

Key Management for Secure Multicast Group Communication in Mobile Networks

Thomas Kostas, Diane Kiwior, Gowri Rajappan, and Michel Dalal

Abstract—This paper describes the implementation of a hierarchical key management system to provide secure multicast communications in mobile network environments. By using this hierarchical system, both efficiency and security are improved and a highly scalable system is created.

Index Terms—secure multicast communications, secure mobile communications, hierarchical key management.

I. INTRODUCTION

The future of military communications is rapidly converging on visions of connectivity such as the Global Information Grid (GIG) and Joint Tactical Radio System (JTRS). In these visions, combinations of wireline and wireless communications are used to provide many to many collaborative communications between troops and sensors in theater, and command and control systems located thousands of miles away. Large numbers of entities participating in these communications warrant using efficient protocols such as multicast in order to reduce network congestion.

When members of a multicast group need to receive the same information securely and are allowed to dynamically join or leave the group, security entails not only distribution of a secret among many but may also be concerned with confidentiality of information as the membership changes. Military communications require different levels of security based on policies governing the shared information. When considering strict secrecy policies, it is important that a new member to the multicast group not be able to decode previous information that was transmitted (backward confidentiality) and a current member who leaves (or is ejected) not be able to decode future information that will be transmitted (forward confidentiality). As the level of confidentiality is relaxed, the amount of forward and backward confidentiality is also eased. Secure group communications also seek to prevent collusion,

in which a set of members exchange information to gain additional unauthorized access. Another feature important to secure groups is containment; compromise of one member should not compromise the entire group.

Approaches that work in unicast transmissions, such as SSL and VPNs do not extend to a multicast group. VPNs do support multicast but only by unicasting the data to each wireless VPN client, effectively removing the bandwidth efficiency of multicast.

The problem of secure group communication has been the subject of much recent research with both an IRTF research group, the Group Security Research Group (GSEC) and an IETF Working Group, Multicast Security Working Group (MSEC) addressing the issue. Most of the work in these groups has been directed toward wired networks but the issues therein identified also apply to wireless networks. Features that have been identified as necessities of a key management system for secure mobile multicast groups include scalability, data confidentiality, data integrity, source authentication, forward and backward confidentiality, collusion resistance, and compromise resistance.

While the following features enhance the performance of any group security scheme, they are particularly important to compensate for the constraints of mobile wireless networks: minimal messaging bandwidth usage, minimal security related computation, storage efficiency of keys, and low latency for rekey messages.

[WCetal] classifies secure multicast protocols into three categories: centralized flat schemes, distributed flat schemes, and hierarchical schemes. [DMS] considers the security and scalability issues of each category with the following analysis. Centralized flat schemes do not scale since one change affects all members, known as the ‘1 affects n’ scalability problem. Distributed flat schemes are vulnerable to collusion attacks. Hierarchical schemes using a hierarchy of keys also suffer from the ‘1 affects n’ scalability problem. However, protocols with a hierarchy of nodes responsible for key distribution, but not data distribution, address the scalability and security risks of the other schemes.

II. SYSTEM OVERVIEW

[HCM] proposes a hierarchy of nodes in a two-tier hierarchical secure multicast protocol described as follows. The domain is divided into administratively scoped areas. Confidentiality of multicast data is provided by encrypting the

This work was supported by DARPA under contract N66001-00-C-8011.

T. Kostas is with Northrop Grumman Corporation, 55 Walkers Brook Drive, Reading, MA 01867 USA (phone: 781-205-7581; fax: 781-942-0636; e-mail: tkostas@northropgrumman.com).

D. Kiwior, is with Northrop Grumman Corporation, 55 Walkers Brook Drive, Reading, MA 01867 USA (e-mail: dkiwior@northropgrumman.com).

G. Rajappan is with Nevelex Corporation, Cambridge, MA 02138 USA. (e-mail: gowri@nevelex.com).

M. Dalal is with Nevelex Corporation, Sunnyvale, CA 94086 (e-mail: dalal@nevelex.com).

data packets with a group key. Distribution of the group key is handled by a Domain Key Distributor (DKD) and multiple Area Key Distributors (AKDs). The DKD and AKDs constitute a multicast group, the All-KD-group. The DKD generates the group key and multicasts it to the All-KD-group. Each AKD, in turn, distributes the group data key to all group members in its area. The data path need not pass through the AKDs, Although the DKD remains a single point of failure, the second tier AKDs address the problem of scalability for a single distributor.




Figure 1: Multilayered hierarchical key distribution control plane.

We extend this notion of a hierarchical secure multicast system to include multiple layers of key distributors as shown in Figure 1. Multiple layers of AKDs allows placement of AKDs around wireless links. This increases confidentiality by decreasing the time necessary to rekey nodes located in an area where a new member has joined or left the group, or where a breach in security has occurred.

A data encryption key, referred to as a Traffic Encryption Key (TEK), is used to encrypt the data as it is sent from a source. To ensure forward confidentiality and backward confidentiality, each time a member joins or leaves, a new TEK must be generated and distributed securely and efficiently to all current authorized group members. This provides data confidentiality over wireless channels with strength of security related to the key size and encryption algorithm. To distribute the TEK securely, it is encrypted with an area local encryption key, known as a Key Encryption Key (KEK). In Figure 1, this is represented by KEK1 through KEK5.

When AKD and end node (also called mobile host (MH)) mobility are considered, the TEK is changed periodically. This makes it unnecessary to propagate requests to change the TEK to the DKD. Therefore, nodes joining the group are able to decrypt data only after they receive the next updated TEK. Additionally, nodes leaving a group are able to decrypt data until the next TEK is updated. The periodicity of the changing of the TEK is based on policy governing data confidentiality.

The set of nodes participating in the secure group communications control plane is segregated using clustering techniques. Two groups are formed within each cluster, one

group is composed of nodes that are eligible to become AKDs and the other is composed of nodes that are allowed to elect the AKD. The selection of different AKDs is accomplished using an asynchronous leader election algorithm which converges to the election of a single leader although the membership of the cluster of nodes may be changing due to their mobility. Additionally, new leaders may be selected periodically to prevent denial of service attacks from disrupting the key management infrastructure.

Finally, various mobile node handoff mechanisms are considered to account for asymmetries in information, handoff efficiency, and different security policies. These include parent key distributor controlled handoffs, child node controlled handoffs, child node assisted handoffs, and parent key distributor assisted handoffs. Each of these handoffs can exchange either a single or multiple nodes between AKDs. In a parent key distributor controlled handoff, the parent key distributor determines that the child node should be associated with another key distributor based on attributes the parent key distributor is monitoring. In a child node controlled handoff, the child node determines that it should be associated with another key distributor based on attributes the child node is monitoring. In a child node assisted handoff, the child node reports attributes to the parent key distributor. Then, the key distributor determines that the child node should be associated with another key distributor based on these attributes. In the parent key distributor assisted handoff, the parent reports attributes to the child node. Then, the child node determines that it should be associated with another key distributor based on these attributes.

III. IMPLEMENTATION

There are three components to the implementation. At the core is the Communication Module (CM) used for key and data distribution. The Authentication and Authorization Protocol (AAP) is an application that uses the services of the Communication Module. The third component is a Policy Engine (PE) that is responsible for maintaining the logical hierarchy of the system, leadership election, and handling dynamic events in the system. In this section we will describe these three components in detail, and also provide a description of the demonstration platform.

The CM is responsible for providing a reliable and transparent end-to-end transport infrastructure. It provides Application Programming Interfaces (APIs) for utilizing its communication services. The API is used, for instance, to multicast packets to a group of nodes, or broadcast packets over a directed tree. The API also facilitates insertion of special headers for application specific control messaging or performance measurement. This multicast communication services infrastructure is implemented over IP multicast. The API services need to be transparent to the underlying physical infrastructure (within the performance constraints imposed by the physical links). For instance, a system with a mix of wireline, short range wireless, satellite, and line of sight optical

links can be envisioned.

AAP uses the APIs provided by the communication module for node authentication and key distribution. It abides by the policy decisions of PE on node admission, handoff, and teardown. It handles authorization events. Upon PE directive, AAP enables Dynamic Coalitions (DC) and hence helps reduce messaging overhead. DC is a migration event of a group of nodes from one AKD to another without having to trigger a local rekeying.

PE is responsible for the logical system integrity. It is responsible for maintaining the system hierarchy. It queries AAP about trustworthiness of the nodes when it needs such information for policy decisions. By use of the CM APIs it is capable of setting up and tearing down a logical hierarchy of nodes for key and data distribution. PE is also responsible for policy decisions on Area Key Distributor (AKD), Domain Key Distributor (DKD), rekeying, node admission and teardown.

The demonstration platform consists of a number of computers connected over wireline links and numerous mobile nodes. The wireless links and node mobility are simulated over conventional wireline links.

IV. PERFORMANCE

A number of performance measures are possible – we are particularly interested in three: (1) Efficiency (2) Security, and (3) Coverage. Efficiency is a measure of the messaging requirement to enable a specified level of security. Security is hard to quantify. We use a proxy measure that covers one aspect of security – rekeying delay. Specifically, if there was a known breach in the system, the amount of time it takes for the system to recoup and isolate the offender via rekeying is a measure of resilience of the secure infrastructure. Coverage is the proportion of time a participating node is a part of the secure infrastructure. Low coverage is symptomatic of a number of issues – mobility and roaming, selection criterion for AKD, individual link characteristics (high latency links) etc. These three performance measures are not mutually exclusive and hence any system specification involves a trade-off between these measures. We are interested in useful performance regions for realistic operating scenarios.

In baseline rekeying, any change in the system hierarchy triggers rekeying of the entire system. The messaging requirement for baseline rekeying scales linearly with mobility. In comparison, the messaging requirement for some of the investigated alternate tactics does not increase quite as rapidly. Hence these alternates are more efficient as a function of mobility. Efficiency is also influenced by factors such as rekeying delay requirements, and coverage requirements. The demonstrations will characterize the specific trade-offs.

Security is a three-tiered issue. (a) Strength of admission and authentication policies, (b) Strength of authorization policies (authenticated nodes should not be able to perform actions that exceed their authority), and (c) Robustness (detection of breach and rapid recovery from a known breach). (a) and (b) are determined directly by the type of

authentication and authorization policies selected. But (c) is influenced by the proposed rekeying strategies. Hence we are interested in studying (c). In the demonstration we will study rekeying delay, and hence the time to recovery from a known breach, as a consequence of the specific rekeying policies and techniques chosen. Breach detection will be saved for a different line of investigation, owing its qualitative nature.

Coverage is a crucial issue. Desire to maximize coverage has implications on such system level issues as AKD selection, rekeying delay specification etc. So this performance metric will be studied in the context of Efficiency and Security.

V. CONCLUSIONS

We have developed a key management system for secure multicast group communications in mobile network environments. A demonstration platform is implemented. We use the demonstration platform to study useful tradeoffs of the system performance metrics. The flexibility of our system allows it to be more efficient, scalable, and secure than alternatives. Future areas of research include incorporation of non-repudiation (e.g., through the use of digital signatures) in the system. Also of interest is the inclusion of multiple DKDs to support highly secure heterogeneous collaborative environments.

REFERENCES

- [DMS] L. R. Dondeti, S. Mukherjee, and A. Samal, "Survey and Comparison of Secure Group Communications Protocols," Technical Report, University of Nebraska-Lincoln, 1999.
- [HCM] T. Hardjono, B. Cain, and I. Monga, "Intra-Domain Group Key Management Protocol," Internet draft, draft-irtf-smug-itragkm-00.txt, September 2000, Work in Progress.
- [WCetal] M. Waldvogel, G. Caronni, D. Sun, N. Weiler, and B. Plattner, "The VersaKey Framework: Versatile Group Key Management," JSAC Special Issue on Middleware, 1999.