

A Commentary on IPv6 Address Distribution Policies

In November 2004 the Director of ITU/TSB, Houlin Zhao, proposed consideration of an address distribution function using country-based address assignments. This was proposed as a parallel service function to that managed by the Regional Internet Registry (RIR) system [1]. A merit of this proposal was suggested to be that "people will have a good choice", implying a positive benefit from the establishment of competitive parallel distribution systems. The document indicated that details and constraints associated with such a proposal were yet to be investigated.

This document will examine some of these details and constraints, related to the established address distribution models, and the model suggested by Mr Zhao.

1. The Theory and Practice of IP Addresses

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 approaches to this three part objective in terms of the architecture of identity schemes.

One approach is to use distinct identity sets for each of these three functional requirements. This is similar to a human use system, where my name identifies myself, but not my location, and a description my location does not reveal my identity, and a set of instructions to travel between two locations does not reveal the identity of the individual at either location. 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, and a pair of locations to be mapped to a path specification.

A simplifying approach is to combine identity and location into a single identity scheme, so that a token drawn from this space not only uniquely identifies a communication endpoint, but also has semantic content which can be used to locate the endpoint within the network. This approach has the advantage of not requiring a mapping system between endpoint identity and current network location, so that with knowledge of an endpoint's identity a communication can be initiated without requiring any additional information. This simplifying assumption 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 attachment point to the network as well as a location identifier for packet transmission.

1.1. Geographic and Provider-based Address Distribution Systems

With any addressing system, there are a number of approaches to the mechanisms of address token distribution. One approach is to divide the token space into a set of distinct token sets, assign each grouping of address token values to geographically distinct administrative entities 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 token space in such a way that the token sets are distinguished by service provider, with a consequent objective of distribution of identity tokens 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.

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 behind a geographic-based numbering plan 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 token sets that are aligned by communication service provider. This is evident in the distribution of public address tokens in the context of the Internet in both IPv4 and IPv6. In this address realm address tokens are assigned to individual entities by the service provider, and the address token value used is drawn in every case from the provider's own token set. Here a set of token values with a common prefix are associated with a common provider, and the common 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.

2. 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.

2.1 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 network integrity, and to ensure 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 aggregation of address sets each router would need to maintain a comprehensive collection of each current end entity's address, and the routing protocol would need to provide comprehensive distribution of updated switching decisions for each end entity 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 10^{**8} to 10^{**9} entries. The related memory requirements to support routing and forwarding tables would imply high speed forwarding memory of some 10^{**9} to 10^{**10} bytes, and a routing database which would consume 10^{**10} to 10^{**11} bytes, configured across some 10^{**8} routers. In addition to the hardware consideration is the further consideration of the load on the routing protocol when the full routing table is some 10^{**10} to 10^{**11} bytes in size. Real time operation of a routing protocol with performance parameters in terms 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 use one basic tool to reduce the size of their databases 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 manipulation of the single 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.

2.2 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 or customer of the other. To date there have no effective mechanisms devised that would create other inter- provider relationships that 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. [2]

The typical starting point for carrier-to-carrier settlements 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. The originating provider then settles with those other providers who have taken some kind of role in providing the service. With such cost-distribution, 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, but one of partial path service. And there is no price component that covers the complete path. The Internet can be viewed as a collection of bilaterally funded 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.

2.3 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. 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.

Furthermore, suppose provider A uses provider T as an upstream transit provider. A is now a customer of T. 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.

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.

3. Geographically Aggregated Address Schemes

It is now possible to analyse 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. 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. As per the previous section, this situation is 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.

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.

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 two providers cannot provide service to the same geographic area.

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.

4. Provider-Based Address Aggregation

In concluding that geographic address aggregation is not a practical approach as a tool for scaling routing and forwarding in the 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 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 end sites this is a complex and expensive task [3]. IPv6 has improved this situation to some extent, but there remain issues with the synchronization of the domain name system, the fate of long-held sessions and the logistics of triggering a renumbering event on each end point within the site. Connecting to more than one provider, or multi-homing, presents significant complexities if provider- based aggregation is to be honoured. In this case each local entity has 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.

With the IETF Multi-Homing in IPv6 Working Group continues to develop approaches in this space the overall observation is that multi-homing support and provider- 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 (who am I talking to) from provider-based endpoint location identification (where are you).

Such an evolution of the Internet architecture would address these noted 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. However such work would do nothing in terms of addressing the shortcomings of geographical-based addresses.

5. Country-Based Address Distribution Structures

In terms of the details and constraints of a country-based address distribution scheme, each national address administration would still need to be constrained to be aligned to provider-based addressing.

This then raises the issue that if the address distribution function is one that is constrained to be aligned to provider-based network topologies and the RIR model is attuned to undertaking this function along such lines, then what would be the role, function and value of a parallel national- based address distribution function?

A related uncertainty within this proposition is how could we ensure that a national administration would adhere to such a constraint? The issue here is one similar to the tragedy of the commons in that the aberrant actions of an individual national administration in operating a non-aggregateable address distribution system may not necessarily imperil the viability of the Internet's routing system. It is, however, a more certain prediction that if every national administration followed the same path then the cumulative outcome would be failure of Internet routing and failure of the Internet itself. The mechanism of recommendations adopted by the ITU provides no reasonable level of assurance of adherence to the common constraint of effective aggregatibility of addresses.

Given that such a parallel function would also have its costs of operation, then what benefits would the end consumer, who ultimately bears such costs, derive from this duplication of function? If the end user is the intended beneficiary of the function, the competitive supply of a commodity is one possible remedial measure when there is the situation of a sole incumbent operator operating the distribution function with the inclusion of monopoly rentals added to the price of the goods and services. However competition is not a panacea for all possible markets associated with the distribution of all forms of goods and services. The RIRs structure is based on the form of a mutual membership association where the function itself is operated on the basis of the cost of the service, rather than as a for-profit enterprise. A 'competitive' national enterprise would not be able to operate at the same economies of scale of operation, but would have to fulfil the same set of functions and operations as the RIR. If this function were operated on the same basis as the RIRs then there would appear to be no particular cost advantage to be derived at all. If it were to be operated at a cost to the public purse, then the national constituency would inevitably have to bear the incremental cost of operation of this duplication of function.

This should be contrasted to the way in which national administrative entities already operate within the RIR framework. The national entity operates in the local time zone, using local staff and local languages. The

address distribution framework is identical to that to the RIR itself: one of provider based address allocation and provider-based aggregation of addresses in the routing system. Within this model the national entity operates under the same set of constraints as policies as all other RIR- serviced entities, ensuring a common outcome of viability of the Internet's routing system and a uniformity of distribution practices. This ensures that there is very limited opportunity for 'registry shopping' for an entity to obtain advantaged outcomes by using a particular registry, and allows the underlying address system to be fully portable in terms of their deployment across the provider's network. One characterization of this model is that of an 'agency' function, where the national entity presents a functionally equivalent interface to providers as that of the RIR.

6. Conclusions

Within the current constraints of the Internet's architecture and 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 to limitations in the flexibility of the inter-provider interconnection models available to industry players.

Address distribution functions need to be aligned to this constraint of provider-based address aggregation, and this model is one that forms the basis of the RIR function.

Nationally based address distribution functions are used in some regions within the RIR model, but in this framework there are some critical features that are at variance with the ITU proposal. The national entities do not have absolute policy autonomy and cannot determine to use radically different mechanisms of address distribution. This ensures that the overall constraint of address aggregation is achieved with a uniformity of outcome in terms of address aggregation and routing viability, without creating imbalanced and untenable pressures on the routing system. It is also the case that such national operations are not in direct competition with the RIR mechanisms, and are structured, like RIRs, as industry service entities.

The overall objective of the RIR system is to ensure that the address distribution function operates in the service of a working network, is efficient, accessible, fair and balanced, and operates impartially in terms of the application of its policies.

The proposed introduction of a second mechanism of address distribution to national entities, which could then operate autonomously with neither collaboration with the Internet community nor adherence to a common outcome, introduces significant elements of risk to the future of a viable and operational Internet. In this respect it would be appropriate to offer the advice to the ITU-T to reconsider its proposal and, with those communities that already have had considerable experience in this area of activity, undertake further careful investigation into effective address distribution mechanisms for the Internet.

References

- [1] "ITU and Internet Governance", input to the 7th meeting of the ITU Council Working Group on WSIS, 12-14 December 2004, H. Zhao, Director of TSB/ITU, Geneva, 30 November 2004. <http://www.itu.int/ITU-T/tsb-director/itut-wsis/files/zhao-netgov02.doc>
- [2] "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
- [3] "Renumbering Needs Work, B. Carpenter, Y. Rekhter, RFC 1900, February 1996. <http://www.potaroo.net/ietf/idref/rfc1900/index.html>