



Secure IP Multicast

Internet Multicast Security Overview of Issues and Technologies

Classification: **Unclassified**
Distribution: **Unrestricted**

Doc No: NU-R111
Release: C
Date: February 8, 1999
Owner: Amit Kleinmann & Shlomo Kipnis
Writer: Fern Levitt

© 1999 NDS Ltd.
All rights reserved.

All information contained in or disclosed by this document is considered confidential and proprietary by NDS Ltd. NDS Ltd. reserves the right to use this design in other projects without reference to the recipient. By accepting this material, the recipient agrees that this material and the information contained therein will be held in confidence and in trust and will not be copied or disclosed in whole or in part to any third party.

Internet Multicast Security: Overview of Issues and Technologies

Table of Contents

1. SCOPE.....	1
1.1. PURPOSE	1
1.2. STRUCTURE OF THE DOCUMENT	1
2. ABSTRACT.....	3
3. INTRODUCTION	4
3.1. UNICAST, BROADCAST, AND MULTICAST	4
3.2. BASICS OF IP MULTICAST TECHNOLOGY	4
3.3. TYPES OF MULTICAST APPLICATIONS.....	6
3.4. MULTICAST SECURITY	7
4. COMMUNICATION SECURITY	8
4.1. GENERAL SECURITY THREATS.....	8
4.2. ADDITIONAL MULTICAST SECURITY THREATS.....	9
4.3. COSTS OF DEFICIENT SECURITY	10
5. MULTICAST CHARACTERISTICS	12
5.1. MULTICAST SESSION CHARACTERISTICS	12
5.2. ONE-TO-MANY MULTICAST SESSIONS: GROUP CHARACTERISTICS.....	13
5.3. FEW-TO-FEW MULTICAST SESSIONS: GROUP CHARACTERISTICS.....	14
6. SECURING MULTICAST TRAFFIC	16
6.1. SERVICE REQUIREMENTS.....	16
6.2. SECURITY REQUIREMENTS.....	17
6.3. SERVICE MODELS	19

7. CURRENT APPROACHES TO SECURITY ARCHITECTURE.....	20
7.1. PAIR-WISE VS END-TO-END MULTICAST SECURITY.....	20
7.2. SERVER-CLIENTS VS MULTIPLE-PEERS MULTICAST SECURITY	21
7.3. THE IPSEC APPROACH	21
7.3.1. General Description.....	21
7.3.2. IPSEC and Multicast	22
8. KEY-DISTRIBUTION PROTOCOLS	24
8.1. SECURITY MANAGEMENT REQUIREMENTS.....	24
8.2. UNICAST KEY MANAGEMENT PROTOCOLS	24
8.3. MULTICAST KEY MANAGEMENT PROTOCOLS.....	25
9. SOFTWARE AND HARDWARE SECURITY	27
9.1. SOFTWARE SECURITY.....	27
9.2. HARDWARE SECURITY	27
9.3. SMART CARDS.....	28
10. NDS' END-TO-END INTERNET MULTICAST SECURITY SOLUTION.....	29
10.1. SECURITY AT THE APPLICATION LAYER.....	30
10.1.1. Client Side.....	30
10.1.2. Server Side	30
10.2. NDS' SECURE IP MULTICAST SYSTEM FEATURES:.....	31
10.3. NDS – A WORLD LEADER IN SECURITY SOLUTIONS	32
11. REFERENCES.....	33

1. Scope

1.1. PURPOSE

This document focuses on and analyzes issues and technologies that underlie Internet Multicast security.

The document is intended to provide the reader with an understanding of issues related to Internet Multicast, in general, and to multicast security, in particular.

NDS' Secure IP Multicast system is a modularly designed solution for securing Internet Multicast which can be tailored and customized to address the security concerns of the customer. Security concerns vary depending on the customer's business model and the sensitivity of the multicast data and content.

NDS' Secure IP Multicast system provides server and client middleware to secure content distribution and protect multicaster revenues. The Secure IP Multicast system is designed to mesh easily with the existing and emerging multicast protocol standards, while requiring low memory and storage overhead. NDS' Secure IP Multicast system is described in more detail later in this document. For further information about the NDS solution to Internet Multicast security, please contact NDS.

1.2. STRUCTURE OF THE DOCUMENT

Section 2 contains an abstract summarizing the paper.

Section 3 introduces the differences between unicast, broadcast and multicast, and gives further information about multicast technology, applications, and security.

Section 4 discusses data communications security, describing general threats, additional threats which relate specifically to multicast, and the costs resulting from inadequate security.

Section 5 presents parameters by which multicast sessions are characterized, and provides typical values for the two types of multicast sessions: one-to-many and few-to-few.

Section 6 lists service and security requirements for multicast traffic and several service models through which multicast data can be delivered to authorized clients.

Section 7 describes currently used approaches to multicast security architecture, such as pair-wise, end-to end, server-clients, multiple-peers, and IPSEC.

Section 8 discusses security requirements for key-management protocols, and discusses unicast and multicast key management protocols.

Section 9 discusses aspects of software and hardware security and how smartcards may be used to enhance security.

Section 10 contains a full description of the NDS solution for end-to-end Internet Multicast security.

Section 11 contains the list of references which are referred to throughout the paper by numbers in square brackets [].

2. Abstract

Internet Multicast is a technology for delivering data packets from a sender to multiple receivers in networks using the Internet Protocol (IP). Over the last few years, Internet Multicast has become widely used as a means for distributing data to groups of receivers and for communicating among multiple participants. Examples of applications based on Internet Multicast include transmission of corporate data to employees, communication of stock quotes to brokers, streaming audio and video content, online conferencing, collaborative work, and caching and replication of databases and web sites in multiple locations.

The basic Internet Protocol, in general, and Internet Multicast, in particular, suffer from lack of security. Security issues for multicast are extremely important, since multicast communication may involve multiple participants and may be spread over many networks. Furthermore, due to the nature of the basic Internet Multicast Protocol, receiver-participants can join and leave multicast sessions at will without notifying the multicast sender or other participants. While much effort has been invested in developing security solutions and mechanisms for point-to-point, or unicast communication, minimal attention had been paid to security issues in multicast communication.

This paper discusses the special nature of security issues for Internet Multicast. It focuses on the important technologies and architectures proposed to handle security in multicast environments, and it highlights the advantages and shortcomings of many of these approaches.

NDS has developed an end-to-end security solution for Internet Multicast which addresses many of these issues and can be customized to the needs of the customer.

3. Introduction

3.1. UNICAST, BROADCAST, AND MULTICAST

Traditionally, data communication in networks consists of two types of services – unicast and broadcast. Unicast (also called point-to-point) communication involves delivering data from one specified sender to a single specified receiver in the network. Broadcast, on the other hand, involves delivering data from one specified sender to all the receivers in the network. In Internet Protocol (IP) networks, the address of a receiver is specified in a unicast communication packet, and a unique address is used for specifying a broadcast communication packet.

Unicast and broadcast alone do not provide a full spectrum of efficient services in a communication network. Many Internet applications require delivering data from one sender, or from several senders, to multiple receivers in a network, but not to all the receivers in the network. Such applications include transmission of corporate data to selected employees, communication of stock quotes to brokers, streaming audio and video to interested individuals, online audio/video conferencing, collaborative work, and caching and replication of databases and web sites in multiple locations. For such applications, unicast is wasteful, in that it generates excessive traffic, while broadcast generates traffic on parts of the network that may not require it.

To address the need for another network service, a third technology, Internet Multicast, has emerged in the form of the Layer-3 Internet Multicast Protocol standard [1] as a means of delivering data from one sender to multiple receivers simultaneously over LANs and WANs. Internet Multicast is an efficient new inter-network service. It enables one-to-many and many-to-many datagram distribution services. It reduces both the sender transmission overhead, the network bandwidth requirements, and the latency observed by the receivers. These properties make Internet Multicast an ideal technology for communication among large groups of principals.

3.2. BASICS OF IP MULTICAST TECHNOLOGY

IP Multicast has been very successful in extending several IP Unicast mechanisms to provide an efficient, best-effort data-delivery service for large groups of receivers. IP Multicast uses a special range of IP addresses (Class-D addresses) designated for network multicast sessions. Multicast, at its core, is a connectionless protocol, much like the IP and the UDP protocols. However,

there are several major differences between the unicast and the multicast communication models.

One important difference between unicast and multicast is that, whereas unicast is a sender-based concept, multicast is a receiver-based concept. The sender is not consulted about the addition of a multicast receiver. Receivers can join and/or leave a particular multicast session, whether or not it is currently active, at will. Multicast traffic will continue to be delivered to all members of the multicast session group by the network infrastructure. The multicast sender does not need to maintain a list of the receivers. It is the job of the network routers to create and distribute copies of the multicast packets to receivers who join a particular multicast session group.

Multicast groups may range in size from a few nodes to thousands of nodes without affecting performance. Basically, the IP Multicast mechanisms work as follows:

1. The multicast sender selects the desired time and the content of the multicast.
2. The multicast sender obtains a Class-D IP address for the multicast session from a Multicast Address Allocation Server (MAAS) using the Multicast Dynamic Host Configuration Protocol (MDHCP) [2, 3]. The sender then submits multicast packets with this address to the network.
3. Interested receivers can request to join (and/or leave) certain multicast sessions by using the Internet Group Management Protocol (IGMP) [4].
4. Multicast routers in the network cooperate to deliver multicast session-addressed packets to receivers that have requested to join a multicast session group. They do so by running IP Multicast routing protocols such as the Protocol Independent Multicast (PIM) [5] and the Border Gateway Multicast Protocol (BGMP) [6].

Many features of unicast have proven to be very difficult to extend to multicast in a scalable manner. For example, multicast address allocation, multicast reliability, multicast flow control, and multicast congestion control all continue to be areas of active research, and proposed solutions are not mature. Research continues even in areas of IP Multicast where standards have been established. New versions of these standards are continually being developed. Examples include: group setup protocols such as the Session Announcement Protocol and the Session Description Protocol (SAP/SDP) [7], routing protocols such as the PIM and BGMP, and

management protocols such as IGMP and the GARP Multicast Registration Protocol (GMRP) [8].

3.3. TYPES OF MULTICAST APPLICATIONS

The spread and deployment of systems and products based on IP Multicast has grown substantially in the last few years, and the growth is expected to be even more dramatic in coming years. IP Multicast is considered to be one of the key technologies that will enable economical distribution of data over corporate Intranets, as well as distribution of data to private consumers over the public Internet. IP Multicast is, in fact, an infrastructure, on top of which many higher-level applications and services can be built. Four generic types of multicast applications are discussed below.

Non-Realtime / Non-Reliable. Examples include push applications that distribute non-critical data to computers at prescribed times, or based on the occurrence of certain events.

Non-Realtime / Reliable. Examples are file transfers and software distribution. Data loss rates on the Internet are about 5%-10%. Unreliable multicast (without error correction and congestion control) may actually increase the communication bandwidth rather than reducing it. Different reliable multicast protocols may be needed to provide good performance for different networks and applications. Several Internet Drafts on reliable multicast have been submitted, and in March 1997 a Reliable Multicast Research Group (RMRG) was formed within the Internet Research Task Force (IRTF) to advance standards for reliable multicast. Some known protocols for reliable multicast are the Multicast option of the Trivial File Transfer Protocol (TFTP) [9], the Starburst Multicast File Transfer Protocol (MFTP) [10], and the Intracom IoS Protocol [11].

Realtime / Non-Reliable. Examples include streaming audio/video applications. Transmitting real-time multimedia content over a network presents entirely different challenges from sending other kinds of data (e.g., text, files, graphics, bitmaps, etc.). The latter data types, while they may vary widely in their bandwidth requirements, necessitate the transmission of bursts of data. This means that they can withstand short and inconsistent periods of delay (on the order of several seconds) between packet transmissions. In contrast, the real-time nature of audio/video traffic mandates that the data be “streaming”, that is sent in a continuous flow rather than in bursts. In real-time traffic of this type, delays are typically measured in milliseconds.

Realtime / Reliable. Examples include reliable services that guarantee to the sender that all the packets are received by all intended receivers. Reliable delivery may be required by real-time applications such as data conferencing, Web services, and data broadcasting (e.g., distribution of real-time finance data). Several vendors have developed proprietary realtime and reliable IP-based multicast protocols designed for different applications and network environments. For example, Globalcast offers three families of such protocols: Reliable Multicast Protocol (RMP), Scalable Reliable Multicast (SRM) and Reliable Multicast Tree Protocol (RMTP) [12].

3.4. MULTICAST SECURITY

Data and network security are recognized as vital to the design of modern communication systems. Much work has been invested in developing security solutions and mechanisms for unicast communication. For multicast communication, the issues and problems are substantially more complex, since multicast involves multiple participants that may be spread over many networks and may join and/or leave dynamically.

Although the deployment of IP Multicast has been picking up, the area of multicast security has not received much attention. The few proposals for securing multicast communication do not address the unique requirements arising from the multicast group communication model. In addition, there are no reasonable schemes for key distribution that scale to large groups or to groups with highly dynamic membership. Until those issues are solved, the deployment of IP Multicast is expected to be limited.

A Secure Multicast Research Group (SMuG) was recently formed in the IRTF to advance standards for multicast security.

4. Communication Security

Basic data communication protocols incorporate very few security features, if any. Typically, security is handled at the level of the system that uses the communication protocols [13]. Most systems at least distinguish between two types of communication modes: (1) Insecure Communication, and (2) Secure Communication. Many systems offer more than these two basic security modes. Additional and more refined security modes may include authenticated communication, confidential communication, private communication, registered communication, different data protection levels, etc.

Insecure Communication modes are used when data may be available to all users and there is no fear of malicious attacks on the system. Secure Communication modes are used when data should be available only to a limited group of authorized users and/or there are fears of malicious attacks on the system. The notions of insecure and secure communication apply to all communication services (e.g., unicast, broadcast, and multicast).

4.1. GENERAL SECURITY THREATS

Communication protocols and the applications based on them, if not used cautiously, are subject to a multiplicity of security threats. Below are some of the common threats and types of attack.

Eavesdropping – One user can eavesdrop on the traffic destined for other users. For example, in a LAN environment, a machine can be tuned to promiscuous mode, thereby picking up any traffic on the LAN. Another case would be a machine, located along the path between the sender and the receiver(s), which can intercept traffic passing through it.

Impersonation – Communication protocols do not include source or receiver authentication. Any user can impersonate the sender or the receiver(s) of the data.

Data Manipulation – Packets flowing in the network can be modified en-route. This presents opportunities for an active attacker to change data as desired.

Denial of Submission – Communication protocols do not provide proof that messages were actually sent from the sender to the receiver.

Denial of Receipt – Communication protocols do not provide proof that messages sent by the sender were actually received by the intended receiver(s).

Repudiation – Communication protocols do not provide proof to a third party that the correct data were submitted to and/or received by the intended receivers.

Replay Attacks – Packets captured during a communication session may be resent later (played back) either to the sender or to the receivers, thereby damaging the state of the communication between them.

Denial of Service – Packets can be dropped or the network can be flooded with traffic, thereby denying delivery of the packets from the sender to the receivers.

Theft of Service or Content – In cases where payment is required for services or for content, there are risks that the services or the content could be supplied illegally without payment.

For unicast (point-to-point) communication, satisfactory security solutions for the threats listed above have been developed. For the most part, these solutions are based on the use of cryptography to sign messages, encrypt messages, submit authorized receipts, block unauthenticated traffic, etc. These solutions are based on Shared-Key Cryptosystems, Public-Key Cryptosystems, or combinations of both Shared-Key and Public-Key Cryptosystems.

For multicast communication, there are almost no straightforward solutions to these general security threats.

4.2. ADDITIONAL MULTICAST SECURITY THREATS

In addition to the above general threats to security, multicast communication is subject to some additional security threats and types of attack.

Uncontrolled Multicast Group Membership – IP Multicast protocols provide no means to specify, control, or limit the membership of a multicast session group.

Leakage of Security State – In multicast environments, the security state of a multicast session may be shared among multiple participants, thereby increasing the risk of security state leakage.

Security State Revocation – There is no simple mechanism to revoke the validity of a security state and to notify all multicast participants of the change. This may be necessary in order to restrict access of a multicast group member that has left the group.

Security Management – Multicast protocols do not provide mechanisms for managing the security attributes of multicast group members. In particular, the problems of key distribution and of re-keying a group are of major concern. Another significant problem is the lack of synchronization among the group members.

Multi-Point Attacks – In multicast, it is easier for an attacker to pose as one of many users, or to attack at several points in the network at the same time, thereby increasing the vulnerability of the system.

Multicast Session Publicity – Multicast sessions are usually well advertised (e.g., through the SAP/SDP protocols mentioned previously), thereby leaking some data in advance and making it easier for an attacker to target an attack.

Wide Impact – In an attack on a multicast environment, a large number of users may be affected, and the scope of the damage is usually unknown.

There are several proposals for enhancing the security of IP Multicast by introducing scope-limited multicast groups. However, in many situations, these mechanisms are not sufficient to address all the security threats described above. To thoroughly address the problem, an architecture is required which provides cryptography-based security services (authentication, encryption, etc.), augmented with specific security protocols for multicast environments.

4.3. COSTS OF DEFICIENT SECURITY

In developing security mechanisms and solutions for communication systems, one must bear in mind the costs associated with not providing security or with providing inadequate security measures:

1. Stolen, lost, or damaged data, which may have negative implications far beyond an assigned monetary value.
2. Error recovery and the risks involved in it. Perfect recovery may be impossible at any cost.

3. Damage to the reputation of a “hacked” system. Because this involves subjective perception, the losses resulting from lost credibility may far exceed the immediate costs of recovery.
4. Administrative and maintenance procedures. Overhead such as this may be handled more effectively and efficiently through a proper security solution.

5. Multicast Characteristics

Internet Multicast is a technology on top of which many higher-level applications and services can be developed. As such, Internet Multicast can be used in a variety of settings and environments. When developing architectures and protocols for securing multicast communication, one needs to consider the specific settings and environments in which the multicast will be used.

This section discusses several important parameters that can be used in defining the settings and environments in which the multicast operates. This classification is based, in part, on taxonomy definitions by Canetti and Pinkas [14]. Of particular significance among these parameters is the number of active senders in the multicast. This section concentrates on the “one-to-many” broadcast model and on the “few-to-few” conference model.

5.1. MULTICAST SESSION CHARACTERISTICS

Dissemination Scope – The geographic extent of the multicast session. This parameter is often measured by the “scope diameter”. Two extreme values for this parameter may be: (1) intra-domain scope (such as LANs, Intranets, or last-mile-switched WANs), and (2) inter-domain scope (such as WANs or the Internet Backbone).

Participation Scale – The size and the nature of the multicast session group. Important parameters are: (1) the number of active multicast senders (e.g., one or many), and (2) the number of receivers.

Sender Identity – The identity of the multicast sender. Some important considerations are: (1) whether the sender is known in advance, (2) whether the sender is a member of the multicast group, and (3) whether multiple senders multicast sequentially or simultaneously.

Group Members Identity – The level of identification of the multicast group members to the multicast sender. This information is extremely relevant in determining the security architectures and protocols to be used.

Membership Dynamics – The nature of, and rate of change to, multicast group membership. Possibilities include: (1) static group membership, in which the members are permanent and are typically known in advance, and (2) dynamic group membership, in which membership changes over time. Other parameters include the type of membership changes (join and/or leave), the frequency of membership changes, the distribution of

membership changes over time (in bursts or at a steady pace over time), etc.

Session Duration – The expected duration of a multicast session. This parameter is important in determining which protocols to use to set up and to dismantle the multicast group.

Traffic Volume and Rate – The amount and the rate of multicast traffic expected in a session. This parameter is important in determining which protocols and algorithms to use to protect the content.

Traffic Type – The nature and type of multicast traffic. Options include: real-time vs. non-real-time, reliable vs. non-reliable, synchronized vs. non-synchronized, interactivity level, etc.

Member Characteristics – Parameters describing the configuration of the member station. Important considerations are: (1) computing power of the station, (2) available memory at the station, and (3) the attention level of the station (e.g., when the station is on-line).

5.2. ONE-TO-MANY MULTICAST SESSIONS: GROUP CHARACTERISTICS

The one-to-many multicast model is most widely used today. In many multicast applications, one sender (or a small and fixed set of senders) submits content to a large multicast group. Examples of one-to-many multicast sessions include broadcasting a lecture, disseminating stock market results, or multicasting a file. In these examples, one sender disseminates data to multiple recipients who have joined an appropriate multicast group. Particular instances of this model include “Publish and Subscribe” (P&S) and Multicast FTP

Dissemination Scope – The dissemination scope may range from local scale (in Intranets) to wide scale (over Internet or satellites).

Participation Scale – The number of senders is one (or a small and well-coordinated set of senders). The number of receivers may range from tens or hundreds (in Intranets) to thousands or millions (over the Internet).

Sender Identity – The sender is typically a strong server capable of multicasting several channels to many receivers. Alternatively, the sender can be split among several servers working in cooperation. The sender must be known in advance, but it does not need to be a

member of the multicast group. Authentication, secrecy, and service levels must be maintained. Security may possibly be handled by additional servers.

Group Member Identity – In order to restrict and secure the multicast traffic, the identity of the receivers must be known in advance to the sender (or to the security servers it maintains). This is similar to the broadcast model (for example Pay-TV).

Membership Dynamics – This model typically assumes that members have long-term relationships with the multicast sender. However, group membership may be very dynamic, with members joining and leaving at high rates. This may result in a high volume of sign-on and sign-off requests. Membership revocation should be performed within seconds or minutes from the time it is submitted.

Session Duration – The multicast sessions are usually long (at least on the order of several minutes, if not hours).

Traffic Volume and Rate – The volume and the rate of the multicast data may vary considerably. With multimedia content, volume can be high and very little latency will be allowed. With textual content, volume is low and latency requirements are quite relaxed.

Traffic Type – As in the previous item, different applications will have different requirements for traffic type.

Member Characteristics – Receivers are typically low-end or medium machines with limited resources. Consequently, security solutions should be optimized for efficiency at the receiver side.

5.3. FEW-TO-FEW MULTICAST SESSIONS: GROUP CHARACTERISTICS

Another model of multicast that receives attention is few-to-few multicast sessions. Examples include audio/video (A/V) conferences among several participants, distributed committees, interactive lectures, and multiparty games. In such scenarios, each participant in a session can send data to and receive data from other participants. In these situations, there may not always be a natural group owner that can serve as a trusted center, and it may sometimes be beneficial to distribute trust among the participants.

Dissemination Scope – The dissemination scope is typically wide (otherwise the group members might meet in person).

Participation Scale – The number of participants may range from several tens to potentially hundreds of peers. Each participant may be both a sender and a receiver.

Sender Identity – Since each group member may be a sender, typically the multicast is sent directly from each group member. Sometimes, however, there may be a small set of members that generate most of the bandwidth. Security of source and data is critical.

Group Members Identity – In order to be able to restrict the scope of the multicast traffic, the identify of each receiver must be verified. This could be done by using the services of a centralized authority or by trusting several of the participants.

Membership Dynamics – The group membership is typically static after the group has been initiated, although there may be cases of members joining or leaving the group during a session. Furthermore, even if a member leaves the group, it may not always be necessary to revoke his membership. Therefore, a low volume of sign-on / sign-off requests is expected. When necessary, member revocation should be performed within seconds or minutes from the time it is submitted.

Session Duration – The multicast group is often formed for a specific event and is relatively short-lived (minutes or hours).

Traffic Volume and Rate – The volume and the rate of the multicast data may vary considerably. With multimedia content, volume can be high and very little latency will be allowed. With textual content, volume is low and latency requirements are quite relaxed.

Traffic Type – The requested communication latency may vary from application to application, but it should typically be small in order to facilitate the simultaneity and interactivity of virtual conferences. In some scenarios, data secrecy may be crucial, while in others it may not. In some scenarios, maintaining anonymity of members may be crucial.

Member characteristics – Typically, group members' machines have similar computational resources. Machines can be high-end A/V terminals, or they can be medium or low-end machines.

6. Securing Multicast Traffic

This section covers service and security requirements which must be satisfied in securing multicast traffic. The last part of the section describes several models through which multicast data may be securely delivered to authorized receivers.

6.1. SERVICE REQUIREMENTS

Service availability – Maintaining service availability against malicious attacks.

Group Management – This includes the following:

Access control – Controlling group membership and keeping records of the amount of usage of each member.

Key management – Creating, changing, and distributing keys, synchronization between encrypted data and the associated keys, and tracking keys.

Logging/Audit – Maintaining necessary records of user and system activities.

Error and exception handling – Manage and recover from illegal system, user, or administrator actions.

Performance requirements – This includes the following:

Bandwidth – Managing the use of multiple channels and ensuring that the available bandwidth is not exceeded by individual or collective channels.

Work overhead per data packet – The allowed latency.

Initialization and termination: the overhead associated with initialization and termination should be minimized. It can be divided into the following:

- **Group** – The overhead of initializing/terminating a group session.

- **Sender** – The overhead of a sender when it starts/stops transmitting to the group.
- **Member** – The overhead of member addition/deletion

Key generation and distribution overhead – Ensuring that keys are handled within the required time limits.

Congestion – Especially around centralized control services at peak sign-on and sign-off times, e.g., when many customers join a group right before the multicast begins and leave right after it ends.

Resume overhead – The worked incurred when a group member becomes active after being dormant (e.g., off-line) for a period of time.

6.2. SECURITY REQUIREMENTS

To secure multicast traffic, numerous requirements may need to be addressed. Some of the most important ones are listed below.

Controlled Membership – The multicast server should be able to specify, control, and restrict the membership of any multicast session group.

Source Authentication – Multicast sessions and data should be protected for authenticity of the multicast server.

Client Authentication – The authenticity of every authorized receiver of any multicast session should be verified.

Data Integrity – Data should be protected against unauthorized changes either at the multicast server or en route.

Data Confidentiality – Only the authorized receivers of a multicast session should be able to understand the contents of the transmission. In some cases, preventing access to crucial parts of the data may be sufficient.

Flow Confidentiality – Only the authorized receivers of a multicast session should be able to detect the existence of the transmission.

Client Privacy – Maintaining the privacy of the selections of a multicast group member, so that selections are not known to outsiders or other group members. Protection from traffic analysis.

Source Anonymity – Keeping the identity of a multicast message sender in confidence, or even anonymous.

Client Anonymity – Keeping the identity of a multicast group member in confidence, or even anonymous.

Proof of Submission – The multicast server should be able to provide proof that multicast sessions and data were submitted properly.

Proof of Receipt – The multicast server should be able to provide proof that multicast sessions and data were received properly by their receivers.

Non-Repudiation – The multicast system should be able to provide proof to a third party that certain transactions were completed in full.

Security Policies – The multicast system should support different policies with regard to authentication, authorization, privacy, etc.

Security State Updates – The multicast system should be able to update the Security State for groups of receivers promptly and efficiently.

Key Refreshment – The need to change the key during a lengthy multicast session, in order to foil cryptanalysis and other methods for comprising the key.

Guaranteed Service – The multicast system should provide protection against lost data packets and against attempts to flood the network.

Perfect Forward Secrecy – The compromise of an client should not cause data sent before the compromise to be accessible.

Local Secrecy – The compromise of an individual client should not cause data sent to other clients to be accessible.

Replay Resistance – The multicast system should be able to resist replay attacks. This may be particularly relevant in reliable multicast services where some data may need to be resent to certain receivers.

Payment – The multicast system should be able to support the collection of payment for content or data consumed by receivers, according to the established service models (e.g., Subscription, Article, Volume, and Time).

6.3. SERVICE MODELS

Security of multicast sessions may be required in order to protect a broadcaster's revenue for the material which has been received by a client, or to enforce secrecy of sensitive data. Either case requires a security system which ensures proper authorization for access to protected data.

There are several models by which potential participants in the network may become authorized to receive a secure multicast.

Subscription – In this model, an interested participant subscribes to a continuous multicast stream by ordering it in advance with the multicast source. If the service requires payment, it is negotiated and arranged in advance. This model is similar to long-term subscription to Pay TV or Cable TV channels, lasting weeks, months, or years.

Article – In this model, an interested participant orders a particular article (e.g., a software product, game, or video clip) or package of articles from an ongoing multicast service. The article or package is characterized by its start and end points in the multicast stream. An order may be placed and paid for in advance with the multicast source, or it may be handled immediately prior to multicasting the article. This model handles the one-time purchase of articles out of a content stream.

Volume – In this model, an interested participant requests some volume of data to be consumed out of a multicast stream (e.g., articles on a particular topic or trickles of updates to a database). The order can be placed in advance or immediately prior to the content consumption. Payment can be handled before, during, or after the content consumption. This model is useful for purchasing just-enough data.

Time – In this model, an interested participant requests the allotment of multicast bandwidth for a specified or measured amount of time. This model is useful in renting bandwidth for user applications (e.g., lecturing or conferencing). The order can be placed in advance or immediately prior to the content consumption. Payment can be handled before, during, or after the content consumption. This model is similar to the manner in which phone calls are charged by phone service providers.

7. Current Approaches to Security Architecture

We may divide up the multicast universe based on two parameters: whether security protects communication on a pair-wise or end-to-end basis, and whether the server-clients or multiple-peers model applies. Security architectures must provide the appropriate solution to the problem as determined by the combination of these two parameters

7.1. PAIR-WISE VS END-TO-END MULTICAST SECURITY

Security architectures for data communication, in general, and for Internet Multicast, in particular, can be classified into two types: Pair-Wise and End-to-End.

Pair-Wise security architectures protect the communication on a link-by-link basis. Specifically, packets are authenticated / verified and encrypted / decrypted at every router, firewall, and proxy along the path. Communication in Pair-Wise security architectures consists of two stages: Setup and Classification. In the Setup stage, each pair of communicating stations coordinates a security state to be used to protect subsequent communication between them. In the Classification stage, arriving packets are matched with the relevant security state and are processed accordingly.

There are several advantages to Pair-Wise security architectures:

- Link layer security mechanisms are well understood and are robust.
- The local security structure simplifies the management of the system.
- High scalability is achievable due to the distributed nature of the protocols.

End-to-End security architectures protect communication between end points in the network. Authentication/verification and encryption/decryption are done at the server and at the clients, and not at intermediate nodes. Protected data packets need to be able to cross routers, firewalls, and proxies. Communication in End-to-End security architectures consists of two stages: Registration and Streaming. In the Registration stage, each client registers with the server and coordinates a security state to be used to protect

subsequent communication. In the Streaming stage, packets are protected using the pre-established security state between the server and each client.

There are several advantages to End-to-End security architectures:

- The server can identify and control each and every client.
- Data is never exposed on the path from the server to the client(s).
- Security can be implemented at various levels of the applications.

These two generic types of architectures assume different models of control for secure communication. These security models have implications on the service models. For example, questions of who monitors the network and who charges for the services are addressed differently by each of these architectures.

7.2. SERVER-CLIENTS VS MULTIPLE-PEERS MULTICAST SECURITY

Architectures for Internet multicasting can also be classified into Server-Clients and Multiple-Peers. (We have coined the term “Server-Clients” to refer to a situation where one server communicates simultaneously with multiple clients.) In Server-Clients architectures, one distributor sends the multicast content to many consumers. The distributor is responsible for clients’ authentication and authorization, and for key management and distribution. Clients may be required to submit multicast content through the distributor.

In Multiple-Peers architectures, each participant can be both a receiver and a sender of multicast content. Each participant must be able to manage security states and perform authentication / verification and encryption / decryption against all other participants in the system. Unless participants have very strong processing capabilities, the levels of performance and security in such systems are very limited.

7.3. THE IPSEC APPROACH

7.3.1. General Description

IPSEC [15] is a standard Multiple-Peers security architecture being developed by the IETF, which can be used in either pair-wise or end-to-end scenarios. The goal is for IPSEC to provide various security services for traffic at the IP layer, including, but not limited to, IP Multicast traffic. IPSEC provides these

services at the IP layer by enabling a system to select the required security protocols, to determine the algorithms to use for the services, and to establish the cryptographic keys required to provide the requested services. IPSEC can protect one or more paths between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

IPSEC uses two protocols to provide IP traffic security.

1. Authentication Header (AH) [16]. The AH protocol provides connectionless data integrity, data origin authentication, and anti-replay service.
2. Encapsulating Security Payload (ESP) [17]. The ESP protocol provides privacy and limited traffic flow confidentiality. It provides connectionless data integrity, data origin authentication, and anti-replay service.

IPSEC specifies IKE [18] – a public-key approach for automatic key management. IPSEC defines the term Security Association (SA) as a “simplex connection” that provides security services to the traffic carried by it. Each SA is uniquely identified by a triplet consisting of a Security Parameter Index (SPI), an IP destination address, and a security protocol (AH or ESP) identifier. Both AH and ESP use SA, and a major function of IKE is the establishment and the maintenance of the different SAs. RFC 1825 [19] provides details on IP SA and SPI definitions.

7.3.2. IPSEC and Multicast

The SA-management mechanisms in IPSEC are currently defined only for unicast. For multicast traffic, a system or person will need to select SPIs on behalf of the multicast group and then communicate the information to all of the legitimate members of that multicast group via mechanisms not yet defined [19]. Support for multicast requires defining and managing a dynamic SA, that is, an SA that can be changed dynamically and frequently in response to changes in the multicast group membership. Such changes may involve some “heavy-duty” tasks in IPSEC:

- Authentication and approval of a join/leave request.
- Distribution of new group keys (especially for a leave request).

Future versions of the IPSEC standard may provide architectural details for IPSEC Multicast support. However, at the present time, no solutions have been proposed. For example, the latest Internet Draft of the IPSEC security architecture (May 1998) [15] does not specify mechanisms for coordinating the selection of the SPI/s and for communicating a multicast group’s IPSEC information to all valid members of the multicast group.

Another concern with IPSEC for multicast is the lack of a definite authentication mechanism for each member of the multicast group. For example, IPSEC specifies that multiple senders in the same multicast group should use a single SA to secure the traffic in the group, when using symmetric-key encryption and authentication algorithms. In such situations, although receivers of the multicast will be able to authenticate the validity of the multicast group, they will not be able to authenticate the specific senders that submitted the multicast traffic.

8. Key-Distribution Protocols

8.1. SECURITY MANAGEMENT REQUIREMENTS

Key distribution protocols are used to establish common key states at participants. These protocols are part of the system security management framework. Below are several security management requirements that are related to key-distribution.

- Support for multiple authentication and encryption algorithms.
- Negotiation of SA, including the algorithm, key, label, and services.
- Revocation of SA, including defining the lifetime of a session.
- Support for key- and session-recovery mechanisms.

8.2. UNICAST KEY MANAGEMENT PROTOCOLS

For unicast communication, several key management protocols have been proposed, including SKIP [20], ISAKMP [21], Oakley [22], SKEME [23] and Photuris [24]. These protocols negotiate the parameters of the SA and establish a secure channel.

ISAKMP provides only a framework for authentication and key-exchange protocols, but does not define them. It is designed to be independent of the actual key-exchange protocol. Oakley describes a series of key-exchange stages called “modes”, and it details the services provided by each mode (e.g., perfect forward secrecy, identity protection, and authentication). SKEME describes a versatile key exchange technique that provides anonymity, repudiability, and fast key refreshment.

IPSEC specifies the Internet Key Exchange (IKE) protocol [18]. The description of this protocol is based partially on Oakley and partially on SKEME, under the ISAKMP framework. IKE presents different key-exchange stages as modes that operate in one of two phases. These modes are used to obtain authenticated keys to be used by ISAKMP and by other security associations, such as AH and ESP.

8.3. MULTICAST KEY MANAGEMENT PROTOCOLS

In many systems, the key distribution function is assigned to a central network entity, sometimes known as the Key Distribution Center (KDC). However, this method does not scale well for wide-area multicasting, where group members may be distributed across many networks and a wide-area group may be densely populated. Even more complicated is the problem of distributing sender-specific keys in a scalable manner. (Sender-specific keys are required when data is to be authenticated on a per-sender basis.) Pair-Wise key-management protocols and Key Distribution Centers do not provide scalable solutions for the multicast key-management problem.

Completely automatic protocols for multicast key distribution are currently not considered mature enough for standardization. For small multicast groups, manual key distribution or multiple invocations of a unicast key distribution protocol (such as authenticated Diffie-Hellman) appear adequate. However, for large multicast groups, new scalable techniques and protocols are needed.

Some recent works on multicast key-distribution protocols include the following:

- Group Key Management Protocol (GKMP) [25]. This protocol generates and maintains symmetric keys for the members of a multicast group. Each multicast group has a dedicated Group Controller (GC) which manages the group keys. The GC has a Mutual Unique Shared Key (MUSK) with each of the group members. The GC generates the group keys together with a selected group member. Then, the GC contacts each group member, validates its permissions and sends it group keys encrypted by the MUSK.
- SKIP Extensions for IP Multicast [26].
- Multicast Key Management Protocol (MKMP) [27].
- Scalable Multicast Key Distribution (SMKD) [28]. This protocol is based on the Core-Based Tree (CBT) routing protocol and provides secure joining of a CBT group in a scalable manner. The core of the tree operates as the GC. As routers join the delivery tree they are delegated the ability to authenticate joining members and provide them with the group key.

However, some of these protocols are not scalable, others are not secure enough, and the remaining ones are not practical in the existing infrastructure. For example, the key-revocation mechanism of GKMP assumes that the user revokes the key (which is generally unacceptable). The fact that a single entity, the GC, in the GKMP protocol, is responsible for sending the keys to all group members makes this approach non-scalable. As another example,

SKMD requires some modifications to the IGMP protocol. Also, since every router in the delivery tree obtains the same keys as the GC, the scheme does not provide a high level of security against corrupt routers in the group tree. Another major drawback is that SMKD assumes that the Core Base Tree (CBT) multicast routing architecture is being deployed, which is not the situation at the present.

9. Software and Hardware Security

Every security system has certain “secure areas.” These areas are typically used to store the long-term secrets of the system, through which ongoing operational security keys and states are derived. There areas can be physically secure (e.g., the master-key databases in Key Distribution Centers, the administrative stations in an organization, or the smart cards that people use to keep their secret keys). Alternatively, these areas can be secured by software only (e.g., root-protected directories in the Unix operating system, the Microsoft Windows Registry, or the Java operating system).

9.1. SOFTWARE SECURITY

Security systems built using software components only are subject to attacks. There are many examples of attacks by computer viruses, Trojan Horses, or other rogue programs. Standard operating system files can be copied, run-time memory can be monitored, and long-term storage can be modified. Security keys and states as well as the security algorithms themselves can be bypassed or even cracked.

There are some efforts to develop tamper-resistant software. Such software is harder to reverse-engineer and is hooked to low-level operating systems primitives in ways that are almost impossible to bypass without damaging the system. NDS’ Secure IP Multicast system key distribution protocols and key stores may take advantage of tamper-resistant software.

9.2. HARDWARE SECURITY

Highly secure systems always base their security on hardware components. There are clear advantages to using tamper-resistant hardware devices. Such devices can run small and highly secure operating system kernels on proprietary CPUs, their long-term storage is protected, and their memories are extremely difficult to monitor.

Numerous hardware devices and tokens have been developed, including secure PC boards, secure chips, PC dongles, authentication tokens, and smart cards. NDS has been working with physically secure cryptographic smart cards for over 10 years. It is important to stress that it is not enough to secure hardware devices alone. Many times, secure data can leak by monitoring the I/O behavior of the secure device. NDS also has extensive experience with secure smart card readers and access protocols.

9.3. SMART CARDS

A smart card is a removable plastic card with an embedded microprocessor chip (as defined in the ISO-7816 standard). The microprocessor is usually based on a proprietary design, has secure storage, and cannot be normally monitored. In NDS' Secure IP Multicast system, a smart card is used in the CPU for a number of functions:

- To store long term secret keys (symmetric and asymmetric).
- To run the authentication protocols and negotiate secret session keys.
- To encrypt and decrypt highly secure data.
- To store certificates and service entitlements.
- To sign and verify transactions (e.g., payments, key revocations, etc.).

10. NDS' End-to-end Internet Multicast Security Solution

This paper has described the security issues in Internet Multicast.

NDS, leaders in safeguarding pay TV content and billions of dollars of broadcasting revenues, have created a Secure IP Multicast system to provide a solution for end-to-end multicast security. This solution provides:

- Advanced multicast security that is modular and futureproof.
- Cost effective, simple implementation.
- Proven security technology used by leading broadcasters worldwide.

NDS' Secure IP Multicast system is designed for corporations, data broadcasters, content providers, multicast service operators, and ISPs who want to manage, control and secure their multicast content distribution.

NDS' Secure IP Multicast system is a complete, secure multicast distribution system. It can be used as a stand-alone system or integrated with existing multicast distribution systems to provide a wide array of security services, ensuring that confidential information remains that way. In addition to providing security for existing systems, NDS' Secure IP Multicast system can itself efficiently deliver data over your LANs and WANs.

As mentioned previously, current and evolving IP multicast communication protocols incorporate very few management, control, and security features. NDS' Secure IP Multicast system solves this problem. It provides the server and client middleware to secure your content distribution and protect revenues. The system is designed to mesh easily with the existing and emerging multicast protocol standards, while requiring low memory and storage overhead.

NDS' Secure IP Multicast system is based on proven NDS conditional access technology which is used by more than 10 million subscribers and protects billions of dollars of broadcaster revenues annually. NDS knows how to protect data with a solution that is modular, growing with your needs, and futureproof, protecting your investment.

NDS' Secure IP Multicast system protects multicast systems against these threats using proven cryptographic mechanisms. It also gives you the tools

you need to manage the security characteristics of the multicast content, the security attributes of your users and the easy distribution and maintenance of decryption keys.

NDS' Secure IP Multicast system provides authentication of multicast source and recipients, data integrity, authenticity, and secrecy, and accountability for multicast content distribution and consumption.

10.1. SECURITY AT THE APPLICATION LAYER

NDS' Secure IP Multicast system is an application-level solution. It uses information about the security characteristics of multicast content to increase protection while minimizing communication overhead.

NDS' Secure IP Multicast system is an open, comprehensive architecture for securing all types of multicast content. It provides easy-to-use security APIs for multicast applications, shielding application developers from complex security concerns and mechanisms.

NDS' Secure IP Multicast system includes a series of modular components that allow the system to keep pace with your needs:

10.1.1. Client Side

- NDS' Secure IP Multicast Security Middleware – Handles client security and manages multicast. Receives secured multicast content, verifies and deciphers content, and manages client-side keys.
- NDS' Secure IP Multicast Client Catalog – Provides a tool for selecting from available multicast content.
- Smart card (optional) – Gives you a tamper-resistant, secure, and portable device for storing and calculating keys, certificates, and entitlements.

10.1.2. Server Side

Server components can all run on one computer, or can be distributed over several machines:

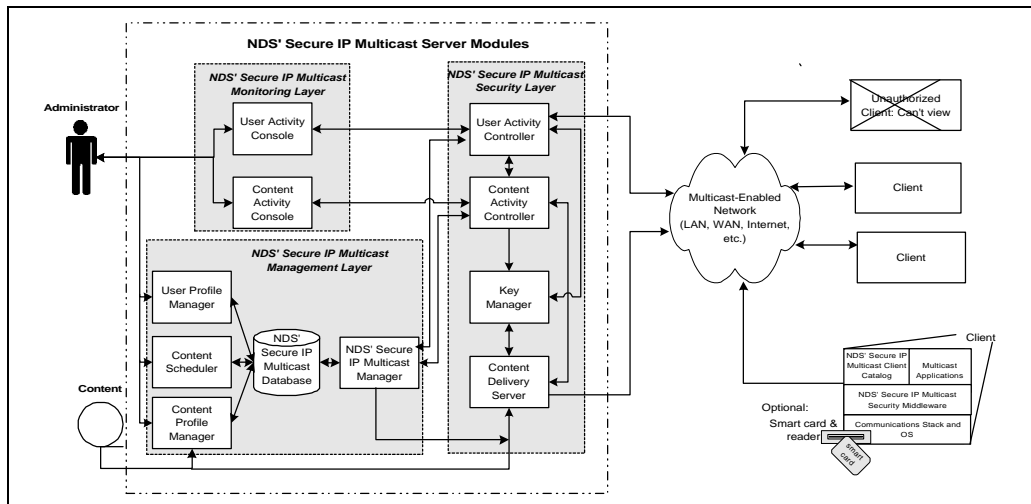
- NDS' Secure IP Multicast Manager – Controls distribution of secure multicast content using a flexible administrator GUI.

- Content Scheduler – Provides a tool for determining multicast times and channels.
- Multicast Distributor – Multiplexes and distributes protected content and control information.
- Security Manager – Authenticates and encrypts multicast content.
- Content Manager – Determines and packages multicast content.
- User Manager – Manages user-profile database and controls user access.
- Key Manager – Issues, stores and manages cryptographic keys.

10.2. NDS' SECURE IP MULTICAST SYSTEM FEATURES:

- Complete server and client middleware solutions, with documentation.
- Comprehensive solution for all types of multicast content.
- Distributed and scalable architecture and algorithms.
- Software solution with optional smart card security enhancements.
- Application-layer solution, enabling tailoring of security services to content type.
- Runs over any IP multicast infrastructure, whatever the switching and routing protocols.
- Easy to integrate into existing infrastructure and applications.

The following diagram illustrates the relationship among the basic components of NDS' Secure IP Multicast system architecture:



NDS' Secure IP Multicast System Components

10.3. NDS – A WORLD LEADER IN SECURITY SOLUTIONS

NDS, a world leader in data broadcasting and in smart card-based conditional access systems, has applied its expertise to Internet multicast. NDS expertise includes security, data broadcasting, content management, programming guides, and subscriber management systems. Our commitment to security is reflected in the ongoing efforts of hundreds of experts in the areas of security, broadcast, multicast and interactive technologies.

NDS is an executive member of the IPMI (IP Multicast Initiative) industry consortium and a leading participant in the Multicast Security Working Group of the IRTF (Internet Research Task Force). NDS' Secure IP Multicast system complies with emerging standards for Internet multicast.

11. References

- [1] S. Deering, “Host extensions for IP Multicasting”, RFC1112, 1989.
- [2] B. Patel and M. Shah, “Multicast address allocation extensions to the Dynamic Host Configuration Protocol”, draft-ietf-dhc-mdhcp-03.txt, November 1997.
- [3] M. Handley, D. Thaler, and D. Estrin, “The Internet Multicast Address Allocation Architecture”, February 1997.
- [4] B. Cain, S. Deering, and A. Thyagarajan, “Internet Group Management Protocol, Version 3”, draft-ietf-idmr-igmp-v3-00.txt, November 1997.
- [5] 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”, RFC2117, June 1997.
- [6] D. Thaler, D. Estrin, and D. Meyer, “Border Gateway Multicast Protocol (BGMP): Protocol Specification”, draft-ietf-idmr-gum-02.txt, March 1998.
- [7] M. Handley and V. Jacobson, “SDP: Session Description Protocol”, RFC2327, April 1998.
- [8] “The Multimedia World according to GMRP”, 3COM, 1997.
- [9] A. Emberson, “TFTP Multicast Option”, RFC2090, February 1997.
- [10] K. Miller, K. Robertson, A. Tweedly, and M. White, “StarBurst Multicast File Transfer Protocol (MFTP) Specification”, draft-miller-mftp-spec-03.txt, April 1998.
- [11] J. Vlontzos, “Internet over Satellite: IoS specification”, White Paper, Intracom, June 1996.

-
- [12] Globalcast Communication Inc, "Reliable Multicast Protocol", White Paper.
- [13] Gregory W. White, et al, "Computer System and Network Security", CRC Press Computer Engineering, 1995.
- [14] R. Canetti, B. Pinkas, "A Taxonomy of Multicast Security Issues", draft-canetti-secure-multicast-taxonomy-00.txt, May 1998.
- [15] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", draft-ietf-ipsec-arch-sec-05.txt, May 1998.
- [16] S. Kent and R. Atkinson, "IP Authentication Header", draft-ietf-ipsec-auth-header-06.txt, March 1998.
- [17] S. Kent and R. Atkinson, "IP Encapsulating Security Payload (ESP)", draft-ietf-ipsec-esp-v2-05.txt, March 1998.
- [18] D. Harkins and D. Carrel, "The Internet Key Exchange (IKE)", draft-ietf-ipsec-isakmp-oakley-07.txt, March 1998.
- [19] R. Atkinson, "Security Architecture for Internet Protocol", August 1995, RFC1825.
- [20] A. Aziz, T. Markson, and H. Prafullchandra, "Simple Key-Management for Internet Protocols (SKIP)".
- [21] D. Maughan, M. Schertler, M. Schneider, and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", draft-ietf-ipsec-isakmp-09.txt, March 1998.
- [22] H. Orman, "The OAKLEY Key Determination Protocol", draft-ietf-ipsec-oakley-02.txt, July 1997.

[23] H. Krawczyk, "SKEME: A Versatile Secure Key Exchange Mechanism for Internet", IEEE Proceedings of the 1996 Symposium on Network and Distributed Systems Security.

[24] P. Karn and W. Simpson, "Photuris: Session-Key Management Protocol", draft-simpson-photuris-17.txt, November 1997.

[25] H. Harney and C. Muckenhirn, "Group Key Management Protocol (GKMP) Specification + Architecture", RFC2093 & RFC2094, July 1997.

[26] A. Aziz, T. Markson, and H. Prafullchandra, "SKIP Extensions for IP Multicast".

[27] MPMP.

[28] A. Ballardie, "Scalable Multicast Key Distribution", RFC1949, May 1996.