Connectivity within the Internet - A Commentary

Geoff Huston
Australian Academic and Research Network
GPO Box 1142, Canberra ACT 2601 Australia
G.Huston@aarnet.edu.au

Elise Gerich
Merit Network Inc.
Ann Arbor , Michigan, USA
epg@merit.edu

Bernhard Stockman
SUNET / NORDUnet
Royal Institute of Technology, Stockholm Sweden
boss@sunic.sunet.se

June 1992

This paper was presented at the Internet Society's INET'92 Conference at Kobe, Japan, June 1992.

Abstract

The global telephone system and the Internet share a major common attribute of value to all users of their respectives services - that of comprehensive connectivity.

As the Internet grows the major problem we face is that of preserving such connectivity while addressing pressures from the technology itself and pressures from matching the Internet capabilities to the desires of policy determination and management. It is predicted that the number of providers of component networks within the Internet will increase dramatically in the near future. What is not possible to predict is whether effective connectivity can be maintained under such pressures of growth.

This paper explores the issues of connectivity within the Internet, but does not attempt to assert that there are instant solutions at hand - more effort is required at both engineering and management levels to address these issues. The continued operation of the Internet is now critically dependant on the requirement to develop effective and commonly accepted policies relating to the interconnection of

component networks.

Most importantly the conclusion that is becoming apparent is that the Internet itself is not capable of supporting an increasing spiral of arbitrary interconnections and arbitrary requirements for determination of diverse policies associated with service provision. Perhaps the major issue is whether the Internet community is capable of asserting self regulation of interconnectivity for the sake of its own survival or whether the pressures of growth have already made the global Internet too unwieldy to effectively manage as a cohesive communication tool.

Introduction

At regular intervals on the TCP/IP mailing list the question of connectivity within the Internet is raised. The format of the question often starts with the observation that a trace of network reachability between two geographically adjacent points reveals a data path that traverses much of the globe, followed by a query as to why this situation has arisen. Often the situation is compounded by the existence of local connections between the two end points that is evidently not being used by the Internet routing environment.

The general response to such queries goes along the lines that each individual network operator has attempted to install internetwork connections that are intended to meet the majority of perceived requirements. Although this leads to anomalous overall situations at times, the position is that each network provider has attempted to optimise the individual provider's role with respect to cost and functionality. There is an obvious dichotomy within this situation, where the individual actions of optimisations accumulate into an overall situation that is far from optimal (and indeed can be seen as detrimental) for the Internet as a whole.

This illustrates a critical issue within the Internet environment: from the perspective of the Internet as being a single system, large components are now being patched together in a seemingly arbitrary and sometimes nonsensical manner. As the Internet continues to grow, the engineering complexities this situation generates are already eroding the Internet's single major asset - that of ubiquitous connectivity.

The current state of the Internet is that end to end connectivity does not uniformly extend over the domain of the connected Internet. While this may have been due to policy constraints in the past, it is disturbing to note that recent additional constraints on connectivity within the Internet have been based on issues associated with the scale and complexity of routing problems, often based in constraints of policy determination.

It must be stressed that if comprehensive connectivity cannot be maintained through this period of growth, then there is a credible case to be made that there is little left of intrinsic value within the Internet. In the same way that a telephone handset is useless until it is plugged into the global telephone network, the Internet is an instrument designed ultimately to allow the end user to accomplish various tasks that are reliant on an efficient and ubiquitous communications infrastructure.

A Definition of the Domain of Connectivity

In considering the task of defining "connectivity" within the Internet it is necessary to consider what is actually under examination here. A common view has been to base the examination upon the underlying physical structure of the Internet, and focus on the physical links that bind the Internet. However there are two fundamental questions that arise from this methodology:

- What constitutes a "link"?

- What links should be considered within the scope of "coordination of connectivity"?

While a view of connectivity as being simply the domain of physical links between geographically separated locations is relatively straightforward exercise, it does not withstand close scrutiny. It is possible to identify new or near-future technologies, such as Frame Relay and SMDS, which provide a real distinction between point-to-point transmission facilities (physical links) and protocol layer 3 links (logical network links) [Reference 1]. Equally important is the difference in viewing a link as being a connection between two geographically separate locations and viewing a link as being a connection between two distinct organisations (or Autonomous Systems[footnote 1])

The exercise of "coordination of connectivity" is one which is carried out between autonomous Administrative Domains (AD's), not by attempting to determine how the internal structure of an AD is configured. This view of the domain of "connectivity" is therefore that of the network layer interaction between AD's.

A couple of examples highlight this perspective:

a) If an AD within Europe and an AD within the United States provide half-circuit funding for a common trans-Atlantic facility, then such a link would be one which requires consideration of inter-AD connectivity (as indicated in Figure 1).

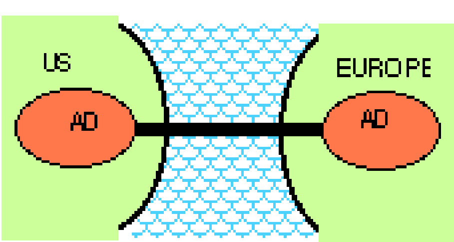

Figure 1 - An Intercontinental Inter-AD link

b) If an AD within the Pacific installed a link, fully funded by the Pacific AD, which was terminated within the United States, then this link is one which is internal to the Pacific AD. If the AD then wished to connect with an AD within the United States, then the link which implements this interconnection is the element which requires consideration of inter-AD connectivity (as indicated in Figure 2).

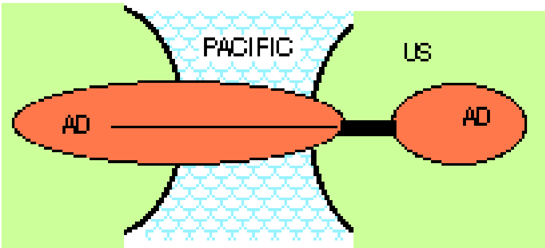

Figure 2 - A local Inter-AD link

To restate this in order to clarify what is being considered here, the impact of connectivity on the Internet is one of the process of defining policy at the point of interaction of the boundaries of the constituent ADs. The coordination process is therefore one of coordinating the network layer interconnections between boundary gateways of two or more networks as a mechanism supporting the determination of the policies of each AD.

Although it would appear intuitive to consider connectivity issues within a bounded domain of physical point-to-point inter-AD links, it is necessary to consider the more complex problem of inter-AD hierarchy routing and associated issues of policy determination at the points of interaction.

Such issues of policy determination include issues such as defined stub and transit AD's, acceptable direct interconnection arrangements between ADs and issues involving acceptance of network routes from other ADs. Such a configuration is as indicated in Figure 3.

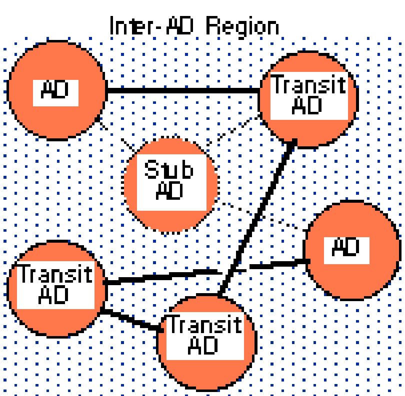

Figure 3 - The Interconnection of AD's with Policy-based Connectivity: Stubs and Transit ADs

The uncomfortable aspect of this shift in perspective is that consideration has moved from an essentially bounded problem to an unbounded problem where determination of the judgements of "positive" or "negative" interconnections (from the perspective of whether the Internet infrastructure is being enhanced or destabilised) is a far more challenging task to undertake. Despite this somewhat negative observation there is also the view that this perspective has in fact identified the key aspect of interconnection, and it allows one to consider the issues of connectivity based on a more realistic view of the process of interaction between the various participants.

The major conclusion is that connectivity within the Internet is subject to a number of very real constraints, most of which are based on considerations of inter-AD routing technologies and capabilities within the Internet. The overall motivations for enhancing connectivity (direct connectivity, enhanced reliability through multiple paths, improved performance through tight interconnection of major service/ client activities and implementation of particular policy objectives) must be balanced against the considerations of viability, manageability and stability which affect the Internet as a whole.

Inter-AD Routing within the Internet Architecture

The initial structure which supported the Internet architecture was the ARPANET. The ARPANET was a single switching domain with connected hosts located at its periphery. Connectivity within such an environment was determined by the ARPANET and within the ARPANET "cloud" host addresses were routed uniformly.

It was a short move from this structure to one which attached networks to the ARPANET core, and from there to a model were a group of related networks, or autonomous Administrative Domains (AD's), were each connected to the core. The routing paradigm used was that of a simple hierarchy. This entailed partial routing information being passed within the interior gateways of each AD (explicit routing information for member networks of the AD, plus a "default" route pointing to the core gateway), and a reachability exchange using the Exterior Gateway Protocol (EGP) between the attached ADs and the core. Within the core itself complete routing information was exchanged between the core gateways, tagging each network with its associated AD number, and performing routing based on the AD gateway [Reference 2].

The resulting model of connectivity was simple: a network either attached directly to the core and used an EGP routing exchange, or a network attached to an existing network domain and the attached networks were advertised as being members of the AD to the core via the EGP exchange. In terms of routing algorithms the model was also simple: if a packet was addressed to a host on a network within the AD then the interior gateways perform the routing of the packet to the destination network, otherwise the packet was passed to the core gateway and then handed to the core for routing (the "default" path).

There are two architectural assumptions within this model: the core maintains complete network reachability information for all attached networks and each network can only be advertised to the core through a single AD. There is no doubt that this hierarchical model of connectivity and associated routing provides architectural simplicity and a certain degree of scaleability[footnote 2], and the resultant configuration can be configured to be a stable platform under the impact of growth.

However the presence of a core is not just a tool for engineering simplicity. The Internet address space is not hierarchically structured, which implies that every router within the Internet must maintain complete and consistent information regarding the relative location of every network address, or know of the relative direction of at least one router which maintains such information. The core model places the responsibility for maintenance of full routing information within the infrastructure of the core itself. Without such a core the picture becomes considerably more complex.

The evolution of the Internet from the original ARPANET has been distinguished by two relevant factors: the single core has been removed, and direct inter-AD (back door) connections have been implemented. The hierarchical routing model of the Internet has been replaced by a model of autonomous regions with various interconnections. Today there is no core (as a physical or a logical entity), and reachability information is passed within the Internet along lines determined more by the accumulation of AD policy rather than overall engineering for optimality and stability.

*IXs, *BONEs and *EXs

To address the problem of interaction and interconnection within the United States, the emerging US Federal Agency networks introduced the concept of interaction across an exchange point (or in this particular case two distinct exchange points). The Agency backbone networks currently interconnect with each other and participate in EGP peer sessions at two points in the US, called Federal Internet eXchanges (FIXes). They are FIX-W at NASA Ames Research Centre in the San Francisco area and FIX-E at SURANET, College Park, Maryland in the Washington DC area. Each Agency can configure their routers to peer with other selected FIX routers in order to achieve the various policy objectives of the Agency. While this provides "open" points of exchange for the participants of the FIX, the transit between the two FIX points is provided by multiple paths which are individually subject to each Agency's constraints.

Perhaps "open" is not the most appropriate word here. The entire US FIX structure operates within US Federal Government policy constraints, so that the use of the FIX as an exchange point for non-Agency traffic is constrained by government policy, in the form of Appropriate Use policies and Agency requirements.

Figure 4 illustrates these interconnections.

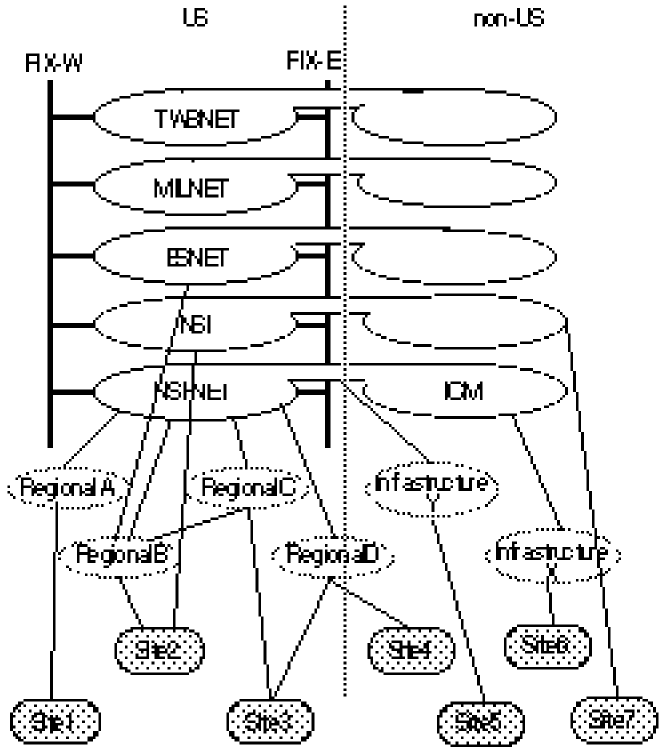

Figure 4 - The US Federal Agency Interconnection Structure October 1991 (after Topolcic [Reference 3])

It should be noted that this (complex) arrangement is not the complete picture within the US infrastructure, and with the emergence of commercial Internet carriers there has been the implementation of a third distinct exchange point, the Commercial Internet eXchange (CIX). The CIX implements a similar function to the FIX, and serves a set of clients some of whom also connect into the US Agency infrastructure (the regional networks). The point to be made here is that this third exchange location adds considerable complexity to the routing problems and it was noted recently that some US regional networks, although they are members of the CIX, have not implemented a CIX connection simply because the routing required to support the multiple *IX connectivity is at this point considered too unstable.

Thus while the concept of using *IXs as boundary points for AD's is a positive direction in terms of implementing connectivity, too much of a good thing (ie. too many distinct *IXs) is extremely detrimental!

This notion of an exchange point to allow interconnection within a bounded domain has been extended by the European community.

The architecture of the proposed EBONE 92 program is that of a backbone interconnecting a number of distinct exchange points. In terms of acceptable use of the EBONE backbone facilities the intention is to have no restrictions on traffic. It will be up to participating networks to restrict traffic according to their own requirements. The EBONE interface is based on a routing exchange with EBONE equipment, with policy being implemented as entry conditions into the regional network.

There are two major aspects to note here: the open nature of the EBONE itself in interconnecting points of interaction with regional networks, and the active role of the EBONE itself as a layer 3 entity.

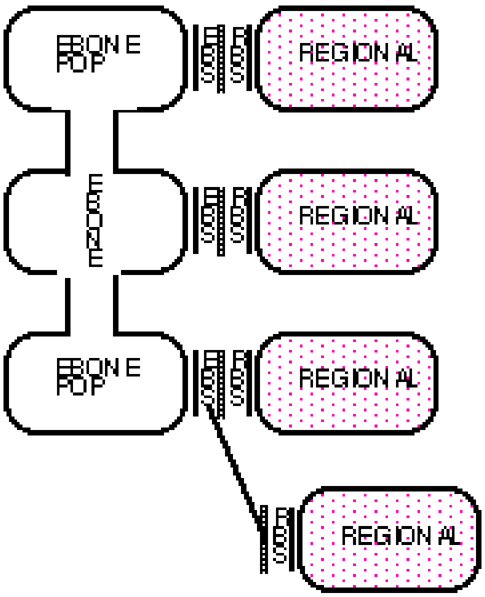

Figure 5 - The EBONE Interconnection Structure:

EBONE POP - Physical Point of Presence of EBONE

EBS - Network Layer EBONE Boundary System

RBS - Network layer Regional Boundary System

To contrast these two approaches it should be noted that in the case of the *IX structure the *IX itself is not an active layer 3 entity, but operates at layer 2 in the protocol stack (the *IX can be an Ethernet segment, an FDDI ring or any other appropriate layer 2 structure). In this way the *IX participants do not interact with the *IX itself - but interact with each other directly.

The positive aspect the *IX architecture is that each *IX participant can explicitly configure their connectivity by enumerating the AD's (and associated networks) with which it wishes to peer. As in the traditional telecommunications model of interaction such a structure supports connectivity based on each entity executing direct bilateral exchanges with those entities with which it wishes to interoperate. The negative aspect is that the core has been collapsed into extinction, and there is no ability to configure a default path within AD's connected to such a *IX. The result of this is that the set of networks which must be explicitly configured within the participating AD's has the potential to grow at the same rate as the Internet itself. The March 1992 announcements of MILNET network filters are indicative of the fact that such a direction does not allow for significant growth in the size of the Internet.

To run the risk of over generalisation the *IX structure can be seen as implementing an explicit inclusion model of connectivity for the connected entities.

The EBONE structure can be viewed as being much closer to the original core model where the core takes on the role of a reachability and routing arbiter. As an active layer 3 entity the EBONE becomes a network routing EXchange (an example of a *EX). Here EBONE connected entities can configure a default path towards the EBONE itself, and then use exclusion filters to implement particular policy objectives. As long as the policy objectives of the AD are not overly restrictive (and in order to allow for maximal connectivity to be passed into the AD this is the typical case) these filters need not grow at the same rate as the Internet itself. In this model the pressures of growth of the Internet are effectively bounded within the EBONE itself as long as the EBONE operates as an open routing domain (here "open" is being used in the a more natural sense; as being without any form of policy constraint on routing and traffic exchange), and are not passed into each participating entity.

The negative side of such a structure is that there is no possibility of direct bilateral AD exchange agreements this structure, nor can an EBONE client determine a specific path through the EBONE[3]. Effectively this implies that inter-AD routing protocols which allow policy-based input cannot be implemented in the exchange between the EBONE and the regional, as the EBONE cannot accept policy constraints on routing exchanges due to its open nature.

Again risking over generalisation, the model implemented within the EBONE structure is that of exclusion at the client level.

The more general question is perhaps why is the US *IX structure of connectivity so much more complex than that proposed within the EBONE model? The answers lie in the fact that policy constraints on connectivity impose significant demands on the underlying technology. In an architecture where individual packets traverse paths as determined by the routing configuration, the only mechanism available to implement mechanisms of policy control on connectivity is based on available routing tools. In an environment where policy objectives determine that multiple entities have a mandate to provide direct connectivity services to a geographically dispersed set of client organisations, then the interaction at the client organisation level, where there may be multiple entity connections, and the interactions of the entities themselves quickly become highly complex. As the size of the routing domain grows, such complexities quickly tend towards instability.

From the perspective of the global Internet the picture is not complete without adding the connectivity issues of the Asia / Pacific region. Here policy objectives and link tariff constraints dictate that there is little international regional infrastructure within the region (unless you consider that Hawaii and the US itself are an integral part of the Asia / Pacific region.).

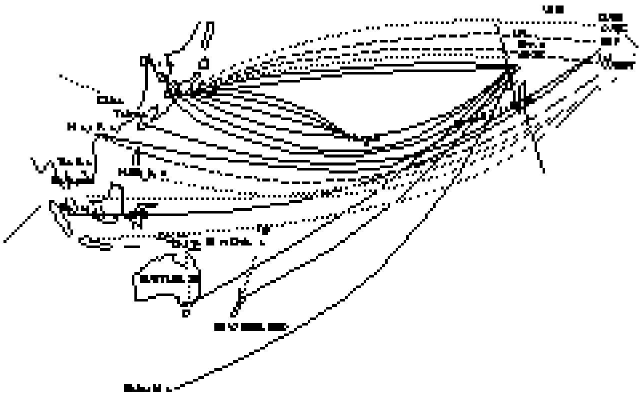

Figure 6 - The Asia/Pacific Interconnection Structure (December 1991)

The overall structure of regional connectivity here is that each national entity (and in some cases more than one entity within a nation) implements its primary connectivity through a link into the US infrastructure, as being an implementation of their primary policy objective[footnote 4]. There are no regional *IX, *BONE or *EX structures in place, as a consequence (in part) of these policy objectives and tariff constraints.

The current picture of connectivity within the region is there are direct connections into Hawaii, the US west coast and even the US east coast, and a large proportion of inter-regional traffic transits portions of the US infrastructure as a consequence. At this point in time there are no direct infrastructural Pacific regional connections to the European infrastructure. While from a routing management perspective this may be considered to be a reasonable position, when one takes into account the fact that there are at present no open transit paths within the US itself, the end result is that inter-Pacific and Pacific / European connectivity is constrained by the policy provisions of the various US entities who are in a position to undertake a transit role for such traffic. Unless this situation is altered in the near future a natural consequence of the growth in connectivity requirements will be a number of Pacific regional entities making direct connections to the European infrastructure, adding even further to the routing management complexities and further eroding the levels of stability of the overall Internet connectivity structure.

The Inter-AD Model of Connectivity

This section examines briefly the factors which are leading the Internet in the direction of increasingly complex interconnection as the support structure for connectivity.

When contrasting this position to that of the telephony system or the global X.25 PSDN structure the fundamental difference lies in the underlying Internet architectural model of connectionless datagram packet exchange.

Within such a model there is no ability to perform virtual circuit setup in response to a virtual connection request, and no ability to perform virtual call accounting. Thus traditional telecommunications interconnection connection arrangements based on a "sender pays all" settlement model between carriers is not a feasible proposition, nor is the establishment of virtual circuits through the switching matrix based on inter-carrier agreements and associated policies.

Within the Internet architecture there is no readily definable concept of the originator of a sequence of datagram exchanges (or a "connection" in the sense of the connection-oriented network architectures), and the application environment of a mixture of direct end-to-end datagram exchanges, and exchanges via forwarders adds an additional layer of complexity to establishing the identity of the datagram transmission originator. To add further to this settlement concepts become even more complex when each datagram may be discarded by any of the network carriers along the transit path, requiring retransmission on the part of the sender of the datagram.

The general model employed by the Internet to address these issues is that of a peer AD structure. For each AD the administrative entity is responsible for the costs in provisioning connectivity within the AD. In terms of inter-AD exchange the general principle of peering applies. Under this principle each AD accepts traffic from its peers and undertakes end delivery or undertakes inter-AD transit to a connected AD, as appropriate. In general no money is passed between AD's within this structure, as the peer principle infers that any such traffic exchanges are notionally decreed to equally balance in both directions between connected AD pairs[footnote 5].

In principle such a notional peer arrangement would infer that all inter-AD connections should be equally funded by the two (or more in the case of a *IX) connected AD entities. In the case where the AD boundaries coincide (as at a *IX for example) then this is typically the case[footnote 6]. Where the costs of the interconnection are non-negligible (trans-oceanic high capacity links are a typical example here) then policy considerations typically dictate that a half circuit funding model only applies where each entity perceives equal strategic benefit in implementing the connection. When the relationship is perceived as being unequal it is no surprise that the lesser partner generally includes the interconnecting link (and the associated costs) within their AD, and implements the inter-AD boundary at the edge of the other AD. A snapshot of current intercontinental Internet links would reveal a mix of both funding models, broadly conformant to the balance to the perceived benefit of the connection to each party.

The Role of Policy Determination within the Internet

It should be noted that further complexities to inter-AD connectivity are added by including consideration of the tools which are deployed to manage traffic flow between (and transiting across) ADs. Here the Internet architecture defines a hop by hop destination based packet forwarding control mechanism, where the determination of the next hop is governed both by applicable relative metrics of "distance" to the destination address and the factor of policy determination as based on source and/or destination addresses[footnote 7].

In looking at policy determination it is appropriate to highlight that the Internet architecture is a direct expression of its genesis within the US research sector. The Internet technology was designed to serve a homogenous user population, so there was no requirement to include heterogenous policy considerations within the packet header format, the routing tools, or mesh such considerations within the associated upper layer application suite. There has been some consideration of the use one of the Type of Service (TOS) bits within the IP packet header as a Research / Commercial bit to aid the effective deployment of policy-based network connectivity, but such a direction would simply allow two policy domains, whereas the scale of the Internet now implies the existence of a broad spectrum of policies which directly impinge on considerations of connectivity.

Even if it is proved possible to efficiently encap-sulate such diverse policy considerations as tagged information within each packet, so that comprehensive policy-based routing configurations were deployable, it can be argued that such a direction is a negative step for the Internet.

In essence such a direction places the role of ultimate determination of policy onto the user base, making the user directly responsible for tagging interactions with the user's local network (in the first instance), and from there to the Internet, with an applicable policy label. As can be noted in a brief study of the evolution of the wide area communications environment over the last three decades, the most useful (and pervasive) wide area technologies are those which seamlessly interact with constituent Local Area Networks, rather than those which attempt to interact separately and directly with the end user clientele in a fashion which is distinct from the LAN interaction.

Thus it is asserted that the most appropriate location for policy determination is as a transaction between ADs, and as a consequence ADs must present a single policy position to all other ADs within the Internet. It should be noted that in this situation the policies that an AD can present to other ADs are not based on the motivations of constituent users and notions of acceptable traffic content within the AD itself, but on the issues of acceptable settlement arrangements and acceptable operation of the connection between AD pairs.

Given such a caveat on the role of policy determination within the inter-AD space, the current efforts in developing robust inter-domain policy routing protocols [Reference 4] can be seen in a reasonable context.

A Structure of Connectivity?

It needs to be asked whether either of the two currently deployed mechanisms of connection described here (the *IX and the *EX/*BONE structures) are capable of supporting the global Internet connectivity requirements.

The *IX structure on a global level would conceivably be a single point of exchange where participating ADs (or grouped ADs with similar policy positions) would place AD boundary routers onto a common layer 2 structure.

The major asset of such a structure is that this would allow each AD to implement exchange policies with peer ADs on a direct bilateral basis, exploiting the major asset of the *IX structure.

However there are a number of engineering and policy problems with such a model. The *IX structure would have to cope with an exchange of not just the current 5,000 connected networks, but also have to scale to exchange some 100,000 networks within a near term timeframe[footnote 8]. The total exchange across such an IX scales by a factor of the number of participating ADs, so that the total exchange requirement would be anticipated to be significantly greater than the number of networks.

Other considerations include the practical impossibility of physically locating such a structure at a position in the globe which allows fair access to all potential ADs, or encompassing the engineering challenges of implementing a distributed IX as a layer 2 wide area structure.

At a policy level the concern with such a model is that the operator of the IX itself is in the position of being a monopoly provider of connectivity - a position which is unacceptable for many Internet ADs who could not then undertake open procurement of connection services.

However perhaps the major problem with such a model is the creation of a single point of potential failure within the global Internet. Within the FIX structure this has been partially addressed by implementing two FIXes, but such a solution is not feasible when considering the complexities of engineering multiple Internet-scale IXs.

The EBONE structure, as an example of a EX, is also not a clear candidate as a feasible structure for supporting comprehensive Internet connectivity.

From an engineering viewpoint such a structure is implementable, as this is in effect an exercise of re-engineering the Internet core as a wide area routing domain. There are issues of scaling of such an EX structure, but it must be noted that such matters of scale are confined to the interior of the EX, and as such an EX would implement no policy constraints internally, there is no associated multiplicative factor on the potential number of connected networks advertised to the EX.

Policy considerations do cause considerable problems regarding the feasibility of the EX. Firstly the connected ADs cannot execute direct bilateral agreements with peer connected ADs - such exchanges occur at layer 3 with the EX itself and admit no policy constraints per se (implementation of policy constraints occur within the AD at the boundary with the EX as an entry constraint according to the EX model,). In practical terms this implies that a settlement model of "sender keeps all" (SKA) is the only practical uniform policy model for inter-AD settlement in this environment.

However the major policy problem of the global EX is one which is shared with the IX model - any potential EX operator is placed in the unacceptable position of being a monopoly provider of Internet connectivity. Within Europe and the EBONE this issue has been effectively addressed by operating the structure under the broad umbrella of a commonly acceptable international organisation (RARE in this case). In a global domain the task of nominating such a potential umbrella EX management organisation is a problem which rapidly approaches intractable proportions.

Thus while an Internet core, implemented as a global routing exchange, offers the potential to provide robust and reliability Internet connectivity, the various policy issues extant across the Internet imply that such a model of structured interconnection is not a uniformly acceptable option at this point in time, and without comprehensive participation the EX structure only adds to the problem rather than addressing it.

In searching for solutions for Internet connectivity today we appear to be left with the single option of attempting to agree upon some basic bounding agreements on the current anarchy which prevails within the inter-AD space.

A Policy for Connectivity?

It is perhaps appropriate to now ask the question as to what factors would constitute the basis of determination that a connection enhances the overall structure of the Internet, and what constitutes a connection which (in some fashion) "degrades" the Internet when examining the inter-AD connection space.

One way to examine this is by examining the motivations which are driving such connections. Here there are two major factors which are common: that of an increase in the overall population of the Internet (and an increase in the scale of the resource requirements from this total population), and an increase in the diversity of the Internet (here "diversity" is used in the sense of the diversity of charter of the network provider, such as open commercial use, publicly-funded support mechanisms for academic and research programs, and so on).

Enhancing the Internet infrastructure is quantified by the objectives of:

* extension of reachability

* enhancement of policy matching

* localisation of connectivity[footnote 9].

* backup arrangements for reliability of operation

* increasing capacity of connectivity

* enhanced operational stability

* rationalisation in the number of ADs being positioned within the connection environment to

allow scalable structuring of the address and routing space to accommodate orderly growth.

From a global perspective this implies a shift in the perspective of policy determination as being one not principally concerned with matching of appropriate use policies per se, but one of determination of inter-AD connection policies which further enhance the infrastructure of the Internet as a whole.

However such an observation does not naturally lead to an enumeration of specific policies of connectivity which would sit naturally within the heterogenous environment which is the Internet today.

Conclusions

It is not proposed that there are acceptable solutions to this issue of connectivity at this point in time. The currently deployed structures of interconnection, such as IXs, EXs and similar do not provide comprehensive answers to the diverse policy objectives which drive the ever increasing complexity of interconnection within the Internet. Even now there are increasingly few locations on the Internet which have comprehensive visibility to the complete Internet structure[10].

This appears to be a critical point we have now reached within the evolution of the Internet. The natural reaction of the various network entities will be to increase the complexity of the structure in order to preserve direct connectivity requirements. Thus today we are in the uncomfortable position of each AD having to provide direct connections to a number of peer ADs in order to achieve even a reasonable level of connectivity. The result of such an increasingly complex mesh of inter-AD connections is overall instability, but in terms of meeting critical immediate objectives, such warnings will not act as an effective deterrent to these actions.

There is little doubt that this results in the situation where the inter-AD space is the critical component of the Internet. This space can be correctly viewed as the demilitarised zone within the politics of today's Internet, and in the absence of any coherent policy, or even a commonly accepted set of practices, the lack of administration of this space is a source of paramount concern.

We have reached the stage where Internet technology is an off the shelf commodity. Thousands of internet protocol capable computing platforms are manufactured and sold every day. The burgeoning environment of network entities operating on a basis of regional, national, international domains is a natural consequence of the shift in this technology from a research project into industrial infrastructure. However the central component of the current Internet is still very much a part of the publicly funded domain, and the shift towards an environment which would see competitive private offerings is exacting a high cost in terms of stability and manageability of the Internet.

Unlike the global telephone system, or the international X.25 PSDN, or the OSI environment, there is no apparent regulation of the overall Internet space by any of the national or international bodies who are actively involved in regulating the deployment and use of related communications technologies. One view is that this lack of imposed regulation is simply a time lag, and that such regulation will inevitably happen, given a recognition of the increasing reliance on this technology as an integral component of each nation's total communications infrastructure. Another view is that the Internet will burn itself out, and that the continued exponential growth in the Internet, as expressed in an exponential growth in connectivity complexity, will cause a fatal balkanisation of the Internet.

Secondly there is no visible international forum where the issues relating to agreed common practices of connectivity could be discussed in an environment where all consituent components of the Internet could have the opportunity to participate.

The most critical function of the Internet Society, and of the various bodies which are active within the Internet domain, is that they immediately direct both engineering and management resources towards defining a structure for the governance of the Internet as a whole, and do so in an open manner with due attention to seeking a reasonable consensus position.

To add a sense of urgency to this plea to the Internet Society for action, the current situation of an effective anarchy within the Internet is not one which can cope with the growth pains anticipated within the forthcoming 12 months, let alone a more comfortable period of some years.

Acknowledgements

Much of the structure used in this paper for the consideration of these issues is the outcome of extensive investigation of these matters by members of the Intercontinental Engineering and Planning Group (IEPG), although the authors would hastily add that the views expressed in this paper are those of the authors, and may not necessarily be shared by the IEPG itself.

References

[1] Bradley, Brown , Malis, "Multiprotocol over Frame Relay", RFC-1294, January 1992

[2] Comer, D., "Internetworking with TCP/IP, Volume 1", Prentice Hall, 1991

[3] Topolcic, C., "IP Routing in the US Federally funded Infrastructure", Briefing paper prepared for the Intercontinental Engineering and Planning Group (IEPG), October 1991.

[4] Breslau, L., Estrin, D., "Design and Evaluation of Inter-domain Policy Routing Protocols", Internetwotking Research and Experience, V2 (4), pp 177-198, December 1991.

---

footnotes

[1] One of the clearest definitions of an autonomous Administrative Domain is that offered by Comer in

[Reference 2].

[2] The qualification here is that although there are issue of scaling, these are confined to within the core in this architectural model. Attached ADs do not have direct visibility of the growth of peer AD's.

[3] Given the absence of policy constraints on the EBONE this would have to be considered a questionable request.

[4] European readers may have a sense of deja vu when reading this, as the situation is not all that different from trans-Atlantic connectivity over the past decade.

[5] While the actual level of inter-AS traffic is a small proportion of the total intra-AS traffic (which is typically the case) any differences between a notional balancing and actual relative traffic levels of inter-AS exchange are regarded as being acceptable compromise between the costs and complexities of undertaking exhaustive traffic accounting and acceptable operational practice. In the case of backbone providers, where all traffic is essentially inter-AD transit traffic, the policy situation becomes considerably more complex.

[6] The usual model used is that each AD connects to the *IX as a part of the AD infrastructure, and contributes a portion of the fixed costs in managing the *IX itself.

[7] While it may be more appropriate to include consideration of other factors in addition to the source and destination addresses within the scope of policy determination, such information is not a component of the Internet packet headers. It can be cogently argued that widespread deployment of policy-based routing tools across the Internet is an inappropriate direction, given that most AD's policies are based on Appropriate Use guidelines which are an expression of acceptable traffic content and acceptable end user motivation rather than end point addresses of datagrams.

[8] And presumably much more in the medium term future - by some orders of magnitude.

[9] While the term may be somehwat obtuse here, the intent is to state that two geographically close points within the Internet should attempt to connect locally in preference to involving a large proportion of the rest of the Internet in their traffic exchange.

[10] The author is compelled to add the note that he is not one of these privileged few!