

Network Working Group
Request for Comments: 1124

B. Leiner
RIACS
September 1989

Policy Issues in Interconnecting Networks

Status of this Memo

To support the activities of the Federal Research Internet Coordinating Committee (FRICC) in creating an interconnected set of networks to serve the research community, two workshops were held to address the technical support of policy issues that arise when interconnecting such networks. Held under the auspices of the Internet Activities Board at the request of the FRICC, and sponsored by NASA through RIACS, the workshops addressed the required and feasible technologies and architectures that could be used to satisfy the desired policies for interconnection.

The purpose of this RFC is to report the results of these workshops. Distribution of this memo is unlimited.

Preface

This report documents the results of two workshops held at the request of the Federal Research Internet Coordinating Committee and under the auspices of the Internet Activities Board. As such, this report represents the work of a large number of people (listed in Section 7). Without their efforts, these results would not have been possible. The author (really more of an editor) would like to acknowledge their efforts and contributions, and thank them for their cooperation in making the workshops a success.

1. Introduction

Computer networking has become pervasive and basic to the conduct of scientific and academic activities. To provide the needed networking support to these activities, each of the agencies funding research has proceeded to establish one or more agency funded computer networks.

Recognizing the importance of such networking support, the Office of Science and Technology Policy (OSTP) working with the appropriate personnel from the research-funding agencies on the Federal Coordinating Council on Science Engineering and Technology (FCCSET) Committee on High-Speed Networks developed a set of recommendations for the evolution and enhancements of scientific and academic networks. These recommendations are described in three phases. The first phase addresses the interconnection of the various agency networks into a ubiquitous networking capability serving several hundred universities and research institutions with a backbone network operating 1.5 Mb/s. The second phase involves upgrading the network backbone to 45 Mb/s and connecting additional universities and other research institutions. The third phase involves the development and installation of a high bandwidth (Gb/s) networking capability.

The motivation for the first two phases are to achieve good performance in a cost effective manner. The scientific and academic community is best served by an interconnected ubiquitous networking capability rather than a set of partitioned networks supporting only subsets of the community. Costs can be reduced and performance improved through sharing of resources and using cross-support (e.g., using one agency's network to serve an institution for another agency's purposes rather than having to connect each institution to every network.)

To accomplish these objectives, the Federal Research Internet Coordinating Committee (FRICC) was formed. Consisting of representatives from the key research agencies (NSF, DARPA, NASA, and DOE), this ad hoc group has been developing strategies for interconnection of networks and evolution of the Internet in accordance with the OSTP recommendations for Phases 1-3. In the process of developing such plans, it became apparent that a set of issues needed to be addressed concerning the various agency policies for their research networks in light of the desire to interconnect such networks.

This report documents the results of a series of two workshops (18-20 June 1988 at NASA Ames Research Center and 8-10 November 1988 at MIT) held to address these issues. Held under the auspices of the Internet Activities Board (IAB) at the request of the FRICC, and sponsored by NASA through RIACS, the workshops addressed the required and feasible technologies and architectures that could be used to satisfy the desired policies for interconnection.

The issues were divided into four categories, and working groups established within the workshops to address each area. The first working group addressed the policies themselves. Working with the members of the FRICC, the initial statements

of agency policies were refined so that the rest of the workshop attendees could better understand the desired and required policies. The second working group addressed issues associated with access control to network resources. The third working group addressed the techniques required to support the sharing of networking resources in accordance with agreed upon policies. The fourth working group focussed on the end-to-end services required to support an interconnected set of networks.

Each of the working groups prepared summary reports of their deliberations. These reports are contained in Sections 3-6 of this document. The report of the policy working group attempts to summarize the existing policies of each of the agencies, particularly with respect to interconnection with other networks. The other three working groups focussed on the technology issues needed to be addressed in light of those policies. In each case, the working group report discusses the issues and develops an evolutionary capability with the goal of fully addressing the agency policies. Summaries of these reports are contained in the next section.

It is hoped that the results documented in this report will help the FRICC and the rest of the research community in achieving this exciting objective: a national research networking capability.

2. Workshop Summary

Driving the workshop were the policies of the individual agencies and a desire to interconnect the networks in a way that was satisfactory to those agencies. A prime policy driver appeared to be OMB Circular A130, which states that appropriate mechanisms must be used to assure some level of accounting for the use of the various networks. Another important policy driver was the need for agencies to assure that sharing of networks did not adversely impact the support of the individual agency users on their specific networks. This led in some cases to the need to be able to dedicate a portion (sometimes all during a specified time period) of an agency network to supporting its own users. Finally, the need to provide appropriate supporting end-to-end services, including security issues, led to the need for coordinating such services.

To facilitate the discussion of the technology issues and the presentation of results, it was decided to describe the evolution of capability in four phases. Phase 0 represented currently deployed and available capability. While not necessarily being currently used for the support of the policy issues, the capabilities of Phase 0 were viewed as being currently available and could be used starting today. Phase 1 consisted of capabilities that were developed and deployed at a limited number of sites. Thus, the issues involved in using such capabilities involved mainly those of widespread deployment (plus perhaps some limited amount of development associated with, e.g., porting of software). Phase 2 represented capabilities that were relatively well understood (little research required) but would require development activity before they could be used to support the policies for interconnection. Phase 3 capabilities require research to achieve, and thus represent the most future capability.

While these phases of capability represent evolution in availability, they should not be viewed as evolution in starting time for action. In all cases, research and development activities would have to start today in order that these capabilities be available in a timely manner.

As the working group on access control discussed the required technologies and mechanisms, it became clear that an important technology driver was the need to label packets with the appropriate information to make determinations of routing and resource allocation internal to the interconnected networks. For example, if certain links in a NASA network was to be restricted to use only by NASA users (even if accessing the network through an NSF network), it would be necessary to provide such labelling information in the packet. The report of the working group discusses the information that needs to be carried in such labels, requirements for authentication, and some potential experiments and development that should be carried out to achieve the required capability.

The working group on resource sharing focussed on the technologies that would allow fair sharing of resources between the participating agencies. The key issue that emerged from the discussions of this working group was the need to develop global

algorithms that permitted sharing and prioritization of the use of resources. As an example, it is relatively easy for an agency to block low-priority traffic from traversing its network during a period of high internal requirement. It is not so easy to do so and assure that the external users still can receive the resources they need from the interconnected internet.

The working group on end-to-end services focussed on those services that are required from a user's perspective from the overall system, and need to be coordinated across the interconnected networks. For example, directory and security services must be provided across the interconnected system. The key element emerging from the group discussions was the need to establish a consistent set of mechanisms to interconnect the various end-to-end services. These must be provided in a secure manner to assure that the security services fulfill their function.

The working groups identified the need to carry out supporting experiments and analysis to carry forward the interconnection of the networks, e.g., to make decisions about the need for stream versus transaction support. Each group developed a set of possible experiments and activities in accordance with the phases of development discussed above. These are summarized in Tables I-III.

A number of possible follow-on activities were identified to be passed on to the various Task Forces of the IAB. These are shown in Table IV.

In summary, the workshop identified a number of critical issues and identified areas where further research and experimentation is required. It is hoped that these results help provide a "road map" for how to satisfy agency policies and requirements in the interconnection of networks.

Table I
Access Control Projects

Phase 0	Access Control based on source/destination access matrix (for traffic not transiting network)
Phase 1	<i>Statspy</i> experiment to determine and define requirement for transactions “ESnet hack” for limited access control based on source/destination addresses. “Xerox hack” for limited access control based on source/destination addresses.
Phase 2	Coloring of stream packets Simple colors/labelling Route filtering for access control using source/destination addresses Incorporate “Xerox hack” into other gateways Authentication and signature architecture
Phase 3	Use of complex credentials Use of policy gateways in route computation

Table II
Resource Sharing Projects

Phase 0	Simple route filtering
Phase 1	Run <i>Statspy</i> to determine source/destination traffic flows (to comply with A130 traffic monitoring requirements)
Phase 2/3	50/50 resource management for link sharing Color packets and observe behavior to improve traffic monitoring Fast encryption of route and certificate packets, to secure traffic monitoring and control Fast mapping from source/destination to packet label/color Demonstration of gateway using soft state Define and support policy source routing Synthesis of source route Management controls and protocols Composition of policy terms Define and structure route set-up protocols

Table III
End-to-End Services Projects

Phase 0	User/process authentication using passwords (origin authentication)
	Mail relays for both function and system isolation
	Name domains system for host name to address mapping
Phase 1	User/process authentication using challenge/response or some other protocol (origin authentication)
	Secure-ID or other authentication technologies
	Challenge/response technologies (overlaps with the previous line)
Phase 2	Kerberos (authentication server)
	Authentication using certificates
	Integrity (MACs, checksums) and labelling
	Key distribution and management
	Secure mail (see RFC 1113)
	Certificates (see same RFC)
	Security of distributed white pages
Phase 3	Integrity labelling, tools (MACs, checksums)
	Distributed white pages for the entire Internet
	Use of VISAs
	Certification across peer domains
	Distributed computation
	National file system
	Trusted accounting
	Firewalls for end-to-end services
	Integrity of data across international boundaries with agreed upon cryptographic technologies
	Use zero-sum knowledge to have a third party to assure integrity without secrecy for such cases

Table IV
Projects for IAB Task Forces

ETETF	Handling of quality of service in gateways
ANTF	Phases 2 and 3 of resource sharing activities
IETF	Policy routing
Privacy	End-to-end privacy services
???	End-to-end services

3. Working Group on Interconnection Policies

Working Group 0 Members

Steve Wolff (Chair)	NSF
Guy Almes	Rice
Matt Bishop	Dartmouth
Brian Boesch	DARPA
Scott Brim	Cornell
Phill Gross	NRI
Dan Hitchcock	DoE
Russ Mundy	DCA
Tony Villasenor	NASA

Network resource sharing is encouraged by the potential for economies of scale both in communication link acquisition cost and in provision of value-added network services (the latter not yet demonstrated in the Internet, but consistent with telephone company experience); it is suggested by the Congressionally-ordered network study that resulted in the OSTP report *A Research and Development Strategy for High Performance Computing*; and it is mandated by OMB Circular A-130. Technical forces in the same direction include the additional connectivity each agency provides to its clients (actual or potential) by acquiring the use of nets belonging to other agencies at little or no additional cost, and the robustness afforded by the sharing of redundant paths or other forms of “excess” capacity.

The agencies represented on the FRICC, however, have differing missions and requirements, and these differences are reflected in differing rules and procedures for network usage. WG0 was created to explicate the rules for network use of the FRICC agencies, for those rules -- particularly the differences among them -- form the foundation upon which the technical specifications of “policy-based routing” must be built. This report, therefore, is the primary input to the technical Working Groups WG1, WG2, and WG3.

Making all FRICC agencies’ network use rules the same is NOT a goal of WG0. Each FRICC agency has more-or-less well-formulated rules for the use of its network in the absence of explicit interconnection with other networks and the attendant “foreign” traffic. These rules are given below. Currently, no agency has rules for interconnection with:

- networks of other FRICC agencies,
- networks of other countries,
- commercial networks, or
- “sensitive” networks (e.g., SDInet, NASA mission-critical nets);

consistent formulation of such rules will be discussed in future FRICC meetings.

It was however noted that, in dealing with subordinate (not peer) networks, NSF has required traffic presented to the NSFnet backbone to conform to NSF rules of acceptable use; DoE on the other hand is tending to the more liberal policy of carrying any traffic that meets the rules for acceptable use of the agency network offering the traffic.

3.1. Existing Policies, Summarized

The following is a summary of the existing policies for network usage of the FRICC member agencies.

NSF (draft, summarized):

- Purpose is to support scientific research and other scholarly activities.
- Use to support research or instruction at not-for-profit institutions of instruction and/or research is acceptable, whether all parties to the use are located or employed at such institutions or not.
- Activities in direct support of acceptable use are acceptable.
- Use for research or instruction by for-profit institutions may or may not be acceptable, and will be reviewed case-by-case.
- Commercial use by for-profit institutions is generally not acceptable.

DoE (draft, summarized):

- Use in which at least one party is supported by Energy Sciences funds is acceptable.
- Use by persons at DoE sites is acceptable, even if they are not supported by Energy Sciences funds.
- Advertising or promotional activities are not acceptable.
- Use in direct competition with commercial services is not acceptable.

NASA (draft, summarized):

- Purposes are to support NASA space science programs, to support collaborating science activities (e.g., with ESA, NOAA, USGS), and to support NASA contractors (e.g., those involved in building scientific sensors and spaceborne hardware).
- Other activities may be supported on a case-by-case basis, provided there is no impact to the NASA programs.
- No Eastern bloc access.
- Shared use of network facilities must be controllable and annually accounted for.
- NASA networking facilities may be made available for other uses and users on a cost-reimbursable basis.

- Direct competition with commercial services is not acceptable.

DARPA:

- Purpose is to support network research and other DARPA research objectives.
- There may be “forbidden routes” for some traffic.

DDN (excluding ARPANET and the proposed DRI):

- Use is for DoD business only, unless otherwise approved by JCS.
- All connections to other nets strictly regulated by mailbridges (now) or trusted guard gateways (future).
- Facilities must comply with DoD Security Architecture and with DoD Directive 5200.28 which requires C2 certification for sensitive unclassified information.

3.2. Refined Policy Statements

As a result of the first workshop discussions on policy, Dr. Cerf met with the various agency representatives to refine the policy statements. The results of these meetings were as follows. Note that these statements are those of the workshop and do not represent official agency policies. Each policy is represented in Clark's Policy Term (PT) notation¹ and then described in English. The standard Clark Form for PTs (Hsrc,ARsrc,ARent)(Hdst,ARdst,ARexit){UCI}{Cg} FRICC={DOE,NASA,DCA,NSF} where H=Host, AR=Autonomous Region, src=source, dst=destination, ent=entry (previous hope), exit=exit (last hop, F=Federal Agency Net, Re=Regional, U=University, Co=Commercial Corporation, and Cc=Commercial Carrier. All PTs are assumed to be symmetrical in these examples.

NSF

NSF1: (*,*,{F/Re})*(*,*,{F/Re}){research,support}{unauthenticated UCI, no-per-pkt charge}

i.e., NSF will carry traffic for any host connected to a F/Re network talking to any other host connected to a F/Re via any F/Re entry and exit network, so long as it is being used for research or support. There is no authentication of the UCI and no per packet charging. NSFnet is a backbone and so does not connect directly to universities or companies. Thus the indication of {F/Re} instead of {F/Re/U/Co} as ARent and ARexit.²

NSF2: ({User svcs, Expert Svcs}, {NSF},{F/Re})*(*,{F/Re},{F/Re})

i.e., NSF will carry traffic to user and expert services hosts in NSF Autonomous Region (AR) to/from any F/Re AR, via any F/Re AR. These are the only things that

¹D.D. Clark, "Policy Routing in Internet Protocols," Version 1.1, May 19, 1988.

² Note: I can't actually decide whether it should be as stated above or (*,{F/Re},{F/Re})*(*,{F/Re},{F/Re})

directly connect to NSFnet.

DOE

DOE1: (*,DOE,-)(* ,*){research,support}{unauthenticated UCI, no-per-packet charge}

i.e., DOE will carry traffic to and from any host directly connected to DOE so long as it is used for research or support. There is no authentication of the UCI and no per packet charging.

DOE2: (* ,*,{F/Re})* ,*,{F/Re}){ }{unauthenticated UCI, no-per-pkt charge}

i.e., DOE will carry traffic for any host connected to a F/Re network talking to any other host connected to a F/Re via any F/Re entry and exit network without regard to the UCI. There is no authentication of the UCI and no per packet charging. (In other words, DOE is more restrictive with its own traffic than with traffic it is carrying as part of a resource sharing arrangement.)

NASA

NASA1: (* ,*,*)(*,NASA,-){NASA-research,support}{unauthenticated UCI,no-per-packet-charge}

i.e., NASA will accept any traffic to/from members of the NASA AR, but no transit. No UCI authentication and no per packet charge.

NASA2: (* ,{F},*)(*,{F},*){research,support}{per-packet accounting, limited to n% of available BW}

i.e., NASA will carry transit traffic to/from other federal agency networks if they are for research and if the total use of available BW by non-NASA Federal agencies is below n%.³

NASA3: (* ,{Co},*)(*,{F/R/U},-){NASA research,support}{not authenticated UCI, no per packet charge}

i.e., NASA will carry commercial traffic to federal, regional, and university ARs for NASA research or support, but it will not allow transit. The particular entry AR is not important.

NASA4: (* ,*,-)(*,*,-){ }{per-packet-charge to recoup cost, limited to n% of available BW}

i.e., On a case by case basis, NASA will consider non-NASA traffic on a cost-reimbursed basis. It will not carry transit traffic on this basis.

³ Note that this non-interference policy type needs some more work in terms of integrating it into the routing algorithms.

DARPA

DARPA1: (*,*,*)(*,DARPA,-){research,support}{unauthenticated-UCI, no per packet charge}

i.e., DARPA will carry traffic to/from any host in DARPA AR from any external host that can get it there so long as UCI is research or support. No UCI authentication or per packet charge.

DARPA2: (*,*,{F/R/U/Co})(*,*,{F/R/U/Co}){research,support}{unauthenticated-UCI, no per packet charge, non-interference basis}

i.e., DARPA will carry traffic for any host connected to a F/Re/U/Co network talking to any other host connected to a F/Re/U/Co via any F/Re/U/Co entry and exit network, so long as it is being used for research or support, and the network is not heavily congested! There is no authentication of the UCI and no per packet charging.⁴

DCA

DDN1: (mailbridge,DDN,-)(*,{F/Re},{F/Re}){research,support}{unauthenticated UCI, all incoming packets marked, per-kilopacket charge}

i.e., DDN will not carry any transit traffic. It will only accept and send traffic to and from its mailbridge(s) and only from and to hosts on other F/Re nets.

An Example Regional⁵

Regional1: (*,{F/Re/U},{F/Re/U})(*,{F/Re/U},NSF){research,support}{unauthenticated UCI, no-per-packet charge}

i.e., The Regional will carry traffic from/to any directly connected F/Re/U network to any F/Re/U network via NSF if it is for a research or support UCI. (NSF requires that all Regional networks only forward to it traffic that complies with its, NSF's, policies!)

Regional2: (*,{F/Re/U},{F/Re/U})(*,{F/Re/U},Cc){unauthenticated UCI, per-kilopacket charge}

i.e., The Regional will carry traffic from/to any directly connected F/Re/U network to any F/Re/U network via a commercial carrier regardless of its UCI. In this case, the packets are charged for since the commercial carrier charges per kilopacket.

⁴ Note: DARPA would like to say something about the need to enter the DARPA AR at the point closest to the destination, but I don't know how to express this.

⁵ Note: No interview was done for this one. This is just a guess.

4. Access Control for Network Switching and Transmission Resources

Working Group 1 Members

Steve Kent (Chair)	BBN
Guy Almes	Rice
Bill Bostwick	Los Alamos
Marsha Branstad	DoD
Vint Cerf	NRI
Deborah Estrin	USC
Tony Hain	Livermore
Dan Lynch	ACE
Russ Mundy	DCA
Anita Holmgren	Unisys

4.1. Introduction

This report reflects discussions among the members of working group with regard to network access control for the National Research Internet (NRI). The NRI will be composed of network resources contributed by various organizations (primarily agencies of the Federal government). The operational model for the NRI is that of a collection of autonomous, administrative domains (referred to as “domains” within this report), each of which manages a collection of network transmission and/or switching resources. (Other, higher level resources also may be shared across domain boundaries, but these are not the focus of the access controls discussed herein.) Some of these network resources are owned or leased exclusively on behalf of the administrative domain responsible for the resource, whereas other resources may be jointly paid for and administered.

There is a perceived requirement that a domain provide access control for the network transmission and switching resources that comprise it. This form of access control is distinguished from measures oriented toward controlling access to subscriber resources, e.g., workstations, file servers, etc. Rather, these measures are intended to apply to communication paths which transit gateways, circuits, networks, etc.

There are several motivations for introducing network resource access controls. The organizations which will contribute network resources or funding for shared resources to the NRI need to be satisfied that sharing of these network resources can be controlled in such a fashion as to accord priority to designated users or groups of users and to account for resource usage in accordance with OMB guidelines. It may be necessary to bill for usage of some resources, especially commercial facilities connected to the NRI. Some organizations have adopted policies that prohibit transport of data from certain classes of users across their networks.

This report examines various aspects of network resource access control measures in the NRI context, including bases for making access control decisions (policy inputs), communication scenarios to be supported, mechanisms for enforcing access control policies, and assurance issues associated with enforcement. Formulation of specific access control policies is outside the scope of this report and is addressed by the report of Policy Working Group.

This report has been prepared by the members of the working group as a result of discussions that took place at workshops sponsored by NASA on June 15-17, 1988 and November 8-10, 1988. Additional inputs have been prepared by working group members during the interval between these workshops and co-ordinated by the chair.

4.2. Access Control Policy Issues

4.2.1. Policies and Models

Any discussion of access control measures should begin with a characterization of the policies which the measures are to enforce, and a definition of the model that underlies the policies. There are various ways to characterize access control policies, one of which (ISO 7498-2) considers two axes: 1) the basis on which access control decisions are made (rule-based or identity-based), and 2) the entity who defines the policy (user-directed or administratively directed). For the NRI environment, we anticipate the policies are all administratively directed since they represent constraints imposed by organizations which contribute resources to the NRI, not individual subscribers.

Discussions with organizational representatives suggest that both identity-based and rule-based policies may be employed. For example, in some circumstances an access control decision will be made based on the identity of the user (or a class of which the user is a member) requesting access. In many cases, possession of a token indicating agency authorization for resource use, perhaps coupled with time and day of week inputs, will form the basis for the access control decision. These two examples illustrate identity-based and rule-based policies and policies that combine both policy bases are also possible.

The security access model we assume for the NRI environment is a traditional one involving subjects and objects. Subjects are active entities (e.g., processes) which are accorded some access privileges with respect to objects. The processes execute in various subscriber equipments (hosts, workstations, servers, etc.) either acting on behalf of users (individuals or groups) or acting as entities independent of any specific, human user. Objects in this context are typically data paths through the NRI, and thus they implicitly entail the use of transmission and switching resources. (Alternatively, we could consider these resources individually as the objects and the paths as compositions of the component parts.)

4.2.2. Policy Inputs

A refinement of policy characterization is provided by considering the range of inputs on which access control decisions will be made. These inputs can be divided into two categories (somewhat arbitrarily): 1) data implicitly available to the enforcement entities, e.g., time and date or utilization and connectivity status, and 2) data explicitly provided by subjects, e.g., in packet headers. Note that this characterization does not specify whether the explicit inputs are provided in every packet or only in some packets, how the inputs are validated, etc. These details are critical components of an architecture, not just an implementation, and thus the final form of this list should take into account these considerations as well as the rationale provided below.

Based on inputs from agency representatives present at the workshops, it appears desirable that information on local resource utilization and global connectivity be major implicit inputs in access control decisions. The rationale is that many agencies appear to be adopting policies which permit sharing of resources by “outside subjects” on a “non-interference” basis. This requires that the enforcement mechanisms be cognizant of the resource utilization status (congestion measures) so as to determine what constitutes non-interfering sharing.⁶ It also requires some explicit identification of subjects to determine whether the non-interference criteria should be applied. More refined sharing policies could take into account relative priorities for various subjects, type of service (TOS)-based routing decisions, etc. The Resource Sharing Working Group is focusing on routing issues which take into account quantitative measures related to TOS. In contrast, this group has focused more on policies in which such quantitative measures are not primary inputs to the access control decision. This suggests that a combination of the architectural proposal from both groups will be required to address some of the access control policy requirements described at the workshops.

Data that might be explicitly required from a subject was the topic of much discussion. A list of candidate data items was developed and is discussed below. Although not all administrative domains might require all of these inputs for an access control decision, it has been suggested that the list be universally agreed upon among all domains. The argument is that global routing determinations are affected by local access control decisions and that it is desirable to enable subscribers (or their local policy route servers) to calculate permitted routes before initiating transmission of data along a path. In order to perform such calculations, each domain must publish its access control policy and the inputs to the policy must be universally interpretable. Thus there is a strong motivation to define a minimum set of explicit inputs to these

⁶There is a potential conflict here in using local congestion measures as inputs to an access control decision. It is desirable for a remote subject (e.g., policy controller) to determine in advance if a specified transmission resource can be used in constructing a (policy) route between two points in the NRI, for reasons elucidated by Dave Clark in his policy routing paper. Thus the conflict arises if either the remote subject cannot obtain the necessary local congestion measures or if these measures are very dynamic.

policies.

At one point in the discussion it was suggested that any inputs to access control decisions that were not universally interpretable could be accommodated by allowing for “domain specific” data items. Such data items would be interpreted by only a few domains (perhaps only a single domain) along a route. However, we note that this concept does not seem to be in concert with the principle cited earlier (and discussed in Clark’s paper), i.e., subjects should be able to predict access control decisions for any domain through which they might construct a route. Thus the concept of a domain-specific access control data item as an “escape” mechanism for including additional inputs to access control decisions may not be appropriate. Recall that no domain is required to employ all the supplied inputs in making an access control decision and thus inclusion of a data item in a widely known collection need not impose on domains that do not wish to make use of the data item.

Since the administrative domains often represent federal agencies (e.g., DOE, NASA, NSF), it was perceived that there should be some means of representing an agency’s granting authorization for resource use to the subject. This might be a hierarchic data item, specifying both an agency identifier and further defining the subject’s privileges as granted by the agency. For example, an agency such as DOE might grant somewhat different privileges to its employees, to its grantees and their staff, and to other individuals engaged in work that is viewed as supportive to the agency mission (though not necessarily funded by the agency). This effect might be achieved by issuing to each of these subjects credentials that specify some form of affiliation with the agency in question but with different qualifiers, depending on the nature of the affiliation. Thus we envision a compound access control data item that will specify an AGENCY AFFILIATION INDICATOR, consisting of an AGENCY ID and AFFILIATION CLASS.

It is anticipated that some form of accounting for use of resources will be required in many circumstances within the NRI. OMB regulations requires this accounting at the agency level, and thus it might be sufficient to rely on the agency affiliation data to satisfy this requirement. In other cases, an orthogonal account identifier might be required and so we allow for inclusion of a BILLING CODE⁷ as part of the explicit access control data. This may prove especially important in contexts where commercial facilities are employed.

In the most extreme cases it may be necessary for an individual subject to be identified, either for accounting or for access authorization. Although details for such an identifier were not discussed, it seems likely that a hierarchic data item would be appropriate, with a domain identifier used to specify the authority that vouches for the subject’s identity, plus a subject identifier that is unique within the domain. Even if users need not be identified as individuals, groups of users may be identified for

⁷Note that this item may enter into the decision process or may be employed only for accounting.

authorization purposes. Hence we expect to see a SUBJECT ID compound data item consisting of a DOMAIN ID and a USER ID, where this later data item may represent a group of users rather than a single individual.

The (ultimate) internet layer (IP or CLNP) source and destination addresses associated with a packet, possibly including protocol identification data, are also viewed as legitimate inputs to access control decisions, but for different reasons than the other data items described above. Use of addresses provides a convenient means of prohibiting access by specific devices or groups of devices (e.g., entire LANs) should it become necessary to revoke access at this granularity. Also, one can imagine simple access control policies that might be employed initially in the NRI and which would be based only (or primarily) on these values. Finally, we note that these data items are already included in every packet and are examined in the course of effecting the routing decisions which are the heart of the internet switching system and which are thus intimately related to the objects being protected. Thus even if these data items are not used in formulating an access control decision, they play an important role in the enforcement of the policies. It is worth noting that the preceding discussion of data items which are candidates as explicit inputs to access control decisions does not address how or when these data items are created, distributed, validated, or transported in subscriber traffic. These are important architectural issues, some of which are addressed in later portions of this document.

4.3. Communication Scenarios

4.3.1. Connection-Oriented Communication

Different types of communication scenarios may impose differing requirements on access control mechanisms. We observe that fine-grained access control mechanisms for connection-oriented communications are better understood and easier to implement than corresponding mechanisms for connectionless communication. The rationale behind this observation is that connection-oriented communication implies some connection establishment procedure. This procedure is a natural place to perform access control checks and to terminate the procedure if the checks fail. Moreover, the processing and bandwidth overhead associated with connection establishment procedures makes the added burden of transporting and processing access control information less onerous. In contrast, additional processing and bandwidth for access control applied to individual packets is much more likely to result in an unacceptable overhead if comparable levels of assurance and granularity of enforcement are sought.

The NRI is expected to provide (lower layer 3) connectionless service as its basic interface. Many proposed designs for IP or CLNP switches for this network environment introduce a notion of "soft-state" for connectionless traffic which is roughly analogous to treating this traffic as though it were connection-oriented. This soft state is usually cited as a prerequisite for providing better congestion control facilities in the Internet and for supporting more sophisticated routing, e.g., type of

service (TOS) routing with support for bandwidth guarantees.

We anticipate that designated IP/CLNP switches in the NRI will act as enforcement mechanisms for the transmission and switching access control policy, an assumption that matches Clark's policy routing model. The switches, designated "policy gateways" in Clark's paper, are ideal candidates for this role as they provide the interfaces between domains and thus have direct control over packet transport at domain boundaries. Based on these observations, it seems reasonable to pursue access control mechanisms which assume that some form of connection abstraction can be imposed on most (though perhaps not all) communications. The intent is that the soft-state database could be augmented to include additional data required for access control enforcement.

Throughout this report we shall employ the term "connection" in this broad sense when discussing path establishment procedures, even if the internet and transport layer protocols employed by the end points do not provide a true connection service. Only when the characteristics of a communication activity cannot be effectively modelled as a connection in this soft state sense (as would be the case in many brief, transaction-oriented communication scenarios) will we use the term "connectionless" to describe the activity.

This orientation is further motivated by the relative ease with which one can devise mechanisms for communication scenarios in which there is a well defined "initiator" of a "connection" and this initiator can be called upon to supply inputs to the access control process. For example, traditional virtual terminal communication involves establishing an actual connection, in real time, between two processes. The initiator of the connection is required to supply authorization data to the target of the connection before access is granted to the computation resources at the target (though this occurs after the connection itself is established). The same holds true for traditional file transfer scenarios, even though 3-way file transfer facilities have been defined which may not precisely fit this model.

4.3.2. Variations on Connection-Oriented Scenarios

When the scenario does not embody the concept of an initiator, then it may become more difficult to devise simple mechanisms for acquiring the authorization data prior to authorizing transmission of data on the connection in question. The example of simultaneous connection initiation by two TCP instances was cited as an example of this sort of deviation from our simple connection establishment scenario. The concern here is not an access control issue per se, but rather that two simplex connections would be separately routed instead of one duplex connection, a situation which could lead to anomalous behavior (in terms of performance). Note also that ISO transport protocols (TP0-4) do not support such simultaneous connection initiation and so the criticality of supporting such "dual initiator" situations is not clear.

Another concern was voiced over situations in which the initiator of a connection is readily identified but permission to traverse a path is a function of the authorization

of the computing resources being accessed, not of the subscriber initiating the connection. The assumption underlying this concern is that the initiator of the connection would not be capable of supplying the necessary, validated authorization data to the satisfaction of the policy gateways because such inputs would be available only at the destination. However, if the host being accessed could distribute appropriate credentials to the user prior to his access, the simple initiator scenario might suffice.

These two examples indicate how discussion of access control in the context of specific communication scenarios can be highly dependent on underlying assumptions about details of enforcement mechanisms. Many such discussions cannot take place without a straw man architecture for such mechanisms, and the straw man must address assurance issues, etc. Nonetheless, it is worthwhile to characterize the range of communication scenarios which need be supported in order to establish a reference for evaluating such straw men. Thus we will continue exploring communication scenarios and postpone enforcement mechanism discussion until the next section.

4.3.3. Electronic Messaging

Electronic mail poses something of a problem for connection-oriented access control models for several reasons. First, the initiator of a connection established for mail transfer is generally not the message originator and may not even have any relationship to the originator or a recipient. In fact, staged delivery of mail permits relay points which have no affiliation with the message originator or any recipient. This decoupling raises concerns with respect to assurance of access control inputs. Second, identifying a single subject for access control purposes becomes difficult in this context as multiple message originators may be served by a single mail transfer connection. Third, if traffic destinations are included in an access control decision, the multi-recipient characteristic of many messages further complicates the process.

We could accommodate mail transfer by treating mail transfer agents (MTAs) as subjects, and according to them a set of privileges appropriate to ensure mail delivery throughout the NRI, though that may not translate into allowing every MTA to access every other MTA directly or via any possible network path. This approach sacrifices fine granularity access control, and possibly efficiency of mail transfer, for simplicity. The fact that mail generally does not require the low delay paths⁸ (which we anticipate will be the most scarce resources) may make this approach more palatable. If commercial paths are employed and fine grained billing is required, this approach delegates responsibility for per-user billing to the message handling system (as envisioned in X.400 recommendations). This approach is analogous to the access control technique typically adopted for end-system access control with regard to mail.

⁸If electronic mail offered priority service categories which imposed stringent limits on delivery delays, then this general comment might not hold.

4.3.4. Transaction-Oriented Communication

Various brief, connectionless interactions will take place between servers. Interactions are so brief, and may be so dispersed over time that they do not fit the connection abstraction noted above. Nonetheless, some form of access control must be allied to all traffic if the access control facilities are to be effective (complete mediation). Such interactions may best be accommodated by not requiring any connection-like authorization procedure, but rather by requiring the access control enforcement points to recognize such interactions (perhaps based on source/destination addresses) and permit them on the basis of fairly static authorizations. This “special case” treatment for connectionless traffic is likely to be acceptable only if the resulting traffic volume is fairly low. Some form of auditing of these traffic flows would still be necessary⁹ to support the accounting requirements cited in section 1 and would provide a basis for detecting anomalous patterns that might be indicative of misuse.

File server interactions may not fit this profile, despite the fact that they are transaction-orientated communications. If the quantity of data returned in response to a small query is quite large, e.g., an entire file or directory, then the traffic volume would likely be too large to treat as above. Fortunately, most file server interactions would likely be local and thus not subject to the access controls we are discussing, i.e., the transfers would not cross domain boundaries. However, a homogeneous collection of file servers in different geographic locations might generate significant amounts of traffic in response to user commands. This poses the potential problem of large data transfers initiated from hosts which employ connectionless protocols and which operate on behalf of (non-resident) users. The first aspect of this problem could be addressed by requiring use of connection-oriented protocols for such transfers (a not unreasonable suggestion for other than local transfers anyway). The second aspect of the problem either requires enforcement mechanisms which support such “proxy” operations or adoption of policies which do not require fine grained access control (so that identification of the file server rather than the specific user is sufficient).

4.3.5. Multicast Communication

One other class of communication was very briefly discussed which was also not well represented by our simple connection-oriented model, i.e., multicast communication. At least some of the concerns about support for multicast seem to have arisen in conjunction with discussion of the need to factor in the authorization associated with the destination of a packet as well as its source. Again, the underlying assumption seems to be that the destination might be required to provide some authorization information data which only it would possess and acquiring this data would become even more complex in scenarios where the packet is addressed to multiple destinations.

⁹If the volume is sufficiently low, the traffic might be considered part of the “noise floor” for the NRI and not explicitly accounted for, as would be the case for routing updates, etc.

One can distinguish two classes of multicast communication: transaction-oriented and stream-oriented. The latter has been typical of conferencing communication while the former is typical of server location queries, etc. Transaction-oriented multicast communication might be accommodated by the static, address-based access control mechanisms discussed in section 4.3.4. Stream-oriented multicast typically involves some form of stream establishment procedure prior to transmission of user data and it does not seem unreasonable to augment such procedures to accommodate authorization data transfer. Thus multicast communication may not be so difficult to accommodate as originally suggested.

4.4. Access Control Architectures

Access control policies can be examined independent of enforcement mechanisms and architectural details, but there are limitations to such isolated examination, as noted in section 4.3. There are several reasons for adopting a (straw man) architecture in which to consider such policies. First, one must identify the transmission costs, e.g., in terms of processing overhead or bandwidth reduction, associated with enforcement mechanisms in support of policies. Second, one must understand how policies' representations and authorization data are managed in order to estimate the infrastructure costs (additional servers and databases, dissemination of authorization data, human management for the databases and equipment, etc.) associated with such policies. Third, one must understand where trust is vested in the architecture in order to gage its social acceptability and establish the level of assurance that might be accorded the resulting access control system.

In this section, we discuss how operating system security principles might be applied in this access control context.

4.4.1. Analogies with Operating System Security

In discussing mechanisms for network resource access control, it is useful to compare them to some of the enforcement precepts generally applied to operating system access control mechanisms. In the context of computer systems (subscriber resources), the concept of a "reference monitor" is widely used. A reference monitor mediates all accesses by subjects to objects. (For any reasonable degree of implementation assurance the reference monitor must itself be protected from tampering so that it cannot be circumvented.) Before any object is accessed, the authorization of the subject to access the object, and to operate on it in the fashion requested, is checked. This a priori checking is deemed essential if the reference monitor is to prevent the unauthorized release or modification of data. Despite the use of reference monitors, even in relatively high assurance operating system implementations, there are usually covert channels via which data can be released to

unauthorized subjects at relatively low data rates.¹⁰ Complete elimination of these covert channels is usually deemed impractical except in the most sensitive applications. Auditing of object accesses is often performed in addition to the access control enforcement described above and post access analysis may be carried out. However, this analysis is best viewed as a damage control measure and a possible means of detecting anomalous usage patterns, not a primary enforcement mechanism.

In the context of network resource access control, neither disclosure nor modification of subscriber data is at risk. (Recall that traffic analysis is not a service considered here, but rather is a subscriber security service considered by the End-to-End Working Group). Instead, the primary concern is transmission of packets via paths which are not unauthorized, i.e., unauthorized consumption of resources. A major failure of these controls could result in denial of service for authorized users, but minor failures result only in some small amount of "theft of service". The impression provided by the report of the Policy Working Group is that such minor violations would be acceptable in the context of most, though not all, of the articulated access control policies for switching and transmission resources.¹¹

This suggests that it is appropriate to adopt enforcement mechanisms which are resistant to attacks which would result in major violations of the access control policies, but that perfect control of traffic flows is not essential (analogous to information disclosure via covert channels in the operating system context). It also suggests that post access auditing is appropriate as a damage control measure and to verify that authorized subjects have not engaged in usage patterns which call into question their trustworthiness. Thus we suggest adopting a reference monitor-like approach for our access control policies, but with the understanding that perfect access mediation is probably infeasible and unnecessary.

4.4.2. Clark's Policy Routing Model and Access Control

We adopted as a strawman architecture the design presented by Dave Clark in his paper on policy routing.¹² Many of our discussions were influenced by the concepts and mechanisms proposed in the paper. In this section, we review those aspects of the design which are relevant to our access control concerns, discuss areas which were not completely specified in Clark's paper, and explore some modifications and extensions to this design.

Clark's paper defines three new entities in the Internet which participate in policy routing and thus network resource access control. Enforcement of policy route constraints is the responsibility of policy gateways. These gateways are present at the

¹⁰ Data rates on the order of 1-10 bits per second are typical for covert channels in this context.

¹¹It is clear that some access control policies would not be satisfied by inherent limitations of the type suggested here and thus would not be accommodated by the architectures proposed herein. For example, NASA is unlikely to trust such architectures to enforce a non-interference policy for network resources critical to shuttle operations during a mission.

interfaces between domains¹³ and thus are capable of controlling the flow of all traffic into or out of a domain. Within each domain are one or more policy servers.¹⁴ These devices serve several functions and are, in many respects, the heart of the access control system proposed by Clark. A policy server serves as the repository for and the management interface to inter-domain access control policies for its domain. Thus it provides representations of these policies to policy servers in other domains and it acquires from them policies applicable to their domains. A policy server responds to queries from subjects on hosts within its domain, synthesizing valid routes based on the subject's communication requirements, the PS's knowledge of current internet connectivity, and of applicable inter-domain access control policies. A policy server provides the selected policy route(s) to the subject, along with authorization and billing data, cryptographically sealed by the policy server. This operation is best viewed as a digital signature process.

A central feature of this proposal is that it requires the policy gateways to trust the policy servers that represent a domain, but does not require this trust to be extended to each subject within the domain. Clark assumes that domains are mutually trustworthy to the extent that the policy gateways rely on the source policy server to have correctly evaluated the subject's authorization to make use of a given policy route. Since domains in the NRI represent organizations (e.g., Federal agencies), there may be a reasonable basis for assuming that the individuals managing a policy server on behalf of a domain can be relied upon to operate in a responsible manner. (The trustworthiness of the hardware and software upon which a policy server is implemented is a separate concern.) Note that the means by which a policy server ensures that a validated route is properly bound to an authorized subject within the domain is a local matter, not specified by the architecture.

Signing of this collection of data serves several purposes. As noted above, the policy server for a domain is vouching for any identification and billing data and is also stating that it has selected a route which is allowed by the access control policies provided by other domains. Clark notes that this does not preclude checking of route validity by policy gateways, but it does allow mutually trusting domains to rely on these checks performed by the originating domain's policy server. It is advantageous that the signature be generated using asymmetric cryptography so that the policy gateways have a non-repudiable record of these claims by a policy server (which might prove useful should disputes arise or in isolating faults). Since only policy servers generate the signatures, the task of managing keys for signature validation becomes manageable.

¹²"Policy Routing in Internet Protocols," Version 1.1, May 19, 1988.

¹³Clark employed the term "Administrative Region" but we adopted the term "Administrative Domain" to avoid any implications of geographic locality.

¹⁴Clark designated these devices "Policy Controllers" but we have adopted our current designation to avoid confusion that might from use of the acronym "PC."

Clark proposed that an initial packet include an IP option consisting of signed policy route data (including billing and authorization information), but that subsequent packets contain only a short form of the policy route option with a "handle" from the option in the original packet. The handle would be generated by the policy server in the source domain and would uniquely identify the current route (based on the combination of the domain identifier and the route identifier). The policy gateways would cache the policy route using the handle as a search key and subsequent packets would be validated by determining if the handle was present in the cache and by processing the packets according to the policy route associated with the cache entry.

This approach to individual packet validation differs from others which have been proposed, e.g., Estrin's VISA schemes,¹⁵ in that it does not assume a crypto checksum binding authorization data to packet contents. Thus it is possible to copy a valid header from a legitimate packet and prepend it to a packet content not associated with the valid header. Clark argues that this is an acceptable vulnerability since the access control afforded here only applies to transmission and switching resource utilization, not information disclosure. The utility of "appropriating" valid packet headers is limited so long as the policy gateways match source and destination addresses against those held in the cache (as specified in the signed, policy route option). However, in circumstances where use of resources results in actual bills, unauthorized transmission of packets using copied, valid headers or forgery of valid headers could result in spurious charges to legitimate users.

In his paper, Clark proposes inclusion of a 16-bit signature and a handle composed of a 16-bit domain identifier, and a 16-bit route identifier unique within the domain in the policy route option. It was not clear if the short form of this option would also contain a signature, though most of the working group membership believed this might have been implied. We observe that a 16-bit signature is probably insufficient to preclude forgery; a more appropriate size quantity would be on the order of 128 or 256 bits. It is critical that the policy route option be unforgeable and thus the extra overhead implied by the larger signature is justified.

On individual packets traversing an established route there is a diminished need for short form option integrity and authenticity, except to prevent malicious, spurious charges. As noted above, if policy gateways check the source and destination address in the packet against that recorded in the cache, there is relatively little to be gained from forging a short form option. Since it is already possible to copy a legitimate short form option from a valid packet, it isn't clear how much additional assurance is provided by incorporating authenticity measures in short form options.¹⁶ Perhaps a prudent safeguard is for policy servers to adopt a process for selecting route identifiers

¹⁵"VISA Scheme for Inter-Organization Network Security," D. Estrin and G. Tsudik, Proceedings of the 1987 IEEE Symposium on Security and Privacy.

¹⁶We also note that the computational overhead of validating a crypto-seal (or reasonable size) on every packet is probably prohibitive.

so as to minimize the likelihood that they can be guessed, e.g., using a pseudorandom process. We do recommend that the policy route option be expanded to include some indication of lifetime, either measured in time or in number of packets or both. This limit on the lifetime of a route further reduces its vulnerability to exploitation by unauthorized subjects and a packet quota could provide an additional means for detecting misuse.¹⁷

4.4.3. Clark's Architecture in Retrospect

Now that we have reviewed the architecture presented in Clark's paper and made some local observations and suggestions, it is useful to view the architecture in the context of our previous discussions. For example, the architecture described in this paper supports both identity-based and rule based, administratively-directed access control policies. It adopts a security model in which the objects are routes through the Internet (which correspond to use of switching and transmission resources) and the subjects are processes executing on behalf of users or groups of users and, hosts or groups of hosts (perhaps entire domains).

Clark's architecture embodies the connection-oriented (single originator) access control model discussed in section 4.3.1 above and thus this class of communication is especially well served by this architecture. Communication scenarios that deviate from this model must be examined to determine how they can be accommodated. For example, electronic messaging would probably be handled by viewing the MTAs as subjects rather than trying to control access on the basis of individual message originators, as suggested in section 4.3.3. Stream-oriented multicast communication could be accommodated as described in section 4.3.5.

Transaction-oriented communication, whether point-to-point or multicast, may not be served very well by this architecture, i.e., it may be difficult to amortize the cost of policy route options in these communication scenarios. However, if cache entries in policy gateways can include "wild card" entries for addresses, then it might be possible for a policy server to seed routes for access to commonly accessed collections of servers, etc. on behalf of all (many?) of the hosts in its domain and pass out the identifiers for these routes to members of the domain.

The remaining deviant case involves dual-initiator connections, a scenario of undetermined criticality. The source and destination hosts could discover that different route identifiers were assigned to a single transport layer connection and co-operate to use only one of the routes (using some unambiguous criteria such as comparing route identifiers as unsigned integers and selecting the larger value route identifier). However, this solution may be viewed as being outside of the architecture in that it does not involve the policy gateways, policy servers, etc. Another aspect of support for some communication scenarios which generated some concern is also outside the

¹⁷If a packet quota were imposed on a route and the route were used by an unauthorized subject, the authorized subject might detect this if the route were to become invalid due to exhaustion of the packet quota.

scope of the architecture, i.e., the need for proxy authorization. The possible need for such a facility was noted in conjunction with file server communication on behalf of users, e.g., transfer of a file between two file servers. It appears that the architecture in Clark's paper could support such communication authorization, but the means by which the initiating policy server determines that the communication is on behalf of a specified user, rather than the file server itself, is a local matter not part of the architecture.

In section 4.3.2, a concern was raised about supporting route establishment when permission for a route was dependent on authorization of the destination, not the initiator. In Clark's architecture, this case would not be treated any differently since it is the initiator's policy server which evaluates the access control policy and makes the decision, and all the inputs required to make the decision are available to that policy server. For the most part, the architecture assumes the policy gateways trust the initiating policy server to interpret the access control policies correctly at the time it generates the sealed route option and supplies it to a subject in the local domain. Intermediate policy gateways can review the data provided in the policy route to confirm the decision, but the paper seems to suggest that this independent confirmation would not usually be carried out during route establishment, for reasons of efficiency, though the signature should be checked.

4.4.4. Trust Implications and Possible Remedies

In Clark's architecture, the ability of policy gateways to validate an access control decision is limited because the authorization data included in the signed route option does not incorporate any independent validation mechanisms. For example, the policy gateways must trust the initiating policy server to have verified the user ID, agency affiliation, etc., because there is no means for the policy gateways to verify these access control inputs directly. The route verification that can be performed by policy gateways is based on checking the signature (thus verifying the integrity and authenticity of the route) and on matching the supplied access control inputs against the policy in effect. Rather, the assumption is that access control policy terms and conditions are distributed and that the data items against which the policy terms and conditions can be matched are all locally validated quantities, i.e., they are vouched for solely by the initiating domain through its policy server. Thus the architecture relies on mutual trust among domains, non-repudiable (signed) policy routes, and post-hoc auditing to reconcile conformance.

If this level of mutual trust proves unacceptable in the NRI, it is worth exploring how one might extend the architecture to incorporate independently verifiable "credentials". First we need to identify which credentials might need to be independently verifiable. One candidate is the AGENCY AFFILIATION INDICATOR. If a connection is initiated with a policy route that claims an affiliation for which the initiating domain is not the certifying domain, then it might be reasonable to require that the AGENCY AFFILIATION INDICATOR be

independently verifiable.

A BILLING CODE might require independent verification if the code is one which does not somehow imply charges to the initiating domain.¹⁸ An analogy can be made with long distance telephone charging. A direct dialed call from a home number is assumed to be legitimate, whereas a similar call from a pay phone or hotel room requires an independently verifiable account number unless the charges are borne locally (via coins or billed to your room). Thus BILLING CODEs also appear to be good candidates for independent verification, at least in some circumstances.

Finally, the other major credential considered for inclusion in policy routes was the SUBJECT ID. Again, the circumstances in which independent verification is likely to be of interest are those in which the subject's domain differs from the initiating domain. Since the SUBJECT ID already includes an indication of the domain which vouches for the subject's identity, it is easy to determine if independent verification is required. Thus in all cases the motivation for an independent verification facility arises only when the certifying domain for a credential differs from the initiating domain for the connection.

In order for a domain to certify a credential for independent verification, the resulting data should be bound to a subject (or class of subjects) so as to render it useless to other subjects. This is easily accomplished by including the subjects (subject class) to whom the credential is issued as part of the signed credential. Note that this also allows the issuer to distribute the credentials directly to subjects, not only through domains, if that proves useful. Thus a domain such as DOE might issue a BILLING CODE and AGENCY AFFILIATION ID to a researcher at a university, binding it to his SUBJECT ID. The researcher could present the credentials to his local policy server for consideration in selecting routes and that policy server could include the credential along with the policy route option.

Policy gateways could verify that DOE had granted permission to use the BILLING CODE to this subject and that the subject was affiliated with DoE by verifying the seal on the credential and matching the included SUBJECT ID against that in the policy route. As above, it might not be feasible for every policy gateway to perform this independent verification prior to processing packets for the connection, but the option would exist and post hoc auditing is feasible. These credentials should contain a validity date range to constrain their lifetime, and some form of hot list would also need to be maintained by each issuing domain and distributed to policy servers and gateways to revoke credentials, e.g., upon termination of affiliation.

This technique would reduce the level of trust accorded the policy server at the university since it could not forge the credential. This binding does not ensure that the subject and the source address are correctly paired. However, if the SUBJECT ID

¹⁸Clark suggested that such codes might incorporate an AD identifier which would explicitly establish the requisite binding. However, he was concerned that a strict requirement for a billing code to be bound to the initiating AD would unduly restrict mobile users.

indicates that the initiating domain is the certifying domain for the subject, then one must ultimately rely on that domain to correctly maintain subject-address bindings. If the subject is foreign to the initiating domain (as might be the case for a mobile user), the incremental assurance offered by independently verifiable credentials seems fairly small. It is not clear what form of credential binding would be useful for mobile users. The "home domain" for a mobile user could certify that he was temporarily associated with another (specified) domain, thus lending credence to a claim by the initiating domain that the "foreign" user was in residence. If the logistics of generating and transferring some sort of travel credential ("hall pass"?) could be made acceptable to users, this might prove to be a viable means of addressing this problem. For these credentials, even more than most, validity dates should be included to limit their lifetime.

5. Resource Sharing

Working Group 2 Members

David Clark (Chair)	MIT
Guy Almes	Rice
Bob Braden	USC-ISI
Scott Brim	Cornell
Jon Crowcroft	University College London
Deborah Estrin	USC
Steve Goldstein	Mitre
Phill Gross	NRI
Bill Jones	NASA/Ames
Dan Nessel	NMFECC
Ari Ollikainen	RIACS
Mike St. Johns	DCA
Tony Villasenor	NASA HQ

5.1. Introduction

This working group was asked to consider the question of mechanism necessary to insure “fair” sharing of resources, in particular bandwidth.

The group proposed, as a starting position, that to permit sharing of resources, such as networks or links, among agencies (for example), the following questions must be answered.

- What sorts of service classes will be required? Which are possible?
- How must the users of the resources be categorized?
- What sort of accounting for the resources are required?
- What levels of assurance are required?
- How global is the impact of various sorts of service classes?
- What management tools are required to control multi-agency policy mechanisms?

Two ideas are central to the discussion: service class and category.

5.2. Service Class

The idea of service class is that in order to provide a controlled sharing of a resource, it is necessary to define how the sharing will be measured. The measurement represents a way of specifying a service class.

In the workshop, most service classes related to policy concerns were defined in terms of relative bandwidth. The following examples were often proposed:

- A link is shared by two (or more) service classes, each of which gets a guaranteed fraction of the link capacity under overload.
- A link is shared by two (or more) service classes, some of which may not interfere with others. That is, they are excluded from the resource if demand is excessive.

An example of a service policy requirement not directly related to bandwidth is mutual aid: two agencies that agree to carry the other's traffic if the resources of the one is down. Half of the mechanism necessary to support this is easy: one could define a service class for traffic belonging to the other agency, and define the service constraint for that class. The hard part of the mechanism is to define how the switch is to know that the other resource is down, so that the usage by that class should be permitted.

In the discussion of service classes, the following comments arose:

- Outside the arena of policy control, there are much broader requirements for service classes, in order to support new sorts of applications. For example, some applications require control of delay. This broader problem is usually called the "Type of Service" or TOS problem (also called quality of service or QOS in ISO). In this respect, the mechanism required of the switch for specifying and measuring the services classes is just a subset of that required for support of multiple classes of service to support applications.
- Some (non-policy) examples of service classes are very difficult to support, e.g., those for real-time speech, or variable rate encoders (that can adjust to changing bandwidth allocation, but must KNOW what rate they are being offered).
- We believe it is not difficult to provide commitment of resources to simple service classes. For example, a gateway could be constructed that would take packets in two service classes, and ensure that under overload each class received equal access to a link. The problems in doing this are to control the overhead in the gateway, which would have an impact on high-speed networks, and to understand the global impact of such guarantees (see below).
- The definition of service classes must be understood globally.

5.3. User Categories

In order to ensure that some user receives some service, it is necessary to identify the packets associated with that user. This is a very hard problem, perhaps harder than supporting reasonable service classes.

Current IP packets do not have user names in them, just source and destination Internet addresses. But a single machine might support users with different privileges, or a user wanting to use different privileges at different times.

In the discussion of user categories, the following points came up:

- To support the sorts of requirements that were offered as examples (e.g., put all NASA packets in service class X), it will be necessary to have some explicit tag in the packet to indicate the packet category. This is a new IP level mechanism.
- The level of "user granularity" is not clear. Would one tag for all of NASA be sufficient, for example?
- It might be necessary for a packet to carry more than one tag, to permit a user with multiple privileges to use them at the same time. Perhaps tags could be approximate, and could resolve in different manners in different parts of the net.
- The level of trust needed for the tag is unclear.
- If a tag is abused, the use must be traced back to an accountable entity, which ought to be a human.
- A very hard problem is multicast: one packet going down several paths that might require different user privileges.

5.4. Additional Discussion

The following comments were made about the other points in the list above.

5.4.1. Accounting for usage:

A clear requirement was that the usage of resources by different user categories be accounted. However, the details of the requirement were not clear. It does not seem too hard to provide a simple measure of total bytes or packets used by each class. As noted above, the hard part is defining the classes, and inserting the class information into the packet.

If a more dynamic accounting for usage is required, then a mechanism can probably be defined to account for usage by any pre-defined measure, but arbitrary measures will be real hard.

5.4.2. Levels of assurance:

There seem to be two obvious levels of assurance as to enforcement of service classes and user categories.

- Separation of traffic into classes, and enforcing and accounting for the usage of each class, will be performed properly so long as the switch elements belonging to each agency operate properly.
- Proper separation and accounting must occur even if the switches of one agency are mis-programmed or malicious.

The latter would be required (probably) in a network operating in hostile circumstances; it corresponds to mechanisms to prevent denial of service. It is a level of assurance that is hard to achieve.

The former level of assurance is much easier. It corresponds roughly to the operation of the Internet today. If one set of gateways is not operating properly, there

may be bad global effects that the other gateways cannot prevent. The problem is cured, not by robust dynamic algorithms, but by detection and correction (e.g., by humans) of the problem.

For many circumstances, e.g., conformance to OMB regulations, the weaker form of assurance is probably sufficient. But DARPA, for example, expressed an interest in as robust an assurance as possible.

5.4.3. Global effects:

The problem of global effects of policy is a very serious issue, the impact of which does not appear to be sufficiently appreciated.

Certain resource constraints, most obviously non-interference (a service class that is excluded when a resource is overloaded), cannot be implemented except in the context of a global routing algorithm that knows about the constraint.

The problem is the following. At the moment, the Internet supports the idea that for any destination address, there is one route out of a switch. If we now support two service classes going to that destination, then each will be sent by the same route, given the current routing algorithm. If one of these service classes is now blocked from a congested resource, there is no mechanism to reroute that class to another resource. The result is that the service class is totally disabled.

In other words, today if a gateway makes a local decision to discriminate against certain users, those users perceive a global disruption of their service.

The problem of propagating and responding to local controls is not impossible. While this section stresses the need to understand the problem, we believe that solutions exist. It will be necessary, however, to contemplate a major adjustment to the current philosophy of Internet routing. In particular, most of the promising approaches are based on some form of source routing.

Above it was asserted that it was not difficult to build a gateway that would make simple resource guarantees. The difficulty is propagating the knowledge of that local guarantee. There are some guarantees that could be enforced in today's Internet without the necessity of global knowledge. For example, if a gateway provided equal sharing of a link under overload to each of two classes, then the global impact would be that of a link whose capacity changed by 50%. A fluctuation of this magnitude could not be globally distinguished from other current forms of congestion. So there are some local controls that can be applied safely in today's Internet, and others (such as non-interference) that can only be contemplated in the context of a global architecture.

5.5. Conclusions

The problem of making a local modification to a gateway to enforce a bandwidth usage limit to a identified category of users seemed reasonable.

Associating a user category with a packet is very hard. The actual requirements are not clear (are one or several categories required, what is the level of assurance that the specified category is legitimate, and so on). In addition, the mechanism is not obvious. This matter is addressed in the report of working group 1.

The problem of level of assurance is also very hard, again because the actual requirement is not clear.

Accounting for usage is probably not too hard.

The hardest problem is redefining the routing algorithms of the Internet to correctly propagate and respond to the impact of local policy controls.

There are several hard and interesting research questions:

- How do service guarantees compose?
- Is it possible to build multi-region systems that are resistant to attack by malicious third-party regions?
- How could user categories be managed? Are they multi-valued, hierarchical or flat?
- How can fault isolation and service assurance be performed?
- What is the relation between statistical resource allocation and possible guarantees of access?

To avoid solving too general a problem, several questions should be asked of the agencies.

- What level of assurance is required?
- What sort of user categories will be required?

5.6. Recommendations

The group proposed a number of experiments and changes that could be undertaken at once, to better understand the problems of policy routing and resource control, and to provide operational facilities toward these goals.

These goals are organized in three categories, things that could be done at once using existing tools, projects with a short time frame, to provide better capabilities and understanding quickly, and finally, projects that would require longer to complete.

5.6.1. Instant projects

Statspy

Although source and destination addresses are not a precise indicator of service class, they do provide much useful information. The so-called *statspy* tool has been used in the past to collect a matrix of traffic sorted by source/destination address. This information could be collected for shared links today to provide a first cut at accounting for the resource.

Route filtering

Route filtering provides a way to instruct a gateway to believe only part of an incoming routing packet, or to change parts of that incoming data, e.g., the cost metric of a proposed path. This capability, available in most commercial gateways and in the gated software for Unix, provides a way to control which destinations are reached by which paths. It cannot separate service classes, but can be used for very rough divisions of traffic based on destination address.

5.6.2. Short-term experiments

These are experiments that could be undertaken at once, with the expectation that they would yield results in the short term. They are not thought to contain high-risk research questions. They might provide some increase in operational capabilities in one to two years.

Simple resource guarantee

A gateway could be programmed to sort incoming packets into two service classes (based on some simple if unrealistic characteristic of the packet, such as addresses or TOS flags), and then divide the use of a link fairly between these classes. That is, in underloaded conditions, each could operate without constraint, but in overload each class would have a fair share of the link.

This would be a first demonstration of allocation of resources to service classes, and would provide a practical way to share a link.

Observe tagged packets

Above, it was noted that the *statspy* program could be used to count packets based on source and destination addresses. One could define a simple IP option, which carried a user identification, and then use the same *statspy* to count these packets. A simple use of this option would be to tag the packet with an indicator of which agency had “sponsored” the packet.

Putting a new IP option into a packet is not hard; some systems like Unix 4.3 BSD provide the hooks to do this today. A simple and general way to find the proper value of the option field would be to implement a very simple form of “Policy Server”, which could be a user process on a Unix system. One would send a packet

to the server with the source and destination addresses, the name of the sponsoring agency, and other credentials. In return, one would get the suitable IP option, which would just be inserted into the packet.

This would provide a more accurate accounting of shared resources, and a first demonstration of the concept of the policy server.

Fast encryption of the policy information

In order to ensure that policy routes, authentications and so on are not forged, it will be necessary to seal them in some way. The obvious technology is encryption. A demonstration is needed of a sealing technique that runs at tolerable speeds. This would permit the introduction of a high level of trust into the accounting.

Demonstration of “soft state” in gateway

Several propositions for management of resources in gateways require that the gateway remember some aspect of the packet sequences passing through it. The idea of “soft state” has been proposed to capture the idea of cached information in the gateway which can be reconstituted if lost without terminating the higher level connection.

A first project is to program a gateway to show that this sort of state can be managed effectively, with acceptable overhead. The information stored in the state could initially be rather simple, for example the resource guarantees mentioned above, or logging of packet tags, or enforcement of source/destination access control.

Demonstration of policy routing with Loose Source Route

Once we have demonstrated the tagging of packets, we have all the pieces of a first demonstration of policy routing. A Policy Server module can be programmed to take the source/destination addresses, sponsor and so on, and receive in return a Loose Source Route IP option. This could be placed in the outgoing packet to achieve controlled routing of the packet.

5.6.3. Longer-term experiments

The following are experiments that have a longer term focus. They deal with harder problems, will take longer, and yield an increased functionality. They represent steps that can be undertaken now, and should be if increased functionality is to be achieved in the next few years.

Define and support Policy Source Route option

Above we described a simple demonstration based on the IP Loose Source Route. While this represents a useful first demonstration, the LSR is not suited for real policy routing, because it binds the route to specific gateways, which is too concrete, and because it has no fields to carry policy information.

What is needed is a new IP option to define a Policy Source Route, a more abstract form of source route containing policy information. There is general agreement on the need for this class of mechanism and the general form it would take. A detailed design is now needed.

Tools for Synthesis of PSR

The Policy source route described above would be generated using information exchanged by the various Policy Servers and Policy Gateways. Algorithms for this have been proposed; a concrete design should now be undertaken.

Define protocols for control interaction

To provide the information for the routing algorithm, it will be necessary for policy gateways, policy servers and hosts to exchange information. Protocols for these exchanges must be designed.

Management Tools for Policy Controls

Current experience teaches us that we must develop suitable management tools for a mechanism at the time that we develop the mechanism itself. The problems of policy control are complex, and can be expected to lead to complex management problems. We must begin the design of a management architecture for policy mechanisms.

Analysis of composability of local policies

We assume that an administrator of a region will express policies reflecting the local concerns of that region. These various local policies must be composed to provide an end to end service. It is necessary to ensure that the various local policies do indeed combine to permit a reasonable global service. It would be nice to have some formal understanding of what sorts of local policies can be composed, and some tools for checking that the actual proposed local policies are reasonable.

Architecture for signatures and sealing

To ensure the needed level of assurance, an overall strategy must be devised to define the trust that holds between the different components of the system, and the mechanism needed to insure the integrity of Policy Routes and related messages.

6. End-to-End Security Services

Working Group 3 Members

Dennis Branstad (Chair)	NIST
Matt Bishop	Dartmouth
Brian Boesch	DARPA
Anita Holmgren	Unisys
Barry Howard	Livermore
James Morrill	Sparta
Dan Nessett	NMFECC
David Peters	NASA
Steve Wolff	NSF

6.1. Introduction

This section deals with end-to-end security services for the National Research Internet (NRI). As described previously, the NRI consists of multiple, autonomous, mutually-suspicious, administrative domains. The NRI is an open environment with a dynamic security perimeter. Each domain may have its own security policy and offers a unique set of security services to its own community. However, if secure interoperation is desired across domains, these security policies must belong to a set of hierarchical, consistent policies, and certain cross-domain agreements with respect to security are needed. Working Group 3 focused on the nature and content of such inter-domain cross-agreements.

A security architecture for the federally-funded research networks (which make up the NRI) was proposed. The architecture consists of security services, where they are needed, example mechanisms, and the implied common technologies and common policies necessary to support interoperation.

First we offer the strawman architecture. Next, we introduce the concept of a “security domain”; we discuss multi-administrative higher-level security services in detail; then, using the workshop model (of phase 0-3 technologies), suggest a phased approach to making the architecture a reality.

6.2. Multi-administrative Security Architecture

We define security to include, not only protection from unwanted disclosure, but also, protection from unwanted modification and prevention of denial-of-service. This working group suggests that a small number of security services are necessary, and that these security services need to be repeated at various layers in the protocol and system architecture. The following chart illustrates some candidate security services such as: confidentiality, integrity, authentication, access control and service assurance; suggests placement in the architecture such as: user-level, host-level, gateway; and suggests common technologies and common policies that are needed to support these security services across domains.

Security Services in a Multi-Administrative Domain Environment			
<i>Security Services</i>	<i>Example Mechanisms</i>	<i>Common Technologies</i>	<i>Common Policies</i>
Origin Authentication			
-user/process	secure-ID card	Key Distribution	global ID
-host	certificates	(common protocols	conventions
-gateway	certificates	and standards)	
-realtime/deferred	challenge/response	Directory Services	
- certificates	(object registration)		
Origin Access Control			
-user	login	can we use policy	global ID
-host	visa	servers?	conventions
-gateway	policy routing		
Object Integrity			
-msg	MACs		
-file	MACs	common format for	global ID
-datagram	MACs	integrity labels	conventions
-connection	MACs		
-field	MACs		
Object Confidentiality			
	protected wire	Encryption-	Key Distribution
		(common protocols	agreement
		and standards)	
Service Assurance			
	routing	Byzantine Robust	Multi-domain Network
		Management	agreement

The International Organization of Standards has recently adopted an International Standard Security Architecture (IS 7498/2) that specifies five security services in the Open Systems Interconnection model of computer networks. The five services and a short definition of each are:

- Authentication: verifying the identity of communicating entities (e.g., computer, software programs) in a network;
- Access Control: restricting access to the information and processing capabilities of a network to authorized entities;
- Confidentiality: preventing the unauthorized disclosure of information;
- Integrity: detecting the unauthorized modification of information;
- Non-repudiation: preventing the denial of transmitting or receiving certain information.

A security label is security relevant information that is attached to other information to assist in providing the above named security services. The U.S. Department of Defense (DOD) has specified the format of a security label to be used at the Internet Protocol (IP) layer of the DOD suite of protocols. This label is used primarily to state the classification of the information in an IP packet. The security mechanisms then use the label to control the routing of the packet through the network (based on the security of alternate routes) and the confidentiality protection to be provided to the packet.

6.2.1. Security Domains

Security needs to be considered from an end-to-end perspective. Secure interactions across administrative domains, a security perimeter must be defined. A hierarchical set of "security domains" could be established for the research internet. A global security domain could then have a security policy and a set of security services that would be enforced and supported throughout the internet. Each sub-security domain could then have additional security services. Security interfaces between security domains would then be defined. Rules for data to cross these interfaces would need to be established and enforced by "interdomain gateways".

6.3. Higher-Level End-to-End Services

In this section, we discuss services in terms of "administrative domains", which are collections of machines and supporting hardware (nets, etc.) controlled by a set of people who have the (recognized or assumed) power to choose what services that set of entities will offer to other entities. We assume that entities in different administrative domains are mutually suspicious but wish to provide some set of services to each other. Note that the managers of each domain will define their own policies towards the provision of services, so the entities must interact in light of the relevant policies. These policies must be consistent; however, this is not a great restriction, since the policies will either be imposed by an authority encompassing both administrative domains or (more likely) by bi- or multi- lateral agreements or adherence to a mutually agreed upon standard.

We describe a set of supportive services designed to provide the basis for other, productive services visible to the users; we also suggest some useful productive services. The distinction between the two is crucial; supportive services, invisible to the user, are essentially a set of library routines designed to provide security and integrity functions in a manner dictated by the administrative domain. Two domains must decree some format for the interchange of information such as user IDs or file checksums, but (for example) the NASA administrative domain may require use of the File Transfer Protocol (FTP) be allowed only to authenticated individual users, whereas the Dartmouth administrative domain may allow any user from an authorized host to access files using FTP. In this case, the supportive services (authentication of the source of the FTP request) for NASA must support per-user authentication, whereas Dartmouth need only support per-host authentication; however, if NASA

allows FTP access by users in the Dartmouth administrative domain, some accommodation must be made by policy (either by NASA, to accept per-host authorization when users from entities at Dartmouth FTP, or by Dartmouth, to enable per-user authentication when dealing with FTP requests to entities in the NASA administrative domain). Productive services simply request of the supportive services whether some condition is met (is the user allowed to use the service, has the file been altered in transit, etc.) and proceed on that basis.

We describe the supportive and productive services separately.

6.3.1. Supportive Services

Supportive services supply the basis for an entity in one administrative domain accessing the services supplied by another entity in another administrative domain. To this end, they provide access control, authentication, integrity, and confidentiality checking.

The first class of supportive services is origin authentication. There are several subclasses. A policy may require per-process (i.e., per-user) authentication, using mechanisms such as SecureID(tm) cards; this will require some common technology for key distribution among the co-operating domains. A policy may require authentication at the host or gateway level, using certificates; here, a set of directory services such as an object registry must be common to co-operating domains. Note that there are really two flavors of authentication here, real-time authentication in which the origin must identify itself immediately (possibly using a challenge/response protocol), and deferred authentication, in which the origin need only identify itself at some time, the identification being preserved using certificates. Finally, regardless of the type of origin authentication done, all administrative domains must have some global object identification convention that all domains respect.

The second class of supportive services provides access control based on origin. For example, access to a user account might depend on the identity of the requester; on 4.2BSD UNIX systems, access is controlled by the .rhosts file in the target account, with each line of that file specifying a user/host pair authorized to access the account. The system assumes authentication has already been done, and controls access strictly based on the user/host names of the requestor. Similarly, if one host needed to access services on another, it might present a VISA or a service-specific certificate entitling it to use that service. A policy might allow or deny access to networks based on the source or destination of a packet (policy routing). In any case, as with the first class, this class of supportive services requires a global object identification convention. The technology which must be shared by administrative domains co-operating to provide these services is not clear; perhaps policy servers would suffice.

The third class of supportive services provides object integrity. A policy might require that the integrity of any (or all) of messages, files, datagrams, fields, etc., be verifiable, possibly using MACs or other integrity checking mechanisms. In this case, administrative domains enforcing this policy must agree on a common format for

integrity labels as well as a common set of mechanisms.

The fourth class of supportive services provides object confidentiality, for example by encrypting files or protecting the network wires. If cryptography is used, some key distribution mechanism must be agreed upon in order that keys for objects in one administrative domain be available to authorized clients in another. The administrative domains must also agree on the encryption algorithms to be used and some common technology for making keys available is necessary.

The fifth class, non-repudiation, will simply ensure that a requestor (or user) of a service cannot deny that that user made the request (use) of the service. Again, the administrative domains must agree on what types of requests are to be subject to this service, and on the mechanism to be used for inter-domain non-repudiations. Further, the granularity of the non-repudiation records must be decided; this impinges on accounting. For example, NASA may bill on a per-project basis, so if a request came from Dartmouth and the non-repudiation mechanism ensured non-repudiation only in that the request came from Dartmouth, the mechanism would be insufficient for NASA's purpose; again, this must be settled by inter-domain multi-lateral agreement or decree from a higher authority.

In terms of the four phases used to characterize the evolution of capability, at phase 0 is process (user) authentication with passwords; at phase 1 is process (user) authentication using other technologies such as challenge/response protocols; at phase 2 are authentication using certificates, integrity checking mechanisms such as MACs, integrity labeling, methods for non-repudiation, and issues of key distribution and management. Phase 3 issues include the use of VISAs for policy routine and certification across peer administrative domains.

6.3.2. Productive Services

Differing administrative domains provide varied services, but most will want to allow entities at other administrative domains to use one or more of the following services on one or more entities in the local domain. This list is by no means exhaustive; we have simply discussed the more common currently-provided productive services. Undoubtedly, equally or more important ones will arise in the future, or inter-domain policies and agreements will require new ones.

Remote job execution will be essential within domains and given the advances in the use of collaborative support services and distributed computations, important in inter-domain support. Currently, mail transfer by far dominates this area, with file transfers coming a close second. Both raise issues of inter-domain use of remote resources such as disk space and CPU time, as well as confidentiality and integrity issues (can only those authorized to read the file/mail do so? can the file/mail be altered?) Further, authentication of the sender/author (was the letter telling me I got my raise a forgery?) and access control will also be essential. Some of these issues are being addressed by Steve Kent's privacy task force (see RFC 1113), which has been examining secure and private electronic mail for some time. Finally, non-

repudiation of mail is important when electronic mail is used to make agreements or convey sensitive information that the sender may wish to deny having sent. Extensions to more sophisticated forms of collaborative support, such as multi-media mail or electronic "whiteboards", will require the same level of supportive services. (Note that the "support" service is a production, rather than a "supportive" service. This terminology is confusing, to say the least, but it is also standard.)

Remote access of computers (e.g., via Telnet) and distributed computations, the other forms of remote job execution, will all require similar supportive services -- that is, authentication, access control, integrity, and confidentiality. In all remote job execution schemes, if the execution is done inter-domain, the administrative domains must use a mutually agreed upon set of control protocols; this may be established either by multi-lateral agreements or by some superior authority (for example, an act of Congress dictating a protocol to administratively-independent agencies.)

Remote access comes in many forms; some computers will simply supply services such as directory services and not allow other forms of remote access. These services will require the usual supportive services, but will also require that the client be able to authenticate the server so the client can be sure it is connected to the intended directory and the server can be sure the client is authorized to access the information. Note that this need not be necessary for non-directory services, since if access is made through a directory server and a session key is obtained, should the client then access a bogus (non-directory) server using the session key the bogus server will not be able to respond. Similarly, user authentication as a productive service will be essential when dealing with certificates designed to be used in a productive service. For example, the use of laptop computers will require the availability of user authentication at this level.

Another resource requiring distributed use of computers would be a "national" file system, allowing remote hosts throughout the country to access a shared set of files; it will require not only mechanisms for the usual supportive services but also a common interface protocol and a common file exchange protocol to allow systems with very different file accessing semantics to use the national file system.

Due to Office of Management and Budget (OMB) constraints at the federal level, and bookkeeping concerns in other agencies, businesses, and institutions, accounting for resources used in and by other administrative domains will be required; since (for example) the Dartmouth administrative domain will not trust the NASA administrative domain to account for the use of electronic mail sent from Dartmouth to NASA, both NASA and Dartmouth would undoubtedly track such mail and check the relevant bills. Non-repudiation of use of service is at this point essential.

Key distribution in support of secure mail, authentication mechanisms, and other services will require protocols and standards agreed to by different administrative domains. Such services may be integrated with directory servers but this is a matter of policy.

Finally, as different administrative domains communicate, network management and control information will have to be passed between administrative domains, raising issues of object integrity, confidentiality, and access control.

In terms of the four phases used to characterize the evolution of capability, at phase 0 is mail relaying, transfer, and name domains. Phase 1 technologies are authentication technologies such as secure-ID, challenge/ response protocols, and authentication servers such as Kerberos. On the border between phases 1 and 2 are the distributed white pages for the entire Internet. Phase 2 mechanisms such as secure mail and key distribution and management mechanisms are currently under development by the IAB Task Force on Privacy; other phase 2 items are certificates, and security of distributed directory servers (white pages). Distributed computation protocols and controls for a national file system, and accounting mechanisms are phase 3. Also phase 3 are "firewalls" for end-to-end services, so that if the services fail over a portion of the Internet the rest of the Internet may continue to rely on the service being correct and functional (this would limit the damage of incidents like the Internet worm of November 1988) and also the integrity of data across international borders, since most nations restrict the transborder use of cryptographic algorithms that can be used for secrecy, which is true of the base algorithms used in the computation of cryptographic checksums for integrity. Hence a solution requires the development of a cryptographic algorithm that can be used for integrity and authenticity, but not secrecy. One possibility is to use zero-sum knowledge mechanisms to have a third party assure integrity without secrecy, might be feasible. Such a solution is Phase 4 (very long range research).

6.4. Projects

The above suggests several projects that the FRICC or some constituent agency should pursue:

- End-to-end private mail is currently in the experimental phase; encryption is done using the DES, and authentication involves certificates built using RSA. The mechanism allows both privacy and integrity of sent mail.
- A national file system will raise issues of access control, authentication, confidentiality, and integrity.
- Directory services should provide white pages for mail and multi- domain object registration; issues to be addressed include registration of services, distributed list service, and authenticity.
- Finally, questions of multi-domain network monitoring and control are at the heart of interconnected network operations and raise issues of access control, authentication, and integrity.

Some common or interoperable approach to authentication, integrity, and access control, as well as the tools and services to be provided, is necessary; note the policies may differ across administrative domains, but the mechanisms must be able to

communicate with one another. They need not rely on each other, however; that is a policy issue. Whether or not these inter-domain mechanisms can be built with common facilities, the specific protocol base (such as OSI or TCP/IP) that these projects are to be conducted, how results are to be transferred into GOSIP and a European context, the role of vendors as opposed to researchers, and the IETF, IAB, and other such organizations, and which agency or agencies shall take the lead, are all issues that can be resolved in the longer range.

Notes: Reference for the use of productive and supportive services is the ECMA (European Computer Manufacturers Association) Security in Open Systems, A Security Framework document, ECMA TR/46, July 1988.

7. Workshop Attendees

Guy Almes	Rice
Matt Bishop	Dartmouth
Brian Boesch	DARPA
Bill Bostwick	Los Alamos
Dennis Branstad	NIST
Hans-Werner Braun	Merit
Scott Brim	Cornell
Ross Callon	DEC
Vint Cerf	NRI
David Clark	MIT
Mike Corrigan	DoD
Jon Crowcroft	UCL
Richard desJardins	CTA
Deborah Estrin	USC
Steve Goldstein	Mitre
Phill Gross	NRI
Tony Hain	Livermore
Jim Hart	NASA
Jack Haverty	BBN
Dan Hitchcock	DoE
Anita Holmgren	Unisys
Barry Howard	Livermore
Bill Jones	NASA
Steve Kent	BBN
Larry Landweber	Wisconsin
Jim Leighton	Livermore
Barry Leiner	RIACS
Dan Lynch	ACE
Sandy Merola	Lawrence Berkeley Labs
James Morrill	Sparta
Russ Mundy	DCA
Dan Nessett	Livermore
Ari Ollikainen	RIACS
David Peters	NASA
Nachum Shacham	SRI
Henry Sowizral	RIACS
Mike St. Johns	DCA
Paul Tsuchiya	Mitre
Tony Villasenor	NASA
Steve Walker	TIS
Jil Westcott	BBN
Steve Wolff	NSF

Lixia Zhang

MIT

8. Glossary

AR	Autonomous Region
CLNP	Connectionless Network Protocol
DARPA	Defense Advanced Research Projects Agency
DES	Data Encryption Standard
DoE	Department of Energy
ECMA	European Computer Manufacturers Association
FRICC	Federal Research Internet Coordinating Committee
GOSIP	Government OSI Protocol
IETF	Internet Engineering Task Force
IP	Internet Protocol
ISO	International Standards Organization
LAN	Local Area Network
MTA	Mail Transfer Agent
NASA	National Aeronautics and Space Administration
NRI	National Research Internet
NSF	National Science Foundation
OMB	Office of Management and Budget
OSTP	White House Office of Science and Technology Policy
PS	Policy Server
PT	Policy Term
RSA	Rivest Shamir Algorithm
TAC	Terminal Access Controller
TOS	Type of Service
QOS	Quality of Service

Security Considerations

None.

Author's Address

Barry Leiner
Research Institute for Advanced Computer Science
National Aeronautics and Space Administration
Ames Research Center
Mail Stop 230-5
Moffett Field, CA 94035

Phone: (415) 694-5402

EMail: LEINER@RIACS.EDU

TABLE OF CONTENTS

	Page
1. Introduction	1
2. Workshop Summary	3
3. Working Group on Interconnection Policies	9
3.1. Existing Policies, Summarized	10
3.2. Refined Policy Statements	11
4. Access Control for Network Switching and Transmission Resources	14
4.1. Introduction	14
4.2. Access Control Policy Issues	15
4.2.1. Policies and Models	15
4.2.2. Policy Inputs	16
4.3. Communication Scenarios	18
4.3.1. Connection-Oriented Communication	18
4.3.2. Variations on Connection-Oriented Scenarios	19
4.3.3. Electronic Messaging	20
4.3.4. Transaction-Oriented Communication	21
4.3.5. Multicast Communication	21
4.4. Access Control Architectures	22
4.4.1. Analogies with Operating System Security	22
4.4.2. Clark's Policy Routing Model and Access Control	23
4.4.3. Clark's Architecture in Retrospect	26
4.4.4. Trust Implications and Possible Remedies	27
5. Resource Sharing	30
5.1. Introduction	30
5.2. Service Class	30
5.3. User Categories	31
5.4. Additional Discussion	32
5.4.1. Accounting for usage:	32
5.4.2. Levels of assurance:	32
5.4.3. Global effects:	33
5.5. Conclusions	34
5.6. Recommendations	34
5.6.1. Instant projects	35
5.6.2. Short-term experiments	35
5.6.3. Longer-term experiments	36
6. End-to-End Security Services	38
6.1. Introduction	38
6.2. Multi-administrative Security Architecture	38
6.2.1. Security Domains	40

6.3. Higher-Level End-to-End Services	40
6.3.1. Supportive Services	41
6.3.2. Productive Services	42
6.4. Projects	44
7. Workshop Attendees	46
8. Glossary	48