Running on Empty: the challenge of managing Internet addresses

William Lehr (MIT)[1]

Tom Vest (Eyeconomics.com)[2]

Eliot Lear (Cisco)[3]

ABSTRACT

Over a quarter century has passed since IP addresses from the Internet's now globally ubiquitous (IPv4) address pool first began to be distributed to and used by operators of the 28,000 and counting independent routing domains that constitute today's Internet. At current demand rates, the total pool of IPv4 addresses is expected to be fully allocated within the next 3-5 years. Despite the looming exhaustion of the address pool, many Internet network operators have exhibited reluctance to transition to the successor address resource pool (IPv6). In light of the looming scarcity of IPv4 addresses, there is growing interest in the potential for market-based allocation mechanisms to induce more efficient use of the IPv4 address space and to smooth the transition to IPv6. Some of these proposals include allowing existing allocations to be reclaimed or traded to supplement the supply of IPv4 addresses to meet continuing demand.

This paper will provide an appropriate contextual framing, suitable for a multidisciplinary audience, of the ongoing discussion of how market forces might be introduced to enhance the management of the Internet address space. We discuss some of the proposals currently under consideration and highlight the key points of contention. While we conclude that market-based incentives can play an important role, how they are implemented needs to be carefully considered in light of the requirements for Internet routing, the politics of global Internet governance, and the likely implications on the transition to IPv6. For example, we believe inappropriately allowing unrestricted trading of IPv4 addresses could have the perverse effect of increasing the aggregate costs of Internet routing, retarding the migration to IPv6, and thereby adversely impacting Internet growth and architecture. On the other hand, while the existing framework for Internet address management has served reasonably well to date, retaining the status quo is likely to result in a burgeoning black/grey market in IPv4 addresses and similar unfavorable outcomes with respect to aggregate costs and Internet architecture. Further research and collective discussion is needed to identify how best to forge a middle road toward improved Internet address management. The discussion provided herein is intended to help frame and seed the debate.

---

[1] William Lehr, MIT, email: wlehr@mit.edu

[2] Thomas Vest (Eyeconomics.com) is a consultant to RIPE NCC Science Group, email: tvest@ripe.net. The author gratefully acknowledges the support of RIPE NCC in making this research possible. Observations and opinions expressed herein are the author's own, and in no way reflect the policy or positions of the RIPE NCC, its member institutions or staff.

[3] Eliot Lear, Cisco, email: lear@cisco.com

Table of Contents

1 Introduction[4]

The Internet is now essential communications infrastructure for our global economy. The space of IP addresses that uniquely identify devices/nodes on the Internet, enabling packets to be successfully routed from source to destination nodes, is a core component of this infrastructure.

The existing address space based on IPv4 is nearly fully allocated (81%).[5] This raises a question of potential scarcity for Internet operators, equipment and application designers, and end-user communities that want additional Internet addresses. The most widely accepted long-term solution to this problem is to transition to a much larger address space based on IPv6 addressing. However this transition is still in the early stages with no assurance of success while the tank of available unallocated IPv4 addresses is nearly empty.

Addressing the challenge of IPv4 scarcity has important implications for the evolution of Internet infrastructure and technology (e.g., maintenance of the end-to-end principle, network routing architecture, and the transition to IPv6), the allocation of benefits/costs among Internet stakeholders (i.e., incumbents v. entrants, ISPs v. end-users v. equipment/application providers), and public policy on a global scale. Moreover, these normative, distributive, and institutional considerations can only be weighed in light of the essential material limitations of the Internet's "routing system," which handles the exchange of data traffic between the hardware and software processes (e.g., access, content, and services) that constitute the Internet from the end user's perspective. Internet addresses are valuable because they can be routed. Because Internet routing depends on the structure of the IP address, IPv4 address scarcity and routing slot cost are interrelated. As a result some address management approaches may be unworkable no matter how appealing their non-technical characteristics might be.

Developing appropriate solutions to the IPv4 exhaustion challenge is inherently multidisciplinary because it engages the political dynamics of the existing Internet governance mechanism (i.e., the set of Regional Internet Registries that administer the IP address space), the industry economics of the value chain (i.e., the distributed set of global ISPs, upstream/downstream suppliers of equipment and applications, and end-users that ultimately will pay for and benefit from whatever solution is adopted), and technology (e.g., what next generation Internet infrastructure ends up looking like).

One possible response to the problem of scarcity is to turn to market mechanisms. Building on the discussions of participants at a recent Cisco-sponsored workshop[6] and the growing literature that has contributed to this debate,[7] this paper seeks to frame the discussion of the IPv4 scarcity problem for an appropriate multidisciplinary audience.

---

[4] The authors would like to acknowledge the support of Cisco Systems, and specifically David Meyer of Cisco, for their support and effort in organizing a workshop *Address Allocation Models for the Internet*, March 3-4, 2008. Additionally, Dr. Lehr would like to acknowledge the support of the MIT Communications Futures Program (http://cfp.mit.edu). The opinions expressed herein are those of the authors alone.

[5] http://www.iana.org/assignments/ipv4-address-space as of March 27, 2008. According to one analysis, the pool of unallocated IPv4 addresses will be exhausted by 2011 (see Huston, 2008).

[6] *Workshop on Allocation Models for Internet Addresses*, San Jose, CA, March 4-6, 2008.

[7] For example, see Perset (2007) for a thorough and extensive background report on the history of the IP address debate. Also, see Edelman (2008), Elmore, Camp, and Stevens (2008), and Mueller (2008) for analyses of the role of market mechanisms for Internet address management.

In Section 2, we discuss technical fundamentals necessary to understand the challenges of Internet address management and describe the history and current practices used to allocate addresses. In Section 3, we provide a generalized taxonomy for classifying management regimes and consider the potential benefits and risks of introducing market-based transfers of IPv4 addresses. In Section 4 we describe and analyze proposals currently under consideration. In section 5 we discuss an on-going technical development with the potential to disrupt the analysis presented here. And, in Section 6, we summarize our conclusions and suggestions for further work.

2 Understanding Internet Address Allocation

In the following sub-sections, we provide contextual background to understand why Internet addresses are useful, important trends impacting their use, and how their allocation has been managed to date.

2.1 What is an IP address and why is it valuable?

An IP address is assigned to every network interface that touches the Internet. The current IP version 4 ("IPv4") generation of addresses are 32-bit numbers with some embedded structure.[8] Without an IP address, a host cannot communicate with other hosts on the Internet. IP addresses are not valuable in themselves; indeed they have no meaning or use value in isolation. However, in the context of other IP addresses, they have tremendous use value: IP addresses are used to identify host computers to route packets from source to destinations across the Internet. Thus, as the Internet has grown, so too has demand for IP addresses.

With a fixed length of 32-bits, there are only about 4 billion unique IPv4 possible addresses.[9] When this format was selected, the Internet was still only an experiment running on a few hosts and it was already understood that should the Internet evolve into a large production sized network, a larger address space would be needed.[10] Nevertheless, many years later, with close to a billion users, we are still using the same 32-bit IPv4 addresses. Unfortunately, today, we are nearing exhaustion of the pool of unallocated IPv4 addresses, and current forecasts predict exhaustion by 2011.[11]

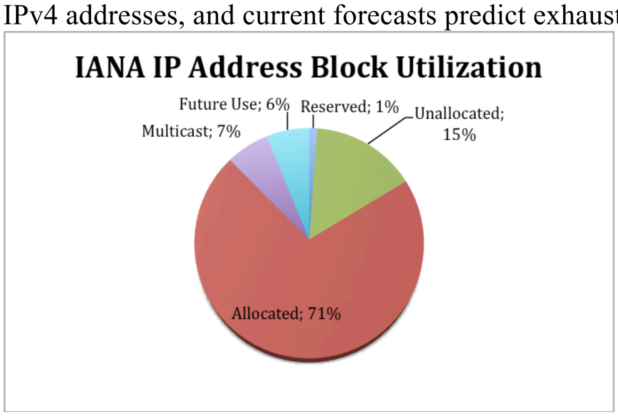

Figure I: Source: Internet Assigned Numbers Authority (IANA), May 2008.

---

[8] Postel, J., *Internet Protocol*, RFC 791, September, 1981.

[9] IPv4 addresses are 32-bit (4-byte wide) addresses of the form x.x.x.x (where x is integer in range 0 to 255). There are $2^{32} = 4{,}294{,}967{,}296$ potential addresses.

[10] Email exchange with Vint Cerf on June 4, 2008.

[11] See Houston (2008), "Projected IANA unallocated address pool exhaustion," available at: http://www.potaroo.net/tools/ipv4.

As mentioned earlier, an IP address contains structured information. A hierarchical allocation model on bit boundaries allows for the Internet routing system to aggregate individual hosts and small networks into larger blocks so that a single routing entry can represent many thousands of devices. The model is similar to that of a postal address, where distant postal services know only of the existence of a country or city, while only the local postal workers need to know of street addresses.

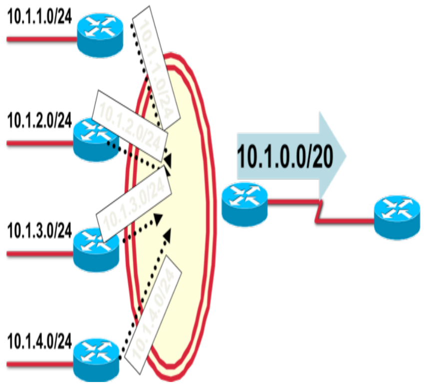

Figure II: Routing Aggregation Example

Figure II shows how a single route can represent multiple networks. In the example, four networks are announced through a routing protocol by respective routers, and a single router aggregates the announcement. The notation "10.1.0.0" represents an IP address of four eight-bit values. "/20" indicates the number of bits that should be applied to that IP address for purposes of routing. The fewer the number of bits, the larger the network represented.

Because of this hierarchical allocation model, routers that process packets on the Internet make a routing decision on a subset of the address. However, as allocations become finer grained in order to conserve address space, more routing entries are created on many routers on the Internet. Put in other terms, there is an economic cost to such routing entries that is borne by service providers and only indirectly passed back to end customers.

According to one estimate, the cost of routing the current number of Internet routes across all of the Internet routers costs on order $2B per year.[12] Thus, preserving the ability to aggregate Internet addresses is important for limiting routing table growth, which has a direct impact on Internet operating costs. In addition to these direct costs, there is the potential adverse impact of slower routing (as tables get bigger) or loss of Internet connectivity if ISPs, confronted with rising routing costs, decide to economize by refusing to route all announced addresses. Such rising costs could also exert considerable influence on the shape or structure of the Internet access industry (e.g., if the cost or complexity of routers capable of supporting routing services across the full Internet becomes so prohibitive that only the largest commercial enterprises can afford them).

Furthermore, because the estimated cost of announcing another route is relatively small for the party who announces the route (say, $0.04/route/router/year), yet large in aggregate (say, $6,200/route/year for all

---

[12] See Bill Herrin, "What does a BGP Route Cost?," February 2008, available at: http://bill.herrin.us/network/bgpcost.html (This analysis is cited in Edelman (2008)). Herrin estimates the cost conservatively at $0.04/route/router/year which suggests the cost is at least $6,200 per year for each announced route, and maybe as much as $12,000 for a rough average total per-route cost of $8,000. This estimate assumes there are approximately 245k routes to be announced and 150k routers that need to be updated. Interestingly, he argues that because IPv6 routes consume more resources (they are longer!), the cost for an IPv6 announced route might be $16,000 per year.

routers),[13] there is a large negative externality associated with inefficient routing of the address space. This provides an important justification for coordinating address management and Internet routing.

2.2 IPv4 Address Space Exhaustion

While the pool of unallocated IPv4 address blocks is dwindling, no one is sure precisely when it will be exhausted. Current estimates for when the Internet Assigned Numbers Authority's (IANA's) pool will be exhausted range from 2008 through 2012.[14] However, because IANA allocates blocks to Regional Internet Registries (RIRs) well in advance of those addresses actually being allocated to service providers, and because RIRs allocate addresses to service providers well in advance of those addresses being allocated to end customers, there will be a lag of uncertain time before the various pools of unallocated IPv4 addresses are empty and this may vary by region. Moreover, in anticipation of exhaustion, it is reasonable to expect that ISPs and end-users may change their behavior (e.g., by migrating to IPv6).

2.3 Mechanisms and trends impacting the pace of IPv4 exhaustion

The relative "scarcity" on IPv4 addresses is also impacted by a variety of business and technical innovations that have both increased demand for and opportunities to conserve address space. Some of these trends are discussed in the following sub-sections.

2.3.1 Private Address Blocks

In 1994 when the growth problem became readily apparent, IANA made available three blocks of address space for private use. Private use means that such addresses could not be routed onto the public Internet, but were instead for use within an end-user network for private purposes.[15] The envisioned use at the time was manufacturing, human resources, and other administrative computing services that would not normally be accessible on the Internet. However, even at the time there was an understanding that such divisions could not last, and the move has been considered controversial ever since.[16] Very quickly technology advanced to allow private address blocks to be masked with a small number of public addresses.[17] The boxes that effect such translation are called Network Address Translation or "NAT boxes." The price paid of deploying such a solution, however, is that computers that make use of private addresses are not capable of receiving unsolicited data (such as service requests) from the public Internet, and lower layer Internet functions are also unable fulfill one of their intended functions of providing for redundancy without service interruption. Never the less, private address blocks can be found nearly everywhere, and their use has assuredly retarded overall IP address utilization growth.[18]

The proliferation of NAT boxes, firewalls, and other boxes on the end-to-end path between host computers is regarded as a threat to the end-to-end (e2e) architecture of the Internet. However, others argue that such violations of the pure form of the e2e principle are practical and desirable, and would arise in any case as a consequence of growing concern to protect end-users from SPAM, security threats, and the like from a fully open Internet. We do not take a position on whether NAT boxes represent a good or bad trend for the future of Internet architecture. However, the need for such translation would diminish

---

[13] *Ibid.*

[14] http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html. We explain the role of IANA and the RIRs further below.

[15] Rekhter, Y., et al., *Address Allocation for Private Internets*, RFC 1597, March 1994.

[16] Lear, E., et. al., *Network 10 Considered Harmful (Some Practices Shouldn't be Codified)*, RFC 1627, July 1994.

[17] Egevang, K., Francis, P., *The IP Network Address Translator,* RFC 1637, May 1994.

[18] It is impossible to account for every use of private IP addresses, since there is no need for a central assignment authority, and their use cannot easily be detected within the Internet.

if the migration to IPv6 were successful. Thus, changes to mechanisms that alter the allocation of IP addresses will impact the future architecture of the Internet since they will impact the need for NAT boxes.

2.3.2 Dynamically Allocated Addresses

In the late 1990s, as demand for mass market Internet access expanded, two technologies were developed for assigning IP addresses to host computers so as to simplify access configuration: Point-to-Point protocol (PPP)[19] and Dynamic Host Configuration Protocol (DHCP).[20] A byproduct of their deployment has been a more efficient allocation of addresses. This is particularly true for service providers who allocate address space only to customers who are actively connected to their infrastructure. Dynamically allocated addresses have proven immensely important for mobile computing devices, such as WiFi-capable cell phones. However, because more devices can come and go with an individual than just a PC, and because workforces are more mobile, previously developed allocation models for IP addresses may no longer be appropriate.

2.3.3 IP Version 6

In anticipation of IPv4 address space exhaustion, an effort was undertaken to develop a new version of IP that supports a larger address space. The result was IP version 6 (IPv6), which is based on 128 bit addresses.[21] To date, while most computing platforms already support IPv6, adoption by ISPs has been quite limited, with approximately 1,400 prefixes appearing in global routing tables,[22] as compared to over 275,000 IPv4 prefixes.[23] Relatively few Internet resources are reachable for standalone operation of IPv6, making it a poor substitute for IPv4 today. Indeed, as long as most of the content, applications and users remain on IPv4, early adopters of IPv6 will incur added costs of maintaining interoperability with the IPv4 world, while deriving little benefit from IPv6. While most Internet architects may agree that conversion to IPv6 is both desirable and inevitable, there are significant interests that dispute both points. Indeed, one industry consultant characterized IPv6 as “the wrong technology and the wrong time” as indicated by the industry’s slow progress toward effecting a transition.[24] At the same time, a number of large ISPs and several governments have made IPv6 a priority by requiring its implementation and deployment.[25]

In this paper, we do not engage this debate, but assume that migration to IPv6 is socially desirable, while noting that there is a sizable contingent that may oppose this migration.[26] This assumption has several practical implications for our analysis. First, any proposal to modify the mechanism for allocating IPv4

---

[19] Perkins, D., *Point-to-Point Protocol: A proposal for multi-protocol transmission of datagrams over Point-to-Point links*, RFC 1134, November 1989.

[20] Droms, R., *Dynamic Host Configuration Protocol*, RFC 1531, October, 1993.

[21] With 128-bit addresses, there are potentially for $2^{128} = 3.403 \times 10^{38}$ unique addresses. For further discussion see, Deering, S. and R. Hinden, *Internet Protocol, Version 6 (IPv6) Specification*, RFC 1883, December 1995.

[22] http://bgp.potaroo.net/v6/as2.0/index.html

[23] http://www.cidr-report.org/

[24] See Todd Underwood, Renesys, “Realities of IPv6 as the Future Network Layer,” presentation slides to TelcomNEXT conference, March 21, 2006 (see http://www.renesys.com/tech/presentations/pdf/telecomnext-underwood-ipv6.pdf).

[25] For example, Comcast has pressed to adopt IPv6. Some of the national governments which have adopted strong IPv6 conversion policies include China, Korea, Japan, and to a lesser extent, the United States.

[26] Their opposition may be based on their opinions as to the technical merits of IPv6, their private interests, or their forecast of what markets will do.

addresses needs to consider the likely impact on the speed of migration to IPv6. Second, assuming an appropriate mechanism is adopted, the issue of IPv4 scarcity will become less important over time as resources migrate to IPv6. This does not preclude there being an incidental demand for IPv4 for the foreseeable future to satisfy the needs of certain legacy systems and specialty needs. For example, it is not reasonable to expect that there would be a long-lived active trading market in IPv4 addresses if the transition to IPv6 were successful. Third, whatever reforms are adopted to manage IPv4 addresses ought to be considered in light of their possible relevance to the management of IPv6 addresses since institutional reforms are seldom easy and, by assumption, the future is IPv6.

2.3.4 Multihoming and Its Impact

While from its earliest days the Internet has been designed with survivability in mind, the cost of that survivability has been in the dissemination of alternate route information to a destination. Multihoming is the process of connecting a computer interface to two or more networks. While providing multiple paths to reach a computer enhances availability, multihoming also increases the number of routes that must be published to routers across the Internet.

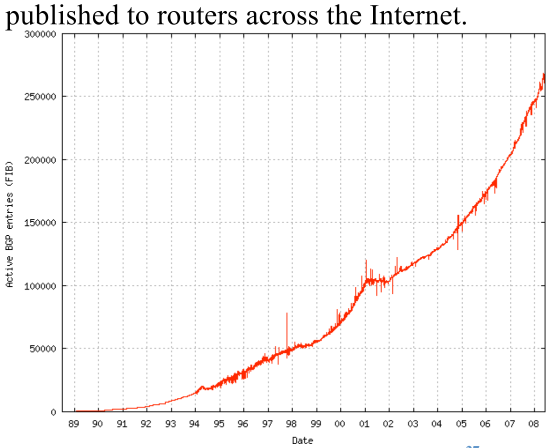

Figure III: Routing Table Growth over Time[27]

Figure III shows the growth of routing information from 1989 through June of 2008. What it says is that for the millions of people who currently use the Internet there exist just over 250,000 entries in the global routing system. While this represents remarkable efficiency, one could be reasonably concerned that if small amounts of address space were allocated and then advertised largely due to multihoming, the number of distributed routes would dramatically increase.

Similarly, were a customer to take address space assigned to them by one service provider and simply move to another service provider, a similar impact would be felt. For this reason, address space received from service providers is generally not portable.

It should be noted that the version of IP (4 or 6) does not change the fundamental scaling properties of routing technology used today as well as most systems under development for the future.

What may change is the nature of routing itself. Work continues within the Internet Research Task Force (IRTF) on new routing paradigms that would separate end system routing from core routing. By doing so these systems would provide for a means to make routing largely independent of addressing. While we do not discuss such work in detail here, technology itself is always a wildcard factor with economic analysis. Certainly the widespread deployment of such new routing protocols would impact RIR policies and this analysis.

[27] Source: Geoff Huston, www.cidr-report.org, June 2008.

2.4 History of IP Address Allocation

In the earliest days of the Internet, a single authority assigned IP addresses and distributed the information through the RFC document series. Initial assignments were made to ease the conversion from earlier protocols, and then later as the potential for eventual scarcity was recognized, assignments were conditioned on a needs assessment. Initially the needs based allocation was based on the faith and credibility of the individual making an allocation request. Over time, more formal approaches were adopted, including the mandate that entities demonstrate use of substantial portions of smaller allocations prior to receiving larger allocations. In addition, over time, responsibility for managing the allocation of IP addresses was distributed to geographically demarcated regional entities known as Regional Internet Registries (RIRs).[28] Today, the Internet Assigned Numbers Authority (IANA) manages the aggregate pool of IP addresses. IANA distributes allocations to the RIRs who then allocate them to ISPs (also known as Local Internet Registries, or LIRs) or in some cases to National Internet Registries (NIRs), as well as to individual network institutions.[29]

Within this system, policies governing resource eligibility criteria and distribution details for both IPv4 and IPv6 addresses are developed through a formal policy development process that encompasses deliberation via email lists and debate in periodic policy meetings. These policy deliberations are open to all interested parties, and typically last for months, but occasionally extend over years in cases of especially important or contentious policy proposals. This institutional arrangement and deliberative process was loosely defined in several RFCs, or standards documents produced by the Internet Engineering Task Force,[30] but has evolved substantially since its inception as a result of intervening technological advances and subsequent policy mandates proposed and adopted within the associated regional communities. As a result, each of five RIR regions now possesses subtly different policy development processes, as well as considerable variation in substantive IP resource-related policies, reflecting the distinctive technological, demographic, and economic realities faced by the "RIR community" – i.e., the IP number resource users who constitute the direct stakeholders -- within that region.[31] That said, the broad contours of the IPv4 allocation process are fairly similar across all regions.

Entities seeking IPv4 addresses submit a needs-based request (demonstrating that the allocation they are seeking is needed to meet their reasonable expectations of growth over the next six months to two years using good address management practices). These requests are usually processed with a minimum of bureaucratic hassle or expense. Network operators that receive resources typically sign a contract that binds them to maintain accurate WHOIS[32] contact information and to abide by other policies defined by the broader RIR community. The non-profit RIRs charge modest membership fees, which entitles members to apply for address resources and which helps to subsidize a variety of other address-

---

[28] The RIRs were established in the 1990s as follows: RIPE (Europe, 1989); APNIC (Asia-Pacific, 1993); ARIN (North America, 1997); LACNIC (Latin America and Carribean, 2002); and, AfriNIC (Africa, 2005). (See Perset, 2007).

[29] NIRs were first introduced as an intermediate address delegating structure in the Asia Pacific Region, in conjunction with the founding of the APNIC regional registry. This model was later adopted in the Latin America region when LACNIC was established in 2002. A present, there are no National Internet Registries within the regions covered by ARIN, RIPE, or AfriNIC.

[30] Hubbard, K., et. al, "Internet Registry IP Allocation Guidelines", RFC 2050 (also BCP 12), November, 1996.

[31] Policies for RIRs are posted on their web sites (e.g., http://www.arin.net, http://www.ripe.net, http://www.apnic.net, etc.)

[32] WHOIS is a protocol for querying the official Internet database that tracks ownership of IP addresses, domain names, or autonomous system numbers on the Internet.

management related activities and operating costs (e.g., maintenance of online WHOIS databases and directories that allow users to identify domain names based on input of an IP address number).[33]

Technically, this contractual relationship invests RIRs with contractual remedies to enforce address usage policies, however, none of the RIRs possess any enforcement or even monitoring capacity apart from those that are available to any third party (e.g., by analyzing open news sources and publicly archived routing tables). Thus, in practice such contractual remedies are almost never used, even when they are explicitly enumerated in RIR "Registration Services Agreements."[34] As a result, most of the actual normative power available to RIRs depends on their ability to evaluate ongoing compliance if and when growth compels an IP address resource recipient to return to the RIR to seek a subsequent allocation of IPv4 addresses. Inasmuch as extending routing services to new customers invariably requires additional IP addresses, successful ISPs that continue to grow are obliged to return to the RIR on a regular basis to obtain more address space.[35] On each such occasion, the RIR is able to reconfirm the completeness and accuracy of WHOIS information, and review the supplemental needs-based request for evidence that previous IP number resource allocations have been assigned in ways that are consistent with established address conservation principles and practices.

This arrangement of incremental allocation and subsequent review, together with the RIR practice(s) of maintaining online policy document repositories, hosting multiple dedicated mailing lists for policy discussions, and mounting regular (typically semi-annual) open policy forums – many of which are scheduled in conjunction with network operator training and "best practices" conferences – have created a low-cost, lightweight, but arguably effective quasi-regulatory regime. Address resource recipients that believe that they are likely to require additional number resources in the future, and by implication know that they could be affected by policy developments that might change the terms of their availability are naturally motivated to remain engaged in the RIR policy development process and comply with community norms.[36] Their on-going engagement also provides a mechanism for sharing information publicly about address usage and network planning/design that otherwise might remain confidential.

Distributing management across multiple RIRs offers several natural benefits. First, distributing the allocation process makes IPv4 more readily accessible to all user populations, and helps to make local participation in policy development more tractable. Second, allowing the flexibility for different RIRs to adopt slightly different rules facilitates opportunities for comparative innovation. Third, regionally distributing address management makes it easier to indirectly incorporate market information by taking into account such local phenomena as trends in ISP growth, divergent architectures, or service offerings

[33] Mockapetris, P., "DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION", RFC 1035 (also STD 13), November, 1987.

[34] The legal basis is even weaker with those holding legacy address allocations (or those assigned prior to the establishment of the RIR system).

[35] In principle, this subsequent allocation review is also applicable to the other category of RIR-held IPv4 address space known as *provider independent*, which is allocated to institutions exclusively for their own internal use. Recipients of provider-independent (or PI) address space – typically large institutions from other (non-Internet) sectors but which possess all of the prerequisites required to "in-source" their own networking and interconnection requirements -- do not provide third-party routing services, have no other enterprise customers of their own, and are thus ineligible to sub-delegate IPv4 addresses to third parties. Although PI recipients rarely need to return for a subsequent allocation, in cases where internal growth creates demand for additional IPv4 addresses, these institutions are subject to the same kind of compliance review described above.

[36] The critical role played by periodic review as a deterrent against noncompliance, and recurring policy education and outreach as a reinforcement tool in self-governance regimes is highlighted in recent empirical research (see Short and Toffel, 2008).

which might impact differentially local demand for addresses. Finally, distributing this function reduces the risk that a single administrative entity would fall capture to a narrow group of interest.

Because the subsequent allocation process is based on prior history, new entrants and slower-growing incumbents tend to receive smaller increments of address space than larger, faster-growing incumbents. In some cases, subsequent allocations may be larger than the initial allocations received by new entrants. Once addresses become scarce, this pattern might be perceived as biased in favor of larger incumbents.

Arguably, this broadly defined self-governance system -- with RIRs in the central coordinating role -- has been successful to date, incurring relatively low administrative costs while facilitating substantial global growth of the Internet. Nevertheless, the system has always had detractors, since even before the time of its establishment, including prominent advocates of market-based approaches to allocating IP number resources.[37] These critics do not dispute the successes as defined above, but argue instead in favor of an entirely different (and superior, they claim) paradigm, complete with different criteria for defining success. In this sense, it is fair to say that the administrative, policy-defined, eligibility-based method for distributing IP addresses never ceased to be a source of contention in some quarters, even after alternative approaches were rejected by the Internet developers and operator's community in 1994 in favor of the current system.[38] While this unresolved controversy does not change the fact of impending IPv4 address exhaustion, it does represent a significant source of tension underlying many perceptions and assertions about the future of IPv4 address management.

In any case, if and when the RIR-held pool of unallocated IPv4 addresses is exhausted, by definition there will be no subsequent allocations of IPv4. Under current arrangements there will also be no more policy compliance reviews, and thus no more natural or induced incentives to share technical information about address usage practices, or to conform to conservation rules or other IPv4-relevant policies – except where purely voluntary self-disclosure and unobservable compliance advances the private interests of individual network operators.[39] Although presumably, the RIRs will still administer the unallocated pools of Autonomous System Numbers (ASNs) and IPv6 addresses, the management of these resources is unlikely to sustain the kind of quasi-regulatory policy review, updating, compliance, or registry data maintenance dynamics that have made self-governance as successful as it has been under the current regime.[40] Since at least some of these functions are essential, new mechanisms must be identified to fulfill the most critical of these functions. Philosophical and ideological disputes notwithstanding, the identification and placement of these new functions (whether in the context of the existing IANA/RIR framework or in some new framework) lies at the core of every serious debate on this issue that is now underway.

3 Approaches to Reform Internet Address Allocation

From the above it is clear that demand for IPv4 addresses will continue beyond the date when the pools of unallocated IPv4 addresses are exhausted. The slow pace of IPv6 adoption suggests that even if this transition does complete eventually, it will happen well after the pools of unallocated IPv4 addresses are

[37] Huston, G., "Observations on the Management of the Internet Address Space", RFC 1744, December 1994.

[38] Rekhter, Y., "Address Ownership Considered Fatal" , RFC 1518, July 1995.

[39] See Short and Toffel (2008) at page 40: the authors observe that the proactive, pro-compliance stance of enterprises that are subject to infrequent regulatory review and regular policy education outreach efforts rapidly erodes once the reviews and outreach efforts are halted.

[40] With ASNs and IPv6 address allocations, those seeking resources are far less likely to need to return for subsequent allocations.

exhausted; and, if the IPv6 transition does not occur, then IPv4 demand will continue indefinitely.[41] Since the universe of IPv4 addresses is fixed, the only source for meeting the on-going demand may eventually be the pool of addresses that has already been allocated. Thus some sort of "transfer mechanism" to reclaim "excess" IPv4 addresses to either replenish the pool from which subsequent allocations may be made or to directly transfer addresses among users is necessary. Moreover, the impending exhaustion of the IPv4 pool poses a challenge for the on-going smooth working of the self-regulatory process that characterizes the current regime.

In light of the looming "scarcity" of IPv4 addresses, it is perhaps natural that there is renewed interest in expanding the role of market mechanisms. Economists such as Coase (1960), Demsetz (1967), and Williamson (1975) have identified the benefits of markets in resolving resource allocation decisions when interests diverge, the stakes get higher, and informational complexity increases. Markets, however, do not exist in a vacuum and their prospective benefits (and risks) are hardly beyond dispute. While markets offer one mechanism for transferring IPv4 addresses, they are neither the only nor obviously preferred mechanism. To properly evaluate the role for markets, they need to be considered in relation to the alternatives. Although it is clear that change is needed in the way that IPv4 addresses are allocated, the wrong change could make matters worse.

3.1 A taxonomy for thinking about address allocation mechanisms

The space of possible allocation mechanisms may be described along two dimensions: (1) whether the mechanism is centralized or decentralized; and (2) whether the process is administrative or market-based (see Table I).

	Administrative	Markets
Centralized	Legacy centralized allocation	Centralized market(s) with market-maker (maybe IANA/RIRs)
Distributed	Current framework of IANA/RIRs	Decentralized, bilateral trading

Table I: Taxonomy of Market Mechanisms

"Administrative" processes encompass frameworks where decision-making is based on rule-based group discussions, where market forces are indirectly and endogenously reflected via the interests of participants.[42] "Market-based" processes involve allocations resulting from the aggregate behavior of decentralized supply/demand decisions mediated via an exchange market. Either type of process may be more or less distributed.

In the context of an administrative process, there may be a single institution (e.g., IANA) that might make the decisions (and in the theoretical construct, by assumption, a single dictator who does not suffer from the need to reconcile strategic interests) or there might be multiple decision-making entities (e.g., IANA/RIRs) with peer-like authority (and in the theoretical construct, may require consensus-based decision-making to reconcile divergent interests). The latter case approximates the address allocation framework used by the IANA/RIRs currently.

41 See Elmore, Camp, and Stevens (2008).

42 The participants in the current administrative regime are employees of ISPs and other vested interests and the RIR system has evolved over time in response to changing market conditions.

Market mechanisms may similarly be characterized as being more or less centralized. For example, at one extreme we might imagine a single market (or a few RIR-managed regional markets) that clear all address transactions. Those wishing to relinquish IP addresses sell them to the market where they become part of the allocable pool. Those wishing to purchase IP addresses buy them from the pool. There is no explicit matching of buyers and sellers, but there is the presumption of a market maker who ensures that there is an orderly sharing of information (i.e., potential buyers/sellers may observe the market price). At the other extreme, trading may be fully decentralized with buyers and sellers negotiating private, bilateral trades.

In reality, most processes will lie somewhere between these extremes. For example, markets are often constrained by administrative rules to ensure their orderly functioning. Such rules may include restrictions on the size of transactions (address block size) or holding periods to control destabilizing speculation or attempts to corner-the-market. Or, the extreme model of decentralized trading might be augmented with some element of centralized information sharing. Similarly, administrative processes may be more open to market-based transfers or may use price signals (e.g., by varying address licensing fees) to recover address blocks that are deemed "excess" from holders of existing allocations.

Finally, if a transfer market for IPv4 addresses is implemented, it need not encompass the entire universe of IPv4 address space. For example, market mechanisms that applied to either or both initial allocations and address resource transfers could be complemented/supplemented by an IANA/RIR managed pool for allocating addresses. That is, there could be a reserve pool created either from the remaining unallocated pool of IPv4 addresses or from reclaimed addresses (e.g., given new RIR rules and resources, reacquired either via an administrative process or via purchases from the transfer market). Maintaining a reserve pool for allocation by the IANA/RIRs could prove useful as a check on market efficiency to help stabilize/publicize[43] prices and to provide a vehicle for subsidizing address acquisition by target communities (e.g., new entrants seeking small allocations that may be regarded as disadvantaged in a decentralized transfer market).[44]

3.2 Potential Benefits and Risks of a Transfer Market

Before examining more specifically the merits of the various proposals currently under consideration for introducing transfer markets, it is worthwhile considering more closely some of the notional benefits and risks from introducing market mechanisms.

3.2.1 Benefits of a transfer market

As noted, we expect demand for IPv4 addresses to persist into the indefinite future. Expanding the supply of addresses to meet on-going demand is the principal benefit promised by each of the proposed transfer mechanisms. While market mechanisms are not the only conceivable vehicle for meeting this on-going supply requirement, there are a number of reasons why a market might offer a preferable alternative to an administrative process that may be posited, including:

[43] Assuming a "reserve pool" of IPv4 addresses could be created and sustained, the IANA/RIRs could leverage it to act as market-makers to help control volatility and to help assure market participants that the market would be orderly. Additionally, keeping a hand in the allocation process would provide continuity with the past (which is reassuring) and a mechanism for monitoring market supply/demand and pricing and publicizing the information to reduce information asymmetries. How complete the IANA/RIRs information of what is going on in the market will depend on how a possible transfer market works.

[44] Using the reserve pool in this way would provide a mechanism for a sort of "universal service" or "carrier-of-last-resort" style role for the IANA/RIRs to play. This would help address concerns that transfer markets would provide unfair access, potentially foreclosing some classes of ISPs or end-users seeking additional addresses.

3.2.1.1 Allocative, Productive, and Dynamic Efficiency

Markets are often viewed as offering an efficient mechanism for allocating scarce resources in complex environments where there are conflicting interests and lots of asymmetric information. In the absence of these features, it is often preferable to rely on a simple administrative process. For example, if there is no scarcity and there are no conflicting interests, then anyone who wants addresses should be able to receive them. In the earlier days, when IPv4 addresses were abundant, this was approximately the case and there were few constraints on supplying demand for IPv4 addresses. Later, the confluence of ISP "subsequent allocation" requirements, RIR-administered address policy compliance reviews, and community-driven policy development helped to reduce information asymmetries that otherwise might have accelerated and exacerbated conflicts of interest. So long as these interlocking processes continued to sustain a consistent definition and application of eligibility criteria for RIR-allocated resources, and adequately maintained the concept of "need" – i.e., as neither so slack as to permit much address waste or surplus, nor so taut as to force owners of address blocks to make disruptive internal rationing choices about their own IPv4 addressing requirements -- the administrative allocation system arguably provided a satisfactory alternate mechanism for the efficient distribution of IP address resources.

Over time, however, as the exhaustion of the IPv4 pool has loomed ever larger and RIR community policies and administrative practices have imposed incrementally stricter constraints on the supply of IPv4 addresses, IPv4 addresses have become increasingly scarce.[45] This increasing scarcity raises the likelihood of conflicting interests and the potential that addresses will need to be rationed (that is, not all demand will be met). To date, Internet developers and RIR communities have sought to minimize such conflicts primarily by modulating the size of allocations, rather than by tightening eligibility criteria.[46] Eventually, however, either because eligibility rules become so restrictive that few institutions would qualify for RIR-held address space, or because the size of "need-based" allocations will no longer permit service providers to assign IP addresses in a manner that is consistent with widely accepted operational best practice, rationing will cease to be a perceived risk and instead be widely recognized as an undeniable reality.

At that point, accurately prioritizing who should get limited allocations so as to maximize collective surplus would require detailed information about true needs for an allocation of additional addresses. Those requesting the allocations are likely to have the best information about their true needs (e.g., expectations for growth, usage of their current allocations, and their firm-specific costs for conserving address space through the use of NAT boxes or other strategies). Some of this information is likely to be

[45] Scarcity is *not* a phenomenon that depends on the unallocated pools being exhausted. The prospect of exhaustion and expectations of mitigation strategies which may be employed by the Internet community prior to exhaustion and thereafter have been influencing behavior and will continue to do so before the date when the pools are finally empty.

[46] Expansion of the original, early 1970s ARPANET address space from eight bits (sixty-four devices) to sixteen bits (256 devices) is described in John Day (2008), pages 143-146. "Classful" IPv4 addressing, which leveraged the familiar IPv4 address space to support up to 127 very large independent networks, 16,384 medium-sized independent networks, and over two million small networks was codified by RFC 791 in September 1981 (http://tools.ietf.org/html/rfc791). In September 1993, RFC 1517 (http://tools.ietf.org/html/rfc1517) defined Classless Inter-Domain Routing (CIDR), which permitted great flexibility in sizing IP address segmentation after that time. Thereafter, IP address segment sizes were primarily determined based on "need" -- i.e., by matching documented, near-term requirements with the nearest size address block -- or through community-driven policy development at the RIR level, with individual RIRs defining and redefining the size of "initial allocations" and "subsequent allocations" over time, in response to changing technology and institutional circumstances. One good example of the latter can be found in the RIPE NCC 2000 Annual Report, which links the availability of a new address-conserving technology ("name-based web hosting") with a 50% reduction in the size of initial allocations for new RIPE members (http://www.ripe.net/ripe/docs/ripe-218/).

strategically sensitive and those requesting addresses may be unwilling to share it fully or truthfully with any third-party, even a "neutral" confidentiality-bound third party like an RIR. Furthermore, entities requesting addresses may have an incentive to misrepresent their demand if they think it will influence their likelihood of receiving their requested allocations. In such situations involving large numbers of potential buyers, complex and asymmetric information, conflicting strategic interests[47] and a scarce resource that needs to be allocated, markets are often regarded as providing a preferred alternative.[48]

Ideally, with a well-functioning transfer market (i.e., one approximating the performance of "perfect competition"), there would be a market-based price for buying or selling IPv4 addresses which would help aggregate the decentralized and asymmetric information on supply and demand conditions. Economic theory suggests that this will help facilitate economic efficiency in a variety of forms:[49]

- *Allocative efficiency*: goods go to the highest value uses first (i.e., to those that can offer the highest price for the resource). This means that if addresses need to be rationed, that higher-value uses, as defined by the highest bidders, are served first.

- *Productive efficiency*: goods are produced with the least cost consumption of resources. In this context, this implies that IP addresses would be both assigned and *routed* efficiently, and thereby facilitate the interaction of more users, content, and online services than would be possible under non-market allocation arrangements.

- *Dynamic efficiency*: investment and consumption are efficient over time. In the present context, this would imply that the investment required to foster the Internet's continued growth and dynamism – including specifically investment required to permit the transition to IPv6 – will happen faster than would be possible under alternate circumstances.

3.2.1.2 Induce strong incentives to use IPv4 addresses efficiently

A scarce resource has a direct or indirect opportunity cost associated with its use (i.e., the value of the resource in its next best alternative use). As opportunity costs increase, users have an incentive to conserve resources, or equivalently, to use them more efficiently. Active resource markets that allow users to buy *and* sell the resource can make the opportunity costs directly observable and thereby strengthen incentives to use the resource efficiently. While a market process is the most common way to induce these direct opportunity costs, it is not the only way. For example, an administrative process may charge resource fees that are set so as to induce a direct opportunity cost even in the absence of market transfers.

Under the current regime, the opportunity costs for IPv4 addresses are artificially low and asymmetric. That is, the fees charged by the RIRs today are designed only to recover administrative costs and are not intended to induce direct financial incentives to conserve addresses. The opportunity costs are asymmetric because of differing growth prospects, differing costs associated with conserving IPv4 addresses (which includes substituting to IPv6), and past allocation decisions. There is also an asymmetry arising between new and past address allocations because there is no approved mechanism for selling or transferring

---

[47] For example, parties may care not just how many addresses they receive but also how many addresses competitors receive. Those with adequate allocations may prefer to see their competition face address scarcity constraints.

[48] See Coase (1960) and Demsetz (1967).

[49] Depending on one's perspective, some of all of these notions of economic efficiency apply equally well to all allocation systems, including non-market mechanisms.

existing allocations of IPv4 addresses.[50] With respect to new allocations, the RIRs' needs-based assessment process is intended to induce incentives to use the addresses efficiently. However, because current rules do not allow those with existing IPv4 addresses to sell those addresses to others, the opportunity cost associated with potential "excess" addresses is also artificially depressed.[51]

A market – assuming it works appropriately (which is a risk we will consider further below) – offers the potential of balancing opportunity costs and enhancing incentives to use addresses efficiently. If there were a transfer market for IPv4 addresses then there would not only be an opportunity cost for acquiring new addresses, but also an opportunity cost for retaining an inventory of IPv4 addresses.[52]

Furthermore, using financial incentives to induce an opportunity cost is sometimes viewed as preferable to administrative approaches (e.g., specifying specific best practices such as a target address utilization rate) because it provides the participant greater flexibility to choose among resource conservation strategies.[53] However, as already noted, because the use of address space needs to be coordinated to manage Internet routing and there are large negative externalities associated with using the address space inefficiently, the calculus of private/public costs/benefits might differ systematically. For example, the private behavior induced by stronger financial incentives via a market for IPv4 addresses could adversely impact the overall public interest by undermining support for technical standards for address use that contribute to the overall security and stability of Internet routing.

3.2.1.3 Competition in Mechanisms

The existence of a transfer market for addresses does not preclude there also being a RIR-administered allocation process. Indeed, we believe that the RIRs can and should play a useful role in offering a safety valve if the market does not behave as intended. We also accept the proposition that RIRs should play a role in preserving basic access to the universe of IPv4-mediated resources for as long as IPv6 addressing alone cannot provide a reasonable substitute means for achieving such access, or at least until such time as a market demonstrably fulfills this critical function.

Having both an RIR administrative process and a market-based transfer market may provide competition in mechanisms and a safety valve if one process worked better than another. The competition in mechanisms may provide a basis for identifying best practices and inducing incentives for improvement. Alternatively, competition in the mechanisms may result in one supplanting the other and the best mechanism need not be the winner. For example, enabling a transfer market may allow those with little interest in converting to IPv6 or otherwise complying with RIR policies to bypass the RIR process, thereby undermining the viability of the RIR system (and, for example, incentives to maintain the WHOIS database). Whether by design or by accident, the potential adverse impact of creating a transfer

[50] Addresses may be transferred through merger & acquisition activity, and based on a case-by-case basis, pending review by the RIRs. Addresses sometimes are also transferred illegally via an unapproved "black" market. This can occur through such means as the sale of shell companies, or just a simple change of assignment.

[51] Of course, even if users with excess addresses could sell them, as long as potential buyers have the option of acquiring addresses from an RIR via a low-cost administrative review process, that option will impact the opportunity cost of holding on to excess address inventory.

[52] Individual "opportunity costs" would include both the direct (market-based pricing) and indirect costs associated with managing usage of the addresses. The latter may include the search costs of identifying a potential buyer/seller, the transaction costs associated with managing the transfer of the resource, and other idiosyncratic costs that may vary by participant. Thus, opportunity costs need not be equal across participants.

[53] In its extreme form and overly simplistically, this is sometimes touted as a benefit of decentralized planning under market capitalism as opposed to centralized planning under socialism. See Arrow and Hurwicz (1960) for a more analytic discussion.

market on the existing regime of Internet self-governance needs to be adequately considered lest the creation of this new option make matters worse.

3.2.1.4 Redefining Fairness

In some contexts, markets are viewed as more fair than administrative processes. However, this perception is hardly universal, and there are clearly examples at both extremes (e.g., entrenched monopoly markets or pure "Athenian-style" democracies). Those that might view markets as more fair may point to the fact that appropriately working markets favor "survival of the fittest" which to them is favorable to adhering to an administrative process that may be captive to legacy interests or over-weights minority interests. Critics of market-based mechanisms may argue that markets are unfair because they are unlikely to work appropriately or will favor the market power of incumbents (and current interests over those of future Internet participants). Thus, whether markets or administrative processes are perceived as more or less fair depends in large measure on the eye of the beholder.

In any case, the perceived fairness of markets is likely to be enhanced if precautions can be built in at the time of market design to ensure that a transfer market remains neutral over time with respect to potential buyers and sellers. Fairness will also be strengthened if adequate protections and remedies are established to protect against the exercise of market power, so that aspiring new entrants might have the same degree of freedom to "choose autonomy" that was enjoyed by all of the institutions that will populate the "sell-side" of any future address transfer market. Finally, the perceived fairness may also require addressing the possibility of windfall profits.[54]

3.2.1.5 An approved, well-functioning market is preferable to the alternatives

If the transition to IPv6 is delayed too long (perhaps forever) and in the absence of a transfer market or other liquidity-promoting measures, IPv4 addresses will become increasingly scarce as pent-up demand grows. As already noted, an opportunity cost will be induced by this scarcity and if it gets high enough on the buy side, it could encourage a variety of responses, including grey/black market dealings (i.e., illegal IPv4 transfers), outright IP address hijacking, or the gradual loss of connectivity in the Internet.

At present, IANA has no enforcement capabilities, and RIR powers to promote compliance with community-defined policies are similarly circumscribed by the authority vested in them by those same communities. As a practical matter, any regulatory authority, formal or informal, must command means to influence behavior that are at least as strong as the incentives that they are attempting to counter. If the incentives to deviate from the rules exceed the regulator's ability to deploy countervailing incentives, then deviation is likely to be commonplace. In this case, the lightweight, low-cost, self-regulating mechanisms adopted by the IP address resource community and vested in the RIRs have never been more than "good enough," even under normal conditions. Absent some investment of new powers and authorities in new or existing institutions, these mechanisms are likely to be no match for the incentives created by the unmitigated scarcity of IPv4 addresses, or those created by an official or unofficial resource transfer market.

Confronted with such alternatives, it is possible to view the creation of an approved transfer market as preferable, or at least, the *least bad* alternative. However, to be successful on the terms defined for success to date (e.g., in RFC 2050), an approved transfer market would also have to encompass new mechanisms and incentive structures that promote overall transfer patterns that are consistent with the goals of fairness (however that is defined), efficiency of use (which includes providing better information to the community on how the transfer market is operating), and non-harm to the Internet's overall security

[54] The windfall profits issue might arise if those with legacy allocations are perceived as benefiting excessively by the opportunity to sell excess IPv4 addresses.

and stability. In order to be more than just an optional mechanism within a broader "black" or "gray market", an approved transfer market must also encompass intrinsic benefits that will make it inherently preferable to both IPv4 buyers and sellers.

3.2.1.6 Reduce the risk of government regulatory intervention

As noted previously, we do not perceive retaining the status quo as a viable option. Something needs to be done to address the challenge of IPv4 pool exhaust and the conversion to IPv6. An active transfer market that coexists with the traditional RIR administrative allocation process would provide a form of mechanism competition that could help counter concerns expressed by some that the current regime is inadequate or lacks legitimacy – weaknesses that could increase the risk of direct government intervention. While further administrative reforms may also help, these alone may not be sufficient, especially in light of emerging global concerns -- especially among some outside critics and potential successors -- about the legitimacy of the various Internet governing bodies (IETF, IANA, and RIRs).[55]

The management of IPv4 pool exhaustion is just one of the many looming issues for global Internet governance. Indeed, the IPv4 problem may be the "canary in the coal mine." This heightens the political saliency of this issue and makes questions of how a transfer market is structured, how risks are assessed/addressed, and fairness issues resolved especially contentious.

For the last several decades we have seen a secular global trend toward telecommunications network privatization and deregulation and a general preference for market mechanisms over traditional public-utility style regulation.[56] Creating a transfer market for IPv4 addresses is consistent with that trend.

Finally, the creation of a transfer market with a realistic, durable role for the (perhaps duly re-empowered) RIRs -- e.g., as managers of a residual pool, as a title registry or recorder of resource transfer transactions, or as a market maker -- could provide a graceful evolutionary path to help the RIRs transition into the future in which direct administration of an unallocated IP address pool is unnecessary. Given the fact that Autonomous System Numbers and IPv6 addresses will continue to require uniqueness and thus some form of managerial oversight, and that the IANA/RIR institutions have served the community reasonably well to date (and importantly that there is no current feasible institutional alternative for maintaining critical WHOIS-related number resource registration data) figuring out a role for the RIRs after the unallocated pool is exhausted seems desirable, if not essential. A transfer market represents one obvious means to this broader end.

3.2.1.7 Accelerate and smooth the transition to IPv6

As with fairness, the question of whether introducing a transfer market would accelerate or retard the transition to IPv6 is ambiguous. On the one hand, if the market works appropriately, it will induce an opportunity cost for IPv4 addresses that reflects their true scarcity value. Those with excess address resources that cannot be profitably assigned to current customers or reserved for certain, near-term customers or other revenue-generating purposes, would find it costly to retain their inventory, and thus might be induced to sell addresses to those in need. Those who do choose to sell might use the resulting capital gains to help fund their own transition to IPv6.

---

[55] See, for example, Butt (2005) or Hongqiang (2005).

[56] Of course, this trend has hardly been consistent. Strong regulatory interventions such as unbundling rules or proposals for structural remedies (including asset separation) have been adopted in a variety of regulatory forums in the United States and abroad. The on-going debate over network neutrality demonstrates that sponsorship of on-going regulation of the Internet and its infrastructure is far from dead.

In the absence of a transfer market, service providers with excess IPv4 addresses and no expectation of future growth would have no material incentive to use their address resources conservatively. This "freedom to be wasteful" could reduce or delay pressure they might otherwise feel to switch to IPv6. Conversely, those network operators who need additional IPv4 in the future but cannot acquire them might be induced to migrate to IPv6 prematurely (at higher total cost).

While we are assuming in this paper (*in arguendo*) that accelerating the transition to IPv6 is a good idea, this does not mean that there are no social benefits in having those with lower conversion costs migrate earlier.[57] To the extent that such claims have merit, a transfer market in IPv4 could provide a mechanism for optimizing the overall sequence of IPv6 integration, substitution, and eventual migration over the total population of network service providers, thereby lowering the aggregate costs of the Internet's collective turn to IPv6.

The problem of inducing IPv6 transition may be usefully illuminated by an S-shaped adoption curve, which has been observed to accurately mirror the uptake of new standards and technologies.[58] On this model, early adopters bear an asymmetrically high cost because the benefits of adopting are less when the number of those who have adopted is less (because of positive network externalities) and the costs are greater (because there are less learning and scale/scope economies associated with the new IPv6 technologies, and because the need to bear interoperability costs is greater[59]). Near the end of the transition to the new technology, the roles are reversed with late adopters bearing the costs of interoperability disproportionately.

There is a point along the S-shaped adoption model trajectory when enough early-adopters have switched so that critical mass has been achieved and the market takes off (steep part of the S-shaped adoption curve). Eventually, the market of adopters saturates which may be at 100% (everyone converts to IPv6) or some lower share of the market. Let the inflection point for when mass adoptions occur be T(A) and the inflection point for when the market has effectively saturated be T(B). The impact of introducing a transfer market may be evaluated by its impact on the timing of T(A) and T(B), and the eventual level at which the market saturates (how close to 100% is the conversion to IPv6?).

Today, the industry is still in the early stages of transitioning to IPv6 (i.e., we are not yet at T(A)). While most Internet hardware and operating system software already supports IPv6, network operators have been slow to adopt.[60] The question to ask is whether introducing transfer markets will postpone T(A), or extend the length of the transition T(B)-T(A).

---

[57] To date, claims of super-linear adoption costs have only been made *speculatively* by very large network operators – i.e., operators that have yet to undertake the adoption effort that would be necessary to empirically determine actual adoption costs. Consequently, the objectivity and/or veracity of these assertions cannot be independently ascertained.

[58] See Elmore, Camp, and Stevens (2008).

[59] In contrast to other applications of the S-curve model, direct interoperation between pure-IPv4 and pure-IPv6 networks is technically impossible. Thus in this case, the S-curve's interoperation variable would correspond to the price of obtaining IPv4 addresses, which is borne as a direct cost by new entrants and accrues as a direct revenue opportunity for incumbent IPv4 holders.

[60] See for example http://ipv6.google.com/

During the interval T(A) to T(B) (the transition), every network operator will have to bear some interoperability costs.[61] These costs may be internal or external, direct or indirect.[62] The longer the transition takes, the higher many of these costs will rise. However, so too will the revenue opportunities enjoyed by those networks that are able to service continuing demand for IPv4 addresses and/or demand for "backward connectivity" to any remaining IPv4-bottlenecked content and services. If the transition is inevitable, then one might view late adopters as imposing an externality on early adopters by failing to switch to IPv6, or at minimum, failing to take unilateral steps to make their public-facing, IPv4-based Internet resources openly accessible to pure IPv6 networks. Alternatively, if one views the transition as purely optional, those early IPv6 adopters might be perceived as the source of externalities confounding the emergence of permanent IPv4 scarcity mediated by durable market mechanisms, or perhaps an as-yet unidentified "third way" forward.

The authors of this paper do not agree on the likely impact of a transfer market on the pace of IPv6 transition. Arguably, this suggests that the likely impact is, indeed, *ambiguous*. One author (Vest) believes there is a significant risk that introduction of a transfer market may substantially increase incentives to delay conversion to IPv6, and quite possibly derail it altogether. This prediction is based on the hypothesis that the establishment of a market for IPv4 transfers will present every individual IPv4-based network operator capable of providing Internet services with a mix of incentives that is broadly weighted toward, and cumulatively strengthened over time, by the perpetuation of IPv4's status as critical, non-substitutable (i.e., "bottleneck") input for Internet service delivery. On this argument, individual IPv4-based incumbents facing the certainty of one-time returns resulting from an immediate simple sale and transfer of surplus IPv4, might opt instead to postpone any sale as long as transfer prices appear to be stable or appreciating, with any resulting opportunity costs being offset by the steady stream of recurring revenues that could be realized by making the surplus IPv4 addresses accessible only as an indivisible part of a "bundled" Internet service (e.g., IPv6-IPv4 translation or IPv4 transit). As long as the market for simple transfers and critical IPv4-related services continues to appreciate – or at minimum, simple transfer prices do not clearly depreciate below the medium-term gains achievable in the bundled IPv4 services market – simple sale would be irrational. As long as simple IPv4 transfers remain a sub-optimal strategy for potential sellers, and IPv4 continues to be sought by at least some growing ISPs -- and remains an absolutely non-substitutable requirement for aspiring Internet service providers -- prices are likely to continue appreciating, in perpetuity. In turn, this appreciation would encourage surplus IPv4 to become increasingly concentrated amongst providers of bundled IPv4 services, both through successful bargaining and competition for the remaining liquid IPv4 addresses, and by the gradual adoption of increasingly attractive IPv4 service-centered business models by large, independent IPv4 reserve holders that previously provided networking services only to themselves.[63] The key ingredient in this perverse mix is not the scarcity of IPv4 itself, but rather the continued scarcity of substitutable Internet content and services that require no IPv4 at all – i.e., that are transparently accessible by IPv6-only networks. Knowing this, every prospective IPv4 seller-cum-service provider could face strong individual *dis*incentives to migrate their own online content and services to IPv6, or otherwise make such resources accessible to pure IPv6-based operators. The cumulative result of such individual-level dynamics would

[61] It is presumed that once a critical mass of online content and services are implemented with IPv6, it will not be an option not to interconnect with at least some of those networks; and, as long as a critical mass of IPv4-based content and services endures, interconnection with those will also be needed.

[62] Costs may arise internally associated with running dual-stack network functions or announcing simultaneous/redundant IPv4 and IPv6 routes for those with both IPv4 and IPv6), or externally, associated with the cost of IPv4 addressing or translation services for those with IPv6 resources. The costs may be direct (i.e., the expense of sustaining global connectivity in a bifurcated Internet) or indirect (the lost benefits of foregoing full Internet connectivity).

[63] http://www.iana.org (see "Legacy" network allocations)

be an indefinite postponement of, and active resistance to any IPv6 transition by the same institutions whose strategies, decisions, and investments will ultimately determine when, how, or *whether* such a transition will actually happen.

The other authors (Lehr & Lear) believe that the above scenario is highly unlikely and consequently that the impact of transfer markets on the transition to IPv6 is ambiguous. Lehr & Lear believe that the above scenario is dependent on a view of market power and the ability of market participants with a vested interest in implicitly colluding in maintaining the rents associated with IPv4 access arbitrarily high into the indefinite future. In light of the general competitiveness of the Internet and opposition from end-users (including content and application provides, manufacturers of Internet equipment and devices) from paying such rents, it seems unlikely that the transfer market would, in itself, enable the adverse outcome described. Moreover, high prices for IPv4 (or remaining in an IPv4 world indefinitely) would provide incentives to transition to other alternatives (presumably, IPv6). Finally, if the posited market power exists, it is unclear that manipulation of the Internet address market would be the best way to express such power since cartelization of that market may be easier to detect and hence more likely to attract antitrust intervention.

While the authors may not agree on the likely impact of transfer markets for IPv6 transition or about the magnitude of the risks of market power in a prospective transfer market, they agree that this is an issue that requires additional consideration. All of the authors agree that the impact on the IPv6 transition can be made worse by inappropriately designed transfer markets.

Finally, it should be noted that the expectation of an S-shaped adoption curve has important implications for the intertemporal path of IPv4 prices and scarcity. Early on (before T(A)), one would expect IPv4 addresses to be relatively scarce and the price for acquiring addresses might be relatively high. Later (near T(B)), when the majority of networks have converted to IPv6 and the absolute majority of online content and services are transparently accessible *without* IPv4, the relative scarcity of IPv4 addresses should be much less and consequently transfer prices would be lower. This suggests the role for transfer markets will change over time and mechanism/institutional arrangements (e.g., enforcement tools) that are created specifically to address the IPv4 problem might have a relatively short life of value. However, these mechanisms may be useful for IPv6 management although the need for a transfer market may be less with IPv6 if the supply is, as many argue, effectively limitless. This also suggests that if the transfer to IPv6 could be blocked, that the long-term prices for IPv4 addresses may remain high and be increasing in the future. These divergent expectations regarding the likelihood of conversion to IPv6 highlight the divergent expectations that participants may have over long term IPv4 prices.[64] This reasonable divergence in price expectations will add to incentives to engage in speculative behavior and may contribute to price volatility – suggesting the need to address these concerns in the design of any prospective transfer market.

3.2.2 Downsides of creating a transfer market

In the preceding sections, we sought to provide a balanced view of the potential benefits of introducing market mechanisms to facilitate IPv4 address transfers. That discussion indicated in general that (1) there are many ways to design a transfer market and some are likely better than others; (2) many of the proposed benefits of a transfer market may also be realized via administrative processes, or may be attainable only within a mixed or regulated market context, so market mechanisms may be neither necessary nor sufficient to achieve the desired goals; and (3) optimizing for some potential benefits may necessitate trade-offs, and therefore require a clear delineation and ordering of goods, from absolute non-

[64] If you think IPv6 is inevitable, then IPv4 prices should be initially high and then decreasing. If you think IPv6 conversion will not occur, then IPv4 prices may be expected to increase over time.

optional requirements to advantageous but non-critical benefits. Overall, it is not reasonable to conclude that a transfer market is a good or bad idea *in general* without considering specific proposals, both individually and, where necessary, in combination(s): that is, what is the market design and what are the options if a market is not implemented?

Before turning to our evaluation of the proposals currently under consideration, it is worth considering several additional risks of introducing a transfer market and the options for mitigating those risks in the following sub-sections.

3.2.2.1 Routing table growth

As previously mentioned, inefficient use of the address space may result in excessive growth of router tables. While the cost of announcing another Internet route is relatively small for the owner of the address, the collective impact on the universe of Internet routers that must be updated is not small. This is the negative externality we discussed earlier.

There is the potential that a transfer market might result in the trading of small address parcels, which could lead to fragmentation of the address space, reducing the efficiencies of address aggregation and resulting in the growth in routing tables. Appropriate design of the market will need to consider how to manage transfers so as to preserve address aggregation.

Additionally, but less of concern on its own, an overly active transfer market might require more frequent updating of the routing tables.[65] Frequent updating increases costs and may threaten Internet routing stability. Thus, an appropriate transfer market design must also consider how to avoid excessive trading (including speculation) that would drive up routing costs.

The issue of routing table growth may be substantially altered by technological innovations on the horizon. The current routing technology is based on early 1990s vintage technology. New mechanisms are in development that would separate routing into provisional and operational forms to reduce the amount of information that needs to be updated/shared across the Internet.[66] If adopted, this would have the effect of reducing the routing cost externality. Prospects for adoption of this technology are enhanced because it would offer additional benefits to service providers.[67]

3.2.2.2 Market Liquidity Failure

There is a risk that a transfer market could be established, but fail to attract a sufficient number of buyers or sellers. If there are too few buyers or sellers, the market will not be adequately liquid. This will distort pricing and reduce the likelihood that the market will deliver an efficiency benefits. The lack of an adequate number of buyers or sellers may be because transaction costs are too high (bid/ask spreads), which may be because of underlying costs of engaging in a transaction (accessing the market), the real scarcity of IPv4 addresses, or because of strategic behavior.

Search costs associated with matching buyers and sellers are an important example of the first type of costs. Given the legacy of private information in the Internet (e.g., associated with pricing and terms for transport and peering agreements and other Internet resources) and the large asymmetries in the cost structures of different sized ISPs, it is far from clear how information about IPv4 trades might become

[65] Transfers which only impact administrative ownership need not result in updates to BGP tables.

[66] See Farinacci et al. (2008), Vogt (2007) and Jen et al. (2007). However, these innovations may not be adopted before IPv4 addresses are exhausted.

[67] Four-byte BGP autonomous system numbers.

available to potential buyers and sellers. There is likely to be a large dispersion in pricing that may be largely unobservable to most market participants (and potentially all outside observers). Without an ability even to infer what the price for an IPv4 address is in the market, the signaling benefits of market pricing are largely lost. In light of such difficulties, it is necessary to consider the aggregate costs of acquiring the information via a market to signal efficient resource allocation versus the costs of the current regime with is reliance on non-market-based self-regulatory incentives to share information. A market, by increasing incentives to take advantage of asymmetric information and costs to extract rents, might make things worse.

A second important source of liquidity problems may arise because IPv4 addresses are really scarce. Presumably, because we anticipate demand for IPv4 addresses to persist into the indefinite future, the risk of too few buyers is less of a concern than the risk of too few sellers. The problem of too few sellers may arise because IPv4 addresses are so valuable that there are no sellers with excess addresses at any price that buyers are able/willing to pay. To assess this risk, it is necessary to understand how much slack exists in current allocations. Two metrics have commonly been used to estimate the quantity of surplus IPv4 addresses – and indirectly, the number of potential surplus IPv4 sellers in a transfer market. The first of these is the ratio of unique public IPv4 addresses that are visible across all *current* routing table entries,[68] compared to the *cumulative* number of IPv4 addresses allocated since their introduction in the early 1980s.[69] The second metric matches the ratio of *pingable* IP addresses against the "delegated" total.[70] The first ratio currently stands at just over 2:1, yielding a 50% utilization rate, while the second suggests that as many as 96% (or more realistically, 53%) of all IP addresses may be technically idle.[71] These metrics have prompted some to speculate that, given the opportunity, a lively market for IPv4 transfers might emerge and be self-sustaining over time. However, both of these metrics are severely flawed since neither controls for the IPv4 addresses that are actively used albeit within physically or cryptographically isolated networks. Although technically unobservable by third parties, the quantity of IPv4 resources that is actively used in ways like this is widely assumed to be considerable. In sum, while there is little doubt that an IPv4 transfer market (or in fact, *any* future course of events other than aggressive and immediate IPv6 adoption) would motivate many operators to pursue greater efficiencies in IPv4 utilization over current levels, considerable uncertainty remains as to how much slack exists within the universe of distributed IPv4 addresses today.

A third potential liquidity concern arises from the risk that an inappropriately designed transfer market might be vulnerable to strategic manipulation (capture, abuse of market power) or destabilizing

[68] The canonical source for such data is the University of Oregon Route Views Project, and associated Route Views Archive. http://www.routeviews.org

[69] This data is captured in "delegated" files that are published online by each of the RIRs on a daily basis, generally in the standardized location: ftp://www.<*RIR name*>.net/delegated/latest. Each daily snapshot presents a comprehensive view of all past IP distribution activities.

[70] http:// www.isi.edu/~johnh/PAPERS/Heidemann07c.pdf The gap between the former, raw figure produced by the ISI survey, and the latter, *HD*-adjusted upper limit is explained in n. 76, below.

[71] Note that it is unrealistic to assume that every single IP address is or could be assigned to a device or interface. The need for some amount of slack on each network to allow hosts to be added is a clear example of why IP addresses cannot be 100% efficiently used. The advent of mobile computing has only highlighted the point. Generalizations of these patterns were used to develop the *Host-Density Ratio* (or simply *HD*) as a more technically realistic unidimensional metric for evaluating the efficiency of individual or cumulative number resource number assignments. A HD ratio of 85-86% -- which would result from the entire 32-bit IPv4 address space being used to support between 150-200 million attached devices – would correspond to levels of utilization that were deemed to be "painful" enough to inspire renumbering efforts in the earlier cases. See RFC3194: The H-Density Ratio for Address Assignment Efficiency (http://www.faqs.org/rfcs/rfc3194.html).

speculation. Once again, because we anticipate demand for IPv4 addresses to persist into the indefinite future, the risk of too few buyers is less of a concern than the risk of too few sellers. Sellers might withhold supply because they are too uncertain about the prospect that the market will work as desired, that future transaction costs may be too high (e.g., search costs) or that future prices might be too volatile or increasing over time (so holding IPv4 addresses for future use or sale is believed better than selling today and potentially having to buy again in the future). In a well-functioning market, disparate expectations ought to cancel each other out; however, there is no *a priori* reason participants should expect this market to operate efficiently absent demonstrated experience of the markets efficacy. This is, in part, a chicken/egg problem.

Furthermore, there is the potential that those who *might* have excess IPv4 addresses today (perhaps because they face low costs in converting to IPv6, large legacy allocations, or limited prospects for growth) and little tolerance for participating in the market might quickly divest themselves of their excess supply. Many of the prospective buyers a market is intended to serve are future buyers who, by definition, are not presently in the market. Thus, there is the risk that such excess supply as might exist might rapidly be absorbed by incumbents, thereby increasing the market concentration of IPv4 addresses, and thereby potentially exacerbating the very problem that markets are intended to address.

Even absent market power concerns,[72] incumbents might be inclined to absorb the excess supply precisely because they fear the risk that if they do not buy, someone else might corner the market; or because they view such purchases as provident in light of the risk that the Internet community will fail to migrate to a workable long-term solution to the increasing scarcity of IPv4 addresses (e.g., via migration to IPv6 or to well-functioning IPv4 markets).

To address the liquidity concerns, we believe it is important that the design of the market seek to enhance the observability of transfer market transactions, and facilitate the matching of supply and demand. In principle, there are a number of ways in which this might be done. First, a market design that incorporates certain centralized features could enhance the opportunity for participants to share information about market pricing and supply/demand conditions more publicly, at least relative to a fully decentralized market of bilateral transfers. A centralized market, or perhaps multiple centralized markets (e.g., one corresponding to each RIR) might also operate in tandem with more decentralized transfers, in order to provide a mechanism for collecting information and influencing market prices without unduly constraining the flexibility of buyers and sellers to negotiate trades.

Since the RIRs have successfully administered the initial registration, registration data maintenance, and number-related WHOIS publication functions for many years, one could easily imagine them serving in similar roles supporting a centralized market arrangement. For example, the RIRs could create and maintain a bulletin board of bid/ask pricing to match prospective IPv4 buyers and sellers.[73] However, if the reduction in overall transaction costs created by such services does not substantially exceed the opportunity costs of having to "go to" the RIR, and abide by any rules or disclosure requirements that might be "imposed by" by the RIRs to achieve such transparency objectives, then it is highly unlikely that such RIR services would actually *cause* the market to spontaneously converge on a centralized form.

[72] If market power did arise (and was potentially worsened by the introduction of a transfer market), then that could provide another reason for market illiquidity. This is discussed further below.

[73] As Edelman (2008) suggests, the RIRs could play a more or less active role in matching buyers and sellers, depending on the needs/desires to preserve privacy or to manage the matching of transactions. For example, according to the ARIN proposal (see discussion below), transfers are approved conditional on a needs assessment. This needs assessment could be used to match potential buyers and sellers – assuming of course that something prevents buyers and sellers from simply opting out of the centralized market mechanisms and requirements.

Thus, barring the adoption of additional market controls by the RIRs, protecting the transfer market from the adverse impacts of liquidity problems and ensuring that the market offers appropriate informational signals will assuredly be problematic.

3.2.2.3 Market Power, IPv4 Address Price Escalation, & Foreclosure of Entry Risk

Another concern is that the creation of an IPv4 transfer market would enable market foreclosure and consolidation of incumbent market power. This might occur if the transfer market enabled incumbents to individually or collectively "corner the market" for IPv4 addresses. Shifting from the current administrative process for allocating addresses to one based on markets could contribute to this outcome in several ways.

First, the creation of transfer markets would provide IPv4-holding incumbents with rents from being the primary or possibly exclusive source capable of providing the assets (i.e., fee-simple IPv4 sales) and/or services (e.g., outsourced IPv6-to-IPv4 translation services) required to access IPv4-based resources. Second, incumbents may seek to limit access to IPv4 assets (or raise the cost of accessing such assets) to potential competitors, especially potential new entrants for who access to such assets is essential. Third, the establishment of a transfer market also might facilitate IPv4-holding incumbents opportunities to leverage any market power they may have into related markets by bundling access to IPv4 resources with content, applications, or IPv6-to-IPv4 translation services. Fourth, IPv4-holding incumbents that may derive benefits from their dominant position in the IPv4 world may use their market power to preserve their strategic positions and the status quo. This could adversely distort the trajectory of technical change, including the transition to IPv6 or other innovations.

While the authors disagree on the likelihood of market power being asserted in a prospective transfer market (as already discussed earlier with regards to the impact on the IPv6 transition), they do agree that the risk warrants serious consideration. Even if one believes that the risk is small, the potential for harm to the Internet and all it supports is too great to disregard.

One possible – and in the authors collective judgment -- promising option for reducing the risks of market closure/IPv4 lock-in is the establishment of a neutral market-maker or provider of last resort for IPv4 addresses, specifically for new entrants.

3.2.2.4 Privatization of IPv4, Loss of Common Pool, and Related Policy Mechanisms

Except in (typically brief) moments following their invention, Internet protocols, technologies, and technical coordination mechanisms seldom exist in a legal or regulatory vacuum, and IP addressing is no exception to this general rule. Nevertheless, the conventions and practices by which IP addresses have been distributed and managed to date have been largely determined by ad hoc, open communities of interest, with the largest segment of which typically coming from direct stakeholder groups like IP address "assigners" – i.e., ISPs that own and operate hardware that, in conjunction with IP addresses, can be attached to and thereby contribute to the overall delivery of online content and services that makes the Internet valuable. The term usually used to denote this kind of internal regulatory arrangement is "industry self-governance," or simply *self-governance*.

In order for self-governance to be meaningful in any context, its specific domain of application must be substantially free of other purposive regulatory authorities. The current regime of Internet governance, of which the RIR address allocation system is a part, rests on a delicate balance of a community of mutually accepted shared interests in preserving the Internet. The institutional bodies such as IANA and the RIRs which collectively play a role in managing the Internet lack the authority of sovereign governments or even international treaties (such as those that empower the ITU or international monetary agreements), which makes them vulnerable to disruption. The introduction of transfer markets might threaten the delicate balance on which Internet governance rests today.

One way a transfer market might pose such a threat is if it moved us from regarding IPv4 addresses as a common pool resource to a private property regime with its attendant regulatory constructs (rights, obligations, and enforcement mechanisms). This would not need to happen all at once or as a result of a specific action to "privatize IPv4 addresses." Rather, it could happen incrementally. In all likelihood the gradual sublimation of regional address policies and communities would begin with the material fact of a private exchange of IPv4 addresses for money. If IPv4 continues to be an essential, non-substitutable input for Internet service delivery, then sooner or later transactions like this that are not clearly mediated by a non-interested third party (i.e., something like an RIR) are likely to occur with some regularity. Whereas transactions in the context of such mediation might retain enough of the RIR's former special status to resist attributions of de facto privatization -- whether made by buyers, sellers, or third parties – unmediated transactions may retain few features meriting continued special status or immunity from such claims. For example, will accountants and tax officials regard such transactions as capital expenditures or operating expenses?[74] Will IPv4 purchased in one jurisdiction and assigned in another be subject to tariffs and/or transfer pricing restrictions? Will they be subject to sector-specific direct investment requirements and/or restrictions? Will privacy restrictions preclude the publication of WHOIS data, or alternately mandate the publication of data deemed to be privacy infringing in other jurisdictions? Commercial network service providers might prefer for IPv4 addresses to become private property and may seek to advance any such trend that may emerge. Absent appropriate mediation policies from RIRs that distinguish the terms of use of IP addresses, resolving uncertainties like this would almost certainly result in legal disputes, each of which would likely have to reconcile a variety of latent ambiguities inherent in the novelty of Internet technologies and the opacity and orthogonality that formerly characterized the relationship between distributed Internet production systems and national legal jurisdictions.[75] Any such ruling might establish a precedent for private property status, thus creating uncertainty about how IPv4 should be treated in other contexts.

4 Current Proposals Under Consideration for Reforming Address Allocation

In several of the RIRs, proposals are currently under consideration to modify RIR policies to enable the transfer of IPv4 addresses that have already been allocated. These policy changes would enable the creation of transfer markets, or in some cases, would segregate some address resources for special allocation treatment. Transfer proposals currently under consideration include the following:

- North America (ARIN): ARIN Advisory Council, "Policy Proposal 2008-2: IPv4 Transfer Policy Proposal," March 7, 2008, available at: http://www.arin.net/policy/proposals/2008_2.html ("ARIN 2008-2"). See, also, Durand, Alain, "Policy Proposal 2008-5: Dedicated IPv4 Block to Facilitate IPv6 Deployment," June 6, 2008, available at: http://www.arin.net/policy/proposals/2008_5.html ("ARIN 2008-5).

- Asia Pacific (APNIC): Huston, Geoff, "prop-050-v002: IPv4 address transfers," January 22, 2008, available at: http://www.apnic.net/policy/discussions/prop-050-v002.txt. "APNIC Prop50." See also, Smith, Philip, Jonny Martin, and Randy Bush, "Use of final /8," 15 July 2008, available at: http://www.apnic.net/policy/proposals/prop-062-v001.html ("APNIC Prop 62").

[74] A potential albeit roundabout *benefit* of viewing IP address blocks as assets would be the requirement in at least some jurisdictions that the purchase be disclosed to taxing authorities; information that regulators could in the future access to assess the market on an ongoing basis.

[75] See Lessig (2000).

- Europe (RIPE): Titley, Nigel and Remco van Monk, “RIPE Policy Proposal 2007-08 Enabling Methods for Reallocation of IPv4 Addresses,” October 23, 2007, available at: http://www.ripe.net/ripe/policies/proposals/2007-08.html (“RIPE 2007-08”).

- Latin America and the Caribbean (LACNIC): Presta, Ricardo, “LAC-2008-04: Special IPv4 Allocations/Assignments Reserved for New Members,” 24 April, 2008, available at: http://www.lacnic.net/documentos/politicas/LAC-2008-04-propuesta-en.pdf. (“LACNIC 2008-04”).

Mueller (2008)[76] provides a good introduction to the three most discussed proposals (ARIN 2008-2, RIPE 2007-08, and APNIC Prop50) as follows:

	ARIN	RIPE	APNIC
Trigger for starting	Last IANA block allocated	No trigger	No trigger
Relationship to RIR	Both seller and buyer must be resident in ARIN territory, and the addresses must be used in that territory.	Address space may only be re-allocated from a RIPE NCC member to another member of the RIPE NCC.	The address block must be administered by APNIC and allocated or assigned to a current APNIC account holder.
Eligibility	Any ARIN member	Only LIRs; end users are not eligible	Any APNIC member
Trafficking restrictions	Seller cannot have received any IPv4 addresses from ARIN or from transfers in the past 24 months. Seller cannot request any for the next 24 months. ARIN decides how much supply the buyer gets.	Buyer cannot re-allocate complete or partial blocks of the same address space for 24 months.	Seller cannot receive any IPv4 addresses from APNIC for a period of 24 months. Future requests to APNIC must be justified.
Need assessment of recipient	Buyer must "pre-qualify" to be eligible. Buyer must justify both existing allocations and the amount transferred.	None	No need assessment unless additional requests for IPv4 addresses made
Fees	Buyer pays a transfer fee in addition to normal fees associated with all addresses held	No transfer fee mentioned. Re-allocated blocks no different from the allocations made directly by the RIPE NCC	Buyer pays a transfer fee in addition to normal fees associated with all addresses held
Aggregation	/24 minimum. Detailed regulation of way in which address blocks are downsized	A /21 is the minimum size	A /24 is the minimum size

Table 1: Comparison of RIR IPv4 Address Transfer Policies

Mueller (2008) argues that the introduction of transfer markets is a good direction for Internet management to follow, criticizing the reliance of RIR policy on common pool resource management models as the right mechanism for managing IP address resources. Mueller offers critiques of each of the proposals, but on the whole favors RIPE 2007-08 as coming closest to striking an appropriate balance in favoring simplicity and minimal regulatory oversight (e.g., no needs-based assessment required) but being flawed in prohibiting end-users from engaging in transfers (something the other proposals do allow). Mueller views ARIN 2008-2 as being overly and unrealistically regulatory of both buyer and seller

[76] Reproduced from Mueller (2008), Table 1, page 16. Because the proposals are under active consideration, the list of proposals and their specific attributes may change over time.

behavior. He criticizes APNIC Prop50 for failing to impose any restrictions on the buyer. While Mueller's analysis is insightful, it is incomplete (in part because policy in this space is moving so rapidly). He does not address complementary policies under consideration by several RIR communities (and already approved by one) that are designed to minimize potential adverse impacts that transfer policies might have on the industry's openness to new entrants.

All three of the proposals preserve a role for the RIR in on-going management of the IPv4 address space, even after the unallocated pool is exhausted. At a minimum, this involves preserving the RIR as the equivalent of a title registry to record address transfers among approved buyers and sellers, and as a maintainer of legacy WHOIS records. Since all of the proposals limit approved transfers to participants within the covered region, each also has the effect of preserving some scope for the RIRs to continue serving as a regional coordinating mechanism for policies that are unaffected by the new transfer mechanism, and perhaps for future regional policy development activities. In the absence of such prohibitions on inter-regional IPv4 sales, transfer transactions might have had the politically untenable effect of precipitating an outflow of IPv4 from developing regions and an increased IPv4 concentration levels in the more advanced industrial regions. Thus, these provisions may also be regarded as providing some secondary market-stabilizing protections. Longer term, we would hope there would be convergence on a common set of transfer rules and the possibility of inter-region transfers may be a further refinement that could be examined.

The RIPE proposal mandates only that any prefixes transferred must be equal in size or greater than /21 (the RIPE minimum allocation), whereas the APNIC proposal allows for any block of addresses from /24 and greater to be transferred. These differences may reflect the varying philosophies of respective proponents, as well as private assessments regarding just how much enforcement power the RIRs actually have.[77] Anti-speculation rules are embodied in restrictions on the buyer (and/or seller) from engaging in subsequent transfer transactions for a period of (typically) 24 months. In light of the negative externality associated with routing table growth from ill-managed address allocations and the other concerns noted earlier, the need for such restrictions seems sensible.

Considered in the broadest possible context, it must be observed that only three of the five RIR regions are currently considering an IPv4 transfer proposal of any kind. To date, no such proposals have been considered in the LACNIC or AfriNIC regions, and recent discussions with associated RIR staff suggest that none is anticipated there. Given the fact that many Internet services providers in Africa and Latin America still make extensive use of overseas IPv4-based Internet resources, varying interest levels in resource transfers no doubt also reflect the skewed distribution of aspiring IPv4 transferors in the more developed ARIN, RIPE, and APNIC regions.

Similarly, of the three regions currently considering a transfer proposal, two (ARIN and APNIC) are also considering related, arguably countervailing policy proposals intended to preserve some stock of IPv4 resources to limit the impact of market pricing on new entrants or those transitioning to IPv6. In addition to these, a third reservation proposal on this model has already been accepted by LACNIC, where no resource transfer market is currently contemplated.[78] If approved, these proposals would guarantee that the RIR for the corresponding region remains in business as a distributor of IPv4, albeit perhaps for a

[77] For example, the author of RFC 1744 (1994), which envisioned a commercial IPv4 trading house as preferable to the then-emerging RIR system, is also the author of APNIC Prop50, the most market-oriented of the current proposals.

[78] According to APNIC prop62, "…At the recent LACNIC XI meeting, consensus was reached on the following policy proposal: LAC-2008-04: Special IPv4 Allocations/Assignments Reserved for New Members." Additional administrative reviews may be conducted before final approval. See http://www.apnic.net/policy/discussions/prop-062-v001.txt.

narrower range of eligible recipients (new entrants or those beginning the transition to IPv6). Assuming that both the resource transfer proposals and reservation proposals are approved, this combination would provide for ongoing mechanism-level competition for ARIN and APNIC members only. However, a more modest and temporary form of competition in mechanisms may result from differences in the timing of the commencement of market-based transfers envisioned in different policy proposals. The APNIC Prop50 proposal specifies that transfers will be allowed only after the exhaustion of the IPv4 unallocated pool, while RIPE 2007-08 and ARIN 2008-2 transfers could start as soon as the policies are approved. Because IPv4 addresses are already scarce, it is unclear why the APNIC Prop50 restriction is appropriate. If transfers are deemed appropriate, and measures are in place to prevent the arbitrage-accelerated exhaustion of RIR reserves, then allowing transfers to begin sooner rather than later would provide an additional opportunity for competition in mechanisms to work in the affected regions.

Another significant variance in prospective benefits arises from the divergent approach to ongoing needs-based assessment envisioned by different transfer proposals. At present, only ARIN's proposal preserves the requirement of a needs-based assessment for approval of transfers. There are several motivations for such an approach. First, retention of the needs-assessment for acquiring new addresses is consistent with the current approach and so represents the most incremental innovation towards approving transfer markets of the three RIR proposals. Second, the needs assessment offers a further set of tools for managing a prospective transfer market. A needs assessment may help to (a) limit speculation, (b) enhance the information available to the RIRs about transfer market conditions, and (c) provide the RIRs with a tool for detecting/preventing strategic manipulation of the transfer markets by participants.

For those who believe that markets are a good thing (Mueller, 2008), preserving the needs assessment may seem unnecessarily burdensome. For those who are concerned about the potential for markets to disrupt a well-functioning Internet, this incremental approach may be seen as appropriately prudent. However, despite its intended goal of preserving some measure of the status quo mechanism, the needs-based assessment may have unintended consequences, in that it may elevate tensions between large incumbents, small and medium-sized players, and new entrants. For example, if the transfer market is characterized by a preponderance of legacy address sellers and very large prefixes, then a needs test might effectively leave aspiring new entrants or IPv4-hungry small operators competing for an insufficient supply of smaller prefixes. If prohibitions against de-aggregation hold, sellers of large blocks may choose not to enter such a market at all. Alternately, if only small prefix sellers participate in the market, large prefix seekers may face incentives to amass cumulative address space by buying multiple small prefixes, thereby directly competing and possibly displacing new entrants and smaller operators. If a needs test makes such resources unavailable in the sanctioned market, alternative sources may be sought out.

While nothing may absolutely preclude the most determined IPv4 buyers and sellers from abandoning the approved transfer process in favor of the black market, the "final reservation" proposals under consideration in the ARIN and APNIC regions could substantially offset such risks for some service provider segments. The ARIN 2008-5 proposal is intended to create a mechanism for ensuring that the RIR retains a pool of unallocated addresses into the indefinite future to allow a supply of IPv4 addresses to assist in IPv6 migration. If approved, this proposal would reserve a relatively large block of IPv4 addresses (the size is still under negotiation) to be made available exclusively for purpose of facilitating the inter-operation of new IPv6-based entrants with the existing universe of IPv4-based hosts that constitute the Internet. Accordingly, aspiring recipients would be obliged to justify the very small (e.g., /28 to /24) allocations that the policy affords in terms of a specific IPv6-based operational or migration strategy. The APNIC version of this reservation strategy, embodied in APNIC prop62, takes a slightly different approach by affording one last chance for any interested APNIC member to secure one final /21 allocation from a reserve block of IPv4 to be set aside for that purpose. Significantly, these final RIR allocations are not subject to traditional needs-based eligibility restrictions. Finally the LACNIC proposal, which garnered community consensus at the May 2008 LACNIC XI meeting, envisions a broadly similar goal, but contains no explicit IPv6 requirements. It provides for allocations that are larger than the ARIN proposal but smaller that APNIC's (/24 to /22), with priority given to new entrants an established ISPs that may be obliged to vacate address space provided by upstream service providers.

Each of the reservation proposals may be interpreted in several ways. Each could be viewed as providing a competing mechanism to those envisioned by the corresponding transfer market proposals. In the case of ARIN 2008-5, the clear intent is to create a sharply contrasting alternative mechanism that focuses on the need to smooth the transition to IPv6 rather than allowing IPv4 address transfers. As such, it refocuses the question from IPv4 address scarcity toward IPv6 transition, which seems to be the more relevant long-term issue. While APNIC prop62 might advance the same IPv6 transition-oriented goals, the absence of a needs test or any IPv6-related requirements might also make it a vehicle for channeling any existing incentives for arbitrage to help kick start the regional transfer market. Since there is no transfer proposal under consideration in the LACNIC region, LAC-2008-4 might be regarded as a straightforward continuation of the traditional methods for distributing IP addresses under conditions of increasing IPv4 scarcity.

Alternately, reservation proposals like ARIN 2008-5 and APNIC prop62 could be viewed as complementary to any market-based transfer proposal. To illustrate, ARIN 2008-5 could provide a foundation for creating a market-maker role for the RIR. Creating such a role may offer several benefits. First, as long as the RIRs have a pool of IPv4 addresses, they retain some capacity to influence the behavior of some market actors, which might be useful in a wide variety of future scenarios.

Second, this proposal could enhance market liquidity, if for example the size of the associated IPv4 reservation is large enough to assure a long stream of new IPv6-based entrants, and thereby encourage IPv4 surplus holders to enter the market sooner rather than later. Given continuing commercial apprehensions about IPv6, and the uncertain results that would ensue if ambivalence to IPv6 hardens into absolute resistance, an IPv6-oriented reservation could help to align expectations about the future in ways that will not only reduce market volatility, but also increase the odds of eventual IPv6 adoption. Finally, if a transfer market does not work as expected or if the rules or outcomes are biased against new entrants, the ongoing ability of the RIR to tap a reserve of addresses for allocation to entrants or otherwise disenfranchised participants would offer a safety-valve. To play the latter role, ARIN 2008-5 would need to be modified since allocations made for such purposes would not be limited to IPv6 migration. In order to play such a role on a sustainable basis, it would likely be necessary to provide for an on-going way for the RIR to participate in the transfer market – for example, as a potential "buyer" – to secure addresses for subsequent reallocation. Were such a direction to prove worthy of further exploration, a number of additional details would need to be worked out.

Overall, these considerations suggest that reserving a portion of the unallocated pool of IPv4 addresses for special allocation consideration, as proposed by LAC-2008-04, ARIN 2008-5, and APNIC prop62 represents a prudent step in the right direction, as a hedge against unknown future challenges.

Looking beyond the various, unevenly distributed benefits that might be provided by IPv4 transfer markets, all of the proposals suffer from a few common limitations, which give rise to common risks. Among these, the risk of accelerating routing table growth seems the most likely. However, the authors concur that, in the near term, transfer markets are not likely to greatly accentuate this risk over and above the fact of IPv4 exhaustion itself.

The authors also concur that further review of the RIR transfer proposals is warranted before moving forward. We view the fact that the current transfer proposal advocates have not given careful consideration to the need to assure the continuity of various critical functions currently supported by the RIRs as a major deficit that should be corrected before moving forward. Even if one accepts, for example, Mueller's view that the common pool resource management approach is less appropriate than a market approach for managing IP addresses, this does not imply that these two polar views exhaust all of the possible approaches, much less that the right way to introduce markets has yet been identified. A bad move may be hard or impossible to reverse.

All three authors agree that introducing markets creates new demands and poses new risks for the future of both Internet governance and the Internet's technical evolution, neither of which have been adequately considered by market advocates. For example, none of the RIRs as currently constituted possesses clear authority or adequate means to enforce any proposed rules and restrictions on transfer markets. Once transfers are allowed, moreover, it is unclear what means the RIRs will possess to effectively discipline buyers or sellers who choose to ignore trading restrictions they find unacceptable, or to encourage compliance with any policy by any transfer recipient anytime thereafter. As noted at several points, the current regime has been sustained by the community of shared interests and by the need to return to the RIRs for subsequent allocations of IPv4 addresses. Introduction of a transfer market opens the door to the opportunity to bypass the RIR process, thereby potentially disrupting important components of the self-enforcement mechanism.

This lack of adequate enforcement capabilities raises questions in particular about the ARIN proposal's market participation rules (i.e., needs-based assessment) and product restrictions (i.e., prefix sizes and origins). Proposal advocates point to the self-evident advantages of the official market itself, and to the control that ARIN still exercises over the WHOIS service as factors that are sufficient in themselves for the sanctioned market to attract and contain transfer transactions over time. To the authors these assumptions seem unrealistic. That said, the best response here may not be simply to dispense with the assessment, but rather to think more carefully about what might be done to capture what is desirable about the needs assessment process.

The restrictions on block size, motivated primarily by the desire to preserve address aggregation, and the restrictions on trading limits, motivated by a desire to promote stability and discourage speculation, may have unintended implications for who can participate in the transfer markets. As constituted, the rules would appear to favor participation by larger incumbents. For example, while a large provider might justify a /13 or a /14's worth of additional address space based on existing utilization, a new entrant would only justify a smaller block, such as a /20.[79] Furthermore, sellers of large blocks are limited to selling to large providers because the proposals place severe restrictions on parceling. While such restrictions might facilitate the emergence of convex pricing in a highly disciplined market – and thus might preserve the affordability of small IPv4 prefixes for new entrants and network enterprises – the discipline required to produce this effect seems unrealistic given the opacity of existing Internet input markets, the limited capabilities of the RIRs, and the absence of a credible expectation that third party or parties will emerge to offset these limitations.

In this sense, the RIPE and APNIC proposals would appear to err in preserving too limited a role for the RIRs in a transfer market. Their function is limited to the registration of transfers and the specification of a few rules to control allowed transfers. ARIN on the other hand errs in being too controlling. Although each of these judgments might be substantially altered by the adoption of a IPv4 reservation proposal of some kind in conjunction with a transfer proposal, consideration of the transfer proposals alone would suggest a the following relative mapping, based on the taxonomy introduced previously:

---

[79] See http://www.arin.net/policy/nrpm.html, Section 4.2.

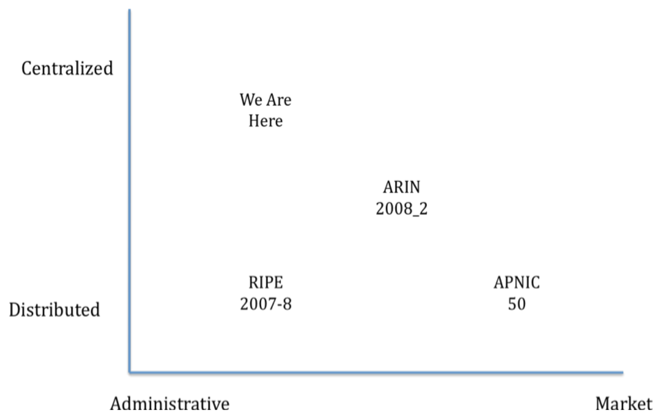

Figure IV: Map of Proposals

Finally, in light of the information asymmetries already discussed, it is unclear how the RIR approved transfer markets will enable information sharing about transfer pricing. In the absence of such market information, the financial incentive imposed by a market-based opportunity cost for using IPv4 addresses inefficiently will be attenuated and the potential risk that transfer markets might be abused by incumbents with market power for strategic reasons may be heightened. Even the ARIN proposal does not fully specify how ARIN might alter its needs assessment process in light of transfer requests or will publicize information about approved transfers. Edelman (2008) suggests it might be appropriate for ARIN to play a role in matching potential buyers and sellers.

Beyond these universally observed risks, consensus among the authors about specific risks is more tenuous. The authors of this paper do not agree among themselves as to the merits of address transfer markets. All authors believe that policies will need to be adapted to accommodate the reality of continuing demand for IPv4 addresses after the exhaustion of the unallocated reserves. Two of the authors (Lear and Lehr) believe that current transfer proposal drafts provide a reasonable starting point for necessary reforms; while none may be perfect, in general the market-based transfer mechanisms, lightweight restrictions, and assumptions about market behavior that they embody provide the best current option for moving forward given the inevitability of IPv4 exhaustion. The dissenting author (Vest) believes that the market-based approach to IPv4 transfers proceeds from a fundamental misunderstanding of the function that IP addresses play in the Internet, and thus while some combinations of policies are better than others (e.g., where transfers are coupled with a reservation policy of some kind), all necessitate acceptance of unnecessary risks, including greater transitional volatility and the possibility of suboptimal but irreversible outcomes. These views are presented elsewhere.[80]

5 Resource Certification: A Technical Wildcard?

Before concluding, there is one more significant variable that deserves mention. Responding to rising global concerns about Internet security and stability, two years ago researchers associated with several RIRs and other Internet technical coordination institutions launched an initiative to develop an incrementally deployable technology to enhance the security of Internet routing. The outgrowth of this effort, often described as Secure Inter-Domain Routing (SIDR), is now nearing the test deployment stage of development, and may be ready for production soon after this paper is presented.[81] SIDR builds on two pre-existing routing security models, with primary inspiration coming from the more lightweight model,[82]

---

[80] See, Tom Vest's website, http://www.eyeconomics.com.

[81] For a recent project update, see http://www.potaroo.net/presentations/2008-03-12-resource-certs.ppt

[82] See http://www.ir.bbn.com/sbgp/ and ftp://ftp-eng.cisco.com/sobgp/index.html.

which enhances security by permitting recipients of Internet traffic to legitimate the origin of that traffic on a packet-by-packet basis. Significantly, validating the origin of traffic entails knowing the true, proper associations between both IP addresses and Autonomous System Numbers (or ASN, a special institutional identifier that serves as an origin, destination, and/or intermediate "hop" in Internet routing), and the institutions to which they are currently assigned. The most secure place to begin assembling that information is the place where each resource began its gradual, stepwise journey down the IP address delegation path – i.e., the RIRs. Thus, if SIDR is ultimately adopted on a widespread basis, there is broad consensus (but not unanimity) that the most logical place to "anchor" distinct regional *validation chains* would be with the RIRs that constitute the top of the address resource delegation hierarchy within the associated region.

Assuming that SIDR does enjoy widespread adoption, and that the RIRs become the preferred site to anchor regional validation chains on an ongoing basis, this new role could reinvigorate RIR capabilities, enabling them to continue fulfilling their traditional role as maintainer and source of both number resources and primary registration data (the public presentation of which is WHOIS). In principle, the RIRs could go even further, leveraging their status as top-most authorities in the SIDR "trust hierarchy" to exercise real power. For example, in their role as the trust anchor, an RIR could threaten or even revoke the top-level certification of an IP prefix or ASN – perhaps for policy non-compliance – which would thereby render it suspect, possibly even prompting some other network operators to refuse to route it. Although the exercise of such a blunt instrument – which would be equivalent to repossessing an airplane in mid-flight – would hardly represent an effective or sustainable policy enforcement model -- concerns about even a potential to exercise this kind of power, no matter how irrational and unlikely it would be, may create enough anxieties to substantially impact SIDR's acceptance by independent network operators. Only time will tell whether and how this development will alter the incentive calculations and related, notional policy recommendations outlined above.

6 Conclusions

The Internet is now critical infrastructure for our global economy. That it continues to grow and operate as well as it does, while relying on technologies and a distributed/decentralized resource management framework that was put in place when the Internet was far smaller and less commercially salient is a testament to the success and robustness of the Internet's design. The IPv4 Internet address framework is an integral part of this structure that has already persisted far longer than was originally anticipated. Today and for the foreseeable future being connected to the Internet will continue to mean being able to route traffic to and from IPv4 addressable resources.

Unfortunately, the number of IPv4 addresses is finite and the legacy framework for meeting growing demand for new address space is approaching a time when it will be unworkable. *The tank of IPv4 addresses that have not already been allocated is nearly empty.* To address this problem of perceived IPv4 address scarcity, it is reasonable to ask whether modifying the existing address management framework so as to more directly incorporate market mechanisms might not offer a significant improvement. A number of proposals are currently under consideration in the RIRs in the North American (ARIN), European (RIPE), and Asia-Pacific (APNIC) regions that would change RIR policies to allow transfer markets in IPv4 address space that had previously been allocated.

The problem these policies seek to address (IPv4 address scarcity) and the proposals under consideration (to create IPv4 transfer markets) have important implications not just for the future of IPv4 address management, but more broadly for the future architecture of the Internet (and the pace of IPv6 transition) and IP governance. With respect to the latter, the transition toward market-based resource allocation would represent a change from the decentralized common pool resource/self-regulatory framework that the Internet currently operates under. Given the importance of these issues, we believe further consideration is needed to appropriately address this mechanism design challenge.

In this paper, we provide context for understanding the key features of the address management challenge and the rationales behind alternative mechanisms for introducing market forces – not just as articulated in the proposals currently under consideration, but more broadly, over the space of potential proposals. Our analysis has lead us to a number of conclusions, but not to agreement on what we think would be the best way to implement address transfers (or even whether such transfers are advisable).

First, we agree that the current regime for address management needs to be reformed. The premise of an on-going supply of unallocated IPv4 addresses with which to meet future demand for address space and as a mechanism for inducing cooperation with collective oversight of address use practices and routing table management via the RIRs depends on the RIRs having a pool of addresses to allocate.

Second, while we agree something needs to be done, we are not clear on precisely what should be done. Two of us (Lear and Lehr) believe that some sort of transfer market is desirable and likely unavoidable, while the third (Vest) agrees that change is unavoidable, but is less sure that market mechanisms represent the best approach. Our internal debates in writing this paper have identified numerous open issues without simple solutions. For example, a transfer market will need enforcement capabilities and may destabilize or render unworkable the current regime of RIRs. None of the proposals or research we have seen adequately addresses these issues.[83] While we have identified problems, we have not offered full solutions.

Third, we agree that the RIRs likely should play an on-going role in the management of Internet addresses in any future in which address transfer markets are allowed. The RIRs provide the natural nexus for registering and sharing information about address assignments, and offer continuity with the institutional arrangements that have heretofore served the Internet community so well. The registration function may be rendered even more important in the future if certain enhancements to the routing infrastructure are implemented.[84]

Furthermore, the RIRs may play a useful role in helping to stabilize the market and provide a provider-of-last-resort “safety valve” function to meet on-going demand for IPv4 addresses into the indefinite future. We believe further consideration should be given to how the RIRs might interact with a transfer market in order to replenish their stock of allocable IPv4 address space, to best share information with market participants to maximize the benefits offered by market processes in enhancing resource conservation incentives, and to help manage an orderly transition to IPv6 and routing table growth. The RIRs may play an important role in helping to manage the large negative externality associated with inefficient use of routable address space. The externality problem is not limited to IPv4 addresses, but will persist and may even worsen with the transition to IPv6 unless appropriate address management (and associated routing policies) can be developed. Solving these challenges will require a mix of technical, market (industry/business model), *and* policy innovations.

Fourth, we believe that inappropriately implemented transfer markets might very well slow the transition to IPv6. However, we do not believe that there is a fundamental inconsistency between supporting the

[83] See the discussion in the preceding section for specific critiques of the each of the proposals currently under consideration. In addition to the websites and listserves of the RIRs that discuss these issues, several recent papers by Perset (2007), Mueller (2008), and Edelman (2008) provide valuable additional analysis and alternative perspectives on the issues considered herein.

[84] For example, LISP or SBGP.

emergence of IPv4 transfer markets and a speedy transition to IPv6. The impact of appropriately implemented transfer markets on the timing of IPv6 transition is ambiguous. To better understand the implications of IPv4 transfers on the pace of IPv6 transition, a more complete analysis of IPv6 stakeholder interests and adoption/transition costs is needed. We sketch out the beginnings of how to think about this as a collective S-shaped adoption process, complementing the work of Elmore, Camp, and Stevens (2008). One direction for possibly extending this work is to embed the private calculus for conversion to IPv6 into a bandwagon dynamic game framework in the presence of externalities.

Fifth, and finally, while economic theory suggests that markets often provide improved performance when it comes to managing the allocation of scarce resources, they are not the only mechanism worth considering and, in any case, do not exist in isolation. The current regime of Internet self-regulation of a common pool resource has served the Internet well to date. Allowing decentralized address transfers may be desirable or inevitable, but it is not a small change. Careful consideration of the structural design features of the transfer market, with the potential of erring on the side of incremental change, seems advisable in light of the risks and potential irreversibility of a decision to move toward a world of addresses that might be traded as private property.

7 References

Arrow, Kenneth and Leonid Hurwicz (1960), "Decentralization and computation in resource allocation," in R.W. Pfouts (editor) *Essays in economics and econometrics in honour of Harold Hotelling*, University of North Carolina Press, Raleigh, North Carolina, 1960.

Butt, Danny (editor) (2005), *Internet Governance: Asia-Pacific Perspectives*, Asia-Pacific Development Programme (APDIP), available from: http://www.apdip.net/publications/ict4d/igovperspectives.pdf.

Coase, Ronald (1960), "The Problem of Social Cost," *Journal of Law and Economics*, Vol. 3 (October 1960) 1-44.

Day, John (2008), *Patterns in Network Architecture: A Return to Fundamentals*, Prentice Hall: New York, January 2008.

Demsetz, Harold (1967), "Toward a theory of property rights," *American Economic Review*, Volume 57, Issue 2 (May 1967) 347-359.

Edelman, Benjamin (2008), "Running Out of Numbers: The Impending Scarcity of IP Addresses and What To Do About It." Mimeo, Harvard Business School, June 10, 2008, available at: http://www.benedelman.org/publications/runningout-draft.pdf.

Elmore, Hillary, L. Jean Camp, and Brandon Stevens (2008), "Diffusion and Adoption of IPv6 in the ARIN Region," paper presented at Workshop on the Economics of Information Security (WEIS2008), June 25-28, 2008, Hanover, New Hampshire (available at: http://weis2008.econinfosec.org/papers/Elmore.pdf)

Farinacci, D., V. Fuller, D. Oran, D. Meyer, and S. Brim (2008), "Locator/ID Separation Protocol (LISP)," IETF Network Working Group, Internet Draft, July 10, 2008, available at: http://tools.ietf.org/html/draft-farinacci-lisp-08.

Hongqiang, Yan (2005), "China Internet Status Quo and International Internet Governance," presentation slides at Internet Protocol Symposium for Asia and the Pacific 29-30 November 2005, Bangkok, Thailand, hosted by United Nations Economic and Social Commission for Asia and the Pacific (ESCAP) (available at: http://www.unescap.org/icstd/events/IP_Symposium/presentation/china.ppt).

Huston, Geof (2008), "Projected IANA unallocated address pool exhaustion," available at: http://www.potaroo.net/tools/ipv4. As of 7/29/08, IANA exhaustion predicted February 2, 2011 and RIR pool exhaustion December 25, 2011.

Jen, D., M. Meisel, D. Massey, B. Zhang, and L. Zhang (2007), "APT: A Practical Transit Mapping Service," IETF Network Working Group, Internet Draft, November 18, 2007, available at: http://www.cs.arizona.edu/~bzhang/paper/draft-jen-apt-01.txt.

Karrenberg, D. and Tom Vest (2008), "Concentration in IPv4 Address Space Distribution," Proceedings of RIPE 56, May 2008.

Lessig, Larry (2000), *Code and Other Laws of Cyberspace*, Basic Books, July 2000 (see http://code-is-law.org).

Li, Victor (2001), "Is Why We Use Money Important?" *Economic Review – Federal Reserve Bank of Atlanta.* Vol. 86, No. 1 (Q1 2001), pp. 17-32.

Mueller, Milton (2008), "Scarcity in IP Addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries," white paper prepared for Internet Governance Project, 20 July 2008, available at: http://internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf.

NIST (2005), "Planning Report 05-2: IPv6 Economic Impact Assessment," report prepared by RTI for National Institute of Standards (NIST) of the U.S. Department of Commerce, October 2005 (available at: http://www.nist.gov/director/prog-ofc/report05-2.pdf).

Perset, Karen (2007), "Internet Address Space: Economic Considerations in the Management of IPv4 and in he Deployment of IPv6," Organization for Economic Co-operation and Development (OECD), Ministerial Background Report, DSTI/ICCP(2007)20/Final, public release March 2008 ("OECD 2007 IPv4/IPv6 report").

Short, Jodi L. and Michael W. Toffel (2008), "The Causes and Consequences of Industry Self-Policing," *Yale Economic Review* (Volume 4, Issue 2: Summer 2008), pp. 32-40

Vogt, C. (2007), "Six/One: A Solution for Routing and Addressing in IPv6," IETF Network Working Group, Internet Draft, November 19, 2007, available at: http://tools.ietf.org/id/draft-vogt-rrg-six-one-01.txt.

Williamson, Oliver (1975), Markets and Hierarchies: Analysis and Antitrust Implications, The Free Press, New York, 1975.

8 Biographies of authors:

Dr. William Lehr is an economist and industry consultant. He is a research associate in the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology, currently working with the Communications Futures Program (http://cfp.mit.edu), which is an industry-academic multidisciplinary research effort focused on road mapping the communications value chain. Dr. Lehr's research focuses on the economics and regulatory policy of the Internet infrastructure industries. He is currently engaged in research on the convergence of the Internet and telecommunication services and the implications for corporate strategy and public policy.
In addition to his academic research, Dr. Lehr provides litigation, economic, and business strategy consulting services for firms in the information technology industries. Dr. Lehr holds a PhD in Economics from Stanford (1992), an MBA from the Wharton Graduate School (1985), and MSE (1984), BS (1979) and BA (1979) degrees from the University of Pennsylvania.

Contact:
William Lehr
Massachusetts Institute of Technology
32 Vassar Street (32-G814)
Cambridge, MA 02139
(tel) 617-258-0630
(email) wlehr@mit.edu

Tom Vest is an Internet technology professional with extensive public, private, and nonprofit sector experience designing, provisioning, building, and interconnecting high capacity production IP networks in Europe, Asia, Oceania, and the Americas. He currently serves as Research Consultant to RIPE NCC, the European Internet Registry, where he reports to the Chief Scientist and collaborates with the newly formed Science Group on strategic and long-range Internet technical and institutional issues. The former Tokyo-based Senior Network Operations Manager for America Online Inc., Tom was responsible for design, deployment, local management, and interconnection/peering for AOL's Transit Data Network (ATDN) in China, Australia, and Japan. Before moving to Tokyo, Tom was responsible for coordinating AOL's international peering and interconnection efforts, and for planning, implementing, and interconnecting remote points of presence and distributed hosting and caching complexes in Europe and South America. He previously served as senior analyst for economic and policy matters at the UC San Diego-based CAIDA initiative, and as Research Program Manager for Packet Clearing House, a nonprofit research institution that pioneered the localization of Internet routing and DNS delivery in developed and developing countries. An ABD Ph.D. candidate at the University of Southern California, Tom has presented excerpts from his forthcoming doctoral dissertation, *Measuring the Wealth of Networks*, at numerous academic, technical, and policy fora around the world.

Contact:
Tom Vest
(tel) +1-703-598-6831
(email) tvest@ripe.net

Eliot Lear is a senior consulting engineer for Cisco Systems. He has authored eight Requests For Comment (RFCs) on various subjects relating to scaling the Internet, is the current chair of the calsify working group as part of the IETF, and is a vice-rapporteur for the ITU-D Study Group 1 Question 22 on Cybersecurity and spam. His current focus includes the economics of IP addressing, as well as the economics of information security. Prior to joining Cisco Mr. Lear was the Internet Architect for a fortune

500 workstation manufacturer. Mr. Lear received his BA in Computer Science from Rutgers University in 1988.

Eliot Lear
Cisco Systems GmbH
Glatt-com 2nd Floor
CH-8301 Glattzentrum
Switzerland
(tel) +41 44 878 9200
(email) lear@cisco.com