Assessing the Supply Curve in the IPv4 Address Market

Lee Howard, Kelly Brooks

Abstract

The scarcity of unallocated addresses in Internet Protocol version 4 (IPv4) has led to a nascent market in addresses. Prior work (Mueller, Howard) considered demand on that market; in this paper, we analyze the supply, identifying four tranches (tiers) in which addresses will come to market.

For each tier, we estimate the quantity of addresses that could potentially be supplied to the market, based on public routing table data (RouteViews Project). Based on interviews with brokers, we estimate the cost of bringing those addresses to the market, which serves as a lower limit to the actual market price.

Introduction

Internet hosts are identified by an Internet Protocol (IP) address, which serves as a numeric locator and an identifier, so Internet protocol packets can be routed from their source to their destination. IPv4 address space includes some 4.2 billion possible addresses; since this is not enough to address Internet growth (Evans), it is slowly being supplanted by the 340 trillion trillion trillion addresses available in IPv6. During this slow transition, many devices and computers may be dual-stacked, that is, run both address families simultaneously.

An IPv4 address typically takes the form 192.0.2.17, where the decimal (base 10) numbers between the dots can be anything from 0 to 255; this is referred to as "dotted decimal" notation. In the binary numbers used by computers and routers, this address would be 11000000.00000000.00000010.00010001. Routers, the devices that determine where an Internet packet should go next to reach its destination, would need to know whether a device was local or remote. Originally, they used the first 8, 16, or 24 bits ("binary digits") to determine how many bits were in the local network: the first two bits indicate whether it's a Class A (01), Class B (10) or Class C (11).

In the 1980's and early 1990's, IPv4 addresses were allocated in "classes," blocks of 16 million (Class A), 65,536 (Class B), or 256 (Class C) addresses. By 1993, a new strategy was required to be more efficient, and Classless Inter-Domain Routing (CIDR) was developed. The address 192.0.2.14, instead of being a Class C network with 255 possible hosts (192.0.2.1 – 192.0.2.255), could be part of a network of, for example 32 hosts, in 192.0.2.0/27. CIDR allows more granular assignment sizes, according to the need of the organization. Instead of using the first two bits, CIDR includes a second number (followed by the slash in the 192.0.2.0/27 example) indicating the number of bits designating the network; remaining bits are specific to the host. Therefore, a higher number indicates a smaller block of addresses. Each IPv4 address contains 32 bits, so the number of addresses in a CIDR block is $2^{(32-n)}$ where *n* is the CIDR notation; e.g., a /18 has 2^{14} or 16,384 IPv4 addresses.

Though CIDR allowed for assigning blocks more appropriately sized to the organization (instead of, say, 16 million addresses for Apple's 50,000 employees), still more addresses were needed. IPv6 addresses use hexadecimal (since they would be too long to write in decimal). A typical IPv6 address, with CIDR notation would be 2001:db8:1234:5678:10bd:84ff:fe95:106e/64. The IPv6 address size is 128 bits, so a /64 is $2^{(128-64)}$ or about 170 billion billion.

IPv4 and IPv6 addresses are allocated from a central registry, the Internet Assigned Numbers Authority (IANA), to Regional Internet Registries (RIRs). The IANA is operated by the Internet Corporation for Assigned Names and Numbers (ICANN), according to principles developed by the open Internet Engineering Task Force (IETF), originally under contract to the U.S. Government but now operated for the public good. The RIRs are roughly continental in scope (Africa, Asia-Pacific, Europe and the Middle East, South America and some Caribbean states, North American and other Caribbean and Atlantic islands). Each has its own open policy development forum and process. The RIRs allocate addresses to organizations in their regions, according to demonstrated need as defined by their policies, to Internet Service Providers (ISPs), data center/hosting companies, or other enterprises.

Several RIRs, however, are already running out of IPv4 address blocks. In February 2011, the IANA allocated its last blocks, and two months later, APNIC (the Asia Pacific Network Information Center) allocated its final address blocks before its last /8, triggering its end-stage policy—organizations can receive a final, single /22 (1000 addresses) and no more—effectively running out. In September 2012, RIPE NCC (Rèseaux IP Europèen Network Coordination Centre) similarly triggered its end-stage policy. ARIN (the American Registry for Internet Numbers), LACNIC (the Latin America and Caribbean Network Information Center) and AfriNIC (the Africa Network Information Center) still have some limited time (see Huston, IPv4).

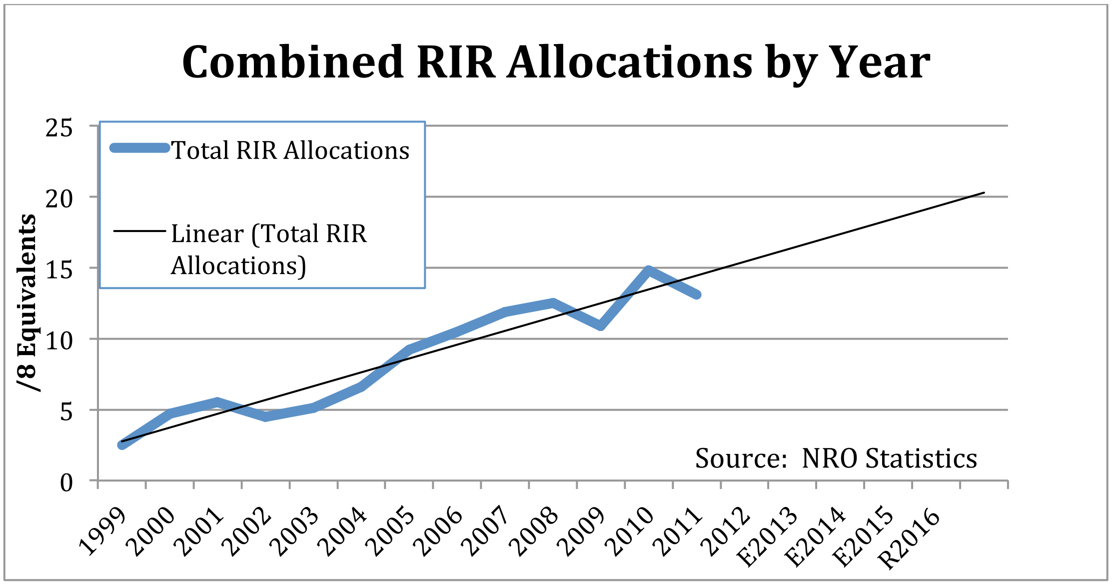

Figure 1: RIR Allocation Trend

In the years prior to IANA's last IPv4 allocation, demand for IPv4 addresses grew worldwide, driven by Internet growth, from both new subscribers and new devices. Once IPv4 addresses are unavailable from RIRs, networks adding subscribers or devices will have a few options: switch to IPv6, deploy address-sharing technologies, and/or buy IPv4 addresses. Since IPv6 is still available nearly free from the RIRs (and with a power of four more addresses, is expected to last a long time), it is one obvious long-term solution. However, organizations are moving their networks, content and products to IPv6 at varying rates, so it is not an effective substitute for IPv4.

Another option is the use of address-sharing technologies, such as carrier-grade network address translation (CGN, also called large-scale NAT, LSN, or NAT444), which allows multiple IPv4 devices to share a single IPv4 address. It may be deployed to use IPv4 addresses more efficiently (to serve more customers, therefore more revenue) (BITAG). However, it is an imperfect technology, interfering with some applications (Donley) and may therefore be an expensive, imperfect substitute for IPv4 (Howard).

As long as IPv4 addresses are available from RIRs, firms can continue to obtain addresses at very low cost. As RIRs run out of IPv4 addresses, however, firms will assess when and whether to transition to IPv6 or to seek IPv4 addresses from other sources. This newly emerging market for IPv4 addresses is immature and generally opaque. Address transfers tend to be made directly between organizations, and transaction prices are not usually made public. This analysis of the potential address supply, in addition to the (Mueller) study of this emerging market, should help increase market visibility.

Tier 0: Available

In Tier 0, the RIRs provide addresses to organizations that can show they need addresses. The requirement to show "need" varies among and within RIRs, according to local policies established by their communities. The "need" requirement is usually documentation showing that previous allocations have been used to a specified level of efficiency and additional documentation describing how the next allocation will be used to the same level of efficiency within a specified period of time. This need-based requirement provides a control on demand.

The pricing structures vary among RIRs, but are very low. For example, ARIN's member categories vary from $500 for 1,000 IPv4 addresses (/22) to $32,000, for members with more than a million IPv4 addresses (/12) (about $0.50 to $0.03, and note that organizations with many millions scale to even less per address).

As noted, IPv4 addresses are no longer available from APNIC (Asia) and RIPE-NCC (Europe and Middle East), except for a single /22 per organization from the RIRs' last /8. LIRs (local Internet registries, generally ISPs) are therefore encouraging enterprises to apply directly to the RIR, both to conserve the LIRs' space, and to provide enough addresses for the enterprise. Other organizations, whose perceived need exceeds the rules set by the local RIR, including those in other regions, are in the Tier 1 market to buy addresses.

The remaining RIRs (North America, South America, Africa) have about eight /8 equivalents remaining, 512 million addresses.

From 2008 to 2011, the RIRs were issuing about 12 to 15 /8 equivalents in total (Figure 1) (NRO). The rate of allocation had been consistent with the growth of the routing table (Huston, BGP), as firms began routing their newly assigned address blocks. As of middle 2013, there are fewer than 3 /8 equivalents remaining in the pools of ARIN, AfriNIC, and LACNIC, respectively. Since the RIRs are generally restricted to allocate within their regions, the rate of runout for these three RIRs may vary.

Tier 1: Abandoned

Space that has been allocated, but is no longer in use for any purpose can be considered abandoned. The analysis by (Mueller) shows several transfers in bankruptcy, at prices from $9-$12 per address, for various block sizes including fairly small /24s, often facilitated by a broker. We can supplement that work by analyzing the potential supply.

It is possible to find the number of addresses that are being used on the public Internet, by analyzing public resources such as the (RouteViews). There are no abandoned Class A blocks, and this analysis excludes privately routed and reserved blocks (IANA).

An analysis of Class B space, undertaken by searching the routing table (RouteViews) for every possible Class B address (128.0 – 191.255) shows nearly 7,000 unrouted Class B's, the equivalent of twenty-seven /8's. This is nearly half a billion addresses, an estimate consistent with broker statements (Burns, Fried, Thimmesch), though (Sterchi) says that only a fraction (maybe 20%) of this space should actually be considered supply as far as a market is concerned, due primarily to the frictional costs of completing transactions. Many of these may be used for private networks, not connected to the public Internet. Some address space may have been assigned to firms that have since merged or been acquired without good accounting of the address space; as firms look to grow in IPv4, or to sell addresses, this space may be recovered and put to use.

The cost of introducing these blocks to an IPv4 address market is relatively low. If the current registered holder of addresses wants to sell—or is instructed to by bankruptcy court (Ryan)—he or she may try to find a buyer (for instance, by listing the availability on ARIN's Specified Transfer Listing Service), or by approaching a broker. This market is currently active (Mueller), and address space is already listed.

Commonly, buyers or brokers identify potential address blocks and contact the current owner to gauge their willingness to sell. This call is often an address holder's first time hearing of the market; brokers describe the time it takes to convince a holder that addresses can be sold as "sounding like a Nigerian prince scam" (Thimmesch), since no value had been assigned.

RIRs keep records for each address block in their region, in a database called "Whois." Each block will have an organization record and a point of contact record. In many cases the RIR's records are outdated, and about 30% of the time the "title" isn't completely clear, such as when company names change such as when the company name changes, it was acquired or was spun out (Sterchi). The RIR will require extensive paperwork to update records, to make sure they don't register a block with the wrong organization. Documenting the provenance of the numbers to the satisfaction of the RIR and of prospective buyers is necessary to confirm a sale. None of this work is particularly difficult, but it does require some professionalism and can be time-consuming. Brokers are vague about how much labor this work costs but all describe it as significant, and (Sterchi) said, "At a typical commission, we break even on a /16. For blocks around that size, it would hopefully be part of a broader business relationship, but not necessarily a business in and of itself." At a 5-10% commission (Burns) on a $650,000 transfer, we infer costs approaching tens of thousands of dollars. The profitability varies significantly, as (Fried) notes, "A/16 at $10 on 10% commission is a lot of money for an individual doing 5-6 per year. For four people doing conferences and doing marketing and doing commission-sharing with a network of finders that's not a lot."

If a broker makes 5-10% (above) of a $9-12 sale price (Mueller), and says that 65,000 addresses (Class B) is the smallest profitable block (above), we find profits of $29,000 to $78,000. Cost to market could be lower if the seller's RIR records are current, and they choose not to use a broker. Many /24 blocks (legacy Class C, or other small blocks issued after CIDR) registered to defunct organizations, points of contact are unaware of the possibility of an address market, so many will not find their way to market. For a Class C (255 addresses) to yield $29,000 - $78,000 on the same 5-10% commission, the price per address would have to be well over $1,000, a price no one expects. Even a /18 (8,000 addresses) would be over $70 to yield that commission.

The cost to recover a block of addresses is the same, regardless of the number of addresses in the block. If a broker is satisfied with $29,000 - $78,000, then they could theoretically be satisfied with the same commission for a Class A. In practice, sellers and brokers have established an expectation of price and commission and are unlikely to discount quite that much. Brokers are well informed, and are likely to encourage sellers to expect higher prices (Thimmesch).

Many sellers are not highly motivated to sell, commonly saying, "I've had it for 20 years" (Thimmesch) and often hold addresses rather than take less than the perceived market rate. Some have said they will wait for $100, or until the three largest RIRs (APNIC, RIPE-NCC, ARIN) are completely out, not just on their austerity policies (Thimmesch). The established price may even be arbitrary, "plucked out of the air by some lawyer from Microsoft" (Burns) buying addresses from Nortel in bankruptcy (Mueller).

Transactions to date show a significant number of /24 transfers (Mueller), but more addresses are transferred in large blocks. For a buyer, a /24 might be enough to implement CGN, but would not be enough for most growing Internet companies. Buyers may prefer to have a single aggregate block: firewall rules, filters, and routing configurations are all simpler with a single aggregate entry than with multiple disaggregate blocks.

As the price rises in this tier, more sellers are likely to become aware of the value of their address space, and work to make them available, partly as more brokers make contact. ISPs may be buyers, up to about $29-40 USD (Howard), when CGN becomes a viable substitution. Content providers may be buyers, and may have less ability to substitute CGN, but may be able to substitute IPv6. As the available quantity of tier 1 space ebbs, tier 2 space will become available.

If one assumes a continuing demand trend as in (Figure 1: RIR Allocation Trend), one can see a short time period for this segment of an IPv4 market, probably less than two years.

Tier 2: Underutilized

Organizations with blocks significantly larger than they are currently using are considered to be underutilizing their address space. For enterprise organizations, ARIN requires 50% utilization of a previous IPv4 assignment to justify an additional assignment; therefore, this threshold is considered sufficiently utilized. A firm using more than 50% would logically request more address space while it was in Tier 0. An organization using less than half of what is needed for additional space may thus be considered "underutilizing." In this tier we include the partially routed blocks from Figure 2 (about six /8 equivalents).

Some early registrants include several governmental agencies. However, financial incentives are unlikely to persuade them to renumber and release their addresses. Although it is possible for a legislative action to force return of some of these address blocks to the free pool, there are no current efforts to do so, so those blocks will be excluded from this consideration.

A survey of the Class A space (IANA), among the earliest allocations, shows that a significant portion is not being routed on the public Internet (RouteViews).

Prefix	Designation	%Routed
003/8	General Electric Company	100
004/8	Level 3 Communications, Inc.	100
006/8	Army Information Systems Center	5.8
008/8	Level 3 Communications, Inc.	100
009/8	IBM	0
011/8	DoD Intel Information Systems	0
012/8	AT&T Bell Laboratories	100
013/8	Xerox Corporation	12.2
015/8	Hewlett-Packard Company	100
016/8	Digital Equipment Corporation	100
017/8	Apple Computer Inc.	100
018/8	MIT	100
019/8	Ford Motor Company	0
020/8	Computer Sciences Corporation	9.0
021/8	DDN-RVN	0
022/8	Defense Information Systems Agency	0
025/8	UK Ministry of Defence	0
026/8	Defense Information Systems Agency	0
028/8	DSI-North	0
029/8	Defense Information Systems Agency	0
030/8	Defense Information Systems Agency	0
032/8	AT&T Global Network Services	100

033/8	DLA Systems Automation Center	100
034/8	Halliburton Company	1.6
038/8	PSINet, Inc.	100
040/8	Eli Lilly and Company	25
044/8	Amateur Radio Digital Communications	100
048/8	Prudential Securities Inc.	0
051/8	UK Government Department for Work and Pe	0
052/8	E.I. duPont de Nemours and Co., Inc.	24
053/8	Cap Debis CCS	100
054/8	Merck and Co., Inc.	6.2
055/8	DoD Network Information Center	100
056/8	US Postal Service	0.8
057/8	SITA	100

Figure 2: Percentage of Legacy Class A Blocks Routed on the Internet

There are many valid reasons for registered networks being unrouted, including:

1. Classified networks. We might assume (though unlikely) that all unrouted U.S. Defense Department (and related agencies) and U.K. Ministry of Defence address space belongs to this category: networks 11, 21, 22, 25, 26, 28, 29, 30.
2. Large industry private internetworks or extranets. We might assume that the block allocated to Ford is in this category: 19.

Among the twelve Class A's above that are registered but unrouted, an equivalent of nine Class A's could plausibly be sold in an IPv4 market. Blocks for which only a portion is routed may be underutilized, which would bring them into Tier 2. Excluding partially-routed space and the above exceptions leaves only two complete Class A's available for a market (fractions of 9, 13, 20, 34, 40, 48, 52, 54, 56).

Among the remaining Class A holders, several are publically traded companies, who although they are injecting their entire block into the routing table, may not need to entire block. In fact, if each of seventeen companies with Class A's used IPv4 addresses only for Internet access for their few hundred thousand employees, even allowing an extra few hundred thousand for corporate systems and inefficiencies inherent in subnetting, nearly 16M addresses would still be available from each. This would make another some 270 million addresses available on the market in Tier 2.

A similar analysis of Class B space is difficult, because a company of only 20,000 employees may be using a Class B's 65,656 addresses fairly efficiently, while a company of 500 may not be. The variation relative to the size of the address block is significant, compared to Class A space. However, the requirements for acquiring a Class B were to have a network too large for a Class C and too small for a Class A; it is reasonable to assume that many of these blocks are numbered to less than 50% efficiency. For an enterprise network in the ARIN region, greater than 50% efficiency justifies more addresses. Therefore, of the 9,000 Class B blocks not included in Tier 1, one might reasonably conclude that most of them could be more used more efficiently, and that the equivalent of 4000 /16 blocks, or about 15 Class A equivalents (250,000,000 addresses), could be made available.

Combined, the 270 million addresses in Class A and 250 million in Class B space extends the market by another 520 million addresses, consistent with estimates by (Burns, Fried, Thimmesch).

Making those addresses available to a market is more difficult than recovering abandoned space [Carpenter]. In addition to the work required to bring Tier 1 addresses to market, work is probably required by the holder of space to renumber more efficiently. Network devices will have been configured assuming a single large block of addresses; managers will have to spend time updating their firewall rules, DHCP servers, logging systems, and potentially Internet-facing servers, to use only the smaller block they will keep. These tasks can be time-consuming, potentially tens or even hundreds of hours for larger networks, performed by qualified IT staff. For a small office with only a few dozen computers and a /24, updates may be as simple as a firewall entry, a router, DNS, a server or two, some printers, a few computers; a couple of hours' work, maybe a thousand dollars for a consultant. At the high end, many more systems require updates, but automation is more likely to be in place; still, maintenance windows must be planned and reviewed. Brokers report clients quoting $1-2 (Sterchi) to $3 (Fried) per host as a low number.

A 50% utilized Class B, however, requires renumbering 32,000 hosts to release 32,000 addresses. A cost of $1-3 ($32,000 - $96,000) will be included in the price. However, if brokers are unprofitable at $10-12 for fewer than 65,000 addresses, then prices, commissions, or block sizes would be expected to increase. Further, CIOs who were unmotivated to sell at Tier 1 prices may still be reluctant to distract from their primary work (where their bonuses are paid) to renumber without a high level of confidence that they will see a significant return. Therefore, payment must be made or guaranteed in advance, or the expected price must be high enough to be worth the risk of non-payment following renumbering work. With Tier 1 pricing at $9-12 per address, the cost to bring Tier 2 addresses to market will therefore be at least $10-16, and based on the other factors (size, CIOs' motivations) prices could be significantly higher.

As price per address rises, smaller blocks may become available. There are not many Tier 1 blocks of /19 or greater in legacy Class C space, but as smaller blocks become profitable to bring to market, there may be abandoned blocks in the /24 to /19 range that were unavailable during Tier 1. Some may become available in Tier 2, but since common practice is not to route anything smaller than a /24, an enterprise with a /22 would have to renumber a /23 and two /24s to be saleable, and the marginal price would have to be at the high end of the range.

Tier 3: Substitutable

Enterprise networks, seeing the price of IPv4 addresses rise, may find that they can substitute alternatives for their unique addresses, and sell addresses that they had been using. That is, where Tier 2 addresses were under-utilized, Tier 3 are currently in use, but could be replaced. Some alternatives include IPv6 and CGN, potentially in combination.

There are limitations to substitution. CGN, using one IPv4 address for many devices, does not work for all applications; externally reachable systems (such as web and mail servers, but also peer-to-peer applications) are particularly vulnerable (Donley, BITAG).

The other main substitution option, IPv6, is not supported by all services and applications. If a system is reachable via IPv4 but not IPv6, then IPv6 cannot be considered a substitute for IPv4. Broad support for IPv6 makes it more useful, a network effect. As it becomes a substitute for IPv4, demand declines, and organizations may be able to sell off their IPv4, increasing supply.

As a Tier 3 example, with prices sufficiently high, an enterprise with a large number of IPv4 addresses might decide to renumber all of its employees' workstations behind a CGN. This is already common in enterprise networks, and for those covered by Sarbanes-Oxley auditing requirements, is often required by auditors. This renumbering would cost about the same as Tier 2 address space renumbering. Some small block of addresses may still be needed for Internet-facing servers, or an enterprise could sell its entire Class B to a hosting company, who would use a fraction for the enterprise's Internet-facing servers, and use the rest for other customers.

In theory, all IPv4 address space would be in Tier 3.

Eventually, some organizations will find that they can use IPv6, and need no globally unique IPv4 addresses at all. At a profit per consumer of $140 per year (Howard), if prices exceed $70 per address, residential ISPs may begin migrating customers to IPv6-only to resell the IPv4 addresses. At this tier, they determine that the few remaining sites and devices that do not support IPv6 are worth less to them than the price for which they can sell their IPv4 addresses. A sufficiently high price may convince an enterprise that they are better off selling their few remaining addresses (no longer reserved for NAT44 or NAT64) and in a few cases might pay the remote end to transition. Significant payments are unlikely, since anyone who wants to be reachable will need to have IPv6 anyway.

A final pressure for substitution is that if the only IPv4 addresses remaining in the market are fragments of what were once aggregate blocks, routing tables may increase, with an entry for each fragment; increasing the costs to use the addresses (for both the buyer and their network peers). Concern about the routing table was a consideration in allowing a market to emerge (Li), with estimates on costs offered (Herrin). Brokers run the gamut of expectations on the degree to which address space will become fragmented.

Exceptions and Caveats

As noted, some of the unrouted (abandoned or underutilized) address space may be in use on private networks. These private networks may also substitute IPv6 as IPv4 prices increases, but since the networks are private, is difficult to estimate the difficulty. It is also possible that some of the unrouted address blocks are generally used on the Internet, but were not visible at the time tested, or from the location tested.

In addition, many blocks of addresses outside of classful address space may be abandoned, or underutilized, or substitutable. This paper analyzed unrouted Class B space and Class C space that was assigned before CIDR, but a similar analysis is needed of all remaining address space, for every possible /18 (assuming smaller blocks are uninteresting, and skipping Class B space) for Tier 1 and down to /22 for Tier 2. Because more recent assignments are less likely to be abandoned, and because they are smaller, additional space will not fundamentally change this analysis.

Continuing Demand

IPv4 addresses in Tier 0 supply will soon be unavailable. Once IPv4 addresses are only available in Tier 1, costs to ISPs will increase, which may increase the price of Internet access for consumers and businesses. Tier 1 addresses, being easier and cheaper to bring to market, may well be nearly exhausted before Tier 2 addresses are available. Similarly, addresses from Tier 3 are less likely to be available until Tier 2 approaches exhaustion.

Extrapolating from the rate of demand prior to IANA runout (as seen in Figure 1), Figure 4 shows how long the total supply available (in Tiers 0-3) might last. This assumes demand continues, either on the same linear growth, or at the volume seen the year before IANA runout, or based on 2012 allocations (once APNIC and RIPE-NCC stopped allocating at regular rates). It is not clear what organizations in Europe and Asia did for addresses in 2012, but if they had hoarded addresses before runout or received addresses from an intermediary, the perceived demand may return once those opaque supplies are exhausted.

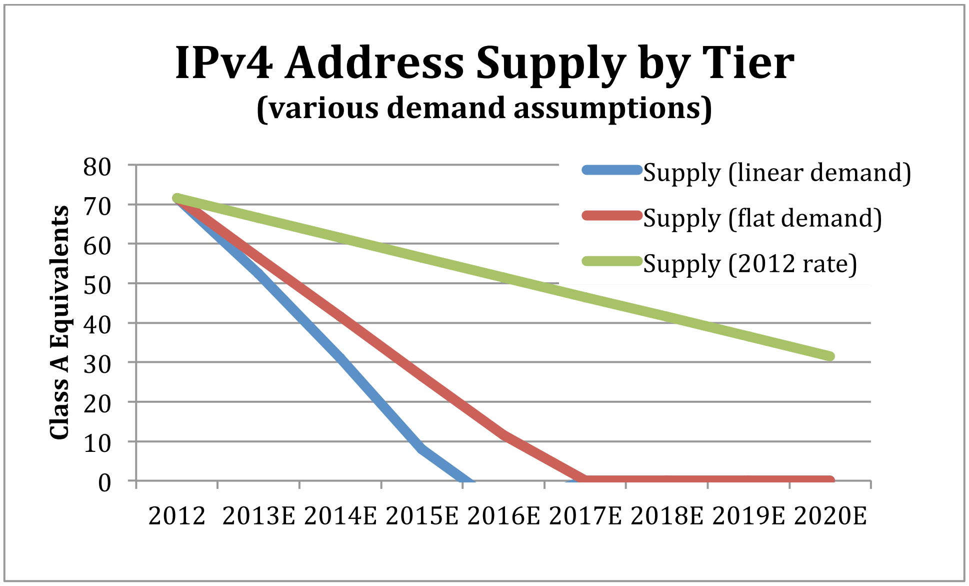

Figure 3: Address Supply for Several Trends

No pricing information is provided here, so this cannot be substituted as a demand curve. When prices are included, demand may shift, as CGN (and eventually, IPv6) is considered a substitute at some prices. Rather, it shows how relatively small the potential supply is compared to historic demand. Similarly, Figure 4 shows the relative supply from each tier, assuming a draw down at the linear extrapolation rate.

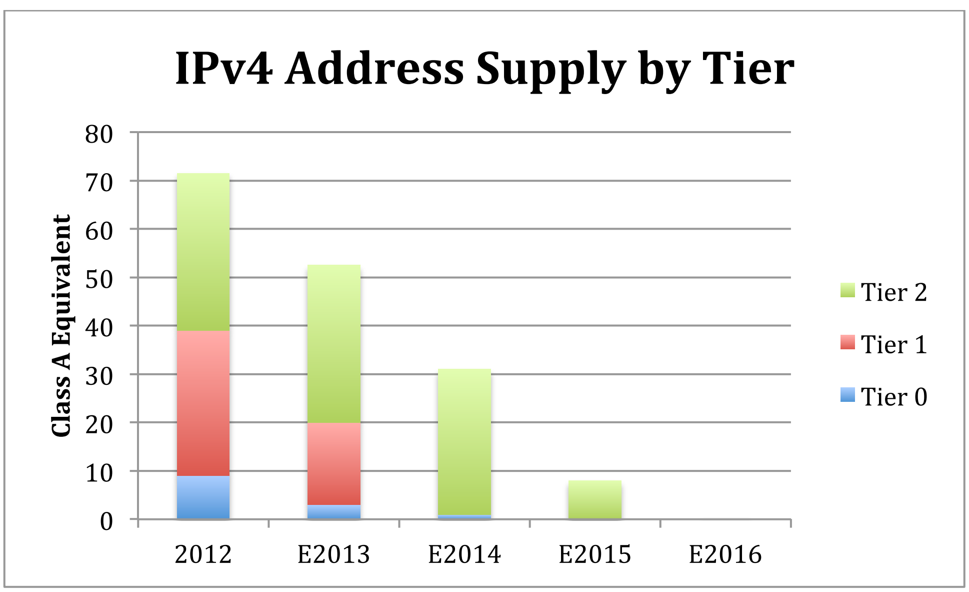

Figure 4: Addresses over Time

Conclusions

Based on this analysis, a billion addresses could supply a market for a few more years. Prices are likely to rise significantly as each tier runs out, but there is insufficient data to know whether prices will rise proportionately with cost.

Tier	Summary	Cost per Address	Addresses Available	Years (inelastic demand)
0	Remaining RIR space	$0.03 - $4	144,000,000	2012 - 2014
1	Abandoned	$9 - 12	480,000,000	2012 – 2016
2	Underutilized	$10 – 16	520,000,000	2016 - 2017
3	Substitutable	Up to $x00	All IPv4	2017 - 2025

Figure 5: Summary of Tiers, Pricing, Timing

Importantly, demand was greatly reduced in 2011 and 2012, so the linear projection may not continue. It is possible that organizations had hoarded a few years' worth of address space, and demand will rise again, or that demand has been permanently suppressed. Again, the "Cost per Address" is clearly not a market price, but a minimum cost to bring IPv4 addresses to market.

Further work is needed to predict the demand curve, and approaches to minimize disruption through the transition to IPv6.

Works Cited

BITAG, "Implications of Large Scale Network Address Translation (NAT)," http://www.bitag.org/documents/BITAG_TWG_Report-Large_Scale_NAT.pdf March 2012.

Carpenter, Brian, R. Atkinson, H. Flinck. "RFC5887: Renumbering Still Needs Work." http://tools.ietf.org/html/rfc5887, May 2010.

Donley, Chris, Lee Howard, Victor Kuarsingh, John Berg and Jinesh Doshi. "Assessing the Impact of Carrier-Grade NAT on Network Applications." http://tools.ietf.org/html/rfc7021, September 2013.

Edelman, Benjamin. 2009. "Running Out of Numbers: Scarcity of IP Addresses and What to Do About It." *Springer-Verlag Lecture Notes of the Institute for Computer Sciences, Social- Informatics and Telecommunications Engineering* series, Volume 14, 2009, pp 95-106.

Evans, Dave. "The Internet of Things: How the Next Evolution of the Internet Is Changing Everything." *Cisco Internet Business Solutions Group*, April 2011.

Herrin, Bill. "What Does a BGP Route Cost?"

http://bill.herrin.us/network/bgpcost.html.

Howard, Lee. 2012. "Internet Access Pricing in a Post-IPv4 Runout World." http://www.asgard.org/images/pricing_v1.3.docx

Huston, Geoff. (Huston, BGP) "Growth of the BGP Table – 1994 to Present." http://bgp.potaroo.net retrieved 4 Dec 2012.

Huston, Geoff. (Huston, IPv4) "IPv4 Address Report." http://ipv4.potaroo.net retrieved 19 Mar 2013.

IANA, Internet Assigned Numbers Authority. "IANA IPv4 Address Registry." http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml retrieved 17 November 2012.

Li, Tony. 2007. "Router Physics—IP Markets." ARIN XX (20th meeting of the American Registry of Internet Numbers), October 17, 2007.

Milton Mueller, Brenden Kuerbis, and Hadi Asghari. "Dimensioning the elephant: An empirical analysis of the ipv4 number market." Proceedings of the Telecommunications Policy Research Conference, 2012.

NRO, Number Resource Organization, "Internet Number Resource Status Report as of 30 September 2012." http://www.nro.net/wp-content/uploads/nro-stats-2012-q3.pdf, September 30, 2012.

RouteViews Project. http://archive.routeviews.org/oix-route-views/2012.11/oix-full-snapshot-2012-11-12-2200.bz2. For this analysis, a complete routing table was retrieved 12 November 2012.

Ryan, Stephen, Mathew Martel, and Ben Edelman. 2012. "Internet Protocol Numbers and the American Registry for Internet Numbers: Suggested Guidance for Bankruptcy Trustees, Debtors-in-Possession, and Receivers." *BNA's Bankruptcy Law Reporter*, 24 BBLR 32.

Interviews

Burns, Mike, IPTrading.com, Phone interview on market, March 20, 2013.
Fried, Gabe, Hilco Streambank, Phone interview on market, March 19, 2013.
Sterchi, Louis, Kalorama, Phone interview on market, March 20, 2013.
Thimmesch, Peter, Addrex, Personal interview on market, Herndon, VA, March 18, 2013.

Acknowledgements

The author would like to thank Kirk Erichsen and Wes George of Time Warner Cable for their review, and Ben Edelman of Harvard University for suggesting this line of investigation.