Security Services for Multimedia Conferencing

Stuart G. Stubblebine

USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
stubblebine@isi.edu

Abstract

Multimedia conferencing is the use of mixed media such as real-time audio, video, and groupware for group tele-collaboration. Multimedia conferencing may well become as widespread an application on network computers as electronic mail. The demand for security in multimedia conferencing is high because the users' expectation of being monitored is high. In this paper we discuss security services for protecting multimedia conferencing, how well these services scale to large conferences, and what security infrastructure is needed.

1. Introduction

Electronic mail is perhaps the most popular application on networked computers. Several standards for secure electronic mail have evolved, including the Internet's Privacy-enhanced Electronic Mail [L93] and the U.S. government's Pre-Message Security Protocol [B93]. With the exception of electronic mail systems running on private networks protected by link encryption, secure electronic mail has been slow to mature. The slow growth is partially attributed to the difficulty of developing an acceptable infrastructure for public key certificate based authentication [X.509]. It is also attributed to users' lack of awareness to the susceptibility of electronic mail to wire-tapping attacks.

Multimedia conferencing, the combination of real-time packet audio, video, and groupware multiway tele-collaborations, is likely to become as widespread an application on network computers as electronic mail. As such, significant research and development is needed to integrate security into the design of teleconferencing systems. The growing demand for teleconferencing is demonstrated by the regular multicasts of audio and video from the Internet Engineering Task Force (IETF) meetings over the MBONE, a multicast capable portion of the Internet [CD92]. The most recent teleconference included 339 sites in 14 countries. There is a growing demand for security as participants are becoming aware of how easy it is for others to monitor their conferences. This is because architectures that scale well to support large conferences do not restrict the distribution of conference packets. That is, conference sessions are announced via network multicast to a session directory service, and users participate in the conference by simply joining the multicast address associated with the conference. Private announcements via electronic mail or other means offer only limited protection since an eavesdropper can still scan the multicast address space to detect conferences in progress.

In the next section, we discuss security services for protecting multimedia conferencing. We assume the use of trusted computing bases and emphasize the requirements of communication security. In section 3, we discuss the scalability of these security services to large conferences. In section 4, we discuss what security infrastructure is needed to provide these services.

2. Security Services for Multimedia Conferencing

A multimedia conference may consist of several media flows, for example audio, video and groupware application flows. (A groupware application typically makes use of a shared workspace. An example of a groupware application is a shared whiteboard.) The protection required for each data flow in the conference may vary. For example, it may be necessary to protect the confidentiality and authenticity of

Proceedings of the 16th National Computer Security Conference, Baltimore, Maryland, September 20-23, 1993.

an audio stream. However, we may only need to protect the authenticity of the corresponding video stream. Furthermore, the rights of a conference attendee may vary according to the particular data flow. For example, designated attendees may be able to present on an audio flow while others can only receive the audio flow.

The security requirements for protecting multimedia conferences depend on the role of the authorized conference attendee. We will say a conference attendee is a *presenter* when the attendee generates information that is sent to other attendees, whereas a *recipient* receives information. The security services desired by presenters and recipients vary depending on the nature of the conference. A list of some of the most common services from which to select follows. (Note that data confidentiality usually takes precedence over other security requirements.)

Presenter

Data Confidentiality. The presenter’s data is protected against unauthorized disclosure.

Identity Confidentiality. The identity of the presenter is protected against unauthorized disclosure. Identity confidentiality can be maintained with respect to all others (as required in voting protocols) or only with respect to non-participants of the conference (as might be required between a reporter and a source or between a buyer and a seller). Traffic analysis must be considered when protecting the identity of the presenter.

Non-repudiation of Receipt. Unforgeable proof of receipt of data by the recipient. Proof of receipt, including time of receipt, can be proven to a third party such as a court of law.

Repudiation of Transmission. Information presented in a conference can not be substantiated.[1] Laws of many countries preclude the contents of a private conversation to be used against them in a court of law. Therefore, parties to a conference have an expectation that the information they exchange will not be used against them [R93]. (Note that identity confidentiality may provide repudiation of transmission.)

Availability to Present. A presenter is not denied from contributing to a conference in accordance with the conference policy.

Recipient

Identity Confidentiality. The identity of a conference recipient is protected against unauthorized disclosure.

Origin Authentication. A recipient can authenticate the origin of data.

Data Integrity. The recipient can detect an unauthorized modification of data. Strict protection against message-stream modification may be required for some flows (e.g. groupware application flows). Weaker requirements such as a hybrid connection/connectionless integrity service may be adequate for real-time video streams. In the hybrid integrity service, the integrity of individual packets is protected; however, some packet loss in transit is acceptable provided the packets are played back in order.

Non-repudiation of Transmission. Unforgeable proof of shipment. Proof of shipment, including time of shipment, can be proven to a third party. Non-repudiation of a sequence of information may need to be shown.

Availability to Receive. A recipient is not denied from receiving information in accordance with the conference policy.

[1] The need for repudiation of data is attributed to Michael Roe and Russ Housely.

Some of the preceding security services can also be required for protecting flows that control site-specific media devices. Controlling access to conferences is also an important security requirement. We can specify conference access policies using the access matrix model [HRU76].

Access Control

> The access control matrix specifies the *rights* a *subject* may have with respect to an *object*. A subject in a conference consists of a conference site or attendee at a site. Objects can be a data flow or a media device. (An example of a media device is a camera or a microphone.) The rights a subject may have with respect to an object include writing (e.g., presenting information to the flow or controlling a media device), reading (e.g., receiving a flow) and owning the object. A subject that owns an object may participate in assigning rights to other subjects (e.g., add read and write rights to an attendee for a particular flow). Rules for assigning rights may be non-trivial. For example, attendees working for different companies may not be authorized to confer between themselves although they may be allowed to join in the audience. Other complexities may exist in the different models for admitting attendees to a flow. For example, assigning read and write rights might require a consensus from multiple conference owners.

3. Scalability of Security Services

We now discuss how well today's security technology scales in providing the various security services. The design of a protection architecture depends on the size of a conference. Conferences are sometimes described by the number of attendees [S93]. A small conference is highly interactive and consists of only a few participants. A medium sized conference might consists of hundreds or thousands of participants and is typical of an interactive seminar with many recipients but relatively few presenters. A large conference consist of hundreds of thousands or millions of participants; an example is a pay-for-view television broadcast.

Security services that require unique processing of information for each recipient do not scale well to medium and large conferences. This is particularly true for multimedia conferences imposing real-time communication constraints (e.g., bounded communication delays). Distribution of session keys is an example of a per-recipient task. This distribution is required for providing data confidentiality using symmetric cryptography. Schemes for distributing keys *before* a conference can help alleviate this problem. However, these schemes have limitations in responding to changes in access permissions during a conference. For example, when an conference moves to executive mode (i.e., a more restrictive session), revoking read privileges requires a new session key to be distributed to the new set of authorized participants.

Session keys may also be used to perform data origin authentication. However, the use of session keys for data origin authentication has an additional scaling problem since representations of authentication digests need to be processed on a per-recipient basis for each packet sent. (For example, authentication digests may be encrypted using symmetric cryptography under a secret key shared between the sender and each recipient.) This problem can be alleviated by using asymmetric cryptography when preparing an authentication digest .[2] The task of key distribution is effectively distributed to the recipients. Each recipient obtains and verifies the presenter's public key.

4. Security Infrastructure for Multimedia Conferencing

Significant research and development is needed to create the infrastructure for wide scale deployment of secure multimedia conferencing over unprotected packet-switched networks like the Internet and the National Research and Education Network (NREN). Much of the foundation for the security

[2] The designer must choose a redundancy function that has the mathematical properties sufficient to support the integrity

infrastructure either exists or is being developed to support other applications. Some of the components needed to protect multimedia conferencing are as follows:

Trusted Computing Bases. The integrity of conferencing applications is protected by building on trusted operating systems and trusted hardware platforms.

Key Management. A key management infrastructure is needed that will enable the creation of trusted paths between heterogeneous key management systems and heterogeneous trust hierarchies/models.

Conference Control Servers. Conference control servers are needed that will enforce access control policies to information flows and media devices.

Secure Communication Protocols. Communication protocols are needed that will support the protection requirements for real-time communications. For example, work is in progress to augment an experimental Real-time Transport Protocol (RTP) [SC93] to support security features that can be used to satisfy the requirements of a hybrid connection/connectionless integrity service. Secure communication protocols at other layers (e.g. network layer) might also be important.

Network Service Guarantees. Network service providers need to provide guarantees for use of bandwidth. This is because real-time traffic is highly susceptible to communication delays leading to a denial of service. Common carriers operating in a muti-vendor multi-user packet switched environment will use network layer security to enforce service guarantees.

Identity Confidentiality Servers. Identity confidentiality servers are needed to protect the identity of conference participants from traffic analysis. As with all security services, the value of using the service should justify its cost.

5. Summary

The demand for security services in multimedia conferencing will rapidly grow as multimedia conferencing becomes widespread. We discussed security services for protecting multimedia conferencing, scalability issues for implementing these services and some of the security infrastructure that is needed to meet the demands for protecting conferences.

Acknowledgements

We thank Celeste Anderson, Stephen Casner, Kathleen Madden, Clifford Neuman, Michael Roe, Eve Schooler, and Joe Touch for their helpful discussions and reviews.

REFERENCES

[B93] B. Bialick, “Pre-Message Security Protocol (PMSP) Overview”, Federal Digital Signature Symposium, National Institute of Standards and Technology, Gaithersburg, MD, February 1993.

[CD92] S. Casner and S. Deering, “First IETF Internet Audiocast,” *ACM SIGCOMM Computer Communications Review*, Vol. 22, No. 3, July 1992.

[HRU76] M. Harrison, W. Ruzzo, and J. Ullman. “Protection in operating systems,” Comm. ACM 19, No. 8 (Aug. 1976), 461-471.

[L93] J. Linn, “Privacy Enhancement for Internet Electronic Mail: Part I -- Message Encipherment and Authentication Procedures,” Internet Working Group, RFC-1421, February 1993.

[R93] Michael Roe, “Real-time Communications”, PSRG Memo, April, 1993.

[S93] E. Schooler, “The Impact of Scaling on a Multimedia Connection Architecture”, to appear in the *ACM Journal of Multimedia Systems*, 1993.

[SC93] H. Schulzrinne and S. Casner, “A Transport Protocol for Real-Time Applications”, Working Draft, Internet Engineering Task Force, Audio-Video Transport WG, May 1993.

[SCP91] E. Schooler, S. Casner, J. Postel, “Multimedia Conferencing: Has it come of age?,” Proceedings 24th Hawaii International Conference on System Sciences, Vol.3, pp.707-716 (January 1991).

[SG93] S. Stubblebine and V. Gligor, ‘‘Protocol Design for Integrity Protection,” Proc. IEEE 1993 Research in Security and Privacy, IEEE Computer Society Press, Oakland, California, May 1993.

[X.509] CCITT Recommendation X.509 (1988), ‘‘The Directory - Authentication Framework”.