                       Models of Policy Based Routing

1. Status of this Memo

   The purpose of this RFC is to outline a variety of models for policy
   based routing. The relative benefits of the different approaches are
   reviewed. Discussions and comments are explicitly encouraged to move
   toward the best policy based routing model that scales well within a
   large internetworking environment.

   Distribution of this memo is unlimited.

2. Acknowledgements

   Specific thanks go to Yakov Rekhter (IBM Research), Milo Medin
   (NASA), Susan Hares (Merit/NSFNET), Jessica Yu (Merit/NSFNET) and
   Dave Katz (Merit/NSFNET) for extensively contributing to and
   reviewing this document.

3. Overview

   To evaluate the methods and models for policy based routing, it is
   necessary to investigate the context into which the model is to be
   used, as there are a variety of different methods to introduce
   policies. Most frequently the following three models are referenced:

       Policy based distribution of routing information
       Policy based packet filtering/forwarding
       Policy based dynamic allocation of network resources (e.g.,
       bandwidth, buffers, etc.)

   The relative properties of those methods need to be evaluated to find
   their merits for a specific application. In some cases, more than
   one method needs to be implemented.

   While comparing different models for policy based routing, it is
   important to realize that specific models have been designed to
   satisfy a certain set of requirements. For different models these
   requirements may or may not overlap. Even if they overlap, they may
   have a different degree of granularity. In the first model, the
   requirements can be formulated at the Administrative Domain or
   network number level. In the second model, the requirements can be
   formulated at the end system level or probably even at the level of

individual users. In the third model, the requirements need to be
formulated at both the end system and local router level, as well as
at the level of Routing Domains and Administrative Domains.

Each of these models looks at the power of policy based routing in a
different way. They may be implemented separately or in combination
with other methods. The model to describe policy based dynamic
allocation of network resources is orthogonal to the model of policy
based distribution of routing information. However, in an actual
implementation each of these models may interact.

It is important to realize that the use of a policy based scheme for
individual network applications requires that the actual effects as
well as the interaction of multiple methods need to be determined
ahead of time by policy.

While uncontrolled dynamic routing and allocation of resources may
have a better real time behavior, the use of policy based routing
will provide a predictable, stable result based on the desires of the
administrator. In a production network, it is imperative to provide
continuously consistent and acceptable services.

4. Policy based distribution of routing information

   Goals:

      The goal of this model is to enforce certain flows by means of
      policy based distribution of routing information. This
      enforcement allows control over who can and who can not use
      specific network resources.

      Enforcement is done at the network or Administrative Domain (AD)
      level - macroscopic policies.

   Description:

      A good example of policy based routing based on the distribution
      of routing information is the NSFNET with its interfaces to mid-
      level networks [1], [2]. At the interface into the NSFNET, the
      routing information is authenticated and controlled by four means:

         1. Routing peer authentication based on the source address.

         2. Verification of the Administrative Domain identification
            (currently EGP Autonomous System numbers).

         3. Verification of Internet network numbers which are
            advertised via the routing peer.

   4. Control of metrics via a Routing Policy Data Base for the
      announced Internet network numbers to allow for primary
      paths to the NSFNET as well as for paths of a lesser
      degree.

At the interfaces that pass routing traffic out of the NSFNET, the
NSS routing code authenticates the router acting as an EGP peer by
its address as well as the Administrative Domain identification
(Autonomous System Number).

Outbound announcements of network numbers via the EGP protocol are
controlled on the basis of Administrative Domains or individual
network numbers by the NSFNET Routing Policy Data Base.

The NSFNET routing policy implementation has been in place since
July 1988 and the NSFNET community has significant experience with
its application.

Another example of policy controlled dissimination of routing
information is a method proposed for ESNET in [3].

Benefits:

A major merit of the control of routing information flow is that
it enables the engineering of large wide area networks and allows
for a more meshed environment than would be possible without tight
control. Resource allocation in a non-hostile environment is
possible by filtering specific network numbers or Administrative
Domains on a per need basis. Another important benefit of this
scheme is that it allows for network policy control with virtually
no performance degradation, as only the routing packets themselves
are relevant for policy control. Routing tables are generated as
a result of these interactions. This means that this scheme
imposes only very little impact on packet switching performance at
large.

Concerns:

Policy based routing information distribution does not address
packet based filtering. An example is the inability to prevent
malicious attacks by introduced source routed IP packets. While
resource allocation is possible, it extends largely to filtering
on network numbers or whole Administrative Domains, but it would
not extend to end systems or individual users.

Costs:

Policy based routing in the NSFNET is implemented in a series of

configuration files. These configuration files are in turn
generated from a routing information database. The careful
creation of this routing information database requires knowledge
of the Internet at large. Because the Internet is changing
constantly, the upkeep of this routing information database is a
continuous requirement. However, the effort of collecting and
maintaining an accurate view of the Internet at large can be
distributed.

Since policy controlled distribution of routing information allows
for filtering on the basis of network numbers or Administrative
Domains, the routing information database only needs to collect
information for the more than 1300 networks within the Internet
today.

5. Policy based packet filtering/forwarding

   Goals:

      The goal of the model of policy based packet filtering/forwarding
      is to allow the enforcement of certain flows of network traffic on
      a per packet basis. This enforcement allows the network
      administrator to control who can and who can not use specific
      network resources.

      Enforcement may be done at the end system or even individual user
      level - microscopic policies.

   Description:

      An example of packet/flow based policies is outlined in [4]. In a
      generic sense, policy based packet filtering/forwarding allows
      very tight control of the distribution of packet traffic. An
      implemented example of policy based filtering/forwarding is a
      protection mechanism built into the NSFNET NSS structure, whereby
      the nodes can protect themselves against packets targeted at the
      NSFNET itself by filtering according to IP destination. While this
      feature has so far not been enabled, it is fully implemented and
      can be turned on within a matter of seconds.

   Benefits:

      The principal merit of this scheme is that it allows the
      enforcement of packet policies and resource allocation down to
      individual end systems and perhaps even individual end users. It
      does not address a sane distribution of routing information. If
      policies are contained in the packets themselves it could identify
      users, resulting in the ability of users to move between

locations.

Concerns:

The major concern would be the potentially significant impact on
the performance of the routers, as, at least for tight policy
enforcements, each packet to be forwarded would need to be
verified against a policy data base. This limitation makes the
application of this scheme questionable using current Internet
technology, but it may be very applicable to circuit switched
environments (with source-routed IP packets being similar to a
circuit switched environment). Another difficulty could be the
sheer number of potential policies to be enforced, which could
result in a very high administrative effort. This could result
from the creation of policies at the per-user level. Furthermore,
the overhead of carrying policy information in potentially every
packet could result in additional burdens on resource
availabilities. This again is more applicable to connection-
oriented networks, such as public data networks, where the policy
would only need to be verified at the call setup time. It is an
open question how well packet based policies will scale in a large
and non homogeneous Internet environment, where policies may be
created by all of the participants. These creations of policy
types of services may have to be doable in real time.

Scaling may require hierarchy. Hierarchy may conflict with
arbitrary Type of Service (TOS) routing, which is one of the
benefits of this model.

Costs of implementation:

A large scale implemention of packet based policy routing would
require a routing information base that would contain information
down to the end system level and possibly end users. If one would
assume that for each of the 1300 networks there is an average of
200 end systems, this would result in over 260000 end systems
Internet wide. Each end system in turn could further contribute
some information on the type of traffic desired, including types
of service (issues like agency network selection), potentially on
a per-user basis. The effort for the routing policy data base
could be immense, in particular if there is a scaling requirement
towards a variety of policies for backbones, mid-level networks,
campus networks, subnets, hosts, and users. The administration of
this "packet routing" database could be distributed. However,
with a fully distributed database of this size several consistency
checks would have to be built into the system.

6. Policy based dynamic allocation of network resources (e.g.,
    bandwidth, buffers, etc.).

   Goals:

      Flexible and economical allocation of network resources based on
      current needs and certain policies. Policies may be formulated at
      the network or Administrative Domain (AD) levels. It is also
      possible to formulate policies which will regulate resource
      allocation for different types of traffic (e.g., Telnet, FTP,
      precedence indicators, network control traffic).

      Enforcement of policy based allocation of network resources might
      be implemented within the following parts of the network:

           routers for networks and Administrative Domain (AD) levels
           circuit switches for networks
           end systems establishing network connections

   Description:

      Policy based allocation of bandwidth could allow the modulation of
      the circuits of the networking infrastructure according to real
      time needs. Assuming that available resources are limited towards
      an upper bound, the allocation of bandwidth would need to be
      controlled by policy. One example might be a single end system
      that may or may not be allowed to, perhaps even automatically,
      take resources away from other end systems or users. An example
      of dynamic bandwidth allocation is the currently implemented
      circuit switched IDNX component of the NSFNET, as well as the MCI
      Digital Reconfiguration Service (DRS) which is planned for the
      NSFNET later this year.

      Another model for resource allocation occurs at the packet level,
      where the allocation is controlled by multiple packet queues.
      This could allow for precedence queuing, with preferences based on
      some type of service and preferred forwarding of recognized
      critical data, such as network monitoring, control and routing.
      An example can be found in the NSFNET, where the NSFNET nodes
      prefer traffic affiliated with the NSFNET backbone network number
      over all other traffic, to allow for predictable passing of
      routing information as well as effective network monitoring and
      control. At the other end of the spectrum, an implementation
      could also allow for queues of most deferrable traffic (such as
      large background file transfers).

Benefits:

   Dynamic allocation of bandwidth could allow for a truly flexible
   environment where the networking infrastructure could create
   bandwidth on a per need basis. This could result in significant
   cost reductions during times when little bandwidth is needed.
   This method could potentially accommodate real time transient high
   bandwidth requirements, potentially by reducing the bandwidth
   available to other parts of the infrastructure. A positive aspect
   is that the bandwidth allocation could be protocol independent,
   with no impact on routing protocols or packet forwarding
   performance.

   Policy based allocation of bandwidth can provide a predictable
   dynamic environment. The rules about allocation of bandwidth at
   the circuit level or at the packet level need to be determined by
   a consistent and predictable policy, so that other networks or
   Administrative Domains can tune their allocation of networking
   resources at the same time.

Concerns:

   The policies involved in making dynamic bandwidth allocation in a
   largely packet switching environment possible are still in the
   development phase. Even the technical implications of
   infrastructure reconfiguration in result of events happening on a
   higher level still requires additional research.

   A policy based allocation of bandwidth could tune the network to
   good performance, but could cause networks located in other
   Administrative Domains to pass traffic poorly. It is important
   that network resource policy information for a network be
   discussed within the context of its Administrative Domain.
   Administrative Domains need to discuss their network resource
   allocation policies with other Administrative Domains.

   The technical problem of sharing network resource policy
   information could be solved by a making a "network resource policy
   information" database available to all administrators of networks
   and Administrative Domains. However, the political problems
   involved in creating a network resource policy with impact on
   multiple Administrative Domains does still require additional
   study.

7. Discussion

   Both the first and the second model of policy based routing are
   similar in the sense that their goal is to enforce certain flows.

This enforcement allows the control of access to scarce network
resources (if the resource is not scarce, there is no performance
reason to control access to it). The major difference is the level
of enforcement: macroscopic level versus microscopic level control.

Associated with the enforcement for a certain network resource is the
cost. If this cost is higher than the cost required to make a
particular resource less scarce, then the feasibility of enforcement
may be questionable.

If portions of the Internet find that microscopic enforcement of
policy is necessary, then this will need to be implementable without
significant performance degradation to the networking environment at
large. Local policies within specific Routing Domains or
Administrative Domains should not affect global Internet traffic or
routing. Policies within Administrative Domains which act as traffic
transit systems (such as the NSFNET) should not be affected by
policies a single network imposes for its local benefit.

Some models of policy routing are trying to deal with cases where
network resources require rather complex usage policies. One of
scenarios in [4] is one in which a specific agency may have some
network resource (in the example it is a link) which is sometimes
underutilized. The goal is to sell this resource to other agencies
during the underutilization period to recover expenses. This
situation is equivalent to the problem of finding optimum routes,
with respect to a certain TOS, in the presence of network resources
(e.g., links) with variable characteristics. Any proposed solution
to this problem should address such issues as network and route
stability. More feasibility study is necessary for the whole
approach where links used for global communication are also subject
to arbitrary local policies. An alternative approach would be to
reconfigure the network topology so that underutilized links will be
dropped and possibly returned to the phone company. This is
comparable to what the NSFNET is planning on doing with the MCI
Digital Reconfiguration Service (DRS). A DRS model may appear
cleaner and more easy to implement than a complicated model like the
one outlined in [4].

The models for policy based routing emphasize that careful
engineering of the Internet needs to decided upon the profile of
traffic during normal times, outage periods, and peak loads. This
type of engineering is not a new requirement. However, there could
potentially be a significant benefit in deciding these policies ahead
of time and using policy based routing to implement specific routing
policies.

8. Accounting vs. Policy Based Routing

   Quite often Accounting and Policy Based Routing are discussed
   together. While the application of both Accounting and Policy Based
   Routing is to control access to scarce network resources, these are
   separate (but related) issues.

   The chief difference between Accounting and Policy Based Routing is
   that Accounting combines history information with policy information
   to track network usage for various purposes. Accounting information
   may in turn drive policy mechanisms (for instance, one could imagine
   a policy limiting a certain organization to a fixed aggregate
   percentage of dynamically shared bandwidth). Conversely, policy
   information may affect accounting issues. Network accounting
   typically involves route information (at any level from AD to end
   system) and volume information (packet, octet counts).

   Accounting may be implemented in conjunction with any of the policy
   models mentioned above. Similar to the microscopic versus
   macroscopic policies, accounting may be classified into different
   levels. One may collect accounting data at the AD level, network
   level, host level, or even at the individual user level. However,
   since accounting may be organized hierarchically, microscopic
   accounting may be supported at the network or host level, while
   macroscopic accounting may be supported at the network or AD level.
   An example might be the amount of traffic passed at the interface
   between the NSFNET and a mid-level network or between a mid-level
   network and a campus. Furthermore, the NSFNET has facilities
   implemented to allow for accounting of traffic trends from individual
   network numbers as well as application-specific information.

   Full-blown accounting schemes suffer the same types of concerns
   previously discussed, with the added complication of potentially
   large amounts of additional data gathered that must be reliably
   retrieved. As pointed out in [4], policy issues may impact the way
   accounting data is collected (one administration billing for packets
   that were then dropped in the network of another administration).
   Microscopic accounting may not scale well in a large internet.

   Furthermore, from the standpoint of billing, it is not clear that the
   services provided at the network layer map well to the sorts of
   services that network consumers are willing to pay for. In the
   telephone network (as well as public data networks), users pay for
   end-to-end service and expect good quality service in terms of error
   rate and delay (and may be unwilling to pay for service that is
   viewed as unacceptable). In an internetworking environment, the
   heterogeneous administrative environment combined with the lack of
   end-to-end control may make this approach infeasible.

Lightweight approaches to accounting can be used (with less impact)
when specific, limited goals are set. One suggested approach
involves monitoring traffic patterns. If a pattern of abuse (e.g.,
unauthorized use) develops, an accounting system could track this and
allow corrective action to be taken, by changing routing policy or
imposing access control (blocking hosts or nets). Note that this is
much less intrusive into the packet forwarding aspects of the
routers, but requires distribution of a policy database that the
accounting system can use to reduce the raw information. Because
this approach is statistical in nature, it may be slow to react.

9. References

   [1] Rekhter, Y., "EGP and Policy Based Routing in the New NSFNET
       Backbone", RFC 1092, IBM Research, February 1989.

   [2] Braun, H-W., "The NSFNET Routing Architecture", RFC 1093,
       Merit/NSFNET Project, February 1989.

   [3] Collins, M., and R. Nitzan, "ESNET Routing", DRAFT Version 1.0,
       LLNL, May 1989.

   [4] Clark, D., "Policy Routing in Internet Protocols", RFC 1102,
       M.I.T. Laboratory for Computer Science, May 1989.

Author's Address

   Hans-Werner Braun
   Merit Computer Network
   University of Michigan
   1075 Beal Avenue
   Ann Arbor, Michigan 48109

   Telephone: 313 763-4897
   Fax: 313 747-3745
   EMail: hwb@merit.edu