INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT
Int. J. Network Mgmt 2009; **19**: 119–137
Published online 6 June 2008 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/nem.695

QoS experiences in native IPv6 networks

Athanassios Liakopoulos[1,2], Dimitrios Kalogeras[1,2], Vasilis Maglaris[2], Dimitris Primpas[3,4] and Christos Bouras[3,4,*†]

[1]*Greek Research and Technology Network, Athens, Greece*
[2]*Department of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece*
[3]*Research Academic Computer Technology Institute, University of Patras Campus, Patras, Greece*
[4]*Department of Computer Engineering and Informatics, University of Patras, Patras, Greece*

SUMMARY

Deployment of IPv6 technology in research and commercial networks has accelerated in the last few years. Inevitably, as more advanced services take advantage of the new technology, IPv6 traffic gradually increases. Today, there is limited experience in the deployment of Quality of Service (QoS) for IPv6 traffic in backbone networks that support the Differentiated Services framework. As available software and hardware are designed to handle IPv4 packets, there is a need to accurately measure and validate performance of QoS mechanisms in an IPv6 environment. This paper discusses tests and technical challenges in the deployment of IPv6 QoS in core networks, namely the production dual stack gigabit-speed Greek Research and Education Network (GRNET) and the IPv6-only 6NET European test network, 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 to 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.

1. INTRODUCTION

Internet Protocol version 4 (IPv4) is the basis of the Internet, allowing millions of hosts to communicate over diverse networks. Even though IPv4 has been immensely successful, the continuous growth of the global Internet requires that the overall architecture evolves. Internet Protocol version 6 (IPv6) [1] has been developed by the IETF in order to enhance the Internet architecture and overcome the limitations imposed by IPv4. Therefore, IPv6 is designed to support increased numbers of users, new applications and services and to make ongoing Internet expansion possible. It provides practically unlimited address space, built-in mobility and security support, easy configuration of end systems, enhanced multicast features, etc. It is envisaged that IPv6 will allow the deployment of a *ubiquity* network, where end-users will be connected at *any* time, at *any* location, with *any* device with the global Internet.

In such a new environment, the provision of Quality of Service (QoS) is a basic requirement of multi-service networks that support multimedia and virtual collaboration applications. By enabling QoS, an Internet provider can guarantee to its subscribers the performance of the transport services over its network. Today, most of the National Research and Education Networks (NRENs) in Europe support QoS services using diverse technologies, while GÉANT [2], the trans-European research network, offers

*Correspondence to: Christos Bouras, Research Academic Computer Technology Institute, N. Kazantzaki Str., University of Patras Campus, 26500 Patras, Greece.
†E-mail: bouras@cti.gr

Received 10 November 2006
Revised 24 April 2007
Accepted 23 February 2008

a high-priority transport service for transit traffic. In addition, the aforementioned networks have recently included native IPv6 interconnection services to their portfolio. As the IPv6 traffic is gradually increased, network engineers have to validate the QoS techniques currently deployed in their dual stack networks.

For quite a long time, technical forums have been holding discussions about QoS support in IPv6 environments. There is a debate on whether 'IPv6 provides better QoS support than IPv4' and whether 'IPv6 experiences worst performance than IPv4'. In this context, the objectives of our work is twofold: to validate the performance of basic QoS mechanisms with IPv6 traffic on hardware- and software-based platforms and to 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 some conditions are fulfilled.

The paper is organized as follows: Section 2 provides background information and presents the QoS-related fields in the IPv6 header. Section 3 presents the GRNET network and elaborates the current deployed (IPv4) QoS schema. Section 4 analyses the results from performance tests with a mixture of IPv6/v4 traffic on a test-bed that uses GRNET's dual-stack production network. Section 5 is dedicated to qualitative tests conducted in a large-scale network (the 6NET IPv6-only core network). Section 6 presents 'wish-to-have' functionality, while Sections 7 and 8, respectively, summarize our conclusions and define our future plans.

2. QoS BACKGROUND

The Differentiated Services (DiffServ) framework, specified by the IETF [3], is widely deployed in today's production networks as it does not exhibit scalability limitations at high-speed interconnection links. DiffServ treats individual flows with similar quality needs as traffic aggregates and identifies a limited number of service classes to which traffic aggregates are associated. Therefore, DiffServ is designed to provide performance guarantees to traffic aggregates, ignoring the level of services provided to individual flows. The DiffServ framework defines only the basic 'tools' for implementing advanced transport services and, thus, it does not define how these services may be realized. For example, the DiffServ framework does not define how admission control is performed, how application traffic is colored in to different service classes, how QoS mechanisms in the network are configured, etc.

Several research groups have proposed alternative methodologies for providing performance guarantees to end-to-end flows over DiffServ networks. The SEQUIN project [4] designed and implemented a high-priority service, called Premium IP (PIP), by aligning QoS provisioning procedures among the European NRENs and GÉANT network. The Resource Management in DiffServ (RMD) framework [5,6], proposed by the Next Steps in Signaling (NSIS) Working Group of IETF, extends the DiffServ framework with signaling in the control plane for managing network resources. Other architectures [7–9] rely on centralized entities, called Bandwidth Brokers (BB), for performing admission control and configuring edge devices of a network domain using out-of-band signaling. As all the above proposals are based on the DiffServ framework, common mechanisms, such as traffic classification, policing and queuing, are enabled at the core routers.

Even if the transport services offered to subscribers' traffic may differ in each administrative domain, the performance guarantees are assessed using the same set of performance metrics. The IP Performance Metrics (IPPM) Working Group of IETF has defined *one-way delay* [10], which assesses the time interval for delivering a packet from a source to the corresponding destination, mainly composed by transmission and queuing delays; *inter-packet delay variation* or *jitter* [11], which assesses variations of one-way delay due to statistical multiplexing in the outgoing queues; *packet loss* [12], which assesses the portion of packets lost due to buffer exhaustion; and *bandwidth* [13], which assesses the maximum amount of data transported in a unit of time.

DOI: 10.1002/nem

2.1 QoS-related fields in the IPv6 header

The IPv6 header [1] is (re)designed to minimize header overhead and reduce the header process for the majority of 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, the IPv6 header is not a superset—and thus backward compatible—with its 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 five 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 the 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 [14] 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 the IPv6 protocol, in terms of QoS functionality, is neither superior nor inferior to its 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.

3. GRNET CORE NETWORK

The Greek National Research and Educational Network (GRNET) (Figure 1) [15] interconnects approximately 90 universities and research institutes. The core network consists of 12 nodes interconnected with STM-16 *lambdas*, while the subscriber access links vary from 1 Gbps down to 2 Mbps. GRNET currently supports native IPv6 interconnection services. Its core routers are Cisco GSR12400 series with 4 × GE (Cisco 12000 Series 4-port Gigabit Ethernet ISE) and 10 × GE (Cisco 12000 10-port Gigabit Ethernet) line cards [16]. Their 10 × GE (*Eng4+*) cards, also called *Tango*, are mainly used in core links and support (IPv4) line rate switching capabilities. On the contrary, their 4 × GE (*Eng3*) cards, also called *Tetra*, are used in access links, providing 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 switch traffic in software. In addition, GRNET uses the Cisco 7200 series platform to connect and aggregate low-bandwidth connections (mainly 2–100 Mbps). The Cisco 7200 series switch IPv6 traffic in software too.

3.1 QoS model and services

GRNET uses DiffServ 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 Per Hop Behavior (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 appropriate for elastic applications.
- *Less than Best Effort (LBE)* exploits network resources without (negative) impact on other traffic classes. It is suited for specific scavenger applications.

DOI: 10.1002/nem

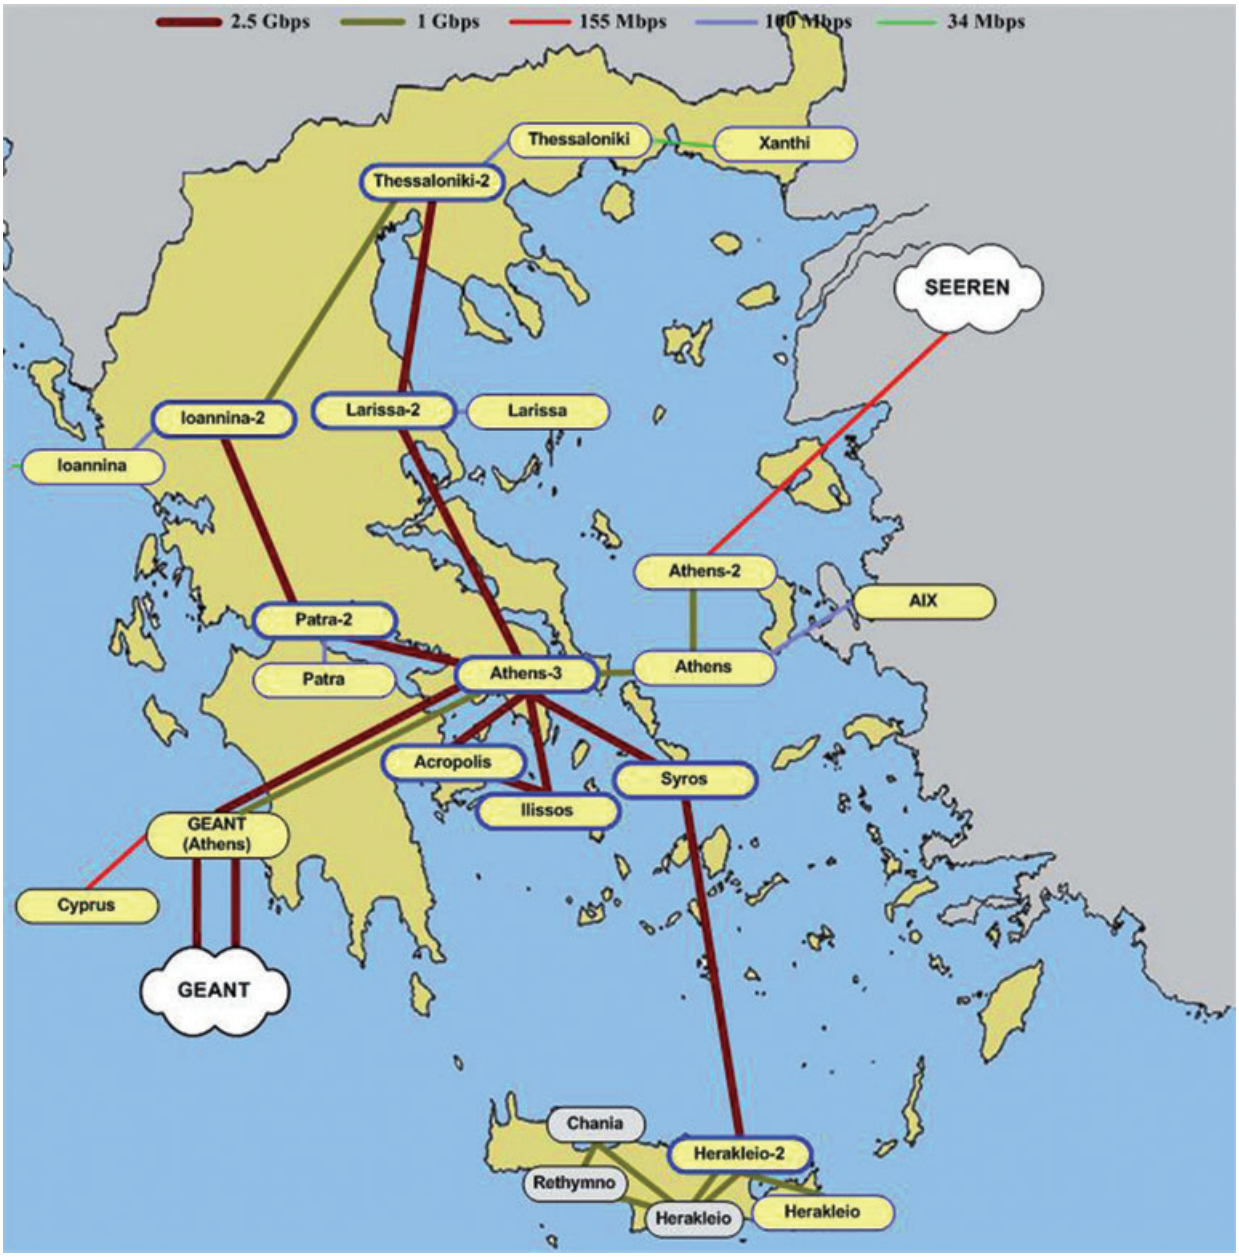

Figure 1. GRNET backbone network (November 2005). A colour version of this figure is available online at www.interscience.wiley.com

The Premium IP class is further divided into three sub-classes: PIP Virtual Wire, PIP for VoIP and PIP Transparent. PIP Virtual Wire is used for traffic exchanged between two well-identified access interfaces and emulates a virtual circuit. Premium IP for 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 20% of the link capacity in order to minimize inter-packet

DOI: 10.1002/nem

delay variation (jitter) and avoid starvation of lower-priority traffic. An automatic provisioning tool is used for performing the admission control and generating the appropriate router configuration [16]. 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 marked with DSCP values 46, 47 and 40, while LBE traffic is marked with DSCP value 8.

3.2 Testing equipment and methodology

The GRNET tests were conducted using hardware-based traffic generators Smartbit 600 [17] with Gigabit Ethernet (GigE) interfaces attached in three different PoPs of the network (see Figure 2). They were connected either directly to 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.

The traffic generators were able to produce in GigE ports a mix of IPv6 and IPv4 traffic up to 1 Gbps. 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. 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.

The testing plan that we followed is split into four main steps.

- Basic tests that will investigate and prove the proper setup of the dual-stack network and the correct operation of basic mechanisms such as IPv6 access list classification and shaping.
- Performance tests that will investigate the impact of IPv6 traffic (in several amounts) on the router's performance in the production network. Also, tests will focus on measuring the packet loss and the

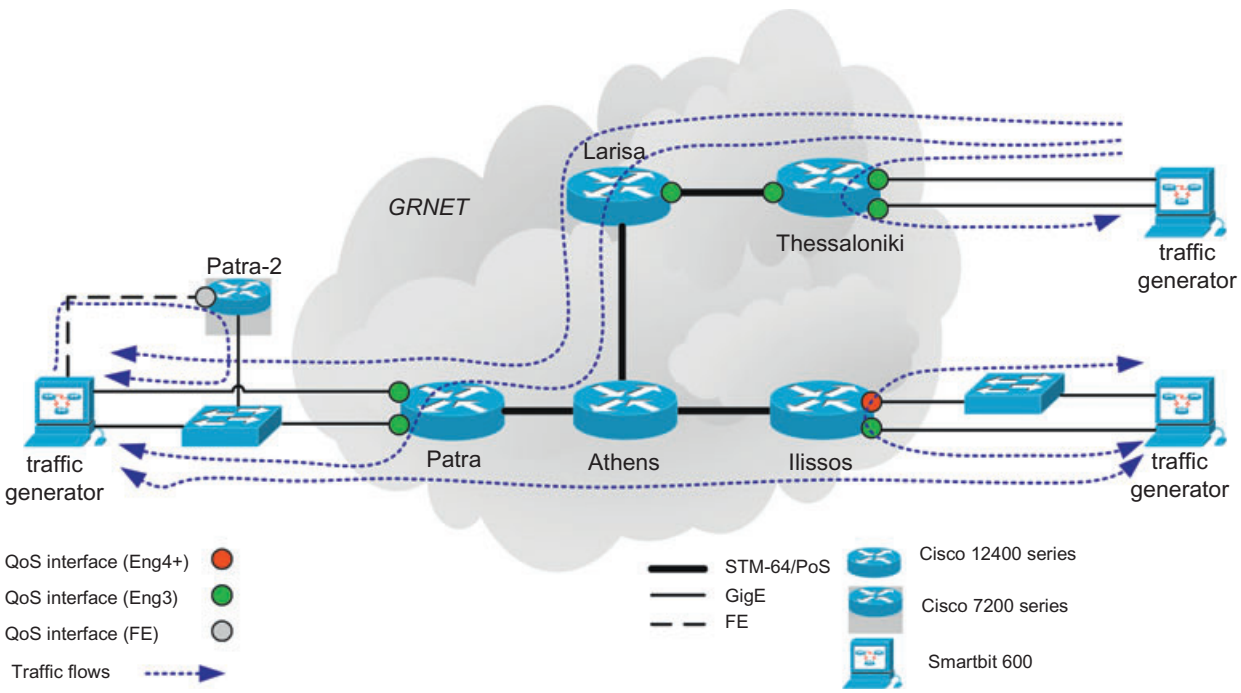

Figure 2. GRNET test-bed topology. A colour version of this figure is available online at www.interscience.wiley.com

DOI: 10.1002/nem

latency of IPv6 traffic and comparison with IPv4, evaluating the IPv6 switching capabilities of our equipment.
- QoS-related tests that will provide results concerning the current implementation of mechanisms for IPv6 QoS (like queuing), their maturity and possible drawbacks. At this point we should notice that all the above tests are conducted on GRNET's production network and therefore possible effects on other production services will also be monitored.
- Finally, our plan is to extend the tests on a large scale (IPv6-only network) in order to measure IPv6 QoS performance and also monitor possible differences (if any) between dual-stack and IPv6-only networks in terms of IPv6 features and performance.

4. IPV6 QOS TESTING IN GRNET

4.1 Basic tests for managing IPv6 traffic

The first set of tests focused on classification mechanisms and access lists at the GigE interfaces. A traffic generator produced 100 Mbps 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 a 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 an average rate of 400 Mbps was shaped at a rate of 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, at approximately 443 ms, 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 μs). As expected, shaping increased the average latency of packets significantly. Generally, those basic tests proved the setup and the operation of the basic mechanisms (classification, policing, shaping) that are necessary in order to provide QoS services.

4.2 Performance tests

4.2.1 Investigation of CPU load

The next set of tests investigated the impact on CPU load of IPv6 traffic switching at gigabit speeds. Bidirectional flows (1 Gbps in each direction) 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. All the scenarios and the results are described in Table 1. More analytically, an initial test with IPv4-only traffic at line rate speed for a 30 min period did not increase the CPU load of the router, which remained approximately at 11%. We repeated the test with a mixture of IPv6 and IPv4 traffic in equal portions and noticed a small increase of 8% absolute value in the CPU load at 1 and 5 min time intervals. The test was repeated with IPv6-only traffic. This time the CPU increased by 11% for 5 min intervals and by 26% for 1 min intervals. Note that an IS-IS routing problem affected production services during the test. The same tests were repeated 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. Therefore, the overall conclusion was that due to the fact that Tetra interfaces support hardware switching for both IPv4 and IPv6 protocols, IPv6 switching is efficient and does not cause any increase in CPU load. On the other hand, in routers that use Tango cards, which switch Ipv6 traffic in software, the existence of IPv6 traffic has a significant effect on CPU load.

DOI: 10.1002/nem

Scenario	description	Average CPU load
1	Normal production traffic (almost 500 Mbps IPv4)	11%
2	Normal production traffic + bidirectional flow (1 Gbps each of IPv4 traffic) with time interval 1 min	11%
3	Normal production traffic + bidirectional flow (1 Gbps each of IPv4 traffic) with time interval 5 min	11%
4	Normal production traffic + bidirectional flow (1 Gbps each of IPv4 and IPv6 in equal portion) with time interval 1 min	19%
5	Normal production traffic + bidirectional flow (1 Gbps each of IPv4 and IPv6 in equal portion) with time interval 1 min	19%
6	Normal production traffic + bidirectional flow (1 Gbps each of IPv6 traffic) with time interval 1 min	37%
7	Normal production traffic + bidirectional flow (1 Gbps each of IPv6 traffic) with time interval 5 min	22%

Table 1. Investigation of CPU load effect by IPv6 traffic

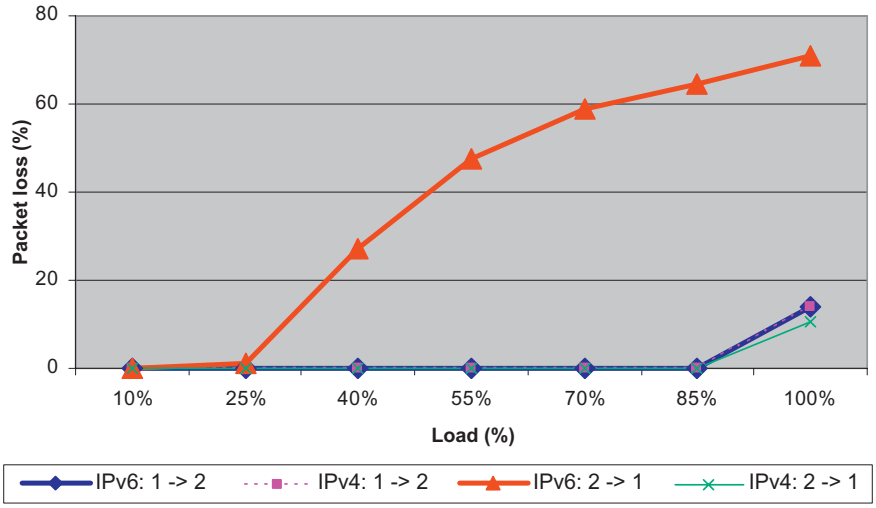

Figure 3. Packet loss for BE traffic. A colour version of this figure is available online at www.interscience.wiley.com

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 to CPU load and the problems were caused by the fact that IPv6 control traffic is not protected in internal CPU queues (at the router's architecture), unlike the IPv4 case.

4.2.2 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 at the Ilissos router. 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% (1 Gbps in each direction) of the port capacity in steps of 15%. Figure 3 presents the results

DOI: 10.1002/nem

regarding the packet loss, where we notice a serious packet loss in IPv6 traffic that entered the Tango card. In particular, packet loss for IPv6 traffic entering the Tango card (direction 2 → 1) is much higher than the packet loss in the opposite direction (1 → 2) and also much higher than the corresponding Ipv4 traffic (direction 2 → 1). 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. Finally, non-zero packet loss (13.88%) is noticed for a 100% utilization in the Tetra case and is explained by the Tetra's architecture, in which the real switching capacity per port is 850 Mbps.

Regarding the experienced average latency, Figure 4 presents the overall results. In detail, IPv6 and IPv4 traffic experience the same latency in Tetra cards (direction 1 → 2). Even when there is packet loss (100% utilization), latency is increased equally for both protocols. On the contrary, latency for IPv6 traffic in Tango cards is twice as high, even with no packet loss. When there is IPv6 packet loss (but no IPv4 packet loss), the difference is increased ~200 times.

4.2.3 Latency and packet loss for BE traffic for different packet sizes

All sets of tests were performed with 128-byte packets, which is the worst case for a router. As real traffic consists of packets of diverse sizes, we repeated the tests with Tango cards, which had serious problems in efficiency of Ipv6 traffic switching when using larger packet sizes. In particular, unidirectional flow of IPv6-only traffic (1 Gbps) was established, using various packet sizes (128, 512, 1024, 1280, 1500 bytes) and we measured the achieved throughput. As Figure 5 shows, the throughput increased for larger packet sizes; however, even for 1500-byte packets the loss is slightly more than 50%.

The same scenario (IPv6 flow in Tango card) was repeated, while the traffic load was now gradually increased from 100 Mbps to 1 Gbps in steps of 150 Mbps. In this case, we measured the packet loss and the latency for all the different cases of packet size (128, 512, 1024, 1280, 1500 bytes). Figures 6 and 7 present the results. An interesting observation derived from Figure 6 is that packet loss showed almost the same pattern for traffic consisting of 512-byte up to 1500-byte packets.

In latency measurements (Figure 7), 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.

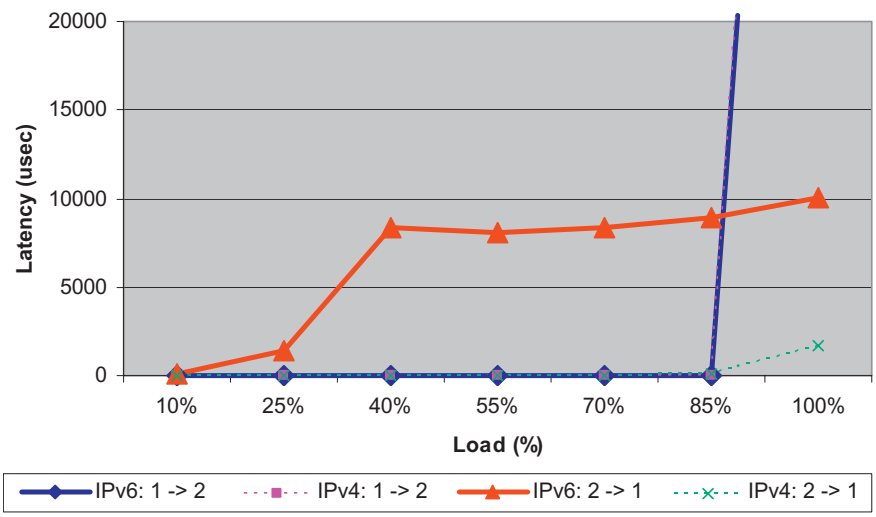

Figure 4. Latency for BE traffic. A colour version of this figure is available online at www.interscience.wiley.com

DOI: 10.1002/nem

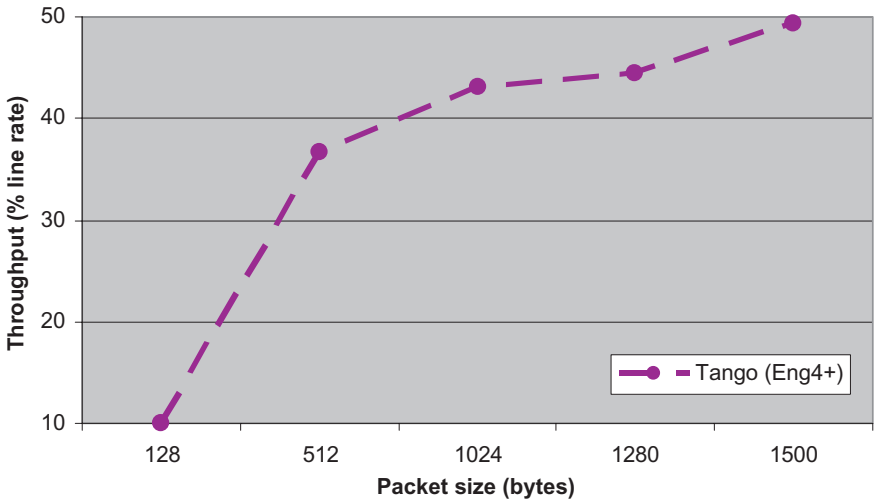

Figure 5. Throughput versus packet size in Tango. A colour version of this figure is available online at www.interscience.wiley.com

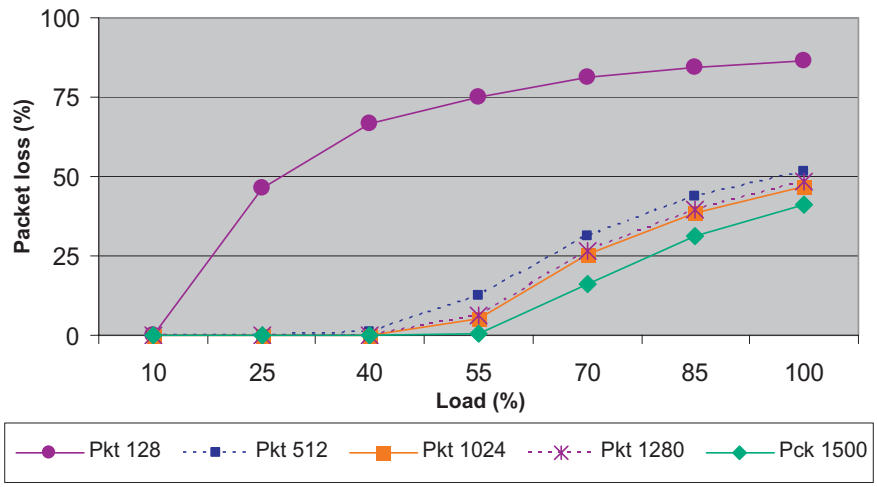

Figure 6. Packet loss for different packet sizes. A colour version of this figure is available online at www.interscience.wiley.com

Generally, the performance tests showed that the Tango cards, which are old and are used in some routers of GRNET's production network, are not capable of switching IPv6 traffic as efficiently as IPv4, which was partly expected as IPv6 traffic is software switched while IPv4 is hardware based. Conversely, the Tetra cards, which are used in all routers of GRNET's production network, provide similar services to both protocols. Provided that today the portion of IPv6 traffic in the GRNET network is approximately 2% 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. Finally, the routing problems that were noticed in heavy congestion using IPv6 is expected to be solved by the vendors in patches to the router's operating systems.

DOI: 10.1002/nem

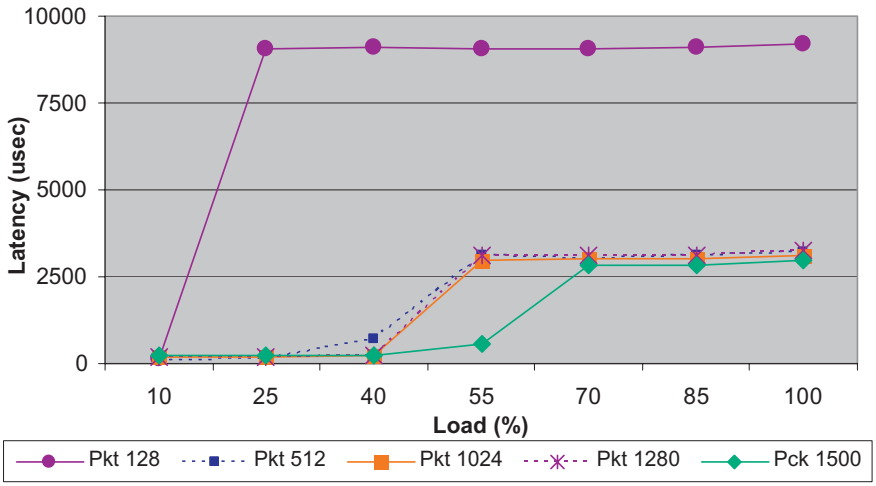

Figure 7. Latency for different packet sizes. A colour version of this figure is available online at www.interscience.wiley.com

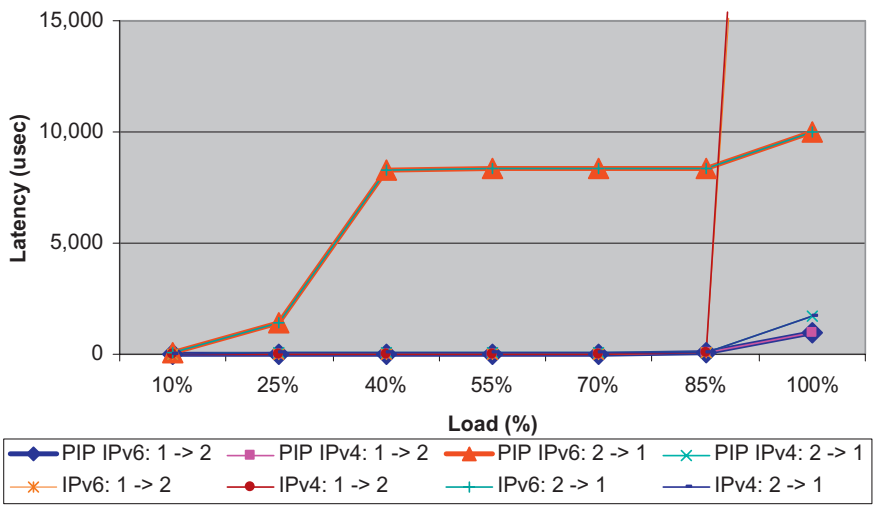

Figure 8. Packet loss for PIP and BE traffic. A colour version of this figure is available online at www.interscience.wiley.com

4.3 IPv6 QoS tests

4.3.1 Latency and packet loss for PIP traffic

The next set of tests investigated the latency and packet loss for PIP traffic, to assess the capability of protecting high-priority data. The priority queue was enabled in the output interfaces, and enqueuing packets were marked with DSCP value 46 (both IPv4 and IPv6). Bidirectional IPv6 and IPv4 flows were established at the traffic generator in Ilissos (Figure 2) and traffic was switched only by the local router. 20% 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%.

Regarding packet loss, the results are presented in Figure 8, where the packet loss for PIP traffic alongside Best Effort traffic is shown. In detail, the packet loss for Premium IPv6 traffic is always zero in Tetra cards, whereas 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

DOI: 10.1002/nem

ignored in the Tango card and, thus IPv6 traffic fails to be enqueued in a queue other than the normal Best Effort one. The latter means that, using the current software (operating system) in the router, GRNET cannot support IPv6 QoS services to customers connected by Tango cards.

Regarding latency, in Tetra cards latency for PIP and BE traffic is the same, provided there is no packet loss (<85% load) (see Figure 9). When there is packet loss (100% load), PIP sharply increases (~20 times) but still remains ~100 times smaller than BE latency. In Tango, PIP latency is at least ~100 times higher than PIP latency in Tetra (100% load).

4.3.2 QoS tests on a software-based platform

The next set of tests focused on full software-based platforms and in particular using a Cisco 7200 series router. This router is used in GRNET in order to aggregate low-bandwidth connections and therefore we extended the tests (see Figure 2) in this platform too. The scenarios was almost the same; thus bidirectional flows of IPv4, IPv6, and a mix of IPv4/6 traffic were established using the traffic generators that were connected on a Gigabit Ethernet port and a Fast Ethernet port, respectively. In each set of tests the traffic load was gradually increased from 70 Mbps to 130 Mbps, while the packet size was set to 512 bytes. In half of the tests, the generated traffic caused severe congestion to the Fast Ethernet interface. The average CPU load, measured during the last 5 s of each test, is shown in Figure 10. The load is increased by 5–9% when IPv6 traffic is present. When priority queuing was activated for handling Premium IPv4/6 traffic, the CPU load was further increased, reaching up to 100%.

The next set of tests assessed the switching capacity of the router. Generated IPv4, IPv6 and a mixture of IPv4/6 traffic was produced using the traffic generator. Consecutive tests of 10 s were performed with traffic rates of 85 Mbps, 100 Mbps and 115 Mbps. The packet sizes also varied from 128 bytes up to 512 bytes. The loss ratio for packet sizes greater than 128 bytes was the same for IPv4 or IPv6 traffic, while the router achieved switching of incoming packets at the maximum theoretical rate through the Fast Ethernet interface, as shown in Table 2. When the packet size was set to 128 bytes, the router exhibited high packet loss under congestion conditions, especially when IPv6 traffic was present.

The next set of tests investigated multiple QoS mechanisms, such as classification, marking, policing, queue scheduling, etc., and no unexpected behavior was noted. In particular, priority queues were enabled on all interfaces, and enqueuing packets marked with DSCP value 46. Bidirectional flows with mixed IPv4 and IPv6 traffic were then set up and 20% of packets were marked with DSCP value 46. The

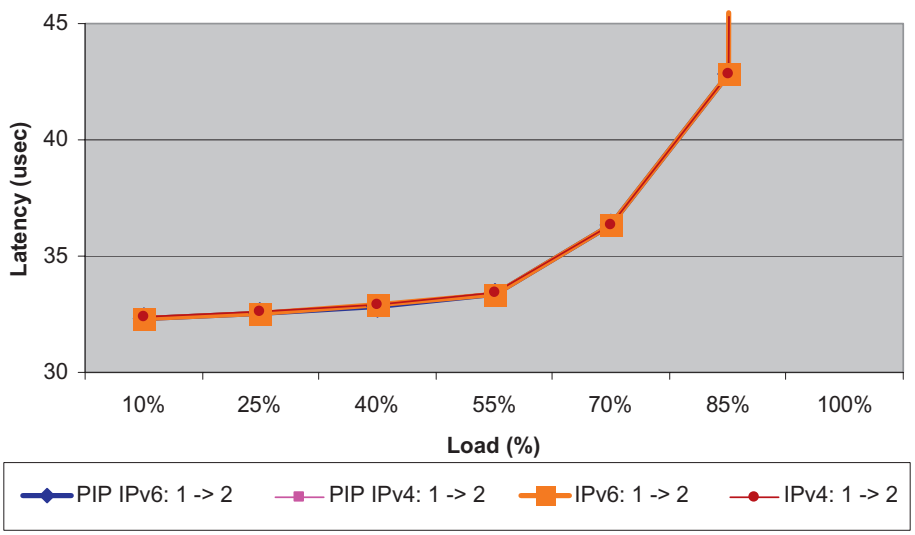

Figure 9. PIP latency in the Tetra cards. A colour version of this figure is available online at www.interscience.wiley.com

DOI: 10.1002/nem

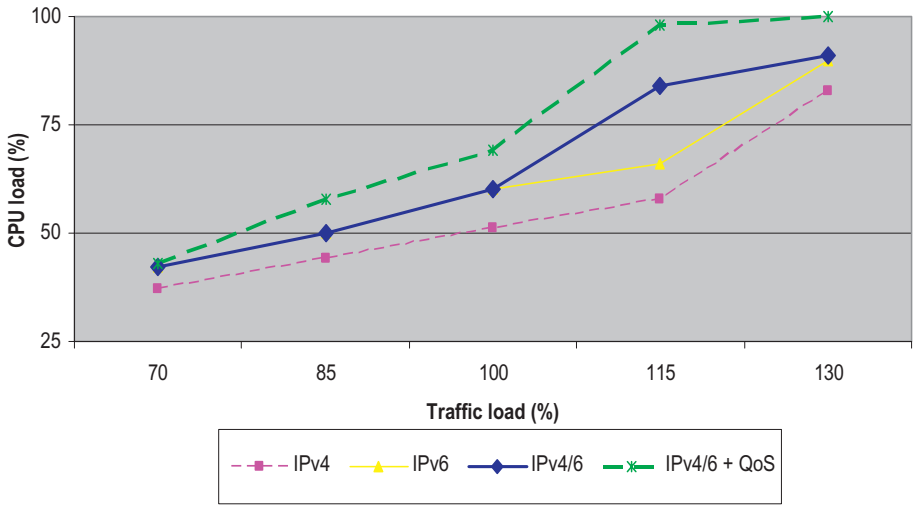

Figure 10. CPU load in software-based platform. A colour version of this figure is available online at www.interscience.wiley.com

Packet size (bytes)	IPv4 traffic (Mbps)			IPv6 traffic (Mbps)			Mixture of IPv4 & IPv6 traffic (Mbps)		
	85	100	115	85	100	115	85	100	115
128	0	0.17	100	0	14.83	100	0	14.99	100
256	0	0	12.55	0	0	12.55	0	0	12.55
384	0	0	12.62	0	0	12.61	0	0	12.62
512	0	0	12.66	0	0	12.66	0	0	12.66

Table 2. Packet loss for various packet sizes and traffic load

traffic rate was from 70 to 130 Mbps in steps of 15 Mbps. As shown in Figure 11, latency for PIP traffic is the same for both IPv4 and IPv6 protocols and, under congestion, the PIP traffic exhibited lower delay than Best Effort traffic, showing that priority queues protected the marked traffic efficiently and provided the relevant guarantees.

Finally, we measured the delay variation (average jitter) for network load equal to 130 Mbps (both IPv4 and IPv6 traffic at a 50:50 ratio), which is presented in Figure 12. In this scenario, 10% of IPv4 traffic and 10% of IPv6 traffic was marked as PIP. We measured the average jitter that Best Effort and PIP traffic experienced during the experiment. As there was heavy congestion in this scenario, the result was expected. The PIP traffic experienced low delay variation, in conjunction with quite high delay variation for Best Effort traffic. No significant differences were noted between IPv4 and IPv6 PIP traffic. On the other hand, IPv6 BE traffic seemed to experience a little higher jitter than IPv4 BE traffic.

5. LARGE-SCALE IPV6 QOS TESTS

After evaluation of the support of IPv6 and the performance of IPv6 QoS-related mechanisms in various platforms of GRNET's production network, we extended the tests in a large-scale environment using 6NET's network [18]. The main goal was to investigate possible differences in a native IPv6-only environment (as 6NET was) compared with a dual-stack GRNET network and also to evaluate IPv6 QoS results

DOI: 10.1002/nem

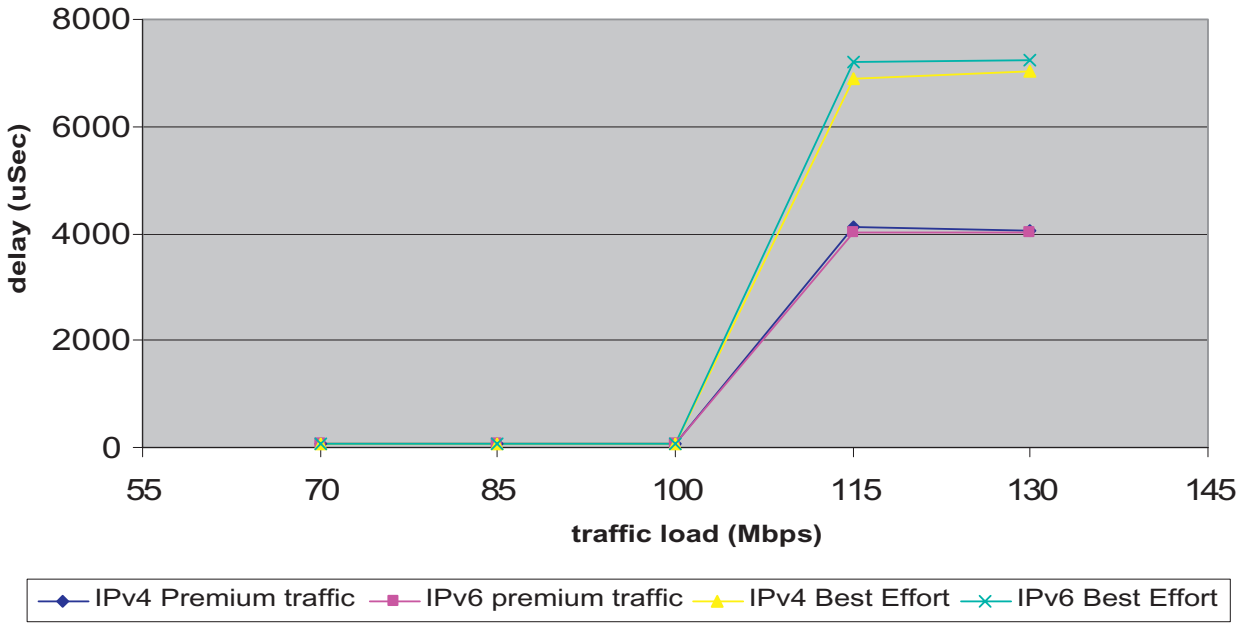

Figure 11. Delay for PIP and BE traffic in software-based platform. A colour version of this figure is available online at www.interscience.wiley.com

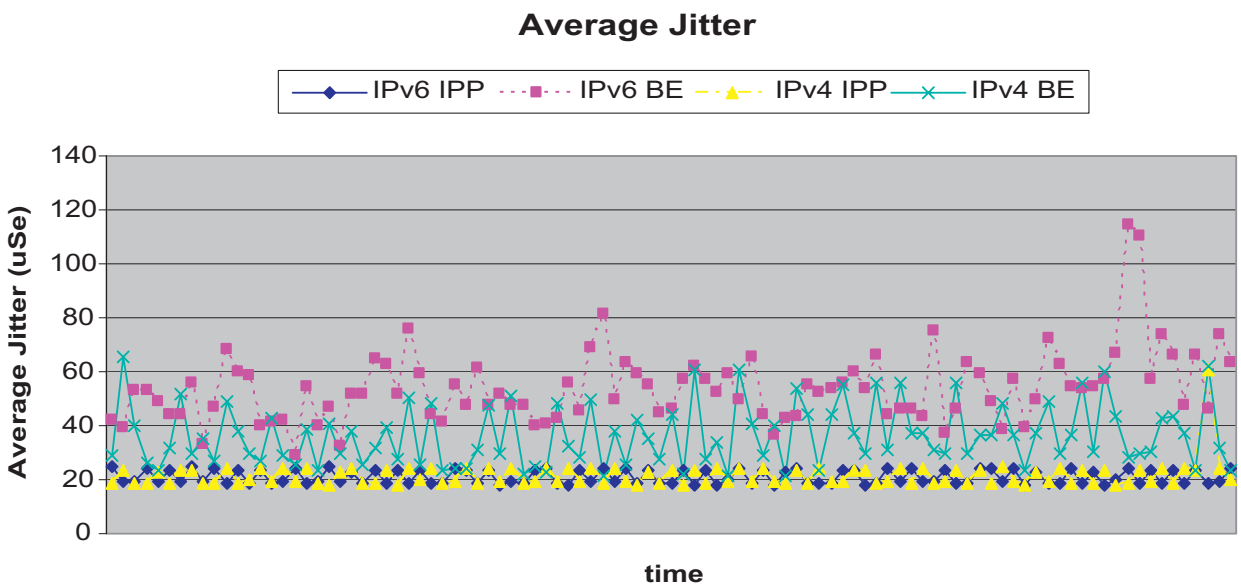

Figure 12. Delay variation for network load equal to 130 Mbps. A colour version of this figure is available online at www.interscience.wiley.com

DOI: 10.1002/nem

from a large-scale network and therefore make more concrete conclusions about the performance of IPv6 QoS mechanisms.

6NET [18] 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 and 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 13, 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 (the same platforms as in GRNET).

The 6NET core network supported 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.

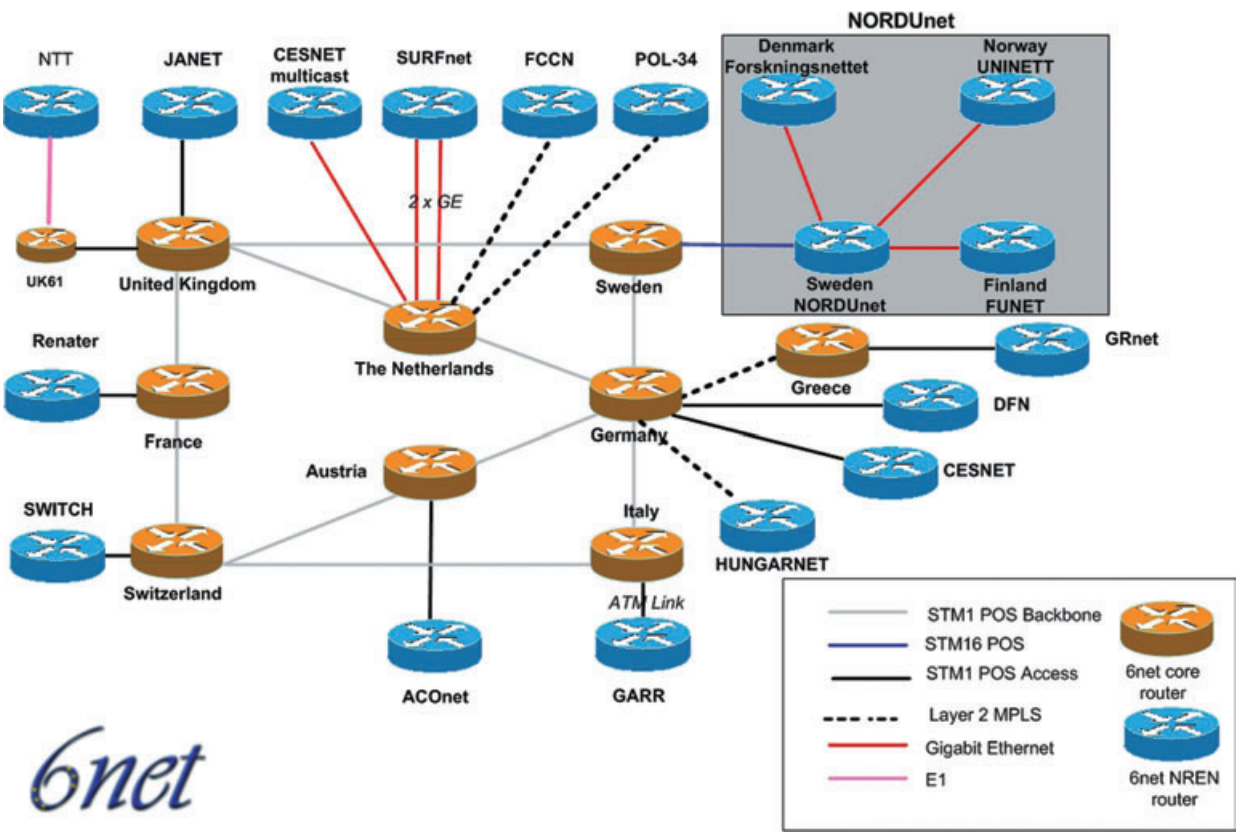

Figure 13. The 6NET network. A colour version of this figure is available online at www.interscience.wiley.com

DOI: 10.1002/nem

Scenario	Best Effort (Mbps)	IP Premium (Mbps)	
1	80	UDP	1.5[a]
2	120	UDP	1.5[a]

[a]Traffic is increased in steps of ~0.5 Mbps.

Table 3. Testing scenarios

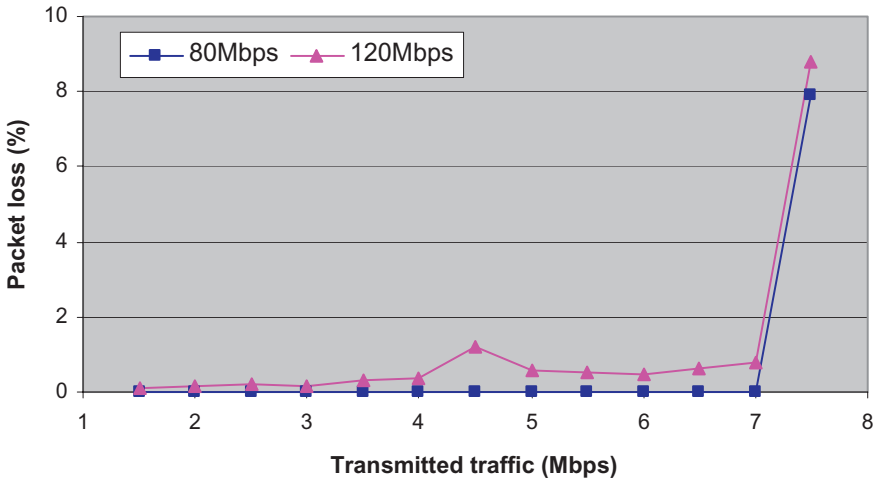

Figure 14. Packet loss for PIP traffic. A colour version of this figure is available online at www.interscience.wiley.com

5.1 QoS tests in IPv6-only environment

The 6NET test-bed consisted of three dedicated PC-based servers connected to Greece, the UK and the Netherlands. The servers generated traffic with *iperf* [19] and *mgen* [20] 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 [21]. We enabled the appropriate classification, policing and queuing mechanisms in core and access routers that allowed up to 5% of the link capacity to be occupied by PIP traffic in high-priority queues. 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 is given in Table 3. Each test was performed with UDP foreground traffic, while the background traffic consisted of a 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 of the UK server (which was connected via a 100 Mbps Fast Ethernet interface), leading to high packet losses.

As shown in Figure 14, PIP traffic experienced approximately zero packet loss with transmitting rates up to 7 Mbps for both scenarios. 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. Conversely, the packet loss for BE traffic was measured as extremely high under congestion conditions (scenario 2). Our results verified the effectiveness of the classification and queuing mechanisms applied at the network interface of the routers.

DOI: 10.1002/nem

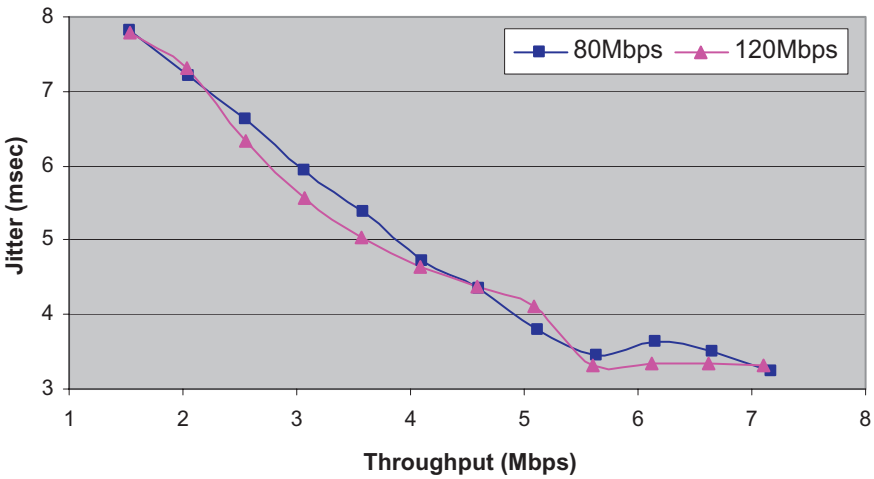

Figure 15. Jitter for PIP traffic. A colour version of this figure is available online at www.interscience.wiley.com

In addition, as observed in Figure 15, jitter experienced by PIP 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 by 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.

6. '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 at the interface level and on classification mechanisms 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) [22] 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 tunneled packets follow the processing switching path (switched by the router CPU).

Additionally, classification criteria in the command line interface did not support flow label criteria and, thus, IPv6 access lists had to be used (IPv6 access lists were the only available tool that allowed classification based on flow label). This workaround solution imposes unnecessary complexity in the definition of classification policies based on the IPv6 flow label field. 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

DOI: 10.1002/nem

hardware and software is quite often deployed in production networks, thus explaining the reluctance of some network providers to migrate to IPv6.

7. CONCLUSIONS

The QoS tests performed in GRNET and 6NET core networks indicated that the 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. Conversely, in older GigE cards, IPv6 is software-switched and experiences worse performance than its IPv4 counterpart, which is hardware-switched. In addition, full software-based platforms (like the Cisco 7200 series router) provided comparable guarantees to IPv4 and IPv6 traffic at most congestion levels, while a small increase in the CPU load was noted when IPv6 traffic was present. Similar qualitative tests in the 6NET network revealed that performance guarantees can be smoothly provided to high-priority traffic in an IPv6-only environment. Also, no differences in IPv6 functionality their performance between IPv6 only and dual-stack networks were noticed. However, in GRNET's test-bed when handling IPv6 traffic under extreme line card congestion, both the Tango and the Tetra cards had a negative impact on routing protocols, due to current internal queue management implementations about control traffic.

Looking at roadmaps of vendors, we notice that new versions of operating systems will give more emphasis to IPv6, and have more stable and error-free implementation of IPv6 features; therefore we expect that most of the 'problems' that we faced and are related to software modules will be solved (partially or completely) and also most of the features in the wish-to-have list will be available. Consequently, as the current portion of IPv6 traffic is significantly low compared to IPv4 traffic, an IPv6 QoS schema can be deployed in research or production networks at gigabit speeds, albeit with 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 the PIP service for IPv6 in its dual-stack gigabit core network. Currently, we completed the QoS setup of the network in order to extend the existing QoS services (IPv4 only), providing to subscribers IPv6 QoS services too.

8. FURTHER WORK

After the evaluation and monitoring of QoS functionality in IPv6 platforms, we already have plans for future work. These plans are divided into two categories: the enhancement of IPv6 QoS with new features, and a management tool for this. We are very interested in testing new features that will be available on network platforms, especially the usage of flow label field and all other issues described in the wish-to-have list. On the other hand, we plan to enhance the existing semi-automatic management tool (which manages IPv4 QoS [16]) in order to support the IPv6 QoS. The goal is to implement the user interface, automatic admission control as well as the relevant configuration for each request. The management tool may also include the handling of flow label value for special IPv6 QoS requests. In addition, we plan to investigate a method for providing real-time statistics. The latter is very important as it can be used in future service-level agreements.

REFERENCES

1. Deering S, Hinden R. Internet Protocol, Version 6 (IPv6). *RFC 2460*, 1998.
2. GÉANT: The Trans-European Research Network. http://www.geant.net [16 March 2008].
3. Blake S, Black D, Carlson M, Davies E, Wang Z, Weiss W. An architecture for Differentiated Services. *IETF RFC 2475*, 1998.

DOI: 10.1002/nem

4. Roth R, Campanella M, Leinen S, Sabatino R, Simar N, Przybylski M, Trocha S, Liakopoulos A, Sevasti A. IP QoS across multiple management domains: practical experiences for the pan-European experiments. *IEEE Communications Magazine* 2003; **41**: 62–69.
5. Bader A, Westberg L, Karagiannis G, Kappler C, Phelan T. RMD-QOSM: the resource management in Diffserv QOS model. draft-ietf-nsis-rmd-04.txt, 2005 (work in progress).
6. Westberg L, Csaszar A, Karagiannis G, Marquetant A, Partain D, Pop O, Rexhepi V, Szabo R, Takacs A. Resource management in Diffserv (RMD): a functionality and performance behavior overview. In *7th International Workshop on Protocols For High-Speed Networks (PfHSN'02)*, 2002.
7. Nichols K, Jacobson V, Zhang L. A two-bit Differentiated Services architecture for the Internet. *IETF RFC 2638*, 1999.
8. Engel T, Granzer H, Koch BF, Winter M, Sampatakos P, Venieris IS, Hussmann H, Ricciato F, Salsano S. AQUILA: adaptive resource control for QoS using an IP-based layered architecture. *IEEE Communications Magazine* 2003; **41**: 46–53.
9. Internet2 Qbone Architecture. Internet2 QoS Working Group. http://qbone.internet2.edu/arch-dt.shtml [16 March 2008].
10. Almes G, Kalidindi S, Zekauskas M. A one-way delay metric for IPPM. *IETF RFC 2679*, 1999.
11. Demichelis C, Chimento P. IP packet delay variation metric for IP performance metrics (IPPM). *IETF RFC 3393*, 2002.
12. Almes G, Kalidindi S, Zekauskas M. A one-way packet loss metric for IPPM. *IETF RFC 2680*, 1999.
13. Chimento P, Ishac J. Defining network capacity. draft-ietf-ippm-bw-capacity-00, 2005 (work in progress).
14. Rajahalme J, Conta A, Carpenter B, Deering S. IPv6 flow label specification. *RFC 3697*, 2004.
15. Greek Research and Technology Network: GRNET. http://www.grnet.gr [16 March 2008].
16. Varvitsiotis AP, Siris VA, Primpas DN, Fotiadis GI, Liakopoulos AC, Bouras CJ. Techniques for DiffServ-based QoS in hierarchically federated MAN Networks: the GRNET case. In *14th IEEE Workshop on Local and Metropolitan Area Networks*, 2005.
17. SmartBits 600 Series traffic generators. Spirent, http://www.spirent.com [16 March 2008].
18. *6NET: Large Scale IPv6 Pilot Network*. IST Programme (IST-2001-32603) http://www.6net.org [16 March 2008].
19. Internet Performance: IPERF. http://dast.nlanr.net/Projects/Iperf [16 March 2008].
20. Multi-Generator Toolset (MGEN). http://mgen.pf.itd.nrl.navy.mil/mgen [16 March 2008].
21. *6NET Deliverable D.4.4.2v2. (2005). Report in QoS Tests, 2nd version*. IST Programme (IST-2001-32603).
22. *Service Assurance Agent (SAA)*. Cisco Systems Inc., http://www.cisco.com [16 March 2008].
23. *Cisco 12000 Series Routers*. Cisco Systems Inc., http://www.cisco.com [16 March 2008].

AUTHORS' BIOGRAPHIES

Dipl.-Ing. **Athanassios Liakopoulos** received the Dipl.-Ing. degree in Electrical and Computer Engineering from the National Technical University of Athens (NTUA), in 1996, MSc with Distinction in Telematics (Telecommunications & Computer Engineering) from the Electrical Engineering Department in University of Surrey (UniS) in 1998. He received also a PhD from NTUA in the field of QoS provisioning at high-speed networks. 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 has awarded for his performance during his academic studies and has published articles in recognized technical journals. He is member of TEE and IEEE.

Dimitrios Kalogeras was born in Athens, Greece in 1967. He received the Diploma in Electrical Engineering from the National Technical University of Athens (NTUA), Greece in 1990 and PhD in Electrical and Computer Engineering from NTUA in 1996. From 1993 to 1995 he worked for the NTUA Network Management Center as Data Network Engineer. From 1995 to 1996 he worked for INTRACOM S.A. as an Engineer in Research & Development. From 1997 to 1999 he worked as Technical Consultant for NTUA NMC and GRNET. From 1999 till today he is a Researcher of Institute of Communication and Computer Systems, Department of Electrical and Computer Engineering.

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

DOI: 10.1002/nem

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 has served 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 obained his Diploma, Master Degree and PhD from the Computer Engineering and Informatics Department of Patras University (Greece). He works in the Research Unit 6 of CTI and on Telematics, Distributed Systems and Basic Services Laboratory of the Computer Engineering & Informatics Department, in Patras University and has participated in numerous R&D projects. His interests include: networks, protocols, Quality of Service and Network applications. He is co-author of 2 books in Greek Language and he has published more than 24 research papers in various well-known refereed conferences and 6 research papers in scientific journals.

Christos Bouras obtained his Diploma and PhD from the Department Of Computer Engineering and Informatics of Patras University (Greece). He is currently and Associate Professor in the above department and a scienfific advisor of Research Unit 6 in Research Academic Computer Technology Institute (CTI). He has published over 200 papers in various well-known refereed conferences and journals. He is a co-author of seven books in Greek. He has been a PC member and referee in various international journals and conferences and he has participated in numerous R&D projects. 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.

DOI: 10.1002/nem