QoS experiences in native IPv6 GRNET and 6NET networks

Athanassios Liakopoulos
Greek Research and Technology Network
Messogion Ave. 56, 11527, Athens, Greece
aliako@grnet.gr

Dimitrios Kalogeras
National Technical University. of Athens
Iroon Polytechneiou, 15780, Athens, Greece
dkalo@noc.ntua.gr

Vasilis Maglaris
National Technical University. of Athens
Iroon Polytechneiou, 15780, Athens, Greece
maglaris@mail.ntua.gr

Dimitrios Primpas,
Research Academic Computer Technology Institute and Computer Engineering & Informatics Department
University of Patras, 26500, Patra, Greece
primpas@cti.gr

Christos Bouras
Research Academic Computer Technology Institute and Computer Engineering & Informatics Department
University of Patras, 26500, Patra, Greece
bouras@cti.gr

Abstract

Adoption of IPv6 technology has been accelerated in the last few years but there is limited experience in the deployment of *Quality of Service (QoS)* for IPv6 traffic in backbone networks. As available software and hardware is designed to handle IPv4 packets, there is a need to accurately measure the performance of QoS mechanisms in an IPv6 environment. This paper discusses tests in the deployment of IPv6 QoS in core networks, namely the production dual stack GRNET and the IPv6-only 6NET networks, using both hardware and software platforms. In either case, we succeeded in delivering advanced transport services to IPv6 traffic and provided different performance guarantees to portions of traffic. The deployed QoS schema was common for IPv6 and IPv4; in most cases both v4 and v6 traffic exhibited comparable performance per class while imposing no significantly different overhead on network elements. A major conclusion of our tests is that the IPv6 QoS mechanisms are efficiently supported with state-of-the-art router cards at gigabit speeds.

I. Introduction

The deployment of IPv6 technology in the research and commercial networks has been accelerated in the last few years. Today, most of the *National Research & Education Networks (NRENs)* in Europe have deployed native IPv6 services, while GÉANT [1], the Trans-European research network, offers IPv6 transit connectivity in its service portfolio. Furthermore, ongoing research is focusing on migration, security, mobility, multicast challenges [2][3][4][5][6] in IPv6 environments, leading to the deployment of advanced IPv6 commercial services.

There has been a long history of discussions in reference of QoS support in IPv6 environments. There is a debate on whether "IPv6 provides better QoS support than IPv4" or "IPv6 experiences worst performance than IPv4". The objectives of our work were twofold; validate the performance of basic QoS mechanisms with IPv6 traffic at hardware- and software-based platforms and identify missing functionality or unexpected performance. The collected results allowed us to conclude that advanced transport services, which have been offered in IPv4 production networks, could also be delivered to dual stack networks provided that some conditions are fulfilled.

The paper is organized as follows: Section II describes the QoS related fields in the IPv6 header. In section III, we present the GRNET network, elaborate the current deployed (IPv4) QoS schema and analyse the results from performance tests with mixture of IPv6/v4 traffic. Section IV is dedicated to qualitative tests conducted in the 6NET IPv6-only core network. Finally, sections V presents *"wish-to-have"* functionality and section VI summarizes our conclusions and defines our future plans.

II. QoS related fields in the IPv6 header

The IPv6 header [7] is (re)designed to minimize header overhead and reduce the header process for the majority of the packets. This is achieved by moving less essential and optional fields to extension headers that are placed after the IPv6 header. Therefore, IPv6 and IPv4 headers are not *interoperable*. Furthermore, IPv6 header is not a superset – thus backward compatible – with

IPv4 counterpart.

The IPv6 header has two fields that are related to QoS; the *traffic class* and *flow label* fields. The 8-bit *traffic class* field is used to distinguish packets from different classes or priorities. The same functionality is provided from the *type of service* (or *precedence*) field in the IPv4 header and, consequently, there is no essential difference among the packet headers of the two protocols.

By definition, a *flow* is a sequence of packets sent from a particular source to a particular unicast, anycast, or multicast destination. In the IPv4 world, flow classification is based on 5 fields; IP source and destination addresses, transport layer protocol type and ports. However, some of these fields may be unavailable due to fragmentation or encryption of packets in the network. In order to overcome such problems, flow classification in IPv6 world is based on the 3-tuple consisting of the *flow label* plus the source and destination address fields, which are in fixed predefined positions in the IPv6 header. The *flow label* field [8] consists of 20 consecutive bits. Whenever the end host wants to identify the packets of a flow, it sets the *flow label* bits to the same non-zero value, which is unchanged throughout the network. Note that currently there is no application or service known to us that takes advantage of the *flow label* field.

It is easily concluded that IPv6 protocol, in terms of QoS functionality, is neither superior nor inferior to IPv4 counterpart. However, the available *flow label* field in the IPv6 header could be a valuable tool for the provision of services in the future.

III. GRNET CORE NETWORK

The Greek National Research and Educational Network - GRNET [9], interconnects approximately 90 universities and research institutes. The core network consists of twelve nodes interconnected with STM-16 *lambdas*, while the subscriber access links vary from 1 Gbps down to 128 Kbps (see Figure 1). GRNET currently supports native IPv6 interconnection services. Its core routers are Cisco GSR12400 series [15] with 4xGE [16] and 10xGE [17] line cards. Their 10xGE (*Eng4+*) cards, also called *Tango*, are mainly used in core links and support (IPv4) line rate switching capabilities. On the contrary, their 4xGE (*Eng3*) cards, also called *Tetra*, interface access links with advanced functionality in Layer 2 VLAN support. The main difference, in terms of IPv6 support, is the fact that *Tetra* cards switch IPv6 traffic in hardware while *Tango* cards in software.

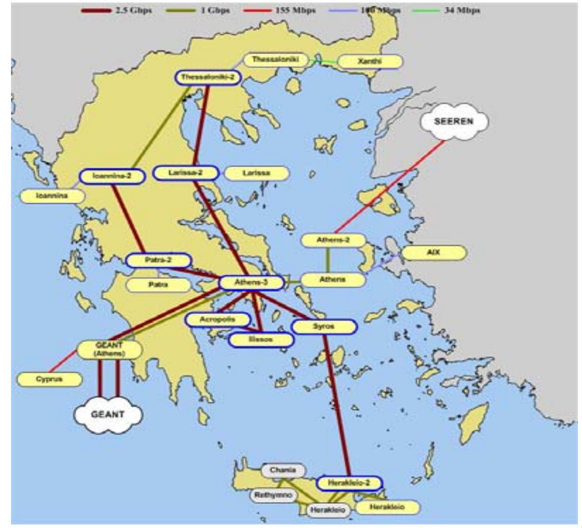

Figure 1: GRNET core network

A. QoS Model and Services

GRNET uses Differentiated Services (DiffServ) [10] in order to support different service guarantees to portions of traffic. The following three classes of service -in descending order of quality- are identified and deployed for IPv4 traffic:

- *Premium IP (PIP)*, based on *Expedited Forwarding PHB (EF-PHB)*, gives absolute priority over any other class and provides low delay/jitter plus negligible packet loss guarantees. It is suitable for real-time applications.
- *Best Effort (BE)* does not offer any qualified guarantees to traffic. It is indented for elastic applications.
- *Less than Best Effort (LBE)* exploits network resources without (negative) impact other traffic classes. It is suited for specific scavenger applications.

Premium IP class is further divided into three sub-classes; *PIP Virtual Wire, PIP VoIP* and *PIP Transparent*. *PIP Virtual Wire* is used for traffic exchanged between two well identified core interfaces and emulates a virtual circuit. *Premium IP VoIP* is used for voice traffic generated in a known source network but heading to an unidentified destination. *PIP Transparent* is used for high priority traffic routed towards GÉANT which is downgraded to BE in the domain borders.

Premium IP traffic is always serviced by output priority queues in core routers. Under stable network conditions, the PIP traffic can occupy up to 10%

of the link capacity in order to minimize inter-packet delay variation (jitter) and avoid starvation of lower priority traffic. A semi-automatic provisioning tool is used for performing the admission control and generating the appropriate router configuration. LBE traffic can potentially occupy all the available network resources and, in periods of high congestion, is granted 1% of the link capacity, which ensures that established connections do not brake. PIP traffic flavors are is marked with DSCP values 46, 47 and 40 while LBE traffic is marked with DSCP value 8.

B. *Testing equipment*

The GRNET tests were conducted using hardware-based traffic generators Smartbit 600 [18] with Gigabit Ethernet (GigE) interfaces attached in three different PoPs of the network (see Figure 2). They were connected either directly to core routers or via gigabit Ethernet switches, allowing us to assess performance of QoS mechanisms in physical and logical ports. The SmartFlow ver.3.0 application was used to control and measure the generated test traffic. GPS receivers were not employed since tests over the WAN did not involve performance measurements.

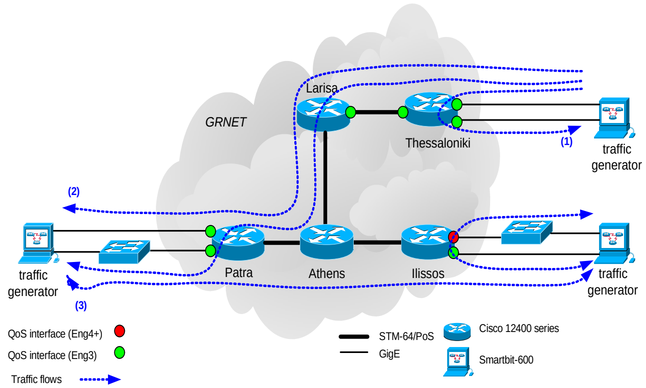

Figure 2: GRNET test bed topology

C. *Tests*

The traffic generators were able to produce in GigE ports a mix of IPv6 and IPv4 traffic up to 1Gbps. The test traffic load could congest the GigE access links but not the STM-16 core links. In all the tests the frame size at the data link layer was set 128 bytes and each test lasted 10sec[1]. Each testing packet was timestamped and counted by the traffic generator. Consequently, collected time-sensitive statistics for traffic generated and consumed in the same traffic generator was extremely accurate.

[1] Tests with different traffic parameters are noted in the text.

1) *Simple tests (classification, policing, shaping)*

The first set of tests was focused on *classification* mechanisms and *access lists* at the GigE interfaces of *Tetra* cards. A traffic generator produced 100Mbps IPv6 traffic with a specific address that was later filtered in the network via an IP-address based access list. Different tests verified the right operation of *input* access list on *physical* ports and *output* access list on *logical* (VLAN) port. Similar tests were successfully executed in the core interfaces (STM-16/PoS). The next set of tests was focused on *policing* mechanisms at GigE ports in *Tetra* cards. The traffic generators produced IPv6 traffic marked as *Premium IP* (EF) traffic, Traffic was policed at 100 Mbps while exceeding traffic was discarded at the output interface of a logical port (VLAN).

Another set of tests investigated output shaping in *Tetra* cards. Bursty traffic with average rate of 400 Mbps was shaped at rate 200 Mbps in the output of a *Tetra* port. Achieved throughput was measured at approximately 21,25% of the port capacity, as expected. Latency for IPv6 and IPv4 traffic was the same, approximately 443 msec and maximum latency was 9% greater than average latency. Without the shaping mechanism, maximum latency could be up to 5.5 times larger than average latency (236 μsec). As it was expected, shaping increased significantly the average latency of the packets.

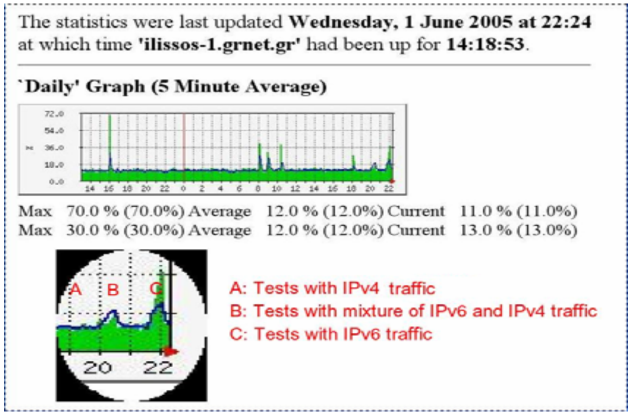

Figure 3: CPU tests

2) *CPU load*

The next set of tests investigated the impact on CPU load of IPv6 traffic switching at gigabit speeds. Bidirectional flows of total rate 2 Gbps were established with the traffic generator at Ilissos (Figure 2). Traffic entered the local router via two separated GigE ports in different cards; one *Tango* and one *Tetra*. In addition, approximately 500 Mbps –mainly IPv4- production traffic was passing through the router. Initial test with IPv4-only traffic at line rate speed for a 30-minute period did not increase the traffic load of the router which remained approximately at 11%. We repeated the test with mixture of IPv6 and IPv4 traffic in equal portions and we noticed a small increase 8% absolute value in the CPU load in 1 and 5 minutes time intervals, as in Figure 3 (B). The test was repeated with IPv6-only traffic. This time the CPU increased by 11% for five-minute intervals and by 26% for one-minute intervals, as in Figure 3 (C). Note that an ISIS routing problem affected production services during the test.

We repeated similar tests at the Patra router (Figure 2) where only *Tetra* cards were in use. In all cases, we did not measure any increase on the CPU load, most probably due to the fact that *Tetra* interfaces support hardware switching for both IPv4 and IPv6 protocols, unlike the *Tango* cards.

In the same set of tests with *Tetra* cards, we noticed that IPv6 BGP sessions in congested GigE ports were always affected and the sessions broke after a while. However, IPv6 BGP sessions in non-congested ports or in IPv4 cases were never affected. We concluded, therefore, that routing problems in the previous tests are not related with CPU load. We suspect that IPv6 control traffic is not protected in internal CPU queues, unlike the IPv4 case.

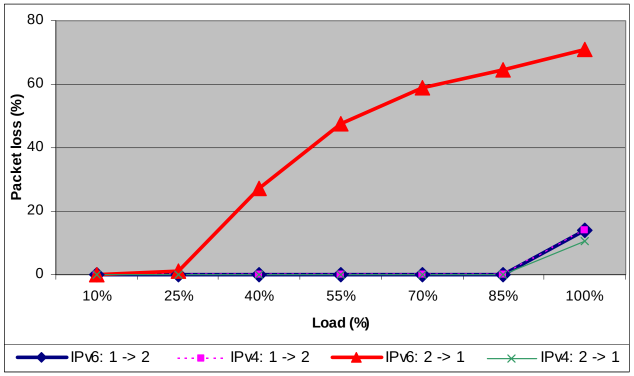

Figure 4: Packet loss for BE traffic

3) Latency and packet loss for BE traffic

The next set of tests investigated the latency and packet loss at gigabit speeds. Once again, bidirectional flows were established with the traffic generator. Traffic entered the local router via two GigE ports in different cards; one *Tetra* and one *Tango*. Traffic load was gradually increased from 10% to 100% (2 Gbps) of the port capacity in steps of 15%.

As shown in Figure 4, IPv6 traffic experiences higher packet loss than IPv4 traffic. Traffic loss for IPv6 traffic entering in *Tango* card (direction 2 ->1) is much higher than the packet loss in the opposite direction (1->2). Therefore, it is easily concluded that the *Tango* card does not support line rate switching of IPv6 traffic, as it does with IPv4 traffic. On the contrary, IPv6 and IPv4 traffic experience the same packet loss in *Tetra* cards under all traffic load conditions. Non zero packet loss (13,88%) is noticed for an 100% utilization in the *Tetra* case.

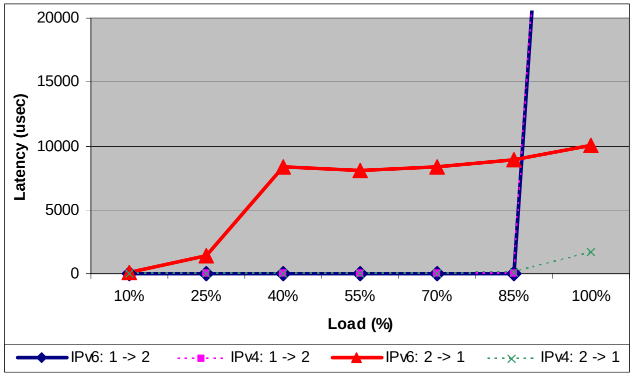

Figure 5: Latency for BE traffic

IPv6 and IPv4 traffic experience the same latency in *Tetra* cards, as shown in Figure 5. Even when there is packet loss (100% utilization), latency is increased equally for both protocols. Similar trends are noticed for maximum latency values (Figure 6) and latency distribution (not shown). On the contrary, latency for IPv6 traffic in *Tango* cards is twice higher, even with no packet loss. When there is IPv6 packet loss (but no IPv4 packet loss), the difference is increased by ~200 times. Similar trends are noticed in maximum latency values.

The last figures show that *Tango* cards are not capable to switch IPv6 traffic as efficiently as IPv4, which was expected as IPv6 traffic is software switched while IPv4 is hardware based. On the contrary, the *Tetra* cards provide similar services to both protocols.

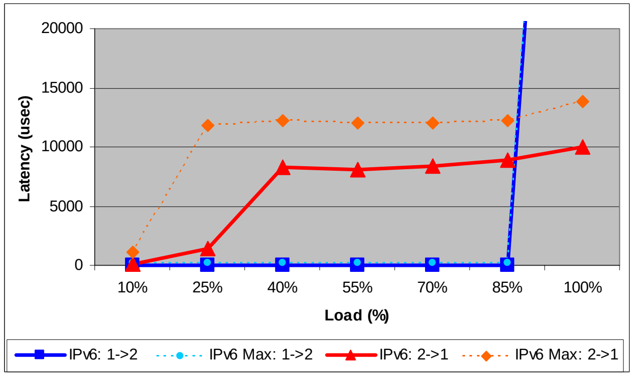

Figure 6: Average and maximum latency

4) Latency and packet loss for BE traffic for different packet sizes

The previous set of tests was performed with 128-byte packets. As real traffic consists of packets of diverse sizes, we repeated the tests for *Tango* cards but for larger packet sizes.

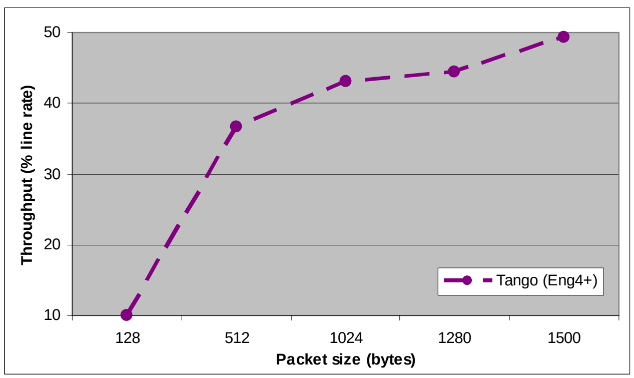

Figure 7: Throughput vs. packet size in *Tango*

As Figure 7 shows, throughput in *Tango* cards increased for larger packet sizes. However, even for 1500-byte packets, the loss is slightly more than 50%. The achieved throughput can be translated to approximately 41 (87) thousand 1500-byte (512-byte) packet per seconds. Provided that today the portion of IPv6 traffic in GRNET network is approximately 0.5% of the combined traffic, we do not foresee any IPv6 packet loss under normal conditions. This might not be the case in temporary IPv6 congestion instances, e.g. caused by DoS attacks.

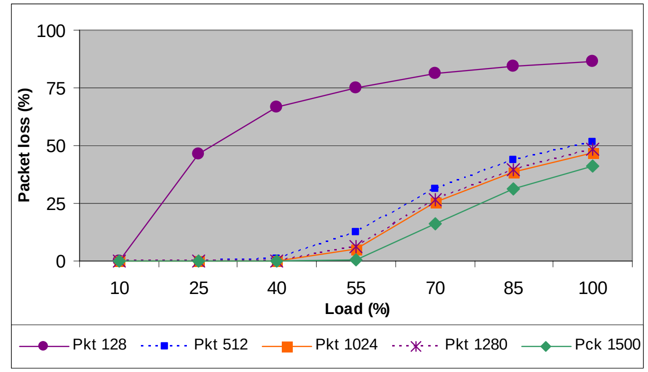

Figure 8: Packet loss for different packet sizes

An interesting observation derived from Figure 8 is that packet loss is almost the same for traffic consisting of 512- up to 1500-byte packets.

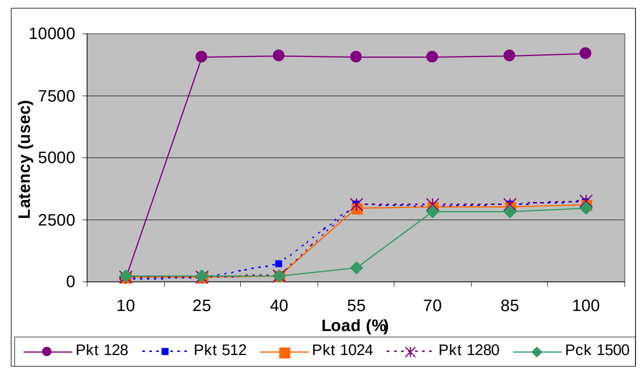

Figure 9: Latency for different packet sizes

In latency measurements, Figure 9, we noticed very small values under zero packet loss. When there is packet loss, latency remains constant, which is probably what a packet experiences while entering a full buffer.

5) Latency and packet loss for PIP traffic

The next set of tests investigated the latency and packet loss for Premium IP (PIP) traffic, to assess the capability of protecting high priority data. Priority queue was enabled in the output interfaces. Bidirectional IPv6 and IPv4 flows were established at the traffic generator in Ilissos (Figure 2) and traffic was switched only by the local router. 2% of the traffic was PIP and the rest was BE. Traffic load was gradually increased from 10% to 100% of the port capacity in steps of 15%.

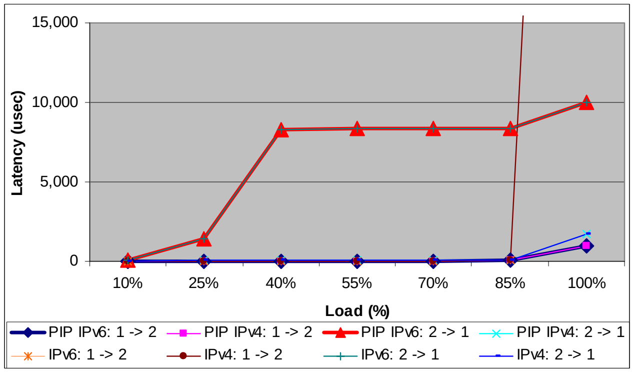

Figure 10: Packet loss for PIP traffic

The packet loss for Premium IPv6 traffic is always zero in *Tetra* cards. On the contrary, in *Tango* cards the Premium IPv6 traffic experienced the same performance as BE traffic. Therefore, packet loss reached up to 72%. Obviously, the IPv6 *traffic class* field in the IPv6 header is not recognized or ignored in the *Tango* card and, thus, classifications mechanism fails to separate traffic into different priorities.

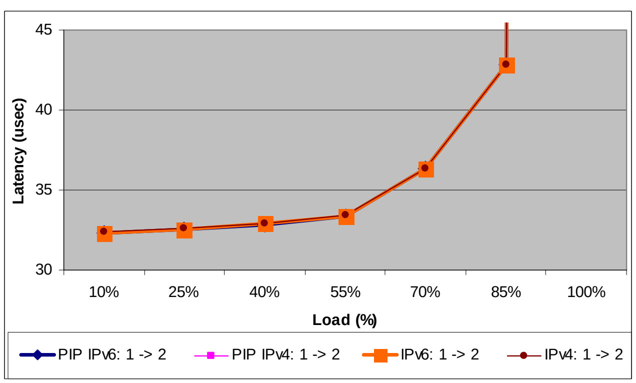

Figure 11: PIP latency in the *Tetra* cards

In *Tetra* cards, latency for PIP traffic and BE traffic is the same provided that there is no packet loss (<85% load) (see Figure 11). When there is packet loss (100% load), PIP sharply increases (~20 times) but still remains ~100 times smaller than BE latency. In *Tango*, PIP the latency is at least ~100 times higher than PIP latency in Tetra (100% load).

IV. 6NET NETWORK

6NET [9] was one of the largest IPv6 research projects funded by the European Commission under the Information Society Technologies (IST) Programme. The project consortium consisted of several partners from industry, European National Research & Education Networks, Universities and Research Institutes. The 6NET network was designed to become a native *IPv6-only* environment for testing new protocols, services and applications and, thus, there were no limitations imposed by existing IPv4 protocols or *IPv6overIPv4* tunnels.

The core network, as shown in Figure 12, extended over several European countries. It consisted of STM-1/PoS core links while the access link speeds and technologies varied; STM-1/PoS, Gigabit Ethernet, ATM or MPLS L2 tunnels, 2 Mbps E1 serial circuits, etc. In the core and access network hardware-based Cisco 12400 and software-based 7200VXR series routers were installed.

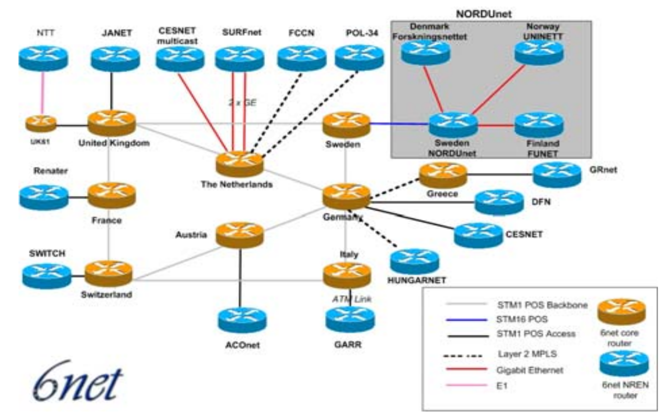

Figure 12: The 6NET network

The 6NET core network supported Differentiated Services (DiffServ) with three classes - *Premium IP (PIP)*, *Best Effort (BE)* and *Less than Best Effort*

(LBE). The implemented QoS schema took into account several aspects related to network dimensioning and resource management, similar to the one described in the previous section for the GRNET network.

A. QoS tests in IPv6-only environment

The 6NET testbed consisted of three dedicated PC-based servers connected to Greece, the United Kingdom and the Netherlands. The servers generated traffic with *iperf* [13] and *mgen* [14] tools and produced throughput, packet loss and jitter statistics. High priority (foreground) traffic was forwarded from Greece towards the UK while low priority (background) traffic generated in the Netherlands caused congestion to the core links towards the UK.

A complete QoS schema was deployed in the 6NET network [12]. Appropriate classification, policing and queuing mechanisms were enabled in core and access routers that allowed up to 5% of the link capacity to be occupied by PIP traffic. The tests allowed us to evaluate promised guarantees to high priority IPv6 traffic in a congested environment, as compared to BE traffic.

A small subset of traffic patterns used in 6NET QoS tests are given in Table 1. Each test was performed with UDP foreground traffic while the background traffic consisted of mixture of TCP (30%) and UDP (70%) traffic. In scenario 1 the network congestion was limited, while in scenario 2, severe congestion was experienced in the access link[2] of the UK server leading to high packet losses.

Scenario	Best Effort (Mbps)	IP Premium (Mbps)
1	80	UDP 1.5*
2	120	UDP 1.5*
* Traffic is increased in steps of ~0.5Mbps.		

Table 1: Testing scenarios

As shown in Figure 13, PIP traffic experienced approximately zero packet loss with transmitting rates up to 7 Mbps. As soon as PIP traffic exceeded the allocated bandwidth, i.e. 5% of the total bandwidth or approximately 7 Mbps (at the IP layer), packet losses sharply increased. On the contrary, the packet loss for BE traffic was extremely high under congestion conditions (scenario 2). Our results verified the effectiveness of the classification and policing mechanisms applied at the input interface of the access routers.

[2] The server in UK was connected via a 100Mbps Fast Ethernet interface, unable to handle the full traffic load.

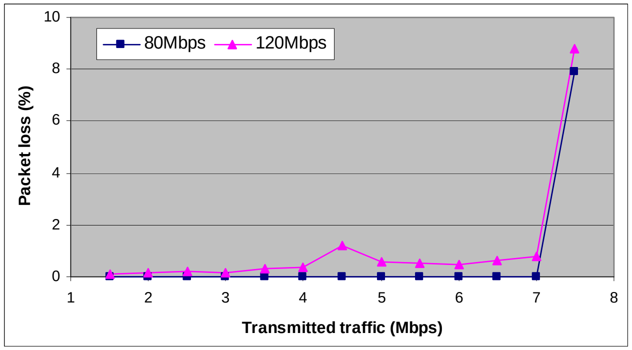

Figure 13: Packet loss for Premium IP traffic

As observed in Figure 14, jitter experienced by Premium IP traffic was the same under different levels of congestion, i.e. scenarios 1 and 2. These results verified that PIP traffic –serviced via the priority queue- was not affected from background BE traffic. In the same figure, it is interesting to observe that jitter is reduced as PIP rate increased. This can be explained by the fact that a higher transmission rate leads to smaller inter-packet delays. As the PIP traffic was served by priority queues in the network, variations in the inter-arrival time decreased.

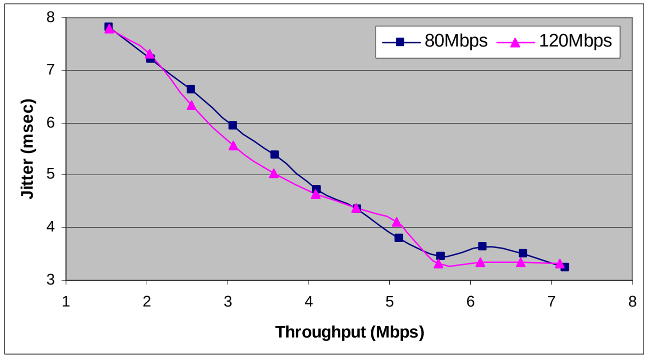

Figure 14: Jitter for Premium IP traffic

V. "WISH TO HAVE" LIST

While performing the tests in the GRNET and 6NET networks, we identified *"wish-to-have"* functions missing from the routers under test. Although command line interfaces (CLI) for IPv6 and IPv4 traffic were identical, allowing us to create a common QoS configuration template for both protocols, some commands were either not supported for IPv6 traffic or different commands existed for IPv6 and IPv4. Secondly, router statistics on interface[3] level do not differentiate IPv6 and IPv4 packets and, thus, it is not easy to count the number of IPv6 packets in a dual stack environment. A work-around solution is to use different sub-interfaces (VLANs) for IPv6 traffic and apply hierarchical QoS policies (per sub-interface). However, this approach exhibits increased management complexity and also requires enhanced functionality (e.g. *Tetra* cards) to be supported in the access ports. Thirdly, it was identified that monitoring functionality for IPv6 traffic was missing. *Service Assurance Agents* - SAAs [19] could not generate IPv6 monitoring packets and, thus, IPv6 performance statistics could not be collected via the routers. A work-around solution would be to use *IPv6overIPv4* tunnels but the accuracy of collected monitoring data would be coarse, as tunnelled packets follow the processing switching path (switched by the router CPU). Finally, it should be noted that during the tests we were able to use advanced hardware (e.g. *Tetra* cards) and the latest versions of the routers operating systems. Obviously, older version hardware lacks IPv6 forwarding capabilities and previous versions of operating systems do not exhibit rich functionality to handle IPv6 traffic. Such hardware and software is quite often deployed in production networks, thus explaining the reluctance of some network providers to migrate to IPv6.

VI. CONCLUSIONS - FURTHER WORK

The performed QoS tests in GRNET and 6NET core networks indicated that gigabit routers under test adequately support QoS mechanisms for IPv6 traffic. Especially in newer router line cards, i.e. *Tetra* GigE cards, performance guarantees achieved for IPv6 and IPv4 traffic were identical. On the contrary, in older *Tango* GigE cards, IPv6 is software-switched and experiences worst performance than its IPv4 counterpart that is hardware-switched. Similar qualitative tests in 6NET network, revealed that performance guarantees can be smoothly provided to high priority traffic in an IPv6-only environment. However, when handling IPv6 traffic under extreme line card congestion, both the *Tango* and the *Tetra* cards had a negative impact on routing protocols, perhaps due to current queue management implementations.

As the portion of IPv6 traffic is currently significantly low compared to IPv4 traffic, an IPv6 QoS schema can be deployed in research or production networks at Gigabit speeds, albeit some limitations of older routing equipment in use. GRNET, based on the results of the tests reported above and the 6NET experience, is expanding provision of Premium IP service for IPv6 in its dual stack gigabit core network.

[3] Counters in classification mechanisms also do not distinguish IPv6 and IPv4 packets.

Acknowledgments

The GRNET IPv6 testing activity is partially funded by the Greek Ministry of Economy, Operational Programme for the Information Society, under the "Development of the Greek Research and Technology Network - GRNET-2" project. The 6NET project (IST-2001-32603) is partially funded by the European Commission under the Information Society Technologies (IST) Programme of FP5. The paper contents do not necessarily reflect the opinions of partners in the 6NET Project

VII. REFERENCES

[1] "GÉANT", IST Programme, http://www.geant.net.

[2] T. Chown; IPv6 Campus Transition Experiences; Proceedings of the 2005 International Symposium on Applications and the Internet (SAINT2005), Trento, Italy, 31 January - 4 February 2005.

[3] 6NET Deliverable D2.3.4v2, "IPv4 to IPv6 transition Cookbook for End site networks / universities", June 2005.

[4] E. Davies, "IPv6 Transition/Co-existence Security Considerations", draft-ietf-v6ops-security-overview-00.txt, work in progress.

[5] "Mobility for IPv6 (mip6)", IETF WG, http://www.ietf.org/html.charters/mip6-charter.html

[6] "6POWER project" (IPv6, QoS & Power Line Integration) IST programme (http://www.6power.org/)"

[7] S.Deering, R.Hinden, "Intenet Protocol, Version 6 (IPv6)", RFC2460, December 1998.

[8] J.Rayahalme et. al., "IPv6 Flow Label Specification", RFC3697, March 2004.

[9] "Greek Research & Technology network - GRNET", http://www.grnet.gr.

[10] S. Blake et al, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[11] "6NET – Large Scale IPv6 Pilot Network", IST Programme (IST-2001-32603), http://www.6net.org.

[12] 6NET Deliverable D.4.4.2v2, "Report in QoS Tests, 2nd version", 6NET project (IST-2001-32603), April 2005.

[13] IPERF tool, http://dast.nlanr.net/Projects/Iperf/

[14] Multi-Generator Toolset (MGEN), http://mgen.pf.itd.nrl.navy.mil/.

[15] Cisco 12000 Series Routers, http://www.cisco.com.

[16] Cisco 12000 Series 4-Port Gigabit Ethernet ISE Line Card, http://www.cisco.com.

[17] Cisco 12000 10-Port Gigabit Ethernet Line Card, http://www.cisco.com.

[18] Spirent SmartBits 600 Series traffic generators, http://www.spirent.com/.

[19] Service Assurance Agent (SAA), http://www.cisco.com.

Author

Dipl.-Ing. Athanassios Liakopoulos received the Dipl.-Ing. degree in Electrical and Computer Engineering from the National Technical University of Athens (NTUA), in 1996 and MSc with Distinction in Telematics (Telecommunications & Computer Engineering) from the Electrical Engineering Department in University of Surrey (UniS) in 1998. Since 2000, he joined the GRNET S.A. and participated in several national and European IST research projects, such as SEQUIN, 6NET, SEEREN, GN2. He is also pursuing a PhD from NTUA in the field of QoS provisioning at

high-speed networks. He has awarded for his performance during his academic studies and has published articles in recognized technical journals. He is member of TEE and IEEE.

Dr. Dimitrios Kalogeras received his Dipl.-Ing. Degree and Ph.D from the Department of Electrical and Computer Engineering, National Technical University of Athens (NTUA), Greece in 1991 and in 1996 respectively. Dr Kalogeras has participated in numerous Greek and European research programs. He has pioneered in the design and development of the NTUA and the GRNET data networks and is a member of the technical and scientific committee of the Greek School Network. From 2000 to 2002 served as a member of the TERENA Technical Committee. He is member of TEE and IEEE.

Vasilis Maglaris is a Professor of Electrical & Computer Engineering at the National Technical University of Athens (NTUA) since 1989. He holds an Engineering Degree from NTUA (Athens, Greece, 1974) and a Ph.D. degree from Columbia University (New York, USA, 1979). Before joining the faculty at NTUA, he held industrial and academic positions in the USA for ten years, all in advanced electronic communications. Apart from teaching and performing research on Computer Networks, he was responsible for developing the NTUA Campus LAN and for the establishment of GRNET (the Greek NREN). He served as GRNET's Chairman from its inception (1995) until June 2004. He served on the board of the Greek National Regulatory Authority on Telecommunications and Posts for two five-year terms (1995 – 2005). He authored more than 100 research papers and regularly delivers lectures on Internet advances.
From October 2004, he serves as the Chairman of the National Research & Education Networks Policy Committee (NREN PC). The NREN PC harmonizes policies amongst the 30 NRENs in the extended European Research Area; it is also responsible for the Pan-European advanced network platform GEANT and its current upgrade GEANT2, the Next Generation Internet initiative of the European research & academic community.

Dimitris Primpas obtained his diploma from Computer Engineering and Informatics Department of the Polytechnic School of the University of Patras on November 2002. Next, he was accepted in the postgraduate program “Computer Science and technology” in the same department and on April 2004 he obtained his Master Degree. Now he continues the postgraduate studies in the same department to receive PhD. He works in the Research Unit 6 of Research Academic Computer Technology Institute and on Telematics, Distributed Systems and Basic Services Laboratory of the Computer Engineering & Informatics Department, in Patras University. His interests include: networks, protocols, Quality of Service, Managed bandwidth services and Network applications. He had worked on GRNET- IP-MBS, GRNET IP-QoS, 6NET (LARGE-SCALE INTERNATIONAL IPv6 TESTBED – IST-2001-32603), ASP-NG (ASP-New Game - IST-2001-35354) project and currently works on Advanced Services GRNET II/VNOC 2 and GN2 (Multi-Gigabit European Academic Network) projects. He has published 3 articles in Journals, 15 papers in well-known refereed conferences and he is co-author of a book in greek language with title “Advanced Issues on Computer Networks and Services”. He is member of TEE

Christos Bouras obtained his Diploma and PhD from the Computer Science and Engineering Department of Patras University (Greece). He is currently an Associate Professor in the above department. Also he is a scientific advisor of Research Unit 6 in Research Academic Computer Technology Institute (CTI), Patras, Greece. His research interests include Analysis of Performance of Networking and Computer Systems, Computer Networks and Protocols, Telematics and New Services, QoS and Pricing for Networks and Services, e – learning, Networked Virtual Environments and WWW Issues. He has extended professional experience in Design and Analysis of Networks, Protocols, Telematics and New Services. He has published 150 papers in various well-known refereed conferences and journals. He is a co-author of five books in Greek. He has been a PC member and referee in various international journals and conferences. He has participated in R&D projects such as RACE, ESPRIT, TELEMATICS, EDUCATIONAL MULTIMEDIA, ISPO, EMPLOYMENT, ADAPT, STRIDE, EUROFORM, IST, GROWTH, 6NET, GRNET/VNOC-2, GN-2 and others. Also he is member of, experts in the Greek Research and Technology Network (GRNET), Advisory Committee Member to the World Wide Web Consortium (W3C), Task Force for Broadband Access in Greece, ACM, IEEE, EDEN, AACE and New York Academy of Sciences.