Access Control and Privacy in Large Distributed Systems

Barry M. Leiner
Matt Bishop

March 7, 1986

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 86.6

Access Control and Privacy in Large Distributed Systems

Barry M. Leiner
Matt Bishop

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 86.6
March 7, 1986

Large scale distributed systems consist of workstations, mainframe computers, supercomputers and other types of servers, all connected by a computer network. These systems are being used in a variety of applications including the support of collaborative scientific research. In such an environment, issues of access control and privacy arise. Access control is required for several reasons, including the protection of sensitive resources and cost control. Privacy is also required for similar reasons, including the protection of a researcher's proprietary results.

This report describes a possible architecture for integrating available computer and communications security technologies into a system that meets these requirements. This architecture is meant as a starting point for discussion, rather than the final answer.

Work reported herein was supported by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Administration (NASA) to the Universities Space Research Association (USRA).

Access Control and Privacy in Large Distributed Systems

Barry M. Leiner
Matt Bishop
Research Institute for Advanced Computer Science
NASA Ames Research Center

March 7, 1986

1. BACKGROUND

As advanced computing capabilities are deployed and made available to large numbers of scientific researchers through both local and remote access, a number of issues of privacy and access control will arise. Among these issues are providing privacy of data (e.g., to protect early "sensitive" technical results), controlling access and use of valuable resources including both supercomputers and networks themselves, protecting remote resources in a shared environment (such as the experimental environment envisioned for the space station), and preserving privacy of communications such as electronic mail.

Systems such as the NSFnet are made up of various workstations, terminals, supercomputers, networks and users. The requirements for privacy and

access control in such distributed systems and the architectural approaches to provide them are not well understood. Many building blocks to provide privacy and access control are either available or will be soon, including link and end-to-end encryption mechanisms, secure computer systems, software encryption mechanisms for both file encryption as well as mail, etc. However, the requirements themselves are not understood; even less is known about how to connect systems in a way that satisfies the requirements.

For example, a scientific team of the future will work in an environment that has a local area network, several workstations with each shared by several scientists and support personnel, a main frame computer, and access to a wide-area network, allowing interconnection to other teams as well as special resources such as a supercomputer. Subteams working on a particular experiment will want to share certain data as well as have remote access to this data via the wide area network in order to coordinate activities between groups at different institutions working on a cooperative experiment. Remote access might also include the need to access and control remote experiments such as those on platforms like a space station. Some information (such as team planning documents) will be shared by the entire team, while other information will have to be kept private for personnel and similar reasons. The scientists will want to use electronic mail and similar facilities to share selected information with collaborators at other sites. They will also want to be able to use special resources when necessary. At those sites with shared resources, individual teams' information

must be kept separate and managed appropriately.

Furthermore, the privacy and integrity of this information must be maintained while in transit over the network. Appropriate access control mechanisms must limit access to the networks, the centralized resources and the local team resources. This all must be done in an environment consisting of the interconnection of a large number of local (probably private), regional, and wide area (probably public to a defined community) networks.

In the above example, we have emphasized the need to keep data private. Other requirements include assuring the integrity of the data (of particular importance in controlling remote experiments) and providing access controls to remote resources. It is clear from this example that, while the building blocks are probably available to satisfy most of the requirements outlined, it is necessary to create a system architecture out of these building blocks to satisfy these requirements.

2. A Possible Architecture

The following sections describe a possible architecture for satisfying the requirements for access control and privacy in large distributed systems. We start by summarizing a number of available technologies that can be brought to bear on this problem. We then describe how these technologies can be combined to satisfy the requirements.

It must be emphasized that the architecture described here is a starting point for discussions and does not represent a final design. We feel that developing such an architecture will require the cooperation and interaction of scientists and technologists in this area, and this paper is meant to start that discussion.

2.1. Available Technologies

There are two areas of technology providing measures of privacy and access control: communication protection and computer mechanisms. As we will see, this division is not clean. We now discuss each of the mechanisms available, and how they relate to an overall privacy and access control architecture.

2.1.1. Protected Wire Distribution Systems

A fairly common method for protecting communications in the past was simply to protect the wires that connect the various components. For example, terminal lines to a central computer can be protected to prevent an unauthorized user from connecting a terminal to the line.

2.1.2. Encrypted Links

Protecting wires is feasible in a small area. For remote access, encryption of the links is often used. As an example, link encryption boxes based on the National Bureau of Standards (NBS) Data Encryption Standard (DES) [1] are available to encrypt all data passing over the wire. These systems require a key to be inserted at both ends of the link, and then one of several modes of the DES [2] is used to encrypt the data. Thus, an unauthorized user is able to neither read nor insert data passing over the link. He might, however, be able to record and retransmit encrypted data unless adequate safeguards are built into the protocols.

2.1.3. End-to-End Encryption

Protected or encrypted links require that the nodes at the ends of the link be secure. Techniques of end-to-end encryption (E^3) have been developed to avoid the need to protect every node in a network, and also to be able to partition information between distinct hosts. These normally use a separate device attached between the network and the hosts or other network being protected. Data flowing through the box is examined to determine the desired destination and other information required for routing; this "network information" is passed through unchanged. The user data (and certain header information) is encrypted using a key known to the similar box at the destination. The data is not decrypted until it arrives at the destination E^3 device.

Two general techniques are available. The first technique establishes a community of devices, and all devices in the community share a single key. Therefore, any device (and its users) have access to all information being transmitted among devices in the community. The second technique uses more sophisticated key distribution algorithms and allows pairs of the devices to be keyed dynamically, with a single key being shared by only the source and destination devices. With such a system, no other devices would have access to the user data.

2.1.4. Restricted Computers

One of the easiest ways to protect a computer is to permit access only to a small set of trusted developers, users and maintainers. This is appropriate for computers dedicated to specific tasks such as packet switching and access control.

2.1.5. Isolated Computers

Similarly, user data can be protected by allowing electronic and physical access only to a set of authorized users. The computer would not be connected to a network and would not have dial-up lines. The physical access points would be protected.

2.1.6. Password Protection

Most computer systems rely on a combination of user identification codes (user names) and passwords to authenticate users. Protection of access to a machine and data within that machine thus relies on protection of passwords. A typical method for protection of passwords is to store them only in encrypted form. Other methods are being developed such as the system by Sytek that allows for a cryptographically verified challenge-response, with the challenge-response pair being changed each time it is used.

2.1.7. Distributed Operating Systems

Considerable research is being carried out in the area of distributed operating systems. Much of this work is in naming and accessing resources in a distributed system. Much of the difficulty with using the above technologies to provide the required access control and privacy lies in providing distributed and authenticated naming.[1] Research in distributed operating systems is expected to contribute to solutions to this problem.

[1]This topic was discussed at the *DARPA Distributed Operating System Workshop* held December 16-17, 1985, at Carnegie Mellon University

2.2. The Problem of Distributed Systems

As seen above, there are a number of technologies that can be brought to bear on the problems described in Section 1. The difficulties lie not in developing new technologies, but rather in developing, installing and using an architecture that can exploit those technologies.

Figure 1 shows a typical distributed system architecture. A user is at his terminal or workstation. In the case of a terminal, he is connected to a local host computer by direct line, dial-up line, or terminal local area network. The local host computer may be a true host (that is, it supports user applications) or may be used solely for the purpose of accessing a network.

Located at various places on the system are server devices. These include mainframe computers, supercomputers, data bases, file servers, and experimental facilities. In some cases, the devices are located on the same LAN as the user and in others they are located remotely.

Furthermore, users may access the system from several places. When in their offices, they may use the same LAN to which their primary host computer is connected. While on travel, they may have to access their primary host computer remotely.

The networks providing long haul communications fall into two basic classes. First, there are commercial networks which base charges on usage. Second, there are special purpose networks designed to support some limited community (such as the Defense Data Network does for the DoD) and which do

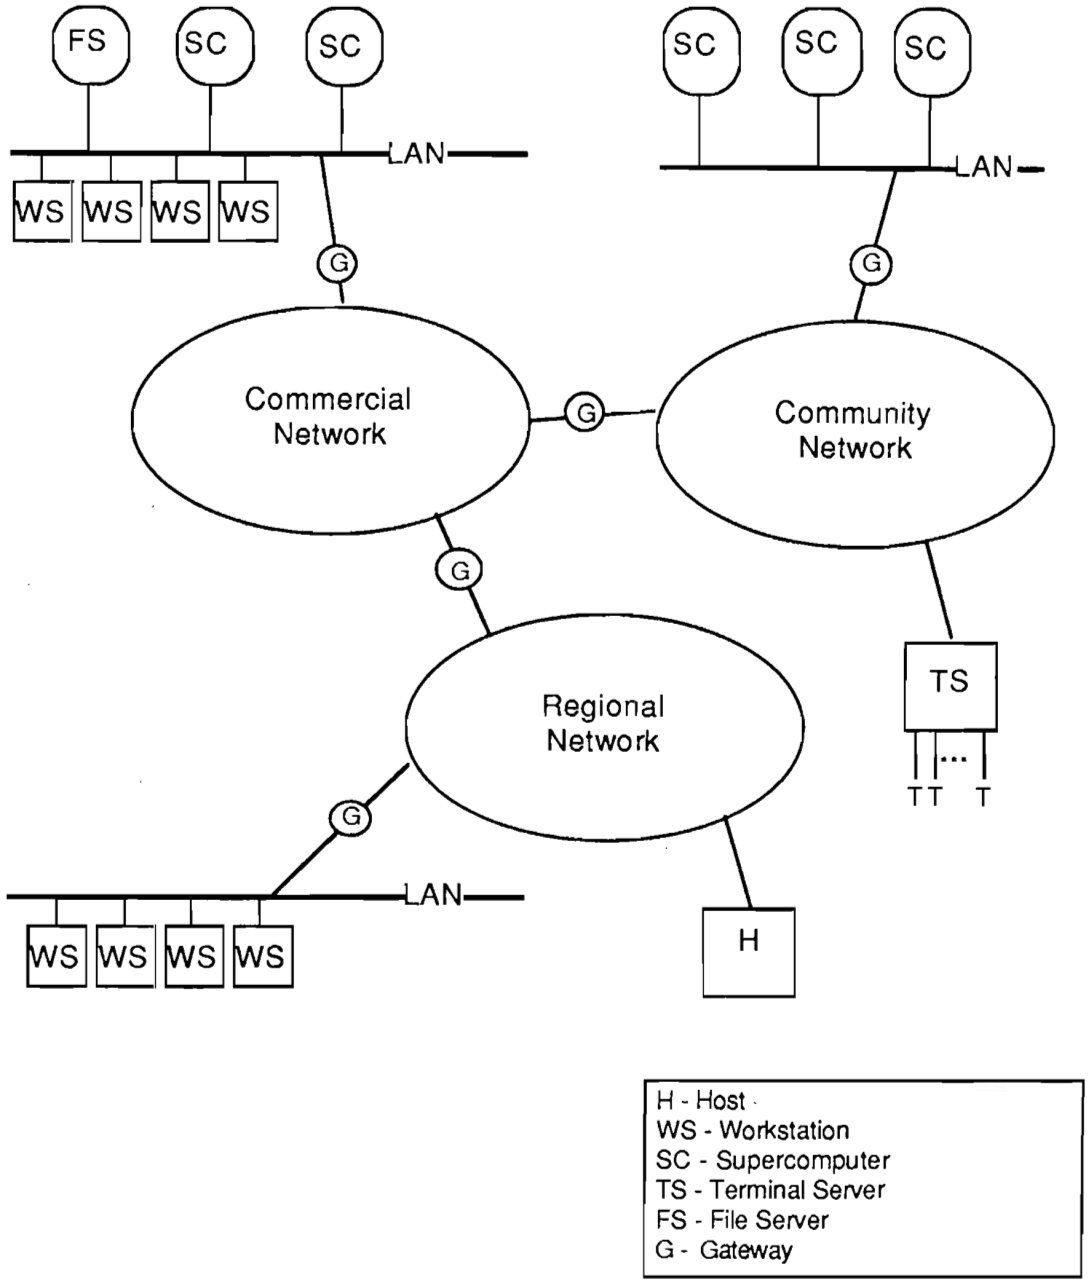

Figure 1: Distributed Systems Architecture

not charge based on use. Rather, the charge is borne by the supporting organization (such as the DoD) as an overhead cost, but access to such networks is limited to those users permitted by the organizations. Although the overall internet architecture currently allows ubiquitous access and connection, such a policy would not work when a large number of networks, each designed to support a portion of the community, are connected to each other. Clearly, access control policies and procedures will be required.

Operating in such an environment in a way that satisfies the access control and privacy constraints requires an architecture that makes full use of the technologies available. It may also require the development of additional technology (although at this time this is not expected to be significant).

Some of the requirements for such a system are as follows:

1. Protect information while in transit between host computers.

2. Protect information while in transit between a user device (workstation or terminal) and a host computer.

3. Protect information on the various host computers and throughout the system (including the various file servers and database servers.)

4. Support private communications between users (confidential electronic mail between collaborators for proprietary results, for example.)

5. Control access to sensitive data.

6. Control access to server resources on the network, including databases, file servers and supercomputers.

7. Control access to networks, allowing only authorized users to use them (even for transit traffic).

8. Control access to private resources, such as experimental facilities or personal files, even when such resources are accessed remotely.

2.3. A Possible Distributed Architecture for Access Control and Privacy

A distributed access control and privacy architecture is required to accomplish the goals described above. Figure 2 shows a possible system architecture. There are two differences between the architectures of Figure 1 and Figure 2. The first is that E^3 devices have been added. The second is that a number of the elements of the system have additional functions, such as the access control func-

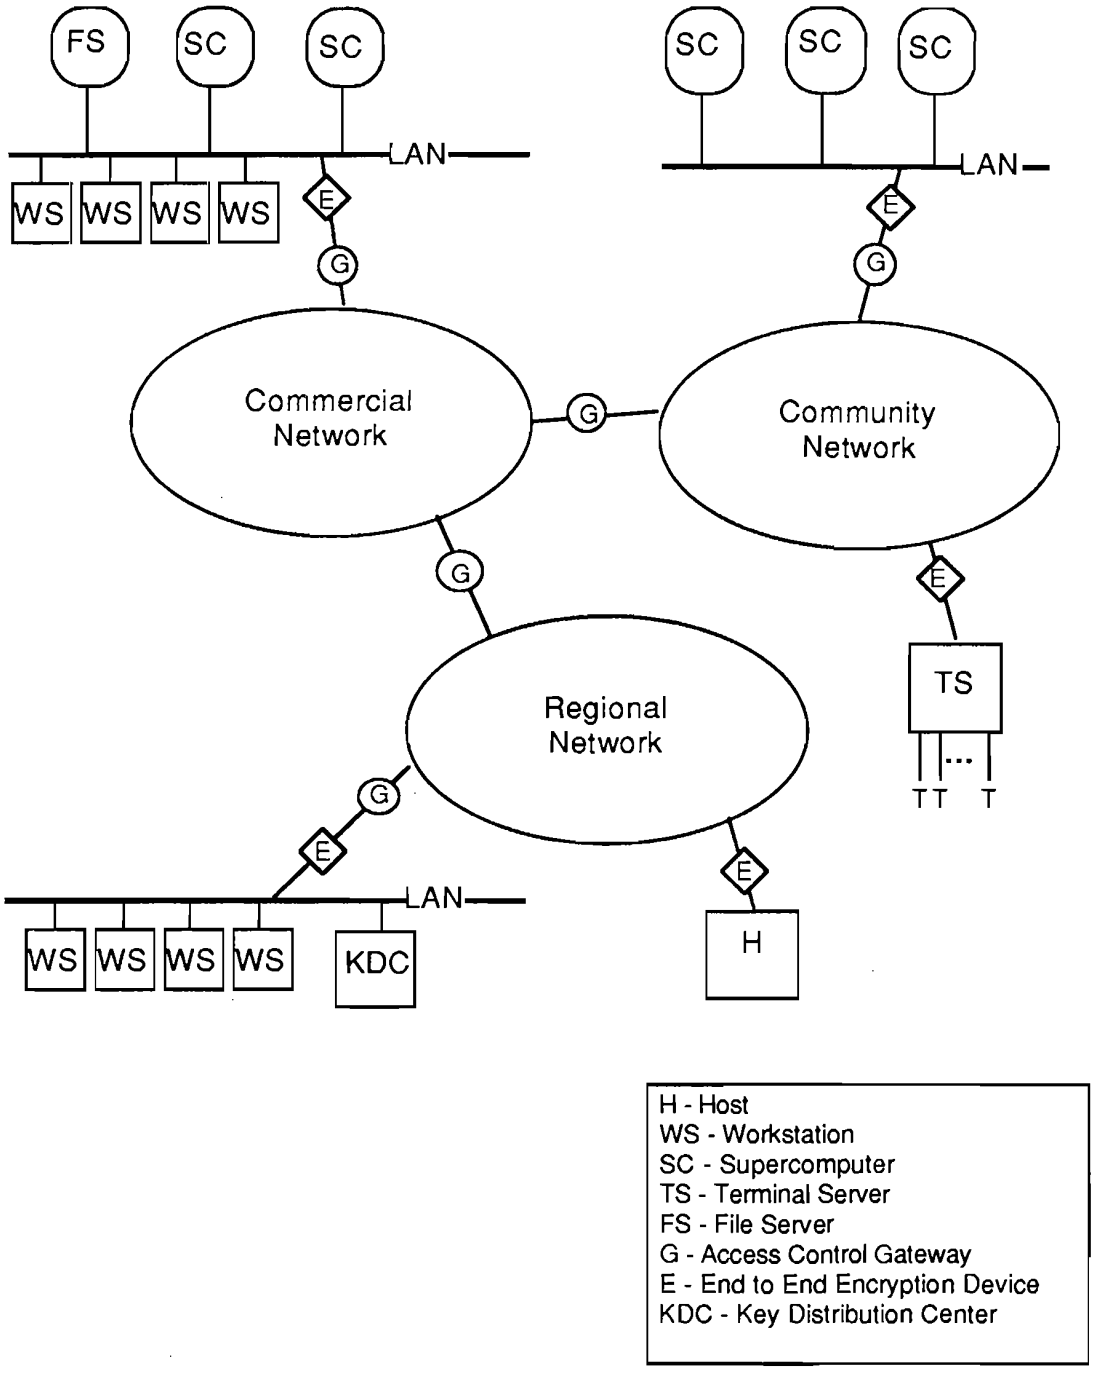

Figure 2: Distributed Access Control and Privacy Architecture

tion added to the gateways.

The architecture is divided into two components, related to the degree of protection in the local environment and the long-haul environment. Typically, a local environment will provide at least a limited amount of protection to both the communications and the data resident in a single machine. Thus, providing protection in the local environment is usually cheaper and simpler than in the long-haul environment. Furthermore, in many environments (such as the NSFnet), the local site is responsible for providing protection in the local environment, while other organizations (such as the NSF) are responsible for protecting information in transit between sites and protecting the long-haul resources themselves.

2.3.1. Local Communications Protection

In the local environment, protected distribution systems guard information and access. This means that the local site is required to provide a sufficient level of protection to the local physical means of distribution by protecting the wires and host computers.

For those users needing greater local protection, the mechanisms used in remote protection can be applied, with the resulting increase in cost and complexity. For example, E^3 devices may be used between workstations on an unprotected LAN.

2.3.2. Remote Communications Protection

The major changes to the architecture are in the long-haul system. E^3 devices are used to protect information in transit between local environments. One or more key distribution centers provide for overall coordination and control of the encryption devices.

Access control gateways will permit use of the network only if the packet is sent by a user in the class of authorized users. To do this, the local environment would be responsible for setting the correct access control flags in the Internet Protocol (IP) header. One possibility is to use Type of Service Options for this purpose. Another is to use the evolving IP Security Option. The important point is that for gateways to deal effectively with the access control, the Internet header must contain the appropriate information.

This implies a certain degree of trust in both the IP header and the gateways. Thus, it must be possible to verify and authenticate the IP header, and to trust the gateways. One way to build such gateways is to use technologies such as the restricted computers described above.

This also implies an effective and trusted method for distributed registration of users and resources. This is most apparent when considering the E^3 devices and key distribution centers (KDC). For a KDC to know that a particular source/destination pair should be permitted to communicate (and therefore given a key), a reliable and trusted method must exist for identifying and authenticating the source and destination. Such a method is also required for

the access control gateways to be effective. Gateways to transit networks (networks not directly connected to either source or destination) must have a method for reliably identifying that a packet associated with a particular pair of users is authorized to use that transit network. Thus, the local environments of those users must be able to authenticate that the users belong to the authorized class, and pass that authentication along to the gateways in a reliable and trusted way. The research in distributed operating systems is expected to shed light on this problem, but further work will be required to provide the necessary authentication mechanisms.

3. Summary

This paper has given a quick sketch of a possible architecture to provide access control and privacy. Clearly, developing and refining a suitable architecture will require considerable additional effort. We hope that this paper will help stimulate the necessary discussions and research to satisfy the important requirement of providing suitable access control and privacy in large distributed systems.

References

1. *Data Encryption Standard,* Government Printing Office, Washington, DC (1977). Federal Information Processing Standards Pub. 46
2. *DES Modes of Operation,* Government Printing Office, Washington, DC (1980). Preliminary copy of Federal Information Processing Standards

RIACS

Mail Stop 230-5
NASA Ames Research Center
Moffett Field, CA 94035
(415) 694-6363

The Research Institute for Advanced Computer Science
is operated by
Universities Space Research Association
The American City Building
Suite 311
Columbia, MD 21044
(301) 730-2656