IPv6 Special

Your Guide to the Latest IPv6 Developments in RIPE, the RIPE NCC and Beyond

RIPE NCC: Taking The Lead On IPv6

The global adoption of IPv6 is one of the key challenges in the Internet industry today, and how successfully we meet that challenge will affect Internet users throughout the world. The RIPE community acknowledged just how important this process is in a community statement issued in October 2007 at the RIPE 55 Meeting in Amsterdam, expressing the community's support for the widespread deployment of IPv6 and urging that "the widespread deployment of IPv6 be made a high priority by all stakeholders".

The RIPE NCC has been providing access to all of its major online services over IPv6 since December 2008, including web servers, FTP servers, DNS, the LIR Portal and email system as well as the RIPE Database and all of the RIPE NCC's Information Services.

As well as leading by example, the RIPE NCC is uniquely positioned to play an important role in educating and advising the various stakeholder groups affected by IPv6 deployment. Working together with the "traditional" RIPE community, the RIPE NCC is reaching out to business, government, law enforcement, regulators and civil society. It is vital that these stakeholder groups work together if we are to ensure that the adoption of IPv6 is as smooth as possible, and it is in everybody's interest that all stakeholders be well informed of the issues that surround IPv6 deployment.

Providing straightforward and accessible information has been a focus of the RIPE NCC's IPv6 efforts over recent months. This includes the recent launch of the IPv6ActNow website (**www.IPv6actnow.org**), a website that explains IPv6 in terms that everyone can understand. The website provides a variety of useful information, including videos from experts, the latest IPv6 statistics, news and links, and is aimed at promoting the global adoption of IPv6.

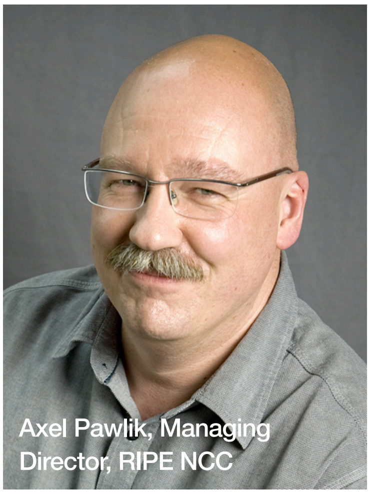

Axel Pawlik, Managing Director, RIPE NCC

The RIPE NCC Training Services Department has also recently launched a classroom-based IPv6 training course as well as an E-Learning module on IPv6.

For several years, RIPE NCC staff have been engaging with specific groups and organisations, such as the OECD, the European Commission and the Internet Governance Forum (IGF), on IPv6-related developments. The RIPE NCC has also recently been working with the European Commission on the IPv6 Deployment Monitoring Survey, which aims to give a clear view of IPv6 deployment in the RIPE NCC service region and to suggest what can be done to ensure the Internet community is ready for the widespread adoption of IPv6. The results of this survey will contribute to a better global picture of current and future IPv6 deployment.

Our hope at the end of this process is to see IPv6 fully deployed in a manner that takes account of the needs of all stakeholders, and that does not endanger the stability of the Internet. By helping to ensure that all groups are well informed, and that any myths or inaccuracies about IPv6 are dispelled, the RIPE NCC is working in the interests not only of its members, but also of the wider Internet community.

Axel Pawlik, Managing Director, RIPE NCC

IPv6: RIPE NCC and the OECD

RIPE NCC Contributes to Work of the OECD

The RIPE NCC has been involved in the work of the Organisation for Economic Co-operation and Development (OECD) since 2007. At the recent OECD Working Party on Communication and Infrastructure and Services Policy (CISP) Meeting, held in Paris on 15-16 June, the RIPE NCC, together with the other Regional Internet Registries (RIRs), submitted a document for the session on IPv6 Deployment. The document contained information on "Measuring IPv6 Deployment" and an overview of IPv6 allocations per country. The document was well received by the delegates and prompted discussion and further work. The contribution to this meeting can be found online at:
www.nro.net/news/cisp-IPv6.pdf

Formation of ITAC

In March 2009, together with the other RIRs and several industry partners, the RIPE NCC formed the Internet Technical Advisory Committee (ITAC), further strengthening its relationship with the OECD. ITAC is a formally recognised OECD committee and is able to comment on and contribute to the OECD's work to develop Internet-related policies. ITAC primarily works with the OECD Committee for Information, Computer and Communications Policy (ICCP) and its specific working parties such as the Working Party on Communications and Infrastructure Services Policy (CISP), the Working Party on Information Economy and the Working Party on Information Security and Privacy (WPISP). More information about ITAC and the organisations involved can be found at:
www.internetac.org

One-Stop Website Explains Everything You Need to Know About IPv6

In May 2009, the RIPE NCC launched the IPv6 Act Now! Website (**www.IPv6ActNow.org**), a one-stop website explaining IPv6 in terms everyone can understand. The website provides a range of information aimed at promoting the global adoption of IPv6.

The website is for anyone with an interest in IPv6, including network engineers, company directors, law enforcement agencies, government representatives and civil society. The website content is regularly updated and includes:

- Education, advice and opinions from Internet experts
- Latest IPv6-related news stories
- Videos and articles from Internet community leaders
- Current IPv4 exhaustion and IPv6 uptake statistics
- The RIPE community's statement on IPv6 deployment, including a list of organisations supporting this statement
- Information on community-developed IPv6 distribution policies
- Useful links to other sources of information about IPv6
- A forum for everyone to share experiences, ask questions and find answers

The website also includes contributions from other Regional Internet Registries (RIRs) and industry partners.

To find out about the latest developments in IPv6, visit:
www.IPv6ActNow.org

If you have and comments or suggestions about IPv6ActNow!, please contact us at:
IPv6actnow@ripe.net

Measuring IPv6 Deployment

By Geoff Huston and George Michaelson, APNIC

This article is an abridged version of a contribution made by the Number Resource Organization (NRO) to the OECD's Working Party on Communication and Infrastructure and Services Policy's (CISP) Session on IPv6 Deployment. The full document, providing a range of statistics and analysis based on the methodologies described below, is available from the NRO's website at: **www.nro.net**.

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?

This article outlines 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. It also examines the utility of such measurements in the context of measuring the status of global IPv6 deployment. The three primary data sources discussed in this paper are the Border Gateway Protocol (BGP) inter-domain routing table, packet captures of Domain Name System (DNS) server queries and access logs from web services.

Measurement of IPv6 Routing

Measurement Using Global IPv6 and IPv4 Routing Tables

This approach compares the number of entries in the IPv4 inter-domain routing table to the number of entries in the IPv6 inter-domain routing table. The rationale here is that the size of the IPv6 routing table will increase in line with progressive IPv6 deployment.

The problem with comparing 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 due to factors of historical legacy, and the routing practices for traffic engineering and multi-homing in IPv4. There is also the consideration of the prevalent use of IPv6 as an overlay network, forcing the route policy determination role into IPv4 routing rather than IPv6 routing. [2]

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 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 useful to examine different routing metrics to measure relative deployment of IPv6.

[1] For more information on IPv4 exhaustion predictions, see: IPv4 Address Report – **www.potaroo.net/tools/IPv4**

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

Measurement Using Autonomous Systems

Another useful approach may be to look at the number of routing entities that are routing IPv6, where each ISP and each distinct 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, and this can be compared to the number of AS's routing IPv4.

These two forms of routing measurements, active IPv6 BGP routing entries and AS numbers in the IPv6 routing table, may be taken at any point in the inter-domain routing space. 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, Firstly, a metric of a network's potential 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. 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. 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. With these caveats in mind, there are, however two measurement approaches that are relatively accessible, namely the server logs from dual-stack DNS and Web servers.

The DNS View of IPv6

Another long-term data set available for examination is usage data from a number of DNS servers. 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. Clients send queries and DNS servers issue responses. The DNS servers we have used for this particular exercise 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.

However, 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 as yet unresolved issues of interpretation.

Web Server Logs: Application Space Measurement of IPv6

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

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, the web server will only record an IPv6 transaction in its logs if all the intermediate components of the connection are configured to support IPv6. Therefore, this approach can be interpreted as a good indicator of the total level of IPv6 deployment capability across all components of the network. The relative deployment metric can be generated as the ratio between the number of unique IPv4 clients in a given period and the number of unique IPv6 clients. As long as clients are configured to attempt an IPv6 connection in preference to an IPv4 connection, then as the levels of deployment of IPv6 increase the relative metric of IPv6 clients to IPv4 clients should rise.

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 Measurement of IPv6 Deployment

There are many potential windows for collecting data on IPv6 adoption across the Internet. This article has investigated but a few of the options available for measuring IPv6 deployment. The authors of this article 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.

It would be interesting to understand 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 article 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 other measurements.

How IPv6 Policy Evolved

It was more than a decade ago, in 1997, when the RIPE community started the first policy related discussion about a new protocol called IPv6. At the beginning, the question was more about the policy framework for the new protocol rather than what the policy itself should be.

By 1999, we saw the very first results of this discussion: a common policy (identified as RIPE Document 196 by the RIPE community) for all three RIR communities that existed then (APNIC, ARIN and the RIPE NCC).

The basic goals of these initial IPv6 policies were not different to today's: uniqueness, conservation, aggregation and registration, with an understanding that conservation is not actually the biggest concern for IPv6. These goals have been central since the beginning of the registry system and, with the exception of conservation in the case of IPv6, are the same for any Internet number resource. However, once you look into the specifics of the old IPv6 policies, you can immediately notice the following differences from today's IPv6 policies.

Firstly, there was an extended initial allocation which included the requirement of peering with three other IPv6 networks. Obviously this requirement was found to be especially problematic at least in the early phases when IPv6 networks were almost non-existent. So the community decided to have a "Bootstrap Phase" which offered relaxed criteria until 100 organisations had received an IPv6 allocation. At this point the original, more stringent criteria, would return.

During this Bootstrap Phase it was required that the requesting organisation had peerings with three other Autonomous Systems (AS) in the default-free zone and had six months of 6Bone (an IPv6 project at the time) experience. Alternatively, the requesting organisation was required to show that that they were a transit provider and they had issued IPv4 to at least 40 End Sites that could meet the criteria of a /48 assignment in IPv6. The requirement of having plans to provide IPv6 services stayed the same in the Bootstrap Phase as well.

The minimum allocation size was a /35 and there was no HD ratio in the initial policy. Instead, further allocation criteria was set, as with IPv4, to 80% usage of the previous space. Minimum End Site assignment size was set to a /48.

This initial policy's mission was to be a starting point to create IPv6 awareness in communities and to get them thinking and talking about this new protocol. Accordingly the policy document, ripe-196 included a clause that called for a revision to the policy three months after it was published in July 1999.

The subsequent discussions were mostly focused on the fixed size for End Sites and whether this was a good idea or not. There was also discussion about whether the minimum allocation size should be changed to a /32 and if more relaxed initial allocation criteria was needed. In May 2002, during the RIPE 42 meeting, all these topics finally reached consensus and a new RIPE Document (ripe-246) was published in June 2002. This time it was called the "Interim Policy".

One can see that among other differences in this new document, the HD ratio was introduced to measure the utilisation of an IPv6 allocation, the minimum allocation size was set to a /32 and a totally new initial allocation criteria was set: an organisation had to be an LIR, not an End Site and had to plan to provide IPv6 connectivity to organisations to which it would assign /48s by advertising that connectivity through its single aggregated address allocation. In addition, LIRs had to have a plan for making at least 200 /48 assignments to other organisations within two years.

In 2002, IPv6 address space for Internet Exchange Points (IXPs) and Root servers was also introduced in the RIPE community with specific criteria.

We saw some other changes after the acceptance of the "Interim Policy". It was the last joint, common policy among the Internet communities. The ARIN community chose to change their criteria as did the new RIRs, LACNIC and AfriNIC, in accordance with the different needs of their communities.

By 2003, it was apparent that in the RIPE community there was demand from people who wanted to experiment with IPv6 and a special criteria for experimental networks was set.

As we reached 2007, the RIPE community witnessed two important changes in its IPv6 policies: one relating to changing mainly the initial allocation criteria by removing the requirement of planning to make at least 200 /48 assignments to other organisations; the other one relating to removing the requirement of fixed assignment sizes, and leaving the assignment size decision to the LIR. With the same policy, the HD ratio was changed and set to a more conservative value.

By April 2009, a long-lasting discussion reached consensus and Provider Independent (PI) resources were introduced into RIPE IPv6 policies. Since then, End Users can receive IPv6 allocation and assignment based on specific criteria.

When the earliest IPv6 policies were made, there was still no significant operational experience with IPv6 and it was hard to make policy given the absence of technology and business drivers for IPv6. Thus the policies that were initially developed, and have been subsequently changed, were mostly focused on making it easy to receive address space so that all types of organisations could receive the resources they initially needed.

However during the last couple of years we have seen an acceleration in technology-driven policy changes in IPv6, as operational experience has increased. One significant example of this is the recent extension of the special policy for Anycasting cc/gTLDs to ENUM operators.

Overall, the policy changes are a good sign that IPv6 has become a primary focus of Internet communities today. As the interest in IPv6 as well as the operational experience increases, policies will continue to change to adapt to the actual developments in the industry, the reality in business requirements and the need of the Internet community in general.

As of July 2009, there were two proposals relating directly to IPv6 policies in the RIPE PDP, one questioning the existence of routing requirements in the policy, the other discussing the IPv6 PI needs of LIRs. More details of these proposals are available at: **www.ripe.net/ripe/policies/proposals**.

Stay tuned to the latest policy developments by following them in the Address Policy Working Group: **www.ripe.net/ripe/wg/address-policy**

RIPE NCC Information Services – Focus on IPv6

With the depletion of IPv4 on the horizon for some years now, discussion about IPv6 at government meetings, industry conferences and various mailing lists has increased dramatically. But with the best estimates suggesting that depletion is just 24 months away, how have the political and policy debates influenced IPv6 uptake on the Internet?

The RIPE NCC's Information Services activities are focused on the measurement and monitoring of various aspects of the Internet, and the network of monitoring probes has supported IPv6 for over five years. Analysing data from our history database allows us to make various observations about growth in IPv6 usage and overall penetration.

Following is a short overview of our Information Services.

Hostcount

www.ripe.net/hostcount

Since 1990, the RIPE NCC's region Hostcount has tracked Internet growth by enumerating A records, and, since 2008, AAAA records, in forward DNS zones. The peak measurement occurred in September 2007 when 87,000 AAAA records were counted alongside 116 million A records with the AAAA total representing 0.075% of the A record total.

Routing Information Service (RIS)

www.ripe.net/ris

Since 1999, the RIS has collected Internet Routing Data over BGP from a total of 15 Remote Route Collectors (RRC) worldwide. IPv6 data has been collected since 2002, and now monitor this from 12 of our collectors. While we observe 307,000 unique prefixes worldwide, just 0.6% (or 1,800) of these are IPv6.

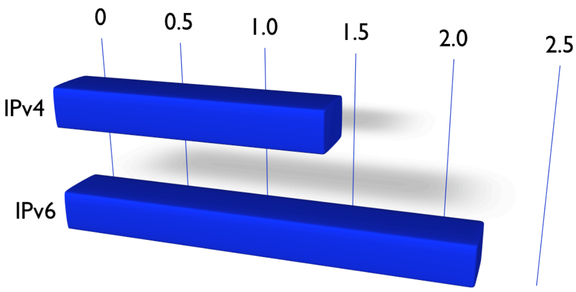

This chart shows the average number of BGP updates per prefix per day for IPv4 (1.49) and IPv6 (2.22), suggesting that IPv6 is almost 1.5 times less stable than IPv4.

Test Traffic Measurements (TTM)

www.ripe.net/ttm

The TTM network, operated by the RIPE NCC since 2000, allows users to monitor their connectivity to other parts of the Internet in near real-time. Of the current 77 worldwide nodes, 36 are dual stacked, meaning that we can compare IPv4 and IPv6 quality between identical points.

One feature of TTM is the tunnel discovery tool, which measures the Maximum Transmission Unit (MTU) over an entire path and attempts to detect tunnels in use. Here the view in 2005 shows widespread tunneling, with around 48% of paths (in yellow) traversing tunnels.

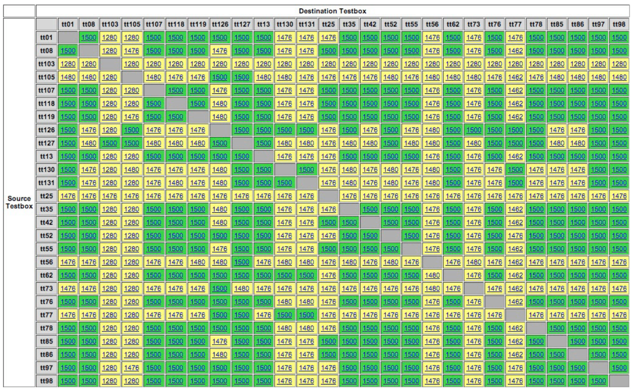

The view in 2009 is significantly different. Aside from more nodes giving rise to more measurement points, the number of native IPv6 links is measured as 88.5%

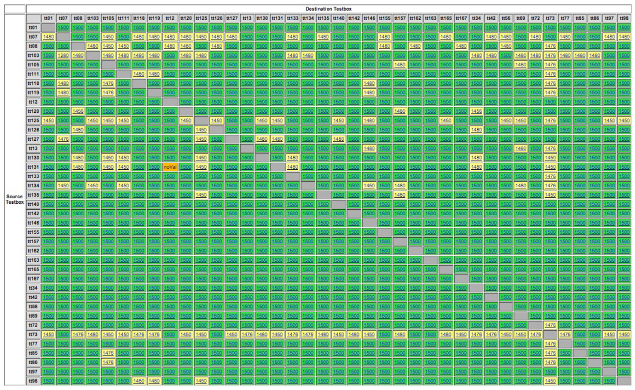

Another function of the TTM network is to measure latency between every node in the global mesh. In 2004, 18 nodes were dual stacked, and the difference in link quality between IPv4 and IPv6 is significant. This chart shows the average median latency in milliseconds, with IPv6 latency 38% higher than IPv4.

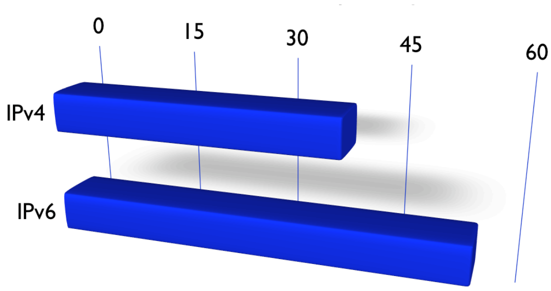

Five years on, the view from 36 nodes is much improved, and the gap has narrowed to 17%.

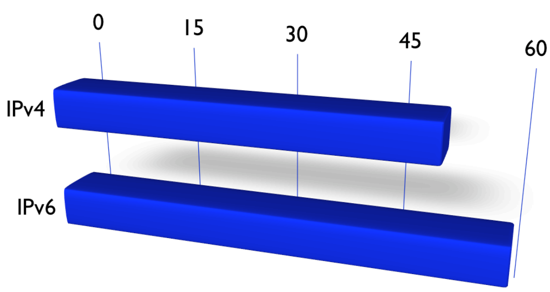

DNS Monitoring (DNSMON)

dnsmon.ripe.net

DNSMON uses a globally distributed measurement network to measure the end-user experience of DNS servers for 38 TLDs and infrastructure domains. Every TLD which is monitored by DNSMON has at least one server available over IPv6, showing a commitment from infrastructure providers to support IPv6.

Although there is support for IPv6, there are fewer IPv6 visible servers (3.4 on average) than IPv4 servers (9 on average), which suggests less traffic arriving over IPv6 at present:

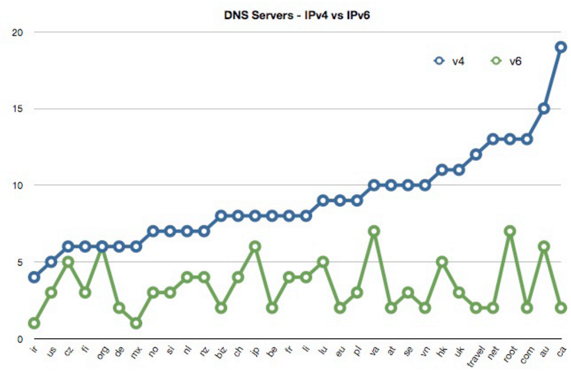

This is just a tiny subset of data collected by the RIPE NCC, but we can still draw some clear and positive conclusions:

- IPv6 link quality is improving – latency is decreasing, and so is the use of tunnels
- IPv6 stability is improving – but with so much testing and development going on, it's still not as stable as IPv4
- Infrastructure providers are showing a commitment to IPv6 but there is still a long way to go.

If you're interested in the work that the RIPE NCC Information Services group does, and want to access our data, or even participate in our measurement projects, you can find out more at **http://is-portal.ripe.net**

RIPE Community Resolution on IPv4 Depletion and Deployment of IPv6

During the Closing Plenary at the RIPE 55 Meeting in October 2007, the RIPE community agreed to issue the following statement on IPv4 depletion and the deployment of IPv6:

Growth and innovation on the Internet depends on the continued availability of IP address space. The remaining pool of unallocated IPv4 address space is likely to be fully allocated within two to four years. IPv6 provides the necessary address space for future growth. We therefore need to facilitate the wider deployment of IPv6 addresses.

While the existing IPv4 Internet will continue to function as it currently does, the deployment of IPv6 is necessary for the development of future IP networks. The RIPE community has well-established, open and widely supported mechanisms for Internet resource management. The RIPE community is confident that its Policy Development Process meets and will continue to meet the needs of all Internet stakeholders through the period of IPv4 exhaustion and IPv6 deployment.

We recommend that service providers make their services available over IPv6. We urge those who will need significant new address resources to deploy IPv6. We encourage governments to play their part in the deployment of IPv6 and in particular to ensure that all citizens will be able to participate in the future information society. We urge that the widespread deployment of IPv6 be made a high priority by all stakeholders.

The original version of this statement is available at:
www.ripe.net/ripe/meetings/ripe-55/presentations/steffan-resolution.pdf

RIPE NCC Represented at IPv6 Related Events

As a Regional Internet Registry (RIR) and an organisation with more than fifteen years experience supporting the technical co-ordination of the Internet, the RIPE NCC and its staff are deeply involved with the latest IPv6 developments. As a result, RIPE NCC staff have given presentations throughout the RIPE NCC service region on various aspects of IPv6 deployment.

Communications and Information Technology Comnmission IPv6 Workshop, Saudi Arabia – February 2009

Axel Pawlik, RIPE NCC Managing Director, and **Paul Rendek**, RIPE NCC Head of External Relations and Communications, attended a workshop focused on IPv6 developments in Saudi Arabia. Axel gave a presentation on RIPE policy development related to IPv4 and IPv6. The workshop was organised by the Communications and Information Technology Commission (CITC) and was held in Riyadh, Saudi Arabia.

LINX IPv6 Workshop, London – March 2009

Arno Meulenkamp, RIPE NCC Technical Trainer, attended the LINX IPv6 workshop in March 2009 in order to share the RIPE NCC's experience of successfully deploying IPv6 and to explain how to get IPv6 address space. His message was simple: any organisation that needs IPv6 address space can get it from the RIPE NCC. The event, which attracted over a hundred industry professionals, was aimed at an audience of network architecture and network engineering staff and focussed on practical IPv6 deployment.

At the same meeting, **James Aldridge**, Senior Systems Engineer, RIPE NCC and **Erik Romijn**, Software Engineer, RIPE NCC, gave presentations on deploying IPv6 for RIPE NCC services and IPv6 operational experiences from RIPE Meetings.

OGF25/EGEE User Forum, Italy – March 2009

Xavier Le Bris, RIPE NCC IP Resource Analyst, gave a presentation on IPv4 depletion and IPv6 deployment at the OGF25/EGEE User Forum, Catania, Italy. Xavier's presentation focused on the IPv4 "endgame" and the need for deploying IPv6 sooner rather than later.

IPv6 Space Odyssey, the Netherlands – March 2009

Nathalie Trenaman, RIPE NCC IP Resource Analyst, gave a presentation on IPv4 depletion and IPv6 deployment at the IPv6 Space Odyssey, Ede, the Netherlands. Nathalie's presentation discussed the issues involved with IPv6 deployment and provided a range of recent statistics on IPv6 allocations by country and over time.

German IPv6 Congress – May 2009

Daniel Karrenberg, RIPE NCC Chief Scientist and Chairman of ISOC, gave the keynote presentation at the first German IPv6 Conference, held in Frankfurt in May 2009. Daniel's presentation focused on the importance of deploying IPv6 now in order to avoid the problems associated with the depletion of available IPv4 address space. The conference, hosted by heise Netze, iX and DE-CIX, featured a range of presentations and tutorials and attracted 200 attendees.

INEX Meeting, Ireland – June 2009

Vesna Manojlovic, RIPE NCC Advanced Courses Trainer, gave a presentation on RIPE, RIPE NCC and IPv6 at the INEX meeting in Dublin, Ireland. The presentation focused on IPv6 including the process for receiving IPv6 address space and current initiatives for supporting, and monitoring, IPv6 deployment.

RIPE NCC Staff at your Event

If you would like to invite the RIPE NCC to speak about IPv6 at your event, please contact: speaker@ripe.net

RIPE NCC and IPv6 Training

The RIPE NCC is the organisation to go to for information on IP addresses, and we are happy to announce that we now also have a training course focused specifically on IPv6.

Ever since the RIPE NCC started registering IPv6 address space, we've included information about IPv6 in our regular LIR Training Course. We treated it as an Internet number resource, much like IPv4 addresses and AS numbers.

However, now that we are nearing the exhaustion of available IPv4 address space, more emphasis is needed on the need to deploy IPv6. With the estimated IANA IPv4 address pool exhaustion less then two years away, Internet organisations can no longer afford to put off thinking about IPv6. It is time to plan and implement now.

Looking at how the RIPE NCC could help with this, we noticed that there were already a number of technical IPv6 training courses available, some offered by RIPE NCC members. We did not want to follow the same approach as these courses, because it wouldn't add new information and also because it is not our role to offer hands-on technical training courses.

Instead, we chose to focus on real-world case studies of those organisations that have decided to deploy IPv6. Part of our training plan is a series of interviews with people from different organisations that have deployed IPv6. Some of these are already available, with more to come, at:
www.ripe.net/training/e-learning

Our goal is to provide RIPE NCC members with a structured approach for creating an implementation plan specific to their organisation. This will allow them to implement IPv6 using a phased approach and should enable members to make a business case for implementing IPv6 in their organisation.

It is important to note that the IPv6 training course will give pointers and an overall structure. It will not tell you which hardware to buy, or which commands to enter, as it is not a technically focused training course. We focus on the need for deploying IPv6, how to get IPv6 addresses and how to plan for deployment. In addition, the course points out some of the common pitfalls that people who have already implemented IPv6 have encountered so that you don't have to make the same errors. By the end of the course, you should be able to identify the steps that are needed to roll out IPv6 services in your organisation.

We would like to extend an invitation to all our members to attend one of our IPv6 training courses. For a list of upcoming courses, and to register, please visit:
www.ripe.net/training

If the locations of these courses are not convenient for you, you could consider hosting a course. More information about this is available at:
www.ripe.net/training/hosting.html

RIPE NCC and the European Commission Conduct IPv6 Deployment Study

The RIPE NCC has been working closely with the European Commission (EC) on an IPv6 Deployment Monitoring Survey. The survey aims to give a clear view of IPv6 deployment in the RIPE NCC service region and to suggest what can be done to ensure the Internet community is ready for the widespread adoption of IPv6. The results of this survey will contribute to a better global picture of current and future IPv6 deployment. These results will be discussed at RIPE 58 in Lisbon, Portugal, in October 2009 and will be published on the IPv6ActNow website at: **www.IPv6actnow.org**.

The survey was developed in consultation with members of the RIPE community, and is inspired by the 2008 survey conducted by ARIN and CAIDA in North America. It is sponsored by the EC, which has actively supported the adoption of IPv6.

The Big Easy: Finding Your Way Around IPv6

By Leo Vegoda, IANA

Internet Assigned Numbers Authority

IPv6 is big. In fact, it's so big that the Internet Assigned Numbers Authority (IANA) has five registries describing different parts of the IPv6 address space while there are just two for IPv4. Because IANA operates so many registries for the IPv6 address space, people can find it difficult to find the assignments or allocations they are interested in.

IANA wants to make sure that people find it as easy as possible to use the registries that it maintains, so IANA staff have developed a new dynamically generated web page which lets you see the whole IPv6 address space tree (see the example below). Users can use it to show and hide unallocated address space and whole branches can be shown and hidden as desired.

The tool IANA has developed is flexible enough to be applied to any finite number space divided between different IANA registries. So in the future IANA will start to provide this for other groups of registries too.

Address Type	IPv4	IPv6
The whole address space	Yes	Yes
Unicast	No	Yes
Multicast	Yes	Yes
Anycast	No	Yes
Special	No	Yes

More information is available at:
www.iana.org/protocols

Combined View of IPv6 Registries

Last Updated

2009-06-09

Description

This web page gives a combined view of all IANA IPv6 registries apart from the IANA IPv6 Anycast Address Registry. They are:

* IANA IPv6 Address Space Registry
* IANA IPv6 Global Unicast Address Registry
* IANA IPv6 Special Purpose Address Registry
* IANA IPv6 Multicast Address Registry

This combined view provides summary information, which can be sorted by column. Detailed information is available from the actual registries, which are linked from each entry in the table. Click on the checkbox at the bottom to hide and show the unallocated prefixes.

Please note that information about anycast address assignments are not included in this view because all unicast subnets contain anycast addresses. Please refer to [IANA registry *ipv6-anycast-addresses*] for details of anycast assignments.

Prefix	Description	Source Registry
☑ **Show unallocated space**		
0000::/8	Reserved by IETF	[IANA registry *ipv6-address-space*]
0100::/8	Reserved by IETF	[IANA registry *ipv6-address-space*]
0200::/7	Reserved by IETF	[IANA registry *ipv6-address-space*]
0400::/6	Reserved by IETF	[IANA registry *ipv6-address-space*]
0800::/5	Reserved by IETF	[IANA registry *ipv6-address-space*]
1000::/4	Reserved by IETF	[IANA registry *ipv6-address-space*]
⊞ 2000::/3	Global Unicast	[IANA registry *ipv6-address-space*]
2000::/16	Unallocated	
⊟ 2001:0000::/23	IANA	[IANA registry *ipv6-unicast-address-assignments*]
2001:0000::/32	TEREDO	[IANA registry *iana-ipv6-special-registry*]
2001:1::/32	Unallocated	
2001:0002::/48	BMWG	[IANA registry *iana-ipv6-special-registry*]
2001:2:1::/48	Unallocated	
2001:2:2::/47	Unallocated	