HARVARD | BUSINESS | SCHOOL

# Running Out of Numbers: Scarcity of IP Addresses and What To Do About It

**Benjamin Edelman**

## Working Paper

**09-091**

# Running Out of Numbers: Scarcity of IP Addresses and What To Do About It

Benjamin Edelman

Harvard Business School
bedelman@hbs.edu

**Abstract.** The Internet's current numbering system is nearing exhaustion: Existing protocols allow only a finite set of computer numbers ("IP addresses"), and central authorities will soon deplete their supply. I evaluate a series of possible responses to this shortage: Sharing addresses impedes new Internet applications and does not seem to be scalable. A new numbering system ("IPv6") offers greater capacity, but network incentives impede transition. Paid transfers of IP addresses would better allocate resources to those who need them most, but unrestricted transfers might threaten the Internet's routing system. I suggest policies to facilitate an IP address "market" while avoiding major negative externalities – mitigating the worst effects of v4 scarcity, while obtaining price discovery and allocative efficiency benefits of market transactions.

**Keywords.** Market design, IP addresses, network, Internet

**Disclosure.** I advise ARIN's counsel on matters pertaining to v4 exhaustion, v6 transition, and possible revisions to ARIN policy. This paper expresses only my own views – not the views of ARIN or of those who kindly discussed these matters with me.

## 1 Introduction

Hidden from view of typical users, every Internet communication relies on an underlying system of numbers to identify data sources and destinations. Users typically specify online destinations by entering domain names (*e.g.* "congress.gov"). But the Internet's routers forward data according to numeric IP addresses (*e.g.* 140.147.249.9).

To date, the Internet has enjoyed an ample supply of IP addresses. The Internet's standard "IPv4" protocol offers $2^{32}$ addresses (≈4.3 billion). But demand is substantial and growing. Current allocation rates suggest exhaustion by approximately 2011 [1].

Engineers have developed a new numbering system, *IPv6*, which offers $2^{128}$ possible addresses (more than three billion billion billion). But incentives hinder transition, as detailed in Section 5.2. The Internet therefore faces the prospect of continuing to rely on the current IPv4 address system even after v4 addresses "run out." v4 scarcity will limit future expansion, hinder some technologies, and impose new costs on networks and users.

This paper proceeds as follows: In Section 2, I present the technology of IP addressing, and the institutions and policies that allocate addresses. I then turn to specific tactics to manage scarcity. In Section 3, I evaluate central planning, and in Section 4 I examine address sharing. In Section 5, I consider IPv6, including factors impeding transition. In Section 6, I explore a market mechanism to reallocate v4 addresses through transfers; I assess key externalities and policy responses.

## 2 The Technology and Institutions of IP Addressing

IP addresses were first distributed by computer scientists at the Information Sciences Institute (ISI). Initially, scarcity seemed unlikely: Computers were costly, few networks wanted Internet connections, and IPv4 offered billions of addresses. But in the interest of good stewardship, ISI attempted to grant address blocks matching networks' needs. The US military, defense contractors, and large universities received "Class A" blocks ($2^{24}$ addresses, approximately 16.7 million). "Class B" ($2^{16}$) and "C" blocks ($2^{8}$) were provided to smaller networks. Early network architecture permitted only these three sizes.

As demand grew, address assignment developed a geographic hierarchy. The Internet Assigned Numbers Authority (IANA) now grants large "/8" (read: "slash eight") blocks of $2^{24}$ addresses to Regional Internet Registries (RIRs). RIRs in turn assign addresses within their regions. Initial RIRs were RIPE NCC (for Europe, the Middle East, and parts of Africa), APNIC (for the Asia-Pacific region), and ARIN (North America and, at the time, Latin America and parts of Africa). Later, RIRs opened in Africa and Latin America.

RIRs seek to satisfy network operators' demonstrated address needs. An interested network submits a request for addresses to its RIR, along with documents showing its need and its exhaustion of any previously-granted addresses. (Documentation often includes equipment receipts, customer lists, or business plans.) RIR fees follow the principle of cost recovery, rather than maximizing RIR revenue or profit. For example, the largest US networks pay ARIN just $18,000 per year.

IANA continues to assign addresses to RIRs. But IANA's *free pool* reveals an impending shortage: As of January 2009, only 34 /8's remain, and RIRs have recently claimed 6 to 12 /8's per year [1]. Even if demand does not accelerate as exhaustion nears, it seems IANA will soon have no more addresses left to provide to RIRs. Projections for IANA's *v4 free pool exhaustion* range from June 2010 [2] to March 2011 [1].

## 3 Relieving v4 Scarcity through Central Planning

In principle, central authorities could ease IPv4 scarcity by requiring that networks, *e.g.*, migrate to IPv6 as presented in Section 5, on pain of losing ongoing RIR services. Networks want to be listed in RIR *Whois* records so that others can confirm their rights in

the corresponding addresses, and networks want RIR *reverse addressing* services so that automated systems can confirm the host names associated with a network's addresses. These services will be increasingly valuable if IP addresses come to be viewed as scarce resources to be safeguarded and potentially exchanged for value.

But in practice, central authorities have limited power to force migration to v6. Once a network receives addresses from an RIR, it does not directly need substantial ongoing RIR service. Whois primarily benefits the larger community by telling others how to reach the network's technical contacts. Thus, withholding Whois would little threaten an existing network. Even if an RIR declared that a network could no longer use its existing addresses, other networks would continue to know the target network by those addresses, so the network could keep the block with impunity. Whois records are more important when a network seeks to change its connectivity, for an ISP typically checks Whois to confirm a network's rights in the addresses it seeks to use. But if address revocations take effect only upon a connectivity change, most networks could ignore revocations for some time, and networks could retain their existing connectivity to avoid losing addresses. In the future, resource certification might link inter-network communications to RIRs' attestations of address ownership, but such linkages are not yet developed [3].

Institutions and norms also constrain central authorities' ability to force migration. Networks control RIRs through periodic elections of RIR directors, so RIRs cannot act contrary to networks' perceived interests. Furthermore, networks have agreed that RIRs serve principally as custodians to assure that resources are allocated uniquely; networks would oppose RIRs forcing use of particular technologies.

Governments are also badly positioned to accelerate v6 implementation. The Internet's worldwide reach defies control by any single country. Furthermore, even large countries struggle to push transition. For example, the US Office of Management and Budget in 2005 set a June 2008 deadline by which federal agency network backbones must support IPv6 [4] – but compliance devolved into installing equipment that need not actually be used [5]. Japanese tax incentives were similarly ineffective in converting users to v6 [6].

## 4 Sharing IP Addresses to Reduce v4 Demand

As new IPv4 addresses become scarce, some network operators may seek to share addresses among multiple devices. Consider the *home gateway* many users today connect to their cable or DSL modems, letting multiple devices share a single Internet connection and a single public v4 address. Through *Network Address Translation* (NAT), a gateway "rewrites" each outbound IP packet so that, from the perspective of outside networks, that packet comes from the single v4 address assigned to the gateway. When an inbound packet arrives, the gateway attempts to determine which device should receive that packet.

In principle, ISPs can implement similar address translation on a large scale. An interested ISP would assign its users private addresses, using NAT to consolidate onto fewer public addresses at the border between the ISP and the public Internet. Just as

many companies offer "extension 101" on their respective phone exchanges, each private IP address may be used simultaneously by many users around the world.

Despite benefits for address conservation, NAT imposes serious disadvantages. For one, NAT is incompatible with certain communication protocols. In general, it is difficult to send a message to a specific computer when that computer is behind a NAT gateway: The gateway does not know which of its users is the intended recipient of a given inbound message. NAT works well for protocols that begin with a user making a request (*e.g.* requesting a web page): The gateway sees the initial request and can route the response to the appropriate requestor. But consider, *e.g.*, IP-based telephone service. A gateway cannot easily determine which user should receive a given incoming call. Indeed, standard SIP VoIP calls do not work if both caller and callee are behind NAT.

More generally, NAT interferes with the Internet's end-to-end principle [7], limiting future communication designs and impeding development of certain kinds of new applications. Of course existing NAT *already* imposes these impediments, requiring most consumer-focused systems to accommodate NAT in some way. (For example, Skype developed a system of supernode relays to transport data among to and from NAT users.) But increasingly widespread use of NAT would further complicate such designs and further constrain some kinds of innovation. (Indeed, supernode system failure caused Skype's two-day outage in August 2007.) Network architects therefore consider NAT a poor architecture for widespread future use.

# 5 IPv6: The Solution to v4 Scarcity?

As early as 1990, engineers recognized the possible future shortage of IPv4 space [8]. A new version of the IP specification, ultimately named IPv6, dramatically expands the numeric address space – offering $2^{128}$ possible addresses. If many networks moved to v6 and ceased to need or use v4, v4 scarcity would disappear.

## 5.1 Transition to IPv6

Transition to IPv6 is discouraged in part by the limited benefits of v6. v6 was designed to improve authentication, security, and automatic device configuration [9]. However, most enhanced v6 features were "backported" to be available in IPv4 also. For an individual network considering transition, v6 therefore offers little direct benefit.

Transition to IPv6 is further hindered by limited compatibility both forward (existing IPv4 devices seeking to communicate with v6 devices) and backward (v6 devices communicating with v4). Because v4 and v6 use different header formats, direct v4-v6 communications are impossible. For example, a v6-only device cannot directly access the vast majority of the current web because most web servers currently support only v4.

IPv4-v6 translators appear to be practical for specific individual protocols. For example, a dual-stack proxy server could readily accept HTTP requests on an v6 interface, obtain the requested web pages via v4, and forward responses to the requesting users via v6. But seamlessly integrating such a proxy adds considerable complexity: Either v6-only hosts must recognize servers they can only contact via a proxy, or DNS servers must intercept v6-only devices' requests for v4-only servers [10]. Furthermore, some protocols defy translation – for example, by embedding IP addresses within their payloads or by encrypting communications in a way that stymies translation (as in HTTPS). Facing the complexity, unreliability, and unpredictability translation inevitably introduces, the IETF in July 2007 abandoned [12] the official design of a general-purpose translator [11].

For lack of robust translation, some software and protocols may not function on IPv6-only devices. For example, a v6-only PC might use a v6-to-v4 proxy to browse the web – yet be unable to play online games or make voice-over-IP phone calls that work fine for v4 users, because no proxy exists (or is correctly configured) for those protocols. Because a separate proxy must be designed for each application, some applications may never work over v6 – especially old systems and custom software developed for a particular business or industry. Thus, even though Windows Vista and Mac OS X support v6, few users are likely to consider v6-only networking a desirable choice in the short run. In trials at RIR meetings, networking experts found that v6-to-v4 translation worked well for the web, but services as common as HTTPS, Skype and iChat were unavailable [13].

Further constraining IPv6 deployment, few tools are available for administering v6 on large networks. Tools for network management and security are currently largely available only for v4 networks, and some categories of tools lack any effective v6 implementations [14]. In principle, market forces could encourage the provision of v6 tools. But with most networks currently operating only v4, developers see a limited market for v6 versions – providing little immediate incentive for developing v6 tools.

### 5.2 Individual Incentives in IPv6 Transition

Transition to IPv6 is hindered by the incentives of individual participants. Consider a network evaluating v6 to reduce its need for v4 addresses. Little web content is available via v6, nor are other important Internet resources available directly to v6-only devices. The network could use v4-v6 translation, but translation adds complexity – inevitable extra costs when applications do not work as expected. Meanwhile, for lack of v6 administration tools, network administrators find v6 more costly and less flexible than v4. The network's deployment of v6 is further stymied by the lack of v6 *transit*: Most ISPs do not provide v6 connectivity [15]. Furthermore, ISPs that provide v6 tend to offer it less reliably than v4, *i.e.* without service level agreements [16,17], with lower reliability, and with greater latency [15]. In short, under current conditions, v6 is an unpalatable choice.

In principle, IPv4 scarcity might spur transition to v6. But here too, individual incentives oppose transition. In the short run, a network can use NAT to let a single v4 address serve multiple computers, as discussed in Section 4. At some cost for internal

renumbering, the network can reassign and reuse any unused or underused addresses it may have. Finally, the network may be able to transfer addresses from others – either an official transfer as discussed in Section 6, or a "black market" transfer prohibited by formal policy. In the long run, these workarounds carry high costs: As discussed in Section 4, NAT adds complexity, impedes flexibility, and remains untested at the scale some networks might eventually require. Similarly, underused addresses will eventually become hard to find – so reusing addresses cannot continue indefinitely. But in the short run, these v4 challenges are easier than implementing v6. Thus, facing v4 scarcity, it seems networks will naturally choose to use v4 more intensively – not to move to v6.

The core hindrance to v6 seems to be lack of end-user demand for IPv6, for lack of v6-specific features that users value. Suppose users *wanted* v6 – perhaps to obtain higher-quality Skype calls, faster Bittorrrent downloads, or more immersive online video games. Seeking such features, users would pressure their ISPs for v6 connectivity. But at present, no such features exist: v6 offers no clear foundations to support such features, and application developers face an overwhelming incentive to make their best features available to v4-only users. Without user demand, the main proponents of v6 are engineers anticipating future design challenges – a less powerful claim on networks' budgets.

Early experience with IPv6 revealed additional disincentives to its use. For one, even when v6 access works, it is often slower than v4: Fewer networks support v6, so v6 data typically flows less directly, often requiring lengthy "tunnel" detours to bypass v4-only networks [15]. Furthermore, v6 malfunctions can make v6-capable services slower and less reliable than those that support only v4. Even if a web site and user are both v6-capable, their connection will fail if an intervening ISP has not set up v6 or has allowed its equipment to malfunction (a more frequent occurrence with v6 than with v4). Furthermore, consider a user who has accidentally enabled v6, whether by hand or through malfunctioning automatic configurations. (For example, some security software enables v6 in order to secure it – incorrectly telling a user's computer that v6 is ready to use.) Initial measurements indicate that misconfigured-v6 users constitute one third of computers currently attempting to use v6 [18]. When any such user attempts to browse a v6-capable site, the user's computer will chose v6 transit – a request which will fail or endure a lengthy timeout before reverting to v4. Meanwhile, affected users can browse all v4-only sites as usual, without delay. As a result, a site *suffers* from enabling v6 – incurring costs such as lost users, slow load times, and user complaints. These incentives have led some early v6-capable sites to *remove* v6 addresses from their servers [19].

Available data confirms the limited deployment of IPv6 to date. For example, Packet Clearing House reports that 78% of Internet exchange points lack v6 support [20] – preventing participating networks from using those exchanges to transfer v6 traffic. Internet routers hold nearly 200 times as many v4 routes as v6 routes [21]. Technical professionals at the APNIC web site still favor v4 by a ratio of 500 to 1 [21].

In short, v6 deployment remains slow and continues to lack the network effects that accelerated deployment of successful Internet standards. It seems unreasonable to expect v6 to succeed on any particular timetable – particularly because self-interest may lead rational networks to prefer v4 (including NAT) over v6 in the short run.

# 6 A Market Mechanism for Transfer and Reuse of IPv4 Addresses

Even if new IPv4 addresses become unavailable from IANA and RIRs, v4 addresses will continue to be held by existing networks. Some networks may have more than they need due to shrinkage, overoptimistic growth forecasts, or address-saving technologies (*e.g.* v6 or v4 NAT). Other networks may have received abundant "legacy" addresses decades ago. These sources could provide at least temporary relief to v4 scarcity.

## 6.1 The Historic Prohibition on IPv4 Transfers

Historically, IP addresses have not been transferable between networks. If an operator no longer needs some addresses, the operator may only return the addresses to its RIR. When one company acquires another, addresses may move with the acquired company [22]. But RIRs have prohibited transactions solely to transfer IP addresses.

## 6.2 Paid Transfers to Achieve Allocative Efficiency

IPv4 scarcity will create strong incentives for transfers. Some operators will have much less address space than they need. (Consider new operators who receive no addresses prior to exhaustion of available v4 addresses from RIRs.) Conversely, other operators will have more than they need, as discussed above. With transfers, those who most highly value addresses would be likely to obtain addresses from those who can provide addresses at lowest cost – creating surplus from the difference between the parties' valuations.

Consider implications for users who cannot readily switch to v6 – perhaps due to custom software that requires v4, applications incompatible with NAT, unusually costly or busy IT staff, or strong customer or partner preferences. Without transfers, these users would be forced to move to v6 promptly, despite their high costs of transition. In contrast, v4 transfers let these users pay *others* to switch (or otherwise vacate addresses) instead.

Conversely, paid transfers of IPv4 addresses create an incentive for networks to offer addresses for others' use. Under current policies, networks have little incentive to return excess v4 resources: The addresses might be useful or valuable in the future, and a network would forfeit such value if it simply returned addresses to its RIR. In contrast, a paid transfer system pays networks for their unneeded resources – thereby encouraging returns, and rewarding networks which vacate v4 space for use by others.

Experience in other markets indicates that trading resource rights can achieve large efficiency gains. For example, tradable pollution rights reportedly reduced pollution for 55% lower cost than ordering across-the-board cuts by all firms [23] thanks to variation in firms' costs of abatement. Networks feature similar variation in their initial address allocations, in the suitability of transition technologies to serve their requirements, and in their staff and equipment costs of migration. Through transfers, networks with low transition costs can move to v6 first – at lower cost than transitions in random order.

But paid transfers threaten other aspects of addressing policy. Subsequent sections consider possible restrictions on transfers to achieve allocative efficiency while avoiding apparent negative externalities.

### 6.3 Hybrid Markets to Prevent Speculation

Experience in other markets reveals a risk of speculation, manipulation, and other market anomalies. Cornering the market would be costly [24] and probably ill-advised, but even the risk of such disruption worries those whose businesses would be affected [24]. These concerns invite evaluation of market rules to discourage speculation.

One possibility comes in the form of a *hybrid market*, requiring that each recipient satisfy two separate tests to receive addresses. First, the recipient would have to satisfy a substantive examination in which an RIR verifies the recipient's eligibility. Then, the recipient would need to pay to receive addresses from a provider – with a market mechanism serving both to set price and to find and motivate counterparts.

Substantive review by an RIR would prevent speculators from participating in the market, for speculators would fail an RIR's assessment of need. Moreover, consistent with current practice, an RIR could confirm the *amount* of each network's need – preventing networks from partnering with speculators or from seeking excess addresses in an attempt to corner the market, increase price, or disrupt competitors.

In other contexts, detailed verification might be rejected as overly costly. But RIRs have long performed such review efficiently and at low cost [25].

In other contexts, detailed verification might raise significant concerns of regulatory error. But RIRs have decades of experience evaluating requests, including experience embodied in staff, software, and procedures. Moreover, incorrect authorization grants leave the system little worse off than a process that omits need-based review.

Prohibiting speculation rejects the price discovery benefits of speculation and arbitrage. As a result, prices would be slower to adjust to new information. But network operators seem to consider this loss acceptable in light of the associated benefits [26].

### 6.4 Preventing Unreasonable Growth of the Routing Table

Paid transfers are in tension with growth of Internet routing tables, and certain transfers create negative externalities that policy might seek to address. I begin by examining key routing characteristics. I then consider and compare policy responses.

#### The Routing System and Routing Externalities

The Internet's routing system determines how best to transport data between networks. Ordinarily, network addresses are aggregated hierarchically: Data destined for any address in a grouping can be sent to that grouping, without requiring that distant devices know the details of a faraway group. Aggregation offers large efficiency benefits: Although the

Internet connects hundreds of millions of devices, routing decisions are several orders of magnitude smaller. Yet routing remains challenging: The Internet's broad reach and complicated structure yield more than 240,000 entries in a full routing table [27]. Moreover, a typical router must forward hundreds of thousands of packets per second, and routes change frequently due to network disruptions and reconfigurations.

Routing table growth imposes substantial costs. Assessing router cost and capacity, network engineer Bill Herrin estimates a cost of $0.04 per route per router per year [28]. Summing over an estimated 150,000 affected routers, each new route costs the Internet community $6,200 per year. Moreover, if routing tables grow rapidly, ISPs might have to replace routers more often than expected – yielding costs above Herrin's projection. Sharply increased growth could even exceed the capabilities of routers reasonably available in the short run [29].

Few incentives currently constrain growth of the routing table. No central authority has meaningful control over route announcements. Nor can ISPs safely reject unwanted route announcements: Route rejections may prevent an ISP's customers from reaching a remote network – prompting customer complaints and unexpected costs. Instead, routing largely follows from address policy: So long as most networks receive large blocks of contiguous addresses, addresses can be aggregated to avoid unnecessary growth in the routing table.

**The Effect of Transfers on Route Disaggregation**

Paid transfers threaten routing because transfers invite recipients to receive small blocks of addresses from multiple providers – requiring correspondingly many routing table entries. For example, suppose a network needs $2^{16}$ addresses. If the network obtains $2^{16}$ contiguous addresses, others' routers can add just a single routing entry. But if the network instead obtains eight noncontiguous blocks of $2^{13}$, eight routing entries will be required. If a network considers only its self-interest, it will choose the eight $2^{13}$ blocks over the single $2^{16}$ any time the former costs less – imposing a negative externality through extra routing costs.

To address this externality, restrictions could bind either side of the market – regulating address providers, address recipients, or both. In the following sections, I suggest that limiting recipients may be the best choice.

**Prohibiting Disaggregation by Address Providers**

Policy could stop disaggregation at its source by disallowing or severely limiting disaggregation *by address providers*. In a complete prohibition on disaggregation, a provider might have to transfer all its addresses to a single recipient – not to multiple smaller recipients. By keeping blocks intact, this approach would make it unlikely that transfers would require additional routing entries. Indeed, if a network had to give up its entire space (keeping none for its own ongoing use), a full prohibition on disaggregation might allow transfers with no effect on routing at all.

However, restrictions on provider disaggregation blunt the benefits anticipated from transfers. It appears that future networks will seek smaller blocks of v4 space than typical current allocations. (For one, many transfers are expected to come from legacy holders, whose allocations are often very large. Furthermore, recipients are likely to use v4 space in ways that need only small blocks of v4 space, *e.g.* to host servers or to provide interfaces to NAT gateways.) If policy substantially restricts disaggregation by address providers, there would likely be a glut of large blocks, yet an inadequate supply of small blocks – preventing many networks from sharing the large resources embodied in the large blocks.

In principle, policy can allow limited disaggregation to increase supply of small blocks. For example, providers could be permitted to disaggregate by, *e.g.* a factor of four. But setting cutoffs adds significant complexity, and it would be difficult to determine optimal values. Furthermore, policy changes invite gaming and delay as providers anticipate possible changes and try to optimize their decisions accordingly.

**Regulating Recipients through a "Full-Fill Rule"**

Alternatively, policy could seek to prevent unreasonable routing table growth by limiting disaggregation requested *by address recipients*. Suppose an RIR's review qualifies a network to receive some specific quantity of v4 addresses. If the network instead requests multiple smaller blocks that sum to the authorized amount, the RIR would reject the request. After all, the network's need could have been satisfied by a single block, reducing the routing burden imposed on others. Recent ARIN discussions call this approach the "full-fill rule" – requiring that a network satisfy its entire need (for some designated period, *e.g.* six months) with a single transfer [30].

Combining the full-fill rule with permissive disaggregation by address providers creates incentives against unreasonable disaggregation. In particular, these rules guarantee that prices will be convex: large blocks will cost at least as much, pro rata, as small blocks. See proof in the Appendix. With convex prices, address providers prefer to transfer their space in as few, large transactions as possible – yielding convexly greater revenue as well as lower transaction costs. Thus, providers attempt to keep blocks intact – a decentralized approach to the negative externality of unreasonable disaggregation.

Moreover, this combination of rules grants v4 space to those who value it most highly. Suppose multiple small recipients are willing to pay more for their joint use of a given block of addresses – exceeding any single large recipient seeking the same total quantity. Then the large recipient is not the highest and best use of those addresses. By allocating the addresses to the various small recipients, the provider can create more value in the sense of Section 6.2. Arguably, such transfers should be permitted: Disaggregation to connect these new networks is *appropriate* disaggregation which usefully enlarges the Internet; it is not the unreasonable disaggregation policy seeks to prevent. The full-fill rule exactly achieves this result, whereas limiting disaggregation by address providers tends to impede even these desirable instances of disaggregation.

### 6.5 Avoiding Transferring Addresses from Poor Regions to Rich Regions

Paid transfer of IPv4 addresses could include transfers between regions. For example, a network in a low-income country might find it profitable to transfer its addresses to a network in a high-income country. From one perspective, this is allocatively efficient in the sense of Section 6.2: The low-income network prefers money over its v4 addresses, while the high-income network needs the addresses more than the money. Furthermore, the low-income network can use the payment to improve other aspects of its service or, via payments to its owners, to otherwise invest in the local economy. So some may conclude that inter-region transfers are laudable and, in any event, ought not be prevented.

v4 transfers may entail important dynamic consequences. If a network implements NAT address sharing rather than globally unique v4 addresses, it may be hindered in the use of new or innovative applications. Some may be troubled by the prospect of such obstacles disproportionately affecting low-income countries. (Perhaps such countries would later suffer second-rate Internet access, further limiting development.) Those who dislike NAT can move to v6, escaping NAT's restrictions. But if v6 expertise and equipment are particularly scarce or costly in low-income countries, v6 may offer little assistance in the short run.

A natural policy response would allow transfers *within* each RIR, but disallow transfers *between* regions. Because substantial wealth variation occurs between RIR service areas, this restriction would sharply reduce likely address transfer from poor countries to rich. That said, a prohibition on inter-region transfers requires careful evaluation. For one, the restriction would prohibit some exchanges that are allocatively efficient in the sense of Section 6.2 – harming both the would-be recipient and would-be provider. For another, the restriction invites circumvention: Large networks might begin to operate (or claim to operate) within each region so they can receive addresses everywhere. Finally, because the largest share of underutilized v4 resources appear to reside in North America (reflecting early Internet usage and generous address allocations to early users), such a restriction might actually keep prices *lower* in North America than elsewhere. To date, there is no clear consensus on this restriction.

## 7 The Decision at Hand

Once RIRs can no longer grant more IPv4 addresses, networks will face an unavoidable choice: Share addresses through NAT gateways? Deploy v6 immediately? Pay to receive v4 addresses from others? With long-term tradeoffs and significant uncertainty, the decision is challenging.

A market mechanism for v4 addresses appears to offer important benefits. By putting a positive price on existing addresses, paid transfers would show existing networks how much their addresses are worth to others – giving those networks a direct incentive to make the addresses available to others if they can do so cost-effectively, and offering

those networks a financial bonus to spur their migration to v6. Meanwhile, by transferring addresses to networks that cannot easily reduce demand for v4, paid transfers can reduce total system costs – helping the Internet continue to grow. v6 may remain necessary in the long run, but in the short run v4 transfers can help both to mitigate the worst effects of v4 scarcity, and to build the incentives necessary for transition to v6.

# 8 Appendix: Full-Fill plus Permissive Disaggregation Guarantees Convex Prices

**Claim**: Suppose address recipients are bound by the full-fill rule, and suppose providers may disaggregate as they see fit. Then prices are weakly convex. That is, if $P(q)$ is the prevailing market price for a block of size $q$, then for any $Q$ and for any $a>1$, it must be the case that $P(aQ)>aP(Q)$.

**Proof**: Suppose not. Then there exists a provider with some quantity $Q$ that could be divided into $Q_1$ and $Q_2$ where $Q=Q_1+Q_2$ but $P(Q)<P(Q_1)+P(Q_2)$. If so, the provider would never transfer a $Q$ block intact, but rather would subdivide that $Q$ into smaller blocks $Q_1$ and $Q_2$, increasing revenue. So $P(Q)$ cannot be the price of a block of size $Q$.

The following graph shows the impossibility of concave prices. The transferor would increase revenue by subdividing its $Q$-sized block into separate blocks of size $Q_1$ and $Q_2$. In particular, notice that $P(Q_1)+P(Q_2)>P(Q)=P(Q_1+Q_2)$.

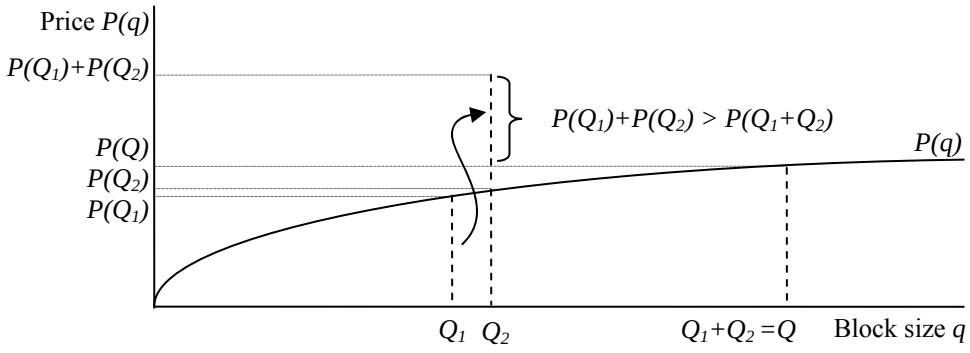

## References


1. Huston, Geoff. "IPv4 Address Report." http://www.potaroo.net/tools/ipv4/, as of May 21, 2008.
2. Hain, Tony. "A Pragmatic Report on IPv4 Address Space Consumption." The Internet Protocol Journal, Volume 8, No. 3, September 2005.
3. Huston, Geoff and Mark Kosters. "Update on Resource Certification." CIDR Report. March 2008.
4. Evans, Karen. OMB Memorandum M-05-22. August 2, 2005.
5. "Federal Government Transition IP Version 4 to IP Version 6 – Frequent Asked Questions." Feb. 15, 2006.
6. Nakamura, Takashi. "IPv6 Deployment Status in Japan." Presentation at APNIC 23.
7. Jerome H. Saltzer, David P. Reed, and David D. Clark. "End-to-End Arguments in System Design." ACM Transactions on Computer Systems 2, 4 (November 1984), pages 277-288.
8. Solensky, Frank. "Continued Internet Growth." Proceedings of the 18th IETF. August 1990.
9. "Internet Protocol, Version 6 (IPv6) Specification." RFC 2460.
10. Durand, Alain. "Issues with NAT-PT DNS ALG in RFC 2766." Internet Draft. Jan. 29, 2003.
11. "Network Address Translation - Protocol Translation." RFC 2766.
12. "Reasons to Move NAT-PT to Historic Status." RFC 4966.
13. "APRICOT 2008 – Lessons Learned." IPv4 / IPv6 Operational Information Collection.
14. Piscitello, Dave. "IPv6 Support Among Commercial Firewalls." Presentation at ARIN XX.
15. Domingues, Mónica et al. "Is Global IPv6 Deployment on Track?" FCCN. October 2007.
16. IIJ Dual-Stack Agreement. http://www.iij.ad.jp/en/development/tech/IPv6/dual/index.html.
17. Bytemark Hosting SLA. http://www.bytemark.co.uk/page/Live/company/terms.
18. Your.org. "Working vs. Broken v6 Clients."
19. Toyama, Katsuyasu, et al. "Clear and Present Danger of IPv6: IPv6/IPv4 Fallback." NANOG 39.
20. "Packet Clearing House Report on Distribution of IPv6-Enabled IXPs." As of January 24, 2009.
21. Huston, Geoff and George Michaelson. "IPv6 Deployment: Just Where Are We?" ISP Column. April 2008.
22. ARIN Number Resource Policy Manual. Section 8.1.
23. Cramton, Peter. "A Review of 'Markets for Clean Air'." Journal of Economic Literature. Vol 38, Sep 2000.
24. Jarrow, Robert. "Market Manipulation, Bubbles, Corners, and Short Squeezes." The Journal of Financial and Quantitative Analysis, Vol. 27, No. 3 (Sep., 1992), pp. 311-336.
25. "IPv4 Transfer Policy Proposal – ARIN XXI Public Policy Meeting Transcript." ARIN XXI. April 7, 2008.
26. "Transfer Policy Proposal Survey – Results." ARIN Advisory Council. August 26, 2008.
27. Huston, Geoff. "Nanogging." ISP Column. November 2007.
28. Herrin, Bill. "BGP Cost." February 2008.
29. "Panel: IP Markets – ARIN XX Public Policy Meeting Draft Transcript." Transcript available at http://www.arin.net/meetings/minutes/ARIN_XX/ppm1_transcript.html#anchor_7.
30. Darte, Bill. "Emergency Transfer Policy for IPv4 Addresses." ARIN Policy Proposal 2008-6.