

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 12th TF-NGN meeting 14-15 September 2003 Jesus College, Cambridge, UK

Issue 1, Kevin Meynell

Attendees

<u>Name</u>	<u>Organisation</u>	<u>Country</u>
Lada altmannova	CESNET	Czech Republic
Wim Barbaix	Alcatel	Belgium
Claudia Battista	GARR	Italy
Artur Binczewski	PSNC	Poland
Mauro Campanella	GARR	Italy
Valentino Cavalli	TERENA	-
Tim Chown	Univ. Southampton	United Kingdom
Peter Clarke	UCL	United Kingdom
Daniel Davids	CERN	Switzerland
Michael Enrico	DANTE	-
Carlos Friacas	FCCN	Portugal
Orla McGann	CNRI-DIT	Ireland
Marian Garcia	DANTE	-
Avgust Jauk	ARNES	Slovenia
Lars Johansen	Wavium AB	-
Dimitrios Kalogeras	GRNET	Greece
Byung-Kyu Kim	ANF/APAN	South Korea
Felix Kugler	SWITCH	Switzerland
Olav Kvittem	Uninett	Norway
Yolanda Lamilla	Cisco	-
Simon Leinen	SWITCH	Switzerland
Roland Leners	Wavium AB	-
Ladislav Lhotka	CESNET	Czech Republic
Marco Marletta	GARR	Italy
Luis Marta	FCCN	Portugal
Kevin Meynell (Secretary)	TERENA	-
Janos Mohacsi	NIIF/UNGARNET	Hungary
Anand Patil	DANTE	-
Victor Reijs	HEAnet	Ireland
Esther Robles	RedIRIS	Spain
Duncan Rogerson	UKERNA	United Kingdom
Roberto Sabatino (Chair)	DANTE	-
Stanislav Sima	CESNET	Czech Republic
Nicolas Simar	DANTE	-
Trond Skjesol	Uninett	Norway
Miguel Angel Sotos	RedIRIS	Spain
Tina Strauf	JOIN/Univ. Münster	Germany
Peter Tomsu	Cisco	-
Szymon Trocha	PSNC	Poland
Bernard Tuy	RENATER	France
Jean-Marc Uzé	Juniper Networks	France
Stig Venaas	Uninett	Norway

Apologies

<u>Name</u>	<u>Organisation</u>	<u>Country</u>
Wilfried Woeber	University of Vienna/ACOnet	Austria

Meeting proceedings

Presentations are available online at:

<http://www.terena.nl/task-forces/tf-ngn/presentations12.html>

1. GÉANT Update

Michael gave a presentation on the current status of GÉANT (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_ME_y4and_update.pdf). The UK-NL and DE-NL links had been upgraded from 2.5 to 10 Gbps, the HU-SI link had been upgraded from 155 to 622 Mbps, whilst the ARNES access port was also now running at 622 Mbps. The Cisco 7000s routers in Luxembourg and Israel had been replaced by a Juniper M20 and M40 respectively.

Upgrades from 2.5 to 10 Gbps were also planned for the IT-ES (22 September), FR-ES (22 September), AT-HU (17 September) and SE-PL (17 September) links, with RedIRIS (September) and NORDUnet (November) also gaining 10 Gbps access.

In addition to the six existing links between GÉANT and Abilene/CA*net4, a 10 Gbps wavelength from StarLight to the NL PoP was being provided by Global Crossing and funded by NSF (for one year). Cisco ONS 15454s were being used to extract the different channels. A 155 Mbps link had also been established between the SE PoP and RBnet, the Russian NREN backbone in July, although NORDUnet already had a connection with RUNnet, the Russian federal university network.

Eighteen NRENs now had IPv6 connectivity, with a further six connected via tunnels. There were currently onward native connections to Abilene, ESnet and CA*net4, and two gateways to 6NET in Amsterdam and Frankfurt. Another five connections were planned.

The overall response to the GÉANT re-tender had been good, with seventeen providers responding. Negotiations were now reaching their conclusion, and final decisions were imminent. Prices in the Baltic states had come down by approximately 50%, and up to STM-4s should be possible. In Slovenia and Croatia, price reductions or upgrades to 2.5 Gbps lambdas/STM-16 were possible, whilst in Portugal dual 2.5 Gbps lambdas were being offered for the same price as the existing STM-1. The current STM-4s in Romania would probably be a bit cheaper, but unfortunately there was little change in either Greece or Cyprus.

Network security had been improved with the introduction of MD5 authenticated eBGP sessions, SSH-only logins, and uRPF checks enabled. Various packet filtering rules had also been implemented at the boundaries of the network.

The connectivity procurement for the EUMEDCONNECT project was nearly complete, and the topology was starting to emerge (probably a star centred on Catania). Most beneficiary NRENs would be receiving a 34 Mbps connection. In the meantime, the connectivity procurement for ALICE (Latin America) was still ongoing, but it was hoped there would be a ring-based regional network with STM-1 or 4 connectivity to Europe.

Finally, the GÉANT Year 4 roadmap was outlined. Some of the Year 3 activities (e.g. performance monitoring, PERT, Premium IP Management, multicasting, and router testing) were still incomplete for various technological and logistical reasons, but these would

continue into Year 4. Proposed new topics were the enhancement of IPv6 monitoring and the extension of IPv6 multicast on the GÉANT backbone, end-to-end and inter-domain provisioning of Layer 2-based connectivity for testing purposes, and the introduction of resource reservation schemes for GRID applications.

Dimitrios asked whether the EUMEDCONNECT and ALICE networks would be part of the GÉANT network. Michael replied they would probably be different domains, but this still had to be decided.

2. GN2

Roberto gave a presentation on the GN2 proposal that would be the successor project to GÉANT (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_RS_gn2.pdf). This would be an Integrated Infrastructure Initiative in the IST 6th Framework Programme, and would be divided into three categories: Networking Activities (NAs) which would provide management, dissemination, outreach and NREN support; Joint Research Activities (JRAs) which would focus on developing new services and testbeds; and Specific Service Activities (SSAs) which would focus on network operations and services. The project would be coordinated by DANTE, and would include TERENA and most of the European NRENS.

Victor remarked that there appeared to be a lot of overlap between the GN2 activities and other IST projects. Roberto replied that the research activities were intended to support and compliment those in other projects.

3. EGEE

Peter gave a presentation on the EGEE project and its relationship with GÉANT/GN2 (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_PK_EGEE.pdf). This aimed to create a European-wide Grid infrastructure, linking existing and new national Grids.

As with GN2, this was an Integrated Infrastructure Initiative in response to the IST 6th Framework call for Grid development. This would include NAs providing user training and support, dissemination, outreach and interface development; JRAs undertaking middleware development and testing, network resource allocation and performance monitoring; and SSAs focusing on core and regional network operations. There would therefore need to be good liaison between EGEE, GÉANT/GN2 and the NRENS to ensure their individual activities complimented each other, and to avoid unnecessary overlaps.

4. PERT

Nicolas provided an update on the performance monitoring initiative (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_NS_performance_monitoring.pdf). This aims to address end-to-end performance problems, develop performance measurement tools for the PERT group, and provide relevant statistics. It will exchange monitored data between domains to help troubleshooting, provide network users with views of edge-to-edge performance, and verify SLAs. Different monitoring architectures will be considered, but individual domains would retain control of their monitoring infrastructure and determine the access policy.

The work has been divided into several working areas including test and metric definition, domain interfaces, path finder, measurement protocols, user representation and statistics, data storage/retrieval/analysis, domain tool architecture, domain tool implementation and

trials. Some of this work is dependent on the IETF IPPM Working Group, but other requires the input of TF-NGN, NRENS and groups of users. The trials will comprise one-way tests using RIPE TTM boxes, with an analyser module also being implemented.

Within the GN2 project, performance monitoring is foreseen as being a Service Activity, and PERT a Joint Research Activity. This will probably start in the third quarter of 2004, and will make around 11 person months available. There is likely to be some collaboration with the Internet2 PiPes initiative, the Intermon project and possibly MBNG.

5. Premium IP Management

Anand gave a presentation on the Premium IP Management being undertaken by DANTE (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_AP_ip_premium.pdf). Premium IP services were being implemented by DiffServ EF PHB, but this service was currently only enabled on the GÉANT network, with admission control and policing being undertaken on the boundary routers. Whilst the operational procedure was already defined (in the context of the SEQUIN project), service management was currently being done manually. This required a lot of effort for set-up and keeping track of parameters, so an application was being developed to automate the process.

This Java-based application utilises Apache, Tomcat and a MySQL database, and features a customisable user interface. This uses XML for communication and allows it to be used with all major browser platforms. Separate access can be granted to the functions and data records, and although this is currently restricted to DANTE, the NOC and APMs, it could be given to other users in future.

The application allows reservations to be managed, reservation levels to be checked on each link, and notifications to be issued. Other functions allow the IS-IS costs to be displayed, router configurations to be checked and updated, reports to be produced, and other system administration tasks to be undertaken. Possible future enhancements may include aggregation between domains, link/interface failure handling, 'what if' analyses, and a multi-domain bandwidth broker capability.

6. Alcatel Router Testing Update

Michael gave an update on the Alcatel A7770 OBX router tests (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_ME_alcatel_testing.pdf). The rationale for investigating Alcatel equipment was the possibility of it being an alternative to Cisco and Juniper, and the fact that it was a European product. This was an important consideration for the IST programme, especially if the NREN community could lend its experience to the developers.

The tests were conducted during week-long sessions in March and July in Antwerp. The goals of the tests were to gain knowledge of the router architecture and functionality, to gain hands-on experience, to perform basic interoperability tests, and provide objective feedback. Four A7770s, one Juniper M40, one Cisco 7507 and Agilent router testers were used for the tests.

Overall impressions were generally good, with both the hardware and software architecture demonstrating resilience. Interoperability appeared to be good, multicasting was stable, and the QoS features were well implemented, albeit more complex than other vendors. On the downside, the CLI was little quirky, and some of the protocols crucial to the current GÉANT network (e.g. IPv6 and IGMP) were still not available.

The next stage is for Alcatel routers to be installed in the GÉANT testbed (probably in the Paris PoP) and for further tests to be conducted over a few months. If these are successful, they will then be moved to the production network, possibly by the end of the year.

7. Internet Deacon

Ladislav outlined his ideas for the Internet Deacon (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_LL_deacon.pdf). This is based on the idea of the Multicast Beacon which is currently used for monitoring IPv4 multicast operation, but aims to improve a number of its limitations. In particular, a more general tool is required that can handle IPv6 unicast and multicast traffic (including SSM), and perhaps even more specific protocols at Layer 4 and above.

The intention is to develop a tool where a server can interrogate and control the probing agents, and with authentication. This requires an agent-manager protocol to be developed which can handle bi-directional message exchanges.

The next stage is to write an overall architecture document, and to define the agent-agent and agent-server protocols. Once this is complete, the actual coding can commence.

More information is available from <http://staff.cesnet.cz/~lhotka/Deacon/>

8. MPLS Layer 3 and Layer 2 VPNs

Dimitrios gave a presentation on MPLS Layer 2 and 3 VPNs (http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_DK_mpls_l2_l3_vpn.pdf). This considered how to establish such VPNs using Juniper and Cisco equipment, and how they interact with routing protocols.

9. IPv6 Update

Tim gave an update on the status of the 6NET project, and IPv6 work in the IETF (http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_TC_ipv6.pdf).

The 6NET project was now at its halfway point, and had deployed a native IPv6-only network. Now that GÉANT had successfully moved to dual-stack operation, the 6NET network could be used for more disruptive tests. The project was now also focusing on applications and services, including management and monitoring tools.

At the IETF, the IPv6 Working Group was heavily debating the proposed deprecation of site-local addresses, and there was still no clear consensus on multihoming within the Multi6 Working Group. The V6Ops Working Group was currently discussing unmanaged network scenarios and 3G protocols, whilst the DNS Working Group was focusing on DNS resolver discovery. However, the DHCPv6 and MIPv6 standards had recently been finalised.

10. Internet2 IPv6 Land Speed Record

Daniel presented the results of the IPv6 Land Speed Record conducted between CERN and Chicago (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_DD_ipv6_lsr.pdf). This requires a minimum of 100 Megabytes to be transferred a minimum terrestrial distance of 100 kilometres, with a minimum of two router hops in each direction between the source and destination nodes. The unit of measurement is bit-metres per second.

The previous Land Speed Record was 5,154 Terabit-metres per second established between ARNES, DANTE and RedIRIS (a distance of 14,800 kilometres) on 9 October 2002. This attempt was under the auspices of the DataTAG project which aims to establish a large-scale intercontinental testbed for data intensive Grids.

Linux PCs (running RedHat 7.3) with dual 2.20 GHz CPUs and Syskonnect Gigabit Ethernet interfaces were used at either end. By using an MTU of 9,000 bytes on the interface, it was possible to obtain 983 Mbps, whilst a TCP window size of 11 MB was also found to improve performance. This resulted in a figure of 6,947 Terabit-metres per second over a 7,067 kilometre distance.

These tests demonstrated that high sustained throughput is possible with TCP over long distances, unlike previous tests where TCP feedback mechanisms would kick-in after 40-60 seconds and degrade performance.

More information is available from <http://cern.ch/ipv6-lsr/>

11. M6Bone

Bernard gave an update on the status of the M6Bone (http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_BT_gtpv6-4.pdf). There were currently thirty connections in several countries, and two separate domains: the M6Bone that used a global RP with RIPng and static routes, and the M6Net that used an organisation RP with MBGPv6. Several applications were now IPv6 multicast enabled including the traditional Mbone tools (vic, rat, sdr and nte), ISABELv6, DVTS and VideoLAN.

Experiments with embedded RP are planned, and there are plans to extent the number of RPs on a regional basis. It was also proposed that the PIM SSM experiments should be restarted.

12. IPv6 Multicast Deployment Experience

Stig gave a presentation on M6Bone deployment issues (http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_SV_m6net_mboned.pdf). This was the first IPv6 multicast initiative and was established as a single domain separated from the IPv6 unicast topology. It uses RIPng for routing and BSR for RP discovery on NetBSD-based routers. PIM Sparse Mode and MBGP are being used for multicasting.

Some problems have been experienced with PIM upstream detection which can cause RPF failures, so it is proposed to implement a new PIM Hello option to announce all interface addresses. Embedded RP is also still missing from the core network, although some testing has been done between Norway and Denmark. SSM has not yet been tested due to lack of MLDv2 host support, whilst bi-directional PIM and MSDP are still unavailable.

13. GÉANT IPv6 Multicast

Marian outlined the plans for introducing IPv6 multicasting on GÉANT (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_MG_geant_v6_multicast.pdf). An IPv6 pilot service had been running since June, and would become a fully operational service in October. Twenty-three NRENs were already connected, and there were onward connections to Abilene, CA*Net4, ESnet, SINET and GX/Telia.

The plan was now to establish an IPv6 Multicasting Taskforce similar to the GÉANTv6 Taskforce, to investigate how multicasting could be implemented on GÉANT. This would define the requirements needed for offering the service, describe the requirements to vendors if further development is required, and investigate the protocols to be used.

The GÉANT testbed would initially be used, which currently comprises three Juniper M20 routers. At a later stage, two Alcatel A7770 routers would be introduced to conduct interoperability tests, whilst the testbed might also be connected to other multicast testbeds such as M6Net.

14. IPv6 Testing within GÉANT

Carlos outlined the IPv6 tests that should be conducted over GÉANT (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_CF_Geantv6.pdf). Whilst much testing was being undertaken in the context of the 6NET project, not all NRENs were connected to the 6NET network. However, just five of the NRENs connected to GÉANT did not yet have either native or tunnelled IPv6, so TF-NGN was considered an appropriate forum for undertaking wider testing.

It was proposed that each NREN maintain a webpage of its IPv6 services (e.g. dns, ftp, http, ntp, smtp and nntp) as these came online, which would itself act as an indicator of IPv6 connectivity. In addition, a loopback address could be configured to act as a test beacon and to measure RTTs between NRENs. Other tests to be defined might include traffic measurement using the RIPE TTM boxes.

15. Optical Networking

Stanislav gave a presentation on optical networking in the Czech Republic (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_SS_Optical.pdf). A national optical fibre network has been established, with two providers offering on-demand fibre laying for the first mile (approximately EUR 3,000 per kilometre). One fibre provider is also offering single fibre leasing for 60% of the cost of leasing fibre pairs. CESNET is currently tendering for a number of new lines, in particular shorter routes between PoPs. They are looking for 19 lines with an overall length of 2,650 kms.

Various tests have been conducted on managed fibre compared to leased dark fibre that demonstrate that a combination of dark fibre with your own equipment becomes more cost effective as speeds increase. However, it was suggested that an exchange of information between NRENs on transmission systems would be useful.

16. Wavium

Lars gave a presentation on optical cross-connect solutions (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn12/20030915_wavium.pdf). This outlined the three different types of Layer 0/1 cross-connect: DXC which uses a combination of time and space switching and can be used with granularities as low as 64 kbps, OEO which utilises electrical space switching, and OOO which utilises optical space switching. The latter two being typically used for 2.5 and 10 Gbps channels, although OOO should scale beyond terabit speeds in future. Having the ability to bypass routers saves on expensive ports, processing in the nodes, and allows control over transmission.

Wavium currently supplies a range of OXC-capable equipment (WX1000, WX4000 and WX8000) that can provide between 16 and 240 ports, at a cost of around EUR 3,000 per STM-16. The WaveMaster software allows easy management of this equipment, and the use

of common standards (XML, SNMP and CORBA) means that it can be integrated with existing network management systems.

Wavium solutions are currently being used in digital cinema distribution networks, and in the Acreo testbed in Sweden.

17. Date of next meeting

The 13th TF-NGN meeting will be held on 22-23 January 2004 in Madrid, Spain (hosted by RedIRIS).

18. Actions from previous meetings

- 11.1 PSNC and GRNET to provide input to the definition of the Performance Monitoring Infrastructure.
 - Ongoing.
- 11.2 Ladislav Lohtka to inform the TF-NGN email distribution list about possible topics for continuation of IPv4 multicast activity.
 - Ongoing.

Open actions