

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 8th TF-NGN meeting

1-2 July 2002,

T-Systems Nova GmbH, Berlin.

Valentino Cavalli, Kevin Meynell,

Issue 2, 21 August 2002

Attendees

Name	Organisation	Country
-----	-----	-----
Andreas Betker	T-System	Germany
Wim Biemolt	SURFnet	The Netherlands
Artur Binczewski	PSNC	Poland
Olaf Bonness	T-System	Germany
Mauro Campanella	GARR-INFN	Italy
Graça Carvalho	Cisco Systems	Portugal
Valentino Cavalli (Secr)	TERENA	The Netherlands
Tim Chown	Univ. of Southampton	United Kingdom
Joel Corral	ENST-Bretagne	France
Dirk Franke	DLR	Germany
Peter Feil	T-Systems/Berkom	Germany
Tiziana Ferrari	INFN-CNAF Bologna	Italy
Bartosz Gajda	PSNC	Poland
Tina Hausmann	JOIN / Univ. of Münster	Germany
Avgust Jauk	ARNES	Slovenia
Joop Joosten	CERN	Switzerland
Andreas Kaatz	DLR	Germany
Olav Kvittem	UNINETT	Norway
Roman Lapacz	PSNC	Poland
Simon Leinen	SWITCH	Switzerland
Ladislav Lhotka	CESNET	Czech Republic
Athanassios Liakopoulos	GRNET	Greece
Kevin Meynell	TERENA	The Netherlands
Christian Müller-Böhm	JOIN/Univ. of Münster	Germany
Julio Orozco	ENST Bretagne	France
Agnès Pouélé	DANTE	United Kingdom
Michal Przybylski	PSNC	Poland
Erik Radius	SURFnet	The Netherlands
Jürgen Rauschenbach	DFN	Germany
Victor Reijs	HEANet	Ireland
Esther Robles	RedIRIS	Spain
Rudolf Roth	Fraunhofer FOKUS	Germany
Roberto Sabatino (Chair)	DANTE	United Kingdom
Rina Samani	UKERNA	United Kingdom
Yves Schaaf	RESTENA	Luxembourg
Christian Schild	JOIN/Univ. of Münster	Germany
Roland Schwaiger	T-Systems/Berkom	Germany
Stanislav Sima	CESNET	Czech Republic
Nicolas Simar	DANTE	United Kingdom
Vladimir Smotlacha	CESNET	Czech Republic
Miguel Ángel Sotos	RedIRIS	Spain
Robert Stoy	DFN	Germany

André Stolze	JOIN/Univ. of Münster	Germany
Matjáz Straus	ARNES	Slovenia
Tomasz Szewczyk	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

Wilfried Wüber	ACONET/ Univ. of Vienna	Austria
----------------	-------------------------	---------

Online presentations

<http://www.terena.nl/task-forces/tf-ngn/presentations8.html>

1. GÉANT update and technology roadmap discussion

Roberto presented the agenda and announced the date of the next meeting then he updated the audience about the status of GÉANT connectivity. Two STM-4 to Zagreb were under test. The STM-1 between Israel and The Netherlands was in place whereas the link between Israel and Italy was still under test. A contract for connectivity to Malta was being finalized. The tender for connectivity to Ireland had been issued and DANTE was busy in evaluating the offers received.

GÉANT has not been affected by Teleglobe and WorldCom, but the KPNQ problems had an impact on commodity IP and North American research connectivity. Roberto said the circuits were still running but there was no guarantee about when they would quit. In order to cope with the issue of connectivity replacement DANTE had established agreements with Deutsche Telecom and Level3 to obtain extra circuits to North America. Telia was replacing Commodity IP services in the UK and Germany, emergency services were provided to ACONET and CESNET, SANET being a third candidate for that.

A few deliverables were still due in 2002. D13.2 describing the GÉANT roadmap for year3 was due in July. Four more deliverables, 9.6 to 9.9, respectively on ipv6, measurement, capacity management and LBE were due in August. Roberto reminded the audience that a new review panel had been established and that an outline of the deliverables needed to be sent to the panel two to three weeks before the submission deadline.

A proposal for new activities on the GÉANT roadmap for year3 was thoroughly discussed. The agreed work items are the following.

- Start deploying an intra/inter-domain Network Monitoring Infrastructure. The goal is not to debug the network, but to understand the requirements and analyse the measurement results. Performance analysis should be done in a user-to-user fashion and also include application monitoring. Poland, Norway, UK and Switzerland said they will contribute.

- Definition and organisation of a Performance Emergency Response Team (PERT). End-to-end performance analysis cannot be managed by the NOC only, the PERT should also involve application people and other entities responsible for portions of the end-to-end path. The goal is to define and provide a sample implementation, working with a limited number of user cases. TF-NGN will not operate a PERT, but will (try to) show how it should be operated by NOCs and other entities as may be defined. This activity should be coordinated at a wide level and liaise with work done by Internet2 and the Grid community. Italy, Switzerland, UK and DANTE were interested in collaborating.
- Optical developments. It was agreed that even if the bulk of work might be done in other projects there is interest in having a forum to discuss things. Italy, Poland, Czech Republic, Germany and DANTE were interested in that. Victor Reijs was not present at the session but the group made the assumption that Ireland would be interested too.
- Multicast. DANTE, France, Czech Republic, Germany, Poland and UK would work on it.
- Many people from different countries in the audience were interested in working on IPv6.
- Support for other projects. GÉANT has defined a Premium IP service which helps some projects, but not all of them. In other cases layer2 networking may be needed (like ATRIUM, 6NET and DataTAG) and there might be other cases which need something even different. It was questioned why does GÉANT need to focus on requirements for connection oriented Layer2 networking and suggested that the projects could access the GÉANT resources at the IP layer. It was observed that GÉANT is committed to support user projects, although the Technical Annex does not clarify the business criteria to decide the nature of the support to be provided. The group said they would work with up to 10 projects. Italy and DANTE wanted to work on that.

The proposed list of items included the development of an AA architecture on GÉANT. It was noted that different AA systems are being developed by different communities, Grid, mobility etc. and there is a need to bridge these different systems. Since this work is carried out by a number of other TERENA Task Forces it was agreed to initially collect information about their activities and report from experiences of other groups every now and then at TF-NGN meetings.

2. Premium IP and LBE

Nicolas gave an update on Premium IP and LBE on GÉANT (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations8.html>). The SPC had collected information for service requests, and had circulated recommendations and guidelines. Questionnaires were being sent out to the various user groups requesting contact information, required topology and timeframes, and requested capacity. In addition, the SEQUIN website (<http://www.dante.net/sequin/>) contained a variety of configuration examples, FAQs and How-Tos.

Two user projects were currently being supported. The AQUILA project was developing an QoS enhanced architecture and was utilising a connection between

Warsaw and Vienna at 2 Mbps. The MOICANE project was developing QoS support in access technologies and was utilising a connection between Pisa and Lisbon at 1.5 Mbps, and between Pisa and Bucharest at 2 Mbps.

The LONG project, which was working on IPv6, would also be supported shortly. This would involve connecting seven sites in Spain and Portugal at 2 Mbps each. Premium IP would be supported across GÉANT and FCCN, but not initially across RedIRIS as the new network was still being set-up.

A few improvements still needed to be made. A central repository for contact and equipment information was desirable. In addition, a tool needed to be developed to monitor OWD, IPDV and OWL, as well as checking available bandwidth.

LBE packet marking on Juniper M Series routers was discussed. The aim was for LBE to be transparent, but this difficult to achieve due to the manner in which the router output queues were configured. Various solutions were therefore suggested.

Tim said the GÉANT deliverable on LBE (D9.9) was due on 31 August 2002, and a decision needed to be made on how to proceed with the workplan. The technical issues had already been discussed at the TF-NGN meeting in April, and a BoF had also been held at TNC 2002 in June.

The participants and application areas now needed to be identified, as the GÉANT core needed to be configured to enable experiments between NRENs and end sites. The original workplan was therefore presented for discussion (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/tfngn8-lbe-v2.ppt>).

Mauro felt too much time was being spent on LBE rather than Premium IP. Simon replied that LBE testing was important as it had actually been shown to improve Premium IP performance.

Roberto asked who besides DANTE would be willing to put effort into developing provisioning tools. GRNET, PSNC and UKERNA volunteered some effort towards this.

Tiziana asked for clarification as to the type of provisioning tools that would be developed. Roberto replied these would be used for changing router configurations.

Tiziana then gave a presentation on LBE testing across GÉANT (see <http://www.cnaf.infn.it/~ferrari/infn-grid-wp5/task-dg/task2/lbe-test-plan.pdf>). The aim was to verify the LBE configuration adopted in the PoPs, evaluate how well best-effort traffic was protected from LBE traffic, and to examine the effect of congestion in LBE queues. In addition, SNMP monitoring of QoS parameters was undertaken.

Mauro pointed that Cisco GSR routers have around 700 MB of buffer memory which potentially introduced delays of 20-50 ms. This needed to be taken into account when conducting the tests.

Esther asked whether these tests would congest NREN connections. Tiziana replied this was inevitable and NRENs would need to agree to this. Agnes suggested 'windows' when tests could be scheduled without causing too much disruption.

3. QoS monitoring discussion

Victor reported about experience with proprietary and NTP based tools for OWD measurement.

An expensive proprietary tool like Chariot, from NetIQ, provides active measurement and has a plug-in to do OWD measurement. A problem with it is that measurement of asymmetric paths is not accurate because clock synchronisation is based on the same path to be measured.

Ipanematech provides passive measurement boxes working on Ethernet, but are not able to provide measurement on GigE. Clock synchronisation uses GPS, Victor was not sure about the accuracy of measurement in case of asymmetric path.

The NTP based OWD measurement method developed by Vladimir Smotlacha provides NTP clients near GPS clocks at domain/country pops. Usually systems next to them are synchronised quite accurately, however there are some limitations due to NTP, like clock adjustment and PC internal clock. Vladimir had written a piece of software to correct these limitations. According to initial tests the accuracy is less than 1 millisecond, but some more tests were still needed. Another method, used in the US, is based CDMA, but it uses a different clock synchronization mechanism and it is not clear if it can be used in Europe.

Victor did not find tools to measure available bandwidth yet.

Victor asked for experience from other NRENs. Steve said that Cisco Internet Performance Monitor is used by UKERNA to measure RTT, showing more packet loss than expected.

Sven Ubik from CESNET made a presentation about the way to achieve and to monitor TCP performance. He studied the behaviour of TCP on long network links and found normally a very low throughput due to small TCP window with any standard off-the-box OSs. It is possible to increase the throughput by changing the Window Scale TCP Option as defined in RFC 1323.

Sven showed UDP throughput between CESNET and UNINETT via GÉANT and Teleglobe, he would like to make tests with other NRENs too. The measurement was done using the QoSplot statistic tool designed in house by CESNET. He found that by increasing the TCP windows the throughput increases significantly. The tests showed a difference in increase on the link UNINETT-CESNET with respect to the one CESNET-UNINETT. Mauro suggested to look at the Linux counter to investigate the reasons.

Sven presented some results of other measurement tools for simulation (ns/2, see <http://sim.cesnet.cz>), emulation (Nist Net), path capacity estimation (pathchar and

pathrate), capture and follow up analysis (tcpdump plus tcptrace) and on-the-fly monitoring of TCP state variables (web100).

Sven plans to carry out some further research on end-to-end performance measurement and to use SACKs in order to avoid slow start after timeout.

Simon updated the audience about recent Flow Measurement developments. The IETF ipfix working group is quite active with its requirements document very close to last call (<http://www.ietf.org/internet-drafts/draft-ietf-ipfix-reqs-05.txt>). He recommended all NRENs who are doing accounting to make sure that they meet the requirements. Other interesting documents available regard application/applicability and architecture model for IP Flow Information Export (see <http://www.ietf.org/html.charters/ipfix-charter.html>).

4. Ipv6

Tim gave an update on the IPv6 activities (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/tfngn8-ipv6-v2.ppt>). There were a number of ongoing tasks within GTPv6, although most NREN effort was understandably going into the 6NET project (<http://www.6net.org/>). This included multicast and interoperability tests over the GTPv6 test network as GÉANT deliverable (D9.6) was due on 31 August 2002.

A workplan was also being developed for running IPv6 over GÉANT. The aim was to migrate to dual IPv4/IPv6 stacks within the next 12 months, so a test programme needed to be defined.

The 6NET core network was now up-and-running. The next stage was to bring specific user projects online via the national networks.

Several useful documents were available on the 6NET website, including a review of MIPv6, a description of the routing architecture, the list of applications to be run on 6NET, and the presentation from the Joint 6NET/Euro6IX Workshop. The transition review for the core network, NRENs and end-sites was still ongoing.

The Euro6IX project was deploying IPv6-enabled Internet Exchanges across Europe. These would eventually be linked at 34 Mbps, but this was not due to happen until August 2002. 6NET would be connected to Euro6IX via the exchanges in London and Paris. 6NET and Euro6IX were also committed to producing a joint deliverable.

There were several other IST projects in the IPv6 Cluster coordinated by 6LINK. These would be able to take advantage of the facilities offered by 6NET. Other initiatives included the IPv6 Task Force and the RIPE IPv6 Working Group.

Tim also mentioned that an Expression of Interest had been submitted for the IST Sixth Framework Programme.

Bernard gave an update on the GTPv6 topology (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/GTPv6-2.PDF>). CERN, CESNET, HEAnet, PSNC, RCCN, RedIRIS, RENATER, RESTENA and SWITCH were currently

connected to the Juniper M5 router in Paris, which was in turn connected to the Hitachi router and the P6R2 Internet Exchange. RENATER was providing support for the M5 router.

Anyone wishing to run tests over the GTPv6 network was asked to submit a short proposal outlining the nature of the activity, the participants and the resources required.

The M6Bone had approximately 20 connections in over 6-7 countries. Tests had already been conducted with PIM SM and multiple RPs, and further tests were planned with PIM SSM when the routers supported this. A mailing list (m6bone@ml.renater.fr) had been created for discussing these tests.

Bernard also took the opportunity to announce an IPv6 Conference that was being jointly organised by G6, RENATER and Aristote. This would be held on 29-30 October 2002 at ENSAM, Paris and would mainly focus on IPv6 applications. In addition, tutorials were planned for 28 October 2002.

Bartosz gave a presentation on the IPv6 network in Poland (http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/IPv6_network_in_Poland_TF-NGN-June.pdf). This was a production-quality network based on Cisco networks, with external connections to 6NET (via NORDUnet), the 6Bone, and several other IPv6 networks. There were also plans to use a separate lambda on the PIONIER network for IPv6 traffic.

A number of services (e.g. DNS and web) were being run successfully, and networking monitoring (e.g. MRTG, Looking Glass and ASPath) had been implemented. Trials were also being conducted with a local cable television provider to provide hundreds of end-users with IPv6.

Tomasz reported on the multicast testing between PSNC, RENATER and ULB (http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/IPv6_multicast_TF-NGN_Berlin2.ppt). This started on April 2002 and utilised routers based on FreeBSD 4.5 running the KAME stack. Several protocols were tested including PIM-SM and RIPng.

The next phase was running multicast over the 6NET network, but this first needed to be supported by the Cisco routers. High performance tests with high-speed multimedia streams and videoconferencing were also planned.

Christian gave a presentation on 6WiN in Germany (<http://www.join.uni-muenster.de/info/slides/schild/tfngn8-Dateien/frame.htm>). The 6WiN core network connects five major cities and Münster (where the JOIN project is based) at 34 Mbps. The Cisco 7206 routers run native IPv6 only, with DFN customers having tunnelled connections via the G-WiN network. There are external connections to the 6Bone via Münster, 6NET, Deutsche Telekom (in conjunction with the IPv6 Showcase project), and other commercial IPv6 providers via DECIXv6.

Finally, the meeting discussed the proposed IPv6 activities for Year 3 of the GÉANT project. The focus would be on IPv6 deployment and many of the architectural

requirements would be identified in the context of 6NET project. However, the 6NET network was designed to be IPv6 only, whilst GÉANT would need to run both IPv4 and IPv6. The group needed to identify what specific aspects needed to be tested, with a view to enabling IPv6 on GÉANT by the end of Year 3.

5. Optical

Roland Leners from PhotonEx <rleners@photonex.com> made a presentation introducing PhotonEx, and its systems for DWDM transmission to TF-NGN. Details of the presentation were confidential and were therefore not minuted.

Michal presented an update about the tests done at PSNC with the Alcatel 1696 metro span (metro DWDM). He briefed the audience about the key features and the PSNC configuration. They tested both the physical layer and the management features, detecting no bit-rate errors in the lab. The equipment is still available for others to test at PSNC for a short time.

Mauro introduced plans about the activity on testing 40 Gbit/s and beyond in TF-NGN. A web page is under development, and comments will be requested when it is ready. The goal is to understand the technology, do some tests and understand monitoring and behaviour of routing protocols. The rationale for the activity lies in increased request for bandwidth. As an example Mauro said that bandwidth bought by GARR users in Italy scaled up a factor of 10 in 3 years, but user traffic in the last 1.5 years increased of a factor of 5. Traffic patterns are changing too, moving fast towards bulk data traffic etc. A simple way to address these phenomena has to be found. To increase the link capacity is one way - multiple 10G links are already available - but you have to be careful at which layer you are multiplexing. Mauro added that packet switching is not going to be there before 5 years from now and he was interested in digital wrappers, to find out a valid alternative to SDH. He said he will investigate further into that and will report at next TF-NGN meetings.

Victor gave an update about the ASTON project. An Expression of Interest was submitted to the European Commission based on the work done before, with the addition of one activity on optical networking, partly including optical switching. He showed the list of work packages, time line and budget as estimated in the EoI.

Jürgen commented about the budget wondering if that was appropriate to the project size. Victor replied that a scalable scheme would be adopted: not many partners would be involved at least in the most research-oriented tasks.

Victor asked whether to involve users in the project or not. It was noted that the project does not aim at testing application and providing a test kit is probably the best thing to do.

Simon remarked that the project needs to focus on the perspective of research networking operators and should not try to compete with Telecommunication companies on the transmission tests. An interesting point, for instance, would be how to integrate optical technology with the current management systems. Victor suggested to work together with Simon in getting the point through.

6. Final session

Lladislav reviewed the status of multicasting on GÉANT (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/berlin.pdf>).

This was now fully operational in the backbone and most NRENs. Users reported that it largely seemed to work, but tended to be unreliable. There was also a lack of attractive applications and content.

The next phase was to implement a Beacon-like infrastructure as well as test other tools. PIM SSM would also be tested when IGMPv3 implementations became available, and integration with the IPv6 multicast was planned. User support also needed to be improved, and some higher bandwidth tests run.

Steve said a number of applications were capable of using multicast if it was available, but they reverted to unicast in its absence.

Simon reported that ETHZ in Switzerland had built a multicast network on their campus to show World Cup matches from four television channels. A few problems had been encountered, but they managed to establish a useable system for over 1,000 users which would have been difficult without multicast. The success of the initiative was demonstrated by the fact that several professors complained that their staff and students were watching the matches!

Roman Lapacz then gave a presentation on the Multicast Visualisation Tool (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn8/MUVI-3.ppt>). This displayed multicast topologies, node information, and MIB details in a graphical format. Planned enhancements included the collection and reporting of statistics, the ability to save topological views, and support for IPv6.

7. Date of next meeting

The 9th TF-NGN meeting will be held at NIIF in Budapest, Hungary on 17-18 October 2002.

8. Actions from previous meetings

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

- ONGOING

5.2 Tiziana to coordinate discussion about loan of Equipment.

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

- ONGOING

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.

- ONGOING

9. Open actions

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

- ONGOING

5.2 Tiziana to coordinate discussion about loan of Equipment.

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

- ONGOING

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.

- ONGOING