TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 13th TF-NGN meeting
22-23 January 2004
Universidad Complutense, Madrid, Spain

Issue 1, Valentino Cavalli

Attendees

Name	Organisation	Country
Monica Aguado	Sycamore Networks	-
David Boyle	Sycamore Networks	-
Mauro Campanella	INFN-GARR	Italy
Albert Cabellos	Univ. Polytechnic Catalonia	Spain
Valentino Cavalli (Secretary)	TERENA	-
Tim Chown	Univ. Southampton	United Kingdom
Larry Dunn	Cisco	-
Michael Enrico (Chair)	DANTE	-
Remco Frijling	Level3 Communications	-
Alexander Gall	SWITCH	Switzerland
Marian Garcia	DANTE	-
Dimitrios Kalogeras	GRNET	Greece
Felix Kugler	SWITCH	Switzerland
Olav Kvittem	Uninett	Norway
Yolanda Lamilla	Cisco	-
Simon Leinen	SWITCH	Switzerland
Roland Leners	WAVIUM	-
Ladislav Lhotka	CESNET	Czech Republic
Pedro Lorga	FCCN	Portugal
Tamás Máray	NIIF/HUNGARNET	Hungary
Luis Marta	FCCN	Portugal
Javier Orellana	UCL	United kingdom
Jordi Palet Martinez	Consulintel	Spain
Hervé Prigent	CRIHAN/RENATER	France
Michal Przybylski	PSNC	Poland
Jürgen Rauschenbach	DFN	Germany
Esther Robles	RedIRIS	Spain
Duncan Rogerson	UKERNA	United Kingdom
Rudolf Roth	Fraunhofer Fokus	Germany
Nicolas Simar	DANTE	-
Trond Skjesol	Uninett	Norway
Vladimir Smotlacha	CESNET	Czech Republic
Miguel Angel Sotos	RedIRIS	Spain
André Stolze	JOIN/Univ. Münster	Germany
Robert Stoy	DFN	Germany
Bernard Tuy	RENATER	France
Sven Ubik	CESNET	Czech Republic
Stig Venaas	Uninett	Norway
Steve Williams	Univ. Swansea	United Kingdom
Wilfried Wöber	ACOnet	Austria
Fuha Yin	ULB	Belgium
András Zákó	NIIF/HUNGARNET	Hungary

Apologies

Chris Cottrell	Sycamore Networks	-
Carlos Friacas	FCCN	Portugal
Victor Reijs	HEAnet	Ireland
Jean-Marc Uzé	Juniper Networks	France

Meeting proceedings

Presentations are available online at:
http://www.terena.nl/task-forces/tf-ngn/presentations13.html

1. GÉANT Update and Year4 Technical Programme

Marian presented the latest updates in relation to GÉANT (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_MG_GEANT_Updates.pdf). In September 2003 the access links of HUNGARNET, NORDUNet and RedIRIS were upgraded to 10 Gbit/s. IUCC, Israel, was upgraded from 155 Mbit/s to 310 Mbit/s in October. In the same period, trunks were upgraded from STM-16 to STM-64 between Austria and Hungary, Spain and France, Spain and Italy, Poland and Sweden. A number of access and trunk upgrades were still pending.

The GÉANT IPv6 service has been fully operational since October 2003. The NOC was currently monitoring the status of BGPv6, changes in number of routes, looking glass and traffic. GÉANT was currently providing support to the DataGrid project, which had installed some Linux boxes in the DANTE lab generating TCP flows. DANTE were studying the behaviour of such flows and packet reordering. Progress was being made in supporting the eVLBI: the first real time data had been sent between a number of locations. A meeting would be held in the Netherlands in the last week of January 2004.

The EUMEDCONNECT project was entering a new phase. A new GÉANT POP would be established in Catania, in the Italian island of Sicily, via an STM-1 link to Milano. The countries to benefit from the first links, expected in February 2004, would be Morocco, Tunisia (both connecting to the POP in Catania) and Algeria (connecting to the GÉANT POP Madrid).

Michael briefed the audience about the progress of the GÉANT Y4 Roadmap (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_ME_y4roadmap.pdf). The activity included enhancements to IPv4 multicast and IPv6 services, testing of transmission equipment, Layer 2 VPNs, resource reservation and allocation. A deliverable would be submitted in 2004 and DANTE was seeking input and contributions from NRENs.

2. Large MTU Support

Simon presented a proposal to increasing the MTU size in research and education networks and discussed the pros and cons of this idea (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_SL_mtu.pdf). The Maximum Transmission Unit (MTU) defines the size of packets that are transmitted on a certain path. Until 1990 the MTU was 576 bytes for off-net hosts. Thanks to Path MTU Discovery (PMTUD), after 1990 it was possible to transmit larger packets. Currently, the majority of traffic on the Internet makes use of PMTU and has 1500 bytes MTU (Ethernet “standard”). However, most backbones run today at 4470 byte MTU and larger MTU size is supported by most Gigabit Ethernet adapters and switches. At Gigabit speed it looks good to support a larger MTU, because larger packets

mean higher throughput or lower CPU load. However there are some issues on host performance and routers. Today it is possible to use 4470 bytes MTU on Gigabit Ethernet hosts. A few caveats are related to the reliability of Path MTU Discovery, which however is being improved: Simon mentioned the availability of two recent implementations (Linux and NEtBSD) providing better reliability. Some attendees remarked that blocking ICMP is an issue, even with 1500 bytes MTU size and that Ethernet Exchange Points (based on Ethernet) will always have 1500 bytes limitations. Favourable remarks were that, on the other hand, Internet2 were moving to 9000 bytes MTU and support Jumbo frames and in Europe most NRENs and GÉANT were already supporting 4470 MTU. So, Simon suggested that TF-NGN could become active in coordinate MTUs between NRENs, GÉANT and Internet2, run tests of large-MTU servers and experiments with different PMTUD methods.

The feedback to Simon's proposal was generally good. It was felt that it might be useful to support large MTU in Storage Area Networks and in radio astronomy applications, which need Jumbo frames anyway, even though some concern was expressed about a potential need for consuming large buffer space.

3. Performance Monitoring Update

Nicolas updated the audience on the performance monitoring activities (see http://www.terena.nl/tech/task-forces/tf-ngn13/20040122_NS_perf_monitoring.pdf). The driver module for RIPE TTM was almost ready and the first trial would start in April 2004. Collaboration with the GGF Network Monitoring working group was progressing: the group was currently working on the review of a generic request schema for bandwidth data and packet-loss, among other metrics, but more feedback was needed from the network side.

Deployment of additional RIPE TTM boxes in Ireland and Italy was under investigation. CESNET was working on IP available bandwidth metrics. Simon plead for some concrete case - he mentioned eVLBI/Renater as an example - for testing the PERT pilot. There was a need to build a database with description of cases and he asked the audience to submit examples.

Collaboration with the Internet2 piPEs was being established in several areas, starting from studying Path Finder tools issues, up to development of tools like OWAMP and BWCTL. The collaboration would be defined further in a face to face meeting with Grid people at CERN 16-17 March 2004.

The GN1 project, currently in its 4th year phase, was planning to undertake some work being defined in the GN2 proposal in order to match early milestones common to JRA1 and SA3. Nicolas would provide more details on JRA1 later in the meeting.

4. The SCAMPI Monitoring Adapter

Sven presented the monitoring adapter being developed in the SCAMPI project (see http://www.terena.nl/tech/task-forces/tf-ngn13/20040122_SU_scampi.pdf). SCAMPI is a two-and-a-half-year IST project developing a scaleable monitoring platform. It runs until the end of 2004 and has 10 partners including TERENA (coordinator), CESNET and UNINETT (see the project web site: http://www.ist-scampi.org).

Passive monitoring is very good compared to other measurement methods, but not at high speed. One goal of SCAMPI was to overcome this limitation. Three types of adapter were being tested in the project, the SCAMPI adapter, based on the COMBO6 card developed by CESNET in the context of 6NET (see the Liberouter presentation later in the meeting), the 10 GigE DAG card developed by ENDACE and Ethernet cards. The SCAMPI adapter was being

developed by CESNET, Masaryk University and other Czech universities. In the first phase a 4-port GigE interface card a timestamp unit and the firmware have been developed. The 2nd phase will add a new motherboard, a 2-port 10GE interface board (XFP) and new firmware. The SCAMPI adapter can be programmed to be able to analyse any kind of frames, not only IP payload or Jumbo frames.

5. OpenIMP Update

Rudolf Roth briefed the audience about a few initiatives and projects related to performance monitoring Fraunhofer Fokus is involved with. These include the OpenIMP project, which would like to collaborate with TF-NGN on the developments of inter-domain monitoring, and other activities within the IETF working groups IPFIX and PSAMP. Finally, Rudolf mentioned MOME a IST coordination action clustering several existing Commission-funded projects in the area of measurement and monitoring (see http://www.ist-mome.org/).

6. GÉANT IPv6 Update

Marian presented an update on the status of the Multicastv6 initiatives carried out on GÉANT (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_MG_v6multicast.pdf). A multicast IPv6 task force was created by DANTE after the TF-NGN meeting in Rome aiming at defining and implementing a pilot service on GÉANT by the end of 2004. At the moment there were participants from DANTE, FCCN, Renater and Uninett, who hosts the web site (see http://www.uninett.no/geantmc6/). A mailing list named multicastv6@dante.org.uk is hosted by DANTE. The group had produced a wish-list of enhancements to be implemented in Juniper routers, it was not known yet when these will be implemented. Tests of multicast v6 on GÉANT as well as on available NREN's testbeds were still being discussed.

In relation to the operational IPv6 service Bernard remarked that a few NRENs were still announcing old 6bone prefixes (3FFE::/16) at the moment even though since IPv6 is in production on GÉANT it was agreed that test prefixes should not be announced any more. The reason was that some university networks, which are customers of those NRENs still use 3FFE prefixes and it is hard to convince them to change to the RIPE prefixes just because the service has become a production one. NRENs were pushing them but could not force them to adopt the RIPE prefixes. It was suggested that if NRENs needed to support them they could do it by announcing their prefixes to other peers, but not to GÉANT. This issue would need to be further discussed in the GÉANT APM meeting.

7. 6NET Update

Tim gave a brief update on the status of the 6NET project and introduced the agenda of the IPv6 session. 6NET had a project review in October 2003. Among the most noteworthy points were the need to demonstrate things that can be done with IPv6 and not necessarily with IPv4, do more work on applications and updating the cookbooks in WP2.

Stig Venaas updated the attendees on the 6NET IPv6 multicast activity (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_SV_v6multicast.txt). Stig reported that almost all NRENs in 6NET were able to demonstrate native v6 multicast. Technical work in 6NET had recently focused on testing embedded RP and scoped BSR. Embedded RP is a very dynamic and interesting technique at the moment only supported by Cisco, but other vendors have plans to support it as well. 6NET would like to test it at large scale on the M6Bone, possibly also adding some application trials. Testing of SSM was planned to start soon. SSM can be seen as a special case of embedded RP, but it is more secure and solves problems with address allocation.

Stig provided details of two recent Linux implementations of MLD proxies, which can effectively be used to have multicast connectivity through routers that do not support PIM. One, by Jeroen Massar (see http://unfix.org/projects/ecmh/) based on ECMH and the other by Christophe Jelger from the 6NET project.

Finally Stig reported on an open source implementation of Reliable multicast (see http://www.atm.tut.fi/mad/), which uses forward error correction and could be used for file distribution. Stig was planning to test it and invited others to collaborate

Lada updated the audience on the Liberouter developments, initially carried out in work package 3 of 6NET (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_LL_liberouter.pdf). The Liberouter project improves forwarding performance of PC routers for IPv4 and IPv6 (NETBSD and Linux) by means of hardware acceleration. The target throughput is 5-10 Gbit/s. Liberouter is based on an open-system design which gave rise to the COMBO6 motherboard already mentioned by Sven and being used in the SCAMPI project. Lada provided some details of the Netopeer configuration system and mentioned an extension to the 6NET multicast activity in order to improve the control plane of multicast and the forwarding plane of the COMBO6 card. More information on the project is available at http://www.liberouter.org.

8. M6Bone Update

Bernard gave an update on the status of the M6Bone http://www.m6bone.net (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_BT_M6bone_update.pdf). The slides provide the maps of the global and European topology. The European M6Bone connects via Renater to NYSERNET in the USA, to Tunisia and Senegal in Africa and to Korea and Philippines in Asia. Via NYSERNET the M6Bone links to Kansas and Latin America (Mexico, Chile and Brazil). Among applications and services running on the M6Bone Bernard gave a special emphasis to VideoLAN, broadcasting MPEG-2 and MPEG-4 video over IPv6 (servers are available at RedIRIS, Renater and SURFnet) and to Vic-gc.

9. Update on Euro6IX and 6POWER

Jordi Palet Martinez from Consulintel briefed the audience on the updated status of the Euro6IX (see http://www.euro6ix.org/) and 6POWER projects (see http://www.6power.org). Jordi's presentations are available at http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_JP_euro6ix.pdf and http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_JP_6power.pdf).

Euro6IX has set up a network of different Layer3 IPv6 Internet Exchanges in Europe thanks to a strong commitment from telecom operators participating in the project. Jordi mentioned a recent announcement from Telefonica and France Telecom who were starting to offer commercial IPv6 services. Euro6IX supports different access means including wireless, PPP and DSL. 6POWER focuses on deployment of power lines network services based on IPv6. With today's technology services can be provided on a 4 Mbit/s network speed, however silicon is available that can reach 200 Mbit/s. Jordi has authored the ISOC Member Briefing no. 13, which addresses the Digital Divide with IPv6-Enabled Broadband Power Line Communications (see http://www.isoc.org/briefings/013/briefing13.pdf).

10. Future IPv6 Testing in TF-NGN

Tim chaired a short discussion about future work on IPv6 in TF-NGN. This could be done in support of GÉANT directly and in support of NREN services, including particularly transition aids, IPv6 enabling services, IPv6 management and monitoring. Incidentally Tim mentioned that Carlos Friacas from FCCN was investigating which NRENs are providing IPv6 services and this study could provide useful input to the planning of future activities. Promotion of IPv6 deployment was felt as one of the possible goals of future task force activities, in particular on applications and devices for campus and home environment. New IPv6 task force activities should definitely have links to research on multicast and mobility.

11. Overview of Intelligent Optical Networking

David Boyle from Sycamore Networks (optical switches manufacturing) provided an overview of Intelligent Optical Networking, focusing on the rationale for ION, the standards being developed and the implementations available (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_DB_ION.pdf).

Intelligent Optical Networks are replacing traditional SDH networks by adding dynamic switching functionality to underlining optical networks to support IP services and applications. There are several standards being developed by different standardisation bodies to support dynamic switching. The ITU is promoting the G.8080 (G.ASON) standard, which is based on the overlay model, whereas the IETF is actively working on GMPLS, OSPF-TE and RSVP-TE - the building blocks of the peer model - as well as on several ASON related standards. OIF, the Optical Internetworking Forum is working on the definition of UNI and NNI protocols, which are also based on the overlay model.

Details of the Broadleaf I-NNI routing and signalling protocol were given by David as an example of ION implementation. Broadleaf, developed by Sycamore, is based on OSPF and MPLS and is used in real networks by Vodafone UK and by LDCOM Networks (France). Lots of demonstrations and test experiments of UNI and NNI from many vendors have been carried out at NGN '02, OFC '03 and ISOCORE.

12. SWITCHlambda Update

Felix Kugler updated the audience on the SWITCHlambda optical network infrastructure (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_FK_switchlambda.pdf). His presentation focused on the network topology, testing of 10 Gigabit Ethernet and deployment of single fibre bidirectional Gigabit Ethernet.

In Switzerland only one university sill receives connectivity via leased lines while other universities are connected to the fibre backbone. Some smaller customers of SWITCH are connected to the national research and education network via regional networks or VPNs on commercial carrier's networks. SWITCH aims to connect more and more sites directly. By the end of 2004 the optical backbone should become a fibre ring by connecting Manno and Lausanne in the south of Switzerland. Sorrento GIGAMUX is the transmission equipment used by SWITCH. It does not require signal regeneration up to 600 km nor amplification for distance less than 80 km, it supports bidirectional transmission on one single fibre and proves to be a simple effective solution.

Felix reported on several sets of experiments. Tests of 10 Gigabit Ethernet BER were carried out in September 2002, September 2003 and November 2003. Overall, the results were very good without FEC, however there was a need to verify link design and fibre quality especially in relation to PMD, because some of the links, like the one to Manno, were still good but quite old.

Lot of attention among the audience was raised by the details on usage of bidirectional GE on single fibre. This is based on a simple set up with Cisco CWDM GBIC cards plus splitters and

OADM (to remove reflection of local lasers). As of February 2004 single fibre bidirectional GE had been deployed on 11 routes with distances lower than 100 km. The solution is a rather convenient alternative to DWM. It adds as little as 3-4000 euro per GBIC card and leaves one fibre free. SWITCH had further plans to experiment with amplifier spans higher than 100 km.

13. PIONIER Update

Michal presented an update on PIONIER (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_MP_pionier.pdf). While still waiting for some links to be built, PSNC had been working on link optimisation basically by removing the need for DCM modules and replacing G.652 with G.655 fibre. The solution was very much dependent on the quality of the single link and involved a higher fibre cost, but allowed a 35% saving in the cost of the transmission equipment.

New network services on PIONIER included support of ATM customers using Juniper CCC VLAN and Bandwidth on Demand, providing specific channels with a defined QoS between edge nodes based on VLAN and 802.1p. Many new tools had been developed by PSNC, including multicast monitoring and Mcast6, a v4-v6 enabled multimedia streaming application.

14. GN2 Update

The GN2 project is structured in a number of networking activities (NA), specific service activities (SA) and joint research activities (JRA). The refinement process is ongoing and should end by the end of February in order to have a proposal ready for negotiation with the European Commission by mid March. The GN2 project is expected to start in September 2004. Michael announced that the procurement had already started with a tender for connectivity and the documentation been already provided to those who expressed an interest.

The technical activities of GN2 included in SA3 and the five JRAs are the most relevant to TF-NGN. Marian presented an updated on the definition of the description of work of SA3 (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_MG_GN2_SA3.pdf). The activity, which is coordinated by Toby Rodwell from DANTE, has three components: QoS provisioning, PERT and extension of basic services, further divided into ten sub-activities. The effort envisaged is about 30 person-years and the equipment budget is 550K euro. The field of activities is initially applied to IPv4 but can be extended to IPv6 as well. QoS Provisioning includes the definition and development of a provisioning team and the definition of policies for the allocation of premium IP. The PERT has a number of phases including definition, demonstration, ticket system, troubleshooting procedures and day-to-day operation.

JRA1 is led by Nicolas Simar (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_NS_GN2_JRA1.pdf). The activity is about performance monitoring and management and has some relation with the operational service provided in SA3. In particular SA3 will choose the monitoring tools to be used and JRA1 will develop them according to the most appropriate metrics. Some tools have to be ready for SA3 since the very beginning, therefore part of the JRA1 work will be carried out in the framework of the GN1 project, now in its 4th year. JRA1 requires an estimate effort of 600 person-months and 230k euro equipment budget. The areas of work will include data consumer, middle layer and tool developments.

Michael presented an update on JRA2, the activity on security which is led by Christoph Graaf from SWITCH (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_ME_GN2_JRA2.pdf). The goal of the activity is to support the GN2 integrity of network and services by defining policies for access, privacy, roles etc. as well as implementing access to GN2 security services.

JRA3 is led by Michael Enrico (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_ME_GN2_JRA3.pdf). The definition of the work plan for this activity was slightly behind schedule because its content matter was partly related to other FP6 proposal which were still being evaluated by the Commission. The current Description of Work (DoW) focused mostly on "circuit-oriented" Bandwidth on Demand (BOD), starting from MPLS layer2 VPNs and moving to (auto) provisioning of "channels", in a multi-domain, heterogeneous environment. The approach of JAR3 is complementary to the connectionless BOD approach of SA3, which is mostly based on IP. The partners were currently discussing about making in-house software development of provisioning components or not. Much sub-contracting was required for bandwidth broker/resource negotiation etc. Because of the confidentiality related to the FP6 proposals under evaluation, only a small group of experts from DANTE, CESNET, PSNC, GARR and GRNET were currently working on the DoW.

Marian gave a brief update on JRA4, the activity on technology and service testing, which is led by Marting Mogensen from DANTE (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_MG_GN2_%20JRA4.pdf). JRA4 aims at building a distributed testbed to support GN2 activities, other FP6 projects and test of new technologies. The estimated effort at the moment is 119 person-months with a budget of 2.6M euro for equipment and connectivity.

JRA5 is led by Jürgen Rauschenbach from DFN, who presented the status of the DoW development (see http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn13/20040122_JR_GN2_JRA5.pdf). The activity focuses on mobility and access of network services by roaming users. It will build on the results of TERENA task forces TF-AACE and TF-Mobility, which will hold a joint meeting on 6 June 2004 in Rhodes to define the terms of collaboration. The main (sub) activities of JRA5 are the design and implementation of an NREN roaming infrastructure and an NREN Authentication & Authorisation infrastructure, both enabling the creation of a Single sign-on, and finally integration of new technologies in the mobility area, including MIPv6.

The audience raised a general question about the relation between NAs and JRAs and what the future of TF-NGN will be. TF-NGN, like other TERENA task forces, will go under NA6. There is no plan to discontinue TF-NGN, which will act as an open forum for wider discussion of developments taking place in JRAs. JRAs need to be very specific for the first 18 months but based on the ideas discussed in NA6 can change or also lead to the creation of new ones.

A second general question was related to the lack of emphasis on either Multicast or IPv6 in GN2. These are implicitly dealt with by GN2 but considered as Business as Usual.

15. Date of next meeting

The 14th TF-NGN meeting will be held on 10-11 May 2004 in Amsterdam, The Netherlands, hosted by SURFnet and the University of Amsterdam.