

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 9th TF-NGN meeting

17-18 October 2002,

University of Technology and Economics, Budapest, Hungary

Valentino Cavalli, Kevin Meynell,
Issue 2, 11 November 2002

Attendees

Name	Organisation	Country
-----	-----	-----
Michael Allenby	UKERNA	United Kingdom
Lada Altmannova	CESNET	Czech Republic
Mauro Campanella	GARR-INFN	Italy
Valentino Cavalli (Secr)	TERENA	The Netherlands
Tim Chown	Univ. of Southampton	United Kingdom
Joel Corral	ENST-Bretagne	France
Daniel Davids	CERN	Switzerland
Rob Evans	ULCC	United Kingdom
Tetenyi Istvan	Hungarnet	Hungary
Avgust Jauk	ARNES	Slovenia
Dimitrios Kalogeras	GRNET	Greece
Joannis Kappas	DANTE	United Kingdom
Andras Kovacs	NIIF	Hungary
Olav Kvittem	UNINETT	Norway
Ladislav Lhotka	CESNET	Czech Republic
Tamas Maray	NIIF	Hungary
David Martinez	RedIRIS	Spain
Octavio Medina	ENST-Bretagne	France
Kevin Meynell	TERENA	The Netherlands
Janos Mohacsi	Hungarnet	Hungary
Julio Orozco	ENST Bretagne	France
Victor Reijs	HEANet	Ireland
Rudolf Roth	Fraunhofer FOKUS	Germany
Roberto Sabatino (Chair)	DANTE	United Kingdom
Yves Schaaf	RESTENA	Luxembourg
Stanislav Sima	CESNET	Czech Republic
Nicolas Simar	DANTE	United Kingdom
Trond Skjesol	UNINETT	Norway
Vladimir Smotlacha	CESNET	Czech Republic
Miguel Ángel Sotos	RedIRIS	Spain
Andre Stolze	WWU Münster	Germany
Robert Stoy	DFN	Germany
Matjáz Straus	ARNES	Slovenia
Andrei Sukhov	Samara University	Russia
Szymon Trocha	PSNC	Poland
Bernard Tuy	RENATER	France
Sven Ubik	CESNET	Czech Republic
Jean-Marc Uzé	Juniper Networks	France
Stig Venaas	UNINETT	Norway
Steven Williams	UKERNA/UW Swansea	United Kingdom
Fuhua Yin	STC-ULB	Belgium

Apologies

Tiziana Ferrari

INFN-CNAF Bologna

Italy

Online presentations

<http://www.terena.nl/task-forces/tf-ngn/presentations9.html>

1. Welcome and introduction of NIIF/Hungarnet, Tamas Maray, NIIF

The meeting was jointly hosted by NIIF/Hungarnet and the Budapest University of Technology and Economics. Hungarnet, the national research and education network of Hungary, is a user association, NIIF is the governmental umbrella organisation providing administration and organisational support to Hungarnet. The Hungarnet backbone is a Gigabit optical network connecting more than 400 member institutes, including universities, high schools, libraries and not-for-profit research organizations totalling up to 5000 end-users. Hungarnet is also running supercomputing facilities and services. Major projects are active in the IPv6 area like TIPSTER6, which started in 2001, and 6NET. Hungary has 6bone PTLA and many connections in place.

2. GÉANT Update, Roberto Sabatino, DANTE

Two STM-4 links to Greece are operational, Malta has a 34 Mbps link in place and CARNET has been upgraded to 622 Mbps. After the Telecommunication industry turmoil earlier in 2002 connectivity to USA is now in place with three 2.5 Gbps links provided by T-System and Level3. The contract for GÉANT connectivity to Ireland has been awarded to Interoute to be operational mid November.

GÉANT had a very successful project review on 18 September in Paris. The auditors recommended to continue IPv6 work and start the deployment phase, to deploy an own testbed facility and monitoring infrastructure, including the establishment of a Performance Response Team and finally urged to do something on network security.

The contract negotiation for the year 3 roadmap is ongoing, some new deliverables are expected in the areas of IPv6, network monitoring and IP/optical technology deployment. By summer 2003 DANTE want to have IPv6 natively available on GÉANT.

3. GÉANT Testbed Facility, Nicolas Simar, DANTE

DANTE plans to set up a GÉANT testbed facility in order to test new services and features, coexistence of new services with existing ones, compatibility of equipment, etc. The facility will be available to NRENs via a booking system similar to the one used in the 6NET project. The NRENs will have access to the routers, will be able to configure them and to get some support from Dante.

The testbed facility will provide three Juniper M20, two Cisco 7000 and one GSR connected to GÉANT via STM-1 plus some PCs and Smartbit traffic generators. CCC from Juniper will be used to provide layer2 connectivity. The routers have been chosen as representative of typical GÉANT and NREN equipment. On request it would be possible to host equipment from other manufacturers.

The facility should be set up before the end of 2002, but the actual date depends on the availability of the equipment, which is being ordered in October.

Renater, Heanet, GRNET and Ukerna expressed an interest in using the testbed facility.

4. GÉANT Year 3 Technical Roadmap, Roberto Sabatino, DANTE

The Technical Roadmap contains three areas of development: technology, network monitoring and services.

The technology area includes IPv6, optical networking and router testing. DANTE will operate and extend the IPv6 overlay network, also adding connection to Abilene. Native IPv6 will be provided on GÉANT. Some static routing with ARNES and RedIRIS have been done already, by setting a new land speed record and showing to the world the commitment to IPv6, the next steps are having dynamic routing on GÉANT by summer 2003, to encourage nrens to use it and pilot to full service in 2004.

Optical developments will include 40Gps and beyond, issues with OOO, integrated IP/optical control plane, alternative service models with carriers. This work will benefit from results of the SERENATE studies (<http://www.serenate.org> see also below in this report). A router-testing programme is already running with Alcatel A7770, DANTE is interested in extending it to the Avici TSR, the Juniper T640 and others. Lab testing will be followed by field trials.

DANTE is planning to deploy a multi-domain monitoring infrastructure by summer 2003 as well as to provide guidelines for establishment of a Performance Response Team by autumn 2003 (see below).

New services to be provided on GÉANT include the Testbed Facility presented in point 3 above, enhancement of multicast and premium IP, as well as Layer2 support for some special projects.

SSM is planned in spring 2003, Per-group Monitoring in summer 2003, 100+ Mbps capacity and IPv6 multicast in autumn 2003.

Premium IP is planned to be enhanced by means of a web-based management interface to collect requirements, automatic resource allocation and router configuration management.

GÉANT will support special projects by understanding their requirements and see what solutions are needed to fulfill them with either high capacity BE, premium IP, L2 VPN or dedicated wavelength. Requirements come at the moment from 6NET, to which L2 VPN function has been delivered, ATRIUM, JIVE, FAIN. Different options for projects to access the GÉANT infrastructure (not for production-like traffic) are still being discussed.

5. TF-NGN update of work items, Valentino Cavalli, TERENA

The audience was asked to provide feedback on the TERENA proposal to update the list of work items in the Terms of Reference of TF-NGN (<http://www.terena.nl/tech/task-forces/tf-ngn/ToR.html>) in order to reflect the activities actually carried out by the task force. It was agreed that the new activities from the GÉANT Year 3 Technical Roadmap would be integrated into the Terms of Reference and the work item 'Policy based networking' would be broadly reformulated in order to ensure that feedback from relevant activities carried out outside of TF-NGN is provided to the task force.

ACTION Valentino and Roberto to update the list of work items in the TF-NGN ToR.

6. LBE session, Tim Chown, University of Southampton (chair)

Tim introduced the agenda of the session and described the content of the GÉANT deliverable D9.9, 'Experiments with Less than Best Effort (LBE) Quality of Service' (<http://www.dante.net/tf-ngn/D9.9-lbe.pdf>), which provides a history of LBE, description of the service and deployment scenarios as well as the GÉANT LBE test programme.

6.1 LBE experiments on GÉANT, Nicolas Simar, DANTE

Tiziana Ferrari from INFN, Italy, has carried out LBE experiments on a GÉANT testbed. Two Smartbit traffic generators connected via STM-16 interfaces in Germany and Italy were injecting traffic to two interfaces from France towards Spain and from Spain towards Italy. The performance analysis showed no big difference between BE and LBE in case of no-congestion. With congestion the results showed an extremely high percentage of out-of-order packets, affecting both BE and LBE. It was possible to significantly reducing packet re-ordering by changing configuration. The conclusion was that enabling LBE queue on GÉANT is possible provided that BE is isolated from LBE in case of congestion. The impact of LBE on BE OWD and IPDV is negligible. Packet re-ordering can be reduced, so that end-to-end TCP throughput is preserved, by assigning lower queue priority to the LBE queue than the one of the BE and EF queues.

6.2 Follow up on LBE, Tim Chown, University of Southampton

LBE should be used to allow throughput of many flows at the same time, but it is questionable whether this is a good metric or not.

At the LBE BoF in Limerick (see report at <http://www.terena.nl/tech/task-forces/tf-ngn/LBEBoF4June2002mins.pdf>) it was agreed to provide DSCP transparency to Internet2 traffic on GÉANT, the discussion now was whether allowing DSCP transparency end-to-end or setting up a "full" queuing policy. There were concerns that enabling LBE at the edge would not help in avoiding packet drop if routers are not properly configured. In order to enable transparent DSCP anywhere and provide a viable service the NRENs have to be involved.

Initial tests with edge router would be held in October-November between UNINETT and the University of Southampton.

A demonstration of LBE traffic between Internet2 and GÉANT is planned at SC2002 in Baltimore on 16-22 November 2002 in the area of bulk data transfer and simultaneous live UDP-based applications. The demonstration would also involve Starlight, TeraGrid, Netherlight and two more sites in Europe: the Czech Republic and Germany or Poland.

7. IPv6, Tim Chown, University of Southampton (chair)

Bernard gave an update on the GTPv6 status (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/tuy.pdf>). CERN, CESNET, HEAnet, PSNC, RCCN, RedIRIS, RENATER, RESTENA and SWITCH were currently connected to the Juniper M5 router in Paris, but there were plans to re-configure the Hitachi GR2000 router with DANTE's IPv4 and 6Bone addresses. This would allow the connections to be shared between the two routers, and bring more than 10 sites on the M6bone into GTPv6.

The M6bone currently had 25 connections in 11 countries and used 6WIND, BSD+KAME and Cisco routers. Most hosts were running an IPv6 stack implementing MLD. The protocols used were RIPng in the tunnels, PIM SM and MLDv1.

Experiments were currently being undertaken with PIM SM and multiple RPs. High bandwidth television (1.5 Mbps) was being continually broadcasted, and the X-Aristote seminar was being broadcast via the v6/v4 gateway. It was hoped to improve the topology by introducing backup links and undertake more experiments using PIM SSM.

The M6bone was being monitored by a pseudo Looking Glass (<http://pasillo.renater.fr/m6bone/>) and a Beacon server (<http://beaconserver.m6bone.pl/>). The intention was to install a Beacon client on every LAN.

Tim mentioned the M6bone had provided a lot of experience with multicasting that would be useful for 6NET. Unfortunately, the Cisco GSRs would not support IPv6 multicasting for another six months, and it may initially have problems when it does appear.

Stig reported interoperability problems between Cisco and FreeBSD routers. He suggested that IPv6 multicasting may not be advisable on large switched networks.

Av gust gave a presentation on the IPv6 'Land Speed Record' that had recently been held over GÉANT in conjunction with Internet2 (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/tuy.pdf>). The participants were DANTE, ARNES and RedIRIS.

The aim was to undertake performance tests for IPv4, IPv6 tunnelled over IPv4, and native IP6, directly between Ljubljana and Madrid (2518 km) as well as via New York (14,800 km). Juniper routers running JUNOS 5.3R2.4 and 5.3R3.3 were utilised, and the hosts used for sending and receiving traffic were 1.1/1.3 GHz Pentium IIIs running Linux.

The first test between Slovenia and Spain via Austria, Switzerland and Italy, yielded an average round-trip time of 62 milliseconds with a 480-490 Mbps transfer rate, and had been accepted as a new Internet2 record. The second test between Slovenia and Spain via Austria, the US, the UK, France, Switzerland and Italy yielded an average round-trip time of 229 ms with a 350-400 Mbps transfer rate, and had been submitted as a record, although it had not yet been accepted [It was subsequently ratified at the end of October.]. It was noted the IPv6 performance was comparable with the IPv4 performance, and no problems had been observed with the dual-stack Juniper routers.

Roberto then outlined the proposed timetable for IPv6 deployment in GÉANT (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/sabatino-v6rollout.ppt>). The preliminary activities such as developing the addressing and routing plan, the peering policy, and the service definition would be undertaken over the next couple of months. The plan was to start operating a pilot service from February 2003.

Janos gave an update on the 6NET activities (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/chown-IPv6.ppt>). The backbone has been operational since May 2002 and is running at 155 Mbps. There had also been a lot of activity with respect to transitional issues (WP2), basic network services (WP3), mobility (WP4) and network management (WP6).

Tim mentioned that research networks had been invited to submit papers to the IEEE Workshop that was being held in February 2003 in Florida. He had decided to submit two papers - one on GÉANT, 6NET, Euro6IX and other IST projects, and one on NREN activities. Unfortunately, the deadline was looming and he needed contributions in the next two or three days.

8. QoS measurement, Victor Reijs, HEANet (chair)

Vladimir gave a presentation on one-way delay measurements that he had conducted between CESNET and HEANet (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/smotlacha.ppt>). This investigated the use of different NTP server configuration when conducting traffic measurements.

Victor gave a presentation on the Ipanematech measurement topology (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/reijs-ipanematech.ppt>). This utilises a central management (ip|boss) station to collect data from a number of clients (ip|engines). Metrics can then be made available for any unidirectional flow between two sites, and sampling can be configured for longer-term graphs. Time synchronisation comes from a GPS receiver or alternatively via ITS (a NTP-derived time protocol).

Nicolas gave a presentation on the multi-domain monitoring infrastructure being developed for GÉANT (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/simar.ppt>). This would offer one-way delay, ipdv and one-way packet loss measurements, SLA verification, and traceroute on a per-link, edge-to-edge or end-to-end basis. It was hoped to have this operational by the third quarter of 2003, working in partnership with the NRENs and with feedback from the community.

Victor then discussed the establishment of a Performance Response Team (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/reijs-prt.ppt>). He proposed that a support structure should be established for end-users to help solve performance issues when using applications over networks. This would have a similar hierarchy to the existing CERT structure with international (e.g. DANTE, Internet2), national (e.g. NRENs), and institutional (e.g. universities) levels.

This initiative would require coordination of people with experience of applications, operating systems and networks, and allocation of queries to the right level for resolution. It should initially be 'voluntary' and concentrate on a particular problem area (e.g. GRIDs, Premium IP) in order to gain awareness of the issues. At a later stage, the structure would be formalised.

The following organisations had expressed an interest in joining the PERT initiative: CESNET, DANTE, DFN, GARR, HEAnet, SWITCH, TERENA and UKERNA. It was proposed to draft some terms of reference, investigate further whether a CERT-like structure was appropriate, develop an end-to-end performance cookbook, and closely coordinate with Internet2.

Sven further outlined the proposal for an End-to-End Performance Cookbook (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn9/ubik.ppt>). The aim was to formulate a set of recommendations to fully utilise the capacity of long-distance high-speed networks, and to contribute to end-to-end performance research. This would be made available online, and would be updated at regular intervals.

Andrei presented a study based on flow measurement to gain knowledge about network quality on an NREN/GÉANT link, to define a working area with minimal number of dropped packet and point of overload.

The percentage of UDP versus TCP flows was discussed, as well as effectiveness of IP over different transport layers.

The main remark was whether the work could be related to the TF-NGN activity on over provisioned-networks performance.

9. Multicast, Ladislav Lhotka, CESNET (chair)

A test plan for SSM has been defined and is available via <http://www.cesnet.cz/tf-ngn/multicast/ssm>. The work done so far regards implementations of IGMPv3 and PIM-SSM, definition of scenarios and applications study. Support for IGMPv3 is available on Linux. A kernel patch is available from Cisco 2.4.19.

Experiments with test utilities have been done between Czech Republic (sender) and Switzerland, Poland and Norway (receivers). The utilities code has been ported to FreeBSD.

The implementation does not support IPv6, however it was noted that a patch is available for the KAME stack, Lada agreed to investigate the possibility to change the implementation to support IPv6.

ACTION, Lada to investigate the possibility to support IPv6 in the SSMtester.

A new version of the Beacon client implementation with improvements on delay calculation is currently available. Szymon provided an update on the development efforts at PSNC. The new features include IPv6 support and an mtrace module. For IPv6, Java version 1.4 is required. The mtrace module supports only IPv4 and uses external mtrace tools. The new client has already been tested by PSNC, others were invited to test it. More information is available at <http://noc.man.poznan.pl/noc/index/strony>

10. Optical Networking

10.1 Long distance fiber connections in NREN, Stanislav Sima, CESNET

Customer owned fiber (COF) networks give NRENs independence from carriers, as well as opportunities to deploy lambda services, BoD etc.

Research is needed on the scale of fiber (lambda) connections on the research and education network segment, both at national or regional levels.

Equipment and services require big investment in optronics and have significant maintenance cost. Issues to be investigated regard leasing or self maintaining fibers.

The equipment needed away of POPs depends on the type of fiber connection: metropolitan, long-haul, etc. in general:

1. dispersion compensator,
2. optical amplifiers,
3. optical regenerators.

The CESNET test programme consists of:

- long-distance GE between Ostrava and Olomouc running since January 2002, 148 Km leased fiber in line, L2 switches used as regenerators
- since May 2002 189 km leased fiber G.652 between Praha and Pardubice, using Cisco GSR 12016 KEOPSIS optical amplifiers – recommended line-end application for 1 and 2.5 Gbps lines
- connecting to Netherlight/ Starlight from December 2002
- Long distance 10 GE fiber lines
- Available Cisco 10 GE cards long-reach (40 km) reach – interested to test equipment of other vendors

10.2 GRNET Activities on Optical Networking, Dimitrios Kalogeras, GRNET

GRNET participates in the OTE-WAVE project providing a testbed to evaluate the status of DWDM equipment, integrated IP and DWDM management, interoperability tests and exploring services like Ethernet over SDH.

The OTE network is a ring of nine OADMs (Marconi PMA-32) used to provide Gigabit services to GRNET.

The test programme includes integrated network management, providing end-to-end MPLS paths over optical SNC, guaranteed bandwidth, configuration, fault and performance management.

Measurements of 10 Gbps DWDM links will be performed paying special attention to attenuation, PMD and chromatic dispersion.

GRNET participated in iGrid2002 with an experiment with intelligent multiplexing equipment using Sorrento Gigamux equipment which proved to be very good in terms of operation but was not free from management problems.

10.3 Progress on ASTON Activities, Valentino Cavalli, TERENA

The ASTON EoI was presented at a Lion/OPTIMIST II meeting in Turin on 15 October together with other EoIs in the the same area. The ASTON work-programme turned out to match the requirements of the next generation of GÉANT, and other proposals from PSNC, Cisco and Alcatel expressed an interest in joining it. At the meeting, the Commission explained that projects for Research Networks Testbeds will be subject to the second IST call in summer 2003, with projects starting mid 2004, and the ASTON representatives there agreed on a plan to carry out the activity within TF-NGN as already planned in view of preparing work and having a basis for the project.

In detail, the TF-NGN optical subgroup led by Victor Reijs will have three activities: Bandwidth on Demand, led by Michal Przybylski, 10 GE over long-distance, led by Stanislav Sima and Transport at 40+ Gbps led by Mauro Campanella. Each leader will find people from NRENs, universities vendors and carriers willing to cooperate and provide a workplan within four weeks after the meeting. Representatives from Juniper Networks, GRNET, DANTE and TERENA attending the meeting agreed to collaborate.

ACTION ASTON Leaders + Jean Marc, Dimitrios, DANTE, TERENA provide work-plans of activities within 15 November 2002.

Besides strengthening the TF-NGN activities the optical subgroup would continue planning the project preparation in view of submission to next suitable FP6 calls.

10.4 BoD Brainstorming Meeting, Victor Reijs, HEANet

A brainstorming meeting on Bandwidth on Demand was held at TERENA on 9 September 2002 and the outcome was widely discussed on the ASTON mailing list. Meeting report and presentations are available at:

<http://www.terena.nl/tech/projects/testbed/meetings/2002-09-09/BoD-meeting.html>

Michal Przybylski prepared a presentation (available via <http://www.terena.nl/tech/task-forces/tf-ngn/presentations9.html>) on the issues and outcome of the meeting and outlined possible tests utilizing national testbeds and GÉANT.

11. SERENATE Project, Valentino Cavalli, TERENA, Roberto Sabatino, DANTE

SERENATE ([http:// www.serenate.org/](http://www.serenate.org/)) is an EC project delivering a set of strategic studies into the future of research and education networking in Europe, addressing the local (campus networks), national (national research & education networks), European and intercontinental levels. The project is led by TERENA with participation from DANTE, the Academia Europaea, the European Science Foundation and the Technical University of Denmark.

In November 2002 DANTE, TERENA and a few experts from European NRENs will hold bilateral meetings with leading manufacturers of optical switches, routers and optical transmission equipment in the context of a SERENATE activity assessing the availability and characteristics of the necessary equipment for large-scale long-distance networks operating at 100 Gb/s or above.

Valentino provided an update about the status of the invitations. Roberto presented an overview of the questions to be used in the interviews and asked the audience to provide feedback. The questionnaire was sent to the TF-NGN list.

12. Update on TERENA Mobility Activity, Valentino Cavalli, TERENA

TERENA is organizing the creation of a task force on mobility with the goal of fostering the creation of an architecture to support "roaming" of students between different NRENs in Europe. The work to be done includes the study of available Authentication & Authorisation techniques (Web-based, RADIUS+802.1x, VPN), the study of next-generation-equipment support for MobileIP (both v4 and v6) and the set up a test bed for inter-NREN Authentication & Authorisation. More information is provided at <http://www.terena.nl/tech/mobility/>

Rudolf Roth mentioned that FhG Fokus is participating in the Moby Dick project that is doing similar things, he would get in touch with the mobility group.

13. Date of next meetings

The 10th TF-NGN meeting will be held in the first week of February 2003 in Rome, Italy, hosted by GARR.

The 11th TF-NGN meeting will be held in June 2003 in Cambridge, UK, Hosted by DANTE.

14. Actions from previous meetings

4.2 Simon and Victor to work on specification of the monitoring system for premium IP service.

- DELETED

5.2 Tiziana to coordinate discussion about loan of Equipment.

Discussion with Juniper Networks was taking place at TF-NGN 7 in Southampton

- DELETED

6.4 Dimitrios to poll the TF-NGN email distribution list to find out how many RIPE boxes are installed in NREN POPs and then define how they can be used in the OvPN performance measurement.

- DELETED

15. Open actions

9.1 Valentino and Roberto to update the list of work items in the TF-NGN ToR.

9.2 Lada to investigate the possibility to support IPv6 in the SSMtester.

9.3 ASTON Leaders + Jean Marc, Dimitrios, DANTE, TERENA provide work-plans of activities within 15 November 2002.