

# TERENA TASK FORCE ON NEXT GENERATION NETWORKING

## Draft Minutes of the 11th TF-NGN meeting

8 May 2003,

Poznan Supercomputing and Networking Centre, Poznan, Poland

Valentino Cavalli, Kevin Meynell,

Issue 1, 25 June 2003

### Attendees

Name	Organisation	Country
-----	-----	-----
Wim Barbaix	Alcatel	Belgium
Artur Binczewski	PSNC	Poland
Valentino Cavalli (Secr)	TERENA	The Netherlands
Tim Chown	Univ. of Southampton	United Kingdom
Bartosz Gajda	PSNC	Poland
Marcin Garstka	PSNC	Poland
Avgust Jauk	ARNES	Slovenia
Marcin Kaminski	PSNC	Poland
Joannis Kappas	DANTE	United Kingdom
Radoslaw Krzywania	PSNC	Poland
Felix Kugler	SWITCH	Switzerland
Olav Kvittum	UNINETT	Norway
Simon Leinen	SWITCH	Switzerland
Roland Leners	PhotonEx	Belgium
Matti Lentinen	Politech. Of Oulu	Finland
Athanassios Liakopoulos	GRNET	Greece
Ladislav Lohtka	CESNET	Czech Republic
Mark Lutz	Fraunhofer FOKUS	Germany
Kevin Meynell	TERENA	The Netherlands
Stefan Piger	University of Hannover	Germany
Wiktory Procyk	PSNC	Poland
Michal Przybylski	PSNC	Poland
Jürgen Rauschenbach	DFN	Germany
Rudolf Roth	Fraunhofer FOKUS	Germany
Stanislav Sima	CESNET	Czech Republic
Nicolas Simar	DANTE	United Kingdom
Trond Skjesol	UNINETT	Norway
Miguel Ángel Sotos	RedIRIS	Spain
Christian Strauf	JOIN-DFN	Germany
Robert Szuman	PSNC	Poland
Jarkko Torvinen	Politech. Of Oulu	Finland
Szimon Trocha	PSNC	Poland
Stig Venaas	UNINETT	Norway
Jean-Marc Uzé	Juniper Networks	France
Fuhua Yin	STC-ULB	Belgium

### Apologies

Mauro Campanella	GARR	Italy
Roberto Sabatino (chair)	DANTE	United Kingdom

### Meeting proceedings

Presentations are available online at:

<http://www.terena.nl/task-forces/tf-ngn/presentations11.html>

## **1. GÉANT Update**

Nicolas Simar presented the latest upgrade of links in Europe, concerning RoEduNet access to GÉANT, and the circuits between Hungary-Slovenia and between Germany, The Netherlands and United Kingdom. He then showed the updates in connectivity to the USA and the latest upgrades of access capacity: GARR was the first NREN upgrading it to 10 Gbit/s.

The services provided on GÉANT include premium IP and IPv6 dual-stack, several NRENs are already connected.

The testbed routers are in place. These include Juniper M20, and Cisco series 7000 and 12000, connected via STM-1 links and can already be used by DANTE and NRENs, initially for multicast.

## **2. Performance Monitoring**

### **2.1 Performance Monitoring Infrastructure, Nicolas Simar, DANTE**

DANTE, HEANet and GARR have been designing a Performance Monitoring architecture to exchanging monitoring data between multiple domains, initially providing users with a view of edge-to-edge performance data, but to be later extended to end-to-end data. A trial is expected to start in September or October 2003. More contributors where sought.

The proposed architecture aimed at ensuring back-compatibility with existing monitoring infrastructures already in place at some NRENs and therefore the test architecture would provide interfaces to RIPE TTM boxes. Nicolas clarified that NRENs can choose different monitoring platforms, but the overall infrastructure must be standardised and able to talk to all of them.

Athanassios Liakopoulos remarked that RIPE TTM sends packets with low frequency and this is not enough to obtain useful information about SLA validation, he therefore objected about the rationale for supporting it. Nicolas stressed that this is an intermediate solution driven by the fact that RIPE TTM are already in place and that support would be limited to the interface layer only: the boxes would collect data and send it to the higher layer monitoring system, which will carry out the actual performance monitoring.

It was generally felt that the definition of interfaces was too immature to start implementing the architecture and a longer-term solution should have been defined. Developing a driver for RIPE TTM was considered not useful. The counterproposal was to aim at an open source measurement solution, develop a primitive measurement point agent and then provide a driver for it. Serious concern was expressed about starting implementation knowing that the tool is not perfectly defined, and more contribution by NRENs would be needed before implementation could start. PSNC and GRNET showed an interest in contributing to the discussion.

**ACTION:** PSNC and GRNET to be involved in the definition of the Performance Monitoring Infrastructure.

### **2.2 PERT, Simon Leinen, SWITCH**

The goals of the PERT activity is to understand how to address the end-to-end performance issue and build a cross-domain/cross-disciplinary team of experts with knowledge about performance at the levels of core and access networks, operating systems and applications.

The team should “own” the problem, which means, assume responsibility for detecting it and pointing the user to the right solution.

The 3<sup>rd</sup> and final version of the “manifesto” documenting the PERT project has been published on the web at <http://www.switch.ch/network/performance/pert/manifesto.pdf>. In order to move the project forward the group needs volunteers in the areas of expertise about campus networks, operating systems and applications. Funding is essential too. The attendees discussed a few opportunities and mentioned national funding that might be available in the UK from the National MBS provided by JANET and in Norway through the Giganet project or also the submission as an SSA to Commission FP6 calls.

### 3. IPv6

Miguel gave an update on the GÉANT IPv6 pilot service. The core routers had been upgraded to support IPv6 in February, and the first NRENs had been connected in March. FCCN, GARR, HEAnet, PSNC, RedIRIS, RENATER and SURFnet currently had native connections, whilst IUCC and LITNET were receiving tunnelled connections. There was also a native connection to Abilene in the US.

Other NRENs wishing to establish IPv6 connectivity needed to ensure their routers were running 5.5R3 (Juniper), or IOS 12.2(13)T and 12.0(23)S1 (Cisco), although there were still a few bugs that needed to be resolved. The routing policy between NRENs, Abilene/CANARIE/ESNET, 6NET and commercial providers (currently GX and Telia) was also explained.

The next stage of the pilot service was to establish DNS and monitoring facilities. The aim was to introduce a production IPv6 service by January 2004.

Christian then demonstrated GnomeMeeting which is an IPv6-enabled VoIP and videoconferencing application. This is based on the Gnome libraries, pwnlib andopenh323, and runs on Linux, FreeBSD, and MacOS X. It can be used with USB or Firewire-based webcams, and a full-duplex soundcard.

IPv6 support was added in April 2003, although it remains backwardly compatible with IPv4. The porting only required 10 lines of code to be changed, and most of the changes were in the supporting libraries. More information can be found at <http://www.gnomemeeting.org/>

Laura gave a presentation on running NTP over IPv6, and whether it affects synchronisation and convergence. Their tests were conducted with two NTP servers at Stratum 1, and an NTP client at Stratum 2. NTP Version 4.1.x was used, and the tests were run over native IPv4, native IPv6, and tunnelled environments.

It was shown that running NTP over IPv6 did not introduce any additional delay when compared with IPv4, and the synchronisation values appear to be better. However, tunnelled environments introduced an additional delay of 53%, and this affected the dispersion values. It might therefore be concluded that NTP peerings over IPv6 are more reliable due to the superior capabilities of the IPv6 stack, although more tests should be conducted to verify this.

Stig provided a brief update on the M6Bone. Native multicast was now being deployed on 6NET, and around 15 organisations had been directly connected. An onward connection to the global M6Bone was being provided by RENATER. In addition, multicast reflectors and an IPv4-IPv6 gateway were currently being investigated.

Finally, Marcin demonstrated a tool that PSNC were developing to enable existing IPv4 network management platforms to monitor native IPv6 networks. This translates SNMP protocol messages, ICMP pings between IPv4 and IPv6 networks. It does not interfere with

SNMP queries, and can be used to gather statistics and monitoring network devices. The tool is managed with a GUI front-end based on Java, SSL and X.509.

#### **4. Alcatel test results, Athanassios Liakopoulos, GRNET**

Athanassios presented the results of tests on the Alcatel A777 OBX equipment, which were done by DANTE and GRNET staff in March 2003 at the Alcatel labs in Antwerp. The tests covered performance issues, routing protocols, QoS and management support, but were also aimed to interoperability with other vendor's routers and getting hands-on experience. Some details, including the roadmap, are covered by NDA, the available results are explained in the presentation at [http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn11/20030508\\_AL\\_alcatel\\_test.pdf](http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn11/20030508_AL_alcatel_test.pdf).

The test will continue for two more months. It is planned to investigate into multicast features, MPLS (VPNs and Traffic Engineering) and accounting (Netflow), but this list can be extended. It would be interesting to get other NRENs involved in the experiments.

#### **5. Multicast**

Ladislav presented the next generation of the Multicast Beacon. He described the Beacon operation, the parameters of probe packets and the information provided by the report packets, which enables monitoring of average delay, jitter and loss rate as well as number of out-of-order and duplicate packets. The essential problems relating to the Beacon architecture need to be addressed and these require some sort of central control/agent coordination, authentication, good documentation of protocol etc. Lada explained the new features that would be implemented and discussed those that are not planned yet but would be desirable.

Radoslaw Krzywania gave a short presentation about MUVI a Multicast network monitoring tool developed by PSNC that allows to discover, monitor and visualize multicast network.

#### **6. GÉANT year-4 roadmap brainstorm**

Nicolas reported about the status of activities from the year-3 roadmap. IPv6 support is ahead of schedule, performance monitoring is on track, the PERT has made slower progress than it was hoped but a test case is expected, the router testing plan had been scaled down to test of the Alcatel A7770 only, optical activities in TF-NGN had been limited to information-sharing, but no joint hands-on experience had been really done, multicast has done little progress, premium IP management was starting in the 2<sup>nd</sup> week of May.

The principle informing the inclusion of research topics in the year-4 roadmap should be: relevance for GÉANT and NRENs, clear leadership, shared work and realistic approach towards achieving the results.

DANTE proposed to carry out the following activities from the year-3 roadmap: performance monitoring in closer cooperation with the Internet2 end-to-end Performance Initiative, PERT and router testing. New topics suggested by DANTE were enhancement of IPv6 monitoring and IPv6 multicast activities, resource reservation and end-to-end provisioning/monitoring of Layer2-based connectivity. The last one might have a two-phase approach, in the short-term it could use CCC or other Layer2 VPN technologies, in the long-term it could investigate into wavelength switching technologies.

The first comment was that all services should coexist and therefore they should be selected on the basis of some synergy.

People agreed on the continuation of Performance Monitoring, but stressed that more feedback must be sought from NRENS. Passive monitoring was felt as quite important and should be deserved further discussion. Full agreement was expressed on continuation of the PERT activity.

IPv6 should be part of the PERT and Performance Monitoring activities, and in general pervasive to all other activities. This could all go into a "advanced IPv6 services" category, where TF-NGN would act as a forum for discussion and coordination of this pervasiveness.

Several NRENS are getting practical experience with optical networking, TF-NGN should encourage joint collaboration and hands-on with DWDM technology, providing interesting links with equipment manufacturers and preparing for the next generation of GÉANT. Leasing fibre across borders between some countries in Europe for a limited time (6 months as an example) might be a viable option worth to pursue. It was questionable whether TF-NGN is the right framework for this, especially concerning where money to acquire the fibre would come from. Another question was about the business case for these tests. It was clearly said that GÉANT next generation would need some lambdas, however some people argued about just using infrastructure already available (or to become available in the short-term) and do some coordination in TF-NGN. Border hopping was considered an interesting issue, but it would need to address the problem of multi-domain multi-NREN management.

Router testing seems interesting, but this requires lots of resources and time. On the other hand routers functionality is improving/changing all the time, so that test results may become useless very rapidly. There was concern about investing resources in this direction. This activity should rather focus on providing requirements to router vendors.

IPv6 Multicast, was found an interesting activity. Stig Venaas said he would be available to lead it. People suggested to carry out testing at multiple NREN domain level. A general issue was whether it should be still part of IPv6 service enhancement or a separate item.

On the continuation of IPv4 multicast, there was some concern whether it is of any interest at all. Ladislav Lohtka said he would think about possible topics and would provide more information to the email distribution list.

**ACTION** Ladislav Lohtka to inform the TF-NGN email distribution list about possible topics for continuation of IPv4 multicast activity.

Resource reservation especially for support to Grid application: Jean-Marc Uzé was interested in participating. Somebody observed this activity might be outside the scope of TF-NGN.

As a further input to the discussion Michal Przybylski presented the slides on the ASTON White Paper that Valentino Cavalli had prepared from the GARDEN meeting held the day before in Brussels. The paper, available at <http://www.terena.nl/tech/projects/testbed/NREN-WhitePaper-v1.0.pdf>, defines lower layer research areas for the next generation of research and education networks. These are Optical infrastructure enabling core networks to operate at 40+ Gbit/s speed, Bandwidth on Demand, Simpler (than SDH) data link transport Routing model, Network monitoring and management, Multi-domain multi-vendor interoperability.

The discussion was largely dominated by the issue of optical switching as alternative to/in combination with packet switching. People argued about scalability of provisioning point-to-point switched connections, either using CCC or Layer2 switching and expressed doubts about their need.

## **7. Date of next meetings**

The 12th TF-NGN meeting will be held on 15-16 September 2003 in Cambridge, United Kingdom, Hosted by DANTE.

### **8. Actions from previous meetings**

NO Actions

### **9. Open actions**

11.1 PSNC and GRNET to provide input to the definition of the Performance Monitoring Infrastructure.

11.2 Ladislav Lohtka to inform the TF-NGN email distribution list about possible topics for continuation of IPv4 multicast activity.