

Is Global IPv6 Deployment on Track?

Mónica Domingues, Carlos Friaças

Foundation for National Scientific Computing (FCCN)
Lisbon, Portugal
{monica.domingues, cfriacas}@fccn.pt

Pedro Veiga

Faculty of Sciences, University of Lisbon
Lisbon, Portugal
pmv@di.fc.ul.pt

Abstract— IPv6 has been the subject of a significant amount of research work in the networking field for more than a decade. The main aim of this paper is to discuss the current status of IPv6 deployment. It is not the aim of the authors to provide any breakthrough discovery which can boost IPv6 usage, but instead focus on what has happened in the past and what is currently limiting the IPv6 global deployment. The 6NET project [1], along with other national and international initiatives produced a vast collection of knowledge about the next generation Internet protocol. This work is still to be used massively in order to evolve the Internet into its next generation.

Keywords: IPv6, Routing, Peering, NAT, Addressing;

I. INTRODUCTION

Mostly due to some academic/research and vendor's efforts one can say IPv6 is mature and a reality on some of today's networks, along with IPv4. However, everyone must keep in mind that IPv6 is still to reach the majority of existent public IP networks around the globe. On early March 2007, the number of visible Autonomous Systems according to the CIDR Report [2] is nearly 24800 on the IPv4 world and only about 700 on the IPv6 world. This clearly supports the idea that IPv6 global deployment is still in its initial phase – i.e. still a lot of work to do!

IPv6 in commercial/enterprise environments is still in a premature stage. The lack of significant new functionalities when compared to IPv4 and the investment needed to upgrade the network infrastructure in order to support it, are often referenced as the two main reasons for the slow evolution. In the scope of the first factor, this missing piece is often referred as the lack of a killer application which would lead to IPv6 mass deployment and adoption. About the second hurdle, the investment to be carried out is not only about hardware and software, but also focused on training people. In this field, the 6DISS project [3] (a spin-off of 6NET) is trying to spread the word about IPv6 in developing regions. IPv6 must be seen as a worldwide need, the Internet cannot be really upgraded and expanded without everyone's effort. It is pretty clear at this point that every IP network administrator's contribution is crucial to accomplish this.

The perspective addressed by this paper is mainly related to network deployment, not application deployment. While both layers are important, the first one is structural for the second. In reality, the application layer will greatly benefit from all problem solving related to the network layer.

II. MYTHS ABOUT IPV6 DEPLOYMENT

Some myths are causing a negative impact on IPv6 deployment status all over the world.

Myth #1: In Europe and USA, there is a strong misconception that IPv6 is well underway in the Far East (mostly Japan, Korea and China). However, some statistics show otherwise.

Despite the lack of traffic data to analyze, the IPv6 address blocks already distributed until this moment, show which world region is in fact leading IPv6 deployment in the Internet: Europe. Frequent updated data and graphs can be found at <http://www.ripe.net/rs/ipv6/stats>. Current status, as of April 2007 is depicted in Figure 1. It is also clear that larger allocations were issued in the European Region.




Figure 1: Global IPv6 Allocations

Myth #2: IPv6 doesn't exist in real live networks, so there is no point in starting to deploy it at our own backyard.

The best example against this myth is academic networks, supplying IPv6 connectivity to a numerous community of teachers and students. Once again, like what happened 30 years

ago on the beginning of the Internet, academic communities are leading IPv6 deployment.

Myth #3: End-Users will ask ISPs for IPv6 connectivity.

This will never happen, simply because end-users only worry about applications and its performance. As an underlying support layer, IPv6 doesn't improve performance dramatically. As a consequence, deployment must occur on the ISP side, after all the majority of IPv6's enhancements only carries benefits to ISPs.

Myth #4: Services/Applications Transition into IPv6 is really hard.

This is currently (2007) not true, although there are still some minor/niche missing gaps. One of the best well-know problems is proprietary solutions, with badly programmed Application Program Interfaces (APIs). Hard coded IPv4 addresses are not only a nonsense barrier to IPv6 deployment, but also a serious flaw in terms of design/architecture.

Myth #5: Some Asiatic countries went ahead with IPv6 because they are getting a hard time obtaining IPv4 address space.

This is an entirely false perspective, and it is explained in detail at the "Global Addressing System" chapter.

III. GLOBAL ADDRESSING SYSTEM

The last myth described above is easy to clarify: the same rules to obtain IPv4 address space apply to all countries in the world (according to self-defined policies in their region), through the existing Regional Internet Registries (RIRs) system. The unused space is a common global resource; it hasn't privileged access rights to any country in the world.



Figure 2: RIRs system map

The global distribution system of addressing space is in fact very simple. A unique entity – IANA (Internet Assigned Numbers Authority) [4], which is currently operated by ICANN (Internet Corporation for Assigned Names and Numbers) [5] – allocates address blocks to five RIRs, which cover all world regions (Figure 2) and also have formed the Address Supporting Organization (ICANN ASO) [6] . The five RIRs distribute the blocks they receive to National Internet Registries (NIRs) - 1 per country where they exist - or directly to ISPs (Internet Service Providers), which are often referred as

LIRs (Local Internet Registries). The NIR concept is only applied in the APNIC region. Figure 3 illustrates how the system works in a top-down approach.

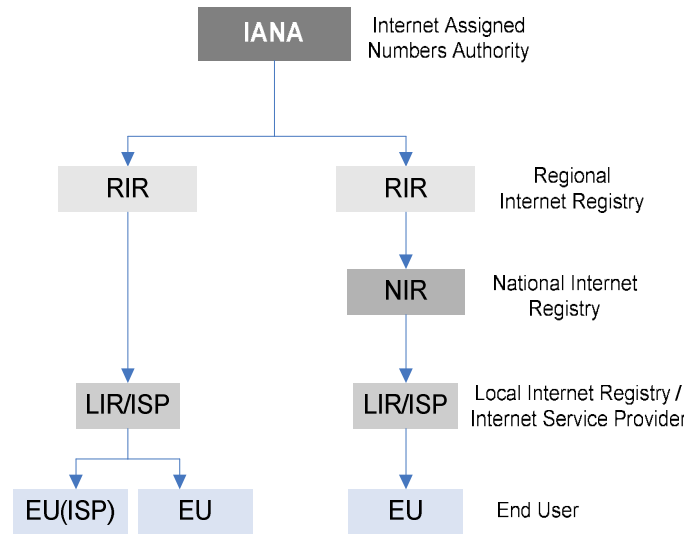


Figure 3: Top-Down Allocation Addresses Process

IV. IPV4 EXHAUSTION

The current shortage of globally unused IPv4 space leads us into doubt about when that available space will in fact be exhausted.

Several authors like Tony Hain¹ from CISCO SYSTEMS and Geoff Huston² from APNIC have different opinions, but using the past allocation data, a timeframe of 2011-2013 can be easily drawn. Geoff Huston's daily updated predictions set from IANA's and the RIRs' real live data is an impressive and useful tool to show people when the limits of public IPv4 Internet will be reached. Yes, it is certainly not going to be tomorrow, but it is going to happen!

While for some specialists IPv6 is inevitable and decisions were taken to do something about it, others feel IPv4 will still be capable of accommodating everyone's needs for the decades to come. The timeframe mentioned above, if correct, implies that the IPv4 Internet will soon lose its ability of expanding. Can the world face this fact and do nothing?

In parallel with the IPv4 address space exhaustion, there is also an historical boundary in terms of identifying autonomous systems, i.e. administrative domains on the public Internet, with its own routing policy. This issue is being timely addressed and, from 2009 on, valid identifiers are going to be from 0.0 to 65536.65536, instead of the current 0 to 65536 range. Of course this expansion-type evolution is not as hard as going from IPv4 to IPv6, because this doesn't impact the application layer, but only the software on the underlying global routing system.

¹http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html

²<http://www.potaroo.net/tools/ipv4/index.html>

V. RUNNING AWAY FROM IPv6?

There are alternative plans to avoid IPv6 usage. However, they are somehow negative. The possible options include: increasing the level at which Network Address Translation (NAT) is used; creating a “black market” of public IPv4 addresses. A third one, although less serious, includes hiding heads in the sand and decide not to act at all.

Using several levels of NAT was referred to be a solution for address shortage/extinction by some, but in practice the world already knows NAT doesn’t scale well, although it has its fair share of success in several scenarios. The current level of complexity is already a serious menace to the development of new technologies. Hardening that complexity would only increase management costs and would certainly reduce performance.

The idea of making money out of stockpiled address blocks is not pleasant for many people around the world. The current distribution system in fact states that address blocks are not property and thus cannot be sold to third parties. If future scarcity of IPv4 addresses helps generate this unfair circuit of trade, the foundations on which the Internet was built in the first place will be seriously damaged. The Internet was envisioned for democracy, and this type of new reality could lead to a system where public IPv4 addresses will only be at the reach of large pockets.

Doing nothing about the emergence of global public IPv4 address scarcity is also another choice. One can always believe the numbers and projections are wrong, and that IPv4 space is endless. When the time for the crude reality arrives, then costs for a hasty transition into IPv6 will rise too. Shortening the transition phase, where everything should be accessible both via IPv4 and IPv6, will only cause havoc and increase costs.

VI. CURRENT OPERATIONAL PROBLEMS

Moving from what is envisaged to happen in the future; this paper also intends to describe problems that are currently happening with IPv6 deployment globally.

It is impossible to have a global IPv6 Internet, without the proper peering between networks and without having adequate transit relationships in place. While in the first case, Internet Exchange Points - IXP (or Network Access Points - NAPs) are somehow promoting IPv6 usage among their membership, the IPv6 commercial offer from global carriers is still very limited. In fact, the effort and the goodwill shown by IXPs which quickly obtained IPv6 address blocks for their own infrastructure is insufficient from a global perspective. The lack of commercial offer from global IP carriers can be observed from CAIDA’s data in Figure 4 and Figure 5. The first one illustrates the light «core» of the IPv6 Internet, whereas the second shows the denser inner centre of the IPv4 Internet.

Looking at Figure 4 and Figure 5, it is interesting to perceive which networks are on the heart of both the «old Internet» and on the «new Internet». Only Global Crossing, Cable & Wireless, Tiscali and Sprintlink (however, using another AS identifier!) are seen in both maps.




Figure 4: IPv6 Internet's core

Carriers like Teleglobe/VNSL, UUNET, AT&T, Verio, Telia and others are found to be still missing. This is an important hurdle to be won. It is not cost wise for an ISP to buy IPv4 transit from Carriers A and B, and IPv6 transit from Carriers C and D. The second half of this option would certainly add extra cost. What should be the turning point on this equation? ISPs are a meaningful piece of the global puzzle, and as customers should start to ask their current IPv4 carriers to provide also for IPv6 transit with the same quality.

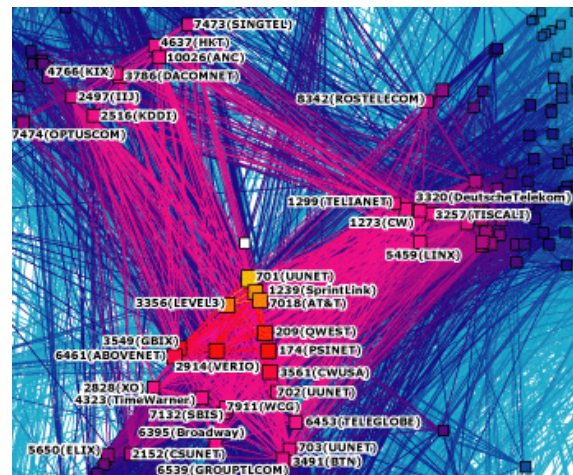


Figure 5: IPv4 Internet's core

The fulfillment of such customer requests should not be really a difficult hurdle for carriers because the most important component (and expensive!) of the IP carrier service, be it IPv4 or IPv6, is its long-haul capability. As an intermediate solution, the offer of dark fiber can also empower ISPs to reach more distant points on the Internet, and would allow them to reach other IPv6 networks at their own discretion. In some way this has been the path for some academic networks in their decision to go ahead without depending on available solutions from traditional carriers.

Other direct conclusions learned from observing Figure 4 and Figure 5 is the number of different networks on each and the degree of connections between them. To perceive a fair

level of IPv6 deployment, we must see in the future the same networks in both cores, as well as an equivalent number of connections. It's essential to generate a solid and reliable IPv6 global infrastructure for end-users and their applications.

Tunnels are still a problem. The 6BONE has formally ended in 6th of June 2006, however many people still insist in using 6BONE address space, or send traffic through long distance tunnels. If on one hand, tunnels are the only solution to connect to the IPv6 world, on the other hand they can be harmful to IPv6 networks' usage experience, if badly drawn. A poor deployment may lead to a negative impact on end-users' perception about the advantages and disadvantages of IPv6. As a consequence, a slower or negative evolution of IPv6 on real networks can happen.

There are some recurrent issues about proper filtering. Excessive strictness can imply unreachability to some networks. What has been diagnosed several times is that a number of network administrators still understand the IPv6 production Internet as being all inside 2001::/16, despite the allocations already made by RIRs. Gert Doring's filtering recommendations [7] are often cited as the de-facto standard for proper BGP filtering in the IPv6 world.

FCCN (AS1930), in the recent past, already witnessed some problems related to IPv6 performance. Reporting the issue on IPv6 latency between our network and a commercial network in Switzerland didn't solve it completely, although with no doubt the complaint effectively reduced IPv6 latency, thus improving quality. The issue was forwarded to one of GEANT2's IPv6 transit providers (Level3), and the IPv6 latency was reduced to half after some fine tuning. We still notice a double value when comparing IPv4 and IPv6 latencies, but this was nevertheless a significant improvement.

Receiving some IPv6 routes is sometimes made impossible when networks generating those prefixes don't register it properly on a Routing Registry database. At a first sight, not being able to instantly receive new IPv6 prefixes through GEANT2 is a negative issue from a NREN perspective, but it's also a positive evolution to see the same policies being applied in IPv4 and IPv6.

VII. KEEPING AN EYE ON LATENCY

One of the planned work areas related to the work described in this paper was the construction of a mechanism to evaluate IPv4 and IPv6 latency. This includes measuring dual-stack latency between FCCN's network and several other IPv4+IPv6 networks around Europe. The first big hurdle in setting this up was finding those networks/hosts in every of the twenty seven European Union countries. A second hurdle was to find networks outside the «national research network family». Neither we did find one target in all the EU countries, nor did we have the choice of not resorting to academic networks. In some cases proper «targets» were found through [8] and [9]. At a first glance we know tests being run towards other research networks (to which FCCN is connected at a high quality standard) are most likely to produce similar values for IPv4 and IPv6 roundtrip times. From the total sample of 28 «targets», 17 traverse the GEANT2 network.

Figure 6 shows the difference on average latency times towards «targets» inside GEANT2, and Figure 7 shows the same values related to external measuring points.

Figure 6 and Figure 7 include values on milliseconds which reflect the difference in readings while running ping and ping6 at the same time for a host in a given country. A negative value states that IPv6 latency was better than IPv4 latency between FCCN's network and the host's network during the period studied. The study period was the month of March 2007. Readings were done every 15 minutes with 5 probes for protocol version, each time. Average values were kept and used to achieve the final averages shown.




Figure 6: GEANT2 related targets

The most interesting difference read during this experiment was to HEANET's Network, the Irish Academic Network. In good true, the paths used for IPv4 and IPv6 between FCCN's network and HEANET are the same in terms of AS path. The usage of IPv6 as a second network protocol clearly showed that the primary protocol's (IPv4) performance can in fact be improved. This also happens regarding the Italian (GARR) and Lithuanian Academic Networks (LITNET). However, the difference seen in the Irish case is a lot more relevant because it represents a possible improvement in IPv4's overall latency bigger than 20%.




Figure 7: (IPv4+IPv6) towards external GEANT2 hosts

From latencies measured towards hosts external to GEANT2, it is easy to perceive there is still a lot of room for improvements. Only the Portuguese dual-stack «target» which is a native peer of FCCN over the local Internet Exchange (Gigapix [10]) shows almost no difference between IPv4 and

IPv6 roundtrip times. That should happen in relation to all other hosts when proper IPv6 deployment becomes a reality.

But we also targeted some dual stack hosts outside the European Union space. Figure 8, which illustrates measurements towards Russia, the Vatican, United States and Japan, we can see that this problem can get worse. The best sign we see from the values collected is from Japan (M-root server), and the Vatican who can in fact be reachable through GARR, making once again use of the GEANT2 infrastructure. Values measured to the US were related to the H-root server and to Freenet in the case of Russia.




Figure 8: Other Non-European Union measurements

Looking deeply into the set of latencies measured towards the Japanese host, it is interesting to take into account extreme values:

- 56.73 ms difference in favor of IPv6 (332 ms against 389 ms);
- 141.57 ms difference in favor of IPv4 (316 ms against 458 ms).

It is important to notice that the M-root server is anycasted in three more locations (Paris, San Francisco and Seoul). From the roundtrip time values, it is safe to understand the Paris mirror instance was never measured against, which cannot in fact be stated in relation to the other two cities with the same degree of certainty.

DNS Root Servers and RIPE Route Collectors (RRC) were also analyzed as possible targets. Unfortunately, only three of the RRC and four of the five root servers [11] that have IPv6 configured are indeed reachable over IPv6.

From this background, we were mainly interested in confirming the quality at which IPv6 is working inside the European Academic Backbone - GEANT2. We were also eager to get a general overview about the quality of IPv6 transit to other IPv6 commercial networks in the European Union. Unfortunately, GEANT2 isn't currently peering with IPv6 commercial networks directly at several cities around Europe.

VIII. SOME BENCHMARKING WORK

Two metrics are being continuously measured by FCCN in the scope of the European IPv6 Task-Force Steering Committee Project [12]: Local Internet Registry's IPv6

coverage, and IPv6 Autonomous Systems usage. The first is monitored through a daily updated list which Europe's RIR - RIPE - makes available, and the second through IPv6 BGP data. While the first parameter of analysis shows that around 13% of European Union LIRs already have an IPv6 address block, looking at the routing system, we can only see about 2.73% of independent networks with one foot in the IPv6 world.

In Figure 9 we can perceive a slow evolution in the last year, and specially note a clean-up of inactive LIRs was made, marginally increasing the percentage of IPv6-ready LIRs.

The percentage of LIRs which already obtained an IPv6 address block must be seen as an encouraging number; despite being only associated with a developed region (Europe), it shows that IPv6 awareness campaigns are somehow working among ISPs, and it unveils that they are starting to prepare for mass IPv6 deployments. However, the Million Euro question remains: will it be enough until IPv4 gets exhausted?




Figure 9: IPv6 Ready LIRs Evolution

The second result of our continued monitoring efforts is a lot less positive. In fact, on April 1st 2007, 24797 different Autonomous Systems (ASes) were to be found on the IPv4 global routing table, as a counterpoint of only 679 in the IPv6 global table. There is also a curious number: out of the 679 different ASes, 33 ASes were only seen in the IPv6 world, not in the IPv4 world. This would adjust the 2.73% figure to just 2.60%.

Looking a bit further into the IPv6 global routing table, at the beginning of April 2007, it is possible to see that from 676 different ASes 345 (51%) are from the RIPE region, 169 (25%) from the APNIC region and 159 (23.5%) from ARIN's. There are also three from LACNIC and three more unidentified (ASes 3, 24 and 2607). There are probably also some more ASes still registered in the three elder RIRs' databases (APNIC, RIPE and ARIN) which in fact belong to African and South American Networks. One issue that cannot be really assessed from these numbers is size. In fact, as each network is represented by one AS number, a small network with few customers is placed at the same level of a bigger nationwide or even of a pan-continental network.

The numbers presented in this section show, again and at some level, that Asian leadership in the IPv6 world is a myth, but also demonstrates there is a rather long to go until IPv6 is globally available. This of course in terms of a new public IPv6

Internet infrastructure, because at an industrial level, the IPv6 Ready Program [13] has clearly identified the fact that most IPv6-enabled products are originated by Asian countries (namely Japan, Korea and Taiwan).

Adding to the work described above, we also frequently look at the top of the DNS system. Having around 90 Top Level Domains supported by dual stack DNS authoritative servers is undoubtedly a good indication. Figure 10 describes the evolution since October 2004, however with an interruption in the readings from May 2005 to March 2006.




Figure 10: IPv6 Enabled Top Level Domains

At the same level, despite the inexistence of IPv6 addresses published in the system for root servers (authoritative servers of zone “.”), there’s also something to be optimistic about the time when 5 of the 13 root servers’ managers claim to have enabled IPv6 [11].

IX. CONCLUSIONS

There is an important item that should be carefully analyzed when deploying IPv6 networks – design it in the same way IPv4 is done. If something works well in IPv4, it should function in a similar way in the IPv6 world - with some few exceptions, and where possible using the advantages inherent to IPv6.

The forward path for the public Internet is the “IP agnostic” concept. With the increased availability of IP agnostic Operating Systems and Applications (which use IPv4 or IPv6 indistinctly, upon availability) the stability on the IPv6 side of each dual-stack network will be fundamental. In fact, inadequately deployed IPv6 can discredit the benefits end-users may feel by having public addressing space available to them.

Some IPv6 advertising is based on claiming that IPv6 addresses should be free to end-users. However, the authors strongly feel this issue could be viewed from another perspective [14], benefiting both end-users and ISPs. ISPs can get extra revenues from assigning fixed IPv6 subnets to customers, and end-users could get an extended public address space for a fair price.

The authors believe that current IPv6 deployment should be in a more advanced stage. Network operators which already

deployed IPv6 need to enhance and fine tune their IPv6 design, in the same way it is crucial that operators still managing IPv4-only networks decide to go ahead and enter the IPv6 era. IPv6 has to be understood as a global effort.

A simpler and quicker answer to this paper’s title would be: NO.

AUTHORS

Mónica Domingues graduated from the Polytechnical Institute of Guarda in 2006. She joined FCCN in 2004, and since then she worked in the IPv6 Portuguese Documentation Project. Mónica is also contributing to 6DISS, GEANT2’s AutoBAHN project and the Service Activity on the Performance Enhancement and Response Team (PERT).

Carlos Friaças is working for the Portuguese NREN since 1998 and dealing with IPv6 since 2001. He has graduated in Computer Science from University of Lisbon in 1999. He has participated in 6NET, 6DISS, the European IPv6 Task-Force Steering Committee and the Portuguese IPv6 Task-Force, while also writing IPv6 related documentation in Portuguese. He is also involved with the Portuguese Internet Exchange Management, and with the daily operation of RCTS – the Portuguese Academic Network.

ACKNOWLEDGMENTS

The work presented in this paper is financed by POS_Conhecimento, in the scope of the DAR-II-ING project (ref. POSC/338/1.3/C/NAC, FEDER). Its main objective is to promote the usage of IPv6 services in the RCTS network, and to produce guideline documents to other organizations concerning IPv6 adoption.

The authors would also like to thank FCCN’s CTO João Ferreira, and to the remaining wide area network team: Pedro Lorga, Ana Pinto and Luís Marta for all the support.

REFERENCES

- [1] 6NET Project site, www.6net.org;
- [2] CIDR Report, www.cidr-report.org;
- [3] 6DISS – IPv6 Dissemination and Exploitation, www.6diss.org;
- [4] IANA Home Page, www.iana.org;
- [5] ICANN Home Page, www.icann.org;
- [6] ICANN ASO Home Page, www.aso.icann.org;
- [7] Examples BGP IPv6 filters, www.space.net/~gert/RIPE/ipv6-filters.html;
- [8] List of IPv6 enabled hosts, <http://ipv6.loercks.net>;
- [9] SIXXS site, www.sixxs.net;
- [10] Portuguese Internet eXchange Point, www.gigapix.pt;
- [11] Root servers site, www.root-servers.org;
- [12] European IPv6 Task-Force Steering Committee Project, www.ipv6.eu;
- [13] IPv6Ready site, www.ipv6ready.org;
- [14] “The New Public Internet & Profits for ISPs within IPv6 Deployment”, www.fccn.pt/files/documents/D5.13.PDF.