Security Considerations

This memo does not address the security aspects of the issues discussed.

Author's Address

Deborah Estrin
University of Southern California
Computer Science Department
Los Angeles, CA 90089-0782

Phone: (213) 743-7842

EMail: Estrin@OBERON.USC.EDU

References

[1] J. Postel, *Internet Protocol*, **Network Information Center, RFC 791**, September 1981.

[2] G. Vaudreuil, *The Federal Research Internet Coordinating Committee and the National Research Network*, **ACM SIG Computer Communications Review**, April 1988.

[3] E. Rosen, *Exterior Gateway Protocol (EGP)*, **Network Information Center, RFC 827**, October 1982.

[4] D. Clark, *Policy Routing in Internet Protocols*, **Network Information Center, RFC 1102**, May 1989.

[5] H.W.Braun, *Models of Policy Based Routing*, **Network Information Center, RFC 1104**, June 1989.

[6] K. Lougheed, Y. Rekhter, *A Border Gateway Protocol*, **Network Information Center, RFC 1105**, June 1989.

[7] J. Saltzer, M. Schroeder, *The Protection of Information in Computer Systems*, **Proceedings of the IEEE**, 63, 9 September 1975.

[8] V. Jacobson, *Congestion Avoidance and Control.* **Proceedings of ACM Sigcomm**, pp. 106-114, August 1988, Palo Alto, CA.

[9] David Clark, *Design Philosophy of the DARPA Internet Protocols*, **Proceedings of ACM Sigcomm**, pp. 106-114, August 1988, Palo Alto, CA.

[10] Gigabit Networking Group, B. Leiner, Editor. *Critical Issues in High Bandwidth Networking*, **Network Information Center, RFC 1077**, November 1988.

[11] D. Estrin, J. Mogul and G. Tsudik, *Visa Protocols for Controlling Inter-Organizational Datagram Flow*, **To appear in IEEE Journal on Selected Areas in Communications**, Spring 1989.

[12] D. Estrin and G. Tsudik, *Security Issues in Policy Routing*, **IEEE Symposium on Research in Security and Privacy**, Oakland, CA. May 1-3 1989.

[13] M. Little, *The Dissimilar Gateway Protocol*, **Technical report**

[14] P. Tsuchiya, *The Landmark Hierarchy: A new hierarchy for routing in very large networks*, **IEEE SIGCOMM 88**, Palo Alto, CA. September 1988.

[15] G. Finn, *Reducing the Vulnerability of Dynamic Computer Networks* **USC/Information Sciences Institute, Technical Report, ISI/RR-88-201** July 1988.

[16] A. Nakassis *Routing Algorithm for Open Routing*, **Unpublished paper**, Available from the author at the National Institute of Standards and Technology (formerly NBS), Washington D.C.

source routing. [12] [13]

9 Summary

Along with the emergence of very high speed applications and media, resource management has become a critical issue in the Research Internet and internets in general. A fundamental characteristic of the resource management problem is allowing administratively ADs to interconnect while retaining control over resource usage. However, we have lacked a careful articulation of the types of resource management policies that need to be supported. This paper addresses policy requirements for the Research Internet. After justifying our assumptions regarding AD topology we presented a taxonomy and examples of policies that must be supported by a PR protocol.

10 Acknowledgments

Members of the Autonomous Networks Research Group and Open Routing Working Group have contributed significantly to the ideas presented here, in particular, Guy Almes, Lee Breslau, Scott Brim, Dave Clark, Marianne Lepp, and Gene Tsudik. In addition, Lee Breslau and Gene Tsudik provided detailed comments on a previous draft. David Cheriton inadvertently caused me to write this document. Sharon Anderson's contributions deserve special recognition. The author is supported by research grants from National Science Foundation, AT&T, and GTE.

[12] Moreover, the source routing approach loosens the requirements for every AD to share a complete view of the entire internet by allowing the source to detect routing loops.

[13] The match between RFC1102 and the requirements specified in this document is hardly a coincidence since Clark's paper and discussions with him contributed to the requirements formulation presented here. His work is currently being evaluated and refined by the ANRG and ORWG.

is that access to an AD at any point in time is contingent upon a local, highly dynamic, parameter that is not globally available. Therefore such a policy term could well result in looping, oscillations, and excessive route (re)computation overhead, both unacceptable. Consequently, this is one type of policy that routing experts suggest would be difficult to support in a very large decentralized internetwork.

- Granularity can also be problematic, but not as devistating as highly dynamic PR contingencies. Here the caution is less specific. Very fine grain policies, which restrict access to particular hosts, or are contingent upon very fine grain user class identification, may be achieved more efficiently with network level access control[11] or end system controls instead of burdening the inter-AD routing mechanism.

- Security is expensive, as always. Routing protocols are subject to fraud through impersonation, data substitution, and denial of service. Some of the proposed mechanisms provide some means for detection and non-repudiation. However, to achieve a priori prevention of resource misuse is expensive in terms of per connection or per packet cryptographic overhead. For some environments we firmly believe that this will be necessary and we would prefer an architecture that would accommodate such variability[12].

In general, it is difficult to predict the impact of any particular policy term. Tools will be needed to assist people in writing and validating policy terms.

8 Proposed mechanisms

Previous routing protocols have addressed a narrower definition of PR, as appropriate for the internets of their day. In particular, EGP[3], DGP[13], and BGP[6] incorporate a notion of policy restrictions as to where routing database information travels. None are intended to support policy based routing of packets as described here. More recent routing proposals such as Landmark[14] and Cartesian[15] could be used to restrict packet forwarding but are not suited to source/destination, and some of the condition-oriented, policies. We feel these policy types are critical to support. We note that for environments (e.g., within an AD substructure) in which the simple-AD-topology conjecture holds true, these alternatives may be suitable.

RFC 1104[5] provides a good description of shorter term policy routing requirements. Braun classifies three types of mechanisms, policy based distribution of route information, policy based packet forwarding, and policy based dynamic allocation of network resources. The second class is characterized by Dave Clark's PR architecture, RFC 1102[4]. With respect to the longer term requirements laid out in this document, only this second class is expressive and flexible enough to support the multiplicity of stub and transit policies. In other words, the power of the PR approach (e.g., RFC1102) is not just in the added granularity of control pointed out by Braun, i.e., the ability to specify particular hosts and user classes. Its power is in the ability to express and enforce many types of stub and transit policies and apply them on a discriminatory basis to different ADs. In addition, this approach provides explicit support for stub ADs to control routes via the use of

Regional B

1. Regional B 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.

   $[RegionalB : (*, \{F/Re/U\}, \{F/Re/U\})(*, \{F/Re/U\}, Cc)\{\}$
   $\{unauthenticatedUCI, per - kilopacketcharge\}\{\}]$

6.3.3 Campus and Private Networks

Similar interviews should be conducted with administrators of campus and private networks. However, many aspects of their policies are contingent on the still unresolved policies of the regionals and federal agencies. In any event, transit policies will be critical for campus and private networks to flexibly control access to lateral links and private wide area networks, respectively. For example, a small set of university and private laboratories may provide access to special gigabit links for particular classes of researchers. On the other hand, source/destination policies should not be used in place of network level access controls for these end ADs.

6.3.4 Commercial Services

Currently commercial communication services play a low level role in most parts of today's Research Internet; they provide the transmission media, i.e., leased lines. In the future we expect commercial carriers to provide increasingly higher level and enhanced services such as high speed packet switched backbone services. Because such services are not yet part of the Research Internet infrastructure there exist no policy statements.

Charging and accounting are certain to be an important policy type in this context. Moreover, we anticipate the long haul services market to be highly competitive. This implies that competing service providers will engage in significant gaming in terms of packaging and pricing of services. Consequently, the ability to express varied and dynamic charging policies will be critical for these ADs.

7 Problematic requirements

Most of this paper has lobbied for articulation of relatively detailed policy statements in order to help define the technical mechanisms needed for enforcement. We promoted a top down design process beginning with articulation of desired policies. Now we feel compelled to mention requirements that are clearly problematic from the bottom up perspective of technical feasibility.

- Non-interference policies are of the form "I will provide access for principals x to resources y so long as it does not interfere with my internal usage." The problem with such policies

Defense Advanced Research Projects Agency (DARPA)

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

   $[DARPA1 : (*,*,*)(*, DARPA, -)\{research, support\}$
   $\{unauthenticated - UCI, noperpacketcharge\}\{\}]$

2. 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 there is 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. NOTE: Darpa would like to say something about the need to enter the Darpa AD at the point closest to the destination...but i don't know how to express this...

   $[DARPA2 : (*, \{F/R/U/Co\}, \{F/R/U/Co\})(*, \{F/R/U/Co\}, \{F/R/U/Co\})$
   $\{research, support\}\{unauthenticated - UCI, noperpacketcharge,$
   $non - interferencebasis\}\{\}]$

Defense Communications Agency (DCA)

1. DCA 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. All packets are marked and charged for by the kilopacket.

   $[DCA1 : (mailbridge, DCA, -)(*, \{F/Re\}, \{F/Re\})\{research, support\}$
   $\{unauthenticatedUCI, allincomingpacketsmarked, per - kilopacketcharge\}\{\}]$

6.3.2 The Regionals

Interviews with regional network administrations are now underway. In general their policies are still in formation due to the relatively recent formation of these regional networks. However, for the sake of illustration we provide an example of a hypothetical regional's network policies.

Regional A

1. Regional A 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 pass it traffic that complies with its, NSF's, policies!)

   $[RegionalA : (*, \{F/Re/U\}, \{F/Re/U\})(*, \{F/Re/U\}, NSF)\{research, support\}$
   $\{unauthenticatedUCI, no - per - packetcharge\}\{\}]$

Department of Energy (DOE)

1. 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.

   $[DOE1 : (*, DOE, -)(*, *, *)\{research, support\}$
   $\{unauthenticatedUCI, no - per - packetcharge\}\{\}]$

2. 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.)

   $[DOE2 : (*, \{F/Re\}, \{F/Re\})(*, \{F/Re\}, \{F/Re\})\{\}$
   $\{unauthenticatedUCI, no - per - pktcharge\}\{\}]$

National Aeronautics and Space Administration (NASA)

1. Nasa will accept any traffic to/from members of the Nasa AD. But no transit. No UCI authentication and no per packet charge.

   $[NASA1 : (*, *, *)(*, Nasa, -)\{Nasa - research, support\}$
   $\{unauthenticatedUCI, no - per - packet - charge\}\{\}]$

2. Nasa will carry transit traffic to/from other federal agency networks if it is in support of research, and if the total use of available BW by non-nasa Federal agencies is below nnon-interference policy type needs some more work in terms of integrating it into the routing algorithms. See Section 7.

   $[NASA2 : (*, \{F\}, *)(*, \{F\}, *)\{research, support\}$
   $\{per - packetaccounting, limitedton\%ofavailableBW\}\{\}]$

3. NASA will carry commercial traffic to federal and regional and university ADs for nasa research or support. But it will not allow transit. The particular entry AD is not important.

   $[NASA3 : (*, \{Co\}, *\}(*, \{F/R/U\}, *)\{NASAresearch, support\}$
   $\{unauthenticatedUCI, noperpacketcharge\}\{\}]$

4. On a case by case basis NASA may provide access to its resources on a cost reimbursed basis. Transit traffic will not be carried on this basis.

   $[NASA4 : (*, *, -)(*, *, -)\{\}$
   $\{per - packet - charge, limitedton\%ofavailableBW\}\{\}]$

for Regional, U for University, Co for Commercial Corporation, and Cc for Commercial Carrier. A hyphen, -, means no applicable matches.

By examining a PT we can identify the type of policy represented, as per the taxonomy presented earlier.

- If an AD specifies a policy term that has a null (-) entry for the ADexit, then it is disallowing transit for some group of users, and it is a transit policy.
- If an AD specifies a policy term that lists itself explicitly as ADsrc or ADdst, it is expressing restrictions on who can access particular resources within its boundaries, or on who inside can obtain external access. In other words the AD is expressing a source/destination policy.
- If ADexit or ADentr is specified then the policy expressed is an exit/entrance path policy.
- If the global conditions include charging, QOS, resource guarantee, time of day, higher level application, resource limit, or authentication related information it is obviously a charging, QOS, resource guarantee, temporal, higher level application, resource limit, or authentication policy, respectively.

As seen below, any one PT typically incorporates a combination of policy types.

6.3.1 The FRICC

In the following examples all policies (and PTs) are symmetrical under the assumption that communication is symmetrical.

National Science Foundation (NSF)

1. 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 there is 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 ADent and ADexit.

   $[NSF1 : (*, \{F/Re\}, \{F/Re\})(*, \{F/Re\}, \{F/Re\})\{research, support\}$
   $\{unauthenticatedUCI, no - per - pktcharge\}\{\}]$

2. NSF will carry traffic to user and expert services hosts in NSF AD to/from any F/Re AD, via any F/Re AD. These are the only things that directly connect to NSFnet.

   $[NSF2 : (\{Usersvcs, ExpertSvcs\}, \{NSF\}, \{F/Re\})(*, \{F/Re\}, \{-\})\{\}\{\}\{\}]$

6.2 Taxonomy of Charging Policies

Stub and transit charging policies may specify the following parameters:

- **Unit of accounting** (e.g., dollars or credits).
- **Basis for charging** (e.g., per Kbyte or per Kpkt).
- **Actual charges** (e.g., actual numbers such as $.50/Mbyte).
- **Who is charged or paid** (e.g., originator of packet, immediate neighbor from whom packet was received, destination of packet, a third party collection agent).
- **Whose packet count** is used (e.g., source, destination, the transit AD's own count, the count of some upstream or downstream AD).
- **Bound on charges** (e.g., to limit the amount that a stub AD is willing to spend, or the amount that a transit AD is willing to carry.)

The enforcement of these policies may be carried out during route synthesis or route selection[4]

6.3 Example Policy statements

The following policy statements were collected in the fall of 1988 through interviews with representatives of the federal agencies most involved in supporting internetworking. Once again we emphasize that these are *not official policy statements*. They are presented here to provide concrete examples of the sort of policies that agencies would like to enforce.

Expressing policies as Policy Terms (PTs) Each policy is described in English and then expressed in a *policy term (PT)* notation suggested by Dave Clark in [4]. Each PT represents a distinct policy of the AD that synthesized it. The format of a PT is:

$$[(H_{src}, AD_{src}, AD_{ent}), (H_{dst}, AD_{dst}, AD_{exit}), UCI, Cg, Cb]$$

Hsrc stands for source host, ADsrc for source AD, ADent for entering AD (i.e., neighboring AD from which traffic is arriving directly), Hdst for destination host, ADdst for destination AD, ADexit for exit AD (i.e., neighboring AD to which traffic is going directly), UCI for user class identifier, and Cg and Cb for global and bilateral conditions, respectively. The purpose of a PT is to specify that packets from some host, H_{src}, (or a group of hosts) in a source AD, AD_{src}, are allowed to enter the AD in question via some directly connected AD, AD_{ent}, and exit through another directly connected AD, AD_{exit}, on its way to a host, H_{dst}, (or a group of hosts) in some destination AD, AD_{dst}. User Class Identifier (UCI) allows for distinguishing between various user classes, e.g., Government, Research, Commercial, Contract, etc. Global Conditions (Cg) represent billing and other variables. Bilateral Conditions (Cb) relate to agreements between neighboring ADs, e.g., related to metering or charging. In the example policy terms provided below we make use of the following abbreviations: Fricc for {DOE,NASA,DCA,NSF}, F for Federal Agency, Re reject a route based on any AD (or combination of ADs) in the route. Similarly, a transit AD could express a packet forwarding policy that behaves differently depending upon which ADs a packet has passed through, and is going to pass through, en route to the destination. Less ambitious (and perhaps more reasonable) path sensitive policies might only discriminate according to the immediate neighbor ADs through which the packet is traveling (i.e., a stub network could reject a route based on the first transit AD in the route, and a transit AD could express a packet forwarding policy that depends upon the previous, and the subsequent, transit ADs in the route.)

- **Quality/Type of Service (QOS or TOS)**
  This type of policy restricts access to special resources or services. For example, a special high throughput, low delay link may be made available on a selective basis.

- **Resource Guarantee**
  These policies provide a guaranteed percentage of a resource on a selective, as needed basis. In other words, the resource can be used by others if the preferred-AD's offered load is below the guaranteed level of service. The guarantee may be to always carry intra-AD traffic or to always carry inter-AD traffic for a specific AD.

- **Temporal**
  Temporal policies restrict usage based on the time of day or other time related parameters.

- **High Level Protocol**
  Usage may be restricted to a specific high level protocol such as mail or file transfer. (Alternatively, such policies can be implemented as source/destination policies by configuring a host(s) within an AD as an application relay and composing policy terms that allow inter-AD access to only that host.)

- **Resource limit**
  There may be a limit on the amount of traffic load a source may generate during a particular time interval, e.g., so many packets in a day, hour, or minute.

- **Authentication requirements**
  Conditions may be specified regarding the authenticability of principal identifying information. Some ADs might require some form of cryptographic proof as to the identity and affiliations of the principal before providing access to critical resources.

The above policy types usually exist in combination for a particular AD, i.e., an AD's policies might express a combination of transit, source/destination, and QOS restrictions. This taxonomy will evolve as PR is applied to other domains.

As will be seen in Section 6.3 an AD can express its charging and access policies in a single syntax. Moreover, both stub and transit policies can coexist. This is important since some ADs operate as both stub and transit facilities and require such hybrid control.

6 Policy types

This section outlines a taxonomy of internet policies for inter-AD topologies that allow lateral and bypass links. The taxonomy is intended to cover a wide range of ADs and internets. Any particular PR architecture we design should support a significant subset of these policy types but may not support all of them due to technical complexity and performance considerations. The general taxonomy is important input to a functional specification for PR. Moreover, it can be used to evaluate and compare the suitability and completeness of existing routing architectures and protocols for PR; see Section 8.

We provide examples from the Research Internet of the different policy types in the form of resource usage policy statements. These statements were collected through interviews with agency representatives, but they *do not* represent official policy. These sample policy statements should *not* be interpreted as agency policy, they are provided here only as examples.

Internet policies fall into two classes, access and charging. Access policies specify who can use resources and under what conditions. Charging policies specify the metering, accounting, and billing implemented by a particular AD.

6.1 Taxonomy of Access Policies

We have identified the following types of access policies that ADs may wish to enforce. Charging policies are described in the subsequent section. Section 6.3 provides more specific examples of both access and charging policies using FRICC policy statements .

Access policies typically are expressed in the form: *principals of type x can have access to resources of type y under the following conditions, z.* The policies are categorized below according to the definition of y and z. In any particular instance, each of the policy types would be further qualified by definition of legitimate principals, x, i.e., what characteristics x must have in order to access the resource in question.

We refer to access policies described by stub and transit ADs. The two roles imply different motivations for resource control, however the types of policies expressed are similar; we expect the supporting mechanisms to be common as well.

Stub and transit access policies may specify any of the following parameters:

- **Source/Destination**
  Source/Destination policies prevent or restrict communication originated by or destined for particular ADs (or hosts or user classes within an AD).
- **Path**
  Path sensitive policies specify which ADs may or may not be passed through en route to a destination. The most general path sensitive policies allow stub and transit ADs to express policies that depend on *any component* in the AD path. In other words, a stub AD could

the complex case, lateral connections must be supported, along with the means to control the use of such connections in the routing protocols.

The different topologies imply different policy requirements. The first model assumes that all policies can be expressed and enforced in terms of dollars and cents and distributed charging schemes. The second model assumes that ADs want more varied control over their resources, control that can not be captured in a dollars and cents metric alone. We describe the types of policies to be supported and provide examples in the following section, Section 6. In brief, given private lateral links, ADs must be able to express access and charging related restrictions and privileges that discriminate on an AD basis. These policies will be diverse, dynamic, and new requirements will emerge over time, consequently support must be extensible. For example, the packaging and charging schemes of any single long haul service will vary over time and may be relatively elaborate (e.g., many tiers of service, special package deals, to achieve price discrimination).

Note that these assumptions about complexity do not preclude some collection of ADs from "negotiating away" their policy differences, i.e., forming a federation, and coordinating a simplified inter-AD configuration in order to reduce the requirements for inter-AD mechanisms. However, we maintain that there will persist collections of ADs that will not and can not behave as a single federation; both in the research community and, even more predominantly, in the broader commercial arena. Moreover, when it comes to interconnecting across these federations, non-negotiable differences will arise eventually. It is our goal to develop mechanisms that are applicable in the broader arena.

The Internet community developed its original protocol suite with only minimal provision for resource control[9]. This was appropriate at the time of development based on the assumed community (i.e., researchers) and the ground breaking nature of the technology. The next generation of network technology is now being designed to take advantage of high speed media and to support high demand traffic generated by more powerful computers and their applications.[10] As with TCP/IP we hope that the technology being developed will find itself applied outside of the research community. This time it would be inexcusable to ignore resource control requirements and not to pay careful attention to their specification.

Finally, we look forward to the Internet structure taking advantage of economies of scale offered by enhanced commercial services. However, in many respects the problem that stub-ADs may thus avoid, will be faced by the multiple regional and long haul carriers providing the services. The carriers' charging and resource control policies will be complex enough to require routing mechanisms similar to ones being proposed for the complex AD topology case described here. Whether the network structure is based on private or commercial services, the goal is to construct policy sensitive mechanisms that will be transparent to end users (i.e., the mechanisms are part of the routing infrastructure at the network level, and not an end to end concern).

expect private data networks to persist for the near future. As the telephone companies begin to introduce the next generation of high speed packet switched services, the scenario should change. However, we maintain that the result will be a predominance, but not complete dominance, of public carrier use for long haul communication. Therefore, private data networks will persist and the routing architecture must accommodate controlled interconnection.

to topology and policy. They contend that in the long term the following three conditions will prevail:

- The public carriers will provide pervasive, competitively priced, high speed data services.
- The resulting topology of ADs will be stub (not transit) ADs connected to regional backbones, which in turn interconnect via multiple, overlapping long haul backbones, i.e., a hierarchy with no lateral connections between stub-ADs or regionals, and no vertical bypass links.
- The policy requirements of the backbone and stub-ADs will be based only on charging for resource usage at the stub-AD to backbone-AD boundary, and to settling accounts between neighboring backbone providers (regional to long haul, and long haul to long haul).

Under these assumptions, the primary requirement for general AD interconnect is a metering and charging protocol. The routing decision can be modeled as a simple least cost path with the metric in dollars and cents. In other words, restrictions on access to transit services will be minimal and the functionality provided by the routing protocol need not be changed significantly from current day approaches.

Complex AD topology and policy model The counter argument is that a more complex AD topology will persist.[10] The different assumptions about AD topology lead to the significantly different assumptions about AD policies.

This model assumes that the topology of ADs will in many respects agree with the previous model of increased commercial carrier participation and resulting hierarchical structure. However, we anticipate unavoidable and persistent exceptions to the hierarchy. We assume that there will be a relatively small number of long haul transit ADs (on the order of 100), but that there may be tens of thousands of regional ADs and hundreds of thousands of stub ADs (e.g., campuses, laboratories, and private companies). The competing long haul offerings will differ, both in the services provided and in their packaging and pricing. Regional networks will overlap less and will connect campus and private company networks. However, many stub-ADs will retain some private lateral links for political, technical, and reliability reasons. For example, political incentives cause organizations to invest in bypass links that are not always justifiable on a strict cost comparison basis; specialized technical requirements cause organizations to invest in links that have characteristics (e.g., data rate, delay, error, security) not available from public carriers at a competitive rate; and critical requirements cause organizations to invest in redundant back up links for reliability reasons. These exceptions to the otherwise regular topology are not dispensible. They will persist and must be accommodated, perhaps at the expense of optimality; see Section 5 for more detail. In addition, many private companies will retain their own private long haul network facilities.[11] Critical differences between the two models follow from the difference in assumptions regarding AD topology. In

[10] Much of the remainder of this paper attempts to justify and provide evidence for this statement.

[11] While private voice networks also exist, private data networks are more common. Voice requirements are more standardized because voice applications are more uniform than are data applications, and therefore the commercial services more often have what the voice customer wants at a price that is competitive with the private network option. Data communication requirements are still more specialized and dynamic. Thus, there is less opportunity for economy of scale in service offerings and it is harder to keep up to date with customer demand. For this reason we

laboratory. They reside in a campus AD along with users who are legitimate users of other AD resources. Moreover, any one person may be a legitimate user of multiple AR resources under varying conditions and constraints (see examples in section 6). In addition, users can move from one AD to another. In other words, a user's rights can not be determined solely based on the AD from which the user's communications originate. Consequently, PR must not only identify resources, it must identify *principals*[8] and associate different capabilities and rights with different principals.

One way of reducing the compromise of autonomy associated with interconnection is to implement mechanisms that assure *accountability* for resources used. Accountability may be enforced a priori, e.g., access control mechanisms applied before resource usage is permitted. Alternatively, accountability may be enforced after the fact, e.g., record keeping or metering that supports detection and provides evidence to third parties (i.e., non-repudiation). Accountability mechanisms can also be used to provide feedback to users as to consumption of resources. Internally an AD often decides to do away with such feedback under the premise that communication is a global good and should not be inhibited. There is not necessarily a "global good" across AD boundaries. Therefore, it becomes more appropriate to have resource usage visible to users, whether or not actual charging for usage takes place. Another motivation that drives the need for accountability across AD boundaries is the greater variability in implementations. Different implementations of a single network protocol can vary greatly as to their efficiency[8]. We can not assume control over implementation across AD boundaries. Feedback mechanisms such as metering (and charging in some cases) would introduce a concrete incentive for ADs to employ efficient and correct implementations. PR should allow an AD to advertise and apply such accounting measures to inter-AD traffic.

In summary, the lack of global authority, the need to support network resource sharing as well as network interconnection, the complex and dynamic mapping of users to ADs and rights, and the need for accountability across ADs, are characteristics of inter-AD communications which must be taken into account in the design of both policies and supporting technical mechanisms.

5 Topology model of Internet

Before discussing policies per se, we outline our model of inter-AD topology and how it influences the type of policy support required. Most members of the Internet community agree that the future Internet will connect on the order of 150,000,000 termination points and 100,000 ADs. However, there are conflicting opinions as to the AD topology for which we must design PR mechanisms. The informal argument is described here.

Simple AD topology and policy model Some members of the Internet community believe that the current complex topology of interconnected ADs is a transient artifact resulting from the evolutionary nature of the Research Internet's history.[9] The critical points of this argument relate

[8] The term principal is taken from the computer security community[7].

[9] David Cheriton of Stanford University articulated this side of the argument at an Internet workshop in Santa Clara, January, 1989.

4 Why the problem is difficult

Before proceeding with our description of topology and policy requirements this section outlines several assumptions and constraints, namely: the lack of global authority, the need to support network resource sharing as well as network interconnection, the complex and dynamic mapping of users to ADs and privileges, and the need for accountability across ADs. These assumptions limit the solution space and raise challenging technical issues.

The purpose of policy based routing is to allow ADs to interconnect and share computer and network resources in a controlled manner. Unlike many other problems of resource control, there is no global authority. Each AD defines its own policies with respect to its own traffic and resources. However, while we assume no global authority, and no global policies, we recognize that complete autonomy implies no dependence and therefore no communication. The multi-organization internets addressed here have inherent regions of autonomy, as well as requirements for interdependence. Our mechanisms should allow ADs to design their boundaries, instead of requiring that the boundaries be either impenetrable or eliminated.

One of the most problematic aspects of the policy routing requirements identified here is the need to support both network resource sharing and interconnection across ADs. An example of resource sharing is two ADs (e.g., agencies, divisions, companies) sharing network resources (e.g., links, or gateways and links) to take advantage of economies of scale. Providing transit services to external ADs is another example of network resource sharing. Interconnection is the more common example of ADs interconnecting their independently used network resources to achieve connectivity across the ADs, i.e., to allow a user in one AD to communicate with users in another AD. In some respects, network resource control is simpler than network interconnection control since the potential dangers are fewer (i.e., denial of service and loss of revenue as compared with a wide range of attacks on end systems through network interconnection). However, controlled network resource sharing is more difficult to support. In an internet a packet may travel through a number of transit ADs on its way to the destination. Consequently, policies from all transit ADs must be considered when a packet is being sent, whereas for stub-AD control only the policies of the two end point ADs have to be considered. In other words, controlled network resource sharing and transit require that policy enforcement be integrated into the routing protocols themselves and can not be left to network control mechanisms at the end points.[6] [7]

Complications also result from the fact that legitimate users of an AD's resources are not all located in that AD. Many users (and their computers) who are funded by, or are affiliated with, a particular agency's program reside within the AD of the user's university or research

[6]Another difference is that in the interconnect case, traffic traveling over AD A's network resources always has a member of AD A as its source or destination (or both). Under resource sharing arrangements members of both AD A and B are connected to the same resources and consequently intra-AD traffic (i.e., packets sourced ***and*** destined for members of the same AD) travels over the resources. This distinction is relevant to the writing of policies in terms of principal affiliation.

[7]Economies of scale is one motivation for resource sharing. For example, instead of interconnecting separately to several independent agency networks, a campus network may interconnect to a shared backbone facility. Today, interconnection is achieved through a combination of AD specific and shared arrangements. We expect this mixed situation to persist for "well-connected" campuses for reasons of politics, economics, and functionality (e.g., different characteristics of the different agency-networks). See Section 5 for more discussion.

3.1 Policy Routing

Previous protocols such as the Exterior Gateway Protocol (EGP)[3] embodied a limited notion of policy and ADs. In particular, autonomous system boundaries constrained the flow of routing database information, and only indirectly affected the flow of packets themselves. We consider an Administrative Domain (AD) to be a set of hosts and network resources (gateways, links, etc.) that is governed by common policies. In large internets that cross organization boundaries, e.g., the Research Internet, inter-AD routes must be selected according to policy-related parameters such as cost and access rights, in addition to the traditional parameters of connectivity and congestion. In other words, Policy Routing (PR) is needed to navigate through the complex web of policy boundaries created by numerous interconnected ADs. Moreover, each AD has its own privileges and perspective and therefore must make its own evaluation of legal and preferred routes. Efforts are now underway to develop a new generation of routing protocol that will allow each AD to independently express and enforce policies regarding the flow of packets to, from, and through its resources[4].[4]

The purpose of this paper is to articulate the requirements for such policy based routing. Two critical assumptions will shape the type of routing mechanism that is devised:

- The topological organization of ADs, and
- The type and variability of policies expressed by ADs.

We make use of the policies expressed by owners of current Research Internet resources and private networks connected to the Research Internet to generalize types of policies that must be supported. This top down effort must be done with attention to the technical implications of the policy statements if the result is to be useful in guiding technical development. For example, some ADs express the desire to enforce local constraints over how packets travel to their destination. Other ADs are only concerned with preventing use of their own network resources by restricting transit. Still other ADs are concerned primarily with recovering the expense of carrying traffic and providing feedback to users so that users will limit their own data flows; in other words they are concerned with charging. We refer to ADs whose primary concern is communication to and from hosts within their AD as *stub* and to ADs whose primary concern is carrying packets to and from other ADs as *transit*. If we address control of transit alone, for example, the resulting mechanisms will not necessarily allow an AD to control the flow of its packets from source to destination, or to implement flexible charging schemes.[5] Our purpose is to articulate a comprehensive set of requirements for PR as input to the functional specification, and evaluation, of proposed protocols.

[4] These issues are under investigation by the IAB Autonomous Networks Research Group and the IAB Open Routing Working Group. For further information contact the author.

[5] Gene Tsudik uses the analogy of international travel to express the need for source and transit controls. Each country expresses its own policies about travel to and through its land. Travel through one country en route to another is analogous to transit traffic in the network world. A traveler collects policy information from each of the countries of interest and plans an itinerary that conforms to those policies as well as the preferences of the traveler and his/her home nation. Thus there is both source and transit region control of routing.

3 Background

The *Research Internet*[2] has evolved from a single backbone wide area network with many connected campus networks, to an internet with multiple cross-country backbones, regional access networks, and a profusion of campus networks. At times during its development the Research Internet topology appeared somewhat chaotic. Overlapping facilities and lateral (as opposed to hierarchical) connections seemed to be the rule rather than the exception. Today the Research Internet topology is becoming more regular through coordination of agency investment and adoption of a hierarchy similar to that of the telephone networks'. The result is several overlapping wide area backbones connected to regional networks, which in turn connect to campus networks at universities, research laboratories, and private companies. However, the telephone network has lateral connections only at the highest level, i.e., between long haul carriers. In the Research Internet there exist lateral connections at each level of the hierarchy, i.e., between campus (and regional) networks as well.

Additional complexity is introduced in the Research Internet by virtue of connections to private networks. Many private companies are connected to the Research Internet for purposes of research or support activities. These private companies connect in the same manner as campuses, via a regional network or via lateral links to other campuses. However, many companies have their own private wide area networks which physically overlap with backbone and/or regional networks in the research internet, i.e., private vertical bypass links.

Implicit in this complex topology are organizational boundaries. These boundaries define Administrative Domains (ADs) which preclude the imposition of a single, centralized set of policies on all resources. The subject of this paper is the policy requirements for resource usage control in the Research Internet.

In the remainder of this section we describe the policy routing problem in very general terms. Section 4 examines the constraints and requirements that makes the problem challenging, and leads us to conclude that a new generation of routing and resource control protocols are needed. Section 5 provides more detail on our assumptions as to the future topology and configuration of interconnected ADs. We return to the subject of policy requirements in Section 6 and categorize the different types of policies that ADs in the research internet may want to enforce. Included in this section are examples of FRICC[3] policy statements. Section 7 identifies types of policy statements that are problematic to enforce due to their dynamics, granularity, or performance implications. Several proposed mechanisms for supporting PR (including RFCs 827, 1102, 1104, 1105) are discussed briefly in Section 8. Future RFCs will elaborate on the architecture and protocols needed to support the requirements presented here.

[2]The term Research Internet refers to a collection of government, university, and some private company, networks that are used by researchers to access shared computing resources (e.g., supercomputers), and for research related information exchange (e.g., distribution of software, technical documents, and email). The networks that make up the Research Internet run the DOD Internet Protocol[1].

[3]The Federal Research Internet Coordinating Committee (FRICC) is made up of representatives of each of the major agencies that are involved in networking. They have been very effective in coordinating their efforts to eliminate inefficient redundancy and have proposed a plan for the next 10 years of internetworking for the government, scientific, and education community[2].

Network Working Group
Request for Comments: 1125

D. Estrin
U.S.C. Computer Science Department
November 1989

Policy Requirements for Inter Administrative Domain Routing

1 Status of this Memo

The purpose of this memo is to focus discussion on particular problems in the Internet and possible methods of solution. No proposed solutions in this document are intended as standards for the Internet. Rather, it is hoped that a general consensus will emerge as to the appropriate solution to such problems, leading eventually to the development and adoption of standards. Distribution of this memo is unlimited.

2 Abstract

Efforts are now underway to develop a new generation of routing protocol that will allow each Administrative Domain (AD) in the growing Internet (and internets in general) to independently express and enforce policies regarding the flow of packets to, from, and through its resources.[1] This document articulates the requirements for policy based routing and should be used as input to the functional specification and evaluation of proposed protocols.

Two critical assumptions will shape the type of routing mechanism that is devised: (1) the topological organization of ADs, and (2) the type and variability of policies expressed by ADs. After justifying our assumptions regarding AD topology we present a taxonomy, and specific examples, of policies that must be supported by a PR protocol. We conclude with a brief discussion of policy routing mechanisms proposed in previous RFCs (827, 1102, 1104, 1105). Future RFCs will elaborate on the architecture and protocols needed to support the requirements presented here.

[1] The material presented here incorporates discussions held with members of the IAB Autonomous Networks Research Group and the Open Routing Working Group.