

April 2008

Geoff Huston & George Michaelson

IPv6 Deployment: Just where are we?

I suppose we all know the plot by now. Back in the early 1990's it was forecast that we would run out of IPv4 addresses well before we were ready for the Internet to stop working, and that we'd need to transition the Internet to use a new version of the IP protocol that allowed for a larger address space if we wanted to avoid the problem. In response to this, the technical community came up with a specification for IP version 6 by the mid 1990's and we were all set.

And then nothing happened.

We didn't run out of IPv4 addresses, and we didn't adopt IPv6. Indeed by late 2002, at the time of the meltdown of the euphoric Internet bust our appetite for additional IPv4 addresses was at a low, and at the time we were consuming IPv4 address rate at some 16 million addresses per year. At this rate we still appeared to have some decades left before we came close to exhausting IPv4's 4 billion address size. Any sense of urgency with IPv6 adoption also receded at the time.

But the wheel has turned some more in the intervening six years, and the pace of growth of the internet has once more picked up. Since 2005 we've been consuming IPv4 addresses at a rate of some 160 million addresses per year, and what's left in the IPv4 unallocated address pool looks like lasting for no more than another 3 years. This situation looks a lot like Peak IPv4, where the rate of expansion of the Internet has never been as high as it has been in the past couple of years, and the IPv4 address consumption rate is also at all-time peak levels in recent months, while at the same time we draw inexorably closer to the exhaustion of the current supply feed from the unallocated IPv4 address pool. So, in 2008, we are visibly running out of IPv4 address space and the timeline to exhaustion is now a couple of years rather than a couple of decades.

To be true to the original transition plan we should be well down the track of uptake of IPv6 at this stage. After all, by the time we reach the end of the IPv4 supply we'd like to have completed the IPv6 transition and have no need to actually use that last IPv4 address.

Well, that was the plan.

How close are we sticking to the transition script? How well are we doing with IPv6 deployment?

We've seen a number of exercises in recent months that look at various parts of this question. One of the more prominent exercises has been a "lights out for IPv4" trial at a number of recent Internet standards and operations meetings, where the IPv4 Internet was deliberately switched off for a period, leaving the meeting attendees with an IPv6 only service. These exercises are useful in terms of assessing some aspects of preparedness for IPv6, both in terms of the capabilities of your local host and in terms of accessibility of services via Ipv6. In other ways they are perhaps a little wide of the immediate operational objective, as the scenario of operating in an IPv6 only network is the end outcome of a potentially rather lengthy transition phase of dual stack operation. Where we are now is in the early stages of this dual stack phase of operation, where end hosts, networks, services and middleware should be in the process of shifting from IPv4 only to some form of dual stack IPv4 and IPv6 support. So if we are now in this dual stack world, just how much dual stack operation is really out there today?

In this article we'd like to look at some measures of the use of IPv4 and IPv6 protocols in today's Internet and see if we can draw any conclusions about just how far down the track we are with the IPv6 part of dual stack deployment.

In this exercise we'll use a number of measurements that have been made consistently since 1 January 2004 to the present, where we can distinguish between IPv4 and IPv6 use in various ways. No single data set can answer the global question of just how much IPv6 is out there in the Internet and how much is IPv6 being used relative to IPv4. Indeed that's something that we can't directly measure and report on, so, instead, we take some small windows that have some bearing on this larger question, and see what we can infer from the data we have to see how that may fit into the larger picture. The primary sources we'll use here are the inter-domain routing table, access logs from a web service and data captured from DNS queries.

The Routing View of IPv6

The first data set we'll use is the inter-domain routing table, and the first plot is the number of entries in the global IPv6 routing table.




Figure 1 – IPv6 BGP Table Size

This figure plots the number of routing entries in the global IPv6 routing table for the period 1 January 2004 until March 2008 on an hourly basis. If this is reflective of the uptake of IPv6 deployment, then the story is one that has some very volatile elements within the IPv6 routing domain. There was evidently a period of steady growth across 2004 and 2005, and then a significant downturn in the first half of 2006. The event in February 2006 was evidently the result of a bug in a BGP implementation where some 150 "ghost routes" were held in the routing system for many months, and then flushed out from the global table when the particular BGP peering session was reset. The second downward correction in June 2006 was related to the formal winding up of the 6Bone network on the 6th June 2006.

It's also evident that interest in IPv6 deployment has picked up from mid-2007, and the number of routing entries has grown 33% from 900 to 1200 in the most recent 9 months. There is a lot of volatility in the data, and the abrupt appearance and disappearance of up to 120 routes in a single event is indicative more of a research and test rig than a fully stable production environment. But overall this data appears to be indicative of a growing level of interest in IPv6.

But how does this compare to the comparable picture in IPv4?

Figure 2 shows the comparable view of the IPv4 inter-domain routing table over the same period. The IPv4 view of the same measurement is much more like the traditional "up and to the right" plots that we've been used to seeing with the Internet, indicating a consistent strong level of underlying growth in the network.

In both the IPv4 and IPv6 cases we've seen a doubling in the size of the routing table over the period in question, with the IPv4 growth picture being one that has been uniformly spread across the entire 51 month period. There is a critical point of difference here, however. The IPv4 routing table is on the whole very stable and while there are individual events that appear to involve one or two hundred routes, this is a relative level of routing volatility of around 0.1% in IPv4, as compared to a comparable measure of routing volatility some 15% or more in IPv6. The observation here is that there are two quite different "signatures" of routing in the two protocol environments.




Figure 2 – IPv4 BGP Table Size

On a relative scale, how do these two sets of numbers compare? What's the relative level of deployment of IPv6, as measured in the routing system, as compared to the deployment of IPv4? Figure 3 plots the relative proportion of IPv6 routing entries to IPv4 routing entries over this same period.




Figure 3 – IPv6 / IPv4 BGP Table Size Ratio

The data indicates that the IPv6 network is some 0.45% of the IPv4 network in size, when averaged across the 4 year period. It also shows that since mid-2007 the IPv6 network has been growing at a faster rate, in terms of number of routing entries, than IPv4. On the other hand, it has to be admitted that a growth from 0.38% to 0.48% in relative terms is not exactly major news!

Of course the number of entries in the inter-domain routing table is not the only metric of deployment size of these two IP protocols. It may be more useful to look at the number of routing entities that are routing IPv6,

where each autonomous routing entity, or ISP or corporate network, is counted as a "routing entity". In this case its not the number of entries in the BGP routing table per se, but the number of unique AS numbers that are contained in the routing table that indicate the number of entities that have IPv6 networks of one form or another that interconnect in the global IPv6 internet.

Figure 4 plots the number of AS numbers in the IPv6 routing table since 1 January 2004, again using an hourly snapshot of the data.




Figure 4 – IPv6 AS Count

This plot shows a more even picture of IPv6 deployment than the number of IPv6 routing table entries would suggest from Figure 1. It shows a larger than doubling in growth in the number of these routing entities, from 400 to 930 over this 4 year period. Also discernable in the data is the acceleration in growth from mid 2007, although its not as prominent a feature of the overall data sequence.

This data indicates a consistent and steady enrolment of new entities that are announcing routes into the IPv6 network, and while there is no spectacular exponential trend in the data so far its also that case that the lpv6 network continues to grow.

Again, this IPv6 data can be compared to the number of unique ASes that were visible in the IPv4 routing table over the same period. Figure 5 shows a comparable plot for the number of ASes in the IPv4 network. In terms of the number of ASes the IPv4 network did not quite double in size over this period, increasing from 16,000 ASes to 28,000 ASes. The relative smoothness of the IPv4 data series indicates a very high level of relative stability in the network. One potential explanation of this data is that the underlying inter-AS topology of the Internet in IPv4 is very stable, and much of the churn in routing may be attributable to the effects of other measures, such as traffic engineering or local efforts to optimise certain aspects of routing policies.




Figure 5 - IPv4 AS Count

These two data series can be compared in the same fashion as we've done with the number of entries in the routing table, and the relative number of ASes that are advertising IPv6 addresses, as compared to the number of ASes advertising IPv4 addresses, is shown in Figure 6.



Figure 6 - IPv6 / IPv4 Relative AS Count Ratio

This metric paints a somewhat different, and more positive, picture of IPv6 deployment than the comparison on the number of entries in the routing table. Here the relative metric of IPv6 as compared to IPv4 is some 2.8% over the period, and the number of AS entities actively routing IPv6 is growing at a faster rate than the IPv4 network.

The reason for the order of magnitude level of difference in metrics between the routing table size ratio and the AS number count ratio appears to have more to do with the inherent fragmentation of the IPv4 routing table. Due to the large scale use of more specific routes as a means of supporting traffic engineering and route policy expression, the IPv4 routing space shows a steadily increasing average number of routing entries per origin AS over this 4 year period. The following plot in Figure 7 shows the average number of advertised entries per origin AS in IPv4. The average number of routing entries per origin AS has grown from around 7.9 to 9.3 entries, and the growth rate appears to have increased sharply in 2008.




Figure 7 - IPv4 Average Routing Entries per origin AS

The comparable picture for IPv6, shown in Figure 8, shows that at this stage there is still strong evidence of provider based aggregate addressing with no major address fragmentation in the routing table.

Compared to an average of around 8.5 routing announcements per origin AS over the period 2004 to the present in IPv4, the IPv6 metric is an average of 1.3 announcements for the same period. This data accounts for the discrepancy in the relative metrics of the number of advertisements in IPv6 vs. IPv4 and the number of unique ASes in IPv6 vs. IPv4.

It is an interesting question as to why these two measurements are different. There would appear to be a number of factors at play here, including the lack of the so-called "legacy swamp" of un-aggregated addresses in IPv6, the lack of provider-independent address assignments in IPv6 until quite recently, the lack of large scale production networking in IPv6 that, in turn, has obviated any need for IPv6 traffic engineering, and the prevalent use of IPv6 as an overlay network, forcing the route policy determination role into IPv4 routing rather than IPv6 routing. So the relatively low number of addresses per origin AS in IPv6 may well be a sign of the current immaturity of IPv6 operational deployment and use, as distinct from any notable effectiveness in enforcing policies of provider-based aggregation into the IPv6 routing environment. The more recent upward trend in the number of entries per IPv6 origin AS may well be an inevitable consequence of a shift in IPv6 from initial small scale experimentation into various forms of pilots of broad scale deployment.




Figure 8 - IPv6 Average Routing Entries per origin AS

It would appear that from the AS origination data in the routing table the current picture is that some 3% of the Internet is IPv6 capable, to one extent or another, in terms of the headcount of distinct entities that compose the Internet. A caveat should be made here that it is not necessarily the case that if an AS originates an IPv6 routing advertisement that its entire network is dual-stack provisioned with IPv6, and that all of its end hosts and customers are dual-stack provisioned. So this 3% metric of IPv6 capability at the AS level is perhaps a reflection of levels of experimental and research interest in IPv6, as distinct from a reflection of the overall picture of IPv6 deployment as a deployed network service platform.

It may be more reasonable to observe that the extent of IPv6 use is certainly no higher than 3% at present, and may well be closer to a value one tenth of this number, or 0.3%.

A Usage View of IPv6

There are a couple of potential issues in the routing table view of the IPv6 Internet. Firstly, a metric of capability of supporting IPv6 in routing is not the same as a metric of actual use of IPv6 in terms of services on IPv6, and IPv6 packets that are sent across the network, and, secondly, this routing view does not take into account the transitional approach used by 6to4, and, more recently, Teredo where IPv6 is tunnelled across the IPv4 Internet and is not directly visible as distinct IPv6 routes in the routing system.

So perhaps we should refine this question of IPv6 deployment from a measure of routing capability of IPv6 to a measure of actual use of IPv6.

How much is IPv6 being used today relative to IPv4?

Has this metric changed in recent years?

And, while we're at it, how much IPv6 use is via the transitional tools of 6to4 and Teredo?

There are a number of ways of attempting to answer this question, and they include gathering long term traffic sampling data from an operational network through to a more controlled experiment using a sample point service point.

Web Server Logs

In this section we'll describe results from the second approach. The data set reported here relates to the use of the APNIC web site, <http://www.apnic.net>. This web site has both IPv4 and IPv6 addresses and has been dual homed on both IPv4 and IPv6 networks for over four years. The approach used here to measuring the relative use of IPv6 to IPv4 is to count the number of unique source addresses visiting this web page each day, and looking at the ratio of the number of unique IPv4 source addresses to the number of unique IPv6 source addresses. This approach hopefully removes out the factors associated with web crawlers (which for this site are still exclusively IPv4) as compared to human-driven web browsers, and also evens out some of the factors of the level of intensity of access and repeat visits to the same site.

Figure 9 shows the daily ratio of Ipv6 to IPv4 source addresses that are accessing the web site since 1 January 2004.




Figure 9 – IPv6 / IPv4 Web Access Ratio

There is some considerable variation in the data from day to day, and a scatter plot has been used here to ensure that the trends in the data are visible as well as the day-to-day variation. It may be the case that if we had access to the web logs of a dual-stack dual-protocol-homed web site with considerably greater volume levels the noise component of day-to-day variation might be lower than observed here.

However, even in this data series there are some interesting observations. The level of IPv6 use, as compared to IPv4, is around 0.2% for the period since 2004. This is a lower figure than was derived from the routing capability ratio based on the number of IPv6 announcing AS's, but does correlate to some extent to the relative number of routing table entries for the two protocol families.

It is also possible to discern a slight falling off in the IPv6 usage rate through 2006, and a resumption of interest in IPv6 at the start of 2007, which has continued to the present.

The anomalous data points that occur twice each year appear to coincide with APNIC member meetings. The web site under examination here is www.apnic.net, and at each of these meetings there has been IPv6 equipped as part of the network infrastructure for the meeting. While the number of sample points is low (8), it does point to a possible explanation that the population of end hosts that are IPv6 capable is generally higher than the number of active IPv6 hosts, and the relative use rate of IPv6 to IPv4 may be strongly influenced by the extent to which IPv6 and 6in4 tunnelling is allowed by access networks and end site configurations, as distinct from dual stack capability in the end host. When the end hosts are reconnected into an access environment that includes IPv6 access, the level of relative IPv6 access to the web site appears to increase notably. The most recent APNIC member meeting in February 2008 also included a period where

the IPv4 network on the meeting site was shutdown, leaving an IPv6 connection as the only possible external connection. The relatively high ratio recorded on that day of 1.2% is probably due to that exercise.

This web service access data also can show the relative use of the 6to4 and Teredo transition techniques within the total V6 usage data, as 6to4 use has a 'signature' source address prefix of 2002::/16 and Teredo has a comparable source address prefix of 2001:0::/32.




Figure 10 – 6to4 / IPv6 Web Access Ratio

Figure 10 shows the relative use of 6to4 IPv6 addresses to access the web site. While 6to4 was in quite widespread use in 2004 and 2005, its relative use was steadily declining from around 30% of all IPv6 access to less than 10% on average. Currently, the use of 6to4 appears to be steady at around 20% of all IPv6 usage. The high variation of the data on a day-to-day basis may be attributable to the relatively low volume of IPv6 traffic to the site, and additional data from sites with higher volumes of access may provide some further clarity here.

6to4 relies on access to a public IPv4 address and does not allow transition across IPv4 NATs. More recently, a number of operating systems have been equipped with Teredo, notably Microsoft's Vista. Teredo can tunnel IPv6 across IPv4 NATs. In Windows Vista and Windows Server 2008 most operating system components support IPv6. When both IPv4 and IPv6 are enabled, Windows prefers the use of IPv6 for applications that can use either IPv4 or IPv6. In the case of Teredo Windows Vista is enabled by default, although the local configuration may disable it, while in Windows Server 2008 Teredo is disabled by default.

Figure 11 shows the relative use of Teredo IPv6 addresses to access the web site. This is the same data series, but in this case we're looking for Teredo source addresses rather than 6to4. It shows that Teredo deployment commenced in late 2006 and its relative use is still around 5% or so of the total IPv6 use.

This web server access data represents an extremely small peephole into the larger Internet, so it could be the case that if one were to aggregate this with data from other sites with larger traffic levels, the picture of IPv6 use may change to some extent.




Figure 11 – Teredo / IPv6 Web Access Ratio

I'd like to update this description of Teredo with a correction provided by Christian Huitema:

"Teredo in Vista is enabled by default in a "dormant" state. That is, it will not be activated unless an application specifically requests it, either by specifically binding to the Teredo interface or by sending packets to a Teredo address. Moreover, contrary to what you say, when both IPv4 and IPv6 are enabled, Vista will prefer native IPv6 by default, but it will prefer native IPv4 to Teredo IPv6. Bottom line, Vista is designed to use IPv4 rather than Teredo when accessing a web site. It will only use Teredo if the web site does not have an IPv4 address."

This implies that our dual stack website will not automatically trigger Vista to use Teredo to access it, and we'll need to use a different measurement approach to see if we can measure the relative level of Teredo deployment.

DNS Name server Logs

Another long term data set that we can use here is usage data from a number of DNS servers. The DNS servers we are using for this data collection are servers for a subset of the reverse DNS PTR zones. These reverse DNS zones map IPv4 and IPv6 addresses back to domain names. What we are interested in here is the relative rate of queries that are made to the *in-addr.arpa* zone, that relate to resolution of IPv4 addresses and the queries that are made to the *ip6.arpa* zone, that relate to the resolution of IPv6 addresses.




Figure 12 – DNS V6 / V4 reverse PTR query Ratio

The relative proportion of these two query types is shown in Figure 12. This data appears to be quite different from the BGP and web server log data that we've reviewed so far in terms of resolution of the relative levels of use of IPv4 and IPv6. This data shows very low relative level of IPv6 queries (less than 0.01%) until 2007, and then some increased level of uptake, with a relative use metric of some 0.17% through 2008. There also appears to be evidence of a number of multi-day experiments in 2006 and 2007 that elevated the relative query rate for IPv6.

Its not entirely clear how to interpret this data. Firstly there is the issue of identifying those applications that perform these form of reverse DNS lookups, and secondly understanding the relationship between original queries for end host and the interaction between the original querying host and the various forms of DNS caches that may lie between the end host and the authoritative name server.

Further Work

There are many potentials windows on this particular aspect of IPv6 transition, as there is much more that could be done here. We would be very interested in learning of other long term data sets that could be used to distinguish some form of relative metric of IPv6 and IPv4 protocol use in the Internet.

More data on the relative use of IPv4 and IPv6 for dual-stack service points would also be helpful, in order to understand the trends in IPv6 usage in service delivery in the coming months, and the impact of host initiatives, such as the use of Teredo in the Vista release of the Windows OS, would also be useful in understanding the overall dynamics of IPv6 transition from the perspective of the balance of end host push and provider pull. So we would be interested to look at other web server logs of dual-stack servers that have been operating a consistent service model over some years to see if this picture of relative access across IPv6 and IPv4 that we've observed for the APNIC web site is also visible for other web sites and other services.

The DNS represents a rich vein of operational data and further iterations of this work could include an analysis of the relative rate of DNS queries for IPv4 address records and IPv6 address records, although the same caveats about the relative roles of DNS forwarders and cached DNS data, as compared to the rates of queries initiated by end hosts, and the queries as seen at the authoritative name servers needs to be factored into this particular DNS perspective of the relative use of the two protocols. In addition data from DNS forwarders may be useful in this context.

Further work in the BGP table could also illustrate the extent to which the IPv6 network is constructed using precisely the same inter-AS topology as the Ipv4 network, or whether the IPv6 network is still constructed as an overlay with a set of Ipv6 inter-AS relationships that appear to have a relatively small intersection with what could be reasonably assumed to be an underlying IPv4 inter-AS topology.

We have not reported on any study of actual data rates in an operational network. It would also be interesting to understand the relative ratio of IPv4 and IPv6 traffic, and the relative amount of traffic in the 6in4 tunnels, on operational networks today. So far we haven't been able to locate open sources for such data that have a long baseline of historical data, but if there are any offers of such operational data we'd be interested to examine it, and see how it correlates to these other measurements that we've studied here.

If you would like to share any such data, please contact us at research@apnic.net

Some Closing Observations

The Internet is facing some quite fascinating pressures in the coming years as the unallocated pool of IPv4 addresses depletes, and it's unclear at this stage just how quickly the Internet will transition to a IPv6 network and how such a transition will be deployed in the network. It's also unclear to what extent the Internet will be able to wean itself of the intensive use of NATs, and the relative pressures to undertake a transition to IPv6 or to persist with a deployment model of use of private IPv4 address space, NATs and various forms of protocol translation to fill the connectivity gaps, or even revert back to the networking model of the 1980s and revert to application level gateways to as an incredibly messy glue for tomorrow's network.


While much of the IPv6 technology set could be described as operationally ready, IPv6 hosts and service delivery platforms are being deployed, and a visible proportion of the network operating entities on the Internet are undertaking various forms of IPv6 deployments, the real level of uptake of IPv6 in the Internet today in terms of services and, by inference, packets, remains quite small. The metric that could be regarded as perhaps the best pointer of the current level of IPv6 usage is the web server access data, and the actual value of the relative rate of IPv6 use appears to be around 0.3% of the IPv4 use, or a relative level of 3 parts of IPv6 per 1,000 of IPv4.

The encouraging observation is that the relative use of IPv6 in today's internet as compared to IPv4 is increasing, so that while the internet continues to grow it appears that IPv6 use is growing at a slightly faster rate.

On the other hand, it also appears that while the relative numbers are increasing, IPv6 is still a very small proportion of the IPv4 Internet, and by whatever metric one would claim is a "critical mass" of IPv6 usage, we've yet to achieve and, and it may still be quite some years off yet.

Postscript

It was later pointed out to us that Google Trends is also a useful metric of IPv6 "interest": <http://www.google.com/trends?q=ipv4,+ipv6> - it shows a somewhat different picture of relative interest in IPv4 and IPv6!



Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

About the Authors

GEOFF HUSTON is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He graduated from the Australian National University with a B.Sc, and M.Sc. in Computer Science. 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 was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

<http://www.potaroo.net>

GEORGE MICHAELSON graduated from York University in 1982 with a B.Sc in Computer Science. His career in the United Kingdom and Australia is in academic and government research in computer science, networking, and systems administration and he is currently the senior Research & Development specialist at APNIC, the Regional Internet Registry serving the Asia-Pacific region. He is a member of the BCS, and a founding member of the Australian chapter of the Internet Society.