

IP Addressing Schemes - A Comparison of Geographic and Provider-based IP Address Schemes

December 2004

Any identity scheme for a communications network must achieve three quite fundamental functional objectives: identity, location and reachability. In more informal terms the network's identity scheme must be able to identify who is attached to the network, where they are located and how to pass a communication element from one location to another. There are a number of ways to do this.

One approach is to use distinct identity sets for each of these three functional requirements. This is similar to what we use ourselves, where my name identifies myself, but not my location, and a description my location does not reveal my identity. In such a multi-part system there is a requirement for a set of mapping functions to allow an identity to be mapped to a location (such as a telephone directory), and a pair of locations to be mapped to a path specification (such as a map).

A simplifying approach is to combine identity and location into a single addressing scheme, so that an address not only identifies who you are, but also allows the network to know where you are. This approach has the advantage of not requiring a mapping system between endpoint identity and current network location. This simplification that binds endpoint identity with network location has been used in a number of communications systems, including the Internet. In the architecture of IP, an IP address is used as the unique identifier of the network attachment point as well as a location identifier for packet transmission.

Geographic and Provider-based Address Schemes

So we know what an address is used for, but how do we get one? There are a number of approaches to the mechanisms of address distribution. One approach is to divide up the addresses by geographic regions, and then allow the use of local policies to further distribute these token values to those endpoints located within the defined region. A somewhat different approach is to divide the address space in such a way that the address sets are distinguished by service provider, with a consequent objective of distribution of addresses in a manner that is quite precisely aligned to the network's reachability topology.

The first approach is evident in parts of the distribution of identity tokens in the context of the public switched telephone network (PSTN), and the ITU-T's Recommendation E.164. Here the established numbering plan distributes this token space among a collection of (approximately) national administrative entity. Sets of token values with a common prefix value are associated with a specific geographic locale.

A database that has collected all the various national numbering plans for the world's telephone network can be found at www.numberingplans.com. There is a script on that web site that performs a reverse mapping of a telephone number to the corresponding geographic location as specified in the geographic numbering plan.

The objectives of such a geo-political number plan include the desire to ensure uniform access to the token space, and to allow national policies to be determined and implemented without undue disruption to other national domains. The essential assumption here is that the addressed entity remains within the confines of the associated geographic region, and that the topology of the network has some basic relationship with such geographical boundaries.

The other approach is to distribute address blocks aligned with each communication service provider. This is evident in the distribution of addresses in the Internet, where Internet addresses are assigned to end users by the service provider. Here a set of addresses with a common prefix are associated with a n Internet Service Provider, and the common address prefix is functionally synonymous with a provider's identity

Of course neither the PSTN nor the Internet exclusively uses one form of addressing system, or numbering plan. The inclusion of deregulation and competition in the PSTN domain, plus the introduction of mobility and various forms of special purpose number blocks has resulted in numbering plans that also have provider-based address blocks (such as, for example, codes +882 and +881 in the ITU-T's Recommendation E.164, which is further delegated on a provider-by-provider).

Some evidence of geographic-based address schemes can be found in the Internet environment. The address distribution structure used in the Internet uses a single original allocation authority, the Internet Assigned Numbers Authority (IANA). IANA allocates blocks of addresses to each of the RIRs as required, and each RIR operates exclusively in servicing a regional of the globe. At this level it could be asserted that there is a form of geographic-based address distribution, although on an extremely coarse scale of international regions. RIRs then allocate addresses generally on a per-provider basis, and providers then undertake further allocations on the basis of their internal topology rather than any common geographic system.

Constraints on Address Distribution

There are some significant propositions and considerations that lie behind the widespread current use of provider-based address distribution systems within the Internet in preference to the use of various forms of geographic-based addresses.

The Necessity of Aggregation

The first proposition is that in the global Internet some form of abstraction through aggregation of address sets is necessary in order to maintain a network integrity. This is required in order to facilitate the viability of the network's routing and forwarding subsystems.

Within the architecture of the Internet each router has to maintain a set of current reachable addresses. The local forwarding decision for each packet switched by the router is based on a best match between the packet's destination address and the set of addresses held in the router's forwarding table. In the absence of a matching entry, the packet is discarded.

Without any form of address aggregation each router would need to maintain a comprehensive collection of each connected end device's address, and the routing protocol would need to provide comprehensive distribution of updated switching decisions for each end device following any change in network topology. While this may be theoretically feasible, it is entirely impractical. Today (late 2004) a reasonable estimate of the number of discrete host routes in the IPv4 Internet would be of the order of 100 million to 1 billion entries. The related memory requirements to support routing and forwarding tables would imply a requirement for very high speed content addressable forwarding memory of between 1 to 10 Gigabytes, and a routing database which would consume between 10 to 100 Gigabytes of memory, and all this would be required in every Internet router. In addition to the hardware to hold such a massive host address table is the further consideration of the load on the routing protocol when the full routing table contains 1 to 10 billion individual entries. Real time operation of a routing protocol with performance targets of global propagation of state change in the order of tens of seconds would consume significant network resources well beyond the capacity of much, if not all, of current transmission infrastructure.

The Internet's routing and forwarding systems uses one basic tool to reduce the size of the space to supportable proportions. This tool is the aggregation of addresses, where blocks of addresses with a common prefix are manipulated as a group through the use of a single address prefix. In geographic address distribution systems it is intended that the common prefix spans multiple providers operating within a single geographic domain. In provider-based addresses the common prefix may span one or more local, national or multi national geographic regions, but is associated with the customers of a single service provider.

Internet Inter-Provider Interaction

The second proposition is that inter-provider relationships within the Internet are generally constrained to the extent that they are based either upon a customer / provider relationship or a mutual peering relationship where neither party is a provider to, or customer of, the other. To date there have no effective mechanisms devised that would create other inter- provider relationships that would rely on some form of calculated financial settlement derived from the inter-provider traffic profile or any other form of net value transfer calculation.

This topic has been the subject of considerable study over the past decade. The basic observation is that with the transmission of an Internet packet from one provider to another it is not evident from the packet itself that there is an associated net value transfer from the sending provider to the receiving provider, or in the other direction, nor what the quantity of any such value transfer may be. Behind this is the further observation that users of the Internet services do not explicitly fund entire complex transactions (such as a "call", or a mail delivery request) from end to end as a unit of a network service transaction. [1].

The typical starting point for carrier-to-carrier settlements, as they are traditionally used in telephony, is that the retail offering of the provider is one of a comprehensive, end-to-end service, in which the originating service provider utilizes the services of other providers to complete the delivery of the service transaction to the customer. The originating provider then settles with those other providers who have undertaken some kind of role in providing the service. With such cost-distribution mechanisms in place, both small and large providers are able to operate with some degree of financial stability, which in turn allows a competitive market to exist. However, the retail model of the Internet is not necessarily one of end-to-end service on a per transaction basis, but one of partial path service related to packet transmission. There is no price component for complex transactions that covers the complete network path. The Internet can be viewed as a collection of bilaterally funded packet path pairs, where the sender funds the initial path component and the receiver funds the second terminating component.

The corollary of this observation is visible in the inter-domain routing table, namely that the overall majority of inter-provider connections can be characterized as either peer-to-peer or provider-to-customer. The implication is that any inter-provider path between two providers can be seen as a sequence of 0, 1 or more customer / provider relationships, following by 0 or 1 peer-to-peer provider relationship, followed by 0, 1 or more provider / customer relationships.

A second corollary is that from the perspective of a single provider all external relationships can be characterized as a collection of relationships with customers, providers, and peers.

The Provider Business Model and Inter-Provider Routing Policies

The third proposition is that in a stable set of inter-provider relationships no provider permits transit traffic without some identifiable funding source.

This proposition is best illustrated by example. Suppose a provider A peers with provider B at one location and provider C at a different location (Figure 1). Provider A will learn both B's routes and C's routes. What routes should provider A announce to each of B and C? In particular, should A announce B's routes to C, and C's routes to B? In both cases the answer is no. Neither B or C are paying A for the peering connections, so the costs associated with the transit across A are unfunded.




Figure 1 Peer-to-Peer Transit

Furthermore, suppose provider A uses provider T as an upstream transit provider. A is now a customer of T (Figure 2). Should A announce T's routes to either of the peer providers B or C? Again the answer is no. T is A's provider, and is not funding A for any form of transit service. Neither B nor C are paying A for transit services, so A has no funding to provide transit between B or C and T.




Figure 2 Peer-to-Upstream Transit

The more general characterization of the inter-provider relationships is that in general a provider has three sets of relationship: customers, peers and upstreams. In order to avoid being used as an unfunded transit provider the routing policies used by providers have a consistent set of constraints such that routes learned from customers are announced to all other customers, to peers and to upstreams, while routes learned from peers are only announced to customers, and routes learned from upstreams are only announced to customers. This is indicated in Figure 3.

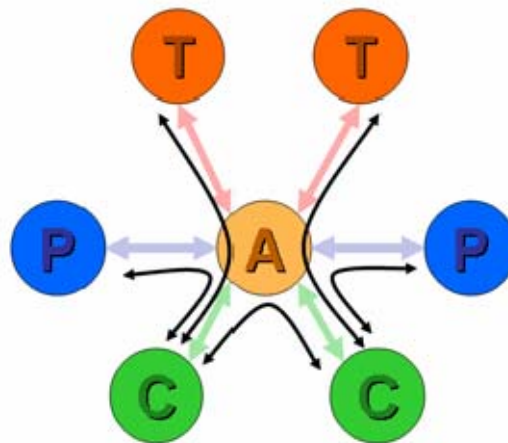


Figure 3 Generic Routing Re-Advertisement Policies

Geographically Aggregated Address Schemes

We can now move on to look at geographically aggregated address schemes and their viability within the context of the Internet. The assumption used here is that geographic address distribution schemes are aggregateable by geographic region, but are not aggregateable by provider.

Any provider providing services into a geographically aggregated service domain can choose either to announce the geographic address prefix as part of its announced routes, or announce the enumerated set of individual explicit customer routes of its customers.

The latter course of action is little different from fully enumerated host routing outlined previously, and suffers from the same issues of inability to scale to any meaningful extent. The former course of action has the provider announcing a single aggregate prefix to its peers and upstreams. Within that single prefix are routes associated with the service provider's customers and routes

associated with customers of other service providers. As long as all other service providers are also customers of this service provider (a geographic monopoly model) then this is a viable structure in terms of routing policies (Figure 4). But where the other providers within the region are peers of this provider, then within the current Internet inter-provider framework, this creates a situation where the provider assumes a role of unfunded transit service provider for certain traffic (Figure 5). As described the previous section, this situations analogous to a provider announcing a peer's routes to its upstream, and incurring unfunded transit traffic as a consequence. This is an untenable business proposition for the service provider.

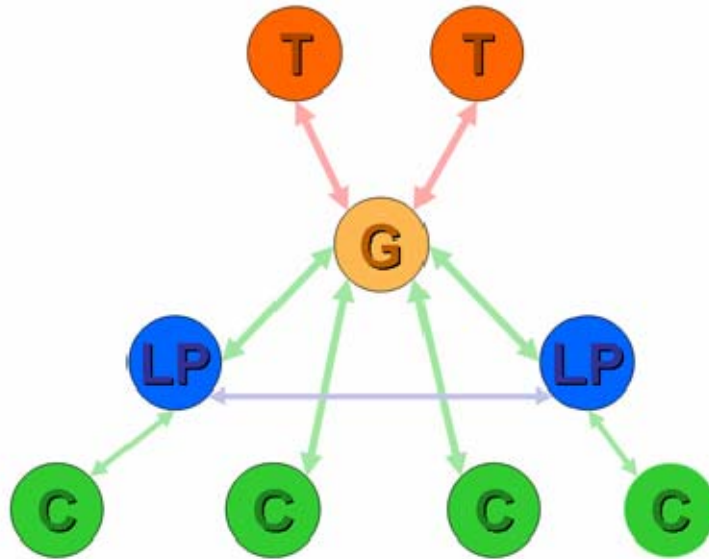


Figure 4 Announcement of a Geographic Aggregate with local monopoly transit provider




Figure 5 Announcement of a Geographic Aggregate with peer providers

Where other providers within the geographic region that do not peer or interconnect directly with each other, then there is a consequent situation of split routing. For example, if providers A and B both provide services to the same geographic region, announce the geographic address aggregate, but do not interconnect, then when A receives a packet from its upstream destined to a customer of B, then A has no option but to discard the packet as it has no more specific information as to how to reach B (Figure 6).

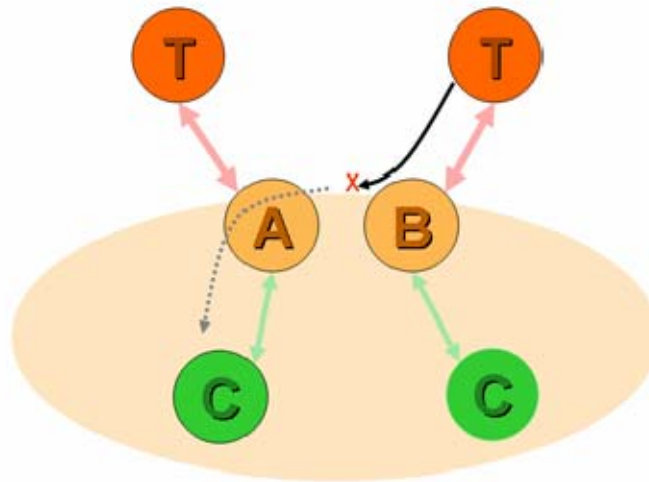


Figure 6 Split Routing of a Geographic Aggregate

Geographic addressing schemes as a means of address aggregation in the routing table are only a viable approach when adopted concurrently with a service provider model of geographic monopoly providers. This could be also be stated as the proposition that in order for geographic address aggregation to be viable the network's topology must match geography, and the provider-based aggregation is then synonymous with geographic aggregation. The inference here is that within the constraints of this model of inter-provider interaction two providers cannot provide service to the same geographic area and both announce transit to the same corresponding aggregate address, and at the same time have the assurance that they are operating within the parameters of a fair and sustainable business model.

In a more conventional multi-provider competitive model of service provision geographic address aggregation is not supportable in terms of inter-provider interaction within the Internet environment. Such a model of address aggregation is sustainable only when, at a minimum, each provider's service quality is directly substitutable with any other, and there is an objective, uniform and fair form of incremental financial inter-provider settlement structure that can provide each provider with the assurance of funding for all forms of transit service provision. Such preconditions are not evident in the Internet today.

Provider-Aggregated Address Schemes

Provider-aggregated address schemes use an address distribution model where a provider receives an address block which is then distributed to end users who are customers of that provider. In announcing its address block to its neighbouring peers and upstream providers the provider is able to announce its role in providing reachability to its customers with a single aggregate advertisement. Provider-based aggregation is an address system where the address plan is aligned to the inter-provider network topology.

In terms of scaling the routing and forwarding system provider-based address aggregation provides significant leverage. Outside of the provider's network, the routing system needs only to learn of the aggregate address block (or collection of blocks) that encompass the provider's address set. At this level the internal details of individual customer assignments need not be visible to the routing system. This form of address aggregation also matches quite precisely the generic model of inter-provider interconnection. The provider's aggregate address block corresponds to the collection of the provider's customers, and any associated service infrastructure operated by the provider. According to this generic route propagation model the provider will announce its aggregate address blocks to all customers, peers and upstreams. The aggregate blocks learned from peers and transit upstreams will only be announced to customers.

It is only within the providers internal routing domain that the individual customer assignments need to be enumerated in the routing system. Dynamic changes to connectivity may generate routing updates and reachability changes within the provider's network, but externally to the provider no change to the provider's aggregate announcement is visible. In this fashion provider-based address schemes provide damping of finer detail of dynamic change in the network.

While the addressing and routing system are constructed as a two tier hierarchy of providers and customers, where providers exchange aggregate address blocks and customers are provided with a smaller block of addresses from their provider, the real world of internet service providers does not align itself so neatly. Customer of providers may also be providers in their own right, creating multiple levels of the provider hierarchy. Providers at each of these levels may interconnect in various ways, and wish to exchange routes that refer only to their customers, and not include all other customers of their respective upstreams.

The challenge in such a model is to define a suitable boundary condition of who is a provider and who is a customer, which, in effect determines whether the provider obtains an aggregate address block that can be announced in the inter-provider routing space in its own right as a 'unit' of inter-domain routing information. This is the subject of many of the policies used by the Regional Internet Registries in guiding their function of allocation of address blocks. The criteria for what defines a provider in terms of these policies can be summarized as a set of technical, environmental and economic criteria that attempt to create a careful balance between meeting the expectations of a very broad diversity of industry players who see themselves as providers in their own right and the constraints of scaling the Internet's routing system within parameters of current capability.

Moving Forward in Addressing

Within the current constraints of the Internet's architecture and the deployed hardware and systems the Internet uses provider-based address aggregation as a means of ensuring that the Internet continues to operate in a viable and cost-effective manner. Other forms of address aggregation, in particular geographic-based address aggregation, cannot provide similar leverage, due, in no small part, to the limitations in the flexibility of the Internet inter-provider interconnection models available to industry players.

In concluding that geographic address aggregation is not a practical approach as a tool for scaling routing and forwarding in today's Internet, this does not imply that the current methodology of provider-based addressing is flawless in practice.

As has been commonly observed, the issues with provider-based addressing concern the customer's portability between providers and the desire of increased service robustness through the capability for a customer to simultaneously obtain service from 2 or more service providers. In the first case the provider is ultimately forced to renumber the devices that are part of their local network from addresses drawn from one provider's address block to addresses drawn from the other provider. For large (and often for smaller) end sites this is a complex and expensive task [2]. The specification of highly capable dynamic address assignment technologies in IPv6 has improved this situation to some extent, but there remain issues with the tight synchronization of the domain name system to the address change, the fate of long-held sessions during an address switch and the logistics of triggering a readdressing event on each end device within the site. Connecting to more than one provider, or multi-homing, presents significant complexities if provider-based aggregation has to be honoured. If we want to steer clear of various forms of edge renumbering technologies then each local device has to have as many alternate addresses as there are providers, and in order to support the desired resiliency that motivates multi-homing, the host must be able to be agile across multiple addresses within a single end-to-end session.

While the IETF Multi-Homing in IPv6 Working Group (Multi6) [3] continues to develop approaches in this space, the overall observation is that multi-homing support and provider-based address aggregation present a significant challenge to the architecture of the Internet itself, and the solution space invariably explores a quite basic shift in the architecture of attempting to decouple endpoint identification from provider-based endpoint location identification.

Such an evolution of the Internet architecture would address many of these issues with provider-based addressing scheme, as both renumbering and multi-homing can be seen as presenting similar functional requirements to the protocol's architecture.

And with this we are back to the basic question of the IP architectural simplification of coupling of endpoint identity with endpoint location. As we try to admit greater flexibility in identification schemes, and at the same time attempt to support an ever-growing network, the pressure to shift the architecture to admit a decoupling of identity and location appears to be increasing. In recognizing the limitations as they relate to issues of consumer choice and flexibility, the provider-based aggregation model is not an ideal outcome. When identity is tightly coupled with location, and both are bound inside a provider's address block, a change in provider implies a change in address, which, in turn, implies a change of IP level identity.

It would appear that that a way forward is through the application of a well known adage of Computer Science, that any problem can be solved with yet another layer of indirection. Perhaps by decoupling identity from location within the IP architecture we can continue to use provider-based location address aggregation in the Internet's routing and forwarding systems, while at the same time preserving a constant identity-based address. In other words, it would be a definite step forward to be able to claim, in the context of the IP architecture that wherever I may be located and whatever network paths we may use to communicate, I am still me!

References

- [1] "Interconnection, Peering and Settlements", Geoff Huston, Internet Protocol Journal, Vol2. No. 2, June 1999.
http://www.cisco.com/warp/public/759/ipj_2-2/ipj_2-2_ps1.html
- [2] "Renumbering Needs Work, B. Carpenter, Y. Rekhter, RFC 1900, February 1996.
<http://www.potaroo.net/ietf/idref/rfc1900/index.html>
- [3] IETF Multi6 Working Group's resource collection, December 2004.
<http://ops.ietf.org/multi6/>

Geoff Huston

Disclaimer

The above views do not represent the views of the Internet Society. They were possibly the opinions of the author at the time of writing this article, but things always change, including the author's opinions!

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is the Executive Director of the Internet Architecture Board, and a member of the Board of the Public Interest Registry. He was an inaugural Trustee of the Internet Society, and served as Secretary of the Board of Trustees from 1993 until 2001, with a term of service as chair of the Board of Trustees in 1999 and 2000. He is author of a number of Internet-related books. He is the Senior Internet Researcher at the Asia Pacific Network Information Centre, the Regional Internet Registry serving the Asia Pacific region.