Comparison of Geographical and Provider-rooted Internet Addressing

Paul Francis ⟨francis@slab.ntt.jp⟩

Abstract

In order to solve the scaling and address depletion problems of IP, the next generation IP protocol will have a larger address. Along with this increase in size is an increase in the number of hierarchy levels in the next generation internet address. This raises the issue of how to assign the new hierarchy levels. This paper discusses the pros and cons of two hierarchical address types—geographic and provider-rooted. It shows that the two address types have different advantages and disadvantages, and that one is not inherently superior to the other. It concludes, however, that because of the unregulated nature of the current internet, provider-rooted addresses are for now the only workable option.

I. Introduction

It has been known for some time that the IP address does not have enough levels of hierarchy, resulting in bad scaling in the routing protocols. Historically, IP has three levels of hierarchy; network, subnet, and host. Because network is at the top of the hierarchy, and because network numbers are assigned at the ratio of one or more per organization, IP scales according to the number of organizations or worse. This level of scaling is unacceptable.

To remedy this situation, and the overall shortfall of IP addresses, the internet community (through the auspices of the Internet Engineering Task Force and its parent organizations) has embarked on a program to improve the addressing characteristics of IP over the near term, and ultimately to define a new internet layer protocol to replace IP. This new protocol, called IPng (for IP next generation) will have larger addresses than IP.

It is universally agreed that the IPng address will have more layers of hierarchy than IP. Thus, there is the question of how to define the hierarchy levels of the IPng address, particularly the layers at the top of the hierarchy—above IP's network number.

This paper explores two basic approaches—geographical and provider-rooted. In the former approach, the top levels of the hierarchy are assigned according to the geographical location of the private network. In the latter approach, the top levels of the hierarchy are assigned according to the provider(s) that the private network connects to.

International telephone numbers (E.163) are assigned geographically [3]. X.121 (the numbering plan of X.25) is geographically assigned at the top, but has an element of provider orientation—from the DNIC [4]. The internet has historically been less regulated than either the global telephone network or X.25 networks. Thus, a pure provider orientation in the addressing is a possibility. This paper explores the pros and cons of geographical and provider-rooted addressing in the context of the global internet.

I.A. Outline

Section II gives some tutorial background on hierarchical addresses. Section III discusses the evaluation criteria for hierarchical addressing schemes. Sections IV and V describe provider-rooted and geographical addressing respectively. Sections VI through IX describe the characteristics of the two schemes in the context of the five evaluation criteria. Finally, Section X summarizes the characteristics and discusses their pros and cons in the current internet environment.

II. Background on Hierarchical Address Assignment

Hierarchical routing and addressing is based on the principal of hierarchical clustering [7]. A cluster is a group of network elements connected together such that there is a path from any network element in the cluster to any other network element in the cluster that traverses only network elements in that cluster. A network element can be a single physical component such as a host or router, a cluster of physical components, or a cluster of clusters.

Because the elements of a cluster are internally connected, once a packet (or call setup) reaches a cluster, it can be delivered to the destination without leaving the cluster. As a result, the cluster can be viewed by the routing protocol as a single element. For this to work, the cluster must be addressable by a single number, typically the prefix of an address. That is, the address prefix identifies the cluster, and the remainder of the address identifies elements in the cluster.

To route a packet to the cluster, only the address prefix is required. Once the packet has reached the cluster, routers in the cluster examine more of the address to further route the packet. Thus, the hierarchical structure of the address directly reflects the hierarchical clustering of network elements. This is the basic premise behind hierarchical addressing.

This having been said, sometimes some levels in a hierarchical address, especially the top levels, are assigned according to a hierarchy of address assignment authorities rather than according to the topological network clustering. For instance, the top level of the NSAP address [6] indicates who the address assignment authority is for the remainder of the address, rather than indicating a network cluster. This is done for administrative convenience, but does not aid routers in routing a packet.

II.A. Internet Topology

While in theory, the only requirement for hierarchical clustering is that the cluster be internally connected, in practice networks often exhibit a natural hierarchical topology. This is the case with the IP internet.

At the bottom of the hierarchy are typically LANs that connect a group of hosts and a few routers. The LAN is a natural cluster. A private network is typically a collection of LANs and point-to-point links connected by routers. Because private networks are usually internally connected, and because the majority of traffic is between internal hosts, a private network is a natural cluster. Private networks are connected to provider networks, which are in turn connected to each other. The purpose of a provider network is to carry traffic from one private network to another, or to other provider networks, which in turn delivers it to a private network. A provider network is a natural cluster.

This natural hierarchical structure, however, is infused with non-hierarchical qualities. Elements in the topology that are not hierarchically above or below each other may still be connected. For instance, two private networks that exchange a lot of traffic may have a link interconnecting them. This link is often referred to as a back door.

In addition, the natural hierarchy isn't a strict tree. A host may be attached to multiple LANs (thus finding itself in multiple places in the hierarchy). More pertinent to this paper, however, is the fact that a private network can be attached to multiple providers, either simultaneously or sequentially over time.

II.B. Geographical Versus Provider-rooted Addresses: An Overview

The relationship between a subscriber and multiple providers (either simultaneously or sequentially) raises some interesting new problems in the IP internet. If the top-level hierarchical address component is assigned to providers, then a subscriber network will get new addresses when it changes providers, and will have multiple addresses if it subscribes simultaneously to multiple providers.

The notion of hosts having a single, static address is deeply engrained in the IP internet. There are no automatic procedures for modifying the addresses of a group of IP hosts, even when all of the IP hosts have the same address prefix and the modification is only to the prefix. In addition, IP hosts generally have little notion of other IP hosts having multiple addresses. For instance, IP hosts generally have no software for choosing among multiple addresses presented to them by directory service, and multiple IP addresses cannot be used to identify a transport connection, even though the multiple IP addresses may identify the same host.

Because of this deeply engrained notion of IP addresses, the introduction of provider-rooted addresses to the IP internet may require significant changes to the operation of the IP internet [9]. While [9] argues that these changes are positive ones, bringing new features and new flexibility to the internet, there is no question that these changes require new functionality and result in added complexity.

An alternative address assignment scheme is that of geographical addresses, such as exists in the global telephone network [3]. Because geographical addresses remove the dependency of address on provider, a subscriber can change providers or have multiple providers without changing addresses. The use of geographical addresses, however, puts an additional burden on providers, in terms of how much routing information they must maintain and on how they must interconnect.

Note that many of the ideas presented here

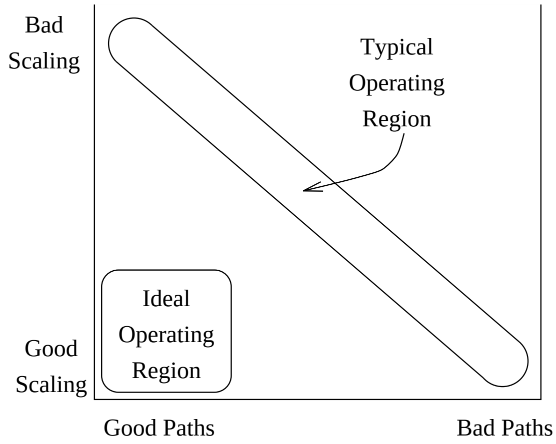

Figure 1: Relationship Between Scaling and Path Quality

were discussed on the big-internet mailing list of the IETF, big-internet@munnari.oz.au.

III. Evaluation Criteria

The assignment of addresses in the internet follows a tree of address assignment authorities. At the root of the tree is the top-level (or level H) address assignment authority. This address assignment authority assigns blocks of numbers to the next level down (level H-1) address assignment authority, which assigns blocks of numbers from the block it owns to level H-2 address assignment authorities and so on. For the sake of discussion, we refer to assigning a block of numbers as simply assigning a number.

The issue is how to assign these numbers so that 1) routing scales well, 2) good paths are found, 3) constraints on the physical topology are minimized, 4) reconfiguration of systems is minimized, and 5) the address assignment process is simple, fair, and politically viable. Consider the graph of Figure 1. In general the "physics" of networking forces operating points on this graph to be along a region extending from the upper left to the lower right. That is, one typically can get good scaling but bad paths, or good paths but bad scaling, or something in between. Depending on the type of address assignment scheme used, however, it is possible to move somewhat towards the lower left (good solutions). This may, however, increase topology constraints or reconfiguration requirements.

Central to the evaluation of any address assignment scheme are answers to the questions 1) what constitutes good scaling, 2) what constitutes a good path, 3) what constitutes unacceptable or costly topology constraints, 4) what constitutes unacceptable or costly reconfiguration, and 5) what constitutes a simple, fair, and politically viable address assignment process. Except for possibly the first question, it is difficult to answer these questions in general terms, partly because the cost of each aspect is borne by different parties, and partly because the cost of each aspect changes over time.

This section generally limits itself to describing the characteristics of the two address assignment schemes, and leaves it to others to determine the extent to which those characteristics are beneficial or detrimental.

IV. Description of Provider-rooted Addressing

The basic approach to provider-rooted addresses is as follows. The top-level address assignment authority assigns numbers directly to providers. This includes both internet protocol service providers and lower-layer (for instance, ATM) protocol service providers. Depending on its size, the provider can either assign the next level internally, or assign the next level directly to its subscribers. The internal assignment would be for clustering groups of subscribers under a single prefix for the sake of internal scaling.

Thus, the address prefixes would be:

provider.subscriber

or

provider.subProvider.subscriber

To understand this notation, consider Figure 2. Shown are three providers with subscribers attached to them. The providers have been given top-level numbers 29, 48, and 14. Provider 29 has given two subscribers next-level numbers 12 and 17. Thus, the upper-left subscriber with assignment 12 has a prefix of 29.12. This means that the field of the address that indicates provider is 29, and the field that indicates subscriber is 12. All host addresses in this subscriber network start with the prefix 29.12.

It is possible that the providers themselves are somewhat hierarchically organized. For instance, there may be long-distance and local-access providers. The subscriber is directly connected to the local-access provider, but may also have a service relationship with one or more long-distance providers to which the local-access provider is connected. In this case, the address prefix could be formed as shown above, or could include both the long-distance and local-access

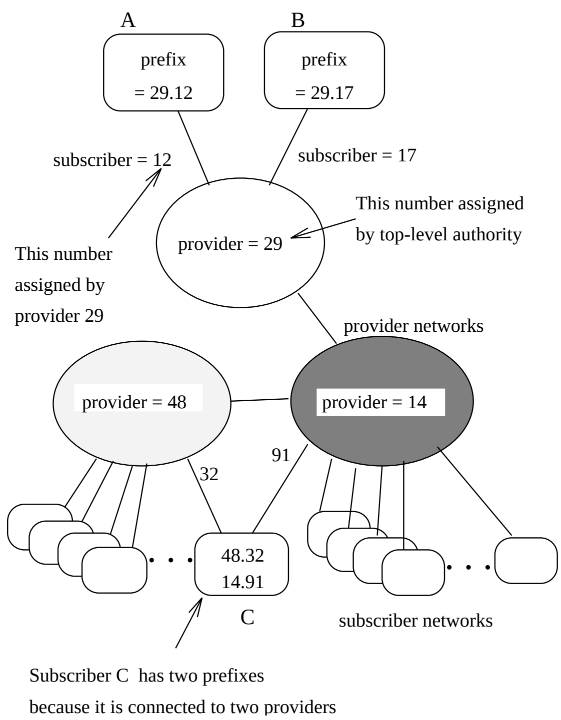

Figure 2: Example of Provider-Rooted Addresses

providers:

LDprovider.LAprovider.subscriber

In either case, subscribers are given an address prefix from each top-level provider through which they derive service. Each host in the subscriber network has one address for each provider through which the subscriber network can be reached—for instance, subscriber C is connected to both providers 48 and 14 in Figure 2. When the subscriber subscribes to a new provider, or unsubscribes from an existing provider, it must change the address prefix for all of its hosts and routers. [1]

Additional levels of assignment happen under the subscriber number for use within the subscriber network. These levels of assignment are not relevant to this discussion.

V. Description of Geographical Addressing

Our working definition of geographical addressing is where the top N hierarchical levels of the address are assigned to geographical regions. Each level of geographical area is completely within the next higher level of geographical area. Three examples of geographical address prefixes are:

country.metro.site
continent.metro.site
continent.country.metro.site

Note that the lowest level of assignment here is to "sites" rather than "subscribers". This is because the notion of a subscriber without a provider doesn't make any sense, whereas the notion of a site within a geographical area does. Both site and subscriber, however, represent private networks that are assigned address prefixes.

Because of the requirement that the elements of a hierarchy cluster must be internally connected, it is necessary in a geographical addressing scheme that all hierarchical elements in a geographical area be internally connected. (Note that this doesn't necessarily require N^2 connectivity, that is, where all N hierarchy elements are directly connected to each other. Rather, it requires that there be some path from any element in an area to any other element in the area that only traverses elements in the area.)

If the geographical clustering is *country.metro.site*, then all metro networks in a country must be able to reach all other metro networks in the country without going through another country. Likewise for all the sites in any metro, etc.

For instance, consider Figure 3. It portrays the providers and subscribers A, B, and C of Figure 2, but shown geographically rather than according to provider. The providers overlap geographical area, so the routers of the providers are shown in Figure 3. The address convention of Figure 3 is *country.metro.site*, where country = 93 and metro = 42. Note that site C (labeled subscriber C in Figure 2) has only one address even though it is connected to two providers. Note also that all of the routers in metro 42 are internally connected by virtue of two (heavily drawn) links between routers of different providers.

Taken to the extreme, the assignment of geographical addresses could be carried all the way to individual hosts. That is, geographical areas could be recursively subdivided until every possible host location in the world (galaxy?) defines a unique address. Clearly this is unworkable. At the local (campus or single building) level, one must assign addresses according to network topology, not some pre-determined geographical partitioning. Thus, at some point in the hierarchy, the addressing must change from geographical to network-physical.

[1] Strictly speaking, the subscriber may not have to change its prefix. This, however, results in worse scaling.

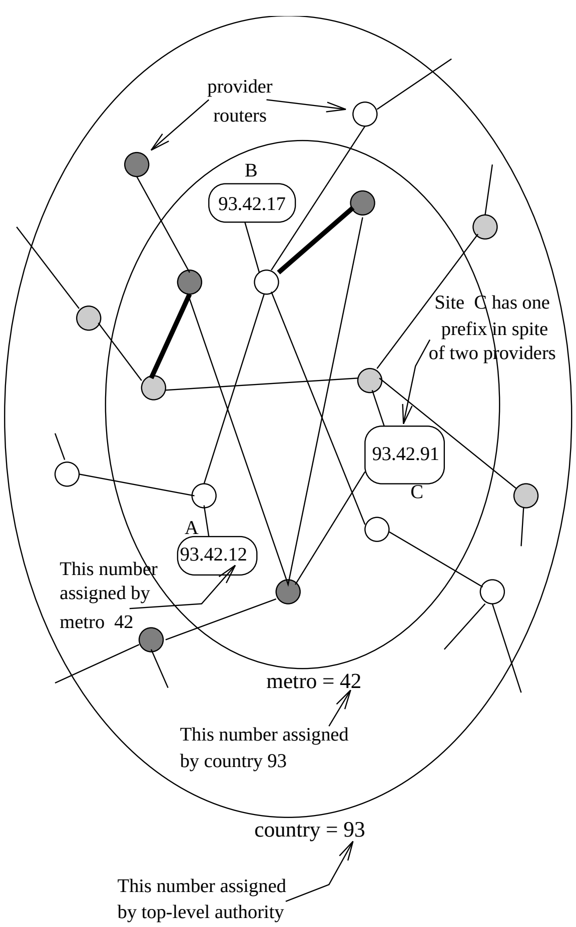

Figure 3: Example of Geographical Addresses

A sensible place to make this change is at the boundary between the private network (or site) and the provider. Within a site, address administration should be completely autonomous and not constrainted by geography (or anything else not within the control of the site). Thus it would not be appropriate to make the change at some level below the provider/subscriber boundary. And, since provider coverage does not necessarily conform to geographical boundaries (some providers are global in scope, and provider coverage areas overlap considerably), it does not make sense to make the change from geographical to network-physical at the boundaries between providers.

Thus, geographical addresses have a geographical part, a site part, and an intra-site part:

geographicalPart.sitePart.intra-sitePart

where each part can have internal layers.

The geographical part for a given site is determined according to the geographical location of the site's connection to a provider. This is where the site "appears" in the global topology. Thus, even though a site may cover multiple geographical areas, if it is attached to a provider in only one geographical area, the whole site will have a geographical prefix indicating that geographical area. More typically, a site that covers multiple geographical areas would be connected to providers in multiple geographical areas.

In any event, the main point is that the specific provider that a site attaches to does not effect the site's address. Thus, a site could change from one provider to another in a given geographical area, or attach to multiple providers in a given geographical area, without changing addresses or having multiple addresses.

VI. Topology Constraints

Provider-rooted addresses place no "unnatural" constraints on the topology. Of course, with provider-rooted addresses, each provider must be internally connected, but a provider would naturally be internally connected, so this represents no real constraint. Provider-rooted addresses place no constraints on how providers interconnect with each other.

Geographical addresses do place an unnatural constraint on topology. That is, they require that the providers that cover a geographical area (that area denoted by the geographical prefix) be connected in that area. While it is natural for providers to be connected to each other somewhere, it is generally (though not necessarily) unnatural to force them to be connected in every geographical area that they cover.

In the current USA Internet topology, provider networks tend to interconnect in a small number places, for instance at FIXs or CIXs (Federal or Commercial Inter-exchange). Thus, requiring connectivity in every metro area, for instance, would require much more interconnectivity that there currently is. On the other hand, the long-distance phone carriers in the USA have connectivity in every geographical area (called LATAs).

VII. Scaling in Routing

In this section, the information needed in routers' forwarding tables for both geographical and provider-rooted addressing is described and compared. The information may vary, depending on the desired quality and flexibility of paths found. This section also describes methods for improving the scaling characteristics of both schemes.

VII.A. Scaling of Provider-rooted Addressing

As stated above, provider addresses are of the form:

provider.subscriber

or

provider.subProvider.subscriber

For addresses of the form *provider.subscriber*, routers in provider networks must, at a minimum, maintain routes (forwarding table entries) for 1) other providers, and 2) subscribers within their own provider network. For addresses of the form *provider.subProvider.subscriber*, routers in provider networks must, at a minimum, maintain routes for 1) other providers, 2) subProvider clusters within their provider, and 3) subscribers within their subProvider cluster.

The number of subProvider or subscriber routes that a router must maintain is within the control of the provider. As a provider obtains more subscribers, it can add internal levels of hierarchy (*subProvider*, *sub-subProvider*, etc.) to keep the number of internal routes manageable. Thus, scaling within a provider network is good.

A provider, however, cannot control the number of other providers for which it must maintain routes. Thus, the size of the forwarding tables at the top of the hierarchy (provider) is open-ended. As a result, the forwarding table size may grow beyond acceptable levels. One solution to this problem is to add another level of hierarchy above the provider level:

providerCluster.provider.subscriber

With this address, multiple providers are clustered within a new top-level identifier, the *providerCluster*. One possible basis for provider clustering (that is, the choice of which providers go into which clusters) is geographical location. In this case, a provider that spanned multiple geographical area would appear as multiple providers, one for each geographical area it appeared in. A perhaps better basis for provider clustering is according to the kind of service provided by the provider. For instance, all ATM providers could form a cluster, all internet providers another cluster, and so on. Another basis could simply be the amount of interconnection various providers have with each other. Providers with a large number of interconnections would naturally be placed in the same cluster.

Another solution is possible under the scenario where a relatively small number of providers are long-distance providers, and the rest are local-access providers. This form of address (*LDprovider.LAprovider.subscriber*) is discussed above. In this case, only the long-distance providers appear in routing tables globally.

Routers in provider networks may also wish to maintain certain information about the internals (*subscriber* or *subProvider*) of another provider. This happens in the case where

1. Two providers are interconnected with each other in multiple places, and
2. the routing policy for one of the providers is to route the packet to the interconnection point closest to the destination (versus simply routing the packet to the nearest interconnection point).

The amount of routing information in this case is also open ended, as it depends on the number of providers with which there are multiple interconnection points (which itself depends on many factors, such as the user traffic matrix), and it depends on the number of subscribers in other provider networks and on how internal clustering is done in other provider networks.

Whether or not it is advantageous for a provider to route a packet to the nearest interconnection point versus route a packet to the interconnection point nearest the destination depends on many factors, not the least of which is the business relationship established between the two providers on how they compensate each other for traffic carried. A discussion of the advantages and disadvantages of this routing policy is outside the scope of this paper.

VII.B. Scaling of Geographical Addressing

As stated above, geographical addresses are of the form:

geographicalN.geographicalN-1 ... site

A router in a provider network must, in the general case, maintain routes for

1. all *geographicalN* clusters, all *geographicalN-1* clusters within their own *geographicalN* cluster, and so on,
2. all sites within the lowest-level geographical cluster that the provider router services.

The number of geographical clusters that a router must maintain routes for is fixed. If the top-level geographical clustering (*geographicalN* in the example above) is continent, then the top-level number of routes is 7 (or so, depending on what constitutes a continent). If the top-level clustering is country, then it is three hundred and something, and if it is metro, then it might be around 10,000 or so. In any event, it is fixed and either does not change or changes slowly and minimally over time. Since the geographical clustering is administratively determined (by whichever address assignment authority has control), the number of routes at the top levels can be set to be something reasonable for current technology capabilities, and thus scales well.

The number of sites within a geographical area, however, is open-ended. Thus, the size of the routing tables at the site level of the hierarchy is open-ended. As a result, the forwarding table size may grow beyond acceptable levels.

One solution is of course to add another level of geographical hierarchy above the site level, resulting in smaller geographical areas. This results in a prefix change for sites, which is counter to the reason for using geographical addresses.

Another solution is to arrange for a packet to visit all providers in a given geographical area, either by putting the packet on a broadcast medium that all providers listen on, or having the packet routed to each provider in turn. Each router that receives the packet knows if the destination is for one of its subscribers, and accepts the packet if it is. Note that the latter solution is generally preferable to the former one, because 1) the broadcast medium may become a traffic bottleneck, and 2) the broadcast medium solution will result in multiple packet deliveries for the case where a subscriber is attached to multiple providers in the geographical area. On the other hand, with the latter solution, there must be a way to prevent a (misaddressed) packet that is not for any of the sites in a geographical area from continuously looping among the providers.

Another solution to this problem is to place a provider layer of hierarchy between the geographical part and the site part:

geographicalPart.initialProvider.sitePart

The extra layer, *initialProvider*, indicates which provider the site initially connected to. Routers in a geographical area, then, must maintain routes for each provider in that area, plus routes for every site that is on longer attached to its initial provider. If most sites remain attached to their initial providers, then the number of routes is greatly reduced.

Routers in provider networks may also wish to maintain certain information about the internals (*subscriber* or *subProvider*) of another provider. This would happen in the case where

1. Two providers were interconnected with each other in multiple places, and

2. the routing policy for one of the providers was to route the nearest interconnection point (versus routing the packet to the interconnection point closest to the destination).

Note that this is the reverse of the routing policy described in the previous section. That is, with provider-rooted addressing, the natural path is to find the interconnection point closest to the source, and with geographical addressing the natural path is to find the interconnection point closest to the destination. In either case, finding the "unnatural" path requires extra forwarding information.

VIII. Address Reconfiguration

This section discusses the conditions under which address reconfiguration in private networks is required for the two schemes.

There are two cases where a private network assigned a geographical address prefix must change that prefix:

1. If the private network changes its provider access location to another geographical area, and

2. If the geographical areas themselves change.

The former would normally happen when a private network moves from one location to another. The latter has happened in the phone network in the USA in the form of area code splits. This happens when the available addresses in an area become depleted, and the area is split in half, assigning a new area code to one of the halves.

Area code splits (or more generally, geographical area splits) can be avoided if the routing supports having multiple (overlapping) area identifiers for the same area. If this is allowed, then a new area identifier can be added to a geographical area if the addresses under the existing area identifier become depleted. Thus, no existing systems need to change address. Routers in an area must still maintain routes to all sites, however.

Another way to avoid area code splits is to simply make the address space in an area large enough to handle all growth. This of course requires a large address space.

There are several cases where a private network assigned a provider-rooted address prefix must change that prefix:

1. If the private network subscribes to a new provider,
2. If the provider has an internal layer of addressing and the subscriber moves to a new location with respect to the clustering defined by that layer, and
3. If the provider modifies its addressing scheme, for instance, by getting a new provider number or adding an internal layer of hierarchy.

Items 2 and 3 for provider-rooted addressing are similar to items 1 and 2 for geographical addressing respectively, and need no further discussion. As emphasized above, the main advantage to geographical addressing is that a subscriber can change providers without requiring a new address. Changing providers is likely to be a fairly frequent event, certainly a much more frequent event than either private networks changing location or geographical areas changing. Just how frequent depends of course on the subscriber, but several changes a year seems feasible.

Because of the frequency of provider changes, it is necessary to have an automatic means of changing all the host addresses in a private network at once. This task is greatly simplified by the fact that it is only necessary to change the provider prefix for each host, and that the change is the same across all hosts. The exception to this would be the case where the new provider prefix is so long that it takes up address space used for numbering in the private network.

There are two basic approaches to automatic private-network-wide prefix reconfiguration. One is to use a general purpose network management device that keeps track of the hosts in a private network and individually updates hosts using a general network management protocol such as SNMP [2].

Another approach is to design a specialized protocol that updates hosts. An example of this would be a modified host/router discovery protocol such as ES-IS [5], where routers periodically broadcast advertisements, and hosts discover the routers by listening to the broadcasts. The broadcasts could contain the prefixes of the private network. In this case, the routers would have to be individually updated to reflect the new prefix, but routers need to be individually configured with addressing information for routing purposes anyway.

Given that a general management facility in a private network is useful for many things, it seems to be a better approach to the prefix reconfiguration problem. Note that the directory service, such as DNS [8], would also have to be updated to reflect the new prefix(es).

It isn't necessarily true that geographical addressing isolates a private network from any per-host administration resulting from provider changes. For instance, consider the case when a private network is connected to multiple providers (or, is connected to one local-access provider but derives long-distance service from multiple providers) and wishes to be able to choose between those providers on a connection-by-connection or packet-by-packet basis. This is called provider selection.

Provider selection is a special case of the more general policy routing [1]. The term policy routing is commonly used to describe the function whereby the source of a packet selects the series of providers that the path traverses. In the case of provider selection, only the providers on either end of the path are selected. In [9] it is argued that the providers closest to the source and destination are the most important, primarily because it is those providers that the source and destination have billing relationships with.

For provider selection to work, the following things, at a minimum, are required [9]:

1. The source must know which providers it is connected to.
2. The source must know which providers the destination is connected to.
3. The source must have enough information about the providers, and possibly how they are interconnected, to make an intelligent policy decision.
4. The source must have a way to indicate in the packet which source-end provider should be chosen.
5. The source must have a way to indicate in the packet which destination-end provider should be chosen.

To do provider selection with geographical addressing, hosts must be configured with information about its connected providers, and directory service must be configured with provider information so that remote hosts can obtain the information. Moreover, this information must be updated when a subscriber attached to new providers. In addition, new mechanisms must be created to cause packets to be routed through the desired providers.

Thus, in order to get provider selection with geographical addresses, the same sort of private-network configuration and packet formatting is required as with provider-rooted addresses. In other words, the network configuration benefits achieved by using geographical addresses are largely lost if provider selection is required.

On the other hand, private-network configuration with geographical addressing is never worse than with provider-rooted addresses. And, if a private network does not require provider selection, for instance because it connects to only one provider, then private-network configuration is easier with geographical addresses in that nothing has to be done if the private network changes providers.

IX. Address and Topology Administration

With provider-rooted addresses, address administration is straight-forward. The top-level address administration authority assigns provider IDs directly to providers. Providers in turn partition the address space as best suits their needs.

Alternatively, the top-level address administration authority can assign blocks of provider IDs to sub-authorities, which can subsequently assign provider IDs to providers in their jurisdiction. For instance, the top-level address administration authority could assign blocks of provider IDs to per-country assignment authorities.

In order to assign geographical addresses, two administrative tasks are required that are not required with the assignment of provider-rooted addresses.

1. The geographical boundaries must be determined.
2. The connectivity between providers within geographical areas must be determined.

It is hard to know the difficulty of these two tasks in the context of the internet. In areas where the establishment of internet providers has been unregulated, it can be imagined that the two tasks are quite difficult. This is because the positioning of geographical boundaries may have an economic impact on providers.

For instance, consider a provider that covers a certain region. If boundaries are drawn such that the provider is completely within a geographical area, then that provider only need to interconnect with other providers in one geographical area. If, on the other hand, boundaries are drawn such that the provider covers parts of several geographical areas, the the provider must interconnect with other providers in each of the geographical areas.

Since it is likely that one of the arrangements (probably the former) will be more advantageous to the provider than the other, the provider will naturally lobby for one set of boundaries over another. This is likely to conflict with the wishes of another provider, resulting in a difficult negotiating process.

Another difficult aspect of address assignment is that of determining how much address space goes to each recipient (either a provider or a geographical area). This is particularly true in the case where the address space is strongly limited, such as is the case with IP.

This aspect of address assignment has both political and technical difficulties. Politically, one organization may object to getting less address space than another. Technically, if not enough address space is allocated, then it is necessary to either renumber or to represent a single entity by multiple prefixes (or both). If too much address space is allocated, then the address space is poorly utilized.

X. Discussion and Summary

Several aspects of geographical and provider-rooted address assignment have been considered. Each technique has different advantages and disadvantages.

Both geographical and provider-rooted address assignments have potential scaling problems. With provider-rooted addressing, the number of providers is open ended. With geographical addressing, the number of sites in a geographical area is open ended. Techniques for improving their respective scaling problems are presented, but the techniques are not attractive.

While it is impossible to predict future growth with certainty, it seems likely that scaling would be worse with geographical than with provider-rooted addressing. The number of large providers is constrained by competitive and economic factors. It is likely that a relatively small number of large providers will dominate. Smaller providers will likely either be merged into the larger providers, or fall under the larger providers in a local-access/long-distance relationship.

With respect to address reassignment, provider-rooted addresses put a larger burden on private networks, since addresses must be reassigned whenever a private network subscribes to a new provider. Since both schemes can have subscriber prefix changes, however, automatic host prefix assignment is desirable in any event, and should be part of the next generation IP. In addition, multiple addresses must be maintained for private networks connecting to multiple providers. This burden, however, can be leveraged for provider selection.

Geographical addressing places more constraints on the topology of the network, since providers must interconnect within geographical areas. Finally, geographical addressing has more administrative/political difficulties, primarily because the geographical boundary locations effect the topology.

Because the two addressing schemes have a different set of advantages and disadvantages, it is impossible to say which is better. Some generalizations, however, can be made. For instance, in general, geographical addressing is better for private networks and worse for providers, where-as the reverse is true for provider-rooted addressing. Also, geographical addressing only works well in a well-regulated or well-organized environment. Because the internet has historically been and still is loosely organized at best, geographical addressing does not seem to be a feasible option at this time. As the internet matures, however, it may obtain better organization, and geographical addressing may become more feasible.

Acknowledgements

I would like to acknowledge Steve Deering's work on geographical addressing. Steve first proposed the use of geographical addresses in the internet, and the discussions of geographical addresses in this paper are largely derived from his work. In addition, many of the ideas in this paper have been discussed on the IETF's big-internet mailing list, big-internet@munnari.oz.au. I would like to mention that most of the work done on this paper was completed at Bellcore. Finally, I would like to thank Mark Garrett for his comments on an early draft of this paper.

References

[1] L. Breslau and D. Estrin. Design of Inter-Administrative Domain Routing Protocols. *Proceedings of ACM SIGCOMM 90*, pages 231–241, September 1990.

[2] J. Case, M. Fedor, M. Schoffstall, and J. Davin. A Simple Network Management Protocol (SNMP). Request For Comments 1157, DDN Network Information Center, May 1990.

[3] CCITT. *CCITT E.163, Numbering Plan for the International Telephone Service*, Blue Book, 1988.

[4] CCITT. *CCITT X.121, International Numbering Plan for Public Data Networks*, Blue Book, 1988.

[5] International Organization for Standardization. *End System to Intermediate System routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)*, ISO 9542, 1988.

[6] International Organization for Standardization. *Information Processing Systems - Data Communications - Network Service Definition - Addendum 2: Covering Network Layer Addressing*, ISO 8348 AD2, March 1988.

[7] F. Kamoun and L. Kleinrock. Hierarchical Routing for Large Networks: Performance Evaluation and Optimization. *Computer Networks*, 1(3):155–174, January 1977.

[8] P. Mockapetris. Domain Names - Implementation and Specification. Request For Comments 1035, DDN Network Information Center, November 1987.

[9] P. Tsuchiya. Efficient and Robust Policy Routing Using Multiple Hierarchical Addresses. *Proceedings of ACM SIGCOMM 91*, September 1991.

Author Information

Mr. Francis has been researching the area of internet addressing and routing for almost a decade, previously with Bellcore and MITRE, and currently with NTT Software Lab in Tokyo. Mr. Francis is active in the IETF, where he is co-chair of the multicast and SIPP working groups. Among Mr. Francis' innovations are Landmark Routing, Kampai Addressing, Shortcut Routing, CBT Multicast, the Pip internet protocol, and aspects of the SIPP internet protocol. Mr. Francis is currently pursuing a PhD with the University College London.