Stardust Forums, Inc.
1901 S. Bascom Ave, #333
Campbell, CA 95008 USA
Tel: 408-879-8080
Fax: 408-879-8081
Web: www.ipmulticast.com

IP Multicast Initiative (IPMI)

IP Multicast Security

From the Stardust Forums State-of-the-Art Series

An overview of the strategies and technologies

Thomas Hardjono
Bay Architecture Lab, Nortel Networks

Scope of this document

This report focuses on the existing and ongoing developments in the area of IP Multicast security from the perspective of practical network engineering solutions.

Rather than providing a comprehensive survey of the relevant works in the broad area of group-oriented security, the current report focuses on identifying and discussing the concepts and issues underlying IP multicast security. Various pointers for readers to follow up upon are sprinkled throughout the report, and in the final section of the report. Advanced issues, such as member anonymity, conference-key computation, fast stream-authentication methods and other application-layer services are not discussed in detail, except in the context of illustration.

To provide a structure for the report, a division of the general problem areas in IP Multicast security is introduced, after we discuss some factors affecting IP Multicast security. The remainder of the report will follow the organization based on the problem areas.

The report acknowledges from the start that many of the issues relating to IP Multicast security are still open problems. To this end, the Secure Multicast Group (SMuG) of the Internet Research Task Force (IRTF) is investigating these problems, with the aim of identifying and developing standard solutions based on a common set of "building blocks" that promote interoperability. You can obtain more information about SMuG at http://www.ipmulticast.com/community/smug.

This report assumes you are familiar with the basic concepts of public-key (asymmetric) cryptography and private-key (symmetric) cryptography, and with the concepts underlying IP Multicast, such as group membership and IP Multicast routing.

Introduction: The IP Multicast model and security

Security is an important concern for today's information age, and more so in today's increasingly internetworked community of people. The three common areas of concern in data security are data *confidentiality* (secrecy), *authenticity* and *integrity*. Confidentiality refers to the desire to have data sent from a sender to a receiver to be available only to the intended receiver. Authenticity refers to the need for the receiver to be assured that the data truly came from the alleged sender. Integrity is concerned with ensuring that data from a sender to a receiver remains intact (unmodified) during its transit. All three desired features are typically achieved today using cryptographic techniques, where confidentiality is achieved using encryption, authenticity with digital signatures and integrity (traditionally) with Message Authentication Codes (MAC) or similar codes. The close relationship between authenticity and integrity often means that a single technique may be used to achieve both requirements simultaneously. For example, certain uses of digital signatures allow a receiver to verify the authenticity of the message and the fact that it has been received intact/unmodified. See [MOV97] for further details on the history and current schemes for confidentiality, authenticity and integrity.

IP Multicast allows packet distribution to many receivers through a *multicast distribution tree*, in which multicast data can be transmitted to the group members (hosts) at the leaves of the distribution tree [D89]. The multicast distribution tree is shaped using a multicast routing protocol (such as [WPD88, B97, M94, and EFH97]). Any host can join a multicast group by using a group membership protocol (such as IGMP [F97, CDT99]) which directs their subnet router to join or be grafted onto the multicast distribution tree.

The anonymous-receiver model underlying IP Multicast is attractive precisely because the distribution tree is easily extendible, subject to the resources available to the multicast routing protocol. Any host in a subnet can join a multicast group without its subnet router passing identification information about the host to other routers upstream in the distribution tree. This allows IP Multicast to scale to a large number of participating hosts.

The extendibility of the distribution tree in IP Multicast makes the IP Multicast model very attractive from the perspective of scalability. However, from the perspective of security, additional mechanisms and services must be built atop the basic IP Multicast model. This decoupling of security from the IP Multicast model is advantageous, since it allows differing security models and architectures to be deployed, without affecting the multicast distribution tree which delivers the multicast data end-to-end. This decoupling is also important from the application's perspective, since each application requires different forms of host information and other security parameters, and may deploy differing user-identification and user-authentication mechanisms.

Factors in securing IP Multicast

There are several interrelated factors or aspects of IP Multicast that influence the approaches and mechanisms used to secure it. Of these, some broad and most relevant factors include:

- the multicast application type
- group dynamics
- scalability issues
- the underlying trust model

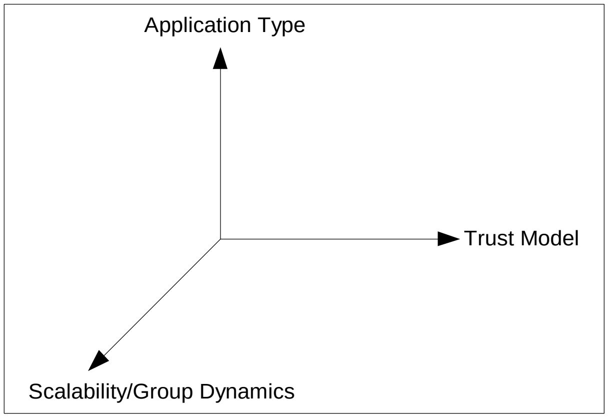

Figure 1: A view of the problem/solution space

Since these factors and others are interrelated, it is difficult to portray their specific relationships and influences. However, Figure 1 displays one view of the problem/solution space, which is made up of these factors (group dynamics falling under scalability). Now we briefly discuss each of these factors.

Multicast application type

IP Multicast commonly views multicast groups as being either one-to-many or many-to-many. This also corresponds to the type of communications occurring among the group members. Together with the value of the data being transmitted, there exists a spectrum of applications of IP Multicast that need to be secured.

As an example, at one end of the spectrum, a subscription service may take the form of a one-to-many multicast from a single source to multiple receivers. Here, the data being delivered may be publicly available (e.g., stock market information). Thus, for this subscriber application, source authentication of the data is more important than confidentiality.

A second example is the pay-per-view (PPV) service where a group of receivers pay a subscription fee for the program being delivered, analogous to the Pay TV scheme. Although the data itself is not confidential, it carries some value, in that the content producer would like to limit access to only the paying subscribers. In this example, encryption of the data may be used to achieve access control, while source authentication may not be as important.

At the other end of the spectrum are the cases that require both confidentiality and source authentication. An example would be a conference call that is implemented over a many-to-many multicast. Here each party in the conference would like to know and be assured of the identity of the source of all transmissions in the conference. Since conferencing events are typically limited in membership and are confidential in nature, encryption must be used to achieve the required confidentiality, while methods for source authentication (such as digital signatures) must also be employed.

Another aspect related to the multicast application type is the frequency and rate of data transmission. This aspect is closely related to the performance of the cryptographic algorithms. Thus, for example, continuous streaming-video may be afforded a different level of security from the infrequent multicast-based delivery of software-update packages, due to the intensive computational requirements of cryptographic operations on streaming-video.

The various applications that may deploy IP Multicast are reviewed in [QA99]. Please refer to that work for further examples of application types.

Group dynamics

Another important factor affecting the security of IP Multicast is the size and behavior of the group. A multicast group may range in size from a few members to tens of thousands of members. The differing sizes affect the mechanisms used to effect security and they also affect the scalability of such mechanisms. Security is also influenced by the behavior of the group in terms of the frequency of members joining or leaving and the average size of the membership change. This, in turn, is related to the application type.

Thus, for example, a multicast group for a PPV service for 100 users may have a different requirements and demands compared to that with 10,000 users. Furthermore, the population distribution of the users and the density of users in certain parts of the Internet may affect both the multicast routing protocol being deployed and the security mechanisms used for the multicast group.

An issue related to scalability is the frequency and number of members that join and leave the multicast group. If encryption is used to protect the value-carrying data transmitted in the multicast instance, then the pattern of group-membership changes will have an impact on the key management for the group.

Scalability issues

In the context of multicast security, scalability refers particularly to the ability of the mechanisms implementing the security features to be extended to cover a larger group of members over a wide physical region without too much deterioration in the level of service and performance of the system as a whole. In general, scalability affects almost all facets of networking. However, in the context of security for IP Multicast, scalability pertains more specifically to the delivery and management of the cryptographic keys, and the propagation and management of security-related policies.

Trust model

When cryptography is employed to provide protection for data, the issue of trust comes to the foreground. The problem concerns the entities that generate, distribute and manage the cryptographic keys and security policies. At the heart of the problem is the need for a *model of trust* underlying the IP Multicast security scheme. This model must address the issues of which entities are to be accorded trust to carry out these functions, the level of trust accorded to them, the source of authority, and other related issues.

General problem areas in IP Multicast security

In order to understand better the various problems surrounding IP Multicast security, we will divide these problems into two categories, which we refer to as the *core* problem areas and the *infrastructure* problem areas. The core problem areas represents issues of pressing concern, where solutions are needed in order to solve the broader infrastructure problems.

The core problem areas covers the issues of:

- IP Multicast group key management
- Methods for IP Multicast data confidentiality and authentication
- IP Multicast security policies

The solutions that are created for the core problem areas will also be applicable and useful for the infrastructure problem areas. The two foremost infrastructure problem areas are:

- Security of IP Multicast routing protocols
- Security of Reliable Multicast (RM) protocols

These five problem areas will be used as the organization of the remainder of this report and will be discussed in the following five sections.

Multicast group key management

Since data related to an IP Multicast group traverses the public Internet and is therefore subject to tapping or copying by non-members of the group, encryption is the method commonly used to provide access control to the data. In the simplest case, shared-key (symmetric) cryptography is used by the sender/source and the receivers, where the data is encrypted by the sender and decrypted by the receivers. This shared key is commonly referred to as the *group-key*, since only members of the multicast group are in possession of the key.

The use of cryptography necessitates the delivery or dissemination of keys, which in this is case is the group-key. Thus, an additional facet to the general problem of multicast security is the method of distributing keys to the appropriate entities involved in a multicast instance and the management of the keys of over given period of time. A *Group-Key Management* (GKM) protocol must not only issue a group-key for a new multicast group, but also update (re-key) the existing group-key under certain conditions and following the prescribed policies, be those general security policies or multicast-specific policies.

GKM requirements

There are a number of requirements that a group-key management protocol must satisfy [CP98, HCD99]. Specifically:

Scalability: Group-key management must be scalable to the scope of the population being catered for in the multicast group, its varying population densities and behaviors, and its wide geographic distribution. The notion of scalability in this context means that events related to group-key management that involve the members of the group should be efficient in resource usage, easily accessible and do not impose delays and other restriction that may affect the usage of the multicast data by its recipients.

Independence: Group-key management must be independent from both unicast and multicast routing [HCD99]. Protocols that implement group-key management must be usable over the various routing protocols available today (and in the future) which may run in different parts of the Internet.

Reliability: The delivery of a cryptographic key must be a reliable event, meaning that there should be no doubt as to the status of the delivered key to a recipient (group member). Members of a group must be able to rely on the group-key management protocol(s) to deliver the group-key to them in a timely fashion.

Security: Group-key management must be carried out in a secure fashion, where the relevant keys are delivered through a secure channel established to the group members. Such a secure delivery method must be resistant against the various possible attacks launched by non-member attackers (and possibly by members of the group itself). Other supporting keys, or *key-management keys* (km-keys) may be deployed to create a safe passage for the important keys used for the multicast data.

Key updates

The updates (re-keys) of the group-key used within a multicast group are affected by the policies governing the multicast transmission, the periodic key-refresh duration, population distribution and dynamics, and other factors.

For instance, a multicast group may be governed by the *forward-secrecy* re-key policy and the *backward-secrecy* re-key policy (or both). The forward-secrecy re-key policy may specify that whenever a member of a group leaves the group, the ex-member must be prevented from having further access to the data in that multicast group. The backward secrecy re-key policy may specify that data transmission to a multicast group previous to the event of a new member joining the group must be unavailable to that new member, even if that new member had been intercepting and storing the data transmissions. In both cases, the method to achieve the aims of these policies is to perform a re-key of the group-key. In effect, the re-keying is triggered by changes in the membership of the group.

In general, changes to the group membership can result from new members joining, existing members voluntarily leaving or existing members being revoked (ejected) from the group.

Several factors may influence this approach to re-keying. These include the costs in terms of the computation cycles and the number of exchanged messages, the frequency of membership changes, the population sizes, the existence (or non-existence) of default periodic refreshes, the value of the data and others. For schemes implementing periodic group-key refreshes to protect against cryptanalysis by an attacker, benefits can be gained by aligning the re-keying, even due to membership changes, to that of the periodic re-keying.

Scalability, domains and key-management keys

Scalability represents an important concern in IP Multicast routing protocols, and often, separate routing domains are delineated in order to ease network management.

The concept of domains is also applicable to group-key management to effect scalability, where members are divided (logically or physically) into domains or subgroups. At least two general types of domains are possible for group key management:

- *Domains according to data encryption*: Here, the domains demarcate regions within which differing group-keys are used to encrypt the multicast data. Thus, each domain is associated with a unique group-key, and "crypto-translations" (decryption using one key, followed by encryption using another key) must be carried out at the domain boundaries. Group-members residing within each domain would be in possession of a unique group-key (per domain). The work of [M97] illustrates this approach. In effect, each domain can be treated independently since each would be associated with a different key.

- *Domains according to key management*: Here, the domains demarcate key management regions, where each region is associated with a different set of key-management keys (*km-keys*) for the express purpose of disseminating the common group-key. Thus, each domain would manage its own km-keys (e.g., different re-key period for km-keys), even though these are used to create safe passage for the common (group-wide) group-key from a key-source (e.g., key server) to each of the receivers residing in differing key management domains.

Combinations of both types of domains can also be deployed, while other interpretations of domains can also be applied.

Architectures for group key management

There are a number of arrangements of keys that are possible for groups of participants in a multicast instance. The work of [WHA98] summarizes three useful architectures for group key management, viewed from the perspective of the logical key arrangements and key-relationships. In order to explain these architectures, we've shown a useful basic physical model in Figure 2. The model consists of a Core (or Root) *Key Management Entity* (KME) serving either the group members directly (centralized key management), or serving other Key Management Entities (distributed key management). Although perhaps more complex, the distributed key management approach lends itself more easily to scalability since KME-to-KME secure communications can be established to deliver keys to furthest KMEs. The three architectures are discussed below.

Pair-wise key arrangement

In this approach the Root KME shares, in a pair-wise manner, a unique key with each (valid) member of the multicast group. The pair-wise secure channel created between the Root KME and each member is then used to deliver a group-key. The Root KME carries out this exchange for each member of the group. Although this approach allows the Root KME to be the single point of trust for each member of the group, the approach is cumbersome and may not scale to large numbers of members.

A possible variation of this approach would be to delegate the exchange process to a number of selected "Subroot KMEs", thereby pushing the computational task to selected group members. This variation, however, requires these selected group members to be trusted. In addition, there is the increased complexity in the task of removing members who are subroots.

Note that the need for a Root KME to share a pair-wise unique key with each group member is crucial in any case, since such a key is the basis for the creation of a secure channel between the two. This in turn represents the point of departure for other more complex solutions, all of which require some form of "boot-strapping" to start the group key management protocol.

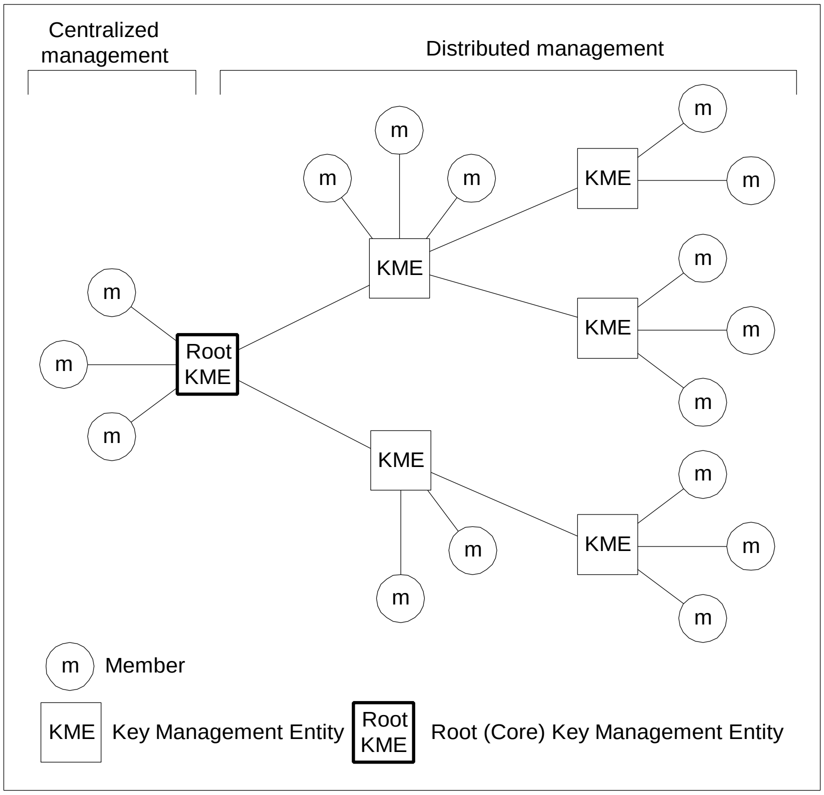
Centralized management
Distributed management
m
Root KME
KME
m Member
KME Key Management Entity
Root KME Root (Core) Key Management Entity

Figure 2: GKM basic model

Complementary keys approach

The basic idea here is to deliver a set of "complementary variables," in addition to the group-key, to the members of the multicast group. Each member is associated with a variable by the Root KME. However, a member's variable is never actually given to it. Instead, a member receives the variable of all the other members (except its own variable). This allows the exclusion of any member in the key-generation process.

When a member leaves the group and a new group-key has to be recomputed, the Root KME will instruct the remaining members to compute the new group-key based on all the variables except the variable of the leaving-member. Assuming all the remaining members obey this instruction, the effect is that the leaving-member is excluded from the key-generation process for the new group-key. The leaving member will not be able to compute the new group-key since it never has possession of the variable associated to it.

This approach is attractive and is reminiscent of secret-sharing schemes [S79, S92]. However, for correct execution, it assumes that collusion will not occur among members of the group. Furthermore, the cryptographic schemes underlying any such complementary-variable approach must be resistant to various attacks to prevent non-members and ex-members from deriving the current group-key using other means.

Hierarchical tree approach

One of the desirable features of group key management protocols is the localization (as much as possible) of the effects of a re-keying event. In other words, a re-keying of one (or a few) members of the group should not affect the other group members too much. To this end the logical division of group members into *subgroups*, arranged in the form of a *logical tree*, represents a promising avenue towards scalable solutions.

The aim of the hierarchical tree approach is for each (logical) subgroup to be assigned a unique *subgroup-key* for the purpose of delivering and updating the global group-key. These subgroup-keys, which are known also to the Root KME, allow the Root KME to address subsets of the groups, by enciphering the messages for a given subgroup using the corresponding subgroup-key. The resulting ciphertext can then be sent via unicast to individual subgroup members, via the multicast group proper (to the entire group), or via a separate subgroup multicast. In any case, only the holders of the corresponding subgroup-key will be able to decipher the ciphertext. This is basis for the solution proposed in [HCM98] based on [WGL98]. Improvements of this approach have been suggested (e.g., [CEK99]).

When implemented in a centralized fashion using a single Root KME, the hierarchical tree approach inherently presents more scalability than the other two approaches since subgroups can be tailored to be of varying sizes following to the population density and membership dynamics. However, for multicast groups with a sparse population spread across wide geographic expanses and for domain autonomy requirements, a distributed KME solution may be preferable to implement (logical) hierarchical tree.

Methods for IP Multicast data confidentiality and authentication

Another core problem area in IP Multicast security concerns to the methods used to ascertain the authenticity (including integrity) of a piece of data and the methods used to establish data confidentiality (secrecy), specifically in the context of voluminous data such as within streaming-video applications.

Where security-related mechanisms are applied to the multicast data, confidentiality and authentication/integrity are typically treated together. That is, since data in a multicast group typically travels end-to-end from the sender to the receiver(s), the cryptographic operations carried out on the data are also typically conducted at the end-points. However, in the context of the multicast applications types, it is also useful to treat the issue of data confidentiality as separate from data authentication. This is because different applications have different requirements. Thus, for example, the publicly-available stock-market data being delivered through a multicast group requires source authentication more than it needs confidentiality. On the other hand, a subscriber-based application (e.g., pay-per-view) requires both source authentication and confidentiality. This reasoning is also useful since confidentiality and authentication/integrity may use differing cryptographic schemes and technologies.

Since IP Multicast traffic, like unicast traffic, traverses the so called "public" Internet, parties that wish to deliver value-carrying data using IP Multicast must deploy mechanisms to control access to the data. One method commonly used to implement controlled access is data encryption. The notion here is that data would be cryptographically enciphered at the source (sender) and the decipherment keys would be available to the intended recipients of the IP Multicast data (namely the multicast group members). Thus, although the IP Multicast traffic (like other traffic) over the Internet can be intercepted by any party, that data would be useless without the decryption key. The same notion also applies for authentication, where only the holders of the authentication key can authenticate the data sent to the multicast group.

Another level of useful distinction is one between source authentication and group authentication. Source authentication is typically achieved using public key (asymmetric) cryptography, where a sender is in possession of a private-key and a public-key. The public-key of the sender is available to the public, through a Certification Authority (CA) that vouches for the relationship between a key and its owner. When a piece of data is digitally-signed using public key cryptography, the digital-signature is verifiable by anyone using the public-key of the signer. Since public key cryptosystems have the property that a for a given digital-signature verifiable using a public-key, only the matching private-key could have been used to generate that digital signature, it follows that only the sender (holding the private-key) could have generated and sent the digitally-signed message. That is, public key cryptography features unique sender (or source) authentication.

This is in contrast to group authentication, which derives from the situation where a group of users share a common shared-key (symmetric key) which is used to "digitally-sign" data by way of generating *a Message Authentication Code* (MAC) using the shared-key and a keyed-hash function [MOV97]. Thus, although only members of the group (holding a copy of the shared-key) can verify that the MAC corresponds to the piece of data to which is attached, they cannot ascertain who among the group members actually generated the MAC. Thus, only group authentication is achieved, namely that the data authentically came from one of the group members.

The relevance of source authentication and group authentication becomes apparent when the performance of public-key (asymmetric) algorithms and shared-key (symmetric) algorithms are taken into consideration, particularly in the context of the high rate of transmission of certain multicast applications. Typically, in software implementations, public-key algorithms are several magnitudes slower than shared-key algorithms. Thus, the choice between source authentication and group authentication must be weighed against the application type, the computing resources available to the group members and the value of the data being delivered through IP Multicast.

The issue of fast source authentication techniques and algorithms for IP Multicast remains an open problem. Intermediate or hybrid solutions, such as the digital signing of hashes of several data packets and stream-encryption methods, will remain attractive until such algorithms become standard and are adopted on a wider scale.

IP Multicast security policies

Similar to other aspects of networking, the correct definition, implementation and maintenance of policies governing the various mechanisms of IP Multicast security are crucial factors. Those which are directly related to IP Multicast security include the policies for key dissemination, for access control, for the re-keying of group-shared keys, and for the actions taken when certain keys are compromised [HH99a, HH99b].

Other policies may be in place to support the mechanisms used to secure the multicast group. Thus, for example, if a member of a group creates an initial secure channel between itself and a key manager (or key server) using IPSec technology (e.g., IKE [HC98]), then policies governing the pair-wise IPSec Security Association [KA98a] and governing the aspects of the key generation must also be in place.

The possible existence and possible interpretations of policies at different levels demands that the designer of any system to secure IP Multicast develop a set of policies which are coherent, free from loopholes and which address the possible scenarios to be met by the system.

Security of IP Multicast routing protocols

Although not directly affecting the security of the contents (data) and key management in IP Multicast, the protection of the multicast routing infrastructure itself is important for IP Multicast as a whole. This is due to the fact that the IP Multicast distribution tree is the packet delivery mechanism that carries the (encrypted) multicast data from the source to the receivers over the public Internet. This problem is a subset of the general problem of routing security, a problem that has received attention, among others, in [MB96, H98, and BBL98] in the context of unicast routing security.

The core of the problem is the authentication of control messages that are exchanged among the entities (such as routers) that constitute the routing infrastructure. Since such control messages inform these entities of the state of the network and impact the routing tables, the information contained in the control messages must be allowed to transit unmodified from the sender to the intended recipient. Related to this is the key management of the authentication keys used by the routing entities.

In the context of multicast routing, the PIM protocol [EFH97] has recently been augmented to include a number of cryptographic keys for the express purpose of authenticating the PIM control messages [W98]. A simple key management approach for the PIM authentication keys have also been proposed in [HC99]. Along similar lines, the authentication of the distribution tree within the CBT protocol [B97] has been addressed in [BC95]. More recent efforts have also been reported in [SG99].

Security of Reliable Multicast protocols

The topic of security in Reliable Multicast (RM) protocols is reasonably broad, and hence we only treat it briefly here. For simplicity and convenience, the security of RM protocols at the transport layer is usually treated separately from the security of multicast at the IP layer. However, such a separation is only artificial since both are closely related and may in fact deploy the same security mechanisms and policies.

When addressing the issue of security of RM protocols it is difficult to find a single solution for all RM protocols, since different RM protocols employ different techniques to provide reliability (e.g., ACK-based, NAK-based, source-retransmission, repair-nodes, etc) and employ different entities (e.g., routers, servers, hosts) to implement the reliability mechanisms. Thus, each RM protocol will require a different solution for their security needs.

Although there are many security issues relating to RM protocols, two core issues which are of immediate importance to all RM protocols are the authenticity of the control messages exchanged between the entities within the RM protocol and the authenticity of the retransmission of lost packets.

First, RM protocols require that all important control messages exchanged between RM entities be authentic. That is, exchanges of control messages should be protected against replay attacks and other freshness-based attacks. Which control messages need authentication is dependent on the specific RM protocol and on how many replay attacks the given RM protocol can sustain without the attacks becoming a denial-of-service event.

Secondly, a RM protocol must specify whether a retransmission entity (i.e., repair node) should apply its own authentication features (i.e., digital signature or MAC) whenever it retransmits a lost packet (assuming the sender has used authentication features on all the data it sent). Again, here the type of entity that performs the retransmission (e.g., source or repair-node) may determine whether that entity can suitably apply authentication.

One security issue related to the retransmission of lost packets is the availability (i.e., usage) of public key cryptography, particularly for source authentication.

If source authentication using public key cryptography is available, then there are two options with regards to the application of the public-key based digital signature. In the first option, source authentication can be provided by the source/sender above (or at) the RM level (transport layer). This has the advantage that a retransmission entity needs only to retransmit lost packets without adding its own public-key digital signature. Thus, in effect, authentication is truly end-to-end from the sender to the receivers, independent of the mechanism to achieve reliability. This convenience, however, comes at the cost of a receiver being open to replay attacks, since the receiver cannot sufficiently distinguish between an original packet being retransmitted from a repair-node and an original packet being replayed. The second option is to apply authentication at the IP layer (e.g., IPSec AH [KA98a, KA98b]) to all packets relating to the RM protocol. This approach would require the retransmission entity to apply its own authentication features to all packets which it retransmits, which can be a burden to some RM protocols. The gain here is that replay attacks can be minimized.

If source authentication using public-key cryptography is not deployed (e.g., for performance reasons), then group authentication via symmetric (private) key is the only remaining viable avenue. Regardless of whether group authentication is applied (via message authentication codes) below or above the RM level, a receiver will only be assured that a data packet was sent by a group member. Hence, an honest receiver will not be able to distinguish a retransmission by the proper retransmission-entity (either a repair node or the actual source/sender itself) from a retransmission (replay) by a dishonest group member who abuses the group-key.

In general, the security of RM protocols share the same underpinning problems as IP Multicast security and thus may in fact use the same solutions. These include key management, security policies, and data authentication and confidentiality. Hence, solutions designed for IP Multicast security should be considered in the larger context of use for RM protocols at the transport layer.

Summary of current efforts

The work of [CP98], [BMS99] and [HCD99] have each surveyed to different degrees the various approaches, protocols and solutions proposed for IP multicast security. In the following section, we briefly refer to the current efforts being conducted in this area, focusing primarily on the practical rather on the theoretical works. The summary is not meant to be comprehensive and the list of cited works not exhaustive. Hence, we encourage you to follow the references in order to obtain more details on each proposal or solution.

Group-oriented security, and more specifically the topic of its key management, has been researched now for more than two decades. Most of the earlier work has focused on cryptographic approaches to manage keys for hierarchic organizations and for conferences (e.g., [ITW82, KO87, BD94, STW96, BD96]). Others have sought different ways of sharing secrets within groups or to create digital multi-signature schemes. Many of these works have the application layer in mind, since they address more complex scenarios and applications.

Recently a number of practical solutions directed specifically at IP Multicast have been proposed, typified by those found in IETF-related efforts. Many of these address the issue of group-key management, since it represents one of the core problem areas in multicast security and the starting point for any security solution.

In general, it is useful to distinguish between solutions that cover the *mechanism* for key dissemination from those that address the *relationships* among keys, although in reality both are needed in order to attain some level of scalability. That is, a key may be derived from other keys, and thus it bears some relationship to them. This relationship may be determined, for example, by a mathematical function that allows the key to be updated (re-keyed) with ease. Examples of this type of work include that of [WHA98], [WGL98], [HH99a] and [BMS99]. However, a vehicle of delivery or dissemination of these keys, at least in the initialization phase, is still required and a number of proposal have also been put forward.

One of the earliest key-dissemination proposals is that of [HMR97a] and [HMR97b], where a group controller is deployed to create and deliver keys to the group members. The group controller also performs checking on the permission of candidate members. A similar approach, based on a centralized key management entity, is described in [CCP99].

Another early effort is that of [B96], which follows from the work of [B97] on the Core Based Trees (CBT) routing protocol. Here, the idea is to employ the core of the tree to distribute keys to candidate members, who must contact certain routers that are connected to the core. These routers then carry out membership checks and key distribution to the candidate members.

The problem of scalability is directly addressed by the framework proposal of [M97], where a hierarchical ordering of subgroups is employed to limit the effects of re-keying. The key management at different levels of the hierarchy is carried out by different controller entities. Thus, when re-keying occurs to a member within a subgroup, only the members in that subgroup will be affected. Although the work of [M97] points to an attractive direction in terms of limiting the effects of re-keying, it requires the decryption/re-encryption of traffic as it enters or leaves a subgroup.

Another effort to address the issue of scalability is that of [HD97], which does so by separating the key generation entity from the key distribution entity. The key distribution entities can be dynamically added by requesting their participation. Authority would then be delegated to such key distribution entities together with access control lists. Members would need to probe for the nearest key distribution entity in order to obtain the key.

A similar effort is reported in [HCM98] based the work of [WGL98], where a multicast region is defined to consist of several areas (subgroups). A member of the group is defined to reside within an area, and administratively-scoped multicast [M98, HTE97] is used to perform key delivery to members residing in a given area. The relationship among the keys follows that of the proposals of [WGL98], although the tree hierarchy of [WHA98] can also be deployed.

The approach is scalable to the extent that administratively-scoped multicast is available in the physical areas within which the members are located.

Conclusions and remarks

The current report has provided a brief coverage on the current developments in the area of IP Multicast security from the perspective of practical network-engineering solutions. It has focused on the concept and issues underlying the problem of IP Multicast security, and provided a brief summary of the current proposed solutions for these issues. Although perhaps not clearly stated, security represents an important aspect of IP Multicast, the lack of which is currently preventing its large-scale deployment. The need of standards in this area cannot be emphasized enough, since such standards are needed for the deployment of multicast over the public Internet.

Currently, the Secure Multicast Group (SMuG) IRTF is working towards standards in this arena. One of the aims of SMuG is to arrive at common building blocks, from which protocols can be developed in such a way that they can inherently interoperate with one another. The building blocks approach has a number of advantages. It allows new multicast security protocols to be built with minimal effort, and it allows new and better building blocks to be introduce to replace older ones. In this manner, interoperability can be designed into the protocols from their initial inception.

Acknowledgements

We thank Bob Quinn for his assistance and insights into the current report. We also thank Brad Cain for comments on certain aspects of multicast routing.

References

B96 A. Ballardie, "Scalable Multicast Key Distribution," RFC 1949, IETF, 1996.

B97 A. Ballardie, "Core Based Trees (CBT) Multicast Routing Architecture", RFC 2201, September 1997.

BBL98 T. Bates, R. Bush, T. Li and Y. Rekhter, "DNS-based NLRI origin AS verification in BGP", IETF, July 1998. draft-bates-bgp4-nlri-orig-verif-00.txt

BC95 A. Ballardie and J. Crowcroft, "Multicast-Specific Security Threats and Counter-Measures", in *Proceedings of the 1995 Symposium on Network and Distributed Systems Security* (NDSS'95), San Diego, February 1995, ISOC.

BD94 M. Burmester and Y. Desmedt, "A secure and efficient conference key distribution system," in *Advances in Cryptology - Proceedings of Eurocrypt'94* (LNCS No. 950), pp. 275—286, Springer-Verlag, 1994.

BD96 M. Burmester and Y. Desmedt, "Efficient and secure conference key-distribution," in Security Protocols (LNCS No. 1189) pp. 119-129, Springer-Verlag, 1996.

BMS99 D. Balenson, D. McGrew, and A. Sherman, "Key Management for Large Dynamic Groups: One-Way Function Trees and Amortized Initialization", dratf-balenson-groupkeymgmt-oft-00.txt, February 1999.

CCP99 R. Canetti, P-C. Cheng, D. Pendarakis, J.R. Rao, P.Rohatgi and D. Saha, "An Architecture for Secure Internet Multicast", IRTF, February 1999. draft-irtf-smug-sec-mcast-arch-00.txt

CDT99 B. Cain, S. Deering, and A. Thyagarajan, "Internet Group Management Protocol Version 3", IETF, February 1999. draft-ietf-idmr-igmp-v3-01.txt.

CEK99 I. Chang, R. Engel, D. Kandlur, D. Pendarakis and D. Saha, "Key Management for Secure Internet Multicast using Boolean Function Minimization Techniques", in *Proceedings of Infocom'99*, March 1999, New York, IEEE.

CP98 R. Canetti and B. Pinkas, "A taxonomy of multicast security issues", IETF, May 1998. draft-canetti-secure-multicast-taxonomy-00.txt (Also published in *the Proceedings of Infocom'99*).

D89 S. Deering, "Host extensions for IP multicasting," RFC 1112, IETF, 1989.

EFH97 D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification," RFC 2117, IETF, 1997.

F97 W. Fenner, "Internet Group Management Protocol Version 2," RFC 2236, IETF, 1997.

H98 A. Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option", IETF, March 1998, draft-ietf-idr-bgp-tcp-md5-00.txt.

HC98 D. Harkins and D. Carrel, "The Internet Key Exchange (IKE)," IETF, March 1998. draft-ietf-ipsec-isakmp-oakley-07.txt

HC99 T. Hardjono and B. Cain, "Simple Key Management Protocol for PIM", draft-ietf-pim-simplekmp-00.txt, March 1999.

HCD99 T. Hardjono, B. Cain, and N. Doraswamy, "A Framework for Group Key Management for Multicast Security," March 1999. draft-ietf-ipsec-gkmframework-01.txt

HCM98 T. Hardjono, B. Cain, and I. Monga, "Intra-Domain Group Key Management Protocol", November 1998. draft-ietf-ipsec-intragkm-00.txt

HD97 D. Harkins and N. Doraswamy, "A secure scalable multicast key management protocol (MKMP)," November 1997. Work in progress.

HH99a H. Harney and E. Harder, "Logical Key Hierarchy Protocol", IETF, March 1999. draft-harney-sparta-lkhp-sec-00.txt.

HH99b H. Harney and E. Harder, "Multicast Security Management Protocol (MSMP) Requirements and Policy", IETF, March 1999. draft-harney-sparta-msmp-sec-00.txt.

HMR97a H. Harney, C. Muckenhirn and T. Rivers, "Group Key Management Protocol (GKMP) Specification," RFC 2093, IETF, July 1997.

HMR97b H. Harney, C. Muckenhirn and T. Rivers, "Group Key Management Protocol (GKMP) Architecture," RFC 2094, IETF, July 1997.

HTE97 M. Handley, D. Thaler, and D. Estrin, "The Internet Multicast Address Allocation Architecture," IETF, December 1997. draft-handley-malloc-arch-00.txt.

ITW82 I. Ingemarsson, D. T. Tang, and C. K. Wong, "A conference key distribution system," *IEEE Transactions on Information Theory*, vol. IT-28, no. 5, pp. 714-720, 1982.

KA98a S. Kent and R. Atkinson, “Security architecture for the internet protocol,” RFC 2401, IETF, November 1998.

KA98b S. Kent and R. Atkinson, "IP Authentication Header (AH)", RFC 2402, IETF, November 1998.

KO87 K. Koyama and K. Ohta, "Identity-based conference key distribution systems," in *Advances in Cryptology - CRYPTO'87* (LNCS No. 293), pp. 175--184, Springer-Verlag, 1987.

M94 J. Moy, “ Multicast Extensions to OSPF,” RFC 1584, IETF, 1994.

M97 S. Mittra, “The Iolus framework for scalable secure multicasting,” *in Proceedings of ACM SIGCOMM'97*, pp. 277-288, ACM, 1997.

M98 D. Meyer, “Administratively scope IP multicast,” RFC 2365, IETF, July 1998.

MB96 S. L. Murphy and M. R. Badger, “ Digital signature protection of OSPF routing protocol,” in *Proceedings of the 1996 Network and Distributed System Security Symposium*, (San Diego), ISOC, 1996.

MOV97 A. Menezes, P. van Oorschot and S. Vanstone, *Handbook of Applied Cryptography*, CRC Press, 1997.

QA99 B. Quinn, K. Almeroth, "IP Multicast Applications: Challenges and Solutions", IETF, February 1999. draft-ietf-mboned-mcast-apps-00.txt.

S79 A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612-613, 1979.

S92 G. J. Simmons, "An introduction to shared secret and/or shared control schemes and their application," in *Contemporary Cryptology* (G. J. Simmons, ed.), pp. 441-497, IEEE Press, 1992.

SG99 C. Shields and J. J. Garcia-Luna-Aceves, "KHIP – A Scalable Protocol for Secure Multicast Routing", in *Proceedings of SIGCOMM'99*. (publication pending)

STW96 M. Steiner, G. Tsudik, and M. Waidner, "Diffie-Hellman key distribution extended to group communications," *in Proceedings of the 3rd ACM Conference on Computer and Communications Security*, (New Delhi), ACM, March 1996.

W98 L. Wei, “Authenticating PIM version 2 messages,” November 1998. draft-ietf-pim-v2-auth-00.txt.

WGL98 C. K. Wong, M. Gouda, and S. Lam, “Secure group communications using key graphs,” in *Proceedings of ACM SIGCOMM'98*, ACM, 1998.

WHA98 D. Wallner, E. Harder and R. Agee, "Key Management for Multicast: Issues and Architectures", IETF, September 1998. draft-wallner-key-arch-01.txt.

WPD88 D. Waitzman, C. Partridge, and S. Deering, “ Distance Vector Multicast Routing Protocol,” RFC 1075, IETF, 1988.