M.I.T. Laboratory for Computer Science

Request for Comments No. 271
February 4, 1985

Usage Control Requirements in Inter-Organization Networks: Implications for Network Interconnection

by Deborah L. Estrin

(I plan to submit a version of this paper to the 9th Symposium on Data Communications.)

Abstract

When two or more distinct organizations interconnect their internal computer networks to facilitate inter-organization interchange, they form an Inter-Organization Network (ION). [3] As organizations establish such connections and extend their networks internally, they require new usage control policies and mechanisms to cope with the increased heterogeneity of the user population. These policies can be enforced by implementing non-discretionary controls in the entry and exit points to each organization's internal network, i.e., the ION gateways. [4]

In order to implement these non-discretionary controls, an ION gateway must determine the organization affiliation of traffic sources and destinations, at least. Typically, this information is not available at the packet level. The source and destination network numbers listed in each packet header usually provide *topological*, not *logical*, information. Consequently, many existing packet-level gateways could not implement the non-discretionary controls proposed.

There are two alternatives to standard packet-level interconnection. The first is to implement the gateway at a higher level (e.g., mail, file transfer, remote procedure call, etc.) so that the logical information needed to map traffic to organization affiliation, and thereby to access privileges, is available. The second alternative is to evaluate affiliation and access privileges at a higher level and then pass the access information to a packet-level gateway along with a means of authenticating the authorized traffic (e.g., an encryption key).

After describing usage control requirements in IONs, the various implementation approaches and associated tradeoffs are discussed in terms of existing IONs.

1. Introduction to the Problem

As organizations establish inter-organization connections and extend their networks internally, they require new usage control policies and mechanisms to cope with the increased heterogeneity of the user population. For example, consider the case of a university computer science department such as MIT's that is connected to the Arpanet. In the past, the department could adequately comply with the Arpanet policy that the Arpanet be used only by computer science researchers because the user population that could access the Arpanet-connected machines was small. However, as MIT extends its computer networks out from the computer science and engineering department to the rest of the campus, the user population that can access the Arpanet-connected machines is no longer small nor homogeneous. Similarly, the university might establish additional external network connections, with local industry, for example. In this case, the potential user population of the computer science department's facilities includes not only members of other departments, but members of other organizations altogether. In this new environment the computer science department may have to introduce control mechanisms to restrict access to the Arpanet gateway or Arpanet-connected hosts in order to adequately comply with Arpanet policy.[1] These control mechanisms need to discriminate between various segments of the user population; in this case these segments are logical groupings of users according to organization or department affiliation.

A second issue that arises when interconnection reduces the homogeneity of a network concerns performance as opposed to policy. Typically, the larger and more heterogeneous is the user population, the less tightly coupled are the applications that operate across the entire network. However, most applications that assume tight coupling among users still operate under the assumption that this tight coupling spans the entire network, even in the presence of extensive interconnection. I will use Xerox's Grapevine system to illustrate this point. Grapevine is a distributed database service that provides mail, naming, and other information to users and applications. Fundamental to grapevine is the manner in which it keeps the distributed data repositories very up to date. The updates can take up significant network resources. If two

[1] A related policy requirement with similar technical implications is that the Arpanet not be used as a transit path between two points, neither of which is itself a legitimate Arpanet node.

organizations that each run Grapevine interconnect their networks in order to support some inter-organization application such as electronic mail, the Grapevine updates will travel across both networks. However, this extremely up to date naming information may be far less appropriate to the loosely coupled relationship of the two organizations than it is within a single organization. And given the gateway bottleneck through which all inter-organization traffic must flow, the updates may exact a significant performance cost. In addition, such inter-organization connections are often more transient than intra-organization connections. Transient connections are the exception internally and the tightly coupled applications may not be designed to adapt easily. Therefore, for performance reasons, it would be useful for broadcast information such as minute-by-minute Grapevine updates to carry information in the packet header indicating that the packet is intended for logically-local destinations only. Logical locality is emphasized since it is what determines the appropriate degree of coupling for this application.

A second example is the use of Address Resolution Packets (ARPs) to locate hosts. Some networks broadcast ARPs over the entire network in order to locate a particular machine.[2] Consequently, when two networks that use ARP are simply interconnected, all ARPs flow across both of them. If these two nets belong to two distinct organizations, it may not be cost effective to broadcast all ARPs across the boundary.[3] However, without information about the logical affiliation of a packet source and destination, there is no easy way to use a broadcast-based search mechanism locally without using it across the ION gateway as well.

In the following sections I focus on the issue of controlling ION flows to meet policy, as opposed to performance, requirements. However, similar mechanisms and issues pertain to the latter concern as well.

[2] CMU is one example. Their use of ARP is described in [7].

[3] In fact, problems related to broadcast of ARPs have been experienced at MIT. MIT's Lab for Computer Science's local network is connected to the AI laboratory's Chaosnet. The AI lab Chaosnet is in turn connected to a local company's, Symbolic's, Chaosnet. Symbolic's Chaosnet is in turn connected to many of its customers' hosts or networks. A bug in the machine of one of these customers caused large amounts of ARP traffic to be generated to the extent that it flooded a MIT LCS local network and caused several network-attached personal computers to cease functioning. Although this same problem would have occurred if the buggy machine had been on the same local network as the personal computers, it is unfortunate that by virtue of interconnecting an organization makes itself so dependent on the correct operation of another organization's machines.

2. Non-Discretionary Controls

An uncontrolled connection between two distinct organizations implies that the organizations are willing to trust one another and all organizations to which each interconnects. Under some circumstances this level of trust is appropriate due to the nature of the organizations (e.g. low risk), their relationship (e.g., not competitive), or existing contract provisions (e.g., liability for violation of connection). However, under many conditions it may be highly inappropriate and inconsistent with other aspects of inter-organization relations and interchange.

My approach to the problem of usage control in IONs is to implement non-discretionary controls in all entry and exit points to an organization's internal network (i.e., all ION gateways). Examples of policies that an organization wants to enforce using such controls are:

1. Accept incoming traffic only if it is from an authorized outside entity and is destined for an internal system or gateway that has been explicitly registered as available to such outside access. Access may be refused to an external user either because of the user's organization or group affiliation, or because of the type of access requested.
2. Forward outgoing traffic only if it is from an authorized internal entity and is destined for an authorized external network. External network access may be refused to an internal user either because of contract provisions that restrict the use of the external network, or because of usage fees charged by the external network. Information flow may also be restricted on the basis of internal sensitivity classifications.

All entry points to an organization's internal network must be treated as ION gateways and equipped with controls. For example, MIT sells time on one of its timesharing systems to a wide range of users--small local companies, international research centers, government personnel, other universities, etc. Because users potentially have access to the MIT network to which the system is connected, the system itself acts as a gateway and must enforce controls consistent with MIT's policies; for example, restricting access to other gateways and certain internal resources (e.g., printers, scarce computational resources, etc.). If controls are desired in connections that cross organization boundaries, all entry points, both full-fledged gateways and hosts such as this one, must be equipped to address incoming and outgoing traffic. The benefit of the approach proposed here is that systems that are used strictly by insiders need not be modified at all, nor made aware of the presence of new interconnections. However, those systems that are made accessible must employ mechanisms to enforce application-specific controls (e.g., which files can be accessed, which programs can be run), and to isolate ION processes from non-ION processes.

To implement non-discretionary controls an ION gateway must have access to certain information about the logical characteristics of traffic; e.g., organization affiliation of source and destination, type of service, amount of resource requested, etc.[4] According to this information the gateway determines which categories of internal information or resources the external entity may access. In other words, in addition to the traditional bindings between user or service and node, node and network attachment point, and network points and path [9], the ION gateway needs a binding between user or service and organization affiliation. Domain naming captures this notion of affiliation. [6] However, as is described in the following section, it is not possible to evaluate the domain affiliation of a packet based solely on the network number that it carries in its header.

If the logical information required for policy decisions is available, then the above method can be applied by assigning category sets to incoming and outgoing traffic according to logical characteristics of the traffic and enforcing invocation and information flow controls accordingly. [4] Below I describe the issues associated with low-level connections (packet-level), for which this information is not always available.

3. Packet-Level Interconnection

An ION gateway, as with any gateway, can be designed to operate at one of several levels. Typically, gateways operate either at the message or connection level or at the packet level. The former implies that that gateway acts as an end-point in a message or connection protocol (such as file transfer, remote login, or electronic mail), whereas the latter implies that the gateway operates at a lower level, forwarding packets between the endpoints of many different applications.

As is discussed below, most packet-level gateways do not have access to the information needed to make ION policy decisions. This is not inherent to this level of connection; rather, it is a result of the competing requirements considered during the design process.

3.1. Network Numbers

The most common requirement for an ION gateway is to identify the organization affiliation of the traffic destination and source. Given this information, the ION gateway can assign categories

[4] For simplicity, much of this discussion focuses on organization affiliation of source and destination and mode of access. Similar arguments apply to other types of information.

and determine the rights of that source and destination.[5] The source and destination in a packet header appear in the form of network numbers. In this section I describe some of the problems of relying on these numbers for identification of organization affiliation.

Networks interconnected at the packet (e.g., Internet Protocol (IP) level in the DARPA TCP/IP family of protocols) level must coordinate the assignment of network numbers in order for packet addresses to be meaningful throughout the internet. In addition, network numbers provide information about proper routing of a packet to its destination, e.g., which subnet on which network a particular host sits. In general, the routing information contained in a network number pertains to the physical location of the destination. When networks cross organization as well as geographic boundaries, *logical* information is desired in addition to *topological* information. In other words, it would be nice to know the organization domain to which a message is being sent, and from whence it came, in addition to the physical locations.

Currently, network numbers in the Internet are allocated to sites by a centralized number czar. Each site may then allocate numbers to hosts and even subnets that lie within its topological network. Most of these hosts and subnets are within the confines of a single organization, but some are not. For example, MIT has direct network connections to several local companies. The network numbers of destinations in these companies look like the network numbers of other MIT subnets because they contain topological information for routing purposes. In order to discriminate between subnets and hosts that are part of MIT's logical network (i.e., actually belong to MIT) and those that lie outside of the logical network (i.e., facilities in the local companies which are accessible but do not belong to MIT), the gateway must be able to bind the source and destination network numbers in the packet header to the organization affiliations.

These issues were not among the many considered during design of the Arpanet/Internet protocols. At that time, the primary concern was to achieve connectivity and transparency and make network boundaries disappear. Therefore, it makes sense that providing information needed to enforce organization boundaries was not a design requirement. Even if it had been a consideration, the number of competing requirements and constraints on the low-level protocol would probably

[5] Many policies may require more information than just source and destination affiliation. But for simplicity I focus on this information to illustrate the argument.

have led the designers to leave such application-specific information to higher levels. In particular, because routing table size is limited, there is pressure to be able to make routing decisions on the basis of a packet's destination subnet number.

One might try to use the network and subnetwork numbers as a hint to organization affiliation. However, because of the decentralized manner in which networks and subnetworks can establish their own interconnections, these topological numbers do not necessarily map into meaningful logical groupings. For example, MIT might implement a filter in the arpanet gateway to reject outgoing messages with source addresses other than MIT's Arpanet network number, 18. However, if the Lab for Computer Science decides to connect some local company to the LCS ring, according to current practices, that company is assigned a subnet number within net 18. Therefore, the filter would not catch transit traffic sent from that company to non-MIT Arpanet sites.

Of course it is possible to identify the various subnet numbers that are assigned to non-MIT entities and add such information to the gateway filter. However, this is not a general solution because given that such interconnections are established in a decentralized manner, there is no good way of keeping track of these exception cases without centralizing the interconnection and number allocation process in some way. One approach might be to establish guidelines that set aside blocks of numbers to be used for non-MIT sites. However, it is hard to know a priori how many such numbers to set aside, and exactly what groupings one will want to be able to distinguish between, i.e., MIT/non-MIT is only one relevant distinction. If connections are centrally managed, or otherwise easy to track, then it may be feasible for the gateway to maintain a list of allowable host and/or subnet addresses and thereby implement packet-level controls.

An example of a packet-level ION gateway that implements usage controls is the University College London (UCL) network connection to the Arpanet. Although the method of control gets around the problems described above, it does not provide a general solution. The UCL network employs two gateways to the Arpanet. One connection forwards packets via a private satellite network to the Arpanet. The second connection forwards packets via an X.25 connection over public packet switched networks. Due to PTT regulations, only Ministry of Defense traffic can be sent via the private satellite path, while civilians (such as university researchers) must send traffic via the public-network path. Because only routing information is available at the IP level, the restriction is enforced by making UCLnet appear as two separate networks, UCLnet and PSSnet.

This is achieved by splitting the namespace in two and assigning addresses to MOD and civilian hosts accordingly. Because there is a small and fixed number of user groups (i.e., two) the mechanism works. A similar mechanism may be employed by MIT to restrict undergraduate student access to the Arpanet. Most undergraduate access to computational facilities and the MIT network will occur via the MIT subnets that belong to project Athena.[6] The Arpanet gateway will simply reject all packets originating from those subnets.

3.2. Visa Scheme

Given that logical information generally is not deduceable from packet headers alone, we can adopt an approach first suggested by D. Reed and documented in [8]. This scheme requires that the source carry out a higher level dialog with a policy server in the destination network in order to authorize a particular conversation (e.g., mail, file transfer, etc.). The policy server passes the authorization information to the packet-level gateway along with a means of authenticating the authorized traffic (e.g., an encryption key). The scheme is referred to as a visa scheme because gateways are analogous to border crossing stations, access control servers to embassies, and the keys to visas.

In this scheme, in order for a source to send a packet or set of packets via an ION gateway, the source must obtain a key from the access control service (ACS) of the network that it wishes to enter; i.e., the owner of the gateway. If the source passes by the ACS's policy filter, the ACS gives the source and the lower-level gateway a key with which to authenticate authorized packets as they pass through the gateway. The key may simply be a ticket appended to each packet header or it may be an encryption key used to calculate the packet checksum. For example, using an encryption key, the ION gateway records which keys correspond to which network numbers. The gateway looks up the key corresponding to the source network-number of incoming packets and calculates the checksum using the key. If the checksum is properly computed then the gateway knows that the packet has been authorized by the ACS of the destination network. If the organization's policy requires that the gateway descriminate according to the destination of each packet in addition to the source, the ACS can include this in the authorization information as well (e.g., the individual destinations that each source can send to, the type of service, the amount of service requested, etc.). In fact, the ACS can be programmed to carry out a wide variety of policies.

[6] Athena is a university-wide experimental project in the use of computers in education.

In summary, in this scheme we offload to the ACS the mapping of traffic affiliation to access privileges by requiring the source to have a higher-level dialog with the ACS before its packets are able to enter the network. Once the ACS provides the low-level gateway with the access and authentication information, the low-level gateway operates as if it were itself equipped with the category information and a way of mapping traffic information to that category information. In effect, by granting a visa to a source-destination pair[7], and informing the gateway of the granting, the ACS wraps a set of individual packets into a logical unit that is then subjectable to policy control in the packet-level gateway.

Such a scheme has been proposed for a dial-up, packet-forwarding gateway to the MIT network. This gateway is connected on one side to the public switched telephone network, and on the other to an MIT local network. Although a single physical gateway is used, MIT would like to apply different access policies to the different groups that use it. Some MIT resources are intended for access by members of the MIT community only (e.g., gateways to other networks, a New York Times clipping service, high-speed printers, etc.). Other resources are intended for access by some non-MIT users as well. In order to implement non-discretionary controls as proposed in [4], the gateway would operate as follows. When a user calls the gateway, the gateway associates the call with a particular port and accepts packets from the user only if they are addressed to an ACS. The external user carries out a high-level dialog with the ACS and authenticates him or herself. After authenticating the user and identifying the internal facilities that the user is authorized to access, the ACS sends the gateway a list of destination addresses to which the particular user should be allowed to send packets; these addresses represent capabilities or categories. In addition, the ACS sends a key to both the user and the gateway. The gateway associates both the list of destinations and the key with the port assigned to the user. The key is used as a *connection-authenticator* for the duration of the connection. In order for the gateway to accept a packet through the port, the checksum of the packet must have been calculated including the connection-authenticator that is currently associated with the port. When the user first dials up the connection-authenticator is zero and the user can send packets to only a single destination, the ACS.

A packet-level gateway together with an ACS can effect higher-level controls. However, if the

[7] The visa may actually be granted to different units of authentication, such as source-destination-service type, for example.

ION is intended to support a small set of higher-level applications, and if performance benefits of packet-level interconnection are not significant, it makes sense to consider interconnection at a protocol level at which the policy-related information is available directly.

4. Level of interconnection

In the introduction I gave two motivations for treating entities on the other side of an ION gateway differently from those within an organization's internal network. First, policy concerns may require that non-local users be restricted from using some internal resources and other gateways. Second, performance concerns may require that information needed for local, tightly-coupled applications, not be broadcast through an ION gateway across which applications are more loosely coupled. In the discussion of network numbers I explained that the logical information needed to implement intelligent filtering in an ION gateways is not available at the packet level. A visa scheme was described that can solve the problem for many simple usage control policies. But, when policy decisions are dependent on higher-level information that cannot be bound to packet-level information, higher-level connections are may be more practical. The advantage is that the information needed to evaluate policy, such as organization affiliation, service type, size of request, etc., is available to higher level protocols; i.e., these protocols deal with aggregated units of traffic that contain more semantic information in the headers and control fields (e.g., electronic mail messages, remote login or file transfer connections, etc.). Even with a higher-level ION gateway, some controls are best implemented in the endpoints themselves; in particular, controls that discriminate according to the content of a message, e.g., the size of a purchase order, or the name of a file requested. In any case, these endpoint applications may require some controls to isolate the ION processes and applications from the non-ION ones.

Higher-level gateways require that higher level protocols be terminated at the gateway. More information about the application of the connection is available at these higher levels. Depending upon the level of connection and application, this information may include the logical affiliation of source and destination, the actual service being performed, and the amount of communication resources requested, for example. Although precautions must be taken to verify the correctness or credibility of this information, the point is that it is available for evaluation. For example, Harvard University is connected to the Arpanet via a packet-level gateway. Harvard would like to allow any university member with a computer account to send electronic mail via the Arpanet gateway, but at the same time it wants to provide file transfer and remote login to select groups of users only.

Currently Harvard is able to control remote login use because for internal resource control purposes it does so anyway within the internal network. Remote login is made a restricted command on all internal hosts and because only certain users can use it internally, only those users can use it through the gateway. However, file transfer is not a restricted command; it is too common and useful a facility to even consider restricting internally. As a result, there are no controls on doing file transfer via the gateway. Because the gateway is an IP level packet-forwarding gateway, such controls would require modifying the kernel of the gateway and has other negative implications such as being incompatible with other sites running the gateway code. If the gateway were an application level gateway, it would be a much simpler task to modify it so as to restrict file transfer use to certain users only.

This discussion of level of interconnection is concerned most directly with what Sunshine referred to as service level and implementation approach. [10] Service level refers to the communication mode supported in the gateway, e.g., datagram, virtual circuit, file transfer, remote login, mail, etc. He classifies implementation approach as either endpoint (where the source and destination each act as an endpoint in the communication mode and each gateway passes lower-level information) and hop-by-hop (where each intermediate gateway acts as an endpoint of the communication mode as well). I refer to the former as lower-level, or packet-level, and the latter as higher-level, or application-level, interconnection.[8] With respect to traditional interconnection concerns alone (not inter-organizational concerns), Sunshine finds the hop-by-hop (higher-level) approach to be more appropriate where backward compatibility of protocols and immediate needs predominate and where user awareness of crossing network boundaries is acceptable. Whereas he finds the endpoint approach (packet-level) preferable when robustness and generality are important and there is more basis for agreement and conformance to standards. Sunshine's conclusions lend support to the argument in favor of higher level (i.e., hop-by-hop) ION connection. When an organization interconnects to the outside backward compatibility with internal protocols and procedures is still of primary importance. Similarly, expediency is often a key criterion for the interconnection and it is often desirable for the connection to be less than transparent so that insiders are conscious of their actions when communicating with outside entities.

[8] Application level refers to even higher protocol levels. For example, where a high-level gateway would forward a file transfer request without looking at the content of the request, an application-level gateway would interpret the request itself.

Two types of higher level gateways can be distinguished--connection based and message based. Those applications that operate on top of connections require that the gateway set up a connection between machines on either network, whereas message based applications require only that the gateway forward messages between two applications. In both cases, the unit of transfer and therefore of control makes accessible more policy-related information than does an individual packet. In addition to being more complex, connection based gateways and applications may introduce more vulnerabilities than do message-based gateways. Connection based gateways establish connections between entities on either side of the gateway and then ship undifferentiated packets back and forth via that connection. Although controls may be added to the connection set-up process, unless controls are applied on a per packet basis, there is more chance for a connection that was approved initially to be used for some undesirable purpose. A message-based gateway applies controls to each message that passes through it while avoiding the cost of applying controls to each packet. If a machine automatically processes the messages, similar attacks may be made by outsiders by embedding executable commands or special control characters within the text of the message. However, the staged delivery of messages makes it easier to guard against such attacks by filtering traffic. On the other hand, message based gateways may not be well suited to applications that are sensitive to delay. The X.400 and X.75 gateways, mentioned below, are examples of a message-based and connection-based interconnection protocols, respectively.

There are many examples of higher-level gateways in use today. A few examples are described below:

* As illustrated earlier, a host connected to two different networks can act as a high-level gateway between them. For example, the host might forward electronic mail between users on either network using the hosts local mail facilities. Similarly, a remote user logs in to the host (perhaps via the internal network and public-network gateway of the user's own organization), and from there might establish another connection to some other host on the internal network. In both cases, the host acts like an endpoint in the mail or remote login protocol.

* X.75 is a CCITT standard for interconnecting X.25 packet switched networks and operates at a connection (virtual circuit) level instead of a packet level. When two X.25 networks are interconnected, each gateway acts as an intermediate endpoint of the inter-network virtual circuits. X.75 implements some controls at this level of interconnection that might be useful to IONs; in particular, a transit bit to indicate transit networks, and closed user group settings to identify limited-distribution dialogs. [12] X.400, an international standard for electronic mail, will define an even higher level of interconnection. [11]

* Cambridge university uses an X.25 gateway to connect their local area network to a public packet-switched network. The gateway implements access controls by checking all connections against a database of authorized users. Similarly, accounting information is collected in the gateway. [2]

* The UUCP based network operates at a higher level than packet forwarding. [5] Electronic mail, mailing list digests, and sometimes files are transferred between hosts or networks of hosts via telephone connections. It is relatively easy to add filters to the forwarding of UUCP network traffic. Filters may determine such things as which mailing lists are distributed, and which types of services are provided to each of the neighboring UUCP sites. Although the UUCP network has several undesirable characteristics such as excessive delays and difficulty in determining routes, it is useful as an example of a network that operates at a higher level than packet forwarding.

Even at higher levels it is possible to blur the distinction between routing information and information useful for determining organization affiliation. If an organization supports transit, and guidelines are not set for the structure of source and destination names, routing information can be confused with logical information. For example, if person Smith at Symbolics Inc. sends mail to DEC via MIT's network, then if DEC receives the source address as Smith.DEC%MIT, DEC must somehow figure out that the logical affiliation is Symbolics and not MIT. At the same time DEC must be able to determine how to return a message, namely via MIT. UCLnet has experienced this problem in its mail connections to the Arpanet. As described earlier, mail from MOD and civilian users must be treated differently. Some mail is forwarded to the Arpanet from other civilian research networks and hosts that are connected to UCL. The addresses assigned to mailboxes on these hosts are constructed so that the mailboxes appear to lie within UCL. The mail gateway relies on a list of registered users to filter mail, in part because the organization affiliation of a user is not necessarily evident from the header. [1] However, in general textual mailbox names carry more semantic information and are taken from a larger namespace than network numbers. Therefore textual mailbox names can be constructed in such a way that the correct affiliation can be interpreted using easy to follow guidelines.

4.1. Example

The following example illustrates an existing environment in which most of the issues described earlier arise. I describe the policy requirements and how they could be addressed; most of these mechanisms have not been implemented to date.

MIT has installed packet-level connections to a few local companies. It is also considering a mail level connection to another firm. MIT is connected to the Arpanet via a gateway and via a separate time-sharing system. In addition, this system serves non-MIT users and is also connected to BITNET and Mailnet. The following controls could be implemented to address various MIT and non-MIT policies.

1. The packet-level connections should allow the firms to use *select* MIT facilities as well as the Arpanet gateway. Since each gateway is used by a single outside organization the binding between service or user and organizational affiliation is fixed. Based on this permanent binding, the gateway could implement non-discretionary controls by simply recording the network addresses of allowable destinations and rejecting packets addressed to all others. Because the destination network addresses are managed internally the binding of destination to organization affiliation is fairly easy to evaluate. However, controlling the flow of outgoing traffic, or the flow of incoming traffic via gateways shared among multiple organizations, is not so easy since the source and destination affiliations are not known by default nor are they deduceable from the network address.
2. The mail connection should forward mail to and from select MIT sites only and no gateways. It is not difficult to implement this policy since the information necessary is available in the mail headers. Authentication mechanisms can be used to increase the trust-worthiness of the header information.
3. The time-sharing system users that are affiliated with MIT have the same rights as any MIT user. Users that are from outside MIT, should be restricted from using packet level Arpanet connections, and even from mail connections. Both restrictions can be realized by controlling access to the respective commands. It is not difficult to implement this policy if organization affiliation is identifiable by evaluating account and mailbox names (e.g., Smith.INC.Computer-services.MIT). This requires that naming guidelines be established and followed by all hosts that act as gateways.

5. Conclusion

In order to implement non-discretionary controls in ION gateways, I have argued that certain logical information must be available. I discussed two types of interconnection and discussed their suitability to ION applications. I conclude that higher level connections are preferable for many ION applications. And where lower-level connections are adopted for performance or generality, a visa scheme can be used to support many simple usage controls policies.

Aside from the implications for network design and management, an interesting aspect of this discussion is that policy issues considered explicitly or implicitly during the design of a technology, shapes not only the technology, but the effects that that technology has on the organizations that use it. For example, given the lack of consideration for control requirements for network protocol developments, the most common way to connect networks across organization boundaries introduces far more boundary penetration than is necessary, or often desirable.

Acknowledgments

Discussions with J.N. Chiappa, D. Clark, D. Reed, and J. Saltzer contributed to the ideas presented here.

References

[1]

Cole, R., Higginson, P., Lloyd. P., Moulton, R.
International net faces problems handling mail and file transfer.
DataCommunications :175-187, June, 1983.

[2]

Dallas, I.
Implementation of a Gateway between a Cambridge Ring Local Area Network and a Packet Switching Wide Area Network.
In Williams, W., editor, *Pathways to the Information Society: Proceedings of the 6th International Conference on Computer Communications*, pages 137-142. North-Holland, September, 1982.

[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]

Estrin, D.
Non-Discretionary Controls for Inter-Organization Networks.
In *Proceedings of the 1985 Symposium on Security and Privacy.* IEEE Computer Society Press, Silver Spring, MD, 1985.
Accepted for Publication

[5]

Horton, M.
Standard for Interchange of USENET Messages.
Network Working Group Request for Comments RFC 850, DARPA Network Working Group, June, 1983.

[6]

Mockapetris, P.
The Domain Name System.
In Smith, H., editor, *Computer-Based Message Services.* Elsevier Science Publishers B.V., North-Holland, 1984.

[7]

Mogul, J.
Internet Subnets.
Network Working Group Request for Comments RFC 917, DARPA Network Working Group, October, 1984.

[8]

Mracek, J.
Network Access Control in Multi-Net Internet Transport.
B.S. Thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, June, 1983.
Supervised by D.P. Reed

[9]

Saltzer, J.
On the Naming and Binding of Network Destinations.
In Ravisio, P.C., Hopkins, G., Naffah, N., editors, *Local Computer Networks.* North-Holland Publishing Company, 1982.

[10]

Sunshine, C.
Interconnection of Computer Networks.
Computer Networks (1):175-195, 1977.

[11]

CCITT.
Recommendation X.400.
CCITT Draft Standards Recommendation, CCITT, 1984.

[12]

CCITT.
Terminal and Transit Call Control Procedures and Data Transfer System on International Circuits Between Packet-Switched Data networks: Recommendation X.75.
CCITT Standards Recommendation, CCITT, 1978.