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

Benjamin Edelman
Harvard Business School
bedelman@hbs.edu

June 10, 2008

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 of available numbers. I begin by presenting the unusual institutions that support the current numbering system, including self-regulatory systems that have managed current resources without detailed government oversight. I then consider the IPv6 standard that would dramatically expand Internet addressing. But I argue that network incentives impede transition to v6 – effectively requiring mechanisms to preserve the current IPv4 numbering system. I consider the possible effects of allowing paid transfers of IP addresses, with special emphasis on rules to ameliorate the worst effects of v4 scarcity, while preserving the core principles of existing regulation and avoiding major negative externalities.

Keywords: Market design, IP addresses, network, Internet

I thank Scott Bradner, Nate Davis, Eric Budish, Peter Coles, John Curran, Lee Howard, Geoff Huston, Mark Kosters, Robin Lee, Scott Leibrand, Bill Manning, Muriel Niederle, Ray Plzak, Al Roth, Steve Ryan, Paul Vixie, Bill Woodcock, and seminar participants at the Market Design Workshop at the Harvard Business School.

Disclosure: I provide advice to ARIN's counsel on matters pertaining to v6 transition, v4 exhaustion, and possible revisions to ARIN's v4 transfer policy.

This paper expresses only my own views – not the views of ARIN or of those who kindly discussed these matters with me.

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

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 when data is sent to its intended destination, the Internet's routers forward the information according to the numeric *IP address* ("Internet Protocol address") of that destination.

To date, the Internet has benefited from an ample supply of IP addresses. The Internet's standard TCP/IP protocol ("IPv4") offers 2^{32} addresses (approximately 4.3 billion). But demand is substantial and growing. At current allocation rates, supply appears sufficient to last only until approximately 2011.[1]

Engineers have developed a new numbering system, *IPv6*, which offers 2^{128} possible addresses. (2^{128} is $3.4x10^{38}$, i.e. more than three billion billion billion.) But incentives hinder transition, as detailed in Section 3.3. The Internet therefore faces the prospect of continuing to rely on the current v4 address system even after IPv4 addresses "run out." The Internet can continue to operate without access to substantial "new" (previously-unused) IPv4 addresses. But IPv4 scarcity will limit future expansion, hinder some kinds of applications, and impose new costs on networks and users.

This paper proceeds in four parts. In Section 2, I present the underlying technology of IP addressing, along with the institutions and policies that currently allocate IP addresses. (Readers already familiar with IP networks may prefer to skip this section.) In Section 3, I consider IPv6, the factors impeding transition to v6, and tactics for encouraging v6 deployment. In Section 4, I examine other likely responses to scarcity of IPv4 addresses, including address translation, paid transfers of v4 addresses, and the challenges that arise in transition to paid transfers of v4 addresses. In Section 5, I reframe the most pressing policy questions and conclude.

My purpose in writing is twofold: First, I seek to present the Internet's unusual addressing institutions to readers who are unfamiliar with this aspect of the Internet's architecture. Second, I hope to use the tools and experience of economics to inform analysis of proposed addressing policy.

2 The Technology and Institutions of IP Addressing

Every computer on the Internet has an IP address. When sending data to a remote computer, the sending computer prepares a *packet* bearing the destination's IP address, as well as the sender's IP (essentially a return address). The computer software that assembles, sends, and receives IP packets is called an *IP stack*.

Standard Internet users run so-called *IPv4* addresses Each IPv4 address is a four byte number, offering a theoretical maximum of 2^{32} different addresses. In practice, at least 10% of addresses have special purposes and are not available for ordinary use.[2]

2.1. The Allocation and Supply of IP Addresses

The allocation of IP addresses tracks the development of the Internet from academic research to mainstream and business use. Originally, IP addresses were distributed by computer scientists at the Information Sciences Institute (ISI), principally Jon Postel. As founder of the Internet Assigned Numbers Authority (IANA), Postel provided TCP/IP addresses to whatever networks requested them.[3]

Initially, scarcity seemed unlikely: Computers were costly, relatively few networks wanted Internet connections, and TCP/IP offered billions of addresses. But in the interest of good stewardship and engineering excellence, if not true preservation or scarcity, Postel attempted to grant each network an address block commensurate with its need.[4]

The First v4 Address Shortage

TCP/IP initially offered only three sizes of address blocks: Large “Class A” networks had 2^{24} addresses (approximately 16.7 million), medium “Class B” networks had 2^{16} (approximately 65 thousand), and small “Class C” networks had 2^{8} (256).[5] Recipients of large blocks included the US military, defense contractors, and research universities that applied early and already operated large computer systems.

By 1992, medium-sized blocks were becoming scarce. It turned out too much space had been reserved for large and small networks: There were many networks with more than 256 devices but less than 65 thousand. Importantly, the initial TCP/IP specification did not allow space to be moved back and forth between network sizes.[6]

With a 1993 change known as *classful interdomain routing* (CIDR), address blocks could be allocated in arbitrary sizes (any power of 2).[7] This flexibility came a cost: Furthermore, routers had to be updated to support the new sizes of address blocks.[8] But the transition came early enough that such upgrades were feasible.

Importantly, computers on existing networks did not need any update to continue to work as usual, even when communicating with devices on networks which used CIDR addresses. Similarly, CIDR-aware devices could communicate as usual with non-CIDR devices. CIDR thus provided both forward and backward compatibility for end-user devices – factors that made the CIDR implementation more straightforward than some changes subsequently under consideration.

Modern Assignment of IP Addresses

As demand grew, address assignment developed into its modern geographic hierarchy. IANA now grants large "/8" (read: "slash eight") blocks of addresses (each comprising 2^{24} addresses) to Regional Internet Registries (RIRs).[9] RIRs in turn assign addresses as needed within their respective continents.[10] Initial RIRs were RIPE NCC (for Europe, the Middle East, and parts of Africa), APNIC (for the Asia-Pacific region), and ARIN (for North America and, at the time, Latin America and parts of Africa). Later, RIRs opened in Africa (AfriNIC) and Latin America (LACNIC).[11]

RIRs assign addresses based on the demonstrated needs of network operators. An interested network operator submits a request for addresses to its RIR, along with documents showing the network's need and its exhaustion of any previously-granted addresses. (Documentation often includes equipment receipts, customer lists, business plans, or other evidence of actual or imminent need.) RIR fees are strikingly low – consistent with the principle of cost recovery, rather than maximizing RIR revenue or profit. For example, the largest US network operators pay ARIN just $18,000 per year.[12]

Some networks received addresses before the RIR system developed. These so-called *legacy* addresses generally lack some aspects of the formality of standard assignments. For example, legacy address-holders may not have a contract with an RIR, need not pay RIR fees, and may have received more addresses than they could justify under the current need-based review.

IANA continues to assign addresses to RIRs. But IANA's *free pool* of available /8's indicates an impending shortage: As of January 2008, there are only 42 /8's left,[13] and RIRs have recently claimed 6 to 12 /8's per year.[14] Even without accelerated demand as exhaustion approaches, it seems IANA will soon have no more addresses left to provide to RIRs. Once the pool of unallocated addresses is depleted, RIRs will be unable to grant further addresses to new or growing networks. Projections for IANA's so-called *IPv4 exhaustion* vary from June 2010[15] to January 2011,[16] but it is clear that unallocated v4 addresses are running low.

2.2. Intelligence and Understanding in IP Network Devices

The intuitive solution to a shortage of IPv4 addresses is simply to "add another digit" to each address, without making more fundamental changes to protocol design. In principle TCP/IP could offer *variable length addresses*, allowing longer addresses without any other changes. But engineers deemed such a design undesirable for its projected increase in complexity and decrease in speed.[17] Thus, adding another digit is infeasible under IPv4 as it stands – far harder than suggested by experience in other networks (e.g. license plates and telephone numbers).

The Internet's structure impedes piecemeal upgrades of the core TCP/IP communication protocols. TCP/IP requires that Internet-connected devices understand the format of IP packets, at least as to source and destination (although not as to the underlying content). That is, an Internet-connected device knows how to send a packet to a desired destination, and how to open a packet it receives.[18] Each Internet-connected device must therefore know the fundamental structure of TCP/IP – including how to produce a packet destined for any other valid address on the network.

That Internet-connected devices understand IP packets is an important design decision: The intelligence of the Internet lies substantially in its endpoints, i.e. users' computers, rather than in the network's core. Compare a computer's understanding of IP with a telephone's lack of understanding of the signaling systems that route phone calls.

Because a computer understands IP, it can transmit any desired data to any desired destination. In contrast, a telephone can only do what its network allows.[19]

The intelligence of Internet-connected devices comes at a cost: Certain fundamental upgrades to the Internet's communication protocols require upgrading end-user devices. Whereas phone system upgrades "only" require upgrading phone company equipment (albeit many thousands of devices around the world), changing certain Internet communication standards requires upgrading every computer. Workarounds can reduce the need for such changes, but the feasibility of a workaround depends on the change.

Fortunately, most Internet upgrades do not require changes to the core communication software on users' computers. The original Gopher system evolved into the World-Wide Web, offering backwards compatibility so that web browsers could contact Gopher hosts. HTTP 1.0 evolved into 1.1, and servers with 1.1-specific features (such as hosting multiple domain names on a single server) could show an appropriate error message to 1.0-only users to encourage them to upgrade. Early 802.11b "Wi-fi" wireless security evolved through several versions, yielding faster and more robust designs that can still connect with older devices when needed. Each of these new versions allows some users to upgrade before others, with reasonable features and benefits resulting from those upgrades even before others joined in. None of these new versions required any changes at the core of the Internet, nor did they require any changes to the underlying TCP/IP protocol. Rather, these upgrades supplemented the stable TCP/IP standard, which serves to reliably transport data from place to place despite developments both "above" TCP/IP (in new and upgraded data formats) and "below" (in new forms of connection-level data transport, such as wireless connections).

2.3. Routing in IP Networks

It would be infeasible to broadcast the precise location of every device on the Internet to every other device all the time. Instead, the Internet's routing system attempts to aggregate nearby addresses hierarchically. Then data destined for any address in a grouping can be sent to that grouping – without requiring that distant devices know the details of individual devices within a group.

Routing system requirements constrain address policy. Suppose local networks often included multiple nonadjacent blocks of IP addresses. Then more routing table entries would be required to explain what data goes where – requiring ISPs to upgrade their routers more often, at considerable expense. When possible, addressing policy therefore seeks to assign large, contiguous blocks rather than multiple smaller blocks.

Despite advances in computing power, routing remains a challenging task. There are currently more than 240,000 entries in a full routing table[20] – reflecting the Internet's broad reach and complicated structure. Moreover, a typical router might need to forward hundreds of thousands of packets per second,[21] and routes change frequently due to network disruptions, reconfigurations, and growth.

Routing presents additional concerns because there seem to be few incentives constraining growth of the routing table. No central authority has meaningful control over route announcements or other aspects of routing policy. At present, some routers refuse route announcements from unknown or distant networks, and router operators can always choose to reject unwanted routing announcements. But if an ISP rejects another company's routes, the ISP's customers may be unable to reach that company's network –

prompting customer complaints and unexpected costs. Route rejection therefore offers a poor check on routing table growth.

2.4. The Regulatory Structure of IP Addressing

The IANA-RIR structure hierarchically assigns unique addresses to interested networks. But neither IANA nor the RIRs can guarantee that the Internet's routing system will correctly transport data to and from such addresses.[22] Routing depends on those who operate the Internet's routers, acting separately and independently from RIRs.

Even within the loosely-coordinated Internet, the IP addressing system is unusual in its lack of centralized regulation. While IANA provides addresses to RIRs, it exerts little policy control over the RIRs. The US government indirectly oversees IANA, in that the US government enters into periodic Memorandum of Understanding agreements with ICANN,[23] IANA's new corporate parent.[24] But neither the US government nor any other has publicly intervened in address policy. Furthermore, when ICANN's structure and function were in dispute, RIRs (through their *Address Supporting Organization*) specifically opposed an enlarged ICANN that might interfere with addressing decisions.[25]

2.5. Sharing IP Addresses to Reduce v4 Demand

As new IPv4 addresses become scarce, some network operators will seek to consume fewer addresses. Available technology offers methods to reduce the number of v4 address a network requires, while preserving compatibility with most applications.

Consider the *home gateway* many users today connect to their cable modems or DSL modems. This device allows multiple devices to share a single Internet connection and a single public IPv4 address. Through *Network Address Translation* (NAT), the gateway modifies ("rewrites") each outbound IP packet so that, from the perspective of the user’s ISP and the Internet at large, all such packets come from a single IPv4 address, namely the address the ISP assigned the gateway. When an inbound packet arrives, the gateway attempts to determine which of the user’s devices should receive that packet.

In principle, ISPs can implement similar address translation on a large scale. An interested ISP can 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 can offer an “extension 101” on their respective phone exchanges, each private address can be used simultaneously by many users around the world.

Unfortunately, NAT imposes some serious disadvantages. For one, NAT is incompatible with certain Internet communications. 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 proper recipient for such an incoming message. For protocols that begin with a user making a request (e.g. requesting a particular web page), this restriction is unimportant: The NAT gateway sees the initial request and can therefore route the response to the appropriate device. But for other protocols, the restriction is more onerous. IP telephony (i.e. voice over IP) is particularly hard-hit: While some proprietary systems, such as Skype, manage to work through NAT, others, including the standard SIP,[26] do not. More generally, NAT interferes with the Internet’s end-to-end principle,[27] limiting future communication designs and probably impeding development of certain kinds of new applications. Of course existing NAT already imposes similar impediments, such that most consumer-focused systems already have to accommodate NAT in some way. (For example, Skype had to develop a system of supernodes and relays to transport data to and from NAT users.[28]) But more widespread use of NAT would further complicate such designs and further constrain some kinds of innovation.

Despite these complications, NAT (or similar address translation) would let ISPs continue operation more or less as they do – providing current service to ordinary consumer users, who need not upgrade their equipment or install new hardware, software, or applications. Each ISP would still need some IPv4 addresses for the public interfaces of its address translation servers. But this public space could be orders of magnitude smaller than existing requirements.

3 IPv6: The Solution to v4 Scarcity?

3.1. IPv6's Features and Design

As early as 1990, engineers recognized the possible shortage of IPv4 space.[29] A new version of the IP specification, ultimately named IPv6, proposed to dramatically expand the numeric address space – while also improving other features including authentication, security, and automatic configuration.[30] Ultimately most of those other features were made available in IPv4 also. At present, the most salient benefit of IPv6 is its expanded address space.[31] But the larger address space, offering 2^{128} possible addresses, would dramatically increase the number of devices that can be connected to the Internet – eliminating the impending scarcity of v4.

Because IPv6 dramatically enlarges the Internet's address space, NAT would no longer be necessary to conserve addresses. To those who emphasize end-to-end connectivity, eliminating NAT is a valuable improvement. But networks deploy NAT not only to conserve addresses: NAT also protects devices from unwanted inbound communications. NAT's inspection of inbound packets effectively serves as a firewall –

only distributing inbound packets that match a prior outbound request, and thereby preventing unexpected incoming messages. In deploying v6, many companies will likely continue similar port-blocking – retaining much of NAT's impediment to end-to-end communications, albeit with greater ability to exempt particular ports when desired. Thus, v6 might change default behavior (allowing inbound traffic, where NAT would not) and might make exceptions easier, yet the practical effect of v6 remains unclear. Moreover, with NAT already widely deployed, it is unrealistic to anticipate rapid transition to a network architecture where unrequested traffic flows unimpeded.

3.2. Transition to IPv6

The transition to IPv6 is hindered by limited compatibility both forwards (existing IPv4 devices seeking to communicate with v6 devices) and backwards (v6 devices seeking to communicate with v4). Because v4 and v6 use different header formats, a v4-only device cannot communicate directly with a v6-only device. This incompatibility has important consequences for deployment of initial v6 systems. For example, a v6-only device cannot directly access the vast majority of the current web because most current web servers only provide data via the IPv4 protocol.

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 IPv6 interface, obtain the requested web pages via IPv4, and forward responses to the requesting users via IPv6. 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.[32] 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). Due to the complexity, unreliability, and unpredictability translation inevitably introduces, the official design for a general-purpose translator[33] was abandoned by the IETF in July 2007.[34] Reviving general-purpose translation creates substantial complexity[35]; so far, such efforts remain in early stages of discussion.

For lack of robust translation, some software and protocols may not work 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 (especially old systems and custom software developed for a particular business or industry) may never work over v6. Thus, even though Windows Vista and Mac OS X support v6 natively, few users are likely to consider v6-only networking a desirable choice in the short run. Indeed, in trials at RIR meetings, networking experts found that v6-to-v4 translation worked well for the web, but services as common as HTTPS (secure web pages), Skype and iChat were unavailable.[36]

Further constraining IPv6 deployment, few tools are available for administering v6 on large networks. Large network use a variety of software and hardware for management, security, and other administrative functions. Yet many of these tools are currently available only for v4 networks, and some categories of tools lack any effective v6 implementations.[37] In principle, market forces could encourage the design of v6 tools. But with most networks currently operating only v4, developers see a limited market for v6 versions – providing little immediate incentive for development of v6 tools.

3.3. Individual Incentives in IPv6 Transition

Transition to IPv6 is hindered by the incentives of individual participants. Consider an individual network considering deploying 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 – meaning inevitable extra costs when applications do not work as expected. Meanwhile, for lack of v6 administration tools, the network administrator would find v6 more costly and less flexible than v4. The network's deployment of v6 is further stymied by the lack of v6 *transit*: Many ISPs do not provide v6 connectivity.[38] Furthermore, ISPs that provide v6 tend to offer it less reliably than their v4 service, i.e. without service level agreements.[39] In short, v6 is not a compelling solution to the network's shortage of v4 space.

In principle, increasing scarcity of IPv4 addresses might spur transition to v6. But here too, individual incentives push in the opposite direction. Consider a network facing a shortage of v4 space. In the short run, the network can use NAT to let a single v4 address serve multiple computers, as discussed in Section 2.5. At some cost for internal renumbering, the network may be able to reassign any unused or underused addresses it may have. Finally, the network may be able to transfer addresses from others – either under a transfer of the sort discussed in Section 4, or in a "black market" transfer prohibited by formal policy. In the long run, these workarounds have major costs: As discussed in Section 2.5, NAT adds complexity and impedes flexibility, and NAT has never been implemented 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 the complicated protocol translation required to implement v6. Thus, the natural individual incentive when facing v4 scarcity is 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 multi-player 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 the v4-only users that constitute substantially the entire Internet. 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. Consider a web site available by both v4 and v6, and a user whose computer prefers v6 access when available. By default the user's computer will attempt to reach the site by v6. If any intervening ISP cannot successfully transport the v6 packets – whether because the ISP has not set up v6, or because the ISP's v6 equipment has malfunctioned – then the user's request will fail. Furthermore, it appears that some users enable v6 by accident, even when their ISPs do not support v6.[40] Such users are able to browse all v4-only sites as usual. But when such users request a v6-capable site, the request fails or suffers a lengthy v6 timeout before reverting to v4. As a result, the seemingly-forward-thinking v6-capable site *suffers* from enabling v6 – incurring costs such as lost users, slow load times, and complaints. These incentives have apparently led some sites to *remove* v6 addresses from their web servers.[41] Moreover, even when v6 access works, it is often slower than v4 – seemingly an artifact of fewer v6 peering relationships, requiring data to take less direct routes from origin to destination.[42] In short, early v6 adopters face penalties without clear countervailing benefits.

Available data confirms the limited deployment of IPv6 to date. For example, Packet Clearing House reports that 79% of Internet exchange points lack v6 support[43] – preventing participating networks from using those exchanges to transfer v6 traffic. The routing table holds nearly 200 times as many v4 routes as v6 routes.[44] The technical professionals who primarily use the APNIC web site still favor v4 by a ratio of 500 to 1.[45]

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

3.4. Interventions to Encourage Deployment of IPv6

In other contexts, a major standards change can be coordinated by a standards body. Consider *Dagen H*, September 3, 1967, the day on which Sweden switched from driving on the left-hand side of the road to the right. But the IPv6 transition is in some ways more challenging than conversion of highway standards: The Swedish government scheduled, coordinated, and funded the major task of modifying that country's highway infrastructure, but there is no clear authority to organize a similar switch in Internet protocols. The Internet's worldwide reach exceeds regulation by any single government. Furthermore, Internet coordinating bodies and standard-setting bodies are relatively weak – lacking mechanisms to require network operators to implement any particular protocol, or to change standards at any particular time.

A few governmental authorities have attempted to encourage or require transition. In 2005, the US Office of Management and Budget (OMB) set a June 2008 deadline by which all federal agencies' network backbones must support IPv6.[46] But as the deadline approaches, progress appears to be slow.[47] Moreover, even if networks' backbones support v6, computers and web server may remain incompatible – making the backbone support useless at least in the short run.

Prior experience in Japan is also cause for caution. In 2002-2003, Japan operated a tax incentive program that exempted v6-capable routers from corporate and property taxes.[48] Japan simultaneously budgeted $7 million for v6 research and development.[49] In response, some Japanese networks (especially university networks) began to use v6. But mainstream Japanese ISPs and users still overwhelmingly run v4.[50]

Seeing the lack of v6-specific benefits to end users, at least one site has proposed a special benefit for v6 visitors: ipv6experiment.com promises (though does not yet deliver) a large quantity of sexually-explicit material, available only to those who access the site by IPv6. The site's operators hope that interested users will ask their ISPs for v6 connectivity – thus building end-user demand for v6. In principle, other sites could copy this strategy – perhaps offering prerelease news stories at Slashdot, reduced advertising at CNET, or free access to sections of the Wall Street Journal Online that otherwise require a paid subscription. Participating sites could justify the associated costs as marketing (reaching new and sophisticated readers) and public relations (building their reputations for innovation). Nonetheless, at present no such offers exist: Even where sites tout their availability by v6 (e.g. Google's ipv6.google.com), the resulting services offer users no apparent advantages over their v4 equivalents.

3.5. Networks where IPv6 More Readily Reduce v4 Address Requirements

Despite the IPv6 deployment impediments set out in the preceding sections, some applications may nonetheless find it workable to use v6 addresses exclusively. Freestanding networks appear to be particularly promising.

Consider the case of mobile phones. Mobile phones generally interact only with each other and with carrier-provided equipment. A carrier can coordinate the upgrade of all of its equipment in order to assure compatibility, without requiring that any other companies make similar upgrades. To the extent that mobile phones need to connect to the existing IPv4 Internet, protocols are limited and well-defined (e.g. email and web). In particular, users generally cannot install arbitrary software on their mobile phones, and users' phones therefore cannot use of unexpected protocols. Protocol translation is thus straightforward. In short, it seems a mobile carrier can probably deploy v6 to its handsets without expecting or requiring any changes by anyone else.

Other freestanding "enclave"-type deployments include internal videoconferencing,[51] building sensor networks,[52] set-top-boxes, and the management interfaces on other network devices (e.g. cablemodems). These devices similarly communicate primarily with internal gateways or other internal devices – and only indirectly, if at all, with the broader Internet. Thus, a single company can readily deploy such devices using v6, without requiring that others coordinate their upgrades.

Yet freestanding networks are unusual on the Internet: Much of the Internet's benefit comes from connections with far-flung and unpredictable counterparts. So while these freestanding networks may serve to demonstrate v6's capabilities and to build v6 experience, they do not serve as a model of v6 deployment more generally.

4 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, overzealous growth forecasts, or early implementation of address-saving technologies (e.g. v6 or v4 NAT). Other networks may have received abundant legacy addresses during a period when networks' needs were not tightly assessed. These sources of addresses could provide at least temporary relief to v4 scarcity – a valuable service if v6 transition is costly, slow, and/or discouraged by unavoidable structural incentives, as suggested by the preceding sections.

4.1. The Historic Prohibition on IPv4 Transfers

Historically, IP addresses have not been transferable between networks. If an operator no longer needs some block of addresses, the operator's only permissible response is to return the addresses to its RIR. If one company acquires another, the acquired company's addresses can flow with the company.[53] But RIRs prohibit sham transactions solely to transfer IP addresses.

RIRs enforce the prohibition on IP address transfers by refusing to update resource allocation databases within their control. After paying for a block of addresses, a network operator would ordinarily want its name listed in the *Whois* database that reports which operators run which networks. By refusing to update Whois, a RIR can prevent many prohibited transfers.

To date, there has been little pressure to allow transfers: RIRs provide addresses to any qualifying registrant, and RIRs' fees pose little impediment to most networks.[54]

RIRs' transfer restrictions reflect the prevailing view that IPv4 addresses are mere numeric identifiers, not intended to be valuable in their own right. IP addresses were never supposed to run out – the Internet was supposed to move to v6 faster.[55] So there was never supposed to be pressure on v4 scarcity, and hence there was never supposed to be substantial reason for network operators to want to "own" or "buy" IPv4 addresses.

4.2. The Prospect of a Paid Transfer System for IPv4 Addresses

Scarcity of IPv4 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 IPv4 addresses from RIRs.) Conversely, other operators will have more than they need, whether because they previously received more than they could justify under current policies (especially operators with legacy addresses) or because their need declined (perhaps from bankruptcy, downsizing, or NAT).

As the prior section explains, current transfer policies prohibit such transfers. But revised policies could allow transfers – unconditionally, or subject to a set of restrictions.

A July 2007 APNIC policy proposal[56] suggested allowing transfers for Asia/Pacific v4 addresses. Restrictions would be minimal: The source and destination of the transfer must be within the APNIC region, and the transferred address range must consist of at least 2^8 addresses. Other terms would be by agreement of transferor and transferee.

Subsequent discussions at ARIN proposed alternative transfer systems – adding restrictions to APNIC's approach in an effort to better serve policy objectives, retain longstanding policy elements, and avoid negative externalities.[57]

Discussions of IPv4 transfers generally feature a single overarching goal: helping Internet users and networks obtain the resources they need. The following sections identify specific objectives, along with policies that might serve these objectives.

4.3. Achieving Allocative Efficiency

The primary benefit of an IPv4 transfer policy is to keep IPv4 addresses available for networks that need them. Some networks can substitute out of v4 with relative ease – whether through NAT (Section 2.5), IPv6 (Section 3), or other redesigns. But other users cannot readily switch – perhaps due to custom software that requires IPv4, applications incompatible with NAT, unusually costly or busy IT staff, or strong customer or partner preferences. For such users, preserving IPv4 is a valuable benefit – a benefit sufficient to justify a monetary payment. Letting the latter group pay the former makes both better off – giving the former money (which they prefer to the v4 space they give up), giving the latter v4 space (which they prefer to the money they pay).

Meanwhile, paid transfers of IPv4 addresses create an incentive for interested networks to offer their IPv4 addresses for others' use. Under current policies, a network with excess IPv4 resources has little incentive to return them: The addresses might be useful or valuable in the future, and the network would forfeit any such value if it simply gave the addresses back to its RIR. A paid transfer system would create a way to pay such users for their unneeded resources – thereby encouraging returns.

Experience in other markets indicates that trading resource rights can produce large increases in efficiency. For example, tradable pollution rights let interested factories cut their pollution sharply (a sensible choice if they can do so at low cost), while others pay to retain the right to pollute (chosen if cutting pollution is costly). By allocating environmental improvements to the firms that can make those improvements most cheaply, tradable pollution rights have reportedly reduced pollution at a cost 55% lower than ordering all firms to cut their pollution by a predetermined amount.[58]

Paid transfer of v4 addresses thus moves towards *allocative efficiency* – transferring resources to those who value them most highly. But addressing policy seeks more than efficiency in making addresses available. Restrictions on v4 transfers might facilitate allocative efficiency while also serving other important policy objectives.

4.4. Preserving the Need-Based Assignment of IP Addresses

To date, IP addresses have been allocated primarily on the basis of need. A network seeking addresses contacts its RIR, demonstrates its need, and receives the addresses it requires. Receiving addresses requires paying a fee, but fees are intended only to cover RIR costs. Indeed, address fees have no particular relationship to the value of the associated addresses. Addresses are allocated based on demonstration of need, not based on who is most able to pay.

Unrestricted paid transfers stand in contrast to the historic approach of low-cost allocations. If paid transfers were allowed without restriction, in principle a transferee could pay to receive a large number of addresses – far more than the transferee could reasonably use. In practice, cornering the market would be costly[59] and probably ill-advised. But even the risk of such an attack is worrisome to those whose businesses depend on v4 addressing. There therefore seems to be significant support within the networking community for continuing to assess each applicant's need for a new or expanded allocation of addresses.[60]

In other contexts, a typical objection to a need-based assessment is that such assessment tends to be costly. However, it seems that RIRs can perform such review efficiently and at relatively low cost.[61]

In other contexts, a typical objection to need assessment is regulatory error – that a central authority might mistakenly grant access to someone who should not have it, or vice versa. But if the authority occasionally wrongfully grants access, perhaps in the face of deceptive requests, the result matches a process that omits a need-based review. Conversely, if an authority wrongfully denies access, the aggrieved party can resubmit its request (albeit with a risk of repeated denial if the authority's errors are systemic). Preserving need-based review therefore seems to present limited risk of regulatory error.

Finally, need-based review of v4 transfer requests offers the additional benefit of preserving consistency with historic RIR practice, as well as with treatment of other number resources (including IPv6 addresses).

4.5. Preventing Unreasonable Growth of the Routing Table

There appears to be a tension between paid transfers and growth of the routing table: Paid transfers might lead transferors to claim small blocks of addresses from many different transferees – requiring multiple routing table entries to provide routing instructions to other networks. For example, suppose a network needs 2^{16} addresses to facilitate its future growth. If the network obtains 2^{16} contiguous addresses, others' routers can address the network using a single routing table entry. But if the network instead buys eight noncontiguous blocks of 2^{13} addresses each, others' routers will have to carry eight entries to reach the network. If the smaller blocks are cheaper, and if a network considers only its narrow self-interest with respect to growth of the routing table, these incentives might produce unnecessary routing table growth.

Routing table growth can impose substantial costs. With data on router cost and capacity, network engineer Bill Herrin estimates the cost of routing at $0.04 per route per router per year.[62] With an estimated 150,000 affected routers, each new route effectively 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 projections. Depending on the rate of growth, some ISPs' needs might exceed the capabilities of routers currently available or reasonably available in the short run.[63]

Prohibiting Disaggregation by Transferors

To prevent unreasonable disaggregation, a natural policy response would disallow disaggregation by transferors. In a complete prohibition on disaggregation, a transferor with (say) a /16 might have to transfer it as such – as a single /16, not as (say) eight /19's.

A complete prohibition on transferor disaggregation would prevent growth of the routing table by keeping existing blocks intact: Each existing block would be transferred *in toto*, and thus would be unlikely to require multiple routing entries.

That said, a complete prohibition on all transferor disaggregation would blunt many of the desired benefits of transfers. It appears that future networks will seek smaller blocks of v4 space than the blocks that have characterized current allocations. (For one, many transfers are expected to come from legacy holders, whose allocations tend to be very large (e.g. /8's). Furthermore, many transferees are likely to use v4 space to offer services, e.g. web servers, or to provide interfaces to NAT gateways or other clusters of end users. These applications call for small blocks of v4 space.) If policy prohibits all transferor disaggregation, then 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.

Rather than completely prohibiting disaggregation, policy could instead allow disaggregation within prescribed limits. For example, a transferor could be allowed to disaggregate by some designated factor – say, by a factor of 16. Alternatively, in anticipation of high demand for small blocks, rules could let large blocks disaggregate more than small blocks. For example, disaggregation down to constituent /16's could be permitted for blocks larger than /16, while smaller blocks disaggregate somewhat less.

The Full-Fill Rule

An alternative policy would prevent unreasonable growth in the routing table by limiting unjustified disaggregation *requested by transferees*. Suppose an RIR's need-based review qualifies a network to receive a /16 block of IPv4 addresses through the transfer policy. If the network instead requests transfer of eight /19 blocks – or, for that matter, two /17 blocks – the RIR could reject the request. After all, the network's need could have been satisfied by a single block, thereby minimizing the routing burden placed on others. So the network arguably ought not use two smaller blocks instead. Recent ARIN discussions call this the "full-fill rule" – requiring that a network satisfy its entire need (for some designated period, e.g. six months) with a single transfer.[64]

A natural objection to the full-fill rule is that it unreasonably increases costs for those who seek a large quantity of v4 space. A transferee who needs a small block can seek offers from any of the many holders of small blocks. But a transferee whose need is larger will have somewhat fewer choices. On one view, higher costs for large blocks are an appropriate hindrance to those whose v4 needs continue to grow quickly. But if requesters of large address blocks are perceived to be unreasonably harmed by the lack of large contiguous blocks, an alternative policy could allow splitting of large requests. For example, a transferee with a need for more than 2^{16} addresses might be permitted to split its request into multiple /16 blocks of 2^{16} or larger. Alternatively, a transferee with a need not met by more than some threshold number of interested transferors might be permitted to split its request into midsized blocks that are more widely available. In short, the special challenges of large requests need not impede policy that prevents excessive disaggregation of smaller network blocks.

Combining the full-fill rule with permissive disaggregation by transferors offers an economic incentive that further impedes unreasonable disaggregation. If transferees are bound by the full-fill rule while transferors may disaggregate as they see fit, then prices will be convex. That is, a large block will cost at least as much, on a pro rata basis, as a small block. See proof in Appendix 6.1, and diagram below.

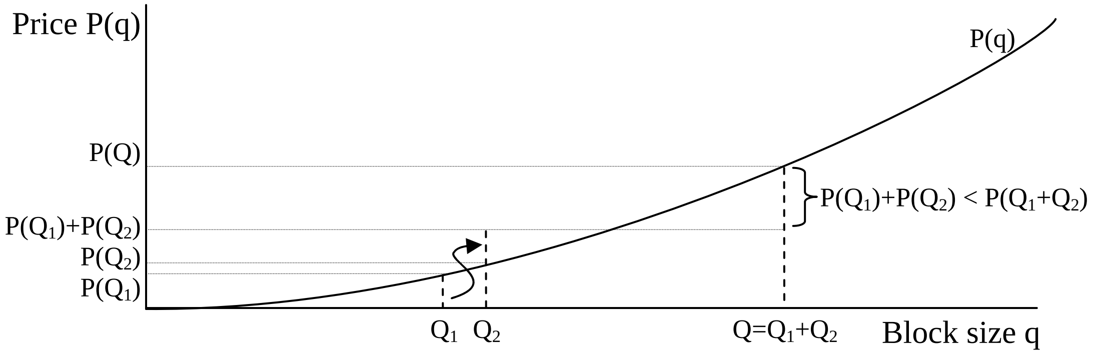

If prices are convex, then transferors prefer to transfer their v4 blocks in as few, large transactions as possible. Larger blocks would yield convexly greater prices, not to mention reduced transaction costs. Thus, combining the full-fill rule with permissive disaggregation by transferors leads transferors to prefer to transfer their v4 space in large blocks – keeping blocks intact, and avoiding unreasonable disaggregation.

Combining the full-fill rule with permissive disaggregation grants v4 space to those who value it most highly. Suppose multiple small transferees are willing to pay more for their joint use of a given quantity of space – beating any single large transferee seeking that same total quantity of space. Then the large transferee is not the highest and best use of that space. By allocating that space to the group of small transferees, the transferor can create more value, in the allocative efficiency sense of Section 4.3. The analysis in 4.3 thus indicates that such transfers should be permitted – for while they may create routing disaggregation, the disaggregation is *useful* (in allowing more bona fide networks to connect to the Internet). Disaggregation to connect these networks is *appropriate* disaggregation; it is not the unreasonable disaggregation policy seeks to prevent.

4.6. Preventing Speculation

In the eyes of some network operators, speculation in IPv4 addresses appears unhelpful and potentially disruptive.[65] Speculators challenge the community view of IP addresses as numeric identifiers, not commodities or stores of value. To the extent that speculators buy and hold v4 addresses in anticipation of future appreciation, speculators drive up prices, at least temporarily. Whether or not speculators succeed in reaping large profits from their efforts, speculators can make prices unpredictable – interfering with businesses that require a reliable source of v4 addresses.[66]

Need-based assessment of transferees tends to hinder speculation. Under a regime of need assessment, a would-be speculator must show an RIR how it would use an address block. A speculator who in fact begins to operate a network, ISP, or other

Internet service is arguably no longer a speculator. But without such a business, the speculator should fail its need-based review.

A minimum holding period could further dampen speculation. Suppose a block of addresses may only be transferred if the would-be transferor has held the addresses for some designated period of time, say two years. Then any speculator would be forced to hold addresses at length – increasing the capital-intensity of speculation, preventing quick profits, and adding considerable risk given the uncertainty of future prices.

4.7. Avoiding a Transfer of 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 high-income network in another country. From one perspective, this is an allocatively efficient exchange in the sense of Section 4.3: The low-income network prefers money over its v4 addresses, while the high-income network needs the addresses more than it needs 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.

But v4 transfers may yield important dynamic consequences. If a network comes to rely on NAT rather than ordinary v4 addresses, it may be hindered in the use of new or innovative applications. Some may be troubled by the prospect of such hindrance disproportionately affecting low-income countries. Perhaps low-income countries would later suffer second-rate Internet access, limiting their ability to develop in other respects. Importantly, the availability of IPv6 serves to limit this concern: Those who dislike v4

NAT can move to v6, escaping NAT's limitations. But if v6 expertise and equipment are costly or scarce in low-income countries, v6 may offer little assistance in the short run.

A natural policy response is to limit transfers to occur *within* RIRs' service regions but not *between* the regions. That is, a transferor and transferee would need to be located within the same RIR service region in order to transfer a block of IP addresses. Some large networks might be able to circumvent such a restriction by claiming to operate everywhere – hence claiming eligibility to transfer addresses within each RIR. But RIRs in low-income regions could implement policies to require that multi-region networks obtain v4 addresses in the same region where those addresses are used – preventing such networks from filling their worldwide address needs through low-income RIRs. RIRs could also publish the listings of consummated transfers – inviting public scrutiny to expose any attempts at inter-region address transfer.

4.8. Avoiding Windfalls

Some networks currently have more v4 addresses than they need. For example, as discussed in Section 2.1, some legacy networks received large blocks without showing the need required under current policies. Other networks justified large requests, but ultimately did not need the space – perhaps due to changed network architecture or business failure. Such networks can offer v4 space if transfers are permitted.

On one view, the availability of underutilized v4 addresses is an opportunity to be celebrated. The more underutilized space available, the greater the assistance available to those who need v4 addresses after IANA and RIRs have no more space to give out. Furthermore, the more networks hold underutilized space, the more intense the likely competition to offer that space – yielding lower prices that better resemble the status quo.

If current network operators receive payments for their v4 space, some may view that as unseemly or improper – an unearned "windfall." Indeed, these networks probably did not expect such payment when they requested the space, received it, or held it. It is unlikely that their past behaviors changed in anticipation of receiving these payments.

To those concerned about possible windfall to legacy holders, one mitigating factor comes in their contribution to the early growth of the Internet. In many instances, legacy networks offered technical assistance, expertise, or equipment useful to the Internet's early development. Even if no monetary payment was ever promised in exchange for such assistance, a payment may not be unwarranted to certain recipients – particularly when the payment also yields benefits to the networks that receive the vacated addresses.

Moreover, discussions of transfer policies offer no realistic alternative to paying networks that return addresses for use by others. Experience indicates that networks are unlikely to freely return unneeded addresses: RIRs have long requested returns of unneeded addresses, but few legacy holders have returned addresses for free. (Notable and laudable exceptions are Stanford University and the US Department of Defense.[67]) As v4 scarcity becomes increasingly widely known, free returns are even less likely.

Finally, it remains unclear how to prevent those with extra v4 resources from receiving windfalls as v4 becomes scarce. With or without an official transfer policy, these networks hold a valuable resource that others seek. Given RIRs' powers and capabilities, it may be unrealistic to expect RIRs to be able to prevent transfer of such resources, even if RIRs wanted to do so.

4.9. A Listing Service to Facilitate IPv4 Transfers

The preceding sections anticipate *allowing* transfers between interested parties. But even if transfers were allowed, networks might face difficulty finding appropriate trading partners. A listing service could assist in making such matches while offering ancillary benefits for safety, security, and convenience.

Market Thickness and a Listing Service

Experience in other markets reveals the importance of market thickness – gathering enough parties seeking to trade.[68] Appropriate institutions can help achieve thickness. In particular, a listing service can help transferors find suitable transferees, and vice versa.

A listing service could show only transfer offers, only transfer requests, or both offers and requests. Showing both offers and requests facilitates market thickness by providing more information to market participants: Each offer would include an asking price, and each request would include an offer price. Seeing others' prices, market participants can determine how far they stand from successful transactions.

In other contexts, auctions often feature a fixed ending time. (Consider auctions at eBay.) But an online listing service need not impose a fixed ending time. Instead, each listing can remain posted until it has been satisfied.

To prevent market congestion, a listing service could (by default) show each participant only the trading partners appropriate for that participant. For example, if a participant has been *prequalified* (based on its demonstration of need) to receive a given size IPv4 block, the listing service could show that participant only those transferors offering blocks of the specified size. Furthermore, the listing service could sort suitable listings based on asking price, giving preferred placement to listings with low asking prices. Pending community interest, the listing service could also show, sort by, and/or filter by additional fields (e.g. availability, accepted payment methods).

Assuring Safety of Market Participants

A paid transfer system creates a risk that some market participants may not do what they promise. A transferee might reasonably worry about a nonperforming transferor. Perhaps the transferor never had the addresses it promised to transfer. Or perhaps the transferor refused to complete the transfer even though the transferee paid in full.

Conversely, a transferor might worry about a nonperforming transferee. A transferee might refuse to pay for the addresses it received (or promised to receive). Or a transferee might tender invalid payment, i.e. a bounced check.

The intangible status of IP addresses lessens some concerns about nonperforming market participants. In an important sense, an IP address transfer occurs only if the corresponding RIR records such transfers in Whois and other records. So an RIR can reverse an invalid transaction – for example, if a transferor satisfactorily demonstrates that a transferee failed to pay.

A listing service web site could further reduce transfer risks. RIR data could assist market participants in authenticating trading partners. RIRs already authenticate address holders to confirm the authenticity of change requests and other communications. Using similar authentication systems, RIRs could verify known parties as they enter a listing service. In particular, an RIR could cross-check each would-be transferror's offerings with the participant's prior address allocations. Similarly, an RIR could compare each would-be transferee's request with its prequalified need. Listing service authentication therefore increases confidence in the identities and qualifications of trading parties.

Some unsuccessful transfers may produce disputes. But these would be standard commercial questions of payment and delivery – matters appropriately resolved through the legal system. Interested listing services could impose standard terms that require parties to submit to arbitration or other alternative dispute resolution. Listing service terms could also attempt to limit the listing service's liability.

Other Benefits of a Listing Service

The use of a listing service offers an opportunity to lower transaction costs for market participants. A single central listing system could offer a simple search for IP address transfers – a useful aid for those who need addresses infrequently.

A listing service could also suggest default contracts to consummate a transaction – making it unnecessary for participants to draft agreements from scratch. Defaults benefit infrequent participants, who would otherwise suffer disproportionate drafting costs.

If transfer policies restrict the size of transfers, a listing service could alert participants to relevant rules and apparent violations. For example, if policies prohibit a transferee from satisfying its need through multiple small blocks, the listing service could warn a transferee that is examining a block smaller than the its prequalified need.

Downsides of a Listing Service

Despite these benefits, a listing service also presents risks deserving scrutiny.

Suppose a listing service suggests a match between some transferor and transferee, but the transaction ultimately turns out badly – leaving one or both parties unsatisfied. The parties might blame the listing service, and the matter might be costly or time-consuming to resolve. A listing service might attempt to avoid such exposure through appropriate language in listing service terms and conditions and through other aspects of listing service design. Listing services would probably seek to emphasize that decisions to transfer, pay, or accept payment are the responsibility of market participants, not the listing service. But market participants may have a different view of the situation, particularly if their transaction goes awry.

Some may oppose an RIR-provided listing service on the grounds that it is an inappropriate or unnecessary expansion of RIR responsibility. In principle, listing service functions could be provided by one or more entities independent of RIRs, rather than by RIRs themselves. Perhaps such separation would better clarify the various parties' responsibilities. Perhaps competition among listing services would offer innovation in structure or design. But these possible benefits require a significant increase in complexity, particularly as to authentication and security. At a listing service not operated by an RIR, it would probably be substantially harder to confirm a trading partner's identity. A natural response to these concerns is to make RIR-provided listing services optional: Transferring parties would be free to find each other however they like – whether through an RIR-provided listing service, another listing service, or ordinary word of mouth. Alternatively, RIRs could provide an API or other automated system to verify a participant's identity, existing allocations, and/or prequalification – granting third-party listing services the benefit of RIRs' data and authentication.

4.10. Challenges of Transition to Allowing v4 Transfers

Allowing paid transfer of IPv4 addresses raises some complications in transition.

RIRs' current contracts with address holders exactly prohibit transfers of v4 addresses. For example, the ARIN RSA specifically provides that an address holder "is not permitted to assign this Agreement or any of its rights or obligations under it."[69]

That said, RIRs retain the right to modify their agreements with users: The ARIN RSA provides that "because of the necessary role that ARIN performs for the Internet community, ARIN reserves the right to modify this agreement at any time."[70] Allowing transfers would grant new rights to each user – namely the right to transfer addresses, consistent with applicable restrictions. RIRs are unlikely to face complaints from address holders when RIRs unilaterally grant new rights without seeking consideration in return.

Allowing v4 transfers also creates the potential risk of accelerated exhaustion of v4 resources. Seeing that v4 addresses will soon be transferable for a fee, users might make a "run on the bank" – requesting more addresses than they need, with an eye to transferring the addresses for a profit later. That said, v4 exhaustion already creates similar risks – users wanting to get addresses while they are still available. Furthermore, existing RIR policies examine the merits of each request – hindering simple mass requests. On one view, paid transfers accelerate exhaustion – providing a new motive (possible future profits) for obtaining scarce addresses. But a paid transfer system would also impose rules to reduce the short-term value of last-minute requests, including a minimum holding period (Section 4.6) and the likelihood of competitive supply of addresses (hence possible lower prices in the future). It is therefore unclear whether paid transfer would worsen the problem of speculative last-minute requests for v4 addresses.

4.11. Could a Thriving v4 Transfer System Impede v6 Transition?

The prevailing view among IP engineers seems to be that v6 deployment will accelerate when v4 addresses will become too expensive, too scarce, or too uncertain in future price or availability. A successful v4 transfer system might prevent those conditions from arising – presenting prices that are low (at least relative to v6 transition costs), confirming the availability of the addresses operators need, and providing a reliable procedure to satisfy future requirements. Indeed, if v4 transfers successfully provide these benefits, there would be reduced impetus for transition to v6.

On one view, delaying v6 transition offers a valuable benefit. The Department of Commerce estimates that IPv6 transition will cost some $25 billion over 25 years, largely in labor costs to implement the switch.[71] Preserving v4 via an effective transfer system would delay these costs – generating billions of dollars of savings. For example, some network infrastructure will have to be replaced – not just upgraded with new software, but entirely replaced – in order to support v6. Delaying v6 transition allows such replacements to occur on their ordinary schedule, e.g. in response to failure or true obsolescence, rather than requiring that such replacements occur earlier than planned.

But suppose a move to v6 is unavoidable – because demand will ultimately exceed even the increasingly efficient allocations produced by v4 transfers. On this premise, the core challenge is *coordinating* the move to v6 – reducing the period during which v4-v6 translation is required. Exhaustion of v4 addresses could provide a coordination benefit: When network operators notice that v4 addresses have become unavailable, they will face a special impetus to move to v6. On this view, a successful v4 transfer system interferes with v6 transition by blunting the news and the consequences of v4 exhaustion.

Even if v6 transition is ultimately unavoidable, preserving v4 is valuable in its own right. The costs of v6 transition are large and would be incurred immediately. The benefits of v6 are further afield and significantly less clear. A well-designed v4 transfer procedure would do much to help those whose networks need more v4 addresses (whether due to legacy software, limited IT staff, the limits of v4-v6 translation, or other good reasons we cannot readily anticipate). An RIR would struggle to force its users to incur the short-run costs of v6 transition when, in the short run, v4 transfers provide similar benefits (i.e. business continuity) at lower short-run cost.

Due to low utilization of many legacy address blocks,[72] a v4 transfer policy could plausibly yield many addresses that have been little utilized under existing policies. Prices might be surprisingly low, and availability surprisingly large. (Consider experience with tradable pollution rights.[73]) If v4 prices are low, it may be possible to delay v6 transition substantially. Conversely, if v4 space carries a high price, then networks will automatically see an incentive to move to v6 to escape the costs of v4.

If v6 later appears to be necessary despite permissive v4 transfers, policy could appropriately limit v4 transfers. An initial step – perhaps sufficient to inspire v6 experimentation if not full deployment – would make the v4 transfer listing service available to v6 users only (i.e. on a web server only reachable by IPv6). Then any interested v4 transferee would have to learn at least enough v6 to connect to a v6-only web site. For a further push towards v6, v4 address transfers could be limited to transferees who have appropriate v6 transition plans, formally authorized by corporate officers. Ultimately, v4 transfers could be allowed through some date certain, but disallowed thereafter. Or v4 transfer fees could increase over time – albeit presenting major questions as to the propriety of such charges, and as to allocation of the proceeds given RIRs' longstanding commitment to cost recovery.

4.12. Alternatives to a v4 Transfer Policy

Instead of allowing v4 transfers, an RIR could choose to retain its current prohibition on IPv4 transfers. The instinct to continue to refuse transfers is clearly present in some RIR discussions.[74] But such a policy risks creating a black market in addresses transferred despite RIR rules.

By refusing to update Whois, an RIR would deter some prohibited transfers. In particular, a transferee contemplating paying for v4 space would ordinarily want that transfer properly recorded in Whois, lest the transferor later renege. But some prohibited transfers would occur nonetheless. The RIR would then have to choose whether to update Whois to reflect the prohibited transfer. The RIR might deny Whois updates to punish the parties that performed the prohibited transfer and to deter further similar transfers. But then the RIR would fail to provide the best possible data to those who depend on Whois – hindering investigations and analyses by network operators, law enforcement, and ordinary users. This unpalatable choice suggests RIRs may be unable to follow through on a strict prohibition on transfers.

Furthermore, an RIR's refusal to update Whois is only powerful if Whois is substantially accurate. If Whois is widely viewed to be unreliable, networks are unlikely to care whether Whois correctly captures the current use of an address block. Refusing to update Whois thus risks an unraveling in which Whois becomes far less accurate and in which RIRs simultaneously lose substantial power to deter prohibited transfers.

Separate from Whois, RIRs might be able to provide useful information to the operators who set routing policy for their respective networks. For example, RIRs could report blocks transferred contrary to applicable rules. Some ISPs might choose to *null-route* such blocks – preventing those ISPs' customers from reaching services within the corresponding networks.[75] Such an action would harm the transferee, perhaps deterring future illicit transfers. Yet this action would also harm the ISP's paying customers, and it might increase the ISP's support costs (e.g. if customers call to complain). So ISPs might find themselves badly positioned to take action even against known violators.

Through *resource certification*, routers could automatically confirm whether an RIR has assigned a given block of addresses to the network attempting to use those addresses. Resource certification would thus discourage prohibited transfers by reducing the likelihood that the resulting addresses would function as expected. But resource certification is not yet operational, and its availability remains uncertain.[76] Furthermore, resource certification adds complexity to routers that are already heavily burdened.

5 Next Steps

Once RIRs can no longer grant more IPv4 addresses to the networks that seek them, networks will face an unavoidable choice: Deploy v6 immediately, to reduce v4 requirements? Pay to receive v4 addresses that others don't need? Share addresses through NAT gateways? Each approach has its benefits. With an intertemporal tradeoff plus substantial uncertainty, the decision is necessarily complex.

Meanwhile, RIRs must decide whether to continue to enforce the existing prohibition on transfers, or instead to implement new policies that allow transfers between interested parties. Allowing v4 transfers would offer immediate operational benefits, particularly given the impediments and disincentives to v6 deployment. Yet preserving v4 would surely dull an important inspiration for v6 transition.

IPv4 exhaustion presents a challenging inflection point in the Internet's growth – a level of usage never expected when the Internet's protocols were designed. A jump to v6 would accommodate orders of magnitude more users – yet with costs and challenges so fundamental that rehabilitating v4 nonetheless demands serious consideration.

6 Appendices

6.1. Full-Fill Rule plus Permissive Disaggregation Guarantees Convex Prices

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

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

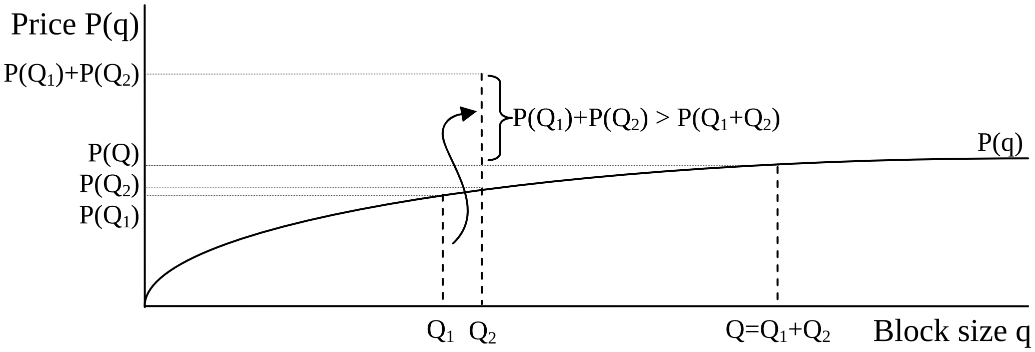

6.2. How Market Rules Can Address Concerns about v4 Transfers

Concern	*Response*
IP addresses should be numeric identifiers, not valuable commodities.	Retain need-based assessment of each transferee.
Prevent unreasonable growth in the routing table.	Require most transferees to fulfill their entire need in a single transfer. Limit the ability of large transferees to resort to a series of small transfers.
Prevent speculation in IP addresses.	Retain need-based assessment of each transferee. Impose a minimum holding period.
Avoid a transfer of addresses from poor regions to rich regions.	Require that transferor and transferee reside within the same RIR service area. Require that addresses be used within the same RIR where they are obtained. Publish listings of consummated transfers.
Achieving market thickness and avoiding congestion.	Provide a suitable listing service to help pair appropriate transferors and transferees.
Assure the safety of market participants.	Limit listing service access to known entities. Cross-check listings with RIR data.

7 Notes and References

[1] Huston, Geoff. "IPv4 Address Report." http://www.potaroo.net/tools/ipv4/, as of May 21, 2008.

[2] "Special-Use IPv4 Addresses." RFC 3330.

[3] "Assigned Numbers." RFC 750.

[4] "Assigned Numbers." RFC 750. However, as discussed in "The First v4 Address Shortage," below, attempts to align allocations with needs were hindered by the few sizes of address blocks then available.

[5] "Internet Protocol Specification." RFC 791.

[6] "Supernetting: An Address Assignment and Aggregation Strategy." RFC 1338. Under the initial TCP/IP specification, an IP-connected device examined its own IP address to deduce the size of its local network (i.e. the range of IP addresses that could be reached without assistance from a router). For example, if the device observed that it was within a large network, it would know that it could directly reach any other device on that same network. Because each device inferred its network size from the numeric value of its IP address, addresses could not be moved between network sizes without disrupting devices' abilities to determine the size of their own networks and hence how and when to reach remote networks.

[7] "Classless Inter-Domain Routing." RFC 1519. CIDR allows each IP-connected device to receive an arbitrary *subnet mask*. This subnet mask tells a device which IP addresses are on the local network and can be reached without passing through a router.

[8] "Class A Subnet Experiment." RFC 1879.

[9] Recall that an IPv4 address is always 32 bits long. The designation "*/n*" (e.g. "*/8*") refers to a network that uses n bits (e.g. eight bits) to designate its network prefix, i.e. the portion of the address common to all addresses within the address block. The remaining 32-n bits (e.g. 24 bits) number the individual addresses within the block. A smaller address block would have a longer prefix, i.e. */16* or */20*, leaving correspondingly fewer bits (i.e. sixteen or twelve) available for designating the 2^{16} or 2^{12} individual addresses within such a block.

[10] IANA IP Address Services. http://www.iana.org/ipaddress/ip-addresses.htm.

[11] "IANA Report on Recognition of LACNIC as a Regional Internet Registry." http://www.iana.org/reports/lacnic-report-07nov02.htm. "IANA Report on Recognition of AfriNIC as a Regional Internet Registry." http://www.iana.org/reports/afrinic-report-08apr05.htm.

[12] See e.g. "ARIN Fee Schedule." http://www.arin.net/billing/fee_schedule.html. See also "RIPE NCC Charging Scheme." http://www.ripe.net/ripe/docs/charging.html.

[13] "Internet Protocol v4 Address Space." http://www.iana.org/assignments/ipv4-address-space.

[14] Huston, Geoff. "IPv4 Address Report." http://www.potaroo.net/tools/ipv4/index.html.

[15] Hain, Tony. "A Pragmatic Report on IPv4 Address Space Consumption." http://www.tndh.net/~tony/ietf/ipv4-pool-combined-view.pdf.

[16] Huston, Geoff. "IPv4 Address Report." http://www.potaroo.net/tools/ipv4/, as of May 21, 2008.

[17] Huston, Geoff. Email communication on file. June 2008. See also Nakajima, Yoshihiro. "Special Interview IPv6 Protocol Stack for BSD Development Project" (interviewing Jun Murai). December 25, 2005. (Detailing the reasons why variable-length addresses were rejected in IPv6.)

[18] A given device need not understand how to interpret the *contents* of an incoming packet. But the device can at least receive the packet, determine whether the packet has reached its intended destination, and pass the packet on to whatever software program is designated to process such packets.

[19] Isenberg, David. "The Rise of the Stupid Network." Computer Telephony, August 1997, pg 16-26.

[20] Huston, Geoff. "Nanogging." November 2007. http://www.potaroo.net/ispcol/2007-11/nanog.html.

[21] See e.g. "Choose a Cisco Router that Best Fits Your Organization's Needs." TechRepublic. June 9, 2005. http://articles.techrepublic.com.com/5100-10878_11-5728608.html.

[22] See e.g. ARIN Number Resource Policy Manual, section 4.1.1. http://www.arin.net/policy/nrpm.html.

[23] "ICANN's Major Agreements and Related Reports." http://www.icann.org/general/agreements.htm.

[24] "Contract Between ICANN and the United States Government for Performance of the IANA Function." http://www.icann.org/general/iana-contract-09feb00.htm.

[25] "RIR Statement on Evolution and Reform." http://aso.icann.org/news/news-docs/rirstatement.html. (concluding that "the RIRs see no value in an ICANN structure that admits the possibility of imposition of arbitrary and potentially capricious policies onto the management of Internet resources")

[26] Cionoiu, Diana. "NAT Traversal for the SIP Protocol." http://freshmeat.net/articles/view/2079/.

[27] 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.

[28] "How Does Skype Get through Firewalls and NAT Routers?" The Register. October 8, 2003.

[29] Solensky, Frank. "Continued Internet Growth." Proceedings of the 18th Internet Engineering Task Force. August 1990. Available at http://www.ietf.org/proceedings/prior29/IETF18.pdf.

[30] "Internet Protocol, Version 6 (IPv6) Specification." RFC 2460.

[31] *See* Durand, Alan. "When Net 10 Is Too Small – Ipv6 @ Comcast." http://www.ripe.net/ripe/meetings/ripe-54/presentations/IPv6_management.pdf, concluding "96 more bits, no magic."

[32] Durand, Alain. "Issues with NAT-PT DNS ALG in RFC 2766." Internet Draft. January 29, 2003.

[33] "Network Address Translation - Protocol Translation." RFC 2766"

[34] "Reasons to Move NAT-PT to Historic Status." RFC 4966."

[35] van Beijnum, I. "Modified Network Address Translation – Protocol Translation: Internet Draft."

[36] "APRICOT 2008 – Lessons Learned." Available at http://www.civil-tongue.net/6and4/wiki/APRICOT2008-Lessons.

[37] *See e.g.* Piscitello, Dave. "IPv6 Support Among Commercial Firewalls," presentation at ARIN XX, http://www.arin.net/meetings/minutes/ARIN_XX/ppm2_transcript.html#anchor_19.

[38] Domingues, Mónica et al. "Is Global IPv6 Deployment on Track?" FCCN. October 2007.

[39] See e.g. IIJ Dual-Stack Agreement ("The service-level agreement policy does not apply to IPv6/IPv4 dual stack Access"). http://www.iij.ad.jp/en/development/tech/IPv6/dual/index.html. See also Bytemark Hosting SLA ("No guaranteed service level is offered for connectivity to IPv6 addresses"). http://www.bytemark.co.uk/page/Live/company/terms.

[40] Your.org. "Working vs. Broken v6 clients." Available at http://www.your.org/v6clients.png . (Reporting that misconfigured v6 users – who attempt to use v6 but cannot – constitute approximately one third of all supposedly-v6-capable users at the Your.org site.)

[41] Toyama, Katsuyasu, et al. "Clear and Present Danger of IPv6: IPv6/IPv4 Fallback." Presentation to NANOG 39. http://www.nanog.org/mtg-0702/presentations/ipv6_katsuyasu.pdf.

[42] Domingues, Mónica et al. "Is Global IPv6 Deployment on Track?" FCCN. October 2007.

[43] "Internet Exchange Point Summary (IPv6)." https://prefix.pch.net/applications/ixpdir/summary/ipv6/, as of May 21, 2008.

[44] Huston, Geoff and George Michaelson. "IPv6 Deployment: Just Where Are We?" April 2008. http://www.potaroo.net/ispcol/2008-04/ipv6.html.

[45] Huston, Geoff and George Michaelson. "IPv6 Deployment: Just Where Are We?" April 2008. http://www.potaroo.net/ispcol/2008-04/ipv6.html.

[46] Evans, Karen. "Memorandum for the Chief Information Officers." August 2, 2005. http://www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf.

[47] Marsan, Carolyn. "Feds Look to Fudge IPv6 Mandates." CIO Today. January 8, 2008.

[48] "Technical and Economic Assessment of Internet Protocol Version 6 (IPv6)." National Telecommunications and Information Administration, IPv6 Task Force, Draft Report. July 16, 2004. http://www.ntia.doc.gov/ntiahome/ntiageneral/ipv6/draft/discussiondraftv13_07162004.htm.

[49] Arano, Takashi. "IPv6 Deployment in Japan." October 28, 2002. http://www.icann.org/presentations/ipv6-workshop-arano-28oct02.ppt.

[50] Nakamura, Takashi. "IPv6 Deployment Status in Japan." http://www.apnic.net/meetings/23/archive/presentations/ipv6-pres-nakamura-jp-deployment.pdf.

[51] Ishiyama, Satoshi. "IPv6." Presentation to North American IPv6 Summit. May 24, 2005.

[52] Arano, Takashi. "IPv6 Deployment Models and IPv6 Solutions." http://www.inetcore.com/material/log/material_060803100723.pdf.

[53] ARIN Number Resource Policy Manual, section 8.1. http://www.arin.net/policy/nrpm.html.

[54] For example, the largest US ISPs need not pay ARIN more than $18,000 per year. "ARIN Fee Schedule." http://www.arin.net/billing/fee_schedule.html.

[55] "Transition Mechanisms for IPv6 Hosts and Routers." RFC 1933.

[56] "IPv4 Address Transfers." http://www.apnic.net/policy/discussions/prop-050-v001.txt.

[57] "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.

[58] Cramton, Peter. "A Review of 'Markets for Clean Air'." Journal of Economic Literature. Vol. XXXVIII, pp. 627-633. September 2000. http://www.cramton.umd.edu/papers2000-2004/00jel-markets-for-clean-air.pdf.

[59] 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.

[60] "IPv4 Transfer Policy Proposal – ARIN XXI Public Policy Meeting Draft Transcript." http://www.arin.net/meetings/minutes/ARIN_XXI/ppm1_transcript.html#anchor_11.

[61] ARIN Budget. http://www.arin.net/about_us/corp_docs/budget.html.

[62] Herrin, Bill. "BGP Cost" http://bill.herrin.us/network/bgpcost.html.

[63] "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.

[64] "IPv4 Transfer Policy Proposal – ARIN XXI Public Policy Meeting Draft Transcript." http://www.arin.net/meetings/minutes/ARIN_XXI/ppm1_transcript.html#anchor_11. "IPv4 Transfer Policy Proposal." http://www.arin.net/policy/proposals/2008_2.html.

[65] "IPv4 Transfer Policy Proposal – ARIN XXI Public Policy Meeting Draft Transcript." http://www.arin.net/meetings/minutes/ARIN_XXI/ppm1_transcript.html#anchor_11.

[66] See e.g. Gerth, Jeff. "Hunts Again Charged in 1979-80 Silver Deals." New York Times. March 1, 1985. (The Hunts' market operations caused the price of silver to increase from $11 per ounce to $50 per ounce, before dropping back to $11.)

[67] "IPv4 Address Space." IANA. http://www.iana.org/assignments/ipv4-address-space. "Recovering IPv4 Address Space." ICANN. http://blog.icann.org/?p=271.

[68] Roth, Alvin E. "What have we learned from market design?" Hahn Lecture, Economic Journal, 118 (March), 2008, 285-310.

[69] ARIN RSA, version 9.1, clause 15(a). http://www.arin.net/registration/agreements/rsa.pdf. See also the RIPE NCC Standard Terms and Conditions, clause 9.1. http://www.ripe.net/docs/ripe-321.html. See also the APNIC statement on Transfer of Membership and Resources, provision 4. http://www.apnic.net/member/member-transfer.html.

[70] ARIN RSA, version 9.1, paragraph 3. http://www.arin.net/registration/agreements/rsa.pdf. See also the RIPE NCC Standard Terms and Conditions, clause 9.4. http://www.ripe.net/docs/ripe-321.html. See also the APNIC Membership Agreement, clauses 5.1.(a),(c). http://www.apnic.net/docs/corpdocs/membership-agreement.html.

[71] "Planning Report 05-2: IPv6 Economic Impact Assessment." National Institute of Standards & Technology, Department of Commerce. October 2005. http://www.nist.gov/director/prog-ofc/report05-2.pdf.

[72] "IPv4 IANA Allocation Report." http://bgp.potaroo.net/ipv4-stats/allocated-iana.html.

[73] Cramton, Peter. "A Review of 'Markets for Clean Air'." Journal of Economic Literature. Vol. XXXVIII, pp. 627-633. September 2000. http://www.cramton.umd.edu/papers2000-2004/00jel-markets-for-clean-air.pdf.

[74] See e.g. ARIN Public Policy Mailing List, discussion entitled "APNIC policy proposal to create a regulated market in IPv4 addresses." August 2007. http://lists.arin.net/pipermail/ppml/2007-August/thread.html.

[75] For example, some networks currently null-route spammers. See e.g. "Mail Server Blocking Policy." PA.net. Available at http://www.pa.net/news/mailblocking.html. Other networks currently ignore route announcements with long prefix lengths. See e.g. Popescu, Alin, et al. "Routing Announcements: Are Small Prefixes Globally Routable?" October 2007. http://www.nanog.org/mtg-0710/presentations/renesys-lighting.pdf.

[76] Huston, Geoff and Mark Kosters. "Update on Resource Certification." March 2008. http://cidr-report.org/presentations/2008-03-12-resource-certs.pdf.