Cost to Renumber and Sell IPv4

Lee Howard

Abstract

Organizations with allocations of IPv4 addresses may be able to sell those addresses. Whether it is profitable to do so depends on the cost to move from IPv4 to IPv6, and the market value of those IPv4 addresses. This paper illustrates how to estimate the cost to renumber into IPv6, to make IPv4 addresses available for sale.

Introduction to IPv6

Every device on the Internet has an address that helps other devices locate and identify it. The format of the address is defined in the Internet Protocol (IP). The format defined by IPv4 included only about 4.3 billion possible addresses, too little for the current Internet. The format defined by IPv6 allows for 340 trillion trillion trillion possible addresses. However, IPv6 does not natively communicate with IPv4: the designers assumed everyone would run "dual-stack," running both IPv4 and IPv6 at the same time, until IPv4 was no longer needed.

IPv6 is appealing for several reasons:

1. Growing companies may be constrained by a lack of IPv4 addresses.
2. Content loads faster for mobile users over IPv6 than IPv4.[i]
3. Companies can sell their existing IPv4 addresses for a profit.[ii]

Organizations with an allocation or assignment of addresses directly from a Regional Internet Registry (RIR)[iii] can sell or designate a transfer of addresses. The exception is AfriNIC, where policy still prohibits sale.

Adding IPv6 is a technical undertaking that is not part of most companies' core competencies. It isn't especially difficult, but can be time consuming and requires coordination.

Most plans call for enabling IPv6 alongside IPv4 (native dual-stack), then later withdrawing IPv4. In most cases, companies will use a translator to provide access to legacy IPv4 web sites and corporate systems. "Native" refers to networks without translation. Network Address Translation (NAT) includes:

- NAT64: an IPv6-only client connecting to an IPv4-only server

- NAT46: an IPv4-only client connecting to an IPv6-only server, often in an IPv6-only data center
- NAT44: an IPv4-only client connecting to an IPv4-only server, often to share IPv4 addresses among many clients.

A network that is nearly all IPv6 but requires some legacy access may deploy IPv6 plus IPv4-as-a-Service, where NAT46/NAT64 provide legacy connectivity. Reducing IPv4 from the network as far as possible will maximize the number of IPv4 addresses for sale. Organizations with more IPv4 addresses than they need can run native dual-stack and sell excess IPv4 addresses to fund other initiatives.[iv]

Costs to Deploy IPv6

Native dual-stack, where IPv4 and IPv6 run in parallel, is widely considered the gold standard of transition technologies. There is some justification for this: with "Happy Eyeballs"[v] implemented, a device will choose the best protocol (or at least not one that is failing). That provides robustness against a failure in a protocol or its routing, and eases anxiety during the deployment of a new technology.

The costs of deploying dual-stack can be split into major efforts[vi]:

- Preparation and Assessment
- Data Center
- User connectivity

The first step should be assigning a cross-functional team, with a project manager and experts in security, applications, server systems, network, user experience, and user support. This team will need a week or two to assess the effects of IPv6 on their areas. This time does not include actual testing, but a first estimate on what will need to be tested, configured, and/or upgraded. Team members should report back on what upgrades are required, what work will need to be completed, how long each piece of work will take, what dependencies exist (determining the order of work items). Often, IPv6 unit testing can be added on to planned projects, as when a system upgrade is being tested. The project manager should be able to develop a draft project plan, and will need the team to help assess the affects on other projects and priorities. Before work progresses, make sure an executive sponsor has seen a cost estimate (like the one shown in Table 1) and a revenue estimate, and that they approve the work.

The team may find legacy systems or custom applications built on IPv4-only platforms. Examples include custom applications using mainframes, SAS, PeopleSoft, or Oracle, or applications that do not support an operating system more recent than 2007.[vii] These systems may work with a simple

translator (such as NAT46) in the server room, or may cost a significant amount of capital and time to upgrade to a modern platform. Legacy systems are legacy for a reason: if translation doesn't work, the cost to upgrade or replace may be significant. Assess the upgrade cost, then compare to current and expected future IPv4 prices; this may be the opportunity to update the system.

It takes little time to ask vendors what their IPv6 capability is; if it is not on their roadmap in time for the business, it takes little extra time to ask their competitors what their IPv6 capability is. Explaining the business reason that IPv6 support is needed (the estimated revenue from sale of addresses) helps set feature prioritization.

Legacy systems notwithstanding, capital costs are generally a fraction of the labor costs; as a rule of thumb, any hardware sold since 2008 can support IPv6. Depending on the organization's accounting rules, the labor may be capitalizable, since it is applied toward selling addresses.

Most organizations spend more time and money on IPv6 training than necessary, and it slows down assessment and testing. For an engineer who understands IPv4 subnetting and ARP (Address Resolution Protocol), it should take less than a day of focused, independent reading to become comfortable with IPv6 architecture. Study topics should include the address format, Router Advertisements (RA), Neighbor Discovery (ND), SLAAC (State-Less Address Auto-Configuration), and IPv6 security.[viii] The rest will look similar to IPv4.

Enabling the Data Center / Server Room

Most administrators enable IPv6 in the data center or server room before enabling user devices, to make sure users can be provisioned and managed and can reach critical systems. Native dual-stack is generally considered best practice; there are fewer servers than users, so there is comparatively little advantage to sharing IPv4 addresses. It is possible to use a translator at the data center edge, as Facebook does (using IPv6 exclusively inside the data center).[ix] This configuration can ease the way for moving from dual-stack to IPv6-only or IPv6+IPv4aaS quickly.

At a minimum, add IPv6 to web servers and externally-facing systems (gaming, chat/IM, downloads, etc.). Even if no other systems are IPv6-enabled, public servers should provide the best connectivity available, which may be IPv4 or IPv6 for any given client.

Enabling IPv6 in the data center requires an address plan, security consideration, testing, routing, switching, and server configuration. The address plan should include a /64 prefix for every LAN or VLAN; there is no

reason to conserve IPv6 addresses by assigning less than a /48 to the data center.

Security should be designed into the transtion plan early. There are few new security problems with IPv6, but a few hours of research and configuration planning can protect against ND Cache Exhaustion attacks, ping-pong attacks, or other simple vulnerabilities. ICMPv6 does require allowing specific message types.[x] Other than that, firewall policy updates may be as simple as adding IPv6 addresses to existing zones/domains along with existing IPv4. If the policies will be kept up to date, add IPv6 addresses to firewall policies all at once, even for systems that may not immediately be updated, so that the protections are already in place once it is time to deploy. Most Intrusion Detection Systems (IDS)[xi] support IPv6, but may need to have IPv6 enabled. Vulnerability scanners must support IPv6, in both testing and production. Log parsing tools must support the larger address format.

Routing and switching require testing, but their configuration is only about as difficult and time-consuming as deploying a new routing protocol, or other major new feature. Only after testing and configuring security systems and routers should IPv6 be enabled with the Internet provider; otherwise, IPv6 connectivity will be uncontrolled. Then, following testing, individual servers or load balancers, or groups of related services, can use IPv6 as they are ready.

Most of the time spent in transition will be spent in testing, since each component must be individually tested, and then tested in a Quality Assurance (QA) environment that closely simulates the production environment. Testing includes making certain clients can connect, and servers can interoperate, even during the transition. Testing the rollout plan will reduce the odds of failure during the deployment, which is especially important over a long deployment time.

Enabling IPv6 to Users

Much of the work done to deploy IPv6 to the data center is required before deploying to users, and much of the work can be reused. Specific work includes an address plan, host security, routing and switching, provisioning, and peripherals. As with the data center, most of the project time will be spent in testing, especially testing users' applications.

Developing an address plan often takes time and causes stress. There are resources and consultants available, but it may be most efficient to develop an address plan that is extravagant with addresses, and modify later as needed. There is no need to be stingy with addresses, as it often causes trouble later. Plan for a couple of hours of labor on the first draft, but plan for multiple reviews and revisions.

Host security includes assessing anti-virus and certificate checking, as well as Group Policy Objects (GPOs) or other policy enforcement. This might include device authentication (802.1x, Network Admission Control, etc.) and guest or temporary network access.

Security and switching are related in a way not found in IPv4. In IPv4, a host (computer, tablet, phone) typically acquires an address via DHCP, and the DHCP log is the record of who connected to the network. In IPv6, default behavior is that routers send Router Advertisements (RA) announcing the network prefix they can route, and hosts can assign themselves an IPv6 address using SLAAC (State-Less Address Auto-Configuration). A good practice is to configure switches to send log messages reporting new entries in the neighbor table (e.g., "neighbor binding logging"); with good log correlation, this will help detect unauthorized connections to the network. DHCPv6 is also often used to provide additional configuration information; decide whether DNS servers will be provided over IPv4 or IPv6.

For security, it is important to complete testing before deploying IPv6 to any routers or switches in the user network. While all current operating systems (including desktop and mobile devices) support and enable IPv6 by default, peripherals such as printers may not. Again, it takes little time to check a vendor's IPv6 support, and little more time to check a competitor's support. Even if devices need to be replaced, the cost may be far less than the revenue from the sale of addresses.

Following IPv6 deployment, some IPv4 renumbering may be required to consolidate IPv4 addresses to be sold into a single block.

Phase	Task	Hours	Hourly Rate	Expense
Plan	Assess	240	$100	$24,000
Plan	Training	80	$150	$12,000
Data Center	Address Plan	15	$150	$2,250
Data Center	Security	40	$150	$6,000
Data Center	Testing	400	$125	$50,000
Data Center	Routing/Switching	20	$150	$3,000
Users	Address Plan	20	$100	$2,000
Users	Security	40	$150	$6,000
Users	Routing/Switching	30	$150	$4,500
Users	Testing	200	$125	$25,000
IPv6-only	Translator	40	$150	$6,000
IPv6-only	Renumber IPv4	40	$100	$4,000

TOTAL	1925	$144,750

Table 1 – Example of Cost Estimate to Enable IPv6 on a Large Enterprise Network

Number of IPv4 Addresses	Price per Address	Revenue
16384 (a /18)	$11	$180,224

Table 2 - Revenue from Sale of IPv4 Addresses

Table 1 includes a sample estimate of the cost of adding IPv6 on a medium-large enterprise or small ISP. It is offered for illustration only; the IT Department or Network Engineering or Operations should be able to provide an estimate specific for their network. A smaller network or one with less formal processes may spend only a fraction of that time.

Table 2 is shown for comparison; an address broker can provide a current estimate of the value of the organization's IPv4 addresses. Again, the network experts in the company should be able to provide a precise count of IPv4 addresses, reserving some for dual-stack Internet-facing systems.

Cost to Turn Off IPv4

With IPv6-only, there is no access to IPv4-only systems. All servers must be IPv6-enabled, or they will be unreachable by IPv6-only clients. For this reason, IPv6-only is not recommended for general-purpose Internet communications. Two exceptions exist: green field networks, and IPv6 plus translation.

For a green field, walled-garden network, where one organization controls the endpoints and the network, IPv6-only is a sensible way to build. For instance, systems that do not need general Internet access, but only need to reach a central system or each other, are perfect cases for IPv6-only. The cost for deploying a new system on IPv6 is theoretically the same as for IPv4, but in practice such systems often require enabling dual-stack for Internet somewhere (for instance, to fetch operating system updates). Still, the effort is incremental, and can provide valuable experience later.

Enterprise networks doing IPv6-only or IPv6+IPv4aaS need to make sure they can push software updates, security policies, anti-virus updates, and Certificate Authority (CA) servers (specifically, CRL or OCSP lists) to client workstations. The servers hosting these services should be IPv6-capable, but must have it enabled before clients can be securely migrated from IPv4; these elements are not a concern for dual-stack migrations.

The process of deactivating IPv4 is not complex: shut down IPv4 DHCP servers, then gradually remove IPv4-specific commands from switches and

routers. Again, testing is recommended, since implementations may exist that assume the existence of IPv4. Clients may continue sending DHCP Discover messages to find an IPv4 DHCP server; this allows for a fall back to IPv4, but in the long term just generates unnecessary broadcast traffic.

An additional warning to the architect of the IPv6-only system: every failure will be blamed on IPv6. For many people, change is intimidating and IPv6 is not well understood (to them).

Conclusions

Many organizations holding IPv4 addresses can make a profit by renumbering into IPv6 and selling IPv4 addresses. Others may use proceeds from selling IPv4 addresses to cover the cost of desired upgrades, such as updates to legacy systems. The complexity of the network and the amount of time spent testing will be the major factors in the cost, and should be assessed early to determine whether the project is financially sensible.

Notes

[i] According to Facebook http://www.internetsociety.org/deploy360/blog/2015/04/facebook-news-feeds-load-20-40-faster-over-ipv6/, and LinkedIn https://engineering.linkedin.com/blog/2016/07/ipv6-at-linkedin-part-i--chippin-away-at-ipv4, and Akamai https://blogs.akamai.com/2016/06/preparing-for-ipv6-only-mobile-networks-why-and-how.html

[ii] $14 each as of this writing, depending on volume; more on this later; http://www.ipv4auctions.com/previous_auctions/

[iii] The RIRs are: ARIN (The American Registry of Internet Numbers), RIPE-NCC (Réseaux Internet Protocols Européens – Network Coordination Centre), APNIC (Asia-Pacific Network Information Centre), LACNIC (Latin America and Caribbean Network Information Centre), and AfriNIC (Africa Network Information Centre).

[iv] For example, MIT adress sale: https://gist.github.com/simonster/e22e50cd52b7dffcf5a4db2b8ea4cce0

[v] Wing, D. and Yourtcheknko, A., "Happy Eyeballs: Success with Dual-Stack Hosts", https://tools.ietf.org/html/rfc6555, April 2012.

[vi] Adapted from Chittimaneni, K., Chown, T., and L. Howard, V. Kuarsingh, Y. Pouffary, E. Vyncke, "Enterprise IPv6 Deployment Guidelines", https://tools.ietf.org/html/rfc7381, October 2014.

[vii] Windows Vista supported IPv6 by default, and all versions of Unix had IPv6 before that.

[viii] A simple web search on each of those terms will provide sufficient reading. Just reading the relevant RFCs will take less than a day and be authoritative. An engineer who does not understand IPv4 (subnetting, ARP) will need remedial study.

[ix] http://www.internetsociety.org/deploy360/resources/case-study-facebook-moving-to-an-ipv6-only-internal-network/

[x] Davies, E., "Recommendations for Filtering ICMPv6 Messages in Firewalls," https://tools.ietf.org/html/rfc4890, May 2007.

[xi] Including related systems like IPS (Intrusion Prevention Systems) and NADS (Network Anomaly Detection Systems)