M.I.T. Laboratory for Computer Science

Request for Comments No. 259
June 14, 1984

Usage Control Requirements in Inter-Organization Networks

by Deborah L. Estrin

Abstract:

Private computer facilities of distinct organizations that are interconnected to support inter-organization interchange (referred to as inter-organization networks) raise interesting issues for network designers, information system managers, and policy makers.

Inter-organization communications via IONs can be qualitatively different from traditional forms of communication (i.e., post and telephone). It can allow a person or device within one organization to cause something to happen automatically within a second organization, without any opportunities for intervening human decision within the second organization. This characteristic has significant implications for the participants as well as for the computer programming needed to provide "judgment". In particular, the connection of computer networks across organization boundaries poses different design requirements from traditional *intra*-organization connection. If a participant does not want to permit unrestricted access to all of its internal facilities, it must discriminate between internal and external users. Consequently, whereas designs for *intra*-organization networks traditionally emphasize connectivity, performance, and transparency with full remote function, ION participants may *insist upon* explicitly limited, remote function. On the other hand, internal requirements for transparency persist and should not be compromised by externally imposed control requirements.

1. Introduction

Private computer facilities of distinct organizations that are interconnected to support inter-organization interchange (referred to as inter-organization networks) raise interesting issues for network designers, information system managers, and policy makers [3, 1, 2, 12, 11, 6]. In this paper, I focus on technical issues; in particular, a new functional requirement of the network technology, *usage control*. In the first two sections of this paper, I introduce a model of IONs and usage control requirements encountered therein. In the third section I describe specific usage control issues encountered in existing IONs. In the fourth section I outline design requirements for usage control mechanisms.

IONs are an interesting research topic for two related reasons. First, inter-organization communications via IONs can be qualitatively different from traditional forms of communication (i.e., post and telephone). It can allow a person or device within one organization to cause something to happen automatically within a second organization, without any opportunities for intervening human decision within the second organization. This characteristic has significant implications for the participants as well as for the computer programming needed to provide "judgment". Second, the connection of computer networks across organization boundaries poses different design requirements from traditional *intra*-organization connection. In particular, individual participants may not want to permit unrestricted access or integration with all of their internal facilities. Each may want to control usage at least by discriminating between internal and external users. Therefore, whereas designs for *intra*-organization networks traditionally emphasize connectivity, performance, and transparency with full remote function, now users may ***insist upon*** explicitly limited, participant-determined remote function.[1] On the other hand, internal requirements for transparency persist and should not be compromised by externally imposed control requirements. In effect, each organization wants to implement multiple logical networks on top of its internal physical network, where some of the logical networks cross organization boundaries and encompass pieces of physical networks belonging to other organizations. Usage control mechanisms are needed to differentiate among classes of users and resources.

---

[1] Additional design efforts will be called for in the area of naming across organization domains, but this is not the focus of my proposed research.

1.1. Definitions and classifications

IONs (and networks in general) can be described on three levels -- physical, logical, and operational (see figure). At the physical level, an ION is the transport mechanism and the supporting architecture (e.g., data format, coding, and exchange protocols) via which data are passed between organizations; this is the level most commonly addressed by computer-communications network designers. The interconnection of the organizations' facilities need not manifest itself in the installation of a physical wire or switch, but only in an agreed-upon protocol for transferring and interpreting data.[2] For instance, travel agents connect via a specialized protocol over dedicated leased lines to the airline's central computer; insurance companies employ a commercial third-party network which is itself an SNA architecture but which is accessed by customers via dial-up or dedicated telephone facilities or via a packet switched network operated by Telenet; the research institutions use packet switched architectures over telephone lines, primarily; and the various customer supplier interchanges use standardized or specialized protocols via dial-up telephone and magnetic tape transfers.

At the logical level, an inter-organization network (ION) is the set of accessible computer resources (e.g., hosts, printers, servers) and applications formed via interconnection of facilities that are owned, operated, and/or used by two or more organizations. Participating organizations typically are most concerned with the network at this logical level. The logical ION excludes human decision making as part of the interconnection process and deals only with automatic procedures. It refers to all processes and applications that are reactive, i.e., that can be invoked automatically from another position in the ION. Currently existing examples of IONs include interconnections between airlines and travel agents, airline companies themselves, banks, insurance companies and agents, research institutions, medical-product suppliers and hospitals, automobile and steel producers, etc. In each of these cases the interconnecting organizations (referred to as participants) desire to enhance certain operations that cross organization boundaries so as to achieve greater efficiency, market certainty, or some other performance criteria.

At the operational level, an ION includes the administrative procedures and policies that govern use of the facilities encompassed in the ION. For example, the types of interchange, patterns of

---

[2]For example, even magnetic tape transfers or automatic processing of telex messages qualify as automatic processing of external transactions; although in the case of tape transfer issues differ because transmission is not automatic.

usage, access rules, and accounting imposed by respective participants, etc. This level is of most concern to the managers of the ION-supported functions within the participating organizations and to the individual end users.

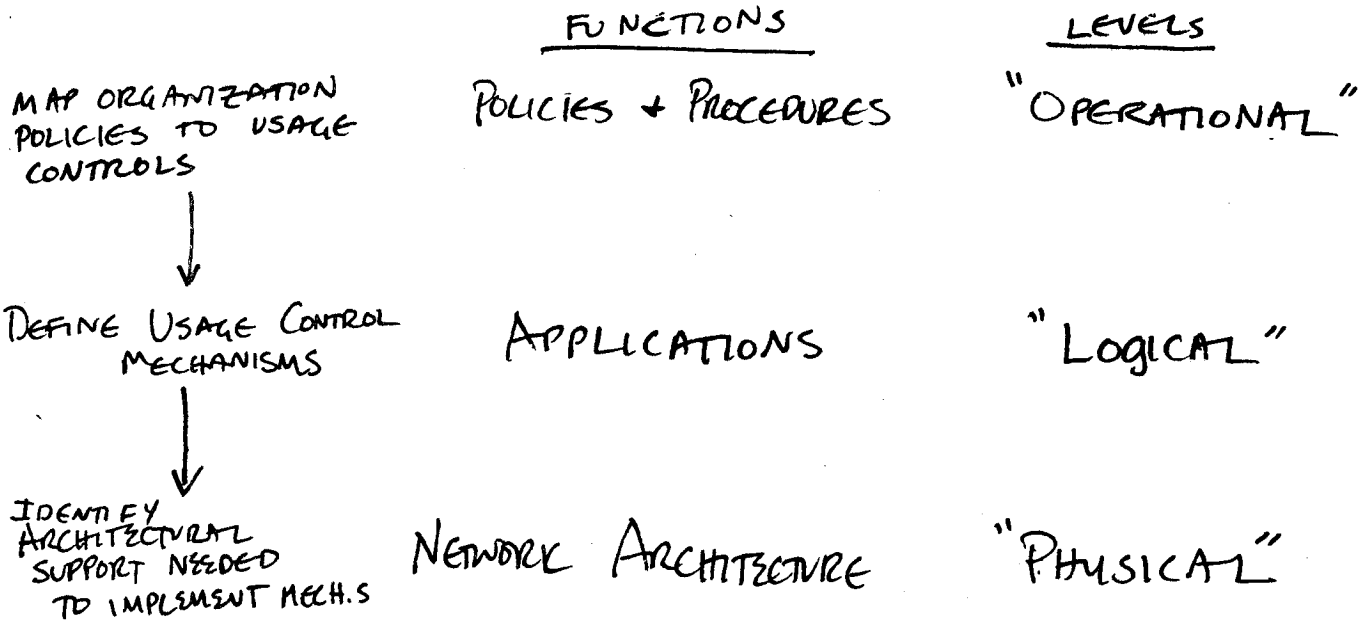

Figure 1: Levels of an ION

I distinguish among these three levels because a given logical network can be supported by any one of a number of physical configurations, and can be operated in a variety of ways, but the design choices made at each of the three levels interact with one another. For example, policy requirements at the operational level imply implementation requirements at the logical level, which in turn imply design requirements at the physical level.

At the *logical level* existing IONs can be divided into two types -- Inter-Organization Links, IOLs (dedicated to a single ION application) and Inter-Organization Communications Networks, IOCNs (multi-purpose, user-defined applications). Some interconnected facilities are dedicated to specific, well-defined, inter-organization interchange functions (e.g., a particular database transaction application such as airline reservations or order/entry); these are appropriately referred to as IOL. Although such inter-organization connection raises significant policy issues for the participants, from a technical standpoint, usage control mechanisms can be treated as an extension of traditional database management and information system security issues. In contrast, IOCNs are composed of facilities interconnected to support generic inter-organization communications (e.g., electronic mail), on top of which a multiplicity of user-defined applications operate. IOCNs are not unified

systems (although the individual internal facilities may be). They arise out of interconnections between a certain set of facilities of two or more organizations. By virtue of this interconnection a range of resources potentially are accessible to persons and machines within the other organization(s). But, the interconnection does not imply that the entire set of resources are intended to form an integrated system or even to be accessible.

This distinction between IOCNs and IOLs can be described in terms of overlap between logical networks. The set of resources that the participants *intend* to make accessible via an inter-organization arrangement forms a logical ION. In addition, each participant has its own logical network(s) used for applications that pertain to internal operations. If the logical ION does *not* overlap with the logical internal networks of the participants, the inter-organization arrangement is an IOL; i.e., the facilities accessed by parties external to the organization are dedicated to that purpose and thus provide a single place where policy can be enforced without imposing on strictly-internal functions. If the logical ION and internal networks do overlap, the arrangement is an IOCN; i.e., the facilities are used for multiple, internal and external applications. See figure 1. Increasingly, organizations will support multiple IOLs on top of an IOCN, and therefore that the distinction between the two will be one of level of analysis, as opposed to discrete types of systems.

The importance of this distinction between IOLs and IOCNs is that in the case of IOCNs internal resources used for internal purposes are used for external purposes as well. As a result, usage controls applied within the system(s) to which external users have access affect strictly-internal uses of the system as well. Therefore, design of IOCN usage controls must take into account the tolerance for changes in internal usage controls (e.g., tolerance for decreased performance, restricted information flow, disincentives to resource sharing and communications). In addition, the IOL interface is dedicated to a single function and therefore has both a simpler set of policies to implement and an easier task of enforcement.

2. Usage Control Mechanisms

In this section I focus on the IOCN environment within which usage control issues are most unique, but within which the answers still vary widely depending upon the technical and organizational characteristics of the interconnected facilities and institutions. I discuss the range of policy requirements encountered in IOCN activities and propose guidelines and criteria for designing mechanisms with which usage control policies can be supported.

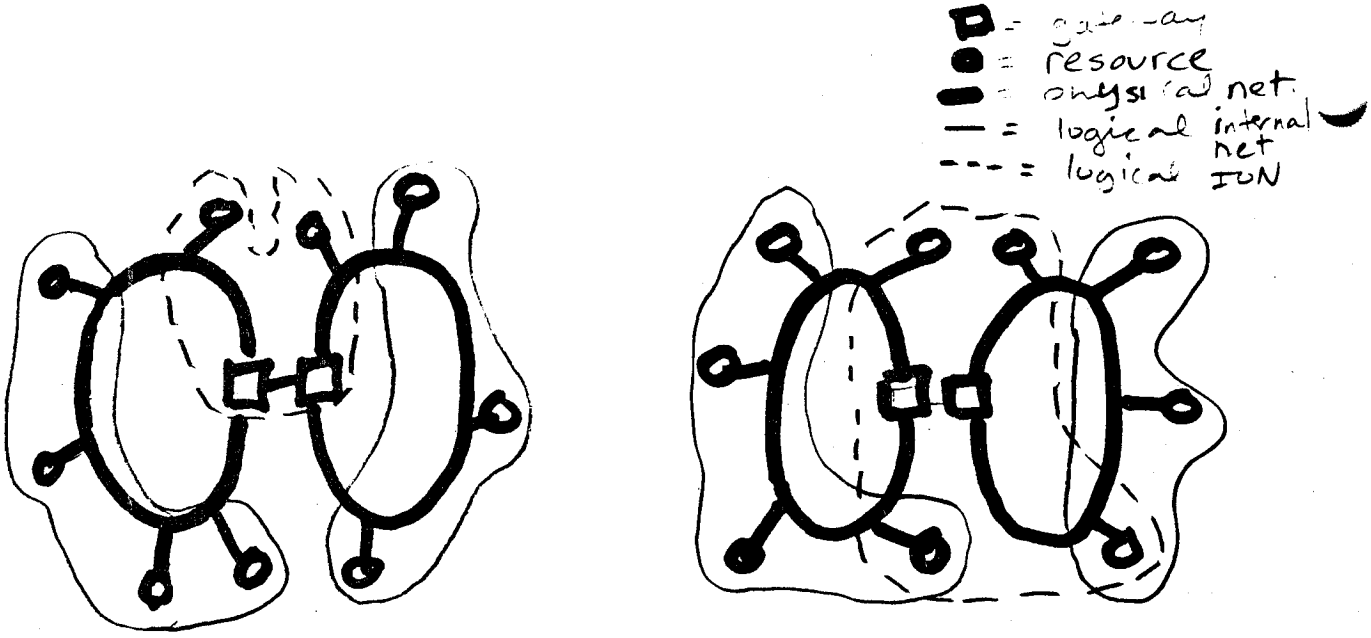

Figure 2: Overlapping and non-overlapping logical networks.

2.1. Usage Control Requirements

For the sake of this discussion I will assume that within an organization's internal network there exists a set of explicit and implicit policies and procedures which are *considered* adequate for the intended environment. namely, the employees of the organization.[3] Typically, the purpose of such internal networks is to facilitate communication and access to shared resources. Therefore, although individual hosts or servers connected to the network frequently will include a protection system to isolate users, many services are treated as internal utilities which have no protection systems because they are perceived as making use more cumbersome with little compensating benefit. I t is even less common for data communications and processing facilities. in particular, network transport or electronic mail, to include logging or accounting systems of any kind. Therefore, when such interconnected internal systems are made accessible to outsiders, there may be no means of treating external users differently from internal users other than by preventing access altogether. In any case. It is fundamentally difficult to convert from an environment composed of networks and resources in which the default is open access to one in which the default is closed; and the difficulty is increased the greater is the decentralization of management control over the resources. In other words. when an organization's internal network is exposed for the first time via an ION gateway,

[3]This assumption is necessary in order to isolate issues regarding inter-organization networks from networks in general.

explicit design effort is needed if resource boundaries[4] are to be preserved in their pre-interconnection state.

As discussed in the introduction, usage control issues differ for IOLs and IOCNs due to the difference in overlap of internal and external logical networks. The well-defined applications of an IOL can imply greater reactivity of internal resources to external inputs due to more concrete automatic processing of external communications. On the other hand, this defined quality can support greater, and more centralized, control over the extent of reactivity than in the case of an IOCN. In the case of IOLs, system security issues are intensified, but usage control policies can be satisfied, for the most part, by adopting or enhancing traditional system-security internally without infringing upon internal operations. In contrast, generic IOCNs raise a broader range of network and resource control issues that differ from traditional notions of security requirements. Each participant may want to implement multiple logical networks, some strictly internal and some that cross organization boundaries. If traditional access controls are implemented within each resource such that all users (both internal and external) encounter equal scrutiny, conflict may arise between internal and external requirements (e.g., tolerance and need for cost and performance overhead of security measures). Alternatively, controls can be implemented in cooperation with other resources on the network so that internal users are treated differently than external ones. The value of the latter, more complicated design, depends upon the value placed on the minimization of usage controls encountered by internal users within each organization.

In general, the function of IOCN usage controls is twofold: isolate non-overlapping logical networks that share a common physical network from one another, i.e., build walls and gatekeepers around each logical network; and maintain the boundaries between overlapping logical networks by implementing usage controls within those resources that belong to multiple logical networks (i.e., resources in an overlap between walled domains) (see figure 2). The first function is the more straight forward of the two and resides in the domain of traditional lower-level network security [10, 8, 13]. In other words, any device that is physically connected to resources outside of a single logical network is responsible for maintaining the boundaries of that logical network.

---

[4]The term resource boundary refers to the dividing line between facilities and information that are owned, operated, and accessed internally, and those that are not.

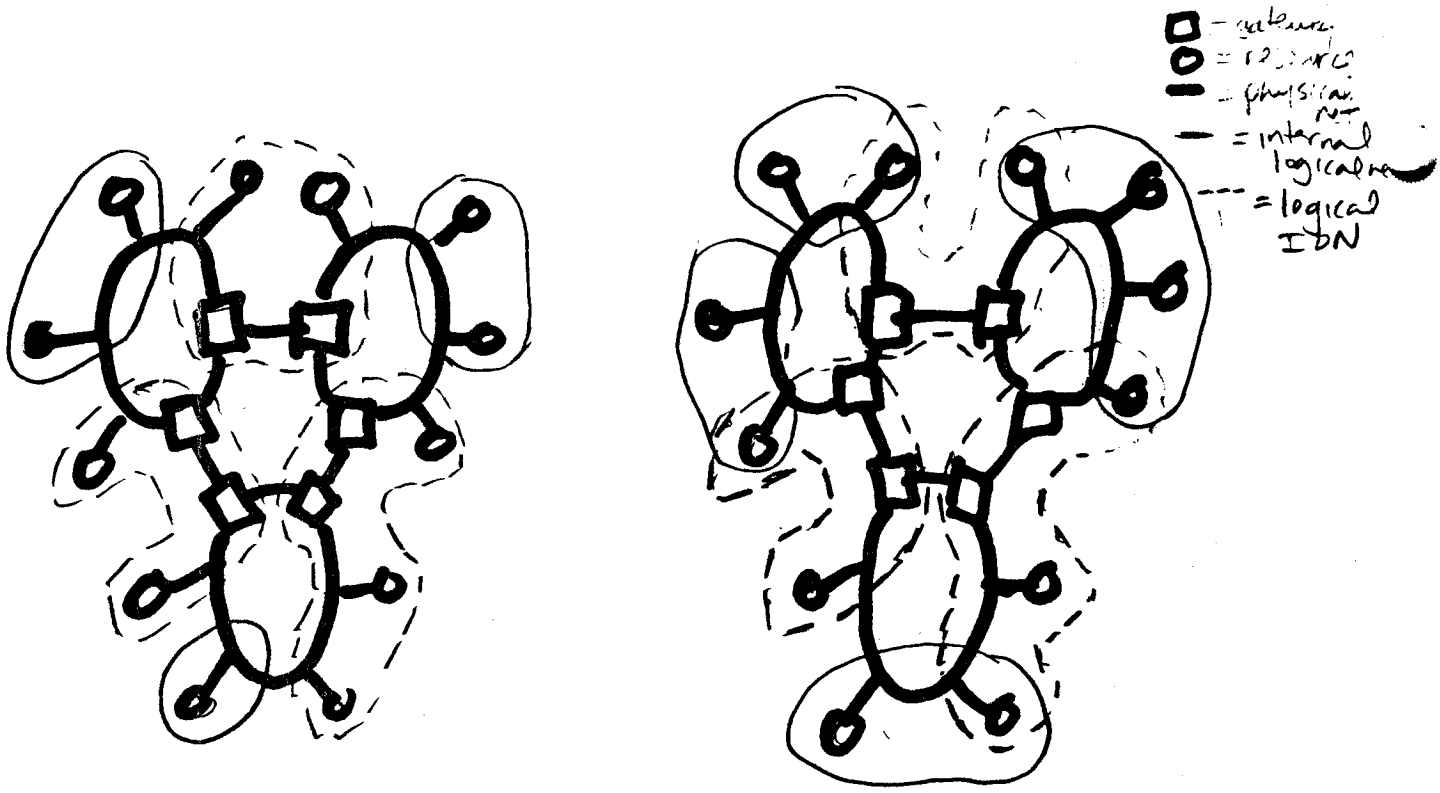

Figure 3: Usage control functions.

A resource that resides within a single logical network can do whatever filtering is desired for the entities in that logical network. But, if a resource resides in multiple logical networks that have different usage control requirements, the resource must be able to discriminate between members of each logical network. For example, a device may define a barrier for one logical network (i.e., no information is intended to flow into or out of the device unless the information is going to and coming from other devices in the same logical network) while acting as a forwarder or shared resource for a second logical network. In order to discriminate in its provision of forwarding services, the shared device must distinguish between the different sets of entities that access it via the common physical network. Traditional protection mechanisms are adequate within a resource dedicated to external functions. But, if the resource is used also for internal functions, performance overheads, restricted information flows, and disincentives to resource sharing may not be tolerable, and new mechanisms are needed.

The usage control requirements that arise in reaction to (or anticipation of) interconnection depend on the nature of the interconnecting organizations (referred to as participants). Formally structured organizations that manage resources conservatively and have proprietary interests to

protect are unlikely to allow changes in external resource-accessibility to occur readily, assuming they are aware of the change. Such organizations are more likely to refrain from interconnection (the ultimate form of usage control) unless or until usage control mechanisms can be implemented to maintain existing resource boundaries. Alternatively, such organizations may adopt new usage controls that impose a change in internal procedures (such as increased internal access control or accounting) in order to accommodate interconnection without effecting a change in resource boundaries. On the other hand, loosely structured organizations that have ill-defined proprietary interests, and that manage resources more loosely, will accept some changes in resource boundaries more readily. In fact, they may be less tolerant of impinging on internal communications than they are of increasing external accessibility. impinge on internal communications. Within the research and development community (which I describe in section 2.2) examples of the former are the industrial labs, whereas examples of the latter are the university labs. Because requirements vary among organizations that are interconnected to one another, ION usage control mechanisms must support a range of participant-defined policies, and different coexisting policies for the participants at either end of the connection. The goal of this research is to define usage control mechanisms that will permit interconnection in such a way as to mitigate undesired changes in both external resource-accessibility and internal usage controls.

Within the internal network of an ION participant, implementation of more flexible usage controls involves two functions, *differentiating* between internal and external users within the network as a whole, and using this information to *discriminate* in the provision of services. In other words, how can one implement multiple logical networks on a single physical network[5] in such a way that minimizes imposition on internal users. In section 2.5 I will explore *tagging* and *filtering* mechanisms and the design issues of their implementation.

This research emphasizes the design of mechanisms that support *articulation* of policies, as separable from certification of a mechanism's security and the information on which it bases decisions. For example, specifying that a type of information is needed to support a particular usage control policy is somewhat separable from questions of who provides the information and whether it is forgeable or trustworthy. The rationale for this emphasis is *not* that certifiable security is

---

[5]The single physical network might itself be composed of multiple local and long haul network facilities. But for the sake of this discussion I will refer to the entire internal facility as a single network.

unimportant; rather, what is most different about IONs from traditional intra-organization networks and systems is the need to articulate and support new policies.

Nevertheless, security per se is addressed in two respects: as a primary motivation for some types of usage control policies (e.g., access control), and as a design parameter of supporting mechanisms. With regard to policy motivation, a significant difference between access control requirements for an ION connection and more traditional requirements is the greater acceptability in IONs of *detection* of abuse as opposed to a priori prevention.[6] The ongoing relationships among ION participants typically are such that there is significant disincentive to abuse the ION facilities, in the presence of detection capabilities, due to resource dependency, legal contract, or cultural standards. With regard to the design of supporting mechanisms, for the most part, these security issues of enforcement and certification are not qualitatively different than they are in the case of internal networks, although the perceived need for such enforcement may be greatly increased. One exceptional issue that is unique to IONs is the absence in IONs of a single, mutually-trusted mechanism to mediate, settle disputes, and provide services such as authentication, key distribution etc.

2.2. Examples of Usage Control Requirements in Existing IONs

The discussion below describes usage control requirements encountered in a number of R&D networks. The examples described in this section are taken from the research community because that is where a number of sophisticated, internal and inter-organization networks are in use. Although the mixture of academic and industrial institutions involved provides some diversity, traditionally thin boundaries between research institutions of all kinds reduces the extent to which we can generalize from these examples to other commercial activities.

2.3. Transit and the pairwise connection problem

CSNET [7] is a network linking computer organizations engaged in computer science and engineering research throughout the US, Canada, and Europe. Membership is open to any university, corporation, government agency, etc., engaged in computer science research or advanced development. CSNET provides electronic mail, gateways to other networks, and a database of

---

[6] The accuracy of this statement varies with the nature of the service type supported, information or resources interchanged, and the perceived threat of malicious attacks.

CSNET entries.[7] Although some CSNET members attempt to constrain the overlap of ION and internal logical networks by forwarding mail to and from authorized registered personnel only, most member institutions forward mail to any mailbox within the organization. If these mailboxes include forwarders and gateways, the logical ION completely encompasses the participants' internal electronic mail networks, as well as other networks to which the internal networks connect.

This cascading of networks, in which the logical ION encompasses the participants' internal nets including gateways, raises a problem of transit. For example, in order to control costs, preserve desired levels of service for CSNET members, and preserve the utility of paying CSNET membership dues, CSNET wishes to limit the amount of traffic that is originated by non-CSNET members and carried over CSNET facilities. On the other hand, it is not desirable to prevent or prohibit all forwarding because the value of CSNET to its members is proportional, in some sense, to the number of institutions that are accessible via CSNET; i.e., gateways are desirable. At a minimum, CSNET does want to prohibit communication between non-CSNET members over CSNET facilities, called transit, since no CSNET member benefits from such use.

Non-CSNET members gain access to CSNET via forwarding by CSNET members. To control undesirable forwarding, the appropriate tactic in this particular (research) community is to produce incentives for CSNET member hosts to not forward non-member traffic. One ad hoc mechanism might be to state the policy and rely on peer pressure, since forwarding is detected when a message is read by the recipient. But this would have little effect on the transit problem since the non-member recipients and sources are not likely to be affected by peer pressure of this sort. A less ad hoc mechanism is to charge per message or set upper bounds on usage for each CSNET member host. This would force users to address the problem of implementing controllable forwarders and gateways so that they forward mail only from authorized machines within the member institutions but do not forward mail from non-CSNET members. In order to comply with such a policy the member hosts need a mechanism to tag and filter non-local from local traffic, i.e., to control the overlap between internal and ION logical networks. One caveat regarding this approach is the tendency to discourage all forwarding of non-CSNET traffic due to cost, difficulty, and imposition of implementing flexible usage controls; even when it makes economic sense on a system-wide basis for CSNET members to receive and send some off-net mail via one another.

---

[7] CSNET also provides, at higher cost and effort, login and file transfer.

Note that CSNET differs from the two research networks, BITNET [4] and UUCP/USENET [9], in this regard. Neither BITNET nor UUCP/USENET charge for services and therefore do not face the burden of protecting their investment as a network.[8] But, although membership is free and largely unrestricted in BITNET and USENET, individual members may want to limit their forwarding burdens due to limited resources, both cpu time and leased line or dial-up capacity. Therefore, to varying degrees, individual members of the networks share CSNET's interest in controlling transit.

A somewhat different perspective on the transit issue is the *pairwise connection problem.* This issue arises when one organization interconnects to two other organizations that do not intend or desire to be connected to one another. Without usage controls in gateways that delineate and isolate logical networks from one another, such interconnection creates a path between the two unrelated organizations by default. One example is the above firm's connection to Arpanet and UUCP/USENET by virtue of its CSNET connection, whereas the firm has not connected to ARPANET or UUCP/USENET directly. Another example is the connection between M.I.T. Artificial Intelligence laboratory and a local company in support of joint research. The AI lab is in turn connected to the rest of the M.I.T. networks and to the Arpanet, and the local company is in turn connected to a number of its customers. Due to the nature of the network interconnections, the logical ION encompasses all internal facilities, and therefore by default a connection exists between all of M.I.T.'s networked resources and this company's customers!

2.4. Insulating participant policies

A single institution is likely to connect to several special-purpose networks, where each network represents an interest group (e.g., CSNET, SCIENCENET, EDUCOM). Membership overlap among the communities results in overlapping logical networks (i.e., a user might belong to more than one of the interest groups), and sharing of facilities within the institution results in multiple

---

[8] Joining BITNET involves acquiring a leased line to a nearby BITNET member and thereby picking up ones portion of the costs directly. BITNET communications software (RSCS) is an IBM product and is available for other types of machines at a small charge. BITNET services, i.e., BITSERVE, are developed in a cooperative manner with a large amount of direction coming from its birthplace, CUNY. Similarly, an institution joins UUCP/USENET not by paying a fee or signing up with any central coordinator, but rather by finding an existing member to connect to and paying the telephone charges to transfer the traffic to that connected host. Again, the communications software is part of standard UNIX software and is available for other types of machines at little charge, and mailing list/bboard services are maintained in a distributed, cooperative manner.

logical networks sharing a single (set of) physical network(s) (i.e., users share an institution-wide network utility as well as special services including name servers and gateways).

Assuming for now that policy is uniform within any single logical network, different institutions will desire different policies regarding facility use, sharing, and gateway access. Similarly, different interest group IONs will desire different policies regarding access, billing, etc. In order for these different institutions and interest groups to support their respective policies without imposing on one another's, they must be able both to share facilities internally and to conform with the controlled gateway desired by an external interest group network. In general, mechanisms should minimize the extent to which the internal institution must adopt or modify internal procedures in order to conform to external policies. For example, if such external traffic is billed on a usage-sensitive basis, an institution should have a means of limiting/controlling outgoing traffic without necessarily implementing accounting internally. An example of this problem is found on CSNET and Arpanet which both have gateways to an X.25 public packet-switched network which charge on a per-packet basis (i.e., Telenet) but neither network implements that kind of accounting internally.

One commercial firm has installed connections of its internal network to two research networks, BITNET, CSNET. Although the majority of the network's members implement forwarders that will accept a message from any accessible source and send to any accessible destination, this firm does not want to support any-to-any communications of this sort. Consequently, it has taken steps to insulate its policies from those of the other ION members. The two gateways both contain access lists of internal authorized users;[9] employees who are not on the list, can not make use of the gateway. The access list defines the hosts and users who are in the ION logical network. In addition, the BITNET gateway requires that the source or recipient on the external side of the gateway be registered as an approved communication partner for the individual within the company with whom communications is desired, i.e., pairs of individuals must be registered. Although a central list is maintained in the gateways themselves, registration is distributed throughout the organization. To be registered, an employee's host must receive the site-manager's approval and the individual must receive a manager's approval.

An additional policy within this firm is to restrict the accessibility of servers to any external

---

[9] In this case a user is a person. But in other applications, the user could be, for example, a program.

communications. Aside from mail forwarding, which can be viewed as a particular kind of server, no servers (i.e., devices that automatically interpret and act upon the contents of an electronic message) can be used via the gateway; i.e., the ION logical network includes mail services only. This policy is implemented by prohibiting the registration of any server machines in the gateways. This strict prohibition is necessary due to the inability of internal servers as implemented currently to differentiate between internal and external requests for service. Even prohibition of this sort is difficult to implement since servers are not always explicitly labeled as such and individual users can quite easily write programs that behave like servers (i.e., automatically interpret messages and carry out actions in response). Currently, the only remedy in effect is education of internal users.

2.5. Mechanisms

This section is a preliminary discussion of issues and approaches to usage control mechanism design. It is intended to delineate a direction, as opposed to define a solution. The overall approach is to identify the mechanisms needed to support observed usage control requirements and to identify the architectural, or network, support needed to implement these mechanisms.

To follow on the discussion begun in the introduction, usage controls delineate and enforce the boundaries between logical networks. Supporting mechanisms can be divided according to the following two functions: differentiating between messages (or whatever the unit of transfer is) that belong to users of different classes, i.e., logical networks; and discriminating in service provision according to origin. In terms of mechanisms, this implies *tagging* external units as they enter a domain and *filtering* on the basis of the tags and other information at the service site, where a service site may be a gateway, host, printer, or other form of computational or communications resource. The following issues are central to the design of tagging and filtering mechanisms for ION usage control[10]:

- ~ Should each application extract the information it needs for usage control or should some network service (e.g., the gateway) do preprocessing and explicitly tag messages; does the end-to-end argument apply?
- ~ What information should be included in a tag?
- ~ From where should the information and criteria according to which a message is tagged be taken?

[10]These issues are the subject of further elaboration in ongoing research.

- ~ From where should the criteria by which a resource should discriminate according to tag values be taken?
- ~ What type of filtering should a given resource do?
- ~ At what protocol level should an ION gateway operate, e.g., packet, mail, transaction?
- ~ What are the internal costs (monetary and performance) of retrofitting mechanisms; what expense is acceptable?

The level at which one should implement a given network function is a familiar tradeoff in network design. In the case of usage controls, the tradeoff is between performing the tagging function at the gateway or a specialized server, and leaving it to the end applications to process the header of each message to determine its origin and other relevant characteristics. On the one hand, it may be difficult to define a single set of information that is appropriate for all internal servers. In addition, tagging in the gateway requires terminating the protocol with the associated performance overheads. On the other hand, requiring each server to do its own processing of headers may be inefficient due to redundancy. The homogeneity of the internal network and policies determines (i.e., are there multiple internal logical and operational networks) whether or not tagging requirements are likely to be the same across internal resources, and therefore, whether tagging is best done in the gateway or the end points..

The question of appropriate level raises a fundamental question of what information should be included in a tag. A tag may be as simple as a one bit flag indicating that a message originated from outside of the organization domain, or it may contain additional information. The information included in the tag may be information available in the header, but in abbreviated form to make filtering more efficient. Alternatively, it may include additional information that can't be determined by any server or host in the network, such as the classification of the source according to some organization-wide criteria or tariff structure. A related design choice is using the tag to represent either the identity of the source or the level of privilege. Tagging information need not be standardized throughout an ION, but information upon which a tag is based must be available in the headers of messages. In other words, the information in a tag determines the range of enforceable policies. In addition, on a given internal network there must be a known set of external gateways which can be relied upon to tag incoming traffic and filter outgoing traffic. Examples of tag-information requirements are Karger's *proxy login* which differentiates between local and remote users of a given host, and filters according to the path via which a remote user

communicates with the host [5]. An example of message-content based access control is filtering out all messages that contain certain control characters to which internal machines are programmed to respond.

Once the information desired in a tag is defined, from where should the seed information be taken or accepted? In any case, it should be explicit where various pieces of information originated from so that individual internal servers can ascertain how to interpret the information and whether or not to trust it.

Filtering is an internally-specified practice which is independent of other participants and of other internal services. Each server, host, or gateway can implement its own filtering policies so long as the necessary information is provided in the message tags. Filtering may be implemented using access lists or ticket mechanisms. In other words, a server may look up a tag value in a table to determine whether or not, or how, to process a message. Alternatively, the tag may be subject to some test to determine whether or not it was provided by the gateway (or some other chosen ticket-giver) and if it was, it can be treated as a ticket whose possession implies acceptability.

Filtering functions are of three types: access control, audit, and accounting. According to the information in a tag, the server may choose to restrict access, account and bill for the usage, and or, log the usage or attempted usage. For different functions, different filtering policies are appropriate. One example of an access control and accounting function in the telephone system is the classification of telephone types and associated privileges. For example, on a number 1 ESS centrex, class-E telephones can not be dialed into from outside the centrex; and some internal phones can only dial within the centrex, and cannot generate billable traffic.

The appropriateness of list or ticket based mechanism depends upon usage patterns and service type, as well as the type of filtering. For example, if the user population is relatively static and small in numbers, a list mechanism makes sense, whereas if users are sporadic and from a large total population, the list would become sparse and inefficient and a ticket mechanism would be more appropriate. A related issue is authorization dynamics. How does a message obtain a ticket? How does a user or class of tags become included in a server's access list?

The level of interconnection of an ION gateway, e.g., packet, mail, transaction, also influences the

nature of tagging and filtering function; both what is feasible within various real-time constraints, and what is necessary in order to implement resource boundaries. For example, in addition to imposing much tighter real-time constraints on filtering and tagging systems, a packet-level connection is more of an exposure and more difficult to control than a staged, mail-level connection since the gateway knows little about the use of the communications. Note that whatever the level of interconnection, all filtering need not be carried out in the gateway itself; e.g., the gateway can pass externally-generated packets to a policy server.[11]

Retrofitting is inherent to usage controls in IONs since we assume that the networks that are interconnected were not themselves intended for such interconnection. What performance degradations should be acceptable in gateways and servers in accommodating tagging and filtering processes? How should this retrofitting be planned and implemented within the organization -- is it the responsibility of the subunit that operates a particular resource to implement filtering or is it the responsibility of the subunit that operates or initiates the interconnection? In any case, in more sophisticated environments there are likely to be a number of functions which have developed over time that are not explicitly thought of as servers but which fit the definition of automatically interpreting and responding to messages.[12] These functions will need to be addressed as they are uncovered but some sort of internal audit for all active servers is in order for organizations with significant proprietary concerns.

Finally, it is essential to refer to the discussion of usage control requirements for criteria by which to evaluate suggested mechanisms. For example:

- ~ Support of autonomously decided upon, dissimilar policies by each participant in an ION.
- ~ Minimize increase in internal usage control.
- ~ Support participant and application specified levels of security for the various mechanisms.

---

[11] This configuration has two additional advantages: destinations need only be able to validate tickets from a single host, and logging is centralized more easily.

[12] For instance, on one known internal network a formatting program used for producing hard copy of on-line messages has the capability of automatically processing system commands that are included in the text of the document to be printed. Thus, a message that is printed using this formatter without having been scanned previously for system messages, can cause any system to behave in server mode.

3. Conclusion

In this paper I introduced the notion of usage control requirements in IONs, including definitions, explanations, and approaches to solving the problem in diverse environments.

Although the bulk of this paper focused on usage control requirements in multi-purpose communications environments, future developments will result in these same resource management issues arising in single-purpose system environments. In particular, as commercial organizations continue to develop highly-connected internal networks, and as they convert increasing numbers and types of external transactions to electronic form, the situation will develop in which multiple IOLs are run on top of an IOCN utility. Consequently, usage control mechanisms will have to address the complexity of the IOCN infrastructure as well as the sensitivity of the commercial transactions typically supported by IOLs.

Acknowledgments

I am very grateful to Bob Baldwin, Jerome Saltzer, and Lixia Zhang for thought-provoking discussions on the subject of this paper and for their comments on a previous draft.

References

[1]

Barrett, S.
A Framework for the Analysis of Automated Inter-Organizational Information Sharing Systems.
Dissertation submitted in partial fulfillment of requirements for degree of Doctor of Philosophy with a major in Business Administration, University of Arizona, June, 1983.

[2]

Cash, J.
Information Systems That Change Company and Industry Boundaries.
Working Paper, Harvard Business School, May, 1984.

[3]

Estrin, D.
Inter-Organizational Networks: Stringing Wires Across Administrative Boundaries.
In Mosco, V., editor, *Proceedings of the Eleventh Telecommunications Policy Research Conference.* Ablex, New Jersey, 1983.

[4]

Fuchs, I.
BITNET -- Because It's Time.
Perspectives in Computing 3(1):16-27, March, 1983.

[5]

Karger, P.
Security in DECnet: Authentication and Discretionary Access Control.
Technical report. TR-121, Digital Equipment Corporation, January 8, 1982.

[6]

Kaufman, F.
Data Systems that Cross Company Boundaries.
Harvard Business Review :141-155, January-February, 1966.

[7]

Landweber, L., Solomon, M.
Use of Multiple Networks in CSNET.
In *COMPCON Spring '82.* IEEE, February, 1982.
San Francisco, California

[8]

Needham, R., Schroeder, M.
Using Encryption for Authentication in Large Networks of Computers.
Communications of the ACM 21(12):993-999, December, 1978.

[9]

Nowitz, D., Lesk, M.
A Dial-Up Network Of Unix™ Systems.
Technical Report, Bell Laboratories, August, 1978.

[10]

Popek, G., Kline, C.
Encryption and Secure Computer Networks.
Computing Surveys 11(4):331-356, December, 1979.

[11]

Stern, L., Craig, C.
Interorganizational data systems: the computer and distribution.
Journal of Retailing 47:73-91, Summer, 1971.

[12]

Veith, R.
Multinational Computer Nets.
Lexington Books, Lexington, MA, 1981.

[13]

Voydock, V., Kent, S.
Security Mechanisms in High-Level Network Protocols.
ACM Computing Surveys 15(2):135-171, June, 1983.