

Goals and Functional Requirements for Inter-Autonomous System Routing

Status of this Memo

This document describes the functional requirements for a routing protocol to be used between autonomous systems. This document is intended as a necessary precursor to the design of a new inter-autonomous system routing protocol and specifies requirements for the Internet applicable for use with the current DoD IP, the ISO IP, and future Internet Protocols. It is intended that these requirements will form the basis for the future development of a new inter-autonomous systems routing architecture and protocol. This document is being circulated to the IETF and Internet community for comment. Comments should be sent to: "open-rout-editor@bbn.com". This memo does not specify a standard. Distribution of this memo is unlimited.

1. Introduction

The development of an inter-autonomous systems routing protocol proceeds from those goals and functions seen as both desirable and obtainable for the Internet environment. This document describes these goals and functional requirements. The goals and functional requirements addressed by this document are intended to provide a context within which an inter-autonomous system routing architecture can be developed which will meet both current and future Internet routing needs. The goals presented indicate properties and general capabilities desired of the Internet routing environment and what the inter-autonomous system routing architecture is to accomplish as a whole.

The goals are followed by functional requirements, which address either detailed objectives or specific functionality to be achieved by the architecture and resulting protocol(s). These functional requirements are enumerated for clarity and grouped so as to map directly to areas of architectural consideration. This is followed by a listing and description of general objectives, such as robustness, which are applicable in a broad sense. Specific functions which are not reasonably attainable or best left to future efforts are identified as non-requirements.

The intent of this document is to provide both the goals and functional requirements in a concise fashion. Supporting arguments,

tradeoff considerations and the like have been purposefully omitted in support of this. An appendix has been included which addresses this omission to a limited extent and the reader is directed there for a more detailed discussion of the issues involved.

The goals and functional requirements contained in this document are the result of work done by the members of the Open Routing Working Group. It is our intention that these goals and requirements reflect not only those foreseen in the Internet community but are also similar to those encountered in environments proposed by ANSI, ECMA and ISO. It is expected that there will be some interaction and relationship between this work and the product of these groups.

2. Overall Goals

In order to derive a set functional requirements there must be one or more principals or overall goals for the routing environment to satisfy. These high level goals provide the basis for each of the functional requirements we have derived and will guide the design philosophy for achieving an inter-autonomous system routing solution. The overall goals we are utilizing are described in the following sections.

2.1 Route to Destination

The routing architecture will provide for the routing of datagrams from a single source to one or more destinations in a timely manner. The larger goal is to provide datagram delivery to an identifiable destination, one which is not necessarily immediately reachable by the source. In particular, routing is to address the needs of a single source requiring datagram delivery to one or more destinations. The concepts of multi-homed hosts and multicasting routing services are encompassed by this goal. Datagram delivery is to be provided to all interconnected systems when not otherwise constrained by autonomous considerations.

2.2 Routing is Assured

Routing services are to be provided with assurance, where the inability to provide a service is communicated under best effort to the requester within an acceptable level of error. This assurance is not to be misconstrued to mean guaranteed datagram delivery nor does it imply error notification for every lost datagram. Instead, attempts to utilize network routing services when such service cannot be provided will result in requester notification within a reasonable period given persistent attempts.

2.3 Large System

The design of the architecture, and the protocols within this architecture, should accommodate a large number of routing entities. The exact order of magnitude is a relative guess and the best designs would provide for a practical level of unbounded growth. Nevertheless, the routing architecture is expected to accommodate the growth of the Internet environment for the next 10 years.

2.4 Autonomous Operation

The routing architecture is to allow for stable operation when significant portions of the internetworking environment are controlled by disjoint entities. The future Internet environment is envisioned as consisting of a large number of internetworking facilities owned and operated by a variety of funding sources and administrative concerns. Although cooperation between these facilities is necessary to provide interconnectivity, it is viewed that both the degree and type of cooperation will vary widely. Additionally, each of these internetworking facilities desires to operate as independently as possible from the concerns and activities of other facilities individually and the interconnection environment as a whole. Those resources used by (and available for) routing are to be allowed autonomous control by those administrative entities which own or operate them. Specifically, each controlling administration should be allowed to establish and maintain policies regarding the use of a given routing resource.

2.5 Distributed System

The routing environment developed should not depend upon a data repository or topological entity which is either centralized or ubiquitous. The growth pattern of the Internet, coupled with the need for autonomous operation, dictates an independence from the topological and administrative centralization of both data and control flows. Past experience with a centralized topology has shown that it is both impractical for the needs of the community and restrictive of administrative freedoms. A distributed routing environment should not be restrictive of either redundancy or diversity. Any new routing environment must allow for arbitrary interconnection between internetworks.

2.6 Provide A Credible Environment

The routing environment and services should be based upon mechanisms and information that exhibit both integrity and security. The routing mechanisms should operate in a sound and reliable fashion while the routing information base should provide credible data upon

which to base routing decisions. The environment can be unreliable to the extent that the resulting effect on routing services is negligible. The architecture and protocol designs should be such that the routing environment is reasonably secure from unwanted modification or influence.

2.7 Be A Managed Entity

Provide a manager insight into the operation of the inter-autonomous system routing environment to support resource management, problem solving, and fault isolation. Allow for management control of the routing system and collect useful information for the internetwork management environment. Datagram events as well as the content and distribution characteristics of relevant databases are of particular importance.

2.8 Minimize Required Resources

Any feasible design should restrain the demand for resources required to provide inter-autonomous systems routing. Of particular interest are those resources required for data storage, transmission, and processing. The design must be practical in terms of today's technology. Specifically, do not assume significant upgrades to the existing level of technology in use today for data communication systems.

3. Functional Requirements

The functional requirements we have identified have been derived from the overall goals and describe the critical features expected of inter-autonomous system routing. To an extent, these functions are vague in terms of detail. We do not, for instance, specify the quantity or types for quality-of-service parameters. This is purposeful, as the functional requirements specified here are intended to define the features required of the inter-autonomous system routing environment rather than the exact nature of this environment. The functional requirements identified have been loosely grouped according to areas of architectural impact.

3.1 Route Synthesis Requirements

Route synthesis is that functional area concerned with both route selection and path determination (identification of a sequence of intermediate systems) from a source to a destination. The functional requirements identified here provide for path determination which is adaptive to topology changes, responsive to administrative policy, cognizant of quality-of-service concerns, and sensitive to an interconnected environment of autonomously managed systems.

a) Route around failures dynamically

Route synthesis will provide a best effort attempt to detect failures in those routing resources which are currently being utilized. Upon detection of a failed resource, route synthesis will provide a best effort to utilize other available routing resources in an attempt to provide the necessary routing service.

b) Provide loop free paths

The path provided for a datagram, from source to destination, will be free of circuits or loops most of the time. At those times a circuit or loop exists, it occurs with both negligible probability and duration.

c) Know when a path or destination is unavailable

Route synthesis will be capable of determining when a path cannot be constructed to reach a known destination. Additionally, route synthesis will be capable of determining when a given destination cannot be determined because the requested destination is unknown (or this knowledge is unavailable).

d) Provide paths sensitive to administrative policies

Route synthesis will accommodate the resource utilization policies of those administrative entities which manage the resources identified by the resulting path. However, it is inconceivable to accommodate all policies which can be defined by a managing administrative entity. Specifically, policies dependent upon volatile events of great celerity or those which are non-deterministic in nature cannot be accommodated.

e) Provide paths sensitive to user policies

Paths produced by route synthesis must be sensitive to policies expressed by the user. These user policies are expressed in terms relevant to known characteristics of the topology. The path achieved will meet the requirements stated by the user policy.

f) Provide paths which characterize user quality-of-service requirements

The characteristics of the path provided should match those indicated by the quality-of-service requested. When

appropriate, utilize only those resources which can support the desired quality-of-service (e.g., bandwidth).

- g) Provide autonomy between inter- and intra-autonomous system route synthesis

The inter- and intra-autonomous system routing environments should operate independent of one another. The architecture and design should be such that route synthesis of either routing environment does not depend upon information from the other for successful functioning. Specifically, the inter-autonomous system route synthesis design should minimize the constraints on the intra-autonomous system route synthesis decisions when transiting (or delivering to) the autonomous system.

3.2 Forwarding Requirements

The following requirements specifically address the functionality of the datagram forwarding process. The forwarding process transfers datagrams to intermediate or final destinations based upon datagram characteristics, environmental characteristics, and route synthesis decisions.

- a) Decouple inter- and intra-autonomous system forwarding decisions

The requirement is to provide a degree of independence between the inter-autonomous system forwarding decision and the intra-autonomous system forwarding decision within the forwarding process. Though the forwarding decisions are to be independent of each other, the inter-autonomous system delivery process may necessarily be dependent upon intra-autonomous system route synthesis and forwarding.

- b) Do not forward datagrams deemed administratively inappropriate

Forward datagrams according to the route synthesis decision if it does not conflict with known policy. Policy sensitive route synthesis will prevent normally routed datagrams from utilizing inappropriate resources. However, a datagram routed abnormally due to unknown events or actions can always occur and the only way to prohibit unwanted traffic from entering or leaving an autonomous system is to provide policy enforcement within the forwarding function.

c) Do not forward datagrams to failed resources

A datagram is not to be forwarded to a resource known to be unavailable, notably an intermediate system such as a gateway. This implies some ability to detect and react to resource failures.

d) Forward datagram according to its characteristics

The datagram forwarding function is to be sensitive to the characteristics of the datagram in order to execute the appropriate route synthesis decision. Characteristics to consider are the destination, quality-of-service, precedence, datagram (or user) policy, and security. Note that some characteristics, precedence for example, affect the forwarding service provided whereas others affect the path chosen.

3.3 Information Requirements

This functional area addresses the general information requirements of the routing environment. This encompasses both the nature and disbursement of routing information. The characteristics of the routing information and its disbursement are given by the following functional requirements.

a) Provide a distributed and descriptive information base

The information base must not depend upon either centralization or exact replication. The content of the information base must be sufficient to support all provided routing functionality, specifically that of route synthesis and forwarding. Information of particular importance includes resource characteristics and resource utilization policies.

b) Determine resource availability

Provide a means of determining the availability of any utilized resource in a timely manner. The timeliness of this determination is dependent upon the routing service(s) to be supported.

c) Restrain transmission utilization

The dynamics of routing information flow should be such that a significant portion of transmission resources are not consumed. Routing information flow should adjust to the demands of the environment, the capacities of the distribution facilities utilized, and the desires of the resource manager.

d) Allow limited information exchange

Information distribution is to be sensitive to administrative policies. An administrative policy may affect the content or completeness of the information distributed. Additionally, administrative policy may determine the extent of information distributed.

3.4 Environmental Requirements

The following items identify those requirements directly related to the nature of the environment within which routing is to occur.

a) Support a packet-switching environment

The routing environment should be capable of supporting datagram transfer within a packet-switched oriented networking environment.

b) Accommodate a connection-less oriented user transport service

The routing environment should be designed such that it accommodates the model for connection-less oriented user transport service.

c) Accommodate 10K autonomous systems and 100K networks

This requirement identifies the scale of the internetwork environment we view as appearing in the future. A routing design which does not accommodate this order of magnitude is viewed as being inappropriate.

d) Allow for arbitrary interconnection of autonomous systems

The routing environment should accommodate interconnectivity between autonomous systems which may occur in an arbitrary manner. It is recognized that a practical solution is likely to favor a given structure of interconnectivity for reasons of efficiency. However, a design which does not allow for and utilize interconnectivity of an arbitrary nature would not be considered a feasible design.

3.5 General Objectives

The following are overall objectives to be achieved by the inter-autonomous routing architecture and its protocols.

a) Provide routing services in a timely manner

Those routing services provided, encapsulated by the requirements stated herein, are to be provided in a timely manner. The time scale for this provision must be reasonable to support those services provided by the internetwork environment as a whole.

b) Minimize constraints on systems with limited resources

Allow autonomous systems, or gateways, of limited resources to participate in the inter-autonomous system routing architecture. This limited participation is not necessarily without cost, either in terms of responsiveness, path optimization, or other factor(s).

c) Minimize impact of dissimilarities between autonomous systems

Attempt to achieve a design in which the dissimilarities between autonomous systems do not impinge upon the routing services provided to any autonomous system.

d) Accommodate the addressing schemes and protocol mechanisms of the autonomous systems

The routing environment should accommodate the addressing schemes and protocol mechanisms of autonomous systems, where these schemes and mechanisms may differ among autonomous systems.

e) Must be implementable by network vendors

This is to say that the algorithms and complexities of the design must be such that they can be understood outside of the research community and implementable by people other than the designers themselves. Any feasible design must be capable of being put into practice.

4. Non-Goals

In view of the conflicting nature of many of the stated goals and the careful considerations and tradeoffs necessary to achieve a successful design, it is important to also identify those goals or functions which we are not attempting to achieve. The following items are not considered to be reasonable goals or functional requirements at this time and are best left to future efforts. These are non-goals, or non-requirements, within the context of the goals and requirements previously stated by this document as well as our interpretation of what can be practically achieved.

a) Ubiquity

It is not a goal to design a routing environment in which any participating autonomous system can obtain a routing service to any other participating autonomous system in a ubiquitous fashion. Within a policy sensitive routing environment, the cooperation of intermediate resources will be necessary and cannot be guaranteed to all participants. The concept of ubiquitous connectivity will not be a valid one.

b) Congestion control

The ability for inter-autonomous system routing to perform congestion control is a non-requirement. Additional study is necessary to determine what mechanisms are most appropriate and if congestion control is best realized within the inter-AS and/or intra-AS environments, and if both, what the dynamics of the interactions between the two are.

c) Load splitting

The functional capability to distribute the flow of datagrams, from a source to a destination, across two or more distinct paths through route synthesis and/or forwarding is a non-requirement.

d) Maximizing the utilization of resources

There is no goal or requirement for the inter-autonomous system routing environment to be designed such that it attempts to maximize the utilization of available resources.

e) Schedule to deadline service

The ability to support a schedule to deadline routing service is a non-requirement for the inter-autonomous routing environment at this point in time.

f) Non-interference policies of resource utilization

The ability to support routing policies based upon the concept of non-interference is a not a requirement. An example of such a policy is one where an autonomous system allows the utilization of excess bandwidth by external users as long as this does not interfere with local usage of the link.

5. Considerations

Although neither a specific goal nor a functional requirement, consideration must be given to the transition which will occur from the current operational routing environment to a new routing environment. A coordinated effort among all participants of the Internet would be impractical considering the magnitude of such an undertaking. Particularly, the issues of transitional coexistence, as opposed to phased upgrading between disjoint systems, should be addressed as a means to minimize the disruption of service. Careful consideration should also be given to any required changes to hosts. It is very unlikely that all hosts could be changed, given historical precedence, their diversity and their large numbers.

Appendix - Issues in Inter-Autonomous Systems Routing

A.0 Acknowledgement

This appendix is an edited version of the now defunct document entitled "Requirements for Inter-Autonomous Systems Routing", written by Ross Callon in conjunction with the members of the Open Routing Working Group.

A.1 Introduction

The information and discussion contained here historically precedes that of the main document body and was a major influence on its content. It is included here as a matter of reference and to provide insight into some of the many issues involved in inter-autonomous systems routing.

The following definitions are utilized:

Boundary Gateway

A boundary gateway is any autonomous system gateway which has a network interface directly reachable from another autonomous system. As a member of an autonomous system, a boundary gateway participates in the Interior Gateway Protocol and other protocols used for routing (and other purposes) between other gateways of this same autonomous system and between those networks directly reachable by this autonomous system. A boundary gateway may also participate in an Inter-Autonomous System Routing Protocol. As a participant in the inter-autonomous system routing protocol, a boundary gateway interacts with other boundary gateways in other autonomous systems, either directly or indirectly, in support of the operation of the

Inter-Autonomous System Routing Protocol.

Interior Gateway

An interior gateway is any autonomous system gateway which is not a boundary gateway. As such, an interior gateway does not have any network interfaces which are directly reachable by any other autonomous system. An interior gateway is part of an autonomous system and, as such, takes part in the Interior Gateway Protocol and other protocols used in that autonomous system. However, an interior gateway does not directly exchange routing information with gateways in other autonomous systems via the Inter-Autonomous System Routing Protocol.

The following acronyms are used:

AS -- Autonomous System

This document uses the current definition of "Autonomous System": a collection of cooperating gateways running a common interior routing protocol. This implies that networks and hosts may be reachable through one or more Autonomous Systems.

NOTE: The current notion of "Autonomous System" implicitly assumes that each gateway will belong to exactly one AS. Extensions to allow gateways which belong to no AS's and/or gateways which belong to multiple AS's, are beyond the scope of this discussion. However, we do not preclude the possibility of considering such extensions in the future.

IARP -- Inter-Autonomous System Routing Protocol

This is the protocol used between boundary gateways for the purpose of routing between autonomous systems.

IGP -- Interior Gateway Protocol

This is the protocol used within an autonomous system for routing within that autonomous system.

A.2 Architectural Issues

The architecture of an inter-autonomous system routing environment is mutually dependent with the notion of an Autonomous System. In general, the architecture should maximize independence of the

internals of an AS from the internals of other AS's, as well as from the inter-autonomous system routing protocols (IARP). This independence should allow technological and administrative differences among AS's as well as protection against propagation of misbehavior. The following issues address ways to achieve interoperation and protection, and to meet certain performance criteria. We also put forth a set of minimal constraints to be imposed among Autonomous Systems, and between inter- and intra-AS functions.

A.2.1 IGP Behavior

The IARP should be capable of tolerating an Autonomous System in which its IGP is unable to route packets, provides incorrect information, and exhibits unstable behavior. Interfacing to such an ill-behaved AS should not produce global instabilities within the IARP and the IARP should localize any effects. On the other hand, the IGP should provide a routing environment where the information and connectivity provided to the IARP from the IGP does not exhibit rapid and continual changes. An Autonomous System therefore should appear as a relatively stable environment.

A.2.2 Independence of Autonomous Systems

The IARP should not constrain any AS to require the use any one specific IGP. This applies both to IGPs and potentially to any other internal protocols. The architecture should also allow intra-AS routing and organizational structures to be hidden from inter-AS use. An Autonomous System should not be required to use any one specific type of linkage between boundary gateways within the AS. However, there are some minimal constraints that gateways and the associated interior routing protocol within an AS must meet in order to be able to route Inter-AS traffic, as discussed in Section A.2.6.

A.2.3 General Topology

The routing architecture should provide significant flexibility regarding the interconnection of AS's. The specification of IARP should impose no inherent restriction on either interconnection configuration or information passing among autonomous systems. There may be administrative and policy limitations on the interconnection of AS's, and on the extent to which routing information and data traffic may be passed between AS's. However, there should be no inherent restrictions imposed by limitations in the design of the routing architecture. The architecture should allow arbitrary topological interconnection of Autonomous Systems. Propagation of routing information should not be restricted by the specification of the IARP. For example, the restrictions imposed by the "core model"

used by EGP are not acceptable.

A.2.4 Routing Firewalls

We expect AS's to have a certain amount of insulation from other AS's. This protection should apply to both the adequacy and stability of routes produced by the routing scheme, and also to the amount of overhead traffic and other costs necessary to run the routing scheme. There are several forms which these "routing firewalls" may take:

- An AS must be able to successfully route its own internal traffic in the face of arbitrary failures of other IGPs and the IARP. In other words, the AS should be able to effectively shutout the rest of the world.
- The IARP should be able to operate correctly in the face of IGP failures. In this case, correct operation is defined as recognizing that an AS has failed, and routing around it if possible (traffic to or from that AS may of course fail).
- In addition, problems in Inter-AS Routing should, as much as possible, be limited in the extent of their effect.

Routing firewalls may be explicit, or may be inherent in the design of the algorithms. We expect that both explicit and inherent firewalls will be utilized. Examples of firewalls include:

- Separating Intra- and Inter-AS Routing to some extent isolates each of these from problems with the other. Clearly defined interfaces between different modules/protocols provides some degree of protection.
- Access control restrictions may provide some degree of firewalls. For example, some AS's may be non-transit (won't forward transit traffic). Failures within such AS's may be prevented from affecting traffic not associated with that AS.
- Protocol design can help. For example, with link state routing you can require that both ends must report a link before it may be regarded as up, thereby eliminating the possibility of a single node causing fictitious links.
- Finally, explicit firewalls may be employed using explicit configuration information.

A.2.5 Boundary Gateways

Boundary gateways will exchange Inter-AS Routing information with other boundary gateways using the IARP. Each AS which is to take part in Inter-AS Routing will have one or more boundary gateways, of which one or more of these boundary gateways exchanges information with peer boundary gateways in other AS's.

Information related to Inter-AS Routing may be passed between connected boundary gateways in different AS's. Specific designated boundary gateways will therefore be required to understand the IARP. The external link between the boundary gateways may be accomplished by any kind of connectivity that can be modeled as a direct link between two gateways -- a LAN, an ARPANET, a satellite link, a dedicated line, and so on.

A.2.6 Minimal Constraints on the Autonomous System

The architectural issues discussed here for inter-AS routing imply certain minimal functional constraints that an AS must satisfy in order to take part in the Inter-AS Routing scheme. These minimal requirements are described in greater detail in this section. This list of functional constraints is not necessarily complete.

A.2.6.1 Internal Links between Boundary Gateways

In those cases where an AS may act as a transit AS (i.e., may pass traffic for which neither the source nor the destination is in that AS), the gateways internal to that AS will need to know which boundary gateway is to serve as the exit gateway from that AS. There are several ways in which this may be accomplished:

1. Boundary gateways are directly connected
2. "Tunneling" (i) using source routing (ii) using encapsulation
3. Interior gateways participate (i) limited participation (ii) fully general participation

With solution (1), the boundary gateways in an AS are directly connected. This eliminates the need for other gateways in the AS to have any knowledge of Inter-AS Routing. Transit traffic is passed directly among the boundary gateways of the AS.

With solution (2), transit traffic may traverse interior gateways, but these interior gateways are protected from any need to have knowledge about Inter-AS routes by means such as source routing or encapsulation. The boundary gateway by which the packet enters an AS

determines the boundary gateway which will serve as the exit gateway. This may require that the entrance boundary gateway add a source route to the packet, or encapsulate the packet in another level of IP or gateway-to-gateway header. This allows boundary gateways to forward data traffic using the appropriate tunnelling technique.

Finally, with solution (3), the interior gateways have some knowledge of Inter-AS Routing. At a minimum, the interior gateways would need to know the identity of each boundary gateway, the address(es) that can be reached by that gateway, and the Inter-AS metric associated with the route to that address(es). If the IARP allows for separate routing for multiple TOS classes, then the information that the interior gateways need to know includes a separate Inter-AS metric for each TOS class. The Inter-AS metrics are necessary to allow gateways to choose among multiple possible exit boundary gateways. In general, it is not necessary for the Inter-AS metrics to have any relationship with the metric used within an AS for interior routing. The interior gateways do not need to know how to interpret the exterior metrics, except to know that each metric is to be interpreted as an unsigned integer and a lesser value is preferable to a greater value. It would be possible, but not necessary, for the interior gateways to have full knowledge of the IARP.

It is not necessary for the Inter-AS Routing architecture to specify which of these solutions are to be used for any particular AS. Rather, it is possible for individual AS's to choose which scheme or combination of schemes to use. Independence of the IARP from the internal operation of each AS implies that this decision be left up to the internal protocols used in each AS. The IARP must be able to operate as if the boundary gateways were directly connected.

A.2.6.2 Forwarding of Data from the AS

The scheme used for forwarding transit traffic across an AS also has implications for the forwarding of traffic which originates within an AS, but whose destination is reachable only from other AS's. If either of solutions (1) or (2) in Section A.2.6.1 is followed, then it will be sufficient for an interior gateway to forward such traffic to any boundary gateway. Greater efficiency may optionally be achieved in some cases by providing interior gateways with additional information which will allow them to choose the "best" boundary gateway in some sense. If solution (3) is followed, then the information passed to interior gateways to allow them to forward transit traffic will also be sufficient to forward traffic originating within that AS.

A.2.6.3 Forwarding of Data to a Destination in the AS

If a packet whose destination is reachable from an AS arrives at that AS, then it is desired that the interior routing protocol used in that AS be able to successfully deliver the packet without reliance on Inter-AS Routing. This does not preclude that the Inter-AS Routing protocol will deal with partitioned AS's.

An AS may have access control, security, and policy restrictions that restrict which data packets may enter or leave the AS. However, for any data packet which is allowed access to the AS, the AS is expected to deliver the packet to its destination without further restrictions between parts of the AS. In this sense, the internal structure of the AS should not be visible to the IARP.

A.3 Policy Issues

The interconnection of multiple heterogeneous networks and multiple heterogeneous autonomous systems owned and operated by multiple administrations into a single combined internet is a very complex task. It is expected that the administrations associated with such an internet will wish to impose a variety of constraints on the operation of the internet. Specifically, externally imposed routing constraints may include a variety of transit access control, administrative policy, and security constraints.

Transit access control refers to those access control restrictions which an AS may impose to restrict which traffic the AS is willing to forward. There are a large number of access control restrictions which one could envision being used. For example, some AS's may wish to operate only as "non-transit" AS's, that is, they will only forward data traffic which either originates or terminates within that AS. Other AS's may restrict transit traffic to datagrams which originate within a specified set of source hosts. Restrictions may require that datagrams be associated with specific applications (such as electronic mail traffic only), or that datagrams be associated with specific classes of users.

Policy restrictions may allow either the source of datagrams, or the organization that is paying for transmission of a datagram, to limit which AS's the datagrams may transit. For example, an organization may wish to specify that certain traffic originating within their AS should not transit any AS administered by its main competitor.

Security restrictions on traffic relates to the official security classification level of traffic. As an example, an AS may specify that all classified traffic is not allowed to leave its AS.

The main problem with producing a routing scheme which is sensitive to transit access control, administrative policy, and security constraints is the associated potential for exponential growth of routes. For example, suppose that there are 20 packets arriving at a particular gateway, each for the same destination, but subject to a different combination of access control, policy, and security constraints. It is possible that all 20 packets would need to follow different routes to the destination.

This explosive growth of routes leads to the question: "Is it practically feasible to deal with complete general external routing constraints?" One approach would allow only a smaller subset of constraints, chosen to provide some useful level of control while allowing minimal impact on the routing protocol. Further work is needed to determine the feasibility of this approach.

There is another problem with dealing with transit access control, policy, and security restrictions in a fully general way. Specifically, it is not clear just what the total set of possible restrictions should be. Efforts to study this issue are currently underway, but are not expected to produce definitive results within a short to medium time frame. It is therefore not practical to wait for this effort to be finished before defining the next generation of Inter-AS Routing.

A.4 Service Features

The following paragraphs discuss issues concerning the services an Inter-AS Routing Protocol may provide. This is not a complete list of service issues but does address many of those services which are of concern to a significant portion of the community.

A.4.1 Cost on Toll Networks

Consideration must be given to the use of routing protocols with toll (i.e., charge) networks. Although uncommon in the Internet at the moment, they will become more common in the future, and thought needs to be given to allowing their inclusion in an efficient and effective manner.

There are two areas in which concerns of cost intrude. First, provision must be made to include in the routing information distributed throughout the system the information that certain links cost money, so that traffic patterns may account for the cost. Second, the actual operation of the algorithm, in terms of the messages that must be exchanged to operate the algorithm, must recognize that fact that on certain links, the exchange may have an associated cost which must be taken into account. These areas often

involve policy questions on the part of the user. It is a requirement of the algorithm that facilities be available to allow different groups to answer these questions in different ways. The first area is related to type-of-service routing, and is discussed in Section A.4.2. The second area is discussed below.

Previous attempts at providing these sorts of controls were incomplete because they were not thought through fully; a new effort must avoid these pitfalls. For instance, even though the Hello rate in EGP may be adjusted, turning the rate down too low (to control the costs) could cause the route to be dropped from databases through timeout.

In a large internet, changes will be occurring constantly; a simplistic mechanism might mean that a site which is only connected by toll networks has to either accept having an old picture of the network, or spend more to keep a more current picture of things. However, that is not necessarily the case if incomplete knowledge and cache-based techniques are used more. For instance, if a site connected only by toll links keeps an incomplete or less up-to-date map of the situation, an agreement with a neighbor which does not labor under these restrictions might allow it to get up-to-date information only when needed.

A.4.2 Type-of-Service Routing

The need for type-of-service (TOS) has been increasing as networks become more heterogeneous in physical channel components, high-level applications, and administrative management. For instance, some recently installed fiber cables provide abundant communication bandwidths, while old narrow-band channels will still be with us for a long time period. Electronic mail traffic tolerates delivery delays and low throughput. New image transmissions are coming up; these require high bandwidths but are not effected by a few bit errors. Furthermore, some networks may soon install accounting functions to charge users, while others may still provide free services.

Considering the long life span of a new routing architecture, it is mandatory that it be built with mechanisms to provide TOS routing. Unfortunately, we have had very little experience with TOS routing, even with a single network. No TOS routing system has ever been field-tested on a large-scale basis.

We foresee the need for TOS routing and recognize the possible complexities and difficulties in design and implementation. We also consider that new applications coming along may require novel services that are unforeseeable today. We feel a practical approach

is to provide a small set of TOS routing functions as a first step while the design of the architecture should be such that new classes of TOS can be easily added later and incrementally deployed. The Inter-AS Routing Architecture should allow both globally and locally defined TOS classes.

We intend to address TOS routing based on the following metrics:

- Delay
- Throughput
- Cost

Delay and throughput are the main network performance concerns. Considering that some networks may soon start charging users for the transmission services provided, the cost should also be added as a factor in route selection.

Reliability is not included in the above list. Different applications with different reliability requirements will differ in terms of what Transport Protocol they use. However, IP offers only a "moderate" level of reliability, suitable to applications such as voice, or possibly even less than that required by voice. The level of reliability offered by IP will not differ substantially based on the application. Thus the requested level of reliability will not affect Inter-AS Routing.

Delay and throughput will be measured from the physical characteristics of communication channels, without considering instantaneous load. This is necessary in order to provide stable routes, and to minimize the overhead caused by the Inter-AS Routing scheme.

A number of TOS service classes may be defined according to these metrics. Each class will present defined requirements for each of the TOS metrics. For example, one class may require low delay, require only low throughput, and require low cost.

A.4.3 Multipath Routing

There are two types of multipath routing which are useful for Inter-AS Routing: (1) there may be multiple gateways between any two neighboring AS's; (2) there may be multiple AS-to-AS paths between any pair of source and destination AS's. Both forms of multipath are useful in order to allow for loadsplitting. Provision for multipath routing in the IARP is desirable, but not an absolute requirement.

A.5 Performance Issues

The following paragraphs discuss issues related to the performance of an Inter-AS Routing Protocol. This discussion addresses size as well as efficiency considerations.

A.5.1 Adaptive Routing

It is necessary that the Inter-AS Routing scheme respond in a timely fashion to major network problems, such as the failure of specific network resources. In this sense, Inter-AS Routing needs to use adaptive routing mechanisms similar to those commonly used within individual networks and AS's. It is important that the adaptive routing mechanism chosen for Inter-AS Routing achieve rapid convergence following internet topological changes, with little or none of the serious convergence problems (such as "counting to infinity") that have been experienced in some existing dynamic routing protocols.

The Inter-AS Routing scheme must provide stability of routes. It is totally unacceptable for routes to vary on a frequent basis. This requirement is not meant to limit the ability of the routing algorithm to react rapidly to major topological changes, such as the loss of connectivity between two AS's. The need for adaptive routing does not imply any desire for load-based routing.

A.5.2 Large Internets

One key question in terms of the targets is the maximum size of the Internet this algorithm is supposed to support. To some degree, this is tied to the timeline for which this protocol is expected to be active. However, it is necessary to have some size targets. Techniques that work at one order of size may be impractical in a system ten or twenty times larger.

Over the past five years there has been an approximate doubling of the Internet each year. In January 1988, there were more than 330 operational networks and more than 700 network assigned numbers. Exact figures for the future growth rate of the Internet are difficult to predict accurately. However, if this doubling trend continues, we would reach 10,000 nets sometime near January 1993.

Taking a projection purely on the number of networks (the top level routing entity) may be overly conservative since the introduction and wide use of subnets has absorbed some growth, but will not continue to be able to do so. In addition, there are plans being discussed that will continue or accelerate the current rate of growth. Nonetheless, the number of networks is very important because

networks constitute the "top level entities" in the current addressing structure.

The implications of the size parameter are fairly serious. The current system has only one level of addressing above subnets. While it is possible to adjust certain parameters (for example, the unsolicited or unnecessary retransmission rate) to produce a larger but less robust system, other parameters (for example, the rate of change in the system) cannot be so adjusted. This will provide eventual limits on the size of a system that can be dealt with in a flat address space.

The exact size that necessitates moving from the current single-level system to a multi-level system is not clear. Among the parameters which affect it are the assumed minimum speed of links in the system (faster links can allocate more bandwidth to routing traffic before it becomes obtrusive), speed and memory capacity of routing nodes (needed to store and process routing data), rate at which topology changes, etc. The maximum size which can be handled in a single layer is generally thought to be somewhere on the order of 10 thousand objects. The IARP must be designed to deal with internets bigger than this.

A.5.3 Addressing Implications

Given the current rate of growth of the Internet, we can expect that the current addressing structure will become unworkable early within the lifetime of the new IARP. It is therefore essential that any new IARP be able to use a new addressing format which allows for addressing hierarchies beyond the network level. Any new IARP should allow for graceful migration from the current routing protocols, and should also allow for graceful migration from a routing scheme based on the current addressing, to a scheme based on a new multi-level addressing format such as that described by OSI 8473.

A.5.4 Memory, CPU, and Bandwidth Costs

Routing costs can be measured in terms of the memory needed to store routing information, the CPU costs of calculating routes and forwarding packets, and the bandwidth costs of exchanging routing information and of forwarding packets. These significant factors should provide the basis for comparison between competing proposals in IARP design.

The routing architecture will be driven by the expected size of the Internet, the expected memory capacity of the gateways, capacity of the Inter-AS links, and the computing speed of the gateways. Given our experience with the current Internet, it is clearly necessary for the scheme to function adequately even if the Internet grows more quickly than we predict and its capacity grows more slowly. Memory, CPU, and bandwidth costs should be in line with what is economically practical at any point in time given the size of the Internet at that time.

A.6 Other Issues

The following are issues of a general nature and includes discussion of items which have been considered to be best left for future efforts.

A.6.1 Implementation

The specification of IARP should allow interoperation among multi-vendor implementations. This requires that multiple vendors be able to implement the same protocol, and that equipment from multiple vendors be able to interoperate successfully.

There are potential practical difficulties of realizing multi-vendor interoperation. Any such difficulty should not be inherent to the protocol specifications. Towards this end, we should produce a protocol specification that is precise and unambiguous. This implies that the specification should include a detailed specification using Pseudo-Code or a Formal Description Technique.

A.6.2 Configuration

It is expected that any IARP will require a certain amount of configuration information to be maintained by gateways. However, in practice it is often difficult to maintain configuration information in a fully correct and up-to-date form. Problems in configuration have been known to cause significant problems in existing operational networks and internets. The design of an Inter-AS Routing architecture must therefore simplify the maintenance of configuration information, consistent with other requirements. Simplification of configuration information may require minimizing the amount of configuration information, and using automated or semi-automated configuration mechanisms.

A.6.3 Migration

In any event, whether the address format changes or not, a viable transition plan which allows for interoperability must be arranged.

In a system of this magnitude, which is in operational use, a coordinated change is not possible. Where possible, changes should not affect the hosts, since deploying such a change is probably several orders of magnitude more difficult than changing only the gateways, due to the larger number of host implementations as well as hosts. There are two important questions that need to be addressed: (1) migration from the existing EGP to a new IARP; (2) migration from the current DD IP to future protocols (including the ISO IP, and other future protocols).

A.6.4 Load-Based Routing

Some existing networks are able to route packets based on current load in the network. For example, one approach to congestion involves adjusting the routes in real time to send as much traffic as possible on lightly loaded network links.

This sort of load-based routing is a relatively delicate procedure which is prone to instability. It is particularly difficult to achieve stability in multi-vendor environments, in large internets, and in environments characterized by a large variation in network characteristics. For these reasons, we believe that it would be a mistake to attempt to achieve effective load-based routing in an Inter-AS Routing scheme.

A.6.5 Non-Interference Policies

There are policies which are in effect, or desired to be in effect, which are based upon the concept of non-interference. These policies state that the utilization of a given resource is permissible by one party as long as that utilization does not disrupt the current or future utilization of another party. These policies are often of the kind "you may use the excess capacity of my link" without guaranteeing any capacity will be available. The expectation is to be able to utilize the link as needed, perhaps to the exclusion of the other party. The problem with supporting such a policy is the need to be cognizant of highly dynamic state information and the implicit requirement to adapt to these changes. Rapid, persistent, and non-deterministic state changes would lead to routing oscillations and looping. We do not believe it is feasible to support policies based on these considerations in a large internetworking environment based on the current design of IP.

Security Considerations

Security issues are not addressed in this memo.

Author's Address

Mike Little
Science Applications International Corporation (SAIC)
8619 Westwood Center Drive
Vienna, Virginia 22182

Phone: 703-734-9000

Email: little@SAIC.COM