

July 2007

Geoff Huston

IPv4 Exhaustion Nears

Funny how some topics seem sit on a quiet back burner for years, and then all of a sudden become matters of relatively intense attention. Over the past few weeks we've seen a number of pronouncements on the imminent exhaustion of the IP version 4 address pools. Not only have some of the Regional Internet Registries (RIRs) and some national registry bodies made public statements on the topic, we've now seen ICANN also make its pronouncement on this topic.

ICANN Board Resolution: **On the Deployment of IPv6**

Whereas, the unallocated pool of IPv4 address space held by IANA and the Regional Internet Registries is projected to be fully distributed within a few years;

Whereas, the future growth of the Internet therefore increasingly depends on the availability and timely deployment of IPv6;

Whereas, the ICANN Board and community agree with the call to action from the Address Supporting Organization and the Number Resource Organization, Regional Internet Registries, the Government Advisory Committee, and others, to participate in raising awareness of this situation and promoting solutions;

The Board expresses its confidence in the Internet community to meet this challenge to its future prospects, and expresses its confidence in the bottom-up, inclusive, stakeholder-driven processes in place to provide any needed policy changes, and;

The Board further resolves to work with the Regional Internet Registries and other stakeholders to promote education and outreach, with the goal of supporting the future growth of the Internet by encouraging the timely deployment of IPv6.

[<http://www.icann.org/minutes/resolutions-29jun07.htm>]

Similar announcements have been made by the RIRs:

ARIN: <http://www.arin.net/v6/v6-resolution.html>,

LACNIC: http://lacnic.net/en/anuncios/2007_agotamiento_ipv4.html

Why the sudden uptake of interest in this topic? I suspect that a small part of this may be my fault! For some years now I've been maintaining a [web page](#) that looks at the accumulated data regarding consumption of IPv4 addresses, and using a relatively simply mathematical model applied to the past 3 years consumption data derives predictions of future consumption. For many years the best fit appeared to be an exponential curve.

The general form of an exponential function is

$$y = e^{f(x)}$$

A simple exponential function, $y = e^{ax + b}$, models an environment of accelerating growth where the total population doubles in fixed time intervals.

During the [RIPE 54](#) meeting in May 2007 it was pointed out to me that this curve didn't look much like the data it was trying to model. The recent growth rates in address consumption appear to be higher than those that correspond to an exponential curve. In response, I tried a few other curves, and found that the recent consumption model appears to be best matched by an O(2) polynomial function.


The general form of an O(2) polynomial, or quadratic, equation is

$$y = ax^2 + bx + c$$

A quadratic function models an environment where the rate of growth increases at a constant rate over time.

A presentation on this change in the IPv4 consumption predictive model can be found at: <http://www.potaroo.net/presentations/2007-05-09-ripe54-ipv4.pdf>, and the associated prediction report can be found at: <http://ipv4.potaroo.net>

The result of this change in the curve function of the predictive model was quite dramatic – rather than looking at an exhaustion date for IPv4 addresses of around 2012 to 2014, which appears to be so comfortably off in the distant future as to be inconsequential to today's Internet industry, the exhaustion date has drawn in to late 2009 / early 2010. This is just a little over two and a half years from today and all of a sudden a rather abstract debate about the viability of various options to cope with this address exhaustion issue is looking uncomfortably real.



What this prediction is saying is that some time between late 2009 and late 2011, and most likely in mid-2010, when you ask your local RIR for another allocation of IPv4 addresses your request is going to be denied. Not because you do not meet the policy criteria, but simply because by then the RIR would've exhausted its pool of available IPv4 addresses. Now the timing of when this will happen may vary a little depending on which region you are working in, and it may vary a little depending how much of a "last chance" panic overruns the industry in the coming months, but the outcome is inevitable in any case. We're running out of IPv4 addresses.

Then what?

If only that were a simple question to answer!

The Grand Plan, as originally envisaged in the 1990's when IPv4 address exhaustion was first predicted, is that we task the IETF to develop a new IP protocol with an address pool that is amply massive and that by the time IPv4 address exhaustion is supposed to become a critical issue the entire internet industry would've already completed a complete transition to this new protocol and the entire issue of IPv4 address exhaustion would be a non-event.

And, at least initially, things went according to this plan. The IETF considered various proposals (Steve Deering's SIP and Paul Francis' PIP with a bit of TUBA mixed in) as the next generation IP protocol, subsequently named IPv6.

The IETF spent some time developing this protocol, and, modulo some slight forms of refinements, the protocol specification was completed in December 1995 with the publication of RFC 1883 (subsequently refined in December 1998 as RFC 2460). So far so good, as by December 1998 we were still only using some 30% of the total IPv4 address pool, and there appeared to be ample time to complete the remaining steps of this plan.

The next step in the transition plan was to enlist vendors of equipment and systems to support this new protocol. In some cases this has gone according to plan, and many of the mainstream vendors support IPv6 in their equipment. Windows XP came with configurable IPv6 support, and Windows VISTA comes shipped with IPv6 enabled as the preferred protocol. Thanks to the amazing achievements of the Japanese [Kame project](#) there was a solid IPv6 implementation for Unix systems. And Apple's MAC OS has had IPv6 enabled "out of

the box" for some time now. The major router vendors all support IPv6. Well in a fashion. These days there are few customers out there who are in a position to require multi-gigapacket per second switching performance for IPv6, and router vendors typically don't over-engineer equipment beyond customers' requirements, so its sometimes that case that the IPv6 switching path through the router may be slower than the IPv4 switching path. But, on the whole, so far so good for this part of the transition plan. What about consumer retail equipment, such as the home DSL WiFi units and similar? Here the story is a little different, and support for IPv6 is not ubiquitous. Again the issue of the market comes into play, and vendors of such equipment in the price-sensitive consumer electronics sector typically provide precisely what the customer requires and no more. So if the customer does not require IPv6, then you won't see it in the box!

At this stage the transition plan envisaged IPv6 deployment in the network, in servers, in clients and anywhere else that you'd find IP. And here is where the wheels appear to have fallen off the plan, because deployment simply has not happened.

To demonstrate the extent of IPv6 deployment to date, a comparison of the basic metrics of the size of the IPv6 and IPv4 networks is illustrative. The IPv4 network as of June 2007 includes some 230,000 routing entries, of which there are 105,000 "root" prefixes, and some 26,000 autonomous system numbers. By comparison, the IPv6 network contains 875 routing entries, of which 731 are "root" prefixes and some 749 autonomous system numbers.

So what are our options?

The first option is that industry actually undertakes a comprehensive deployment of IPv6, supporting "native" IPv6 in the network in servers, in infrastructure such as the DNS, and in all other places where we use IPv4 today, running in parallel to the current IPv4 network. Of course part of the reason why this has not already happened is that such a deployment is neither costless nor completely simple. IPv6 is a second protocol, requiring a second set of forwarding tables in the network's switching elements, maintenance of a second routing domain, the deployment of a second network management domain, and a second operational support domain. None of this reduces the existing IPv4 support workload, so the additional activity is an additional imposition on the service provider. But its unlikely that customers will see a different value proposition when the service provider switches from an IPv4 service to full dual-stack IPv4 / IPv6 support. Its still the web, its still email, it's much the same Internet, and worth no more in terms of retail price than the IPv4 Internet. The expectation is that the costs of transition to a dual stack service network are not accompanied by incremental revenue from customer base. So the observation that IPv6 deployment has failed to materialise so far is not necessarily a condemnation of industry as being short-sighted or ignorant about the issue of IPv6 deployment, but a pointer to a possible failure of the business model where the typical drivers for investment in new infrastructure are not apparent in this transition plan.

And its not as if IPv6 is the only option, particularly in the short to medium term.

NATs have proved to be incredibly prolific in today's Internet simply because the business case is so effective that it has overrun the very real technical shortcomings of the technology. They provide elements of control over connection initiation, as well as the potential to "compress" address use by overloading the UDP and TCP port address fields. Yes, NATs compromise the "end-to-end" coherence of the network, but one can also observe that NATs have already destroyed the last vestiges of ubiquitous end-to-end in the network, and any application that is deployed in today's Internet must factor in NAT behaviours or, in practical terms, the application is undeployable. One possible short term response to the looming IPv4 address exhaustion is increased density of NAT deployment, and the associated deployment of multi-party application architectures that perform more complex rendezvous operations in order to set up the application. A classic case of this approach can be seen in the SIP-related work, with the development of protocols such as STUN and TURN designed to perform dynamic discovery and negotiation of NATs. There is some feeling that we are still a long way from the practical limits of NAT deployment, and from the service provider's perspective, as long as the costs of this option are not ones that they have to bear directly, then this looks like a very attractive short term option.

In the longer term, however, NATs have their limitations. The recent trend of applications that can dynamically discover the presence of NATs along the path and then generate NAT binding keepalive traffic to maintain a semi-permanent address binding through the NAT tends to reduce the efficiency of the NATs. Also the increased deployment of NATs introduces the issues of paths that contain 2 or more NATs, and the techniques of dynamic discovery of NAT behaviours tend to perform poorly under such conditions.

Of course the exhaustion of the RIR's unallocated address pools does not imply the end of the road for IPv4 networks. IPv4 will probably be around for some time yet, and there will be enterprises who want more IPv4 addresses and presumably there will be enterprises who have more IPv4 addresses than they have a need for. From such conditions markets typically emerge, and it's reasonable to surmise that redistribution of IPv4 addresses will happen within some form of market-based regime in the short term. As a short term response to the exhaustion of the IPv4 address pool, there is some merit in supporting such a market. The alternative, that of a black market, has the potential for various forms of distortions and chaos in the area of address management, with its own attendant risks in terms of the security and stability of the Internet during a somewhat critical phase of forced transition.

How long such a market remains useful for the industry is a matter of some conjecture. It may even be the case that such an explicit pricing function being placed on IPv4 addresses provides a far clearer signal to industry about the incentives of transition to IPv6 than any form of public notice or ICANN resolution! In any case a market cannot make the finite infinite, and while a market, with its attendant pricing function, can act as an incentive for more efficient use of the addresses in play, ultimately the situation where demand exceeds supply dictates a response of moving to a technology that can provide significantly larger address space.

Normally, such transitions are undertaken through the realisation of some benefit to the customer – cheaper, faster, more flexible, more functions, more colours, or other attributes that are clearly visible to the customer. But here there are no real changes to the package – IPv6 does not really present any superior performance, price or value to the end customer that is not already available in IPv4, and the pressure to transition is not a case of self interest at work, but more a case of common interest.

There's a "network effect" going on here. There is no benefit whatsoever being the only IPv6 speaker in a world of IPv4. Similarly, there is no residual benefit being the last IPv4 speaker in an otherwise all-IPv6 world. Would an ISP invest in dual protocol support across their entire network for a single additional customer who could only speak IPv6? Would a server operator add IPv6 support for potentially one extra client who could only speak IPv6? Neither situation is likely, as the incremental potential benefit from that single customer is far lower than the incremental cost of dual stack operational support. And the example is somewhat artificial, but the more substantive question is: What's the threshold here? How widely does IPv6 need to be deployed before the business case for IPv6 deployment starts to make some financial sense? As an ISP, or as a server operator, when do you know that this potential benefit threshold has been reached that justifies the incremental costs of IPv6 support?

The entire issue here is that the network effect works only when there is critical mass to drive it, or when there are sufficient folk out there willing to take higher risks with their initial investment. Neither condition is happening right now out there in the industry, which is an undeniable current problem. The Internet has turned into a low margin commodity service trying to work in a highly price competitive market, with extremely low levels of high risk capital. As a commodity utility industry we are now about as conservative and risk-averse as one could get. And as a predominately deregulated industry we are dominated by short term issues which tend to be addressed through perceptions of where self-interest lies on the part of each player. When confronted with a broader problem that actually demands some level of response in concert with considerations of longer term common interests, then we're finding it hard to resolve in a sensible manner.

It appears that this industry has driven itself into a rather messy place right now, and extracting itself from this state in a sane and rational fashion, and without excessive levels of blame shifting as we do it, is indeed quite a challenge!

We've got around two years left of IPv4 address distribution as we knew it, and then that particular world comes to an end.

So what should we be doing now?

This article was originally published as "The End of the (IPv4) World is Nigh!" on www.potaroo.net.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001.

<http://www.potaroo.net>