

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 10th TF-NGN meeting 6-7 February 2003, University of Rome – La Sapienza, Rome, Italy

Valentino Cavalli & Kevin Meynell,
Issue 1, 21 February 2003

Attendees

Name	Organisation	Country
-----	-----	-----
Michael Allenby	UKERNA	United Kingdom
Lada Altmanova	CESNET	Czech Republic
Wim Barbaix	Alcatel	Belgium
Mauro Campanella	INFN-GARR	Italy
Massimo Carboni	GARR	Italy
Graça Carvalho	Cisco Systems EMEA	Portugal
Valentino Cavalli (Sec)	TERENA	The Netherlands
Tim Chown	University of Southampton	United Kingdom
Christian Cinetto	GARR	Italy
Jordi Domingo-Pascual	UPC/CCABA	Spain
Rob Evans	ULCC	United Kingdom
Sergi Figuerola	UPC/CCABA	Spain
Abraham Gebrehiwot	CNR-IIT, Pisa	Italy
Christian Grimm	University of Hannover	Germany
Liviu Iusan	RoEduNet	Romania
Avgust Jauk	ARNES	Slovenia
Dimitrios Kalogeras	GRNET	Greece
Joannis Kappas	DANTE	United Kingdom
Otto Kreiter	RoEduNet	Romania
Simon Leinen	SWITCH	Switzerland
Laura Leone	GARR	Italy
Athanasios Liakopoulos	GRNET	Greece
Tamas Maray	NIIF	Hungary
Kevin Meynell	TERENA	The Netherlands
Gabriella Paolini	GARR	Italy
Lorenzo Pinizzotto	CNR-IIT, Pisa	Italy
Hervé Prigent	CRIHAN	France
Michal Przybylski	PSNC	Poland
Jan Radil	CESNET	Czech Republic
Jürgen Rauschenbach	DFN	Germany
Victor Reijs	HEANet	Ireland
Duncan Rogerson	UKERNA	United Kingdom
Lorenzo Rossi	CNR-IIT, Pisa	Italy
Rudolf Roth	Fraunhofer FOKUS	Germany
Roberto Sabatino (Chair)	DANTE	United Kingdom
Michele Sciuto	GARR	Italy
Yves Schaaf	RESTENA	Luxembourg
Christian Schild	JOIN-DFN	Germany
Nicolas Simar	DANTE	United Kingdom
Trond Skjesol	Uninett	Norway
Vladimir Smotlacha	CESNET	Czech Republic

Miguel Ángel Sotos	RedIRIS	Spain
Christian Strauf	JOIN-DFN	Germany
Peter Tomsu	Cisco Systems EMEA	Austria
Hans Trompert	SURFnet	The Netherlands
Stig Venaas	Uninett	Norway
Chris Welti	SWITCH	Switzerland
Fuhua Yin	ULB-STC	Belgium

Apologies

Xiaoming Fu	University of Göttingen	Germany
Janos Mohacsi	Hungarnet	Hungary
Stanislav Sima	CESNET	Czech Republic
Jean-Marc Uzé	Juniper Networks	France

Meeting proceedings

All presentations are available online at:

<http://www.terena.nl/task-forces/tf-ngn/presentations10.html>

1. Welcome and introduction

The meeting was hosted jointly by the University of Rome "La Sapienza" and by GARR-INFN, the Italian national research and education network. Giuseppe Patroni, from INFN, welcomed the participants and briefly explained the importance of networking and international collaborations for INFN.

The chairman showed the latest updates to the meeting agenda and invited the attendees to a round of introductions.

2. GÉANT Update

2.1 International connectivity, Roberto Sabatino, DANTE

In January 2003 the ULAKBIM (Turkey) access circuit has been set up via a native STM-1 satellite link, the SINET (Science Information Network of Japan) STM-1 link has been put in service and the circuits to the Irish POP have been upgraded to 2.5 Gbps. In December 2002 SANET (Slovakia) access had been upgraded to 200 Mbps. Since October 2002 several circuits have been upgraded, including the links between Germany and Greece, Ireland-UK and Ireland-Germany. In December 2002 GRNET (Greece) and RoEduNet (Romania) had their DWS subscriptions upgraded.

QoS configuration has been completed on most of the GÉANT routers, only a few PoPs partially support or do not support QoS at all. Details on the status of QoS configuration on GÉANT are provided in the slides, which also describe the set up of the new GÉANT POP in NYC. The POP will move from 60 Hudson St. to 32 Avenue of the Americas. In the new configuration the GTRN router will become part of the GÉANT AS.

2.2 GÉANT IPv6 transition, Gabriella Paolini, GARR

Gabriella introduced the GÉANT IPv6 rollout plan, which foresees router configuration by February 2003, first connections of NRENS in March, a pilot service starting to become available in June and the start of operational service in January 2004.

The IGP is designed such that only ISIS is used for both IPv4 and IPv6, whereas external and internal BGP peerings are separate. Details of the separate BGP peerings, addressing scheme and routing policy are provided in the slides.

Currently nine NRENs are connected via native IPv6, five additional NRENs will be connected via tunnels during 2003 and later on.

The open issues are: finalisation of the routing policy with respect to Abilene, deployment, monitoring, DNS and definition of the service.

A task force has been created to address the transition process, which members are DFN, Renater, RedIRIS, GARR, NORDUnet and DANTE.

3. IPv6

3.1 Introduction and other items, Tim Chown, University of Southampton

Tim introduced the IPv6 session by mentioning that several 6NET work packages would be meeting on 12-14 February in Paris. He also provided an update on the ongoing IETF activities.

The IPv6 Operations Working Group (v6ops) had replaced ngrans since September 2002, and had agreed to discontinue the 6Bone as it was no longer deemed necessary. Both DHCPv6 and MIPv6 are approaching draft standard status, although the Multihoming Working Group (multi6) was still stalled. A couple of unofficial groups had also sprung-up called ipv6mh and ipv6flow.

The future of GTPv6 which had previously been used for inter-NREN tests was discussed. It was suggested this could be used to test hardware (and software) that wasn't already being tested on 6NET (e.g. Alcatel, Hitachi and Zebra).

Finally, it was mentioned that the RIPE Test Traffic Measurement system had been ported to work with IPv6. This had been trialed at the University of Southampton, the University of Vienna, HEAnet and the RIPE NCC since the beginning of the year, and was now available to new users.

3.2 IPv6-v4 multicast gateway, Stig Venaas, UNINETT

Stig gave a presentation on the IPv4-IPv6 multicast gateway. The intention was that IPv6 hosts should be able to send and receive data from any IPv4 multicast group. At the present time, a gateway can be deployed and used by an entire IPv6 PIM domain with no modifications to other software. The IPv4 and IPv6 multicast trees joined at this gateway which is an IPv4 leaf node, and an IPv6 RP. This solution allows bi-directional videoconferencing with both IPv4 and IPv6 participants.

3.3 M6bone, Lorenzo Rossi, CNR

Lorenzo gave a presentation on how to join the M6bone. He also provided an overview of the how CNR-IIT in Italy was connected to this.

3.4 Support for IPv6 on the Alcatel 7770 OBX, Wim Barbaix, Alcatel

Wim gave a presentation about IPv6 support in the Alcatel 7770 OBX switch. This was being used in the ATRIUM testbed which had recently been extended to include PSNC and Telefonica via GEANT and NRENs. PSNC would be using it for two GRID projects when the

platform became available in May 2003, and several experiments had been scheduled. Alcatel would also be discussing with DANTE and NRENs about their future IPv6 requirements.

3.5 6NET Update, Ioannis Kappas, DANTE

Ioannis gave an update on the 6NET project. A backup connection to PSNC had been established in early-October via a tunnel to Germany. CESNET had also been connected to Germany in October, whilst Hungary had been natively connected via an STM-1 at the end of January.

8FE cards were due to be installed on the 6NET PoPs in Austria, Germany and Sweden, and new 6NET routers were being delivered to the Czech Republic, Hungary and Poland. An STM-1 circuit between the Czech Republic and Germany was also due to be delivered in February. In addition, most the core routers had been upgraded to IOS 12.0(23)S1, although Austria, Germany and Sweden were still running 12.0(22)S due to MTU problems on tunnels.

Finally, 6NET had established peerings with Euro6IX via JANET, Abilene via SURFnet and NORDUnet, 6TAP and CANARIE via SURFnet, and ETRI via RENATER.

3.6 Dual stack pilot and IPv6 on GARR, Gabriella Paolini

Gabriella gave a presentation on GARR's IPv6 activities. Their IPv6 trial currently involved twelve local partners who were connecting via dedicated ATM PVCs. The plan was to introduce dual stacks on the core routers of their production network, which would provide IPv6 access to all sites on the GARR-G and GARR-B networks. This would allow the existing 155 Mbps ATM network to be replaced by 2.5 Gbps circuits. However, there were a number of technical problems to overcome with the Cisco routers before the migration could be undertaken.

A number of IPv6 training events were being held in Italy. Three had already been held in Florence, Rome and Turin with approximately 100 attendees at each, and two more would be held in Bari and Milan during February and March.

Network storage was being considered as an application that would help promote IPv6. GARR was collaborating with the University of Tennessee in the NSI-funded Logistical Storage Infrastructure project, which aims to provide content delivery on an IPv6 platform.

4. Optical Networking

4.1 The SWITCHlambda Project, Simon Leinen, SWITCH

SWITCH have changed their network backbone at the end of February 2003. All locations are connected via STM-1 links. In addition, since last year they have built a fibre pair between Zürich and Geneva. One fibre runs along the motorway and a second one, which is used to provide redundancy, runs along the railway. Since 2003 a third fibre connects Zürich to the supercomputer centre in Ticino, in the south of Switzerland.

SWITCH has acquired Indefeasible Right of Use (IRU) for the fibre running along the motorway between Zürich and Geneva. The investment looks very cost-effective, considering that the fibre will stay in operation for 20 years and the cost being equal to leasing an STM-1 connection for one year only. DWDM equipment has been deployed with routers installed at five sites along the path.

The SWITCHlambda project will investigate the risks and benefits of connecting all SWITCH customers via fibre in the next five years. The costs of providing dark fibre to small locations far away from the backbone seem very high, but one should consider that with the current

instability of the telecommunication market raise of bandwidth prices can be expected as well as changes of providers this entails new (expensive) local loops. Deployment of dark fibre can be seen as economically effective if considered as a long-term investment with cost depreciation of fibre and DWDM equipment over several years.

The project started with DWDM equipment and will provide CWDM to large backbone sites plus single point-to-point links to small sites. The system is designed to support 10Gbps, but it is not clear yet if single or multi channel.

Gigabit Ethernet is a much better technology than STM-16 SONET/SDH for the access network. It is less complex and less expensive and ensures best compatibility with SWITCH GigaMux equipment from Sorrento Networks (<http://www.sorrentonet.com>). However, 10GE WAN PHY interfaces are extremely expensive, whereas 10GE LAN PHY, though relatively cheap are not supported by SORRENTO yet.

The Sorrento GigaMux DWDM system has been chosen among a number of offers not only by its cost but also because of its limited Telco-style functionality and basic management tools available. The equipment supports 10Gbps transmission capacity and has built-in Chromatic Dispersion Compensation. It is able to provide an all-optical channel between Zürich and Geneva (350km) without electrical regeneration. The Span of amplifiers is up to 100km.

The bandwidth offer to customers will vary from 400Mbps up to Gbps granularity. 10 Gbps single channel has been tested with Force10 (<http://www.force10networks.com>) and Cisco equipment. New tests will be done in 2003 using 10GE LAN PHY cards and with 10GE line cards from Cisco 760x OSR.

4.2 The SERENATE study on equipment, Roberto Sabatino, DANTE

The SERENATE project (see <http://www.serenate.org>) has produced a study, which is going to be available in spring 2003 on the characteristics of equipment for next generation networking. The study is based on interviews to routers, switching and transmission equipment vendors plus two research labs carried out between November and December 2002 by TERENA and DANTE with consultancy from experts from CESNET, PSNC and HEANet.

The results of the study show some controversial conclusions about the availability of 40Gbps and its cost effectiveness, economical factors seem to favour its first appearance in Long Haul systems whereas some vendors predicted that it would be in the metropolitan area. Transmission at 80Gbps, 160Gbps is technically possible, but only in labs so far. Higher transmission capacities have been demonstrated too.

Routers and OXCs (Optical-electrical Cross Connects) will be ready for 40Gbps operation in the next few years. Interdomain interoperability is still questionable. Routers appear to be the main drivers for single channel 40Gbps transmission.

OXCs and PXCs (Photonic Cross Connects) can scale to hundreds of Gbps but have various proprietary implementations and are not interoperable at the moment.

Developments in transmission technology show that it is possible now to cover relatively long distances without need for electrical regeneration. The span of amplifier is also increasing allowing reach with Nothing in Line in some cases, as proven by CESNET. Altogether such developments enable a Do It Yourself approach for NRENs.

The SERENATE study investigates some alternative network architectures being enabled by technology developments and the consequences for network management by NRENs.

Victor Reijs remarked that the usage of terms like OXC and PXC does not match with the one used in the GÉANT/TF-NGN deliverable on optical networking. He acknowledged the reasons for SERENATE to use such terms but asked to make clear the difference.

4.3 i2CAT and CARISMA Projects, Jordi Domingo-Pascual, UPC

i2CAT is a project promoted by the Catalan government and is lead by UPC with partners from universities local and regional authorities, Telcos and service providers aiming at providing an experimental –separate network (GigaCAT, <http://gigacat.ccaba.upc.es>) for broadband applications and services. GigaCAT uses equipment from Nortel (Optera Metro) plus Cisco and Juniper, and provides the platform for projects like "Opera Oberta" and "Gaudi", sending uncompressed HDTV as well as multicast Mpeg-2 over optical fibre.

The CARISMA project is a more experimental optical test bed (dynamic provisioning), which will run until the end of 2004. The project has put in place a CWDM optical ring on the campus of Terrassa, a Catalan University and is doing experiments with Bandwidth on Demand on the backbone to provide optical channels end-to-end. The (static) i2CAT and the (dynamic) CARISMA networks are interconnected via a OXC located at CCABA.

On top of the experimental infrastructure UPC runs several projects including 6ONITE, which provides additional IPv6 support. UPC is willing to join research efforts with other similar experimental projects.

4.4 Optical Networking in CESNET, Jan Radil, CESNET

The presentation contains an update on the CESNET2 topology as of February 2003 and the expected fibre deployment in 2003. One major update since the previous TF-NGN meeting is the connection of CzechLight to NetherLight in Amsterdam via a 2.5Gbps wavelength service procured from Telia. CzechLight provides a "lambda service" by using Cisco ONS 15454 and is operational since January 2003.

CESNET had presented test results of 1GE at the TERENA Conference in 2002 and are now starting a 10GE test based on a similar scenario with Nothing In Line (NIL). They are currently negotiating contract for hardware components, which will include Cisco Catalyst 6503 with ER modules and Keopsys amplifiers (<http://www.keopsys.fr>), and preparing the fibre test bed. More details will be presented at TNC2003.

5. QoS Monitoring

5.1 Progress of PERT activities, Victor Reijs, HEANet

Victor explained that the aim of PERT was to provide a proactive support structure for end-users to help solve performance issues when using applications over a network. It would be similar to a CERT and would have a number of link levels (e.g. international, national, institute and user).

The terms of reference were currently being defined with the aim of initiating a pilot project in the IST Sixth Framework. The most appropriate EC instrument (e.g. in which category the proposal should be submitted) still needed to be identified, but a two-page project description would be available by 28 February 2002 for prospective partners. It was hoped the full project proposal would be ready by 1 April 2003.

5.2 Performance monitoring activity update, Nicolas Simar, DANTE

Nicolas gave a presentation on the performance monitoring activity. This aims to monitor performance (e.g. one-way delay, jitter, packet loss, available bandwidth, etc) within and across domains, with a view to providing different groups of users (NOCs, GRIDs and users) with an overview of network performance. Measurements are collected and stored in each domain and exchanged when requested by another domain.

This meant there needed to be standard way of measuring networks and exchanging data between domains. There were looking at what the IETF IPPM Working Group and the End-to-End Performance Initiative were doing, but he wondered whether there were any other activities in this area.

The types of traffic that needed to be monitored included best effort, premium and LBE over both IPv4 and IPv6, but the data needed to be presented in different ways to the end-user, the network planner, and the NOC. The type of statistics to be collected and the sampling techniques also needed to be defined.

A draft specification of the domain central tool that would retrieve domain measurements and exchange them with other domains, had already been published, and the aim was to finalise this by the end of March. Any input was welcome and a mailing had been established to discuss requirements. The activity was also looking for at least three consecutive domains that had implemented MPs in order to trial the system. GÉANT, GARR, HEAnet and SWITCH had volunteered, whilst the CESNET, GRNET, SURFnet, PSNC and UKERNA also expressed interest in being involved.

5.3 TraceQoS, Chris Welti, SWITCH

Chris gave a presentation on TraceQoS which is an active one-way measurement system for locating QoS failures on an IP path from source to destination. It allows for on-demand tests by measuring OWD, OWDV and OWPL on a segment-by-segment basis, and making concurrent measurements on the subpaths.

5.4 QoS measurement for IPv6-IPv4, Jordi Domingo

Jordi provided an overview of UPC's QoS measurement activities. The objective of MIRA was to develop a flow-based traffic analysis and classification tool which would gather knowledge of how users were using network resources, apply a usage-based billing scheme for cost-sharing, and detect irregular usage. They were also developing a TCL/TK-based interface called NetMeter that allowed remote control of test machines, and compiled collected data (e.g. one-way delay, IPDV, throughput, packet losses, one-way delay distribution, and IPDV distribution) in a graphical format.

They were passively monitoring QoS parameters using optical splitters, traffic copying (available on some routers and switches), and with normal network adapters. However, there were plans to use Endace DAG cards in the future.

There were also plans to provide support for IPv6 and establish a distributed mesh of capture platforms (initially on RedIRIS). In addition, NetMeter would include GPS synchronisation for owd, a new version of mgen, and support for mobile scenarios.

5.5 IST INTERMON and 6QM projects, Rudolf Roth, FhG Fokus

Rudolph gave a presentation on two IST projects related to traffic monitoring. The INTERMOM project (<http://www.ist-intermon.org/>) comprises twelve partners and is developing a scalable inter-domain QoS control architecture with components for discovering topology, monitoring network traffic, traffic modelling based on measurement data,

simulation of network behaviour, data mining and visualisation. The 6QM project (<http://www.6qm.org/>) comprises five partners and is developing a flexible architecture for QoS measurements in IPv6 networks.

6. Other

6.1 Support for Radio Astronomy, Roberto Sabatino, DANTE

The NREN PC has agreed that GÉANT will provide support to JIVE, the Joint Institute for VLBI in Europe. VLBI is a network of European Radio Telescopes providing data to a Correlator located in Dwingeloo, The Netherlands. Currently, observation data is stored on tapes and sent via trucks to the Correlator, where it is processed off line. The whole process includes three sessions per year each lasting about 25 days and would become much more versatile with real-time network connectivity, even though it can be tolerant to some packet loss and re-ordering and transmission delays

The support would provide a best effort IP service between a maximum of five observatories and the correlator with traffic transiting NRENs and GÉANT. The eVLBI Traffic Profile would require 256Mbps, 512Mbps and 1Gbps real time modes of operation per site. In general it would be a single unidirectional UDP flows towards JIVE.

Given the traffic model and the performance requirement is looks like no significant upgrades to GÉANT are initially necessary, however the TF-NGN audience expressed concerns that the current GÉANT links between UK, NL and DE (2.5Gbps) may not have enough bandwidth capacity for 1Gbps flow and therefore might need to be upgraded. Additionally, it was questionable whether such requirement would be symmetric. The required bandwidth for real-time provisioning is still very theoretical and need to be further investigated.

6.2 GÉANT test bed, Nicolas Simar, DANTE

GÉANT is setting up a test bed for NRENs and DANTE to test OS and protocol features. The traffic will flow over GÉANT but separately from the production backbone. Six PoPs equipped with a router and a workstation connected through a Layer2 VPN using LSPs. Requests from NRENs for using the test bed need to be scheduled. The GÉANT NOC will be responsible for maintaining the schedule and restoring the test bed to the initial configuration. A few operational issues are still open at the moment, including procedures for accessing the test bed traffic measurement, assistance from the NOC and changes to the LSP configuration.

6.3 Update on ipfix, Simon Leinen, SWITCH

Simon briefly reported on the current status of the ipfix working group and warned the audience about the last chance to provide comments to the requirements document, which is in the process of being finalised. At the next IETF meeting in March the working group is expected to make a decision about the candidate protocols for an IP flow information export system.

7. Date of next meetings

The 11th TF-NGN meeting will be held on 8-9 May 2003 in Poznan, Poland, hosted by the Poznan Supercomputing and Network Centre

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

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

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

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

9. Open actions

None?