindholm
Ericsson Research
SE-16480 Stockholm, Sweden

Phone: +46 8 58531705
EMail: fredrik.lindholm@ericsson.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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.

Acknowledgement

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