

Observations on the Management of
the Internet Address Space

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This memo examines some of the issues associated with the current management practices of the Internet IPv4 address space, and examines the potential outcomes of these practices as the unallocated address pool shrinks in size. Possible modifications to the management practices are examined, and potential outcomes considered. Some general conclusions are drawn, and the relevance of these conclusions to the matter of formulation of address management policies for IPv6 are noted.

1. Introduction

The area explicitly examined here is the allocatable globally unique IPv4 address space. Explicitly this includes those address groups uniquely assigned from a single comprehensive address pool to specific entities which are then at liberty to assign individual address values within the address group to individual hosts. The address group is handled by the technology as a single network entity.

At present these addresses are allocated to entities on a freely available, first-come, first-served allocation basis, within the scope of a number of administrative grounds which attempt to direct the allocation process to result in rational use of the space, and attempt to achieve a result of a level of equity of availability that is expressed in a sense of multi-national "regions" [1].

In examining the current management policies in further detail it is useful to note that the IPv4 address space presents a number of attributes in common with other public space resources, and there are parallels in an economic analysis of this resource which include:

- the finite nature of the resource

This attribute is a consequence of the underlying technology which has defined addressed entities in terms of a 32 bit address value. The total pool is composed of 2^{32} distinct values (not all of which are assignable to end systems).

- the address space has considerable market value

This valuation is a consequence of the availability and extensive deployment of the underlying Internet technology that allows uniquely addressed entities the capability to conduct direct end-to-end transactions with peer entities via the Internet. The parameters of this valuation are also influenced by considerations of efficiency of use of the allocated space, availability of end system based internet technologies, the availability of Internet-based service providers and the resultant Internet market size.

- address space management is a necessary activity

Management processes are required to ensure unique allocation and fair access to the resource, as well as the activity of continuing maintenance of allocation record databases.

Increasing rates of Internet address allocation in recent years imply that the IPv4 address space is now a visibly finite resource, and current projections, assuming a continuation of existing demand for addresses predict unallocated address space exhaustion in the next 6 - 12 years (rephrasing current interim projections from the IETF Address Lifetime Expectancy Working Group). There are two derivative questions that arise from this prediction. Firstly what is the likely outcome of unallocated address space exhaustion if it does occur, and secondly, are there corrective processes that may be applied to the current address management mechanisms that could allow both more equitable allocation and potentially extend the lifetime of the unallocated address space pool. These two issues are considered in the following sections.

2. Outcomes of Unallocated Address Space Exhaustion - No change in current Address Management Policies

As the pool of available addresses for allocation depletes, the initial anticipated outcome will be the inability of the available address pool to service large block address allocation requests. Such requests have already been phrased from various utility operators, and the demand for very large address blocks is likely to be a continuing feature of address pool management. It is noted that the overall majority of the allocated address space is very

inefficiently utilised at present (figures of efficiency of use of less than 1% are noted in RFC 1466, and higher efficiency utilisation is readily achievable using more recent routing technologies, such as Variable Length Subnet Masks (VLSM) and disjoint subnet routing). Given the continuing depletion of the unallocated address pool, and the consequent inability to service all address allocation requests, it is a likely outcome of interaction between those entities with allocated address space and those seeking address allocation that such allocation requests could be satisfied through a private transaction. In this situation an entity already in possession of a sufficiently large but inefficiently utilised allocated address block could resell the block to a third party, and then seek allocation of a smaller address block from the remaining unallocated address space. The implication is that both address blocks would be more efficiently utilised, although it is the entity which has large blocks of allocated address space which would be the primary beneficiary of such transactions, effectively capitalising on the opportunity cost of higher efficiency of address block use.

Such reselling / trading opportunities which involve the use of the unallocated address pool would in all likelihood be a short term scenario, as the high returns from this type of trading would increase the allocation pressure from the pool and act to increase depletion rates as more pressure is placed to claim large address blocks for later resale once such blocks are no longer available from the unallocated pool.

Following exhaustion of the unallocated address pool a free trading environment in address blocks is a probable outcome, where address blocks would be bought and sold between trading entities. The consequent market, if unregulated, would act to price address space at a level commensurate with the common expectation of the market value of addresses, trading at a price level reflecting both the level of demand, the opportunity cost of more efficient address use, and the opportunity cost of deployment of additional or alternate internetworking technologies to IPv4. It is interesting to note that within such an environment the registry (or whatever takes the place of a registry in such an environment) becomes analogous to a title office, acting to record the various transactions to ensure the continued accuracy of "ownership" and hence acts as a source of information to the purchaser to check on the validity of the sale by checking on the validity of the "title" of the vendor. This impacts on the characteristic features of Internet address registries, which effectively become analogous to "titles offices", which typically are structured as service entities with "lodgement fees" used to fund the action of recording title changes. Whether existing registries adapt to undertake this new function, or whether other entities provide this function is a moot point - either way the function is a

necessary adjunct to such a trading environment.

It is also anticipated that in an unregulated environment the trade in address blocks would very quickly concentrate to a position of address trading between major Internet providers, where a small number of entities would control the majority of the traded volume (market efficiency considerations would imply that traders with large inventories would be more efficient within this trading domain). It is also reasonable to expect that the Internet service providers would dominate this trading area, as they have the greatest level of vested interest in this market resource. This would allow the Internet service provider to operate with a considerably greater degree of confidence in service lifetime expectation, as the service provider would be in the position of price setting of the basic address resource and be able to generate an address pool as a hedge against local address depletion for the provider's client base. There is of course the consequent risk of the natural tendency of these entities forming a trading cartel, establishing a trading monopoly position in this space, setting up a formidable barrier against the entry of new service providers in this area of the market. Such a scenario readily admits the position of monopoly-based service price setting. Compounding this is the risk that the providers set up their own "title office", so that in effect the major trading block actually controls the only means of establishing legitimacy of "ownership", which in terms of risk of anti-competitive trading practices is a very seriously damaged outcome.

Assuming a relatively low cost of achieving significantly higher efficiency address utilisation than at present, then the resultant market is bounded only by the costs of agility of renumbering. Here renumbering would be anticipated to occur in response to acquisition of a different address block in response to changing local address requirements, and the frequency of renumbering may occur in cycles of duration between weeks and years. Markets would also be constrained by deployment costs, where local address trading within a provider domain would have little cost impact on deployment services (as the aggregated routing scenario would be unchanged for the provider and the provider's peers) whereas trading in small sized blocks across provider domains would result in increased operational service cost due to increased routing costs (where efforts to create aggregated routing entries are frustrated by the effects of address leakage into other routing domains).

In examining this consequent environment the major technical outcome is strong pressure for dynamic host address assignment services, where the connection and disconnection of hosts into the Internet environment will cause a local state change in allocated addresses (which may in turn trigger consequent extended dynamic renumbering

from time to time to accommodate longer term address usage trends). It is also reasonable to predict a strengthening market for dynamic address translation technologies, as an alternate client strategy to the purchase of large address blocks from the trading market (this scenario is the use of a private, potentially non-unique address space within the client network, and the dynamic translation of end host addresses into a smaller unique Internet routed address pool to support external end-to-end sessions), and also the strengthened market for firewall boundary technologies which also admit the use of private address space within the client domain.

While it is not possible to accurately predict specific outcomes, it would appear to be the case that increasing overall efficiency of address utilisation will be most visible only after unallocated address pool exhaustion has occurred, as there is then a consequent strong economic motivation for such activity across all the entire Internet address space.

As perhaps a cautionary comment regarding evolutionary technologies for IPv4, it would also appear to be the case that evolutionary technologies will not assume a quantum increase in economic viability simply because of unallocated address pool exhaustion. Such technologies will only lever additional advantage over IPv4 once the marginal cost of increased IPv4 address space deployment efficiency exceeds the marginal cost of deployment of new technologies, a situation which may not occur for some considerable time after unallocated address pool exhaustion.

3. Modification of Current Internet Address Management Policies

The three major attributes of the current address allocation procedures from the unallocated pool are "first come first served" (FCFS) and allocation on a "once and for all" (OAFAs) basis, and the absence of any charge for address allocation (FREE).

As noted above, the outcomes of such a process, when constrained by the finite quantity of the resource in question, ultimately leads to a secondary market in the resource, where initially allocated resources are subsequently traded at their market valuation. This secondary trade benefits only those entities who established a primary position from the unallocated pool, and it is noted with concern that the optimal behaviour while the unallocated pool exists is to hoard allocated addresses on the basis that the secondary market will come into existence once the pool is exhausted. Such a market does not benefit the original address management operation, nor does it necessarily benefit the wider community of current and potential interested parties in the Internet community.

It is also noted that the outcome of a free address allocation policy is the vesting of the management of the address space to the larger Internet Service Providers, on the basis that in the absence of end client address allocation charging policies which have the capability of ensuring an independent address management function, those entities who have the greatest vested interest in the quality of the address allocation and registration function will inevitably fund such an operation in the absence of any other mechanism. The risk within this scenario is that placing the major asset of any communications medium into the sphere of interest of the current entities trading within that medium acts to increase the risk of anti-competitive monopolistic trading practices.

An alternate address management strategy is one of allocation and recovery, where the allocation of an address is restricted to a defined period, so that the allocation can be regarded as a lease of the resource. In such an environment pricing of the resource is a potential tool to achieve an efficient and dynamic address allocation mechanism (although it is immediately asserted that pricing alone may be insufficient to ensure a fair, equitable and rational outcome of address accessibility and subsequent exploitation, and consequently pricing and associated allocation policies would be a normative approach to such a public resource management issue).

It is noted that pricing as a component of a public resource management framework is a very common practice, where price and policy are used together to ensure equitable access, efficient utilisation and availability for reallocation after use. Pricing practices which include features of higher cost for larger address blocks assist with equitable access to a diversity of entities who desire address allocation (in effect a scarcity premium), and pricing practices can be devised to encourage provider-based dynamic address allocation and reallocation environments.

In the same fashion as a conventional lease, the leasee would have the first option for renewal of the lease at the termination of the lease period, allowing the lease to be developed and maintain a market value. Such pricing policies would effectively imply a differential cost for deployment of a uniquely addressed host with potential full Internet peering and reachability (including local reachability) and deployment of a host with a locally defined (and potentially non-unique) address and consequent restriction to local reachability.

It is also observed that pricing policies can encourage efficient address space utilisation through factors of opportunity cost of unused space, balanced by the potential cost of host renumbering practices or the cost of deployment of dynamic address allocation or

translation technologies.

There are a number of anticipated outcomes of a management mechanism which including pricing elements for the IPv4 address space

Firstly current address space utilisation projections (anticipated useful lifetime for the pool of unallocated addresses) would extend further into the future due to the factors of cost pressure for more efficient address utilisation, and the additional cost of issuing a local resource with a globally unique address and the opportunity cost of extravagant use of global addresses with purely local domains.

Secondly dynamic host address binding technologies, and dynamic network address translation technologies would be anticipated to be widely deployed, based on the perceived cost opportunities of using such technologies as an alternative to extensive static host address binding using globally unique addresses. Use of such technologies would imply further extension of the lifetime of the address pool.

Such pricing practices could be applied on a basis of all future address allocations, leaving those entities with already allocated address blocks outside of the lease mechanism. Alternatively such previous allocations could be converted to leases, applying a single management policy across the entire address space and accordingly leveraging the maximal benefit from such pricing policies in terms of maximising the lifetime of the address space and maximising the value of the address space. In such a situation of conversion some level of recognition of previous implicit OAFSA allocation policies can be offset through delay of conversion to lease and also through conversion of such previously allocated addresses to the lease, waiving the lease purchase costs in such cases.

4. Internet Environment Considerations

Pricing for IPv4 addresses as a component of the overall address management framework is by no means a novel concept, and despite the advantages such pricing policies may offer in terms of outcomes of efficiency of utilisation, fair and equitable access, security of allocation and consequent market value, and despite the address pool exhaustion time offsets such policies offer, it is the undeniable case that no explicit pricing policies have been successfully introduced into the Internet address allocation processes to date.

There are two predominate reasons offered in this analysis. The first is the somewhat uncertain nature of the exact origin of primary ownership of the IPv4 address space, and the unallocated address pool in particular. The address pool has been administered according to

policies drafted by the Internet Assigned Numbers Authority (IANA). The policies drafted by IANA are effectively policies which are the outcome of the same consensus seeking approach used within the Internet Standards process, and it is noted that within such an environment unilateral declarations of ownership and related assertions of policy control have difficulty in asserting an effective role within the Internet community and such declarations are generally incapable of gathering consensus support (It can be argued that "ownership" is not a relevant concept within this domain, as the essential attribute of such address elements are their uniqueness within the global domain, and such an attribute is only feasible through common recognition of a coordinated and reliable management environment rather than the historical origin of the resource in question). Secondly there is no formal recognition of the address space as being a shared international resource which sits within the purview of national public resource management policies and administrative entities of each nation, nor is there a recognition of the address space as a private resource owned and administered by a single entity.

Recent policy changes, whereby large segments of the unallocated address pool have been assigned to international bodies on a regional basis, with further assignment to bodies within national contexts, have been undertaken with a constant address allocation policy of FCFS, OAFAs and FREE, and although some effort has been made to increase the deployment efficiency through explicit allocation policy enumeration, the general characteristics of address allocation are unchanged to date (those characteristics being of course FCFS, OAFAs and FREE).

One potential scenario is to speculate that pricing processes imposed by the address allocation agency are not feasible within the current Internet environment to the extent that any such policies could significantly motivate increased address deployment efficiency to the levels required for longer term unallocated address pool lifetime extension. The lack of capability to employ pricing as a managerial mechanism, even to the extent of cost recovery of the allocation and subsequent registry maintenance function has a number of possible longer term outcomes:

- a) such functions will be restructured and operated from duly authorised national administrative bodies for each nation. Here the observation that the address pool delegation sequence within the current Internet environment has not to date been aligned with recognised national public communications resource administrative entities is an expression of the major problem that the unallocated address pool is not recognised as being intrinsically the same public resource entity as the radio

spectrum or the telephone number space. The consequence of this mismatch between existing public resource management structures and IPv4 address space management implies that public operation for this activity on a national basis is not a commonly observed attribute. The competency of such established public resource management structures in managing what continues to be a remarkably vibrant and dynamic technology-influenced domain must be questioned. Potential outcomes may possibly include a rational and equitable address space management mechanism, but would also in all probability include a cost of a heavy damping factor on further technological innovation and refinement of the underlying technology base upon which the address space is sited as a longer term outcome.

- b) such functions are operated (and/or funded) by Internet Service Providers. This is a more common scenario at present in the Internet IPv4 environment, and although such an operational environment does admit the potential for adequate funding for competent administration of the operation, the strong association of these entities who have established interests in the operation of enterprises based on the provision of services across the address space (i.e., strong interest in exploiting the address space) has a natural tendency to express domination of the market by established interests, threatening fair access to the common resource and threatening the open market of deployment of the technology. It is reasonable to suggest that such alignments are undesirable from a public policy perspective.
- c) such functions are inadequately funded to service the level of activity, and / or administrated informally and consequently managed poorly, and the essential attribute of reliable address space management is not achieved.

It is noted that these issues are largely unresolved within the Internet community today, and tensions between established and incoming Internet Service providers over equitable access to the unallocated address space pool are a consequent risk.

5. Concluding Observations

In the absence of the capability to price the management of the Internet address space at administrative cost levels, let alone the capability to set pricing of address leasing at prices which reflect the finite nature of the resource and reflect (even in part) the market value of the resource, as a component of overall common address management practices, the most likely scenario is a

continuation of the FCFS, OAFA and FREE address management policies until exhaustion of the unallocated address pool occurs.

It is perhaps a sad reflection of the conflict of short term objectives and longer term considerations that the evident short term motivations of ready and equitable access to the IPv4 address (which were the motivational factors in determining the current Internet address allocation policies) run the consequent risk of monopoly-based restrictive trade and barrier-based pricing as a longer term outcome of unallocated address space exhaustion.

While free address allocation and the adoption of policies which include pricing components both ultimately produce an outcome of strong pressure for increased address space utilisation efficiency, the removal of the neutral presence of the unallocated address pool does induce considerable risk of open market failure within the Internet itself if free address allocation policies continue until pool exhaustion has occurred.

Further strengthening of the current FCFS, OAFA and FREE address allocation policies, in an effort to induce higher address utilization efficiencies across the remaining address space is not a viable address management strategy refinement, in so far as the trading market will then commence before unallocated pool exhaustion, trading in large address blocks which are precluded from such strengthened address allocation policies.

The most negative aspect of this are is that these processes will erode levels of confidence in the self regulatory capability of the Internet community, such that significant doubts will be expressed by the larger community the Internet process is one which is appropriate for effective formulation of common administrative policy of one of the core common assets of the Internet.

These outcomes can all be interpreted as policy failure outcomes.

The seriousness of these outcomes must be assessed in the terms of the anticipated timeframe of such policy failure. Current expectations of unallocated address pool lifetime of 6 - 12 years does allow the Internet community some time to revisit their methods of administrative process definition, but this observation is tempered by the IPv6 process and by increasing levels of pressure on the address space in terms of growth in address demand through growth of deployment of the Internet itself.

It is perhaps an appropriate conclusion to acknowledge the impediments of existing processes to admit any significant process or policy change that would produce a more efficient and effective

address space management regime.

However it is this policy failure to efficiently utilise the IPv4 address space through inadequate address pool management policies, rather than the exhaustion of the pool per se which is perhaps the driving force to design and deploy an evolutionary technology to IPv4 which possesses as a major attribute a significantly larger address space.

It is also appropriate to conclude that any outside observer of the IPv6 refinement process will look to see if there is any evidence of experiential learning in address management policies. If there is to be a successor technology for IPv4 it would be reasonable to anticipate that associated address pool management mechanisms show a greater degree of understanding of public resource space management capability in the light of this experience. If no such evidence is forthcoming then there is no clear mechanism to instil sufficient levels of consumer and industry confidence in such technologies in such a way which would admit large scale public deployment, irrespective of the technical attributes of the successor technology. Such potential mechanisms may include pricing components irrespective of the actual size of the address resource, given that the number's uniqueness is a resource with inherent market value irrespective of whether scarcity pricing premiums are relevant in such an address space.

It is also appropriate to conclude that continuation of current address space management policies run a very strong risk of restrictive and monopoly-based trading in address space, with consequence of the same trading practices being expressed within the deployed Internet itself.

The immediate action considered to be most appropriately aligned to both the interests of the Internet community and the broader public community is to examine Internet address space management structures which include pricing as well as policy components within the overall management mechanism, and to examine the application of such mechanisms to both the existing IPv4 address space, and to that of any refinement or successor Internet technology base.

6. References

- [1] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1466, Merit Network, Inc., May 1993.

7. Security Considerations

Security issues are not discussed in this memo.

8. Author's Address

Geoff Huston
Australian Academic and Research Network
GPO Box 1142
Canberra ACT 2601
Australia

Phone: +61 6 249 3385
Fax: +61 6 249 1369
EMail: Geoff.Huston@aarnet.edu.au

