M.I.T. Laboratory for Computer Science

Request for Comments No. 234
December 2, 1982

A discussion of policy requirements and technical mechanisms for inter-network access control

by Deborah L. Estrin

In mid-August I spent two days in the Washington, D.C. area speaking with members of the defense community about network access control policies and mechanisms[1]. My purpose was twofold:

1. To examine the policies and mechanisms used by the defense community to control network access across the many administrative and organizational boundaries currently spanned by the Arpanet.

2. To discuss existing and anticipated policies regarding interconnection of individual Arpanet-connected networks with non-government or non-U.S. institutions, i.e., private corporations or international research centers.

As a result of these visits and subsequent discussion with Jerry Saltzer, I have written this memo to elicit comments, suggestions, and discussion of the issues raised with local networking-people (not to be confused with local-networking people). The ideas presented are not new, rather I have attempted to lend some organization and structure to various issues and concepts that I have encountered, and thereby establish a basis for discussion and for more in-depth research.

One goal of my proposed research in inter-organizational networking is to develop mechanisms that will allow two organizations to establish network interconnections via *pairwise* agreements, i.e.,

---

[1] I met with Robert Lyons, Phillip Selvaggi, and Edward Kaine of DCA; Stephen Walker of OSD; and Vinton Cerf of DARPA.

without need for consultation with a third organization to whom one or the other is connected. To do so, there must be adequate controls in these various connections so that a particular organization can be confident that its policy concerns are not being violated by the actions of other organizations to which it is connected; similarly, an organization requires some means of preventing members of its community from wittingly or unwittingly misusing the resources of other inter-connected organizations, for which the originator's organization could be held accountable[2].

I find it interesting to note that mechanisms of access control, accounting, and billing on networks are at a rather primitive stage of development relative to most other aspects of networking technology. I attribute some of this to the origins of the bulk of the networking experience in this country; namely the Arpanet where, for good reason, access was for the most part encouraged through lax access procedures and *free* service. Similarly, relatively little research effort in local area networks has been dedicated to the development of accounting or access-control mechanisms because of the limited geographical and organizational span of most local area networks.

1. Policy problems associated with interconnection

1.1. Organizational boundaries

The nature of the organizations being interconnected largely determines the nature of the policy requirements for internetwork access controls. Therefore, in this section the discussion of interconnection is divided into four parts, according to the nature of the administrative or organizational boundaries that separate the networks. Many instances of interconnection involve multiple types of boundaries, but a stricter, if artificial, classification is helpful in isolating and analyzing boundary-type specific, policy requirements.

1.1.1. Classified networks -- classification boundaries

In order to provide users of classified, military-security networks with access to off-network resources, mechanisms are required at the boundary between the two networks that will deny unauthorized users access to the classified facilities. To date, no mechanisms are available which can do this in a certifiable manner. Currently network security on classified networks relies upon end-

---

[2]One example, discussed at the end of this memo, is the perceived need for MIT to consult with DARPA before establishing a direct connection between one of MIT's local networks and a private corporate network, such as IBM's.

to-end encryption, a solution whose administration is not always feasible across network boundaries (see section 3.2.2).

1.1.2. Tariffed networks -- budget boundaries

When networks belonging to different organizations are connected, problems arise with regard to budget boundaries. Namely, transmission and handling of messages must be billed to the users, or groups of users, and payment must be made to the particular carrier, or carriers, that processed the communications. Some of the accounting functions required among interconnected data networks are equivalent to those used in the existing voice, record, and postal communication networks, while other requirements are unique to data networks. In particular, mechanisms are desired to address the following accounting requirements:

- ~ Allocating payment to all the networks across which a packet travels on course to its destination.
- ~ Charging the recipient for incoming communications when the application is such that it is not appropriate for the originator to incur all or any charges.
- ~ Allowing the recipient to reject unwanted communications without/before incurring the transmission and storage costs that are associated with accepting and handling packets on the recipient's private network.
- ~ Allowing internetworked organizations to maintain dissimilar internal accounting practices.

Some mechanisms which address these requirements are are described in more detail in section 1.2.4.

1.1.3. International interconnection -- regulatory boundaries

Interconnection of networks across national boundaries can raise jurisdictional issues of all sorts because of the sensitivity of international relations. For example: government-sensitive or classified networks which carry information of national security import and are also connected to the public network must not be accessible to foreign users, payments must be negotiated for calls between and through countries[3], and national policies regarding the use of private telecommunications services must be enforced.

---

[3]The international voice and record carriers have significant experience in this area but because the number of carriers involved is generally restricted to one per country, i.e., the PTT, the solutions might not be applicable to data networks where there will be a great number of independent private and public networks in many countries; excepting those countries that forbid private networking.

The latter issue is perhaps the most unique to international networks. Telecommunications regulation varies from country to country [13]. Therefore, interconnection of networks that are located in different countries might result in violation of a country's policies regarding telecommunications usage. For example, some countries stringently regulate private networks and usage of non-PTT supplied electronic mail services. Interconnection of an organization's U.S. local network with its network in some other country could provide the non-U.S. operation with private electronic mail services, in violation of the other country's national policies.

1.1.4. Private networks -- proprietary and policy boundaries

When organizations interconnect their private networks, concerns arise regarding protection of proprietary or private information and resources and protection from maliciousness and snooping. There is a distinction between the interconnection of private networks directly, versus interconnection via a carrier network (public or value added). Namely, the carrier network can act as the gateway between networks, or it can merely transmit bits from one half-gateway to another (see section 3.1). Many access control problems might be more easily addressed in the former case where the carrier can provide a jurisdictional buffer. But, if the carrier does not provide the desired level of service (data rate, switching, security, cost), or if the interconnecting organizations desire a higher degree of *control* over their joint facilities[4], then the later might be preferred by the organizations involved.

Interconnection of private networks might also involve the crossing of policy boundaries, as is the case in many European countries. It is possible that anti-trust concerns in the United States might call for certifiably controlled connections between otherwise competing corporations. In addition, when private networks are connected to government networks, the controls needed in the gateway may be dictated by public policy concerns.

1.1.5. Department networks -- administrative boundaries

Although the term *inter*-organizational is repeatedly used in this memo, controls and restrictions are in many cases desired between networks belonging to different departments or divisions of the same organization. The primary difference in the case of inter-departmental networking is that

---

[4]This is often one of the significant motivations for implementing a private network to begin with, i.e. military and corporate networks.

there exists a central authority to which both departments report, whereas no such central authority exists in the case of most inter-organizational networking efforts. The existence of such an authority can aid the coordination of many implementation issues such as standards and accounting.

The primary concerns across administrative boundaries are protection of proprietary information and privacy.

1.2. Control requirements

Having defined a number of policy boundaries which separate organizations we now discuss the controls that the entities on either side of the boundaries might require.

1.2.1. Protection of valuable resources

Protection of valuable resources entails restriction of access to transmission, computational, and information data-base, resources to unauthorized users.

Transmission resources can be abused through third-party routing, inappropriate usage, and uncontrolled incoming traffic. Third party routing [2], refers to the situation in which a network is used as a transit path between two other networks. The third-party network therefore bears the cost of transmission for the other networks. In those cases in which the third party is not affiliated with either the sender or receiver, this arrangement is not appropriate unless the third party has agreed to do it; even so, issues of accounting, dependence, and liability are raised. Transmission capacity and other overhead imposed by extra traffic are valuable resources which many organizations will want to protect via control of usage.

In an attempt to protect access to valuable transmission resources, incoming traffic from off-net sites might be prevented from traveling over private, enhanced[5] facilities. Unfortunately, such action can impair the ability of in-house network users to engage in valuable communications with off-net sites. Autovon, the voice network used by the defense community, can be used by persons within the physical community (i.e., certain private exchanges) to contact persons outside, via connection to the public switched network. But incoming calls that are not from other Autovon users, are not routed over Autovon facilities and use only the public network. Therefore, although

---

[5] Enhanced can refer to lower cost, higher performance, higher security, etc.

the necessary connectivity is achieved with the public, access to the special network facilities are restricted to outgoing calls of authorized users or to calls within the Autovon community. Autovon thereby solves the policy problem of restricting access to enhanced facilities to authorized users while still allowing authorized users to connect to unauthorized users via enhanced facilities. This solution is not immediately extendable to private data networks because in order to restrict off-net, public, messages from being transmitted via private-network facilities, each host that supports mailboxes must have a direct connection to the public network so that incoming mail may come to the host directly, and not via private facilities which connect the particular host to a shared gateway. The difference between the Autovon voice and the data network case is that all Autovon-user complexes (analogous to hosts on the private network) are in fact directly connected to the public switched network in addition to Autovon facilities, whereas one of the explicit roles of private data networks is to provide shared gateways.

Protection of computational resources and information contained in data bases are important issues both within a single operating system and within a single network. The need for protection of these resources is aggravated when the number of users that might have access to the resources is enlarged, particularly when the additional users are not part of a single community, i.e., not controlled by a common authority. In general, protection of system resources is best left to the system itself. In particular, a network should not be expected to increase the security of a system, so long as it does not lower that security. One approach to controlling access to valuable resources that are not proprietary, is to charge for them (see section 1.1.2. Otherwise, the solution to this problem lies in the domain of system security.

Many of the control requirements discussed are also applicable to protection of users' privacy. On the other hand, some of the methods that are proposed for protecting valuable resources have a side affect of decreasing the level of privacy in a system; for instance, the increased monitoring and accounting of communication flows through gateways.

1.2.2. Appropriate usage of network resources

There is a dichotomy between restricting access to network facilities to authorized users, and restricting usage of the facilities by authorized users to appropriate applications; the first is a matter of control of carriage, the second, control of content. Whereas the control of carriage might lend itself to a priori control policies and mechanisms such as examination of the source and destination

fields in the packet-header, a priori control of content impinges on the privacy of users because it entails examination of message contents[6]. A posteriori content-regulation, i.e. policy statements as to appropriate usage of facilities accompanied by strict enforcement should abuse be detected, also entails examination of message contents in order to detect abuses.

Different organizations will undoubtedly adopt different postures as to the ethics of examining message contents, but given that there is any organization among a group of interconnected organizations that regards it as unethical, it is likely that the only mutually acceptable solution would be to not have examination of contents[7].

1.2.3. Maliciousness

Malicious behavior can take the form of snooping, destruction or damage to valuable resources, or denial of service. The first two examples can be addressed in a similar manner as protection of valuable resources. On the other hand, denial of service through such things as jamming gateways or flooding communications paths requires a rather different approach. In many communities, these malicious threats are too expensive to address given the value and sensitivity of the information and resources involved. On the other hand, because of the greater diversity of the user-community, internetworks must address problems of malicious behavior more seriously than do networks in general.

1.2.4. Accounting

Mechanisms are required to allow local networks to charge for transmission and handling of traffic that originates from outside of the local network. This traffic may be destined for a user on the local network or for a user on a third network, to which the local network is connected[8]. It is not feasible for each network to keep track of the originator of each packet that is transmitted over the network and then bill that user or host directly. Telephone company methods are not directly applicable because of the large number of packets associated with any single data-communications

---

[6]The term *a priori* control refers to prevention of the abuse through advanced detection; *a posteriori* control refers to enforcement of policies once an abuse is detected.

[7]The privacy issues related to monitoring of electronic mail were highlighted by Willis Ware of the Rand Corporation at a recent National Computer Conference [6].

[8]The latter instance was referred to earlier as third-party routing.

session. Accounting is easier if messages are relayed between networks a message-at-a-time, as opposed to a packet-at-a-time, because the message originator is easier to identify. On the other hand, message-at-a-time relaying can degrade the performance of the system [12].

If we can assume that the local networks are equipped with mechanisms for allocating charges to the hosts and users within the local network, then each network could treat each gateway as another host and thereby charge the neighboring network for any traffic that arrives via that gateway. The costs of the communications would eventually be passed back to the originator's local network and then to the originator. If, as is often the case, the local network is operated by a central authority and has had no previous need or mechanism for internal accounting and charging of local network resources to individual hosts on the network, no internal accounting mechanisms can be brought to bear and mechanisms must be implemented in the gateways for recording the communication flows into the network. If a considerable amount of traffic flows among a group of networks (electronic mail carriers, industry-wide networks, etc.), accounting can be handled on a statistical basis, where charges are based on a sampling of the traffic flowing between networks, thereby reducing the amount of monitoring needed in the gateway[9].

Existing communication services, i.e., postal, telephone, and telegram, typically charge the originator of the letter, call, or telegram, as opposed to the recipient. In accord with this model, one might expect that the originator would usually incur the cost of a transmission until the point at which the packet or message enters the destination's network. On the other hand, there exist applications, for which the recipient should incur the charge, i.e., some mailing lists, broadcast messages, etc. In many of these cases the communications charges can be paid by the originator, and the originator can in turn bill the receiver for the *service* provided, i.e., billing is handled at a higher level than the packet transmission level. When this is not feasible, a mechanism similar to collect telephone calls, or pay-on-delivery postal service is desired. In addition to imposing more complex accounting procedures on the gateways (keeping track of reverse-charged packets), such mechanisms do not prevent the originator and transit networks from incurring the cost of transmitting the message or packet in the case in which the recipient chooses not to accept the charges.

---

[9]Such a scheme has been used by airline carriers to allocate payments from passengers.

There are a number of scenarios in which a mechanism is needed to reject unwanted packets at the gateway to a local network, before resources are used for transmission and storage of the packet. Such mechanisms might be needed when the cost incurred by the originator of a message does not discourage the originator from transmitting packets onto the recipients network and using up the recipients resources. Junk mail and malicious flooding are two scenarios in which this might be the case. It might be desirable to allow the receiver to first detect such unwanted mail (if the cost imposed is in fact significant) and then to take some action that is less drastic than just turning off the gateway. This scenario is analogous to private PBX telephone exchanges and internal mail-rooms, with three significant exceptions: both of the latter cases require that the originator incur some cost of delivery via a public carrier which discourages such action; in the case of maliciousness, statutory action can be taken in response to mail and telephone abuse; and, private operators or secretaries and mail-room personnel provide some amount of intervention before the mail or telephone call is passed onto the private facilities.

Another scenario in which such a mechanism might be desired arises when local networks are interconnected via a carrier service (e.g., Telenet, AT&T, Satellite Business Systems), as opposed to via a private gateway. Typically, the carrier service receives a payment from each local network that originates traffic onto the carrier's facilities; the carrier is typically not charged for traffic that it delivers to a local network, nor does it charge for such traffic. Therefore, if network A sends a packet to network C, via network B, where A is connected to B and B is connected to C via a carrier service, then A only pays the charge of transmitting the packet from A through B. More importantly, B incurs the cost of transmitting the packet from B through C and does not receive any payment from A to cover the cost of the carrier service or use of B's internal resources. If no accounting mechanism is implemented then any network *B* will require a means of preventing off-net users from generating traffic through *B* onto a tariffed network.

Dissimilarities among the internal accounting practices of the the interconnected organizations might dictate a need for special control mechanisms at network boundaries. If one of the internetworked organizations does not charge its users on a traffic-sensitive basis (e.g., Arpanet) then access to off-net services that are charged to the organization on a traffic-sensitive basis (e.g., Telenet, TWX, Telex) need to be restricted in some manner [24]. That is, either the off-net service-charge must be passed back to the user, requiring an accounting system at the junction of the two

networks (e.g., long-distance phone calls), or there must be an access-checking mechanism at the junction (e.g., a list of authorized users which is matched with the originator field on the packet header, or a password procedure). In addition, the organization that does not use traffic-sensitive billing mechanisms requires some means of detecting and stopping a malfunctioning host or malicious user from generating large numbers of packets for off-net transmission since there is no cost-incentive to the hosts themselves to prevent such occurrences[10].

1.2.5. Restrictive routing

Restrictive routing allows users to express preference regarding the desired class-of-service (security, cost, etc.) and to avoid untrusted links or networks [5], [19]. The ability to designate restrictive routing is desirable for controlling the flow of information in an internetworked environment where not all network facilities are equally trusted or desirable for all applications, i.e., which packets will be transmitted via which gateways and which nets should get which packets. In addition, if communication charges vary according to the path taken, then users will desire the ability to specify minimum-cost routing.

2. Level of interconnection

The level of interconnection between networks, i.e., mail, terminal-to-host communications, or processor-to-processor communications, largely determines the control requirements for the gateway between the networks[11]. The objective at all levels of interconnection is to check that a user (person, device, or process) requesting access is indeed authorized to do whatever it is that is being requested. This might require identifying and authenticating who the user is, if the user is authorized to use the facilities, and if the particular use intended is appropriate.

---

[10] One example of a budget boundary that is complicated by the dissimilarities in the internal accounting practices of the two organizations is the Telenet gateway on the Arpanet. Currently the Arpanet allows any authorized user of an authorized Arpanet host to generate mail on Arpanet facilities, but only a restricted subset of users, those belonging to a restricted set of hosts, are permitted to send mail via the Telenet gateway.

[11] The levels of interconnection discussed in this section do not correspond to protocol layers [2], but rather to the type of communications applications for which networks are used, i.e., mail, terminal-to-host communications (i.e., transfer (FTP) and remote login (Telnet)), and processor-to-processor communications.

2.1. Electronic mail

Electronic mail forwarding between networks is a very restricted form of file transfer; in particular:

1. In most mail transfer schemes, files can only be *sent to*, not *removed from*, a remote host.

2. No remote-file naming is done by the remote user. Mail is automatically sent to the mailbox file of the specified user.

3. It is not possible to overwrite an existing file on a remote system; mail is appended to the mailbox files.

If the gateway serves only as a mail-forwarder between networks then the control requirements are some subset of the following, depending upon the nature of the organizations interconnected and of the boundaries spanned:

~ Protection of valuable resources in the form of traffic handling and mail storage.

~ "Appropriate" usage of mail facilities, as determined by the operating organization.

~ Protection from maliciousness, in the form of large amounts of traffic or junk mail generated onto the network, i.e., ability to reject incoming traffic.

~ Accounting mechanisms that charge the originator, recipient, or forwarder (in the case of third-party routing), as appropriate.

~ Controlling or accounting for the flow of outgoing mail onto tariffed networks.

~ Ability to indicate class-of-service and routes to be avoided, i.e., restrictive routing.

There are a number of interesting problems that mail-forwarding connections have raised in the Arpanet community and which are likely to occur elsewhere. One example regards mailing lists, which broadcast messages to a roster of destinations all over the network or networks. Although mailing lists, in and of themselves, do not constitute an abusive use of mail facilities, they nevertheless can use up significant network resources and generate significant off-network charges if such mail is forwarded onto tariffed networks. Unless mailing-list packets are specially tagged, there is no way to detect or differentiate this traffic. A second example, discussed in [4], regards the *private mail-forwarder*, i.e., an authorized network-user who takes on the role of a mail-forwarder for otherwise unauthorized network-users. Messages are encapsulated in the headers of the

authorized user and there are no mechanisms to detect such traffic without reading the contents of packets. In fact, any authorized network host can implement an automatic mail-forwarding utility and thereby allow unauthorized users to insert traffic onto the network. One instance of this is a daily occurrence on the Arpanet. Any UNIX system that is equipped with the necessary software, called Uucp, can forward mail onto the Arpanet via one of the Arpanet-host, UNIX systems that runs the Uucp utility [15]; the originating UNIX system can dial up the Arpanet-host directly or can forward mail to it via a third UNIX system. Mail-forwarding is handled automatically by the Uucp software so no explicit permission is needed from the Arpanet-host UNIX system in order to forward mail onto the Arpanet.

2.2. Terminal-to-host communications: remote login

A gateway that supports terminal-to-host communications is both more and less restrictive than a mail-only gateway. On the one hand it is more restrictive because it is only equivalent to a dial-up telephone line and unlike a mail-gateway, it provides no capability to generate traffic onto the network unless the user can get past the security of the particular host system. On the other hand, terminal-to-host communications avails the system resources to a much larger and more heterogeneous community of users. Even if the password mechanisms are adequate to keep unauthorized users from gaining access, the existence of a larger and more diverse community of authorized users intensifies the need for file protection and data-base security mechanisms.

Remote login avails all system resources to the remote user, with the exception of those resources that are explicitly protected by the system security. Therefore, the policy requirements for terminal to host communications via a network or internetwork are to a large extent the same as the security requirements of the host system. The primary difference is in the size and diversity of the community of users that knows how and is able to access the resources.

On the other hand, some commonly-used, system-security methods are not suited to use over a network. For example, a password is typically transmitted in the clear from the user's terminal, over a direct-link or telephone line, to the destination system; in the case of remote login, a password would be transmitted through, and stored in, intermediate host systems on its way to the destination. As a second example, whereas keeping dial-up telephone numbers secret is a relatively effective means of restricting access to dial-up computers, no analogous mechanism exists for

networks. Eventually, trusted, third-party, authentication-servers may be needed to address such complications.

2.3. Host-to-host communications -- file transfer and inter-process communications

File transfer capability limits remote-user access to the sending and retrieving of files. As with terminal-to-host communications, the security of host-to-host applications such as file transfer should rely on the host system security. If such security is adequate, no unauthorized traffic could be generated onto the network or off of the network, in contrast to the mail gateway. On the other hand, if system security is not sufficient the potential loss is much greater than with a mail gateway. If it is necessary to restrict authorized, off-network users from gaining access to some of the valuable system resources, secure login procedures may not be adequate in and of themselves. Since most systems in operation today are not "secure", and robust system security can not be retrofitted, safety-net security-mechanisms will be required in the gateways between some networks.

A particular source of difficulty in inter-process communications is that authentication capabilities for the particular processes, not just the users that initiated the processes, must be passed from one host to the other; permission for disclosure is not adequate, as it is in file transfer. For instance, commands issued by the network control hosts must be authenticated at the gateways in order to preclude impersonations of the network controller by other hosts.

3. Mechanisms

In this section we discuss a few mechanisms which address the problems outlined above. We wish to identify and set aside those mechanisms that can be handled by the individual hosts themselves and concentrate on those issues that are raised or are particularly aggravated by the existence of interconnected networks and multiple network communities and on those mechanisms that must be implemented at the junction between networks as opposed to in hosts within the networks.

The most basic requirement implied by our discussion thus far is the need for mechanisms that can effect a controlled outlet or connection between networks. Without a controlled connection, organizations might not be satisfied with a pairwise agreement between themselves and another network community. That is, organization *A* might require assurance that by interconnecting to

organization *B*'s network, *A*'s resources are not made available users on the other networks to which *B*'s network is connected. For example, should MIT wish to implement a connection between a private corporation's network and one of MIT's local networks, it should not be necessary for MIT to seek the consent of the Arpanet community which is also connected to MIT's networks. Nevertheless, the loose control mechanisms currently in place between MIT local networks and the Arpanetwork make such consent desirable. A number of control functions must be implemented to provide the cooperating organizations with more control over the connections between their networks: access control, authentication, restrictive routing, and accounting.

3.1. Gateways

Although the trend in gateway-design is towards simplicity and high-performance, it is safe to assume that inter-organizational gateways will be somewhat limited in their performance because of protocol conversion. That is, it is unlikely that two, interconnecting organizations will use identical network protocols and therefore they will require some level of protocol conversion in the gateway[12]. We might conclude that the stricter the boundary between interconnecting organizations, the greater will be the need for both protocol conversion and control mechanisms. It also seems plausible that in many instances, the stricter the boundary between the organizations, the lower the traffic requirements will be. Therefore some of the concerns about burdening the gateway may be alleviated. In those instances where network protocols are compatible and both access-control and traffic requirements are high, access controls could prove to be the limiting factor on overall performance.

Nevertheless, currently gateways control transit only at the header level and it is generally viewed as unworkable and undesirable to pollute the gateway by looking at more than the header. It is unworkable because the number of possible protocols increases as you move up in the protocol-hierarchy and therefore the amount of information that the gateway has to have blossoms, the complexity would result in degraded performance capabilities of the gateway. Unfortunately, most of the control requirements described in section 1.2 imply a level of traffic analysis in the gateway which might be intolerably burdensome to the gateway and which might conflict with privacy concerns of the users.

---

[12]Personal communications, Judith Estrin.

Most gateways between networks are composed of two half-gateways, each belonging to one of the networks. Each half-gateway may then be designed and maintained by a single community so long as the format used for communications between the two halves is agreed upon. This independence is not complete because in order for one half-gateway to trust the other one, it must trust that the other half functions correctly and can authenticate itself, etc.

3.2. Some examples

3.2.1. Network Login

One network access control mechanism for terminal access to network facilities is *network login*, i.e., implementation of a login procedure in the gateway upon entrance to a network. In 1975, a TIP login routine was unsuccessfully implemented and soon thereafter removed. The failure was only partially due to technical difficulties (inadequate buffer space in the TIPS); the granting of login passwords was to be administered manually which proved to be impossible because of the number of access requests. Since then these failings have been resolved and implementation of TAC-login procedures is planned for dial-up hosts.

The TIP-login also raised a largely unresolved problem of interconnection between hosts and networks that maintain different levels of security. In the case of TIP-login, the insistence that the user provide two passwords in the login scenario would probably be unacceptable to a higher-security entity. This is because the tendency of most users would be to use the same password for both purposes, lowering the security of either the host or the network to the level of the other.

Network login addresses only a portion of the network control functions that a gateway might be required to implement. In particular, it only authenticates network access and does not control network usage nor is it easily applicable to high-bandwidth or processor-to-processor communications. In addition, as with most password procedures, the password is transmitted in the clear; the potential damage of this is aggravated in the case of networks as opposed to dial-up telephone lines because the shared communication paths provide a much easier snooping grounds than does the switched telephone network. Finally, the password administration tasks are immense.

Access control lists in general are not sufficient in large network environments where each user must be *authenticated* and/or where gateways must be high performance. The need for

authentication is one instance of a more general problem of who trusts and depends on whom at the various levels of the network protocols. Authentication helps to certify a level at which all lower levels can be depended upon.

3.2.2. Encryption

Private users who trustingly use the postal and telephone service, might not express much concern regarding the security of equivalently insecure electronic system, but those commercial entities that do have economically valuable information will eventually demand such protection[13].

Encryption mechanisms in networks can be used to effect: privacy protection (encryption of data storage and in transmission), authentication (two-way and one-way) [14], and digital signatures [17]. But, doubts remain concerning the practicality of efficiently and effectively coordinating key distribution, i.e., using an out-of-bound channel for key distribution, taking care of old messages encrypted in old keys, and administrative complexity.

Encryption might also provide a mechanism for restricting access to gateways, and therefore to particular networks. That is, users that are authorized to use particular gateways or networks would be given a key that is compatible with the key used by the particular gateway; all packet-headers that were not encrypted in this key would be rejected[14].

3.2.3. Source routing

Most dynamic routing algorithms keep track of all connectivities and delays in the network. When restrictive routing is used (see section 1.2.5), each restriction introduced (according to security, cost, bandwidth, etc) forms a new connectivity, and each connectivity in turn requires a separate routing algorithm! In addition, care must be taken to avoid looping as a result of uninformed nodes that repeatedly route packets through unacceptable links [16].

Source routing is one feasible way of introducing restrictive routing. But, there are some unresolved objections regarding the use of source routing methods. Namely, by foregoing centralized control over routing and allowing users or hosts to control their own routing, paths can

---

[13] Currently, this concern is not widely expressed; it might take a number of additional incidents, in which significant monetary loss or inconvenience is experienced, before they are willing to incur the higher costs [6].

[14] Idea suggested by David Reed.

be forced, wittingly or unwittingly, through particular links, gateways, or networks, perhaps against the common good of the network. Although gateways are equipped to refuse packets, should a user try to force packets through an overloaded gateway, the gateway must be able to determine when it is and when it is not appropriate to do so. Another difficulty associated with source routing is that the routing servers need to know about the current topology of the network; Singh has proposed mechanisms for dealing with ever-changing topologies such as are found in a campus-wide network community [19], [21], [22][15].

4. An Example: Arpanet structure and operation

4.1. Current structures

Currently the Arpanet supports communications among Department of Defense (DOD) affiliated universities, research centers, contractors, and government agencies. The university and research center affiliations are administered by the Defense Advanced Research Projects Agency (DARPA), the agency responsible for technical development and operation of the Arpanet. All policy concerning the Arpanet is set by the Defense Communications Agency (DCA). In addition, DCA is responsible for the separate military network used for operational, strictly classified communications, Defense Data Network (DDN).

Currently, Arpanet usage policy is largely a posteriori; with the exception that access to the net is controlled a priori to the extent that access to connected-hosts is controlled. DCA and DARPA have set policy as to what the network cannot be used for, such as for personal gain, but the only enforcement mechanism is detection after the fact. Because of privacy concerns, as well as burdensome accounting and costs, associated with examining the contents of messages, policies that differentiate appropriate and inappropriate usage according to content cannot be regulated in an acceptable fashion, but only through publication of policy statements, peer pressure, and strict enforcement (should abuse be detected through non-invasive means).

A third network, CSNET, is is being developed and operated under the auspices of the National

---

[15] Of the two primary internetwork standards to date only one has included a source routing capability. In the X.75 specification for interconnection of public networks, the user cannot specify any routing restrictions, whereas, in the Internet Protocol, there is a provision for the user to force a routing path [7].

Science Foundation. It is a logical network which uses Arpanet, Telenet, and telephone dial-up connections as its physical infrastructure. Its purpose is to connect non-DOD affiliated university and research organizations that do not have access to Arpanet facilities and are engaged in Computer Science related research. Unlike the Arpanet its purpose is not to experiment and develop network protocols per se, but rather to provide an electronic communications infrastructure. Currently CSNET supports only mail communications. The system is based on Arpanet technology and mail is forwarded between the two networks. There is currently no charge for mail forwarding between CSNET and Arpanet on the assumption that flows would be approximately equal in both directions. The close relationship between the government agencies that operate the two networks makes this lax arrangement feasible, if not realistic.

Mail and electronic data interchange networks are also being developed by federations of educational [8] and commercial organizations that are not directly involved in computer or communications related research or business (e.g., insurance, grocery, banking industries).

4.2. Future

Although we have only a small amount of information regarding future development and evolution of the Arpanet, military networks, and CSNET, we do know that in the future usage will be shifted off of Arpanet facilities and onto other networks, specialized according to function and user-groups. Future interconnection of the various networks and their related domains, is an interesting example of inter-organizational networking which warrants further study.

The structure that currently comprises the Arpanet will eventually be partitioned into two sections, military and research. The military section, the Defense Data Network (DDN), will be composed of a highly classified network and a sensitive, but unclassified network (MILNET). The network will be operated by DCA and engineering support will be provided by Defense Communications Engineering Center (DCEC). The Arpanet will therefore be released of much of its obligation to offer the level of reliable communication-services required by military users; a level which is generally incompatible with the network's role as a testbed for experimental protocols and services. Likewise, the DDN will be used for operational communications and its services and configuration will be considered stable. In the future, DCA will continue to operate and oversee policy for the military network(s) while DARPA will take on policy responsibilities for the research and development network, Arpanet.

There is some discussion of splitting the remaining, non-military, Arpanet users between a scaled down version of the existing Arpanet and a scaled-up version of CSNET. The Arpanet's primary function would become experimentation and development of network protocols and techniques, as opposed to support of message and other types of communications among researchers and their computers. DARPA hopes that the later function would be increasingly served by CSNET; although, in order to take over much of Arpanet's current role, CSNET will have to implement Telnet and File Transfer Protocol in addition to mail. As users shift over to CSNET for communications services, the Arpanet would become a much smaller network and will be used primarily by research groups who are themselves conducting network research and development. Much of this depends upon the financial viability of CSNET, which in turns depends on the willingness of universities and other institutions to actually pay for network facilities. As CSNET grows, more sophisticated accounting and access control mechanisms will be needed, particularly in those gateways or hosts that connect CSNET to the Arpanet[16].

4.3. Issues of interconnection

It is not yet certain what type of connection will be implemented among the three network communities identified. Between the military and non-military networks the issues of access control are severe and were touched upon in section 1. Between the Arpanet and CSNET, the issues and concerns are somewhat more mundane but nevertheless are more representative of the sorts of problems that are likely to arise among other governmental, as well as commercial, organizations; these issues were also discussed in section 1.

The conflict regarding standardization is one manifestation of the problems associated with networking among heterogeneous organizations. Many of the communication standards that potentially affect all three network communities are established by the National Communications System (NCS), which oversees all telecommunications for DOD, and National Bureau of Standards (NBS), which oversees standards for data processing as well as higher level data-communications with the goal of inter-operability for all federal government systems. Because the DOD has far more stringent and immediate requirements for reliability, security, and survivability than do the other communities involved, conflicts arise in regard to accepting national and international

---

[16]CSNET developers note that because of these strict access-control and accounting requirements, CSNET hosts, as opposed to simple gateways, may be used to interconnect CSNET and Arpanet [12].

standards that present trade-offs between cost, flexibility, and international compatibility on the one hand, and security, reliability, and survivability, on the other.

Interconnection of Arpanet-connected networks with other networks *should* not be of concern to DARPA or DCA. But, because the existing gateways that interconnect networks to the Arpanet do not implement adequate control, DARPA and DCA are in fact concerned that they not be held liable for either favorable treatment afforded to some third party that gains access to Arpanet facilities via an Arpanet-connected network, nor for damage caused or traffic generated by users of the Arpanet who are able to gain access to the third party's machine. DCA policy is as follows: "If it is possible to gain access to the Arpanet from another network or from a tributary terminal of a host via the IMP-host connection, it is the responsibility of that host to provide software protection which will permit only authorized Arpanet users to access the network." [1]

References

[1]

Network Information Center.
Arpanet Directory.
Technical Report NIC 49000, Defense Communications Agency, March, 1982.

[2]

Cerf, V., Kirstein, P.
Issues in Packet-Network Interconnection.
Proceedings of the IEEE 66(11):1386-1408, November, 1978.

[3]

Cheng, D.
Interconnection of Existing Electronic Mail Systems.
Master of Science Thesis, Massachusetts Institute Technology, August, 1981.

[4]

Cohen, D., Postel, J.
Internet Mail Forwarding.
In *COMPCON.* IEEE, Winter, 1982.

[5]

Cohen, Danny.
Controlled Routing in the Catenet Environment.
Internet Experimental Note 156, USC/ISI, September, 1980.

[6]

Gillin, P.
Little Interest Greets Ware's Call to Study Electronic Mail, Linkage Security Problems.
Computerworld :18, October, 4, 1982.

[7]

Grossman, G., Hinchley, A., Sunshine, C.
Issues in International Public Data Networking.
Computer Networks 3:259-266, 1979.

[8]

Heller, Paul.
Mailnet: A Convenient Inter-Campus Electronic Mail Service.
EDUCOM Bulletin :3-7, Summer, 1982.

[9]
Horton, M.
How to Read the Network News.
Technical Report, Bell Telephone Laboratories, 1981.

[10]
Karger, Paul A.
Non-Discretionary Access Control For Decentralized Computing Systems.
Technical Report M.I.T./LCS/TR-179, Massachusetts Institute Technology, May, 1977.

[11]
Kern, W.
CSNET Project Outline.
CSNET document 2.1, National Science Foundation, Computer Science Section, March, 1981.

[12]
Landweber, L., Solomon, M.
Use of Multiple Networks in CSNET.
In *COMPCON.* IEEE Computer Society, Spring 1982.

[13]
Mathison, Stuart.
Commercial, Legal, and International Aspects of Packket Communications.
Proceedings of the IEEE 66(11):1527-1538, November, 1978.

[14]
Needham. Roger, Schroeder, Michael.
Using Encryption for Authentication in Large Networks of Computers.
Communications of the ACM 21(12):993-999, December, 1978.

[15]
Nowitz, D.
Uucp Implementation Description.
Technical Report, Bell Laboratories, October, 1978.
Unix Version 7 supplementary documentation.

[16]
Perlman, Radia.
Access Control: An informal discussion.
Internet Experimental Note 58, BBN, October, 1978.

[17]
Popek, Gerald, Kline, Charles.
Encryption and Secure Computer Networks.
Computing Surveys 11(4):331-356, December, 1979.

[18]
Postel, J.
Internetwork Protocol Approaches.
IEEE Transactions on Communications 28(4):604-611, April, 1980.

[19]
Saltzer, Jerome.
Source Routing for Campus-Wide Internet Transport.
Internet Experimental Note 144, MIT, March, 1980.

[20]
Shoch, J.
Inter-Network Naming, Addressing, and Routing.
In *COMPCON Fall 1978*, pages 280-287. IEEE, Fall, 1978.

[21]
Singh, Vineet.
The Design Of A Routing Service For Campus-Wide Internet Transport.
Technical Report M.I.T./LCS/TR-270, Massachusetts Institute Technology, August, 1981.

[22]
Sunshine, C.
Addressing Problems in Multi-Network Systems.
Internet Experimental Note 178, USC/ISI, April, 1981.

[23]
Sunshine, C.
Interconnection of computer networks.
Computer Networks 1:175-195, 1977.

[24]
Taylor, C.
Keeping track of AT&T alternatives.
Data Communications ():, August, 1982.