

Secure Group Management in Large Distributed Systems: What is a Group and What Does it Do?

John McHugh

CERT / SEI
Carnegie Mellon University
Pittsburg, Pennsylvania

J. Bret Michael

Computer Science Department
Naval Postgraduate School
Monterey, California

Abstract

The secure management of groups containing thousands or possibly hundreds of thousands of members with very high rates of membership turnover is claimed to be a critical need for high confidence networking. Among the needs mentioned are the ability to ensure that former group members can no longer obtain access to group materials and to prevent new members from accessing material distributed to the group prior to their entry. Suggestions made in this area exhibit a strong bias towards cryptographic techniques and key management to realize these goals, pointing out the weaknesses in currently available techniques. The purpose of the present paper is to examine some of the assumptions that appear to be implicit in these suggestions. An examination of group function and behavior might indicate alternative ways to manage large groups securely. We note that the call for ever more complex technological solutions to problems that may be sociological in nature continues a disturbing (and largely unsuccessful) trend that began in the pre-TCSEC days and that continues into the present.

Key words: cryptography, key management, network architecture, network management

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. 1999 New Security Paradigm Workshop 9/99 Ontario, Canada © 2000 ACM 1-58113-149-6/00/0004...\$5.00

1 Group Structure and Communication Patterns

*

A recent DARPA white paper titled "Research Challenges in High Confidence Networking" [1] discusses the secure management of groups containing thousands or possibly hundreds of thousands of members with very high rates of membership turnover[†]. Although the available sources imply that the management of large, flat, groups is necessary to support certain critical missions, no examples are given and no rationale is provided for the necessity of the approach.

We begin by examining the rationale for group formation and the ways in which group members may interact. We see two fundamental modes of group organization, hierarchical and flat.

1.1 Hierarchical Group Structures

Hierarchical group structures, by their very nature, avoid much of the large group membership problem because the organizational rules of discourse preclude direct communications that cross many levels of the hierarchy. This is the case with the military chain of command and most corporate or institutional structures. It reaches its limiting case with the cell structure used by some underground, terrorist, and resistance groups. We suspect that external considerations

*Funding for preliminary work on this topic was provided by the Center for INFOSEC Studies Research (CISR), Naval Postgraduate School. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright annotation thereon.

[†]DARPA BAA 99-33[2] contains further requirements in this area, indicating that interest in the topic is still current.

keep the size of any group within the hierarchy manageable by adding additional levels as necessary.

In most cases, controlling membership within a given subunit of the hierarchy is tractable and does not present the kinds of problems associated with the large flat groups discussed below. Although communications from the root of the hierarchy to all members of all subunits of the hierarchy are not unknown, they are typically either rare or do not involve particularly sensitive material. They are seldom both sensitive and so time critical as to require bypassing of the normal distribution channels. We note in passing, that hierarchical groups may be easily organized to contain confidential material since the structure is often made to reflect need-to-know or mission-oriented factors. Properly managed, the cell structure can conceal the size and structure of the group from the membership.

1.2 Flat Group Structures

In a flat group there is a potential for each member to communicate with any or all others. The extent to which this potential is realized depends on the nature of the material being handled and the policies that apply to the group. We posit that there are four primary modes of member interaction within a flat group. We will term these many-many, many-few, few-many, and few-few. The limiting cases of many and few are all and one. We are interested in steady state behavior and will ignore startup and shutdown cases. The behavior of receivers, in particular, their joining and leaving groups, can be modeled as a continuous-time stochastic process.

1.2.1 Many-Many Interaction

In the limiting case of many-many interaction, every member is equally likely to emit a message in a given interval and each message needs to be observed by all other members. An example is group behavior during an auction where any bidder may raise the previous bid. Even in this case, there is a distinguished member of the group, the auctioneer, who must see all the bids and make an award to the winner. We suspect that this kind of interaction is unlikely for very large groups requiring secure membership with instantaneous addition or revocation of membership.

1.2.2 Many-Few and Few-Many Interaction

In a many-few interaction, most members are likely to originate communications, but all the communica-

tions are directed to a few (possibly one) members, for example, the reporting of sensor information to a common repository or processor. In the converse situation, few-many interactions, messages are broadcast from a few sites, possibly one, to the rest of the membership. An example is the dissemination of weather forecasts.

Both of these forms are believed to be relatively common. We can envision a number of cases in which secure group membership is required with these patterns of communication. For instance, some styles of auction, such as the Dutch auction, are characterized as many-few in the bid process and few-many in bid acknowledgment. In this case, the auctioneer can control known rogues by declining to accept their bids and prevent both buyers and sellers from manipulating the auction to their own or someone else's advantage.

1.2.3 Few-Few Interaction

Few-few interactions result when individual group members direct communications to other individuals or to small subsets of the group. Examples are control messages between individual routers in a network or emails among individuals within an organization. It is not clear that cryptographically enforced group membership with the sharing of a single key is appropriate for this communication pattern.

2 Group Membership and Policy Considerations

In all cases, the group membership policy says that communications between group members are permitted and that communications between members and non members are forbidden, at least under the membership key. The policy can be enforced by the sharing of a single, symmetric encryption key among all group members and requiring its use on all communications. A problem arises when the membership of the group changes. This can be viewed in several ways.

The most restrictive view holds that either addition or removal of a member for any reason whatsoever creates a new group with nothing in common with the previous group, necessitating a new key. For example, a seller (i.e., merchant) in a Dutch auction can be removed from the group by the auctioneer if the seller violates the trust of the buyer.

More permissive views are also possible. For example, new members may be allowed (or even required) to observe the history of the group and may be simply given the current key, some or all past keys, and access

to the group archives. Similarly, members departing in good standing may be removed from the group distribution mechanisms and trusted not to attempt to access materials which will continue to be encrypted under a key that they possess until the next scheduled key change. If key changes occur relatively frequently, the risk of exposure may be deemed acceptable.

It is our belief that the purposes and policies surrounding the formation of large groups need to be addressed more carefully prior to describing requirements for dynamically controlling group membership. In addition, the availability and applicability of mechanisms other than cryptography for controlling group membership should be considered. This, in turn may lead to considerations of network architecture and communications mechanisms.

3 Access Control versus Secrecy

We believe that the white paper operates on the assumption that the group key serves as both an access control mechanism and as a secrecy preservation mechanism. This would need to be the case only if all communications were broadcast universally and always available to both group members and non group members. This is unlikely for many systems, and is particularly unlikely to be so for many-few and few-few interactions because universal flooding of the communications would consume excessive resources. It may be approximately the case for many-many interactions, depending on the nature of the delivery mechanism. Only in the case of certain few-many distribution mechanisms such as direct satellite broadcast, the universal observer assumption may be satisfied.

In most cases, access control will be partially provided by the delivery mechanisms. Many-few systems can effectively remove a group member on the many side by ignoring traffic from that member on the few side. Few-many systems can remove a member on the many side by removing that member from the distribution list in cases where a communications mechanism other than broadcast is used. This is discussed further in the next section.

The use of key change, even when done to remove a specific member, is problematic as it assumes that the problem member or members have been identified. For groups of hundreds of thousands, it may be impossible to know with certainty that all members are trustworthy and will not compromise the group key (or material protected by it). If the members are human users or systems from which humans or corrupt software can extract the key, it is quite likely that the

group key will fall into the hands of outsiders and that key change will not be an effective mechanism for preserving confidentiality or for controlling current group membership since anyone with the key is a *de facto* group member. It is even more likely that material protected under the group key will pass from group members to non group members in the clear. We suspect that large groups are inherently risky, especially for broadcast delivery mechanisms and that key confidentiality cannot be easily assured and should not be equated with information confidentiality in any event. For these reasons, we posit that abrupt key changes—via middleware [6], key graphs [4], computational grids [5], or some other type of mechanism—may not be a particularly effective means for either providing confidentiality or managing group membership in large groups.

4 Assurance Issues

For the purpose of discussion of assurance issues, assume that sensitive information will be transmitted across a wide-area network (WAN) via one or more IP multicast mechanisms. Let's first consider assurance issues related to an IP host joining a multicast host group associated with a specific multicast transmission session. In order to join a multicast host group, a host must send a request to the multicast host group, which is then forwarded to the local area network and possibly the routers located between the requesting host and the multicast host group. Authentication of the requester needs to take place for each multicast session: group membership expires at the end of a session. Authentication does not need to take place on the multicast routers: the multicast routers only need to access the list of member hosts for each group for which there is one member on a subnetwork. Thus, due to the dynamic binding of IP host group addresses to interfaces at local area networks, authentication need be handled on a few-many basis (i.e., at the source and destination, not the intermediary nodes). The two addresses that must be authenticated are the hardware multicast address and the group address.

Multicast routers use the Internet Group Management Protocol (IGMP) to learn the existence of host group members. Each host maintains its host group memberships. At the data link layer, the router queries the subnets to which it is linked and receives reports from the IP hosts listing their group memberships. The information exchange is few-few or few-many due to the fact that each host sends back only

one reply per active host group.

Membership in a group is not persistent: membership ends when the multicast session closes. A member can leave a group at any time, or may block incoming multicast datagrams from specific addresses. Due to the short-lived nature of multicast sessions, the use of cryptography as an access control mechanism might be considered overkill. Moreover, the routers prune the spanning tree, eliminating branches on which there are no members of a multicast group. Thus, the branch cannot be reappended to the tree unless the IP host requesting membership from that subnet is permitted to join the group. Once again the nature of the interaction among hosts and among hosts and routers is few-few or few-many, not many-many.

Assumptions about the distribution of multicast group members through the network is more of a concern from the perspective of availability than that of confidentiality. The choice between dense-mode and sparse-mode routing protocols will affect the ability of a network to handle large fluctuations in the level of message traffic, or alternatively, how much bandwidth will go unused.

The formation of multicast groups is also driven by the goals of the members. For example, suppose multicasting is overlaid on a switched multimegabit data service (SMDS). This is an example of a connectionless service, that is, all packets sent on the network are available for reception by all nodes (i.e., receivers). However, an issue of equity will arise in the case in which not all nodes receive all IP multicast traffic. If membership in multicast groups is handled at the IP layer, then all of the nodes in the network will be charged a usage fee for the SMDS, even in the event that one or more of the nodes is not a member of an active multicast group. An alternative is to create multicast groups at the SMDS level, mapping IP multicast groups to the SMDS multicast groups.

Equity and quality-of-service issues can affect policy decisions made about group members in the context of overlaying IP multicast on frame relay, ATM, and other types of network technology. For instance, in a connection-oriented environment, such as frame relay (an example of few-many), cryptography can play an important role in protecting confidentiality, but the focus will be on end-to-end encryption across permanent virtual channels.

5 Thinking in Terms of a Few-Few Paradigm

In this section, we discuss a specific protocol that exemplifies a few-few paradigm, followed by a discussion of the protocol in terms of a proposal for a different security paradigm than that proposed by the DARPA-sponsored working group.

5.1 ATM PNNI Protocol

One of the key technologies at the enterprise level for tying connectionless networks (e.g., Gigabit Ethernet, Token Ring) together is asynchronous transfer mode (ATM). ATM provides for high-speed virtual circuits for which quality of service can be specified explicitly for each circuit.

The ATM Private Network-Network Interface (PNNI) protocol [3] provides for few-few interaction among network switches. The few-few interaction results from an hierarchical organization of groups of switches, called peer groups. A switch only exchanges routing information with the other switches in its peer group, except for the peer group leader which also belongs to the parent peer group in the hierarchy: it acts as an interface across two levels.

There is no need for routing information to be sent directly from the bottom of the hierarchy to the very top of the hierarchy because each peer group maintains a summary of the routing information of the child peer groups. In essence, a peer group functions as a single logical switch.

PNNI supports dynamic exchange of information about the ATM switches (i.e., network nodes). However, information is only distributed by a node to its peer group when there has been a significant change in the state of the PNNI topology: this requires flooding and synchronization, via PNNI topology state packets, within a peer group, not across the entire hierarchy. After the database synchronization among the members of the peer group is complete, the topology information is summarized and passed up the hierarchy in summary format by the peer group leaders, from child to parent node. The scope of any logical group node is its own level and all levels above it in the hierarchy, much the same way scope is defined using the subnet masks of the Internet Protocol.

5.2 A Counter Security Paradigm

The purpose behind the hierarchical, peer-to-peer organization of the switches is to provide the designer

of the network with the ability to manage the complexity of the structure and the size of the composite network (i.e., composition of all of the individual switches). Quality-of-service (QOS) policy for an ATM network can be specified explicitly within the topology database. Network QOS policy encompasses more than just performance—it is also a specification of security policy.

In the security paradigm that we espouse, the switching structure should drive security policy, rather than the other way around. That is, in the case of ATM switching, it is possible but not desirable to construct a many-many relationship by building a one-level (i.e., flat) switching hierarchy consisting of one peer group, even in the case in which an efficient and trusted means exists for both the global and dynamic management of cryptographic keys.

The features of the PNNI protocol are congruent with some of the foundational engineering principles of computer security, such as the use of

- Encapsulation (i.e., data hiding): there is a well-defined interface between peer groups.
- Constrained vocabulary: only specific types of information, that is, information about the network topology, can be exchanged from one peer group to another and hence between levels in the hierarchy.
- Consistency of security policy: network security policy, as a QOS parameter, is propagated up the hierarchy.

Therefore, rather than reinventing PNNI to conform to a baroque architecture characterized by many-many relationships—this is not needed because the intent behind ATM is to provide a high-speed link among hubs in a network—it would be more fruitful in our opinion to develop new security paradigms that address few-few relationships and the tradeoffs among security and non-security (e.g., performance, maintainability) issues of system design.

Our paradigm is by no means “new;” it is based on existing engineering principles that have evolved over time within the computer security community.

6 Conclusion

We have suggested that the management of large secure groups is not as simple as solving a key management problem (which may not be simple at all). In addition to addressing key management, secure group

membership needs to take other factors into account. These include the nature of the group interaction including transmission and reception patterns, the external policies concerning member trust and access to prior or future group material, the infrastructure on which the group is built (i.e., broadcast, multicast, or point-to-point), and the nature of the participants. Assurance issues need to be considered also, both with respect to the participant nodes and the network internals. Some of these factors may ease the cryptographic and key management aspects of the problem while others may complicate them. Group behavior is a sociological issue. We are seeing increasing evidence that security failures are as much attributable to human failures as to technology failures. Future attempts to provide security through technological solutions are no more likely to succeed than past attempts.

7 Comments from the Audience

Two presentations were made based on this manuscript: one at the University of Toronto and the other during the workshop. The following is a list of comments that were germane to the issues we raised in the previous sections:

- Many-to-many communications are, in reality, the composition of many-to-few and few-to-many communications. In addition, if one models members of a group as Von Neumann machines, the group members cannot listen to more than one other member at a time, though they can transmit to more than one. Thus, true many-to-many communication is either (i) non-simultaneous or (ii) analog or parallel.
- The phrase ‘rapidly changing group membership’ connotes scenarios in which members of a group leave and arrive faster than the cryptographic key management operations can support auditing and deleting users.
- The issue may not be *who* is in the group. Rather, the question to be decided may be the following: given that one knows who is in the group, what should be the security policy for the group?

During the workshop, a long discussion on the value and quality of information—and its use in decision making—ensued. However, the discussion drifted away from the issue at hand: what is the nature of communication of sensitivity information in large groups?

8 Conclusion

We have suggested that the management of large secure groups is not as simple as solving a key management problem (which may not be simple at all). In addition to addressing key management, secure group membership needs to take other factors into account. These include the nature of the group interaction including transmission and reception patterns, the external policies concerning member trust and access to prior or future group material, the infrastructure on which the group is built (i.e., broadcast, multicast, or point-to-point), and the nature of the participants. Assurance issues need to be considered also, both with respect to the participant nodes and the network internals. Some of these factors may ease the cryptographic and key management aspects of the problem while others may complicate them. Group behavior is a sociological issue. We are seeing increasing evidence that security failures are as much attributable to human failures as to technology failures. Future attempts to provide security through technological solutions are no more likely to succeed than past attempts.

9 Future Directions

We are exploring the need for secure group management of systems comprised of large numbers of communicating objects, such as intelligent land and sea mines, and MicroElectroMechanical Systems (MEMS). For example, we are exploring the case in which there is a large population and a high degree of redundancy of MEMS. In this case, we believe that changing the key—for example, when a MEMS is captured—is less efficient than destroying the captured MEMS.

Acknowledgements

We thank the Heather Hinton and Steve Easterbrook for providing us with the opportunity to present our work at the University of Toronto prior to the workshop. The feedback we received from the audience during that presentation and at the workshop was used to revise the final manuscript.

References

- [1] Research Challenges in High Confidence Networking. White paper, Defense Advanced Re-

search Projects Agency, July 1, 1998.
www.dapra.mil/ito/research/hcn/problems.html

- [2] BAA 99-33 Proposer Information Pamphlet, DARPA, August 11, 1999.
www.darpa.mil/iso/ia&zs/IASPIP990811Final.html
- [3] Private Network-Network Interface Specification Version 1.0 (PNNI). Technical report, The ATM Forum technical committee, March 1996.
- [4] Wong, C. K., M. Gouda, and S. S. Lam, "Secure Group Communications Using key Graphs," in *Proceedings of the ACM SIGCOMM '98 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ACM, 1998, pp. 68–79.
- [5] Foster, I., C. Kesselman, G. Tsudik, and S. Tuecke, "A Security Architecture for Computational Grids," in *Proceedings of the Fifth ACM Conference on Computer and Communications Security*, ACM, 1998, pp. 83–92.
- [6] Waldvogel, M., G. Caronni, D. Sun, N. Weiler, and B. Plattner, The VersaKey Framework: Versatile Group Key Management, TIK Technical Report No. 57, Computer Engineering and Networks Laboratory, Swiss Federal Institute of Technology Zurich, September 1998.