

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Report on the 19th TF-NGN meeting

3-4 November 2005

NTUA, Athens, Greece

Issue 1.1, Kevin Meynell

Attendees

<u>Name</u>	<u>Organisation</u>	<u>Country</u>
Spiros Antonatos	FORTH-ICS	Greece
Emilie Camisard	RENATER	France
Mauro Campanella	Consortium GARR	Italy
Vassilis Chatziannakis	NTUA	Greece
Tryfon Chiotis	GRNET	Greece
Athanasios Douitsis	NTUA	Greece
Jerome Durand	RENATER	France
Michael Enrico (Chair)	DANTE	-
Hans-Martin Foisel	Deutsche Telekom	Germany
Marcin Garstka	PSNC	Poland
Mary Grammatikou	NTUA	Greece
Vangelis Haniotakis	University of Crete	Greece
Jon Kåre Hellan	Uninett	Norway
Gabor Ivanszky	NIFF/HUNGARNET	Hungary
Avgust Jauk	ARNES	Slovenia
Dimitrios Kalogeras	GRNET	Greece
Theodoros Kostis	University of the Aegean	Greece
Radek Krzywania	PSNC	Poland
Felix Kugler	SWITCH	Switzerland
Angelo Lenis	NTUA	Greece
Athanasios Liakopoulos	GRNET	Greece
Basil Maglaris	NTUA	Greece
Evangelos Markatos	FORTH-ICS	Greece
Kevin Meynell (Secretary)	TERENA	-
Simon Muyal	RENATER	France
Spiros Papageorgiou	NTUA	Greece
Tanya Politi	University of Peloponnese	Greece
Andreas Polyrakis	NTUA	Greece
Vassiliki Pouli	NTUA	Greece
Michael Przybylski	PSNC	Poland
Jürgen Rauschenbach	DFN	Germany
Victor Reijs	HEAnet	Ireland
Stelios Sartzetakis	GRNET	Greece
Laura Serrano	RedIRIS	Spain
Afrodite Sevasti	GRNET	Greece
Christos Siaterlis	NTUA	Greece
Miguel Sotos	RedIRIS	Spain

Jan Späth	Marconi	Germany
Pyrrhos Stathis	TU Delft	The Netherlands
Robert Stoy	DFN	Germany
Chrysostomos Tziouvaras	GRNET	Greece
Sven Ubik	CESNET	Czech Republic
Jean-Marc Uzé	Juniper Networks	France
Stig Venaas	Uninett	Norway
Joseph Vojtech	CESNET	Czech Republic
Chris Welti	SWITCH	Switzerland

Apologies

<u>Name</u>	<u>Organisation</u>	<u>Country</u>
Tim Chown	University of Southampton	United Kingdom
Rob Evans	ULCC	United Kingdom
Carlos Friaças	FCCN	Portugal
Simon Leinen	SWITCH	Switzerland
Szilárd Márton	NIIF/HUNGARNET	Hungary
Janos Mohacsi	NIIF/HUNGARNET	Hungary
Wilfried Woeber	ACOnet-CERT	Austria

Meeting Presentations

All presentations from the meeting are available online at: <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/presentations.html>.

This report records the main discussion items and actions arising during the meeting. Readers should refer to the presentations for detailed information.

1. GÉANT2 Rollout Update, Michael Enrico

Michael reported on the status of the GÉANT2 rollout (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/enrico-geant-rollout.pdf>). They were currently planning the transition of the existing GÉANT network to the new GÉANT2 topology, based on a hybrid infrastructure. There would be twenty-five PoPs (plus another four for NORDUnet) serving thirty countries, a total of 11,600 kilometres of fibre, and more than 140 ILA sites – a significant increase in the number of stations requiring operational management compared to GÉANT(1). At least fifty of the 10 Gbps lambdas would be ‘owned’ (meaning implemented on the GÉANT2 DWDM equipment), with another nine 10 Gbps lambdas being leased. A further eight 2.5 Gbps lambdas plus a handful of lower-speed links would also be leased. A combination of Juniper T640, M160 and M40 routers would be used in the PoPs, offering NRENs access to IP services at up to 10 Gbps and lower-speed (sub 1-Gbps) point-to-point links based on the L2VPN approaches used today. In addition, many NRENs would have access to the co-located optical cross-

connects (NG-SDH switches) to access fully deterministic point-to-point services with guaranteed bandwidth (circuit-based 'lightpaths') at up to 10 Gbps and/or direct access to the DWDM terminals to access conventional 'lambda' services. Finally, there are expected to be 4 x 10 Gbps to the North America (currently under tender) with connections to Abilene, ESnet, CA*net4, SINET, TENET, EUMEDCONNECT, RedCLARA as before, with TEIN2 (Asia-Pacific) being added shortly.

GÉANT2 would continue to support the usual IP services such as Best Effort (BE), less-than-best-effort (LBE), Premium IP (PIP), IPv6 and multicast (v4 and v6), and these would be enhanced with new features as required (e.g. Embedded RP for IPv6 multicast). There would also be improved performance monitoring, more automated PIP provisioning, and implementation of security best-practice. Emulated Layer 2 Point-to-Point services would continue to be supported, with a new portfolio of L1 and L2 point-to-point services (e.g. native GE to GFP VC-4-7v over SDH) being added. Emphasis would be placed on support for Ethernet services with access interfaces being native 1 GE, 10 GE or SONET/SDH (in which case GFP-encapsulation of Ethernet into virtually concatenated VC-4s or STS-3c will have to already have been performed by the NREN).

The transition of the network required that 11 of the current 18 sites containing IP equipment be moved to a different location. This included the decommissioning of 30 leased circuits, and the transport of approximately 60 items of equipment between PoPs. Another 200 ancillary sites also needed to be set-up, and the necessary equipment installed.

The GÉANT2 build and GÉANT(1)-GÉANT2 migration has been planned by DANTE operations to progress in five phases, with Phase 1 (Vienna-Milan-Geneva-Frankfurt) now being mostly completed (with the exception of a few key parts which have been the subject of lengthy, unforeseen delays). The Phase 2 (Frankfurt-Amsterdam-Brussels-London-Paris-Geneva) PoP installations/upgrades and parts of the transport build were also now quite far progressed. Further details of the phases (including details of relevant PoP moves/upgrades, transport builds, router deployments/moves, NREN access migrations, etc...) can be found in the slides referenced above.

It was expected that the entire rollout and migration would now be completed by the end of the first quarter of 2006, as it was currently 2-4 weeks behind schedule since contract negotiations with Alcatel took longer than expected (leading to a delayed start to the rollout/migration). There had also been significant and unforeseen access problems at CERN.

2. GN2 Research and Service Activity Updates

2.1 JRA1 Update, Nicolas Simar

Nicolas gave an update on the status of the JRA1 activity (see <http://www.terena.nl/tech/>)

task-forces/tf-ngn/tf-ngn19/simar-jra1.pdf). The MA prototype had been deployed on GÉANT, Abilene, ESnet and Uninett, and was currently being refined. The LS prototype was mostly implemented, although it was still missing auto-registration of services. The next stage would be the implementation of multi-LS. Work had also started on the MP service, with a development week planned in Dublin in January.

The integration of OWD-IPPM into the MA was ongoing, with an implementation date scheduled for the first quarter of 2006. This included enhancements to the visualization interface and measurement probes. The installation status of AB was similar, with integration into MP planned for December. In addition, various NetFlow tools were being investigated.

A passive measurement system was being developed using two 10 Gbps and two 1 Gbps cards. This was used to measure short-term available bandwidth, user traffic packet loss, tracefile capture, and TCP monitoring. As there was also a requirement to start L2 circuit monitoring by May 2006, JRA1 had started to collaborate with JRA3.

The following demonstrations had been made available:

CNM http://cnmdev.lrz-muenchen.de/cnm-app/def/jnlp/topoapplet_devel.jnlp
EGEE <https://egee.epcc.ed.ac.uk:28443/npm-dt/>
ESnet tool <https://performance.es.net/cgi-bin/perfsonar-trace.cgi>

2.2 JRA3 Update, Afrodite Sevasti

Afrodite gave an update on the status of the JRA3 activity (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/sevasti-jra3.pdf>). The goal of JRA3 was to develop a pilot bandwidth and reservation allocation service to streamline and automate the inter-domain setup of lightpaths. This would be a point-to-point connection-oriented service, supporting multiple L1 and L2 technologies such as SDH with GFP, MPLS L2VPN and native Ethernet.

The characteristics of the service would allow end-user points to be located in different domains, with available capacity depending on local policies and the granularity of the technology used (e.g. SDH). It would also be possible to request point-to-multipoint service as a set of point-to-point services.

An inter-domain manager will simplify the path-finding process and handle the necessary signal and routing. It will ensure adequate end-to-end capacity is available, ensure the technical feasibility of specific links (e.g. a common VLAN id is available), and then initiate the connections. Various implementations were currently being investigated, including the Quagga Routing Suite (<http://www.quagga.net/>), and that used by DRAGON (<http://cni.gmu.edu/dragon/software.htm>).

In conjunction with this, an inter-domain monitoring system was required to provide quality-of-service information, and to troubleshoot any failures. JRA3 was therefore

working closely with JRA1 to develop this, as well liaising with the VIOLA, DRAGON, OSCARS-BRUW, MUPBED and HOPI projects.

2.3 JRA5 Update, Jürgen Rauschenbach

Jürgen gave an update on the status of the JRA5 activity (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/rauschenbach-jra5.pdf>). JRA5 is building a next-generation European roaming infrastructure (eduGAIN) based on the experience of existing roaming areas (e.g. Eduroam). It will pilot federated support for existing authentication and authorization infrastructures (AAIs) for research and education. In addition, where federated AAIs are already available (e.g. Shibboleth, PAPI, Moria, A-Select), eduGAIN will cooperate with them.

JRA5 currently involves sixteen NRENs and around 97 other organizations. It collaborates with TF-Mobility, TF-EMC² and various other external groups to ensure as wide coverage as possible. A glossary of terminology, the roaming requirements, an architecture design, and a review of legislation had already been produced.

There are a number of limitations with the current infrastructure, such as authN and authZ traffic having to flow through the complete hierarchy, a static trust model, and single points of failure at the RADIUS servers. Eduroam was also not flexible enough with SSIDs, ciphers and VLAM mapping, and it was sometimes difficult to identify access points.

There were a number of architectural alternatives such as DIAMETER (RFC 3588), RadSec, DNSSec, and RADDNSec. DIAMETER currently seemed the most suitable option, but it was likely RadSec would be available sooner. At the present time, the various architectures were being evaluated, with wider trials commencing in due course. The fact there were already a number of national federations in operation, would provide an extensive test platform for the eduGAIN architecture.

3. MUPBED: Overview, Interworking Issues and Relevance for GN2 and NRENs

Jan Späth and Hans-Martin Foisel gave a presentation on the MUPBED IST project (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/spaeth-mupbed.pdf>). This aimed to integrate telecommunication operator techniques into a research networking infrastructure, based on a high-bandwidth multi-domain, multi-layer environment. This included on-demand circuit switching techniques and ASON/GMPLS control plane technologies. The service and network requirements of high-end applications would be identified, and an experimental environment developed to assess the solutions.

The project was comprised of 16 partners from 8 countries, and included equipment vendors, telcos, research institutes and NRENs. There were five testbeds located in Italy (ASON/GMPLS), Germany (ASON/GMPLS), Sweden (GMPLS), Spain (IP/MPLS) and Poland (Ethernet), which were interconnected through GÉANT. These offered a variety

of data transport possibilities, and the opportunity to trial different control plane-based services.

The nature of the network meant that it was necessary to use standard mechanisms to ensure interoperability. These included the UNI 2.0 Ethernet specification, the ITU ASON and NG-SDH recommendations, the MEF service specifications, the IEEE Ethernet standards, and the IETF signalling standards. These had recently been demonstrated at OIF World, using a global testbed spanning North America, Europe and Asia.

The MUPBED project was interested to understand the future requirements of NRENs with respect to traffic requirements, services, and the types of technology being deployed. This would help develop guidelines for the introduction of ASON/GMPLS technologies and ultra-broadband services in future European networks.

4. New Hardware Evaluation

4.1 Hands-on evaluation of new routers and switches, Marcin Garstka

Marcin reported on the plans for testing Juniper T640s in a TX configuration (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/garstka-evaluation.pdf>). The TX matrix allows up to four routers to be combined into one more powerful router.

The performance impact of upgrading T640s to the latest JUNOS release would be evaluated, as well as adding them to a TX matrix. These tests would include throughput, latency and jitter tests for IPv4, IPv6 and MPLS traffic. The redundancy and failover mechanisms would also be tested, as well as the intra- and inter-T640 convergence of routing protocols in the control plane. Finally, a few QoS/CoS scenarios would be considered, along with the 10 GE DWDM PIC.

4.2 TX Matrix Platform, Jean-Marc Úze

Jean-Marc gave a short presentation about the Juniper TX matrix platform (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/uze-txmatrix.pdf>). This allows clusters of T640 (up to four at present) to form a powerful routing node. It utilises a distributed switch fabric using a common JUNOS image, and offers full redundancy with no single points of failure.

The TX matrix platform has a passive midplane design which offers expansion capability for future upgrades. There are five non-blocking fabric planes, four of which can provide full line-rate bandwidth at 40 Gbps. It also includes dual routing engines, control boards, connector interface panels, load-sharing power supplies, and cooling systems. The routing engines share control of the system and evenly distribute the processing load.

The 10 GE DWDM PICs on the T640s and M320Ts offer full C-Band tunable wavelength support (from the CLI). 100 GHz spacing for up to 45 individual wavelengths is supported, with an 80 km reach for extended metropolitan deployments.

More information is available at http://www.juniper.net/solutions/literature/white_papers/200089.pdf

5. Transport Protocols, Radek Krzywania

Radek gave an update on the transport protocols activity (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/krzywania-transport-protocols.pdf>). Two new Linux kernels (2.6.13 and 2.6.14) had become available in October that offered high-speed TCP, BIC, New Reno (congestion control), HTCP, TPC-HYBLA, Vegas (queuing and packet loss control) and Westwood (congestion control). By default, the kernel included BIC, but the 'sysctl' function could also be used to change the algorithm as necessary.

A new XCP (eXplicit Congestion Control Protocol) release was now available for FreeBSD 5.3. This could cooperate with non-XCP queues, share bottlenecks with Reno-style flows, and evaluate performance in different environments. It could be obtained from <http://www.isi.edu/isi-xcp>.

A new patch for NS2 (network simulator) was also available that could speed-up simulations by factors of up to 50. It was specifically developed for fast, long-distance network simulation, and achieved its remarkable improvements through better use of event scheduling.

Both HEAnet and the Hamilton Institute in Ireland had joined the testbed. Each had a PC with a 1 Gbps interface connected to the GÉANT2 network, providing an RTT of around 73 ms.

Finally, the proposal for the FASTNESS project had been submitted in the final call of the IST 5th Framework Programme. This had passed the initial evaluation, and had progressed to the next stage.

6. The Logistical Session Layer, Martin Swany

Martin gave a presentation about the Logistical Session Layer (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/swany-lsl.pdf>). This has been developed to allow systems to support service nodes that can support short-term storage and cooperative data forwarding. The aim is to improve throughput in a reliable fashion by creating a more responsive control loop via reduction of signalling latency. This would enable better adaptation to local conditions, and the inherent buffering in the network would mean retransmissions need not come from the source.

LSL has been tested over the Internet2 network, and significant improvements in direct transfers have been observed. Logistical Service Nodes (LSNs) have been placed at strategic locations in the network, which provide data stores and control points. The LSL protocol intercepts transport-layer connections and redirects them to appropriate LSNs based on policy. This somewhat contradicts the traditional approach that networks should be transparent for end-to-end connections, but many assumptions regarding scalability and complexity no longer necessarily hold true anyway.

The LSL client library provides compatibility with current socket applications, although more functionality is available using the API directly. The daemon runs on all Unix-based platforms, although forwarding is slightly better with Linux than with BSD variants.

Phoebus offers a similar technique for switched optical networks that pushed adaption points towards shared-edge networks. The Generic Session Layer (GSL) protocol is a new version of LSL that combines the Phoebus work and offers support for session-layer frames. It also provides a rendezvous mechanism for changing IP addresses.

LSL essentially turns sessions into a composition of network-specific transport layers. There are many cases in which a single transport protocol from end-to-end may not be the best choice, and it may be better to adopt a totally different model rather than try to find short-term fixes for TCP problems. LSL offers the advantages of traditional TCP services at network edges, meaning that it can be incrementally deployed, and without the need to upgrade end-hosts.

7. IPv6 (Moderator: Stig Venaas)

7.1 6DISS Project: IPv6 Dissemination and Exploitation, Jérôme Durand

Jérôme gave a presentation about the 6DISS project (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/durand-6diss.pdf>). This aimed to build on the IPv6 deployment experiences of the 6NET, Euro6IX and GÉANT projects, as well as the liaisons established with the IPv6 Forum, European IPv6 Task Force and IETF. It would allow organisations in developing countries to benefit from this experience, since they have fewer legacy installations and are therefore able to utilise the new technology in an efficient manner.

The 30-month IST project was coordinated by Martel (Switzerland) and included Alcatel, Cisco, FCCN, GRNET, NIIF/HUNGARNET, RENATER, TERENA, UCL and the University of Southampton. It would hold IPv6 training workshops in South-Eastern Europe, Central Asia, Southern Africa, Sub-Saharan Africa, South/Central America, the Caribbean, as well as the Mediterranean and Asia-Pacific regions. It would also hold information exchange workshops in China and India.

In conjunction with this, the project was developing training material that would be made publicly available. IPv6 laboratories would also be set-up at Cisco and RENATER for hands-on training, whilst a 'Tiger Team' was available to answer questions about IPv6 deployment and operation (send mail to 'helpdesk@6diss.org').

Training workshops had already been held in Taipei (Taiwan) in August, and in Port Elizabeth (South Africa) in September. Three further workshops were planned for Egypt, Malta and South America during the first quarter of 2006.

Another goal of the project was to assess the levels of Internet connectivity in the different regions, and to establish contacts with organisations to enable their participation in future IST projects. This would hopefully provide opportunities for European manufacturers to open new, or enhance existing markets, and identify key personnel who can assist with improving ICT infrastructures.

7.2 IPv6 Multicast Specifications: RENATER-4, Jérôme Durand

Jérôme gave an overview of the multicast facilities on the RENATER-4 network (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/durand-multicast.pdf>). PIM-SMv2 was configured on all backbone routers, and customers were able to ask for ASM and SSM, or just SSM. MLD was disabled on every interface where there was no host.

Scope management was utilised for both ASM and SSM, with RENATER managing Scope A. No scope lower than A was visible in the backbone, and no scope equal to or lower than A travelled outside RENATER. Inbound and outbound filtering was imposed at the boundaries of the network to enforce this.

Embedded-RP was being used without BSR or static RPs, and was deployed on all core routers. Sites were also free to deploy their own RPs (using Cisco 7204 or 7206 routers), but an RP service was available for those unable to do so. There were two permanent RPs configured on their own Cisco 7204 router: RENATER scope (A) which accepted registrations from 2001:660::/32 only, and Global scope (E) which accepted registrations from any valid IPv6 address. At the moment, there was no RP redundancy, although Anycast PIM may eventually be a solution. In addition, SAP was not currently supported, so they were working on a web-based interface to manage multicast sessions.

There were a number of external connections to GÉANT, SFINX (via a VLAN), and to the M6Bone through the existing testbed. Both static and dynamic IPv4/IPv6 multicast gateways were provided, but most end-sites were still facilitated through tunnels. The Hexago migration broker was still not IPv6 capable yet, and there was currently no date for when it would be.

They were currently discussing whether to allow IPv6 unicast SAFI for multicast RPF checks. There would then be no requirement to deploy MBGP if the topologies were the same for both IPv4 and IPv6.

The multicast network was currently monitored using a combination of DBeacon, SSMPing, QoSmetrics, AS_Path_Tree, Looking Glass, and NetFlow where possible.

7.3 IPv6 QoS Tests in GRNET2 network, Athanassios Liakopoulos

Thanassis gave an presentation on the IPv6 QoS trials undertaken on the GRNET2 network (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/liakopoulos-ipv6qos.pdf>). The aim was to validate the performance of basic QoS mechanisms with IPv6 traffic, to test different types of hardware supporting QoS, and identify any missing functionality or unexpected performance. A further aim was produce guidelines for the deployment of QoS services in an IPv6 production environment.

The GRNET2 testbed utilises six Cisco 12400 and 7200 series routers (based in Athens, Ilissos, Larisa, Thessalonikini, and Patra (x2), and three Smartbit 600s for generating traffic. The QoS tests were conducted for Premium IP, Best Effort and Less-than-Best Effort using Eng3 and Eng4 GE cards. Bi-directional IPv4 and IPv6 flows were created, with the traffic load gradually being increased.

The results showed that for Best Effort traffic, IPv6 packet loss in the Eng4 cards was much higher than for the Eng3 cards. IPv4 and IPv6 experienced the same packet loss in one direction, but it was higher for IPv4 in the other direction. Latency was the same for IPv4 and IPv6 with the Eng3 cards, whilst IPv6 latency was always higher with the Eng4 cards, even with no packet loss.

For Premium IP traffic (PIP), IPv4 and IPv6 latency using the Eng3 cards was low provided there was no packet loss (<85% load). Where there was packet loss (100% load), latency increased, but was still much better than with Best Effort traffic. In both cases, PIP experienced no packet loss. With the Eng4 cards, IPv6 latency was higher than with the Eng3 cards, although it remained constant for most traffic loads.

Examining the impact of different types of traffic in software-based platforms, it showed that CPU load increased by approximately 5-9% when handling IPv6 compared to IPv4 traffic. Turning on the QoS mechanisms increased this by another 10% or so.

It can be concluded that the QoS mechanisms in IPv6 are comparable to those in IPv4, although certain older hardware lacks some functionality. Common QoS models could be supported for both IPv4 and IPv6, provided these restrictions can be overcome. At the moment, IPv6 does not offer any significant advantages over IPv4, although the flow-label field in the header may yet be useful when provisioning future services.

7.4 asmping, Stig Venaas

Stig gave a presentation about asmping (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/venaas-asmping.pdf>). This is a tool for testing IPv4 and IPv6 multicast connectivity that behaves somewhat like normal ping if a server is running the latest version of ssm pingd (which supports asmping). A client can ping a server by sending a unicast query, and the server responds with both unicast and multicast replies. In this

way, a client can check whether it can receive ASM from the server, along with information on delay and the number of hops etc...

An interesting example of how asmping had been used to track-down problems in the multicast network was shown. This revealed a differing number of multicast hops, and duplicate responses with a high RTT. This was attributed to the interaction between PIM and MSDP which was being used to encapsulate the packets, with the high RTT being attributed to caching.

More information about asmping was available at <http://www.venaas.no/multicast/ssmping/>

8. LOBSTER Project, Evangelos Markatos

Evangelos Markatos presented an overview of the LOBSTER IST project (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/markatos-lobster.pdf>). This was motivated by the 'Research Horizons in Networking' Committee that was formed in 2001 to take a fresh look at networking research. This identified three grand challenges, the first of which was to develop and deploy technology to make it possible to record a day in the life of the Internet. This essentially requires various traffic characteristics to be monitored, which can be used to determine bottlenecks in a network, the route packets follow, why particular flows are missing packets, and the types of traffic using a network.

As part of this effort, the SCAMPI IST project developed the Monitoring Application Programming Interface (MAPI) which provides a standardised API that can be implemented on a variety of hardware. This includes standard NIC cards, Endace DAG cards, and the COMBO series of cards produced by Masaryk University. The MAPI can be used to search traffic flows for specific characteristics which can then be dumped to disk for post-processing. If several passive monitors are available, Distributed MAPI (DiMAPI) may be used to interrogate them as if they were one, thus allowing several links to be monitored simultaneously.

The LOBSTER project aims to expand on the work of SCAMPI by establishing a passive European-wide network monitoring infrastructure. In particular, it aims to highlight suspicious traffic, identify covert traffic, and investigate the vast amounts of unidentified traffic.

All that is needed to get started is a regular PC with an NIC, which is sufficient for speeds of up to 100 Mbps. It can also work at speeds of 1 Gbps although heavy string searching is not really recommended. For speeds beyond that, a high-end PC with a custom interface card (e.g. DAG or COMBO) will be required. This should run either MAPI or PCAP, and your choice of applications for processing the output (some of which are supplied with MAPI).

Information on how to participate in the project is available on the LOBSTER website (<http://www.ist-lobster.org/>), including links to the MAPI software and a variety of monitoring applications. The project would also be holding a passive monitoring tutorial during TNC 2006 in Catania, Italy; the details of which would also be available on the website.

Quite reasonably, there are often concerns about privacy, so LOBSTER is flexible about the type of information it collects. In most cases, monitoring data can be anonymised or different views presented to different levels of user. Even if a site is only prepared to share general statistics, that can also be of use. The project has produced a document that outlines a complete framework for applying anonymisation policies, which can be found at http://www.ist-lobster.org/deliverables/Deliverable_D1.1a.pdf.

9. Service Level Monitoring with Nagios, Athanasios Douitsis

Athanasios gave an overview of Nagios SLA monitoring system (see <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn19/douitsis-nagios.pdf>). Many modern applications (e.g. VoIP) are affected by packet round-trip time, jitter, packet loss and one-way delay, and SLAs are usually based on these qualities. SLA monitoring can be undertaken with network probes that collect data for processing and presentation. Obviously there are requirements for accurate measurements, low installation costs, resilience to outages, and for systems to be standards compliant.

The Cisco Service Assurance Agent (an IOS feature) defines probes that can be used to take appropriate measurements and aggregate the data for statistical purposes. This data is then made available through the SAA-MIB for a 24-hour period, which can be collected and stored in a database for interrogation by SNMP clients.

Nagios offers a extensible network and service monitoring plug-in that is able to monitor the measurement database on a periodic basis (e.g. hourly, daily, weekly, monthly). It can make SLA calculations against pre-defined metrics (e.g. downtime, jitter thresholds), and provide graphical, mail or pager alerts if conditions are not met. It will also facilitate discovery of available probes, and establish SNMP sessions to the target devices.

Experience of deploying Nagios was very favourable as it offers a low administration overhead. In addition, the processing load of the combined collector, database and plug-in was very light which meant that low-cost hardware could be utilised. However, there were a number of bugs in the SAA-MIB, which means alternative MIBs from the IETF DISMAN Working Group should perhaps be considered. In addition, support for the NMWG schema would be advantageous.

10. Next meetings

The next meeting will be held on 13 January 2006 in Cambridge, United Kingdom. This will be held in conjunction with the GN2 meetings scheduled for 9-12 January 2006.

A host is being sought for the following meeting in April or May 2006.

Summary of Actions

ACTION 050113-01	Victor to provide a draft of the term definition document before the next TF-NGN meeting.	Outstanding
ACTION 050113-03	Marcin to provide a list of features to be tested and the test methodology in work item 9.7.	Outstanding
ACTION 050728-01	Stig Venaas to draft test plan for running Embedded-RP on GN1 testbed.	Outstanding
ACTION 040728-02	Kevin Meynell to contact GRNET about arranging alternative dates for the next meeting.	Done
ACTION 050729-03	Michael Enrico and Kevin Meynell to discuss organisation of workshop on dynamic routing configuration.	Outstanding