

Probing Next Generation Portuguese Academic Network

Carlos Friças, Emanuel Massano, Mónica Domingues

Foundation for National Scientific Computing (FCCN), Lisbon, Portugal
e-mail: {cfriças, emanuel.massano, monica.domingues}@fccn.pt

Pedro Veiga

Faculty of Sciences, University of Lisbon, Lisbon, Portugal
e-mail: pmv@di.fc.ul.pt

Abstract

Beyond plain connectivity, one of today's biggest challenges for network managers and designers is guaranteeing network quality at any moment for evolving users' needs. FCCN has been deploying several probes on the Portuguese Research and Education Network during last year. The analysis of data collected from this set of probes as well from other monitoring systems deployed inside RCTS will be considered. In this paper we will present and discuss the evaluation of the collected data by crosschecking results of the several monitoring systems. Our main focus will be the network layer running IPv6, but data from probes operated exclusively in IPv4 will also be presented, as a comparative source of network information.

Keywords

Monitoring, Performance, Measurement, IPv6;

Introduction

National Research and Educations Networks (NRENs) have been in use by a large number of scientists and research projects interested in maximizing the benefit from using advanced networks, namely GÉANT [1], as an important resource for achieving goals in a broad range of scientific fields.

The recent concerns [2][3] about IPv4 address space exhaustion increased the attention given to IPv6 deployment. Despite the constant warnings about the urgency to move to IPv6, coming from several sources over the last years, only now the attention on the need to move faster to deploy IPv6 is starting to rise. It is now clear that everyone needs to enter the «IPv6 era» in order to cope with the Internet's ongoing growth [4].

IPv6 has been the focus of a significant amount of work done in the Portuguese NREN in the last decade. Indeed FCCN has been a pioneer in the Portuguese Internet, continuously working towards its stability and improvement. FCCN clearly understands its role as maintainer of local critical infrastructures, such as the operator of the Portuguese ccTLD .PT [5] and GigaPix [6], the main Internet exchange in Portugal.

In the context of this work, it is necessary to add that IPv6 has been continuously promoted inside the Portuguese Research and Education Network for a long time. But on October 2007 and during the 12th RCTS (RCTS is the Portuguese Research and Education Network) members' meeting [7], the Portuguese NREN issued a challenge to its members: “enable IPv6 support in all basic network services, such as DNS (Domain Name System) servers, Web servers and E-Mail servers of each institution connected to RCTS”. There is an expectation that this initiative will help generate more IPv6 awareness, more IPv6 compliant systems and, consequently, more IPv6 traffic within RCTS and, thus, enrich the set of data obtained by the monitoring probes. At the same time, in order to assure quality for the newly deployed IPv6 services it is necessary to fine tune tools to demonstrate that the IPv6 network is not underperforming when compared with its operation with IPv4's.

The evaluation of collected data will rely on crosschecking results from several monitoring systems that have been installed inside relevant nodes of the network. This method will provide a greater degree of accuracy and understanding about IPv6 performance.

Monitoring infrastructure at RCTS

After a tender process, a QoSmetrics' (acquired by Symmetricom, Inc. [8] on January 2007) proposal was selected in June 2006 to deliver the monitoring solution to be deployed inside RCTS. This monitoring solution architecture comprises high capacity network probes (NetWarrior 301, 1Gbps) and low-end probes (NetWarrior 100, 100Mbps). Both models are Linux®-based. It has been decided that the low-end probes should always be placed at members' premises, with two possible architectures: nearby members' access routers; directly connected to (layer two) equipment owned by FCCN.

In some cases, the probe is located inside the member's network (University, Polytechnical, or Research Laboratory). This architecture provides the ability of monitoring the performance characteristics between the institution RCTS' entry point and the backbone.

The set of QoSmetrics probes provide measurements related to one-way performance metrics, such as One-Way Packet Loss (OWPL) [9], One-Way Delay (OWD) [10], and One-Way Jitter, which are important in the overall monitoring of RCTS' backbone. The time synchronization needed to achieve one-way performance data is based on one GPS antenna, connected to the Lisbon high capacity probe, and on the usage of QTP, the QoSMetrics proprietary time protocol, for the low-end probes. The high-end probe is the time reference for the later.

Currently, twenty-six probes are operational. However only a subset is running IPv4 and IPv6. The high capacity probes have been installed in Lisbon and Oporto, taking into consideration RCTS' dual-star topology [11]. However, only the Lisbon high capacity probe is used for measurements. For the purpose of this work, we consider data collected from ten probes that are running both IPv4 and IPv6. The remaining set of probes could also be running IPv6 over tunnels, but the application that sets the measurements imposes the use of a physical interface (eth0). This restriction implies that the virtual interface (sit1) used to create the tunnel endpoint on a probe itself cannot be used for measurements. The NetWarrior 100 (low-end) IPv6 active probes accountable for this work have been geographically distributed over eight cities along the country, in addition to the central network nodes located in Lisbon and Oporto, as depicted in Figure 1.



Figure 1 Geographical Distribution of IPv6-enabled Probes in RCTS.

The set of selected probes can be divided in two groups in terms of topology: those who connect directly to RCTS' central network node in Lisbon (and measurement point); and those that are attached to the Oporto backbone node. The first set includes two probes in Lisbon (Polytechnical Institute of Lisbon and University of Lisbon) and the probes installed at Coimbra, Tomar, Évora, and Faro.

<i>Probe</i>	<i>#</i>	<i>City</i>	<i>Domain</i>
nwUMinho	1	Braga	Backbone
nwIPB	2	Bragança	Member
nwUCoimbra	3	Coimbra	Member
nwUGuarda	4	Guarda	Member
nwCovilha	5	Covilhã	Backbone
mwTomar	6	Tomar	Member
nwIPL	7	Lisboa	Backbone
nwULisboa	8	Lisboa	Backbone
nwEvora	9	Évora	Member
nwAlgarve	10	Faro	Backbone

Table 1 Probe Domain.

There is also another possible grouping view, related to the «administrative domain» where each probe is installed. Table 1 shows the complete probe list considered in this paper, highlighting if the probe is attached on RCTS' backbone, or inside a RCTS' member network. The probe at Tomar has a small difference towards every other probe: Despite having native IPv6 available in the local link, its gateway (the next hop) is an IPv6 over IPv4 tunnel end. All the other probes have native IPv6 paths to reach RCTS' main node at Lisbon. This detail can be verified by using the `tracpath6` [12] unix tool in order to test each path and see the IPv6 MTU in each hop. The data samples collected from the deployed QoSMetrics infrastructure at RCTS are from 19th March 2008 to 11th April 2008. The term «slightly» indicates a difference below the value of ten percent.

Probe #	Analysis
1	IPv6 packet loss is marginally higher than IPv4, and losses were registered exactly on the same periods for both protocols.
2	IPv6 packet loss is slightly higher than IPv4 in the Lisbon-Bragança direction and slightly lower on the Bragança-Lisbon direction.
3	Only packet loss on the path Coimbra-Lisbon was registered, and IPv6 packet loss was slightly lower.
4	Significant IPv6 packet loss was registered on both directions between Lisbon and Guarda; however, there were no IPv4 packet losses, as no IPv4 measurement data was found.
5	Packet loss was present during the entire sample, and IPv6 loss was slightly lower in the Covilhã-Lisbon path while in the opposite path it was significantly lower.
6	IPv6 packet loss is slightly higher than IPv4.
7	There were no losses on the path from IPL to RCTS' main node, however, in the opposite path some minor losses occurred, with approximate values for IPv4 and IPv6.
8	Significant losses on IPv4 and IPv6 were recorded on the path from the central node to University of Lisbon, with a marginal difference showing more packets lost on the IPv6 network. There is also the detail about the losses time frame: in the last two thirds of the period analyzed packet loss was heavily reduced (from 26th March onwards).
9	IPv4 packet loss is lower on the Évora-Lisbon direction; however packet loss is higher in IPv4 over the Lisbon-Évora path. The amount of packet losses is also very different. It's significant on the second case, and marginal on the first case.
10	IPv6 packet losses were not seen, while IPv4 packet losses were low on the Faro-Lisbon path and significant on the Lisbon-Faro path.

Table 2 Packet Loss Samples Analysis.

Table 2 shows a brief comparison between IPv4 and IPv6 packet loss for every of the ten probes.

Path	Probe #	IPv4 (milliseconds)	IPv6 (milliseconds)
To Lisbon	1	2.635	2.778
From Lisbon	1	2.951	n/a
To Lisbon	2	5.730	6.151
From Lisbon	2	5.581	5.933
To Lisbon	3	3.714	4.121
From Lisbon	3	3.946	4.239
To Lisbon	4	n/a	10.641
From Lisbon	4	n/a	12.892

To Lisbon	5	6.671	7.254
From Lisbon	5	7.466	7.871
To Lisbon	6	7.762	n/a
From Lisbon	6	6.337	9.805
To Lisbon	7	0.487	0.480
From Lisbon	7	0.812	0.935
To Lisbon	8	0.741	0.954
From Lisbon	8	0.633	0.826
To Lisbon	9	1.188	1.177
From Lisbon	9	1.020	1.128
To Lisbon	10	3.477	3.133
From Lisbon	10	3.862	3.300

Table 3 Average One Way Delays.

From our perspective most of the packet loss seen derives from full usage (mostly on the member network to RCTS direction) of bandwidth capacity intervals during certain parts of the workday.

One-way delay (OWD) [10] measurements inside RCTS are a central aspect of this paper. Average values are presented in Table 3. Shaded lines indicate measurements with lower IPv6 average OWD, which is the most uncommon comparison found.

Table cells containing «n/a» indicate it is not possible to establish a comparison, because either the IPv4 jitter [13] or the IPv6 jitter was absent from the data sample. The missing readings may have been caused by ill configured firewalls (related to probes inside a member network), or by QoSmetrics software misconfiguration (in the case of backbone hosted probes).

Table 4 compares jitter (One way delay variation) values found, from the same sample.

Shaded lines indicate probes where higher jitter values depend on which direction measurements are related. Conclusions are hard to draw from Table 4, but two patterns emerge: IPv4 higher averages and standard deviations are mainly seen on the Lisbon-*Probe_Location* direction; even in the same direction, it is possible to find different higher IPv6 average jitter values, and higher IPv4 standard deviation values.

Path	Probe #	Higher Average Jitter	Higher Standard Deviation
To Lisbon	1	IPv6	IPv6
From Lisbon	1	n/a	n/a
To Lisbon	2	IPv6	IPv6
From Lisbon	2	IPv6	IPv4
To Lisbon	3	IPv6	IPv6
From Lisbon	3	IPv6	IPv4
To Lisbon	4	n/a	n/a
From Lisbon	4	n/a	n/a
To Lisbon	5	IPv6	IPv6
From Lisbon	5	IPv4	IPv4
To Lisbon	6	n/a	n/a
From Lisbon	6	IPv6	IPv6
To Lisbon	7	Equal	Equal
From Lisbon	7	IPv6	IPv6
To Lisbon	8	IPv6	Equal
From Lisbon	8	IPv4	IPv4
To Lisbon	9	IPv6	IPv6
From Lisbon	9	IPv6	IPv6
To Lisbon	10	IPv6	IPv6
From Lisbon	10	IPv4	IPv4

Table 4 Jitter values comparison.

IP Performance Metrics

FCCN, as a partner in JRA1 [14] (Joint Research Activity 1) subproject of GÉANT2 (European Research Network), is also hosting inside RCTS four HADES (Hades Active Delay Evaluation System) probes. These devices (also referred as IPPM - Performance measurement in Internet Backbones) were developed by WiN-Labor [15] at RRZE (Regional Computing Centre Erlangen), to provide quality of service measurements in DFN's G-WiN infrastructure and other domains based on metrics developed by IETF's IPPM WG [16]. HADES devices have been deployed in numerous GÉANT2 Points of Presence (PoPs).

Inside RCTS, HADES devices are located in Lisbon, Oporto, Coimbra, and Aveiro. From this set of probes, only Aveiro's is not operating in an IPv6 network. Unfortunately, Oporto's probe and Coimbra's probe have not started to perform any IPv6 measurements. This section and architecture has an obvious difference when compared to QoS Metrics mentioned previously: its scope. IPPM is deployed (and gathers data) on multiple Academic network domains (i.e. autonomous systems). Currently, the IPPM network has thirty-nine nodes; however, only six are configured in order to perform IPv6 measurements. Those six probes are installed in Amsterdam, Athens, Budapest, Lisbon, Sofia and Zagreb as displayed in Figure 2. There is also an extra drawback: the timeframe used is more limited than RCTS QoS Metrics' sample. The period covered by the data in the available sample is from 8th April 2008 to 12th April 2008.

Aligned with the goals of this paper, we only consider measurements over paths including Lisbon on its origin or destination. On the five paths inside Géant2, independent on the direction, OWD values are very similar in IPv4 and IPv6, with slightly higher average values on IPv6 measurements (below 10%). On the Sofia-Lisbon and Lisbon-Sofia paths average OWD values suddenly increase and some hours later decrease to previous values. Both IPv4 and IPv6 data show this pattern. Regarding packet loss, data present on the samples does not show any distinguishable difference between the IPv4 and IPv6 network.



Figure 2 IPv6-enabled IPPM probes.

RIPE's Test-Traffic Measurements

Test Traffic Measurements (TTM) [17] service was launched by RIPE-NCC [18] in October 2000, aiming to provide any Internet Service Provider with data about the quality of its Internet connectivity. The TTM network can be used to diagnose problems involving third-party networks, and to understand issues regarding external/global connectivity. This may enable network administrators to act preemptively in order to avoid serious network problems. The system continuously records one-way delay and packet-loss measurements as well as IP router-level paths («traceroutes») between a large set of probe machines (referenced as «test boxes»). It is also possible to measure bandwidth and delay variation (jitter).

TTM deployment is done across several independent domains, with the advantage of covering not only Academic, but also Commercial networks. It is a third form of comparing IPv4 and IPv6 data involving RCTS, and with focus on the quality provided by our transit providers.

RCTS hosts two TTM boxes, one at Lisbon and another at Oporto, named respectively `tt55.ripe.net` and `tt72.ripe.net`. This pair of probes are also associated each one with a GPS antenna, which enables the probes as Network Time Protocol Stratum 1 servers. Since February 2003 TTM probes support IPv6, and during 2004 the instances inside RCTS started to communicate both over IPv4 and IPv6. However, due to the current RCTS' dual-star topology, which only has «border» points in Lisbon, the pair of probes works mainly as one redundant TTM node.

RIPE-NCC's TTM service has an associated policy with data disclosure; therefore, data presented in this paper only illustrates RCTS' point of view. Specific data about measurements is presented in an aggregated form and never disclosing any third-party network.

From a broader set of TTM boxes (locations), we chose thirty nodes, which are operating both on IPv4 and on IPv6. From this set, sixteen nodes are immediately reachable from RCTS over GÉANT2 (AS20965). The remaining fourteen nodes are reachable over the combination of GÉANT2 with at least one DWS (DANTE World Service) provider, thus implying a longer/worse path.

Table 5 and Table 6 show ratios between IPv6 and IPv4 one-way delay and packet loss, obtained from Lisbon's tt55.ripe.net on 30th March 2008. The values were measured in milliseconds, and the IPv6 average values were divided by average IPv4 values. As it can be easily checked, every ratio obtained is higher than one, which means one-way delay values found (independently of direction) were worse in the IPv6 network.

	IPv6 / IPv4 One-way delay (Incoming)	IPv6 / IPv4 One-way delay (Outgoing)
Academic (16 probes)	1.006	1.006
Commercial (14 probes)	1.377	1.288
Overall (30 probes)	1.237	1.179

Table 5 TTM One way delay ratios.

The one-way delay IPv6/IPv4 values towards TTM nodes located in other Academic Networks are somewhat encouraging. However, it seems outside of this domain the difference between IPv6 and IPv4 communication is still relevant. This matches with the known reality of tunnel usage [19] and a worse degree of peerings between IPv6 independent networks [20].

	IPv4 Packet Loss (Incoming)	IPv6 Packet Loss (Incoming)	IPv4 Packet Loss (Outgoing)	IPv6 Packet Loss (Outgoing)
Academic	0.014	0.075	0.070	0.075
Commercial	0.018	0.243	0	0.131
Overall	0.016	0.170	0.030	0.110

Table 6 TTM Packet loss data.

The data regarding packet loss also shows a worrying trend: packet loss over the IPv6 Internet is more frequent, despite the somewhat low values found.

Crosschecking Analysis

With three different systems (QoSmetrics, IPPM, and TTM) deployed, able to perform measurements both on the IPv4 and on the IPv6 network, we believe that the set of IPv4 and IPv6 operational data, which involves RCTS, is rather interesting. Several variables like the Interior Gateway Protocols used, the Inter-NREN connectivity (provided by Géant2 [1], operated by DANTE [21]) and the quality of IPv6 transit can be used to evaluate the global readiness of RCTS' IPv6 component. Figure 3 describes the three different scopes covered.

The original intention of performing the work presented in this paper with a larger number of measurements and nodes was mainly disrupted by two difficulties: The lack of flexibility on the QoSMetrics system to support measurements over tunnels' virtual interfaces, and the lack of IPv6 measurements on the NRENs scope. While we initially assumed the availability of IPv6 addresses on the probes would enable IPv6 measurements, they were not available unless centrally configured. This was only noticed on early April, and it resulted in a very small set of probes, which undoubtedly reduces gathered data value.

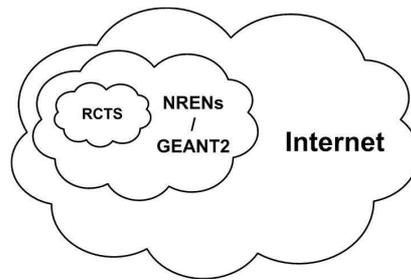


Figure 3 The three Monitoring Scopes.

Apart from the work related with the expanding of the IPv6-enabled probes in the several domains, data from the three domains seems to indicate that the IPv6 network is starting to be as reliable and performant as its IPv4 counterpart. However, there are still improvements to be made on all three domains, especially in the interconnection element with commercial IP networks.

Conclusion

Experience with deploying IPv6 has led us to believe that a poor deployment perspective will cause more harm to the Internet's future than no IPv6 deployment at all. End users will need to feel confidence that the new protocol version (that will enable Internet's continued growth) will keep their usage experience at the levels they already know. From the Service Provider perspective, this leads to the need of continuously collect network performance related measurements. In any transition plan, comparison between two steps is something unavoidable. This is again the case with IPv4 and IPv6.

This paper has two main differences with previous related work from other sources, namely [22]: On the RCTS domain, the networking technology is 100% based on Ethernet; since 2003 (when [22] was published), IPv6 technology has matured despite the low levels of worldwide adoption. These low levels can be confirmed simply by looking at autonomous system numbers present at each global routing system [23][24] (see AS Summary section).

The main goal of this work was to perform an initial approach on the quality level regarding RCTS' both IP stacks. This work based on real network measurements, not on theory. Future work will include the enhancement of measurements, the activation of more probes inside RCTS (which mostly depends on RCTS' member adoption of IPv6), and a continued eye on topology changes intended to improve quality.

The authors also feel that more sets of probes could be beneficial, namely including nodes on third party Portuguese networks, which are currently accountable for a significant portion of RCTS' external IPv4 traffic. The ultimate purpose is to provide RCTS' users the infrastructure to develop and use the most innovative end-to-end applications, without the severe constraints of address translation mechanisms.

Data from RCTS QoSMetrics infrastructure should get richer when measurements become possible towards our second main node (Oporto). There is a known problem with GPS synchronization currently rendering useless all efforts to obtain any measurements with origin or destination at Oporto. We will also focus on driving the IPPM team to enlarge the pan-european set of probes configured to measure IPv6 variables.

Acknowledgements

POS_Conhecimento, in the scope of the DAR-II-ING project (ref. POSC/338/1.3/C/NAC, FEDER), finances the work presented in this paper. The main objective of this project is to promote the usage of IPv6 services in the Portuguese Academic network, and to produce guideline documents to third-party 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, Luís Marta, and Luís Nunes for all the support. Stephan Kraft (DFN) and Nicolas Simar (DANTE) directly contributed to this work by enabling IPv6 measurements on the IPPM network.

References

- [1] Géant2 User Stories, <http://www.geant2.net/server/show/nav.1368>
- [2] Internet Assigned Numbers Authority, <http://www.iana.org/assignments/ipv4-address-space>
- [3] IPv4 Address Report, <http://ipv4.potaroo.net>

- [4] Internet Usage Worlds Stats, <http://www.internetworldstats.com/stats.htm>
- [5] ccTLD .PT, <http://www.dns.pt>
- [6] Gigabit Portuguese Internet Exchange, <http://www.gigapix.pt>
- [7] ECI-12, <http://www.fccn.pt/eci/index12.htm>
- [8] Symmetricom, Inc., <http://www.symmetricom.com/>
- [9] G. Almes, S. Kalidindi, M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC2680, September 1999
- [10] Géant2 PERTKB, <http://kb.pert.geant2.net/PERTKB/OneWayDelay>
- [11] FCCN's network topology, <http://www.fccn.pt/rcts>
- [12] <http://kb.pert.geant2.net/PERTKB/TracePath>
- [13] Géant2 PERTKB, <http://kb.pert.geant2.net/PERTKB/DelayVariation>
- [14] http://monstera.man.poznan.pl/wiki/index.php/Main_Page
- [15] Das WiN Labor Erlangen, <http://www-win.rrze.uni-erlangen.de/>
- [16] IETF IPPM Working Group, <http://www.ietf.org/html.charters/ippm-charter.html>
- [17] F.Georgatos et al, "Providing Active Measurements as a Regular Service for ISPs". In: Proceedings of Passive and Active Measurement Workshop (PAM2001), Amsterdam, April 2001. <http://www.ttm.ripe.net/>
- [18] RIPE.NET, <http://www.ripe.net/>
- [19] X. Zhou, P.V. Mieghem, "Hopcount and End-to-End Delay: IPv6 Versus IPv4", Proceedings of Passive and Active Measurement (PAM2005), Boston, MA, USA, April 2005.
- [20] IPv6 Operations List, hosted by ClueNet, <http://lists.cluenet.de/pipermail/ipv6-ops>
- [21] Delivery of Advanced Network Technology to Europe, <http://www.dante.org.uk>
- [22] S. Katsuno, T. Kubo, K. Yamazaki, H. Esaki, "Measurement and Analysis of Multimedia Application and IPv6 ADSL Internet Access Network", IEEE/IPSJ SAINT2003, January 2003.
- [23] CIDR Report, <http://www.cidr-report.org/as2.0/>.
- [24] IPv6 CIDR Report, <http://www.cidr-report.org/v6/as2.0/>.