41st Meeting of the Working Party on Communication and Infrastructure and Services Policy, 15-16 June 2009, Paris

Introduction

This contribution was prepared by the Number Resource Organization (NRO). The NRO is comprised of the world's five Regional Internet Registries (RIRs) and is a member of the OECD Internet Technical Advisory Committee (ITAC).

The RIRs have collaborated to collect, analyse and make public global statistics on the deployment of Internet number resources (IPv4 and IPv6 addresses and Autonomous System (AS) Numbers) for several years. The RIRs are committed to supporting the OECD's overall goal of encouraging the deployment of IPv6 and look forward to further cooperation on this and other related matters.

For the 41st Meeting, the RIRs have prepared the following two documents for the delegates' consideration:

- Measuring IPv6 Deployment [ID: OECD001]
- Trends in IPv6 Allocations 2005-2008 [ID: OECD002]

These documents provide an introduction to a small part of the statistical analysis jointly carried out by the RIRs. This document is available online at:

www.nro.net

About the Regional Internet Registries (RIRs)

There are five RIRs in the world. Each RIR:

- Provides services related to the technical coordination and management of Internet number resources
- Participates in Internet community meetings and events
- Operates as a not-for-profit, self regulating membership organisation
- Facilitates policy development by its members and the Internet community via open meetings and mailing lists
- Is governed by a member-elected Executive Board
- Publishes publicly accessible documents about its activities

The five RIRs cover the following regions:

- AfriNIC (www.afrinic.net): Africa
- APNIC (www.apnic.net): Asia Pacific
- ARIN (www.arin.net) Canada, United States and several islands in the Caribbean Sea and North Atlantic Ocean
- LACNIC (www.lacnic.net): Latin America, Caribbean
- RIPE NCC (www.ripe.net): Europe, the Middle East and parts of Central Asia

About the Number Resource Organization (NRO)

The NRO was formed in 2003 and is comprised of all five RIRs. The purpose of the NRO is to undertake joint activities of the RIRs, including joint technical projects, liaison activities and policy co-ordination. The NRO serves as a coordinating mechanism for the RIRs to act collectively on matters related to the interests of the RIRs. It also offers a single point of contact for all interested parties to reach the five RIRs collectively. For more information about the NRO, please see:

http://www.nro.net

Statistics and Reports

The five Regional Internet Registries (RIRs) each provide publicly accessible daily updated reports on:

- IPv4 address ranges (IPv4)
- IPv6 address ranges (IPv6)
- Autonomous System Numbers (ASNs)

The reports are linked from:

http://www.nro.net/statistics

Measuring IPv6 Deployment

By Geoff Huston and George Michaelson, APNIC, June 2009

Abstract

This paper examines how best to measure IPv6 deployment in the Internet. A number of potential metrics relate to the IPv6 capability of various subsystems and components of the Internet and this paper looks at their plausibility as effective measures of overall IPv6 deployment. A metric of network capability for IPv6 is proposed, based on a variant of the original end-to-end IP architecture. The proposed approach is based on measurements conducted at dual-stack servers. By analysing server logs and traffic capture, it is possible to build a picture of client capability to conduct client/server transactions over IPv6. This can be compared to IPv4. As this type of measure exploits complete end-to-end communication, it is a strong indicator of *every* component of the Internet's architecture. Use of this metric is felt to be strongly indicative of near-term trends for overall Internet IPv6 support. A continuing programme of measurement for IPv6 is laid out.

Introduction

It is unfeasible to conduct a comprehensive analysis of every connected device, every network switching element, every circuit and every data packet that collectively makes up the Internet. Therefore, to generate meaningful metrics for the entire Internet it is necessary to carefully define the nature of the metrics, identify a bounded subset of the network on which to conduct the experimental observations, and, finally, to understand the broader context of the experiment across the Internet as a whole.

Interest in relative IPv4/IPv6 metrics has been prompted by the prospect of depletion of the remaining pools of unallocated IPv4 addresses in the coming two to three years.[1] A related activity is the tracking of the level of IPv6 deployment across the Internet. Clearly, in the context of an exhausted supply of IPv4, continued growth of the Internet demands deployment of IPv6 or some other technology. Can this use of IPv6 be measured and predicted?

There have been a number of exercises in recent months that have looked at various parts of the IPv6 deployment issue. One of the more prominent exercises has been a "lights out for IPv4" trial at a number of Internet standards and operations meetings. In these trials, the IPv4 Internet was switched off for a period, leaving meeting attendees with an IPv6 only service.[2] These exercises

---

[1] For more information on IPv4 exhaustion predictions, see:

- IPv4 Address Report
  http://www.potaroo.net/tools/ipv4/

[2] Recent experiments include:

- Big IPv6 Switch (RIPE 56)
  http://www.ripe.net/ripe/meetings/ripe-56/report.html
- Ready, set, go IPv6! (APNIC 26)
  http://archive.apnic.net/meetings/26/program/ipv6
- IPv4 Outage (IETF 71)
  https://wiki.tools.isoc.org/IETF71_IPv4_Outage

are useful in terms of assessing some aspects of preparedness for IPv6 in terms of the capabilities of local hosts and the ability to access services via IPv6. In other ways, however, these exercises are perhaps tangential to the immediate operational objective. An IPv6-only network is the end-point of a potentially lengthy transition phase where both IPv4 and IPv6 will exist in a “dual-stack” mode of operation that could exist for decades. The Internet is in the early stages of this dual-stack phase: end hosts, networks, services, and middleware are in the process of shifting from IPv4 only to some form of dual-stack support. If we are already in this dual-stack world, just how much dual-stack operation is actually out there today?

No single data set can answer the question of just how much IPv6 is deployed in the global Internet, nor how much IPv6 is being used relative to IPv4. Instead, it is necessary to take snapshots of smaller data “windows” that have some bearing on this larger question. From these it is possible to make inferences that may shed light on the larger picture.

This paper examines some of the options for measuring the use of IPv4 and IPv6 in today's Internet using data that can be readily gathered at a single point. The paper also examines the utility of such measurements in the context of measuring the status of global IPv6 deployment. The three primary data sources used in this paper are the Border Gateway Protocol (BGP) inter-domain routing table[3], access logs from web services, and packet captures of Domain Name System (DNS) server queries.

Time Series Data Sets vs. Single Measurements

A series of identical measurements taken over time provides a more robust data set for analysis than a single measurement taken in isolation. To assess not only the current state of IPv6 deployment, but to also assess the likely timescale to achieve comprehensive deployment, data about the *rate of change* of the metric is as important as the *current value* of the metric. A time series of data allows various forms of trend analysis to be performed over the data set that, in turn, can be used to generate projections, based on the assumption that future trends can be inferred from historical data.

This paper, therefore, evaluates a number of time series measurements, where the observation has been taken regularly over an extended period of time, to distinguish between IPv4 and IPv6 use in various contexts.

The Routing View of IPv6

Measurement using global BGP routing tables

The first data set evaluated in this section is the Internet's global routing table. On the following page, Figure 1 shows the number of entries in the global IPv6 routing table from 1 January 2004 until March 2009. The data was sampled on an hourly basis.

---

[3] Pages 4 and 5 of *Internet Addressing – Measuring the Deployment of IPv6: Outline* (DSTI/ICCP/CISP(2009)10) provides some additional discussion on global routing tables.

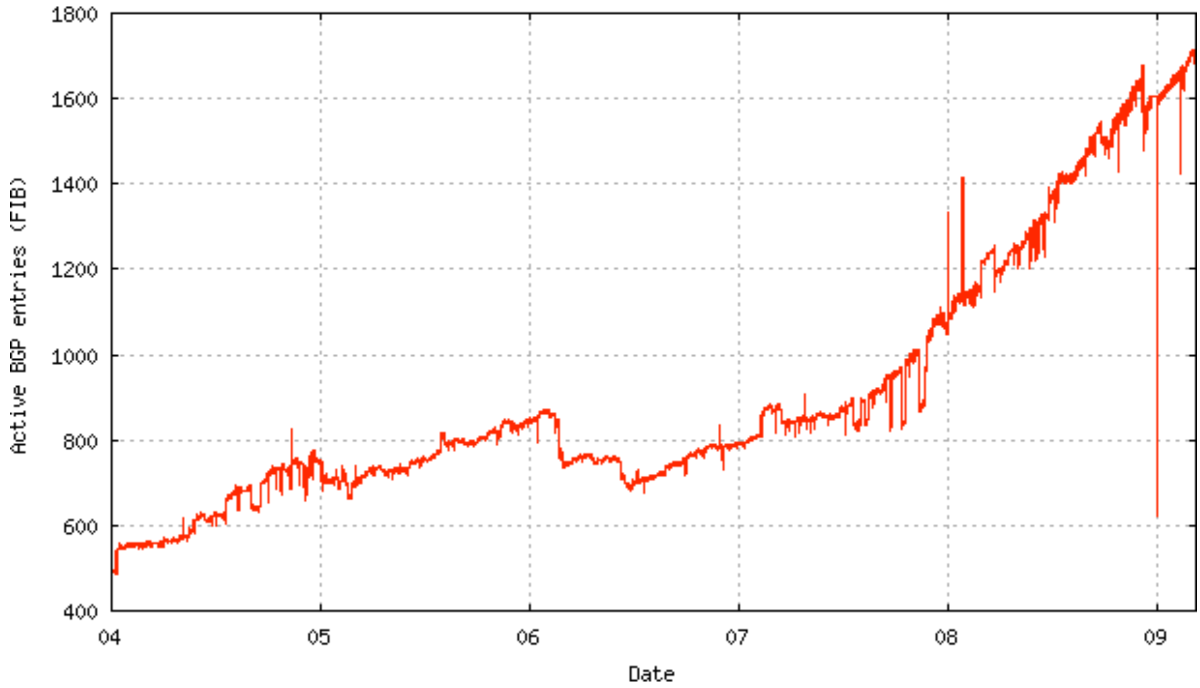

Figure 1 - IPv6 BGP Table Size

If the global routing table in Figure 1 reflects real IPv6 deployment, it shows some very volatile elements within the IPv6 routing domain over time. The graph in Figure 1 shows a period of steady growth across 2004 and 2005, then a significant downturn in the first half of 2006. The event in February 2006 was 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.[4] However, as Figure 1 shows, interest in IPv6 deployment increased beginning mid-2007 and the number of routing entries has doubled in the past 24 months. The data series in Figure 1, therefore, appears to indicate a growing level of interest in IPv6 in more recent times.

In comparison, Figure 2 on the next page shows the view of the IPv4 global 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 of data series that we have been used to seeing with the Internet, indicating a consistent strong level of underlying growth in the network.

[4] For more information on the 6Bone, see:

http://go6.net/ipv6-6bone/

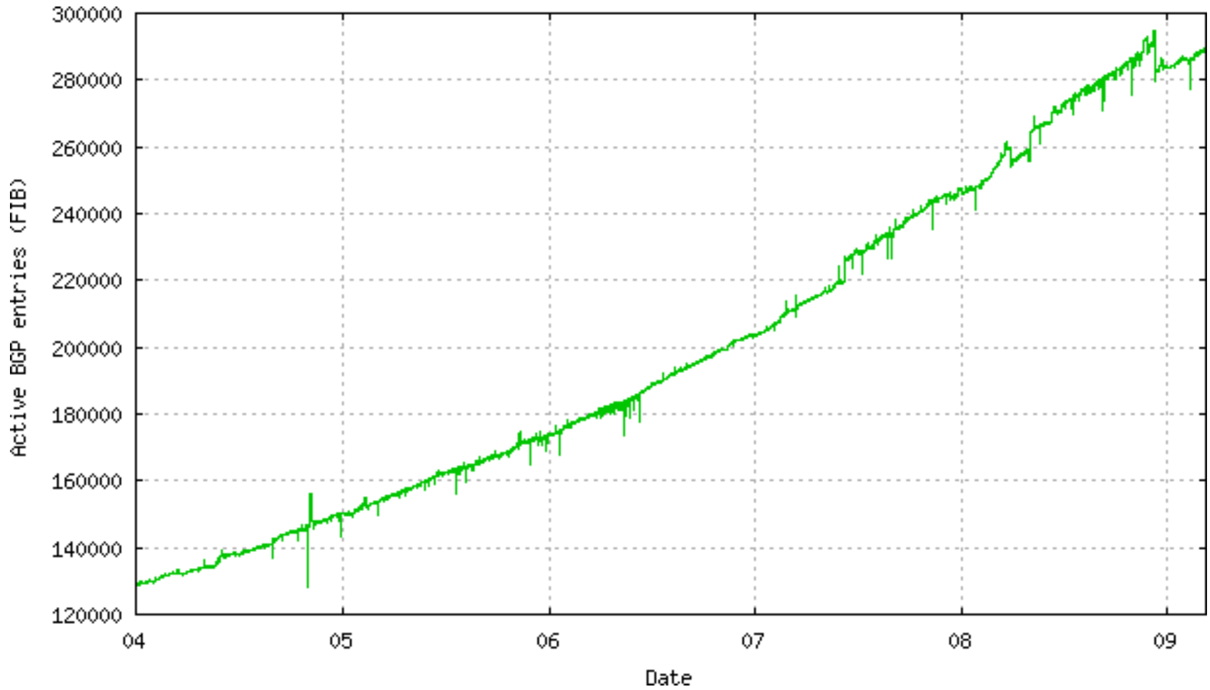

Figure 2 - IPv4 BGP Table Size

In both the IPv4 and IPv6 cases, we have 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. However, there is a critical point of difference. 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.

Next, this paper compares the above two data sets to assess the relative level of IPv6 and IPv4 deployment and to assess if IPv6 is growing at a faster relative rate than IPv4 as measured by advertised routes. Figure 3 below 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 in Figure 3 indicates the IPv6 Internet is 0.6% the size of the IPv4 Internet. The graph also shows that the IPv6 network has been growing at a faster rate, in terms of number of routing entries, than IPv4 since mid-2007. However, a growth from 0.38% to 0.58% in relative terms is not considered significant at this level. Assuming this metric of 0.22% p.a. continues indefinitely, this implies the IPv6 routing table will reach the same level as IPv4 in 452 years from now. Even with a target of an IPv6 routing table one-third the size of the IPv4 routing table, the target is still 149 years away. This extrapolation seems highly implausible and the issue may lie in the nature of the network behaviour being measured using this approach. The problem with these two data series is that they are not actually measuring the same quantity, despite appearances to the contrary. The IPv4 address space is heavily fragmented and the average number of routing table entries per origin AS (Autonomous System) is currently around 9.4 (and has been steadily rising from a low of 8 in early 2004) while the same metric in IPv6 is 1.3 entries per origin AS.

It is interesting to examine why the two measurements of advertisements ("routing table entries") per AS are so different in IPv4 and IPv6. There appear to be a number of factors at play here, including:

1. The IPv4 "legacy swamp"[5] of un-aggregated addresses is not present in IPv6;
2. The lack of provider independent address assignments in IPv6 until quite recently;
3. The lack of large scale production networking in IPv6 that, in turn, has removed any need for IPv6 traffic engineering; and,
4. The prevalent use of IPv6 as an overlay network, forcing the route policy determination role into IPv4 routing rather than IPv6 routing.[6]

The relatively low number of addresses per origin AS in IPv6 could be interpreted as effective enforcement of policies of provider-based aggregation into the IPv6 routing environment, but

[5] IANA (http://www.iana.org/assignments/ipv4-address-space/) states that legacy addresses are IP addresses "allocated by the central Internet Registry (IR) prior to the Regional Internet Registries (RIRs)".

[6] In an overlay network, the underlying "tunnel" provider does most of the actual routing.

considering the factors mentioned above, is more likely to be a sign of the current immaturity of IPv6 operational deployment and use.

This disparity in IPv4 and IPv6 routes per AS raises some interesting questions:

- Are we measuring relative deployment levels in this comparison of BGP entry counts, or measuring relative routing fragmentation?
- Does an increase in the value of the ratio imply more IPv6 being deployed or more IPv6 address fragmentation?
- Is fragmentation of the routing table a necessary component of traffic engineering or an artefact of history?
- If it were a fully deployed IPv6 network today, how large would the IPv6 routing table be without any contribution from historical address fragmentation?

As there are no clear answers to these questions it may be worthwhile finding a different routing metric to measure relative deployment of IPv6.

Measurement using Autonomous Systems (AS)

Another useful approach may be 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, it is not the number of entries in the BGP routing table per se, but the number of unique AS numbers routing IPv6 that indicates how many entities participate in the global IPv6 Internet.

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

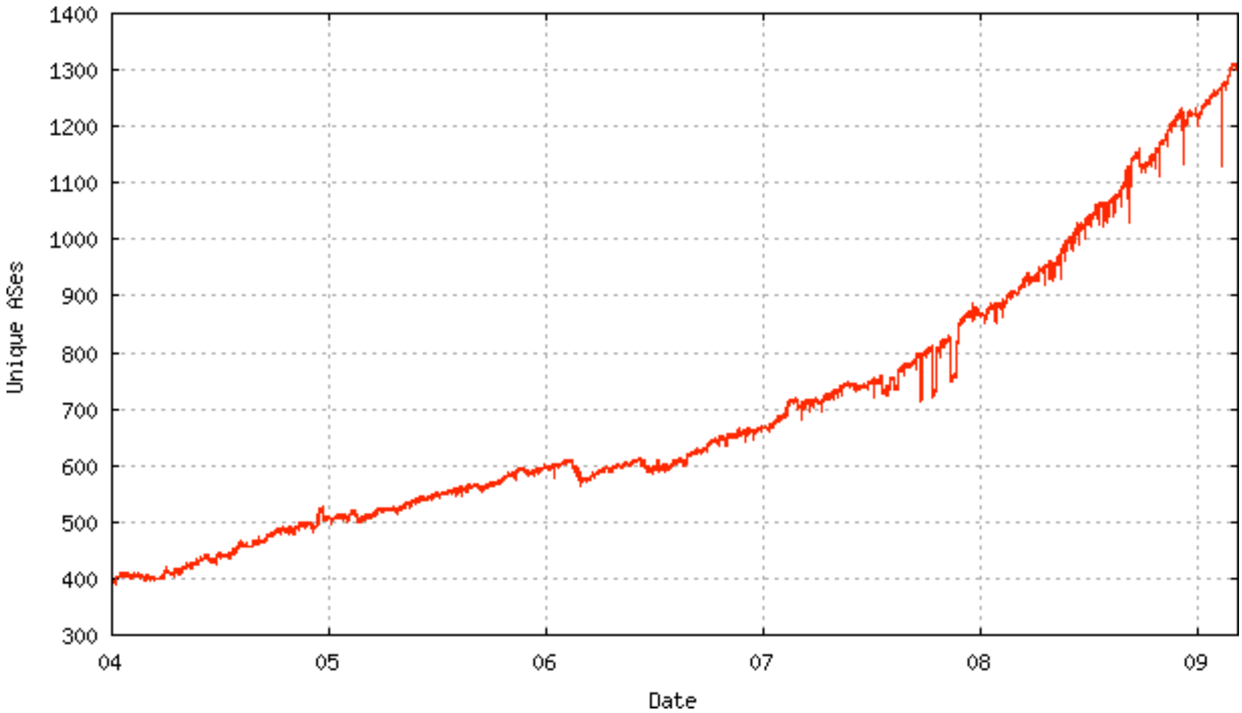

Figure 4 - IPv6 AS Count

Figure 4 shows a more even picture of IPv6 deployment than is suggested by the number of IPv6 routing table entries in Figure 1. Figure 4 shows a larger than tripling in growth in the number of these routing entities, from 400 to 1300 over this 5 year period. Also discernable in the data, is the acceleration in growth from mid-2007. This AS count data indicates a consistent and steady enrolment of new entities that are announcing routes into the IPv6 network. While there is no spectacular exponential trend in the data so far, it is clear the IPv6 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. Next, Figure 5 shows a comparable plot for the number of ASes in the IPv4 network.

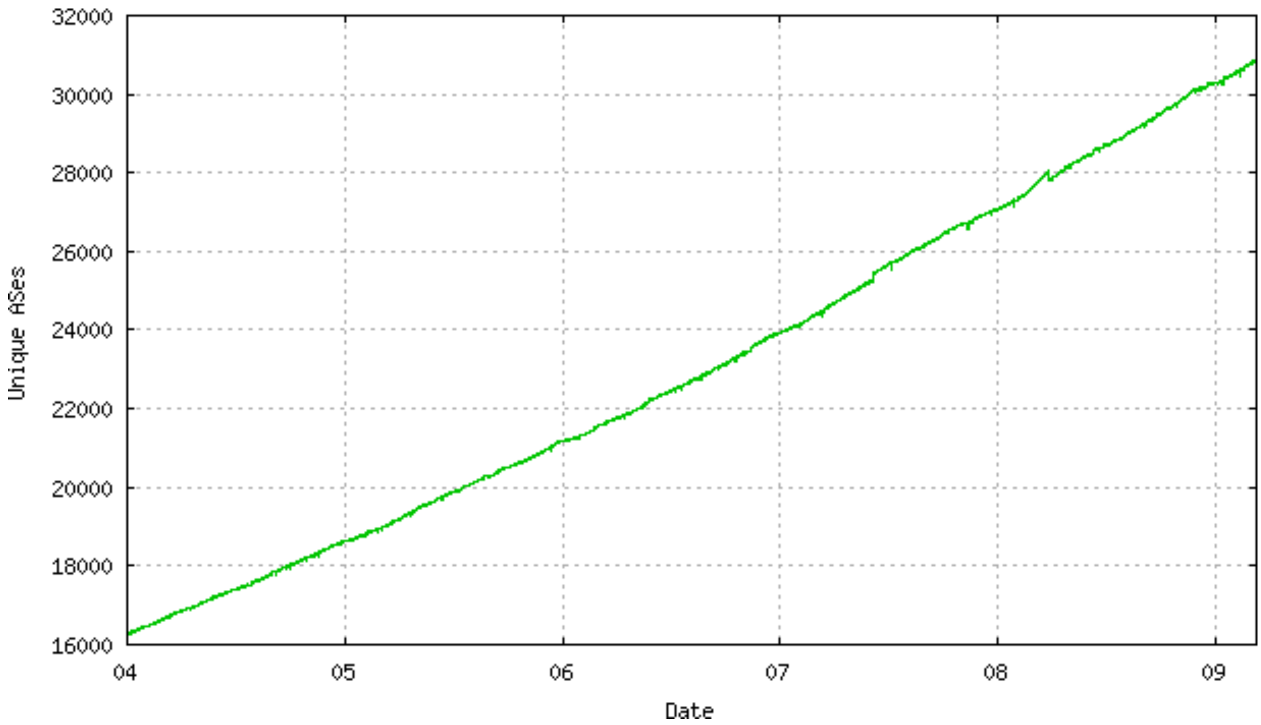

Figure 5 - IPv4 AS Count

In terms of the number of ASes, Figure 5 shows that the IPv4 network did not quite double in size over this period, increasing from 16,000 ASes to 30,500 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 that much of the churn in routing is attributable to the effects of other measures, such as traffic engineering or local efforts to optimise certain aspects of routing policies.

These two data series can be compared in the same manner as the number of entries in the routing table were compared in Figure 3. 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.

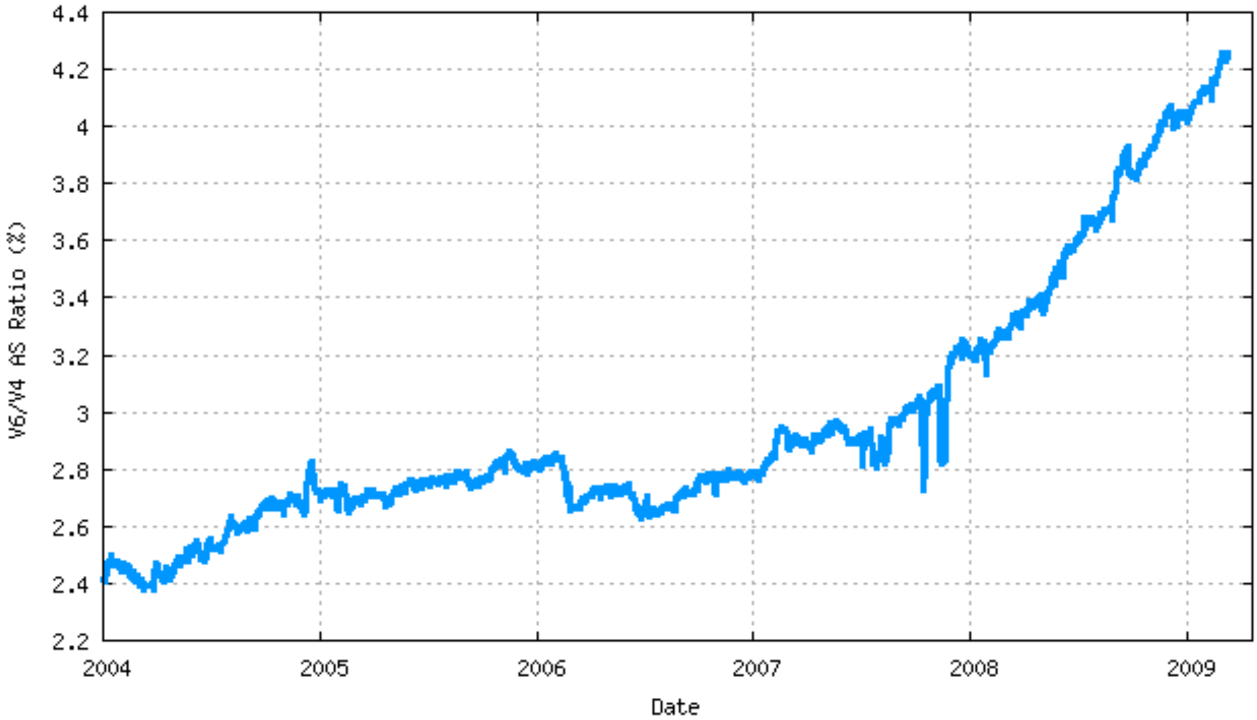

Figure 6 - IPv6 / IPv4 Relative AS Count Ratio

This metric paints a somewhat different, and more positive, picture of IPv6 deployment than the comparison of global routing table entries. Here, the relative metric of IPv6 as compared to IPv4 is 4.2%, and the number of AS entities actively routing IPv6 is growing at a faster rate than the IPv4 network. If this relative rate of 0.8% p.a. were sustained, the IPv6 AS count would equal that of IPv4 in 120 years.

From the AS origination data in the routing table, it appears that some 4% of the Internet is IPv6 capable—to one extent or another—in terms of the population of distinct entities that compose the Internet. One caveat, however, is that it is not necessarily the case that an AS advertising an IPv6 route has its entire network dual stack provisioned with IPv6, nor that all of its end hosts and customers are dual stack provisioned. Rather, this 4% metric of IPv6 capability at the AS level is perhaps a reflection of levels of experimental and research interest in IPv6, rather than a true picture of IPv6 as a deployed network service platform.

Refining the methodology

These two forms of routing measurements may be taken at any point in the BGP routing fabric. The measurement technique is relatively simple and there are a number of data archives that track this data back over many years. However, this class of metrics measure aspects of IPv6 support within a single class of network components. While each component has to support IPv6, such component-based measurements are not overly illustrative of the capability of the network to support IPv6 at the application level. In particular, there are two potential issues in the routing table view of the IPv6 Internet:

1. 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.
2. 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. The next section will examine this measurement option.

A Usage View of IPv6

The class of questions that a usage-oriented view of IPv6 could possibly answer include:

- How much is IPv6 being used today relative to IPv4?
- Has this metric changed in recent years?
- How much IPv6 use is via the transitional tools of 6to4 and Teredo?

There are many ways of attempting to answer these questions, including gathering long-term traffic sampling data from an operational network, through to a more controlled experiment using a sample at a service point. It must be noted, however, that there are a number of issues with traffic sampling of a commercial and legal nature that limit the extent to which traffic sample data sets are made available to the research community. In addition, there are considerations that impact on the appropriate interpretation of such data. For example, when measuring traffic by total volume, the “heavy tail” distribution of traffic flows comes into play where a small class of flows are significant contributors to the total traffic volume. Also relevant is the extent to which IPv6 still uses IPv6-in-IPv4 tunnelling approaches, effectively "hiding" IPv6 packet headers from the outer IP header. Therefore, this paper uses data gathering approaches that are more accessible for study and whose interpretation is a little more straightforward, while leaving the usage view of IPv6 for a future study.

The DNS view of IPv6

Another long-term data set available for examination is usage data from a number of DNS servers[7]. It is important at this point to distinguish between the *configuration* view and the *query and response* view. The configuration view searches the DNS zone files and counts the number of AAAA records that are configured into the DNS. However, while such configuration elements are a necessary precursor to the use of IPv6 for service access, in isolation they are not a useful metric about the extent of deployment in terms of usage of IPv6.

The DNS is, however, also a source of use-related data. DNS is itself a protocol. Clients send queries and servers issue responses. The DNS servers we are using for this particular 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. Of interest here is the relative rate of queries that are made to the *in-addr.arpa* zone, which relate to resolution of IPv4 addresses and the queries that are made to the *ip6.arpa* zone, which relate to the resolution of IPv6 addresses. The assumption behind this metric is that a client may perform a reverse IPv6 DNS lookup in response to a traffic event originated by that IPv6 source address and is unlikely to perform such a lookup under other circumstances. Therefore, the total lookup rate is likely to be related to the number of network transactions that have occurred using IPv6. In addition, the comparison of reverse lookups of IPv4 address to IPv6 addresses is related to the total use of IPv6 for network transactions in relation to the use of IPv4.

The relative proportion of these two query types is shown in Figure 7, with a further breakdown to IPv6 addresses that have been allocated to entities operating in the Asia Pacific region and other parts of the world.

---

[7] Pages 7 and 8 of *Internet Addressing – Measuring the Deployment of IPv6: Outline* (DSTI/ICCP/CISP(2009)10) provides some additional discussion on using DNS to measure IPv6 deployment.

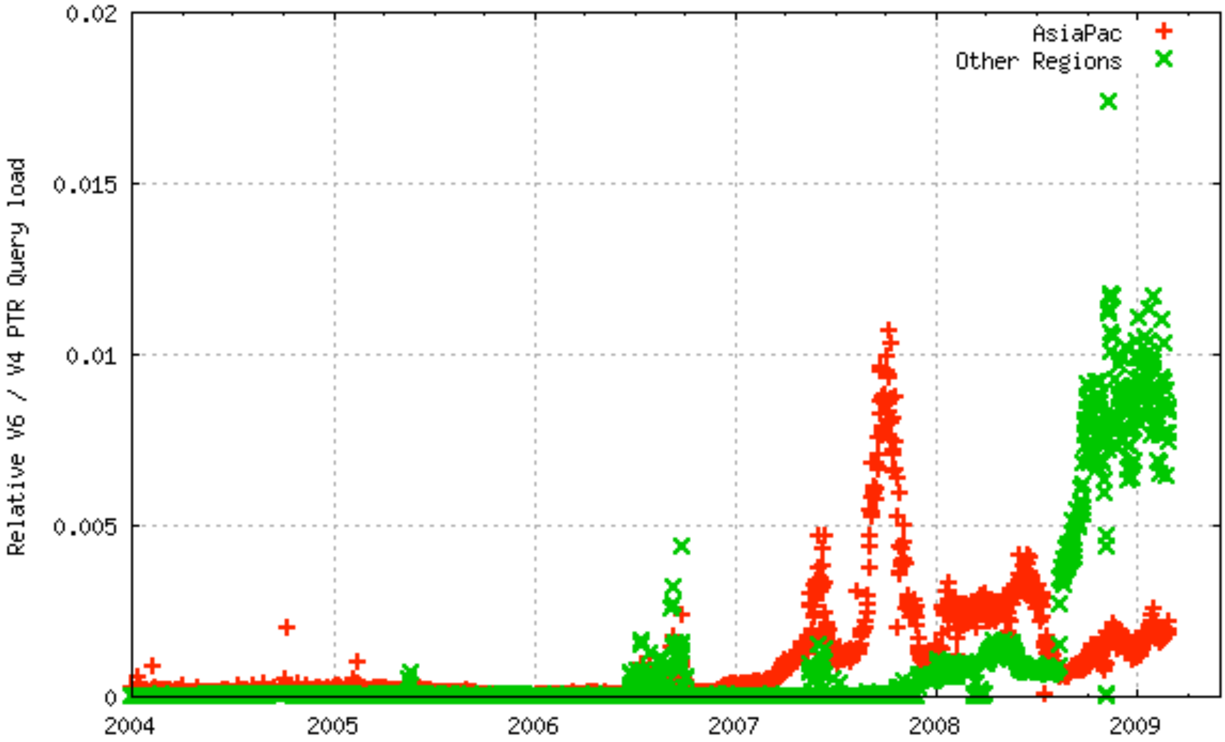

Figure 7 - DNS V6 / V4 reverse PTR query Ratio

The relative levels of use of IPv4 and IPv6 in Figure 7 appears to be quite different from the BGP that we have reviewed so far and web server log data to be discussed later in this paper. The differences between the Asia Pacific region and the rest of the world servers is also apparent in Figure 7, particularly in relation to the sharp peaks in 2007 and the upward sudden movement from below 0.01% to 0.075%, highlight some of the issues in the appropriate interpretation of this data series. The 2007 data appears to be the outcome of a number of multi-day “turn off IPv4” experiments or other events that in the first case spanned a week and impacted the entire DNS server set and in the second case spanned more than a month and was only visible to the Asia Pacific reverse DNS servers. The data for the other regions in late 2008 is similarly anomalous. That a single activity such as these can swamp the long-term trends may be indicative of the current small size of the IPv6 usage in DNS.

It is not entirely clear how to interpret this DNS query data. First, there is the issue of identifying the class of applications that perform these forms of reverse DNS lookups. Second, it is important to understand the relationship between original queries for end hosts and DNS caches that may lie between the end host and the authoritative name server. The third issue in interpreting DNS query data is the relatively low volume of IPv6 queries. This, in turn, makes the data susceptible to bias resulting from individual actions in querying the DNS. At this point in time, therefore, as an indicator to the relative uptake of IPv6 over an extended period, it appears that DNS server data set has a number of unresolved issues before appropriate interpretation can be made.

Web Server Logs

Another approach to measuring the deployment of IPv6 is to measure IPv6 use from the perspective of a dual-stack server.

Application space measurement of tunnelled and native IPv6

The server will record an IPv6 transaction only if all of the following conditions are met:

1. The client has an IPv6 stack;
2. The client's application is configured with IPv6 support;
3. The client's DNS configuration is able to perform an IPv6 address query; and,
4. The client and server can communicate end-to-end using IPv6.

In other words, this measurement will only succeed if all the intermediate components of the connection are configured to support IPv6. Therefore, this metric would be a good indicator of the total level of IPv6 deployment capability across all components of the network.

The data set used in this paper relates to the use of the APNIC web site, www.apnic.net, and the RIPE web site, www.ripe.net. These web sites have both IPv4 and IPv6 addresses and have been dual homed on both IPv4 and IPv6 networks for over five years. The approach used to measure the relative use of IPv6 to IPv4 was to count the number of unique source addresses visiting these web sites each day and to look at the ratio of the number of unique IPv4 source addresses to the number of unique IPv6 source addresses. This approach was used to remove the factors associated with robots and web crawlers (which for these sites are evidently still exclusively using IPv4) and to even out some of the factors of the level of intensity of access and repeat visits to the same site.

Figure 8 below shows the daily ratio of IPv6 to IPv4 source addresses that have accessed the APNIC and RIPE web sites since 1 January 2004. There is considerable variation in the data from day to day. Therefore, a scatter plot has been used in Figure 8 to ensure that the trends in the data are visible as well as the day-to-day variation. It is possible that the noise component of day-to-day variation could be lowered by gaining access to the web logs of a dual-stack, dual-protocol-homed web site with considerably greater volume levels.

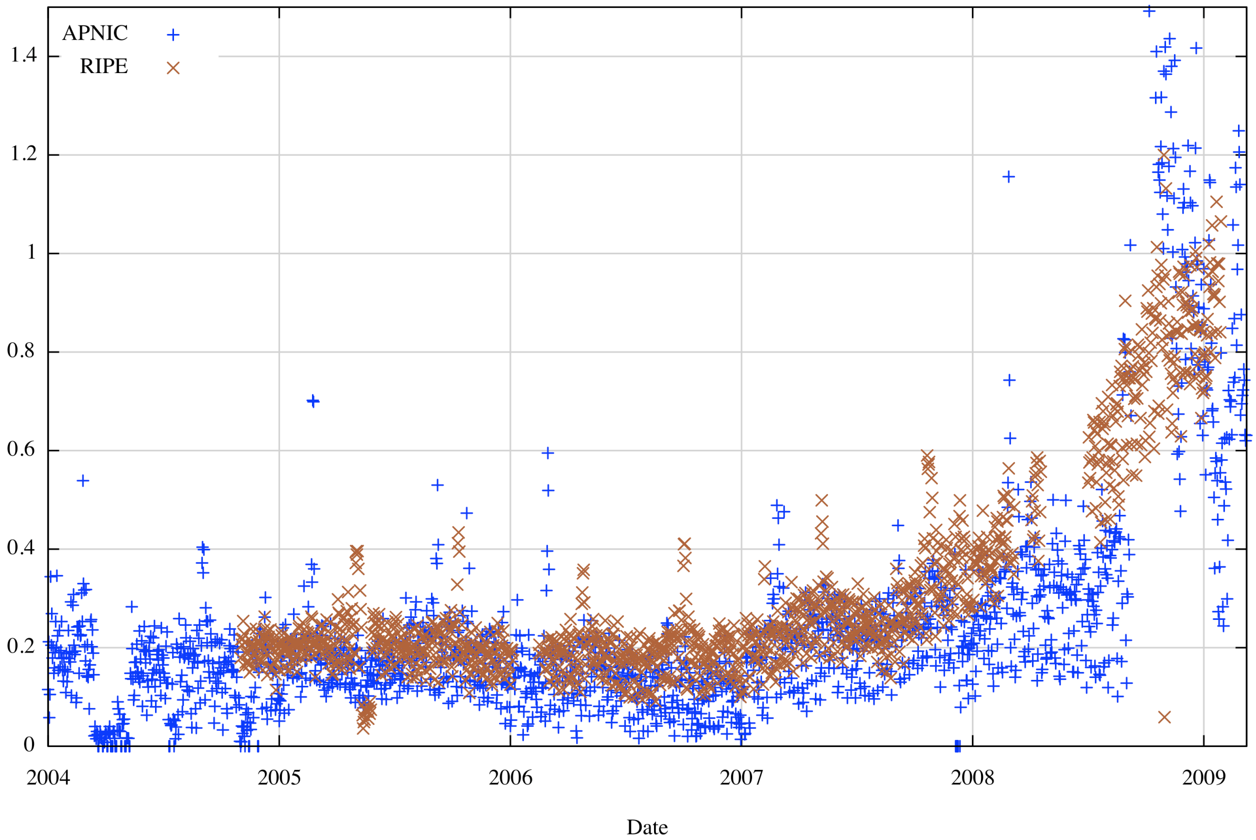

Figure 8 - IPv6 / IPv4 Web Access Ratio

In Figure 8 the anomalous data points occurring twice each year appear to coincide with APNIC and RIPE member meetings, where IPv6 was provided as part of the network infrastructure. While the number of sample points for the APNIC and RIPE websites is low, this spike during the meetings could indicate that the population of IPv6-capable end hosts is generally higher than the number of active IPv6 end hosts. IPv6-capable end hosts are not visible if IPv6 and 6in4 tunnelling is not allowed by access networks and end site configurations. However, when end hosts are in an access environment than includes IPv6 access—such as the APNIC and RIPE meeting networks—the level of relative IPv6 access to the web sites appears to increase notably.

Measurement of IPv6 tunnelling

Web service access data also can show the relative use of the 6to4 and Teredo transition techniques within the total IPv6 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.[8] Figure 9 shows the relative use of 6to4 IPv6 addresses to access the APNIC and RIPE web sites.

---

[8] 6to4 relies on access to a public IPv4 address and does not allow transition across IPv4 Network Address Translators (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 and the relative order of use of protocols stacks is to attempt a connection using IPv6 in native mode, then IPv4, then IPv6 using Teredo, unless the applications specifically initiates a connection using the local Teredo interface. Thus, for Vista, Teredo is invoked only in the event of failure of IPv4 connectivity, so that a dual-stack server would not trigger a Vista host to use Teredo.

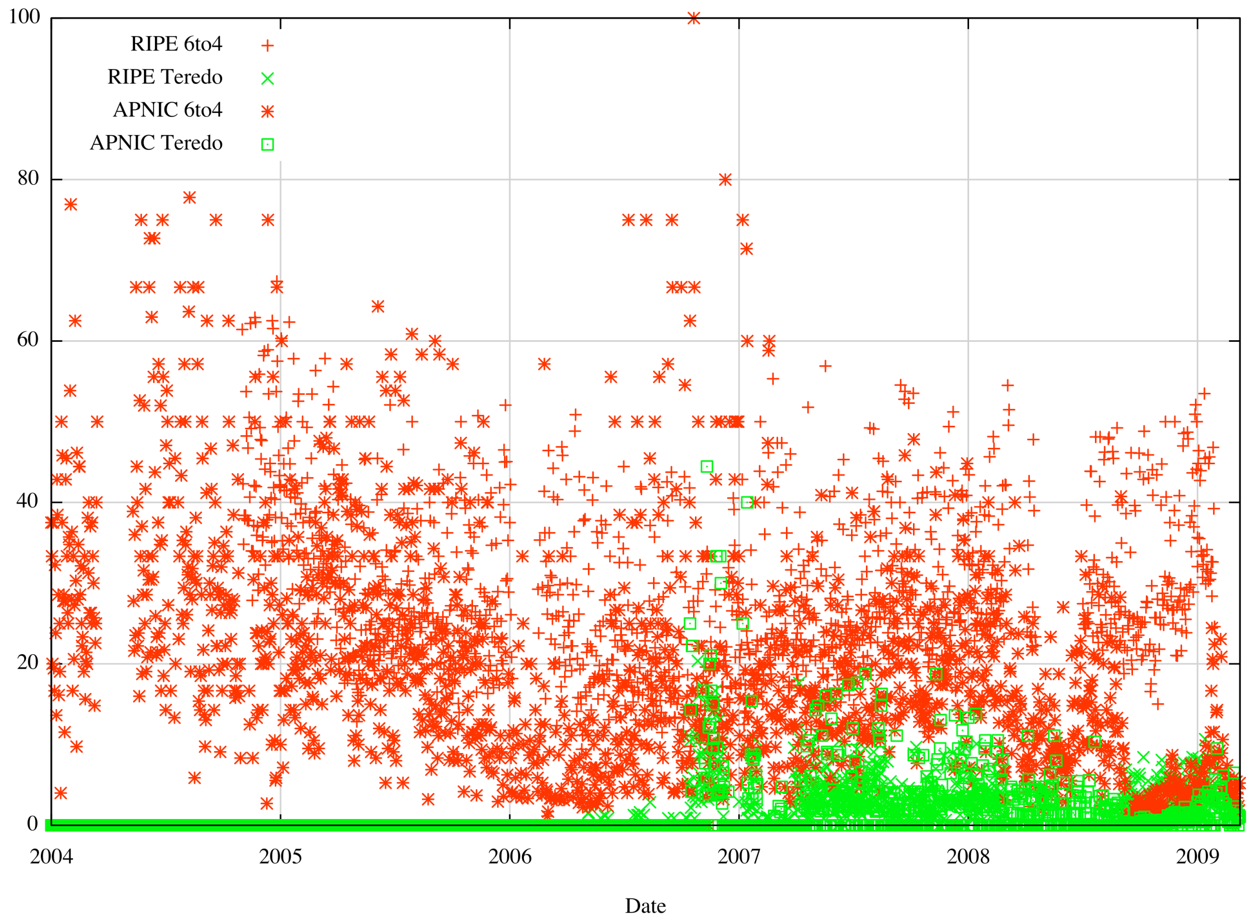

Figure 9 - 6to4 / IPv6 Web Access Ratio

Figure 9 shows that while 6to4 was in quite widespread use in 2004 and 2005, its relative use appears to be declining. Currently, the use of 6to4 appears to be steady at around an average of 20% of all IPv6 usage each day, with quite significant day-to-day variation. 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 APNIC and RIPE web sites. Additional data from web sites with higher volumes of access may provide some further clarity on the extent of IPv6 tunnelling in use.

Figure 9 also shows the relative use of Teredo IPv6 addresses to access the web sites. It shows that Teredo deployment commenced in late 2006 and its relative use is still around 2% or so of the total IPv6 use.

Observations of the web-based methodology

The web server access data reported here represents an extremely small peephole into a significantly larger Internet. If this data were aggregated with data from other sites with larger traffic levels, the picture of IPv6 use may change to some extent. In addition, it must be noted that the APNIC and RIPE sites are oriented towards technologically adept users; more mainstream dual-stack sites may see lower relative numbers for IPv6 access. It is also noted that NATs are widely deployed in IPv4 and not in IPv6, so there is a certain component of under-counting of the host count in IPv4 that also has some level of impact on this metric.

Conclusions

The Internet is facing some quite fascinating pressures in the coming years as the unallocated pool of IPv4 addresses depletes. It is unclear at this stage just how quickly the Internet will transition to an IPv6 network and how such a transition will be deployed in the network. It is also unclear to what extent the Internet will be able to wean itself off the intensive use of NATs. It is unclear what the relative pressures are as networks decide whether to make the transition to IPv6 or to persist in using private IPv4 address space, NATs and various forms of protocol translation to fill the connectivity gaps.

Much of the IPv6 technology set could be described as operationally ready. There is clear evidence that IPv6 hosts and service delivery platforms are being deployed. There is also good evidence that a visible proportion of the organisations that manage the infrastructure of the Internet are undertaking various forms of IPv6 deployments. However, the real level of uptake of IPv6 in the Internet today, in terms of service access, remains very small. The most reliable metric of the current level of end user IPv6 uptake is the web server access data and the observed level of the relative rate of IPv6 use appears to be around 0.9% of the IPv4 use, or a relative level of 9 parts of IPv6 per 1,000 of IPv4.

A more 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.

Global adoption of IPv6 to satisfy foreseeable demand for Internet deployment would require a significant increase in its relative use, in a short space of time. By the measurements explored here, this cannot yet be demonstrated. In particular, IPv6 is not measured as being deployed sufficiently rapidly at present, to offer an “intercept” to the predicated IPv4 exhaustion date. Should a change to the dynamics of deployment be possible, it is believed this methodology can demonstrate such change and stands as a metric. The sensitivity of current deployment to experimental use of IPv6 is noted.

Further Work

There are many potential windows for collecting data on IPv6 adoption across the Internet. This paper has investigated but a few of the options available for measuring IPv6 deployment. The authors of this paper are very interested in learning of other long-term data sets that could be used for relative metrics of IPv6 and IPv4 protocol use in the Internet.

Further work in the BGP routing 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.

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. However, such analysis would require 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, query data from DNS forwarders may be useful in this context.

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 operating system, 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 the authors of this paper 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 observed for the APNIC and RIPE web sites is also visible for other web sites and other services.

This paper has not reported on any study of actual data rates in an operational network. It would be interesting to unders tand the relative ratio of IPv4 and IPv6 traffic, by payload volume, by packet count, and by port addresses, as well as the relative amount of traffic in the 6in4 tunnels, on operational networks today. So far, the authors of this paper have not been able to locate open sources for such data that have a long baseline of historical data. If there are any offers of such operational data, the authors would be interested to examine it to see how it correlates to these other measurements that have been reviewed here.

Trends in IPv6 Deployment 2005 – 2008

This document shows trends in IPv6 deployment over the last five years, per economy.

Information to assist in interpreting the table on the following pages:
To make best use of the data, the NRO recommends that you read the notes below in conjunction with the data.

Economy	ISO 3166-1 currently recognizes 246 different economies. The RIRs use this list for registering allocations and assignments in their whois databases. In addition to the officially recognized codes, the RIRs also use 2 regional designations: Asia Pacific and European Union. The data in this table shows the economy to which IPv4 addresses were originally delegated. It does not show any subsequent change to the economy of IP address delegations that may occur after original delegations have been made. For example, while Serbia and Montenegro now exist as separate economies, the earlier designation as a single economy still appears in the table. Delegations made after Serbia and Montenegro were split into two economies will appear as separate entries under Serbia and, separately, Montenegro. Another implication of this preservation of original delegation information means that some economies appear to have no IP addresses when they in fact do. For instance, in the case of Martinique, the table below shows the economy as having no IPv4 addresses delegated directly by an RIR; however, Martinique does have address space assigned to its ISPs from RIR-delegated address ranges allocated to ISPs in the Netherlands Antilles and France. The addresses in this case, are therefore recorded in the table below as being delegated to the Netherlands Antilles and France.
Total address holdings (measured in units of /64)	Because of the very large scale of numbers allocated in the IPv6 address space, we have not included absolute numbers of individual IP addresses, but have used /64 (or 264 addresses) notation. A single /64 is used to address a single IPv6 subnet where auto-configuration of host addresses is desired. The table also includes a column to show the number of addresses using the standard Internet technical equivalent of a total number of /32s (which represent 296 IP addresses, the standard minimum address block allocated to a network). It is important to note that, due to the hierarchical nature of addressing, not all IP addresses that are delegated to a network can, in fact be used. This is particularly relevant to networks using IPv6, as the protocol's addressing scheme is designed to facilitate more efficient network management and routing rather than efficient use of individual addresses. For a more in depth overview of IPv6, please see: http://www.ripe.net/info/faq/rs/ipv6.html For more information about IPv6 addressing and 'slash notation" please see: http://www.ripe.net/info/info-services/addressing.html
% increase 2007-2008	For economies that have experienced a change from no IPv6 addresses in 2007 to having IPv6 addresses in 2008, it is not possible to calculate a percentage increase. In such cases, the table below shows "NA" (not applicable).

Trends in IPv6 allocations, 2005-2008					
Economy	Total address holdings (measured in units of /64)				% increase 2007-2008
	31 Dec 2005	31 Dec 2006	31 Dec 2007	31 Dec 2008	
Algeria	0	0	0	4,294,967,296	NA
Andorra	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Angola	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Argentina	120,259,084,288	128,849,018,880	146,028,953,600	163,208,822,784	11.76%
Armenia	8,589,934,592	8,589,934,592	8,589,934,592	12,884,901,888	50.00%
Australia	17,635,672,981,504	17,635,673,047,040	35,257,923,862,528	35,361,003,405,312	0.29%
Austria	84,825,669,632	97,710,571,520	110,595,473,408	166,430,048,256	50.49%
Bahamas	0	0	4,294,967,296	4,294,967,296	0.00%
Bahrain	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Bangladesh	0	4,294,967,296	17,179,869,184	25,769,869,312	50.00%
Belgium	23,085,580,288	35,970,482,176	53,150,351,360	83,215,253,504	56.57%
Benin	0	0	4,294,967,296	4,294,967,296	0.00%
Bermuda	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Bhutan	0	0	0	4,294,967,296	NA
Bolivia	4,294,967,296	8,589,934,592	8,589,934,592	8,589,934,592	0.00%
Bosnia & Herzegowina	0	0	4,294,967,296	8,589,934,592	100.00%
Brazil	0	0	0	282,299,610,431,488	NA
Bulgaria	12,884,967,424	12,884,967,424	21,474,902,016	42,949,738,496	100.00%
Cameroon	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Canada	81,604,378,624	98,785,034,240	154,621,247,488	201,867,526,144	30.56%
Chile	12,884,901,888	12,884,901,888	30,064,771,072	38,654,771,200	28.57%
China	73,014,509,568	111,669,280,768	133,144,117,248	244,813,266,944	83.87%
Colombia	0	8,589,934,592	21,474,902,016	34,359,803,904	60.00%
Costa Rica	4,294,967,296	4,294,967,296	8,589,934,592	25,769,803,776	200.00%
Cote D'ivoire	0	0	4,294,967,296	4,294,967,296	0.00%
Croatia (Hrvatska)	8,590,000,128	8,590,000,128	8,590,000,128	17,179,934,720	100.00%
Cuba	12,884,901,888	12,884,901,888	12,884,901,888	17,179,869,184	33.33%
Cyprus	8,589,934,592	8,589,934,592	8,589,934,592	17,179,869,184	100.00%
Czech Republic	53,150,285,824	61,740,285,952	91,805,057,024	169,114,468,352	84.21%
Czechoslovakia	12,884,901,888	12,884,901,888	12,884,901,888	17,179,869,184	33.33%
Denmark	27,380,482,048	27,380,482,048	48,855,318,528	87,510,024,192	79.12%
Djibouti	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Dominican Republic	12,884,901,888	12,884,901,888	12,884,901,888	12,884,901,888	0.00%
Ecuador	0	4,294,967,296	8,589,934,592	12,884,967,424	50.00%
Egypt	4,294,967,296	12,884,901,888	17,179,869,184	21,474,836,480	25.00%
El Salvador	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Estonia	18,790,547,456	31,675,514,880	31,675,514,880	35,970,482,176	13.56%
Fiji	0	0	4,295,032,832	4,295,032,832	0.00%
Finland	49,392,254,976	53,687,222,272	53,687,222,272	66,572,124,160	24.00%
France	35,300,873,273,344	35,335,233,011,712	35,369,592,750,080	35,721,780,199,424	1.00%
Germany	39,921,721,344,000	41,068,477,743,104	41,446,434,996,224	41,742,787,805,184	0.72%
Ghana	0	0	4,294,967,296	4,294,967,296	0.00%
Gibraltar	0	0	0	4,294,967,296	NA
Greece	1,610,612,736	1,610,612,736	10,200,547,328	18,790,547,456	84.21%

Trends in IPv6 allocations, 2005-2008					
Economy	Total address holdings (measured in units of /64)				% increase 2007-2008
	31 Dec 2005	31 Dec 2006	31 Dec 2007	31 Dec 2008	
Guatemala	4,294,967,296	8,589,934,592	17,179,869,184	17,179,869,184	0.00%
Haiti	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Holy See (Vatican City State)	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Hong Kong	21,474,836,481	30,064,771,073	38,654,705,665	64,424,574,977	66.67%
Hungary	10,200,612,864	18,790,547,456	27,380,482,048	53,150,285,824	94.12%
Iceland	4,294,967,296	4,294,967,296	12,884,901,888	25,769,803,776	100.00%
India	47,244,640,256	60,129,542,144	60,129,607,680	85,899,411,456	42.86%
Indonesia	25,769,934,848	51,539,738,624	68,719,607,808	85,899,542,528	25.00%
Iran (Islamic Republic Of)	25,769,803,776	25,769,803,776	34,359,738,368	47,244,640,256	37.50%
Ireland	24,696,127,488	54,760,898,560	71,940,767,744	97,710,571,520	35.82%
Isle of Man	0	0	4,294,967,296	4,294,967,296	0.00%
Israel	17,179,869,184	17,179,869,184	17,179,934,720	21,474,902,016	25.00%
Italy	120,796,151,808	17,717,277,163,520	17,743,046,967,296	17,837,536,247,808	0.53%
Japan	31,204,548,083,714	31,217,432,985,602	35,636,954,398,722	35,735,738,974,210	0.28%
Kenya	0	4,294,967,296	4,295,032,832	17,180,000,256	300.00%
Korea, Republic Of	17,802,639,507,457	22,269,405,495,297	22,295,175,299,073	22,325,240,070,145	0.13%
Kyrgyzstan	0	0	0	4,294,967,296	NA
Latvia	4,294,967,296	4,294,967,296	12,884,901,888	21,474,836,480	66.67%
Liechtenstein	0	0	0	8,589,934,592	NA
Lithuania	5,905,580,032	5,905,580,032	5,905,580,032	10,200,612,864	72.73%
Luxembourg	18,790,547,456	18,790,547,456	27,380,482,048	48,855,318,528	78.43%
Macau	8,589,934,592	8,589,934,592	8,589,934,592	8,589,934,592	0.00%
Macedonia, The Former Yugoslav Republic Of	0	0	4,294,967,296	4,294,967,296	0.00%
Madagascar	0	0	0	4,294,967,296	NA
Malawi	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Malaysia	34,359,738,368	47,244,640,256	60,129,542,144	73,014,509,568	21.43%
Mali	0	0	4,294,967,296	4,294,967,296	0.00%
Malta	4,294,967,296	12,884,901,888	12,884,901,888	12,884,901,888	0.00%
Mauritius	0	8,589,934,592	8,589,934,592	17,179,869,184	100.00%
Mexico	38,654,705,664	47,244,640,256	47,244,640,256	55,834,574,848	18.18%
Moldova, Republic Of	0	0	4,294,967,296	4,294,967,296	0.00%
Monaco	4,294,967,296	4,294,967,296	4,294,967,296	8,589,934,592	100.00%
Morocco	4,294,967,296	8,589,934,592	8,589,934,592	8,589,934,592	0.00%
Mozambique	0	0	0	4,294,967,296	NA
Namibia	0	0	0	65,536	NA
Nepal	0	0	0	4,294,967,296	NA
Netherlands	2,385,854,595,072	2,394,444,529,664	2,437,394,202,624	2,583,423,156,224	5.99%
Netherlands Antilles	0	4,294,967,296	4,294,967,296	12,885,098,496	200.00%
New Zealand	34,359,869,440	55,834,705,920	73,014,837,248	124,554,575,872	70.59%
Nicaragua	4,294,967,296	4,294,967,296	4,294,967,296	8,589,934,592	100.00%
Nigeria	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Norway	1,139,777,011,712	1,144,071,979,008	1,165,546,815,488	1,234,266,357,760	5.90%
Oman	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Pakistan	4,294,967,296	17,179,869,184	21,474,836,480	30,064,771,072	40.00%

Trends in IPv6 allocations, 2005-2008					
	Total address holdings (measured in units of /64)				
Economy	31 Dec 2005	31 Dec 2006	31 Dec 2007	31 Dec 2008	% increase 2007-2008
Palau	0	0	0	4,294,967,296	NA
Panama	4,294,967,296	12,884,901,888	12,884,901,888	12,884,901,888	0.00%
Papua New Guinea	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Paraguay	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Peru	21,474,836,480	21,474,836,480	25,769,803,776	30,064,771,072	16.67%
Philippines	21,474,836,480	30,064,771,072	34,359,738,368	51,539,607,552	50.00%
Poland	84,825,604,096	8,889,508,560,896	8,966,817,972,224	9,001,177,776,128	0.38%
Portugal	35,970,482,176	44,560,416,768	44,560,416,768	61,740,285,952	38.55%
Puerto Rico	0	65,536	65,536	65,536	0.00%
Qatar	4,294,967,296	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Romania	21,474,836,480	21,474,836,480	21,474,836,480	34,359,738,368	60.00%
Russian Federation	24,696,127,488	33,286,062,080	80,530,702,336	222,264,623,104	176.00%
Rwanda	0	0	4,294,967,296	4,294,967,296	0.00%
Saudi Arabia	8,589,934,592	8,589,934,592	8,589,934,592	8,589,934,592	0.00%
Senegal	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Serbia	0	0	0	4,294,967,296	NA
Seychelles	0	0	4,294,967,296	4,294,967,296	0.00%
Singapore	25,769,803,776	25,769,803,776	30,064,771,072	64,424,640,512	114.29%
Slovakia (Slovak Republic)	8,590,000,128	17,179,934,720	25,769,934,848	34,359,869,440	33.33%
Slovenia	12,884,901,888	17,179,869,184	17,179,869,184	38,654,771,200	125.00%
South Africa	8,589,934,592	30,064,771,072	55,834,836,992	64,425,361,408	15.39%
Spain	61,740,285,952	83,215,122,432	83,215,122,432	108,984,926,208	30.97%
Sri Lanka	0	0	4,294,967,296	8,590,000,128	100.00%
Sudan	0	4,294,967,296	4,294,967,296	4,294,967,296	0.00%
Swaziland	0	0	0	4,294,967,296	NA
Sweden	71,940,898,816	84,825,931,776	114,890,768,384	729,071,157,248	534.58%
Switzerland	233,539,043,328	255,013,945,344	280,783,749,120	418,222,702,592	48.95%
Taiwan	485,331,369,984	9,637,906,677,760	9,917,079,552,000	9,917,079,552,000	0.00%
Tanzania, United Republic Of	0	8,589,934,592	8,590,000,128	12,884,967,424	50.00%
Thailand	30,064,771,072	38,654,705,664	55,834,640,384	68,719,607,808	23.08%
Trinidad & Tobago	0	4,294,967,296	8,589,934,592	8,589,934,592	0.00%
Tunisia	0	0	0	4,294,967,296	NA
Turkey	12,884,901,888	17,179,869,184	30,064,771,072	73,014,444,032	142.86%
Ukraine	0	4,294,967,296	8,589,934,592	51,539,607,552	500.00%
United Arab Emirates	4,294,967,296	4,294,967,296	4,294,967,296	12,884,901,888	200.00%
United Kingdom	362,925,326,337	423,054,868,481	4,932,770,660,353	5,091,684,450,305	3.22%
United States	627,068,370,944	837,523,734,528	1,378,733,654,016	63,561,341,272,064	4510.12%
Uruguay	8,589,934,592	8,589,934,592	77,309,411,328	81,604,640,768	5.56%
Venezuela	17,179,869,184	25,769,803,776	55,834,574,848	60,129,542,144	7.69%
Viet Nam	8,590,000,128	8,590,000,128	17,179,934,720	42,950,787,072	150.01%
Zimbabwe	0	0	4,294,967,296	4,294,967,296	0.00%

The Number Resource Organization (NRO)

Formed by the Regional Internet Registries (RIRs) to formalise their cooperative efforts, the NRO exists to protect the unallocated Number Resource pool, to promote and protect the bottom-up policy development process, and to act as a focal point for Internet community input into the RIR system.