

November 2008

Geoff Huston

The Changing Foundation of the Internet: Address Transfers and Markets

All this talk of address exhaustion in IPv4 has prompted some reaction from the address communities in terms of the consideration of various forms of a "transfer policy for IPv4" in the APNIC, ARIN and RIPE regions. IN this article I'd like to look at these transfer policies in a little more detail, and take the opportunity to share some of my thoughts about the more general situation we find ourselves in with IPv4.

I must admit at the outset some level of personal involvement in this topic, as I'm the author of the IPv4 address transfer policy proposal in the APNIC region (<http://www.apnic.net/policy/proposals/prop-050-v003.html>), so I have been following the proposals and their discussions in each of the regions pretty closely to try and understand why the proposals differ, and what are their relative strengths and weaknesses.

The Transfer Policy Proposals

Each regional addressing community operates its own policy development process, and proposed policies are considered independently by each community. Sometimes the same policy proposal is submitted into each regional process. Sometimes the same material comes out as policy in each region, and sometimes different outcomes in terms of the detail of the policy are seen in each region. At other times different policy proposals are submitted into each regional process, and in the developmental process these policies may or may not converge. In the case of the transfer policies, the latter course of action was taken and different proposals have been submitted to each regional policy development process.

APNIC: Proposal-50

The APNIC proposal attempts to take a minimal position in terms of policy-based constraints on transfers of IPv4 addresses. The proposal is that APNIC will enter into its registry the transfer of an IPv4 address block between two parties who are "known" to APNIC by virtue of being members of that organisation. There are few other constraints on the parties to an address transfer, namely that the disposer of the transferred address block cannot subsequently request an IPv4 address allocation from APNIC for 24 months following the transfer, and that is assuming that APNIC still has IPv4 addresses to allocate at that time, and the recipient of the transfer will have this address block counted in their holdings in terms of meeting the utilization criteria for any further IPv4 address allocations, assuming again that APNIC still has IPv4 addresses to allocate at that time. The only other constraint is that the transferred address block must be no smaller than a /24. And that's it.

ARIN: 2008-02 and 2008-06

The IPv4 address transfer policy proposal currently being considered by ARIN (http://www.arin.net/policy/proposals/2008_2.html) has been refined progressively over the past few months. In its original incarnation this proposal contained a considerable set of constraints, related to the underlying concept here that a "transfer" was interpreted as a "return to the registry" by the source of the address block, followed by a "allocation by the registry" to the recipient of the address block. The recipient of the transfer had to meet the prevailing ARIN IPv4 address allocation criteria in terms of the demonstration of need for addresses, efficiency of use of their current holdings. Both parties were constrained for undertaking any further address transfers for an extended period of time. The motivation for these constraints were based around concerns of the potential for distortions in any address market that may arise and concerns related to excessive fragmentation of the routing system arising from unconstrained subdivision of address blocks as a precursor to optimising the disposer's position in terms of value of the transferred addresses. Subsequent discussion in ARIN has seen this proposal drop many of these constraints, but the concept that the acquirer must meet the prevailing IPv4 address allocation criteria, including demonstration of need and efficiency of utilisation, has remained across these all refinements. A more recent policy proposal in the ARIN region (http://www.arin.net/policy/proposals/2008_6.html) drops all of these constraints except that of the recipient of the transfer having to meet the "documented need" allocation policy, the imposition of limits of the extent of subdivision of aggregate address blocks, and a constraint that this policy would not come into force until the depletion of the general use unallocated address pool in the ARIN region.


RIPE: 2007-08

In the RIPE region there has also been some refinement of the original transfer policy proposal, but in the opposite direction (<http://www.ripe.net/ripe/policies/proposals/2007-08.html>). The original proposal was very similar to the APNIC proposal, taking a relatively minimal approach to the conditions applicable to the registration of address transfers, while a more recent version of this proposal has also adopted the constraint on the acquirer of the demonstration of need and address utilisation efficiency, similar to the ARIN proposal.

The Roles of the RIR's

If all three RIR's are facing the same situation, and all three RIR's have adopted approximately similar policy framework so far to undertake their role, then why the divergence of approach to this particular issue of address transfers with the three RIR's policy forums?

One way to answer this is to look at the RIR's current roles.



RIR's perform three roles:

- they are allocation entities who distribute address blocks,

- they are regulators, or at least de facto regulators, in that the allocation of address blocks is accompanied by the imposition of certain constraints and regulations applied to recipients, and
- they act as registry operators, maintaining a public registry describing the disposition of allocated addresses.

The de facto regulatory role is of interest here. The RIR's have embraced a broad agenda that has associated their allocation role with goals that span not only effective stewardship of the address space, but also to encourage industry conservatism in the use of these resources. They have also embraced an agenda of attempting to limit extravagant growth in the routing table through adoption of the principles of provider-based addressing, minimum address allocation sizes, and allocation practices that attempt to preserve aggregatability of addresses.

The RIR's also adopted the position of claiming that addresses are not property, cannot be sold or transferred, and in and of themselves have no intrinsic monetary value. If the address holder has no further use for currently held addresses then they are to be returned to the registry.

The RIR's have operated using the principle that addresses are made available to all within the general constraint of the demonstration of need, preference for use of provider-based addressing, aggregation in routing, and efficient use of all previously acquired addresses, and having addresses be made available at relatively nominal price to the recipient.

The RIR's are also de facto monopolies in this allocation function. Unique address blocks are distributed from the IANA only to the RIR's for use in the Internet, and each RIR exclusively services a particular region of the world. The RIR's do not compete for members and their members do not themselves compete with the RIR or between themselves in terms of the address allocation function.

The abundant pool of addresses and their ready availability at relatively nominal cost, coupled with the regulatory stance that denied the characterisation of addresses as property, the prohibition of general forms of address transfers, and their de facto monopoly role in address allocation effectively suppressed the formation of any form of aftermarket for addresses.

The major enforcement mechanism within this structure was one based on the matter of further allocations of addresses. Each RIR allocated addresses to meet a parties anticipated growth requirements for the following 12 months, effectively ensuring that each active entities address holding would be reassessed at regular intervals, and failure to comply with the RIR's policies would lead to denial of access to further address allocations.

The authority used to undertake this role was by virtue of the open, transparent and consultative process of formation of these policies, resulting in a form of industry self-regulation that was eminently suited to the characterization of the Internet itself as a child of the progressive deregulation of the public communications industry.

It is not the intent in this article to formally to assess the effectiveness of these regulatory constraints, nor to measure the outcomes in terms of address utilization efficiency, routing table size, the extent of low cost availability of addresses across a diversity of economies and the efficiency of the allocation and management function as a common overhead, but an informal assessment appears point generally to its effectiveness on all of these counts.

The Impact of IPv4 Address Exhaustion

The current RIR framework for the distribution of Ipv4 address relies on the relative abundance of addresses.


The prospect of having future requests satisfied implies that there is no potential gain in abusing the process by attempting to request more addresses than would meet the current needs. As addresses cannot be on sold they have no residual value and there is no benefit to hoarding and no speculative outlook of escalation in the value of addresses. In addition, operating counter to these RIR polices imperils the ability of the entity to have access to further RIR address allocations.

The problem with a fixed length address pool is that the size of the pool is finite and, has been explored in many ways in recent times, the current expectation of the continued availability of addresses only spans a further 24 months from today (<http://ipv4.potaroo.net>). Abundance of IPv4 address will draw to a close at that point, and the allocation function of the RIR's will draw to a close at the same time.

As explored in the previous article on this topic (<http://www.potaroo.net/ispcol/2008-10/v4depletion.html>), the prospect of address exhaustion is not unexpected and the entire development of IPv6 was the IETF's response to this envisaged situation. The intended plan was that industry would have undertaken the process of transition to IPv6 before the IPv4 address pool was exhausted, and that the overall RIR operational framework of allocation, registration, and regulation would smoothly shift from IPv4 to IPv6 without undue disruption to the industry, or to the RIR's themselves.

But that has not eventuated, and we are now facing the prospect of continuing demand for IPv4 well past the exhaustion of the unallocated address pool, and well past the termination of the RIR's current IPv4 address allocation function. It is this prospect that is driving the development of these transfer policies within the RIRs' policy forums. The basic motivation of all of these transfer policy proposals is that once the allocation supply has exhausted, the only other source of supply to meet continuing demand will come from already allocated addresses, and the function of allocation will be replaced by the operation of transfers.

The implication here is that the cessation of the allocation function of the RIRs requires a reappraisal role of the regulatory role of the RIRs.



The Evolution of the Transfer Policy Proposals

So why do we need a transfer policy at all?

Even though there is no further allocation function to be performed following the point of cessation of the IPv4 address allocation function, there still remains the registry function, and that function is of critical value to the Internet.

Regardless of how addresses are distributed and what policy framework is used to perform address distribution, addresses must be distributed uniquely. The manner of this unique association between an entity and an address block is a record in the public address registry. Once the allocation function has finished the residual requirements for RIRs' policies and associated regulatory role rests within the need to maintain an accurate, useful and complete registry of current address holdings.

But that observation does not explain why there are such differences in these proposed policies, since, after all, the respective RIR's are really no different to each other in terms of roles and responsibilities.

My personal view of the reason for these differences lie in a belief that there is a certain amount of confusion and uncertainty about the future role of the RIR in a world that has replaced abundance of addresses with acute scarcity.

The correlation between the prevailing address allocation policies and the characteristics of the efficiency of address use, the efficiency of the routing system and the scalability of the Internet itself was recognised

very quickly within the over all time frame of the Intern's development. The RIR system was established not only on principles of fairness of access and efficiency of the allocation function but also on the principle that the allocation system was to be used to support the scalability of the Internet, by ensuring that all the addresses in use were used efficiently, and that the routing system was not going to collapse under the pressures of excessive fragmentation of the address space.

Enforcement of Policies in a Deregulated Environment

While these motivations were the guiding principles of the allocation policy framework, adherence to these policies by users of addresses is by no means uniform. In a deregulated environment of diverse players where there are no coherent, uniform enforcement mechanisms sees players tending to observe only those policies that are aligned to their perception of how to maximise self interest.

Efficiency of address use and comprehensive accuracy of maintained registration data are directly linked to the ability of an entity to obtain further addresses as the need arises. Compliance of these particular policies has been generally observed. However, fragmentation of the routing space has not been directly linked to the further allocation function, and the results of this decoupling of policy with a risk of any negative outcome is clearly evident in the continuing fragmentation observed in the routing space.

The observation is therefore that policies that are effective in this environment are limited to those policies which are clearly linked to the allocation function, since the primary motivation of the entities serviced by the RIR system is access to addresses, and the RIR's form a de facto monopoly in terms of this allocation function.

From Abundance to Scarcity

The forthcoming termination of the address allocation function has led directly to a degree of role and authority confusion. it appears that the RIR policy forums are still strongly convinced of the correctness of the principles used to guide the RIR's policy development process, and they appear also to be also believe in the adoption of associated RIR operational practices that formed the address distribution framework.

For example, "addresses are not property and cannot be purchased or sold," is one of the more common mantras of this community. Similarly we hear, "all addresses should be in use in networks, and any unused addresses should be promptly returned for re-allocation.

The confusion of role arises when these same forums consider what the appropriate policies should be when the allocation function has terminated. Attempting to apply the same principles to the registry function and the associated policies related to making changes to this registry is where the confusion lies.

Here previous practice becomes confused with the generic principles, and the differing characteristics of allocation and registry functions, and the de facto monopoly in one role and not the other, are not clearly distinguished.

For example, the French governmental delegation at the June 2008 ICA yNN meeting loudly proclaimed their policy position at the opening plenary session of abhorring the formation secondary markets in addresses, claiming they were "unfair". Similar critical comments have been made by others on these transfer proposals, based on the strong belief that addresses are not property and their title simply cannot be transferred between consenting parties, and that markets in addresses are an intolerable outcome.

This situation raises a number of quite challenging questions about addresses and their distribution When the unallocated aPv4 address pools are empty, and when there is still strong demand for further IPv4 addresses, what is a "fair" way of meeting that demand? What if demand lies in poorer predominately developing economies while unused holdings may tend to sit in the richer parts of the world? Should the poor be forced to buy from the rich? What is fair? Is the monetisation of addresses and the emergence of a market a fair outcome in an environment of a global communications system? Should we look instead to forced resumption and selective reallocation of addresses? Which body would be in a position to force

resumption in terms of authority in a legal and regulatory sense? Would such a measure fit comfortably within the general deregulated framework of today's communications industry? Who is in a position to identify the most deserving recipients and the least deserving recipients?

There appear to be no perfect and universally acceptable answers here.

This consideration has led to an [at times] reluctant acceptance of a registry transfer policy as being a necessary, and hopefully temporary, measure, and a related concession that such a transfer policy and the registry itself underpins the operation of any associated market, or markets, in addresses by proving clear title to the goods being traded in such a market, namely addresses. But if that is the case why have transfer policies at all? Why not simply allow the registry to record all outcomes of any transfers between consenting parties without further qualification or constraint?

Transfer Policies and Registries

Much of the motivation for transfer policies for registries lies in the observation that most, if not all, efficient, fair and open markets are regulated in some manner. Markets are intended to establish a dynamic equilibrium between supply and demand using a pricing function as the mechanism of mediation, but markets are seldom "perfect". Markets are subject to various manipulative and distortive pressures, including efforts directed at price manipulation, speculative acquisition, fraudulent transactions, misrepresentation, withholding of information, and deliberate leaks of false information.

In the context of a market in addresses all such risks apparently apply, but the means to eliminate [or "address" if you will pardon the painful pun] these risks are not so clearly obvious. It is also the case that such market distortions would impair or compromise the cherished principles of address administration—efficient use, fair and equitable access, an avoidance of profiteering and extortion.

In much the same manner if all you have is a hammer then everything else starts to look like a nail, if all you can regulate is an address registry and if the problem is risks of distortions in an address market, then the temptation to apply your only available mechanism to the problem is often overwhelming.

So what is evident as the basic motivation for the constraints in the transfer policies is the effort to mitigate some of the more odious risks of distortions appearing in a market for IPv4 addresses by placing strict regulatory constraints on access to the address registry.

And that is the form of response in the original ARIN transfer policy proposal (http://www.arin.net/policy/archive/2008_2_orig.html). The proposal was replete with measures designed to deny access to the registry for those transactions that were potentially speculative in nature, and to deny transactions that attempted to corner the market in addresses through a combination of acquisition and withholding. And it is certainly evident that that such an approach could act as a deterrent to those particular forms of market distortions.

The RIPE transfer policy proposal has been edited to head in a similar direction to the ARIN transfer policy proposal, requiring the acquirer of a transferred address block to meet the criteria of a further allocation by RIPE, including provisions of demonstrated need and efficient utilization of existing address holdings.

The RIPE Transfer policy proposal includes the constraint:
"An LIR may only receive a transferred allocation after their need is evaluated by the RIPE NCC, following the policies set for receiving further allocations within RIPE region (see the Section 5.3 Additional Allocations of this document)."

RIPE Policy Proposal 2007-08, version 3.0

However, this approach of placing such constraints on the recording of an address transfer in the RIR-operated address registries raises the fundamental question here of whether this RIR-operated registry function is capable of sustaining such a weighty agenda. In a diverse and deregulated industry sector, such as the Internet, will the imposition of onerous restrictions on access to the address registry impose adherence to the constraints, or prompt the emergence and use of alternate registries that do not impose such constraints? After all, while the address allocation function was a de facto monopoly function, the same cannot be said for the registry function, and the general compliance with registry policies, particularly with potentially onerous registry policies, is not necessarily a certain outcome.

An alternative outcome to the effort to impose such constraints on the operation of one set of registries is that alternate registries appear that do not impose such constraints on the registrants. The reason behind this lies in the observation that the selection of a registry, and the derivation of authority of the registry operator, is more based on common convention than by external imposition.

There is no comparable de facto monopoly in the registry role.

The registry is public, so it is not possible to occlude the contents of the registry from public view. This implies that cloning the registry in some form or fashion is potentially possible at any time. If the recording of otherwise legitimate transactions is excluded from the original registry, the possibility exists that alternate registries may appear to meet that demand. The original registry is left with an information set that is historical rather than current and, in the face of the emergence of variants, is then incomplete and inaccurate. Anyone wishing to understand the actual current disposition of addresses would have to consult these alternate registries in order to understand if the address has subsequently moved to another holder.

In and of itself this could be seen as just another expression of a competitive market for the supply of goods and services in action within a deregulated industry. If one registry operator no longer meets the demand of users of the registry's services than these users are at liberty to use an alternate provider of registry services that is able to meet their needs.

But is such an outcome desirable? Or even tenable for the Internet. Alternate points of authority relating to the disposition of addresses is an attack at the one essential property of addresses that we simply must preserve, namely the uniqueness of address holdings. When multiple registries all have different entries relating to the current holder of the same address, then how actually is the current holder of that address? Where should routing point to? Why? In this case diversity of choice in address registries with deliberate overlap would be a catastrophe for the Internet.

An alternate perspective has motivated the APNIC policy proposal, namely that the registry function is incapable of sustaining a substantive regulatory agenda. The proposal being put to the APNIC policy forum offers a minimal set of constraints that are intended to be as neutral as possible and to ensure the continued acceptance of this particular registry as the region's authoritative address registry. This approach does not necessarily condone poor market behaviours but deliberately avoids attempting to use the registry operation in novel and untested ways to regulate market behaviour.

Registries are, in effect, title offices, and their function is simple: they record who holds what resource. Title offices need to be respected by all parties, and accessible for all parties in order to ensure that the registry itself reflects the true disposition of the resource. Not only a registries essential to sustain the integrity of operation of any market, but in the case of address registries and the Internet, address registries are essential to sustain the coherent operation of the Internet. Registries are not an eternally imposed monopolistic constraint, and registries derive their authority through common acceptance by all parties and by conventional use and accessibility. This is generally best achieved if, like a title office, a registry is neutral in its operation. Attempting to overload the access to a registry with what is in effect a regulatory agenda is a course of action with a high risk of failure.

That does not imply that markets are perfect nor that market distortions should be tolerated simply because the registry is incapable of acting as a regulator of market behaviours. Neither is the case. But the

observation should be made that markets in all manner of good and services have a rich history in human societies, and the role of the regulator is similarly one with a rich history. Market regulators exist in many guises and in many regimes, and can exert influence through various direct and indirect means. There is no need to believe that this issue of transfers and address markets requires a comprehensive solution based solely on the resources, capability, authority and enforcement authority of the RIR's. Once the RIR's are no longer address allocation entities much of their ability to enforce certain behaviours goes with that role, and the residual role, that of operation of an address registry, necessarily has to take a more neutral stance if it is to be a role that is discharged effectively.

The observation I'd like to conclude with is that what seemed superficially to represent the most conservative approach to the issue of exhaustion of IPv4 unallocated addresses, namely the preservation of as much as possible of the existing address allocation policy framework within the transfer policy of a post-exhaustion IPv4 Internet, actually represents a radical departure from a more conventional view of appropriate registry practices, and tends to offer the Internet the highest risks of failure. For registries the critical consideration is that neutrality, accessibility, openness and availability are the essential attributes that ensure its continued relevance and value. Overloading the operation of a registry with constraints that are, in effect, regulatory provisions, creates continuing risks to the acceptance of the registry as a trusted neutral resource that works in the common interest of all parties and in the interests of the Internet itself. A registry is not an appropriate mechanism to enforce regulatory constraints.

After all, when all you have is a hammer, then you really need to appreciate that some problems will get far worse if you try and hit them on the head, and that a fully stocked hardware store is never too far away!

References and Further Reading

The Transfer Policy Proposals:

APNIC: Prop-50 <http://www.apnic.net/policy/proposals/prop-050-v003.html>
ARIN: 2008_2 http://www.arin.net/policy/proposals/2008_2.html
2008-06 http://www.arin.net/policy/proposals/2008_6.html
RIPE: 2007-08 <http://www.ripe.net/ripe/policies/proposals/2007-08.html>

Papers:

Benjamin Edelman, "Running Out of Numbers: The Impending Scarcity of IP Addresses and What To Do About It." Unpublished paper draft, 2008.
<http://www.benedelman.org/publications/runningout-draft.pdf>

Milton Mueller, "Scarcity in IP addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries." 2008.
http://www.internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf

William Lehr, Tom Vest, and Eliot Lear, "Running on Empty: the challenge of managing Internet addresses," 2008.
http://tprcweb.com/files/Lehr%20Lear%20Vest%20TPRC08%20Internet%20Address%20Running%20on%20Empty_0.pdf

Hillary Elmore, L. Jean Camp and Brandon Stephens, "Diffusion and Adoption of IPv6 in the ARIN Region," 2008.
http://tprcweb.com/files/IPv6%20in%20our%20Lifetime9_4.pdf

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

About the Author

GEOFF HUSTON is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. he graduated from the Australian National University with a B.Sc, and M.Sc. in Computer Science. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

<http://www.potaroo.net>