TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Summary Report on the 15th TF-NGN meeting
29 September - 1 October 2004
FCCN, Lisbon, Portugal

Issue 1, Valentino Cavalli

1. Introduction

1.1. Attendees

Name	Organisation	Country
Tiago Camilo	DEI UC	Portugal
Mauro Campanella	GARR	Italy
Valentino Cavalli (Secr.)	TERENA	-
Tim Chown	University of Southampton	United Kingdom
Christian Cinetto	GARR	Italy
Wim Derijnck	BELNET	Belgium
Jerome Durand	RENATER	France
Anders Elgemyr	Wavium AB	-
Michael Enrico (Chair)	DANTE	-
Carlos Friacas	FCCN	Portugal
Andras Jako	NIIF/HUNGARNET	Hungary
Avgust Jauk	ARNES	Slovenia
Olivier Jerphagnon	Calient Networks	-
Loukik Kudarimoti	DANTE	-
Felix Kugler	SWITCH	Switzerland
Olav Kvittem	UNINETT	Norway
Yolanda Lamilla	Cisco Systems	-
Roland Leners	Optovia	-
Pedro Lorga	FCCN	Portugal
Luís Marta	FCCN	Portugal
Chris McGugan	Cisco Systems	-
János Mohácsi	NIIF/HUNGARNET	Hungary
Antonio Pinizzotto	IIT-CNR	Italy
Emmanuel Quémener	ENS de Cachan	France
Jürgen Rauschenbach	DFN-Verein	Germany
Victor Reijs	HEAnet	Ireland
Esther Robles	red.es / RedIRIS	Spain
Rudolf Roth	Fraunhofer FOKUS	Germany
Afrodite Sevasti	GRNET	Greece
Nicolas Simar	DANTE	-
Trond Skjesol	UNINETT	Norway
Miguel Angel Sotos	red.es / RedIRIS	Spain
Bernard TUY	RENATER	France
Ana Tavares Pinto	FCCN	Portugal
Szymon Trocha	PSNC	Portugal
Sven Ubik	CESNET	Czech Republic
Jean-Marc Uzé	Juniper Networks	-
Stig Venaas	UNINETT	Norway
Tom Walsh	Calient Networks	-
Chris Welti	SWITCH	Switzerland

Steve Williams	UKERNA	United Kingdom
Wilfried Wöber	UniVie-ACOnet	Austria

1.2. Apologies

Name	Organisation	Country
Tiziana Ferrari	INFN	Italy
Dennis Paus	SURFnet	The Netherlands

1.3. Online presentations

All presentations from the meeting are available online at the following URL: http://www.terena.nl/tech/task-forces/tf-ngn/presentations15.html.

1.4. Welcome

Several meeting attendees had never been to a TF-NGN meeting before; therefore Michael Enrico provided a brief overview of the task force goals and activities. He stressed that TF-NGN is an open forum to discuss and promote new networking services and testing new lower- layer technologies. The task force is established by the TERENA Technical Committee and traditionally chaired by DANTE staff. It meets 3-4 times per year and its current mandate runs until November 2004.Discussion about its extension would be a major item for the current meeting's agenda.

2. GÉANT and GN2 update, Michael Enrico

2.1. GÉANT

In June 2004 a number of NREN access circuits were upgraded. The primary DFN access port was upgraded to 10Gbit/s POS (STM-64). GRNET primary access port was upgraded to 2.5Gbit/s POS (from half STM-16 to full STM-16). SWITCH backup access was upgraded to 10Gbit Ethernet.

In the context of the ALICE project, providing intercontinental connectivity to Latin America, a new STM-4c link was connecting Sao Paulo to the GÉANT POP in Madrid from 1 September 2004. Not all circuit planned in the project were in place yet, but Chile and Brazil were already sending traffic on that link. Cisco Systems had donated GSR routers to NRENs involved in the project, but these were not in place yet, so they were running their own equipment at the moment.

MPLS Class of Service has been implemented on GÉANT in July 2004 to support specific traffic needs. The first candidate would be the DEISA project, connecting supercomputing sites in France, Italy and Germany and requiring 1Gbit/s EF traffic flows to be implemented over GÉANT.

Michael recalled the discussion in Madrid about large MTU size to support jumbo frames in the research and education community (see the presentation "Increasing MTU in Research Networks" from Simon Leinen at http://www.terena.nl/tech/task-forces/tf-ngn/presentations13.html). In August 2004 the MTU was increased on most GÉANT core links (except in STM-1 links). The MTU was also increased in some access links to the US and Canada. NRENs (SWITCH and BELNET at the moment) were starting to benefit from that.

Wilfried Wöber asked it there was any measurement about the amount of traffic corresponding to Jumbo frames. Michael replied that DANTE had not looked into this yet.

The last update on GÉANT regarded the configuration of Martini Layer-2 connectivity, limited to 1Gbit/s between Poznan and Geneva, to support the ATLAS project.

2.2. GN2

After reminding the audience about the structure of the GN2 project, with its division of (human) Networking Activities (NA) Service Activities (SA) and Joint Research Activities (JRA), Michael informed the attendees that the project had officially started on 1 September 2004.
Most activities were just kicking-off, whereas others had already started a few months earlier, like the procurement of connectivity – started at the beginning of the year and already completed, leading very likely to several fibre-based links – and the equipment procurement. The latter was still ongoing and limited to transmission and switching equipment only, given that the Juniper Networks routers currently used will still be in place in the new network. The equipment that would be required to support the testbed in JRA4 was not included in this procurement.

The new network, which will be called GÉANT2, would support both IP and lightpaths:

- with respect to the current network, the IP capacity is expect to increase up to 10Gbit/s in as many of the smallest trunks as possible, whereas it may exceed 10Gbit/s in a few trunks by 2008;
- the capacity required to support lightpaths on the hottest trunks may reach up to 5x10Gbit/s.

The current expectation was that contracts for the new network would be signed at the beginning of 2005.

Separate, extensive presentations of SA3 and JRAs, were given to TF-NGN attendees in Amsterdam and are available at http://www.terena.nl/tech/task-forces/tf-ngn/presentations14.html (see, Michael Enrico on SA3, JRA3 and JRA4, Jürgen Rauschenbach on JRA5, Jacques Schuurman on JRA2 and Nicolas Simar on JRA1). In Lisbon Michael, Jürgen and Nicolas briefed the audience about the latest updates.

SA3, dealing with end-to-end QoS and led by Toby Rodwell (DANTE), had a kick-off meeting on 9 September in Cambridge. Work on high level specifications for the interface between EGEE/GN2 was almost done and was expected to be ready by beginning of October 2004. A PERT pilot was expected to be available at the end of October based on work previously carried out in GN1-Y4 and TF-NGN. A provisioning System architecture was under consideration.

JRA1 deals with network performance, is led by Nicolas Simar and held its kick-off meeting immediately before the TF-NGN meeting at FCCN. Scope of the activity is to provide monitoring information to be used by the PERT (see SA3 above), various projects (Grids, VLBI etc.) researchers and end-users and make them able to retrieve data from a specific domain. The activity will also improve existing monitoring tools and build specific visualisation tools. More details about the activity including a status report are included in Nicolas presentation at http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn15/20040930_ns_perfmon.pdf.

JRA2 deals with security, is led by Christoph Graf (SWITCH) and had a kick-off meeting in Malta one week before the TF-NGN meeting.

JRA3 focuses on connection-oriented dynamic provisioning (Bandwidth on Demand) in multiple domains. The activity led by Michael Enrico is split in four work items (WI) led by GRNET (requirements), GARR (specifications), PSNC (implementation) and HEANet (test and validation). The work would start in October even though the kick off meeting would not be held before the 1st week of November.

JRA4, led by Marian Garcia (DANTE) will provide an enhanced testbed facility to carry out disruptive experiments. It is expected to be used mainly by other JRAs, but also, in some related

case by GN2 partners in the course of their research activities and by other IST projects. This activity would hold a kick-off meeting back to back with JRA3 in November.

JRA5, led by Jürgen Rauschenbach, held a kick-off meeting in Berlin early in September, back to back with the TERENA task force TF-Mobility. Jürgen listed the activity tasks and mentioned, among others, building a web portal to increase the visibility of the EduRoam pilot service (see http://www.terena.nl/tech/task-forces/tf-mobility/) and working on the Authentication and Authorisation Infrastructure. The AA-RR (Authentication Authorisation Requester-Responder, see details at http://www.terena.nl/tech/task-forces/tf-emc2/aarr.html) system by RedIRIS could be a core model for the Infrastructure. Design of the AAI is due in June 2005 when the testbed implementation would also start.

3. Advances in Optical Amplifier Technologies, Roland Leners

Roland provided an overview about the evolution of optical transmission, and in particular of the advances in optical amplifiers technologies in the past decade.

A number of trends can be observed. Performance increased over the last 10 years, in terms of reach, capacity and OADM support. At the same time, costs have come down dramatically. For instance, since 1996 the cost an EDFA amplifier decreased from more than 5000eur/lambda to less than 1000eur/lambda, whereas a Raman amplifier went down to the same amount in only 5 years.

What is going to happen next? Roland examined two scenarios, based on All Optical Networks and on Hut Skipped Networks. The latter approach is also known in the NREN community as NIL, Nothing In Line, approach.

All Optical Networks – extend the reach of amplifiers. Require very low noise amplifiers, moderate gain, moderate dynamic range, Raman amplifiers only. This technology is particularly suitable (and cost-effective) only for networks requiring very high capacity (>100 Gbit/s per route) and very long circuits (>2000km).

Hut Skipped Networks – no inline amplifiers (get rid of the maximum possible # of amplifiers). Require low noise amplifiers, high gain, large dynamic range, **Hybrid EDFA/Raman amplifiers.** This technology is particularly suitable (and cost-effective) for networks requiring moderate capacity (<100 Gbit/s per route) and shorter circuits (<1000km).

Hybrid EDFA/Raman amplifiers solutions entail about 20 percent reduction of capital expenses on the same distance in comparison with Raman only solutions. They also provide a significant reduction of operating expenses but these vary on the basis of different conditions and are difficult to estimate in general.

The cost reduction in the Optovia solution is obtained by integration of traditional Raman pump components into an integrated block, and by standardisation.

There is progress in the area of interworking transmission components. WDM technology so far was single vendor. It is now possible to plug transmitters/receivers from one vendor into a different vendor WDM system. However, a separate (common) umbrella element management system is needed to make these components interwork. Optovia has chosen to support this approach and is aiming at Open WDM Systems.

Victor Reijs asked if the transmission capacity concerned in the solutions under examination was 10Gbit/s or 40Gbit/s. The answer was that the focus now is on 10Gbit/s, which on multiple 1000 km is quite challenging.

Mauro Campanella stressed the need for a common management system. SNMP appeared to be widely supported and would be the obvious solution.

4. Lightpath definition activity, Victor Reijs

Victor Reijs prefers to use the expression “point-to-point connection” instead of the term “lightpath” because the latter is not well defined yet.

The term “lightpath” has been used in many different definitions, theories and experiences by CANARIE, GN2, Internet 2 HOPI, SURFnet, WINMAN, ASTON and the IETF Pseudowire working group.

In his presentation Victor examined the way a lightpath is meant by each of these organisations. For CANARIE a lightpath is basically a Layer 1 connection providing guaranteed bandwidth in a deterministic fashion. The GN2 definition is not so strict about the deterministic aspect, so, besides Layer 1, the concept can include premium IP, MPLS and other Layer 2 and Layer 3 technologies. The Internet2 HOPI group defines it starting from a different aspect: the main criterion being that applications can be sure that there is no loss in the network. The SURFnet concept is similar to the CANARIE one, whereas WINMAN is really looking at SDH, simply meaning a high bandwidth pipe. The ASTON group of TF-NGN used the term “channel” instead of “lightpath”, meaning basically a link between two points with some defined SLS. A definition by the IETF Pseudowire working group is still being developed.

Victor pointed out an interesting paradigm shift in the HOPI approach: applications in the past had to cater for the “limitations” of the network. They were built and prototyped in LANs and often needed a lot of extra tuning to support users on WANs. With the HOPI definition, applications should not care about network limitations, like congestion, loss, excessive delay and delay variations anymore.

Victor examined the common features of the “lightpath” concept and remarked that a few ones are probably missing from the definitions given above. TF-NGN should evaluate features and understand which ones should be included in its definition.

A comment from the audience (Wilfried) was that the investigation should not be limited to studying the characteristics of the network, but one should start understanding what type of applications/boxes/labs users plan to connect to these point-to-point connections.

5. Performance Monitoring update, Nicolas Simar

Latest developments of the work on performance monitoring (Perfmonit) include implementation of the analyser module (in charge of multi domain communication) and the Pathfinder module. The TF-NGN group has proposed an implementation of GGF NM-WG literal based web services instead of the SOAP RP.

The Perfmonit group has established collaboration with Internet2 piPEs on interoperability but also wants to develop a joint document describing the respective monitoring architectures, defining future joint developments and investigating how to monitor “lightpaths” (this is a tricky

task because it is not yet clear which specific metrics should be monitored). Other collaboration activities include the EGEE project, which is mostly interested in measuring the IP available bandwidth.

Various NRENs have started early work on GN2 JRA1 already in GN1y4. They drafted a requirement questionnaire specifically targeting three different user groups: NRENs, projects and end-users. So far 15 answers were received, plus a number of potential respondents were identified who wanted to provide feedback but did not have time to fill out the questionnaire. DANTE would follow up with them individually. Another task was to identify interesting, state-of-the-art tools and evaluate them. Evaluation guidelines were developed by PSNC regarding specific tools and by DANTE regarding the general monitoring infrastructure.

6. Load map and the Scampi project, Olav Kvittem

Olav Kvittem gave two presentations, one about Network Load Map Systems and one about the STAGER application developed as part of the SCAMPI project.

Three Network Load Map Systems have been developed by UNINETT: NEMO (auto zoom geographical maps), ZINO (schematic load maps) and NAV (auto detect campus topology).

Olav explained what a Network Map system is: a set of tools generating clickable www-maps starting from a network topology description. NEMO is a Java program accessing a SQL database of maps and traffic statistics, which draws maps and animates network characteristics. NEMO can be made available to TF-NGN, but has not been officially released yet.

Olav described in detail the architecture of the NEMO tool and its features and provided a demonstration of the visualisation tool functionality.

STAGER stands for STatistical Accumulating Graphing Extensible Reporting and is a flow reporting tool developed in the IST SCAMPI project (see http://www.ist-scampi.org). The application is running on UNINETT Cisco routers at the moment.

7. Monitoring QoS classes on JANET, Steve Williams

Steve's presentation showed the administrative complexity of configuring QoS in a complex environment. He reported about experiments where an infrastructure running monitoring probes based on Cisco SAAs was put in place. A separate machine was needed to coordinate all this. In the experiment, data was gathered from many sources in the network and graphically visualised so as to show effectively if anything was wrong (e.g. misconfiguration etc.).

The tests were initially running between Swansea and Manchester with various applications including voice and video. More complex tests were carried out at the end, also including the University of Lancaster, the Imperial College, the University of Southampton and UKERNA. More details on the Janet QoS trial are available at http://www.ja.net/development/qos. Related information is available in a guide to reliable campus H.323 networks http://www.video.ja.net/323campus.pdf.

Steve observed that expanding network monitoring to different domains is very valuable and remarked the complexity involved in such a work. This would be particularly the case in JRAs where they aim at operational monitoring tools.

Nicolas added that JRA people are looking forward to obtain feedback from operational people and end-users; he also invited the TF-NGN audience to feel free to express any concern that might arise regarding the suitability of the GN2 activity plans.

8. TF-NGN Charter discussion

TF-NGN will officially terminate its mandate in November 2004. The discussion about new Terms of Reference for the task force had started at the meeting in Amsterdam in May, was followed by a BoF during the TERENA Networking Conference in Rhodes, and continued via email during the summer. This discussion provided the input to a draft of new Terms of Reference for TF-NGN that Michael Enrico and Valentino Cavalli had distributed to the TF-NGN mailing list a few weeks in advance of the meeting.

The list of proposed topics/work items in the draft of new Terms of Reference is:

9.1 Improvement of current multicast service
9.2 IPv6
9.3 IP routing (including issues surrounding the scalability of current IGPs and EGPs)
9.4 MPLS and GMPLS testing
9.5 Transport Protocols (TCP limitations and tuning, fast TCP, new transport protocols)
9.6 Optical Networking (including a "Dark Fibre Forum" in which experiments on fundamental optical transmission and all-optical switching are reported along with news on collaborations with existing testbeds)
9.7 Hands-on evaluation of new router and switching hardware (including performance evaluation of homegrown networking hardware and, where appropriate, in collaboration with established equipment vendors willing to enter into EFT testing programmes)
9.8 New link protocols
9.9 L1/L2 services
9.10 Access technologies (first-mile technologies)
9.11 VPN services (layers 2 and 3)

Most topics in the list are new with respect to the current Terms of Reference: 9.3 IP routing, 9.4 GMPLS, 9.5 Transport protocols, 9.6 Dark fibre forum, 9.7 Hands on HW testing, 9.8 New link protocols, 9.9 L1/L2 services, 9.10 Access technologies, 9.11 VPNs.

The following topics, which are listed in the current Terms of Reference, are now largely covered by GN2 JRAs. These can, in principle, still be discussed in TF-NGN but priority will be given to the topics listed in Article 9 mentioned above.

- Tools for network monitoring and flow measurement (GN2-JRA1)
- Definition, Testing, Implementation and Monitoring of Differentiated IP Services (GN2-SA3)
- AAA development (GN2-JRA5)
- Performance Enhancement Response Team (PERT) development (GN2-SA3)

The procedure leading to the renewal of TF-NGN is the following:

- Agree on the topics/work items and identify people (leaders) who take responsibility for them
- Each leader put more flesh around each bullet above
- Make a final call to the list
- Agree on the final draft Terms of Reference by end October/beginning of November

- Submit the draft ToR to the TERENA Technical Committee (TTC)
- The TTC makes decision in their meeting on 20 December 2004
- TF-NGN meets under new mandate on 13-14 January 2004

There were two sessions about the Terms of Reference in the meeting. The topics were discussed in the second session. In the first one there was a general discussion about the interaction of JRAs/TF and about what people want from TF-NGN.

Traditionally the activity of TF-NGN was tightly linked to testing technologies and services that would later be introduced on GÉANT. Things are different now. These developments and tests will be carried out in the next four years as part of JRAs. This has two main effects for TF-NGN: the task force should have an independent agenda and be a more open forum than it used to be in the past.

Ideally outsiders could be invited to attend meetings of JRAs or collaborate in their activities; however these are closed groups. JRAs need an open forum like TF-NGN to obtain feedback on their developments from experts from peer organisations, universities or research labs in Europe or in other continents who are not GN2 partners, GN2 partners who do not have effort to dedicate to a specific JRA, equipment vendors and generally industrial players.

The majority of people active in TF-NGN nowadays come from the GÉANT constituency. In order to match the goal above, by ensuring that TF-NGN has an own research agenda, while at the same time providing a sounding board for JRAs work, it is essential to pull in people from the outside.

Getting people from institutions other then NRENs is mostly a dissemination issue: TF-NGN would benefit from some "publicity' targeted to the wider European research environment. A few participants were concerned that people from such a wider community might not have funding to attend meetings and participate in task force activity, however, it was observed that most of the time relevant research is funded by some project and this fact might allow people to participate in TF-NGN.

Valentino said if the task of involving more research institutions in TF-NGN is to be taken seriously then the work item leaders should be more active. They should identify projects/people to talk to, test-beds to report about and or to be supported. Work item leaders likely know who is working on these matters in other communities and should try to get them on board, invite them to subscribe to the list, attend meetings, give presentations but also actually collaborate in the activities. Similar things have happened in the past, for instance with various Internet2 groups or with equipment vendors providing equipment on loan for specific tests.

It became clear from the discussion that there are two views about the usefulness of TF-NGN. One opinion, expressed by Olav Kvittem, is that TF-NGN helps participants in getting to know good quality work and to meet people to exchange/compare ideas: there is no need to produce something for each activity. The other, expressed by Bernard Tuy, is that real tests need to be carried out in the task force to be able to raise outsiders' interest in participating in its activities.

There was a general consensus about the need to add a general goal to the charter: the task force should aim at the introduction of the technologies tested in TF-NGN in a production environment; this should be pursued by producing tangible results like How-To documentation.

ACTION 040929-1: Michael Enrico to draft a text in the new Terms of Reference about producing How-To documentation.

Not all work items where discussed during the meeting. Leaders for the following items only were identified:

9.1 Improvement of current multicast services – Steve Williams had volunteered to lead this work item.
9.2 IPv6 – As it emerged from the session on IPv6 on Friday morning (see section 9 below) there is still a lot of potential work to be done in TF-NGN. Tim Chown is prepared to lead it, but asked for a co-chair; Stig Venaas agreed to co-chair this work item.
9.3 IP routing – Mauro Campanella explained the rationale for this activity: most routing protocols work fine on today's networks, but these are changing in many directions: scale, optical fibre, circuit switching at high speed, more meshing, grids etc. All these changes present new challenges and new tasks to be addressed. Mauro volunteered to lead this work item.
9.4 MPLS and GMPLS testing – Jean-Marc Uzé pointed out the potential interest represented by new upcoming MPLS services. János Mohácsi said that HUNGARNET were interested in using MPLS for Layer3 VPNs and also in investigating VPLS technology – János agreed to lead the MPLS activity. Whether GMPLS should be included in this work item or not remains an open issue. Some vendor's representatives at the meeting said GMPLS/ASON is getting a real hype in Asia and there should be more focus on it in Europe.
9.5 Transport protocols – Szymon Trocha said somebody from PSNC would lead this work item. He will let the group know who will be the actual leader. [This happened a few days later, the leader of this activity will be Radoslaw Krzywania]
9.6 Optical Networking – Victor Reijs agreed to lead this work item but said it is important to double check it with CESNET, in particular regarding their plans for a CEF forum.

Points 9.7 to 9.11 were not discussed. It was agreed to send the list of work items to the email list asking the group to provide feedback.

ACTION 040929-2: Valentino Cavalli to make sure that new TF-NGN work items are finalised via the task force email list.

The audience was asked if there were other interesting areas not covered in the proposed Terms of Reference: Jean-Marc said Terabit routing/logical routers are an interesting topic, he would send some details to the mailing list. Additionally Jean-Marc was seeing a trend in the routing industry for much more dynamic configuration, customer- empowerment: he was interested in getting feedback from the TF-NGN community: it was agreed to have this as an agenda item in the next meeting.

ACTION 040929-3: Jean-Marc Uzé, Michael Enrico to organise a discussion on dynamic configuration, customer- empowerment at next TF-NGN meeting.

9. IPv6 (text based on Tim Chown's notes)

The goals of the IPv6 session where the following:

1. To move back from "talking" to "doing", by defining tests that can be carried out within TF-NGN. Bear in mind that for many partners 6NET is finishing Dec'04, even though it has been extended to Jun'05. Also, the 6NET multicast testbed may be withdrawn so the need to have the desired functionality on GÉANT is growing.
2. To determine what NRENs can do to encourage and support IPv6 deployment into the campuses. IPv6 is generally well supported in the backbones, but not so into many

regional and campus networks. The monitoring and end-to-end performance tools may still be lacking, for example: should NREN tunnel brokers be offered, or 6to4, or Teredo, or all three?

GN2 contains work items on Multicast IPv6 and IPv6 Measurement and Monitoring, but IPv6 is largely taken as standard, thus it doesn't warrant special attention in the way it was done in GN1.

As a result, the specific threads for the session were:

a) Discuss the Multicast work to date, requirements and future tests.
b) Identify areas for NRENs to support IPv6: transition tools, monitoring tools, end-to-end performance tools, understanding if/where new protocols need special consideration, e.g. MIPv6.
c) Identify areas where NRENs can help with campus deployment, by both undertaking gap analysis and also promoting interesting applications, such as new peer-to-peer systems.

The aim of the session was to produce an identification of 3-4 key work areas, and an initial idea of useful tests to run to support these areas.

9.1. GÉANT, Michael Enrico

GÉANT plans a pilot Multicast IPv6 service by end of Dec'04.

Tests are running now with JUNOS on the GÉANT testbed (3 x Juniper M20, old Ciscos).

A Multicast IPv6 task force was created with Renater, UNINETT, FCCN and SWITCH.

The testbed is connected to m6bone via a tunnel.

There is a beacon client on the testbed, with beacon agents on m6bone.

GÉANT is currently using JUNOS6.1 on the production network, which includes PIM forwarding, but with RP provided by another party. NRENs are sending /32 - /35 multicast and unicast routes, and all go into one routing table, causing some RPF check problems.

Within a few weeks the production network should use JUNOS6.4, which will fix the multicast routing table problem, and will support the RP, with internal monitoring via a beacon. The RP will not be offered to users, so GÉANT remains transit only.

Beyond that, embedded RP needs to be implemented; this is expected early in 2005, but will need some testing. Someone will need to offer an RP for the community in the meantime.

There will also be a need to handle non multicast IPv6 enabled NRENs.

Michael gave an overview of IPv6 in GN2. It is "implicitly assumed" in JRA1 that all services will support IPv6, e.g. for monitoring QoS, IPv4 and IPv6, and multicast.

In JRA2 (security), the CERTs will need to include IPv6 thinking. This is at a site level as well, and is probably closer to TF-CSIRT. End to end security is an issue here, as is security of transition mechanisms.

JRA3 is on lower layer channels, i.e. light paths. These should support all Layer 3 transmission techniques. There may be some issue with IPv6 in the control plane? Probably more at the routing layer than the switching side.

JRA4 is the new testbed, including new switching and optical facilities. The testbed should be available for IPv6. We need as many vendor equipment types here as possible.

Ubiquity (JRA5) includes a variety of technologies, including MIPv6, but a little bit later into the activity.

SA3 is on end-to-end QoS. Mauro said Premium IP is not defined for multicast yet (IPv4 or IPv6). Multicast and QoS are needed for some applications, e.g. AccessGrid. Premium IP is available for IPv4, so it should be straight forward to define and use a Premium service for IPv6, even though there is not yet any plan for Premium IPv6 on GÉANT. There are some issues like the possibility to do better with the Flow Label and the way to combine Premium IPv4 and IPv6.

9.2. IPv6 SSM, Stig Venaas

Most routers support SSM as they are.

You need MLDv2 support at the edges (as per IGMPv3 for IPv4). Cisco and Juniper support this. MLDv2 can let listeners specify sources to listen to or to not listen to.

ICMPv6 protocol numbers for MLDv2 are now finalised (and in the latest Linux kernels) MLDv2 is needed in OSs beyond Linux, FreeBSD+KAME and Solaris 10.

RFC3678 specifies the Multicast API. Stig has sample Linux code available. SSM applications are still few in number: mcastsend/mcastread, ttcp, dvts, dtms/dtmc, xdvshow, mad-flute.

Strasbourg has produced an ssmsdpifier, which turns a traditional multicast application into an SSM one.

SSM is difficult with dynamic sources. SSMSDP is a bit like an application layer RP. The controller sits on a group, and the (S,G)s are sent from the controller when announced. This of course recreates the RP issues.

ASM is useful for some areas, e.g. service discovery.

9.3. m6bone, Jerome Durand

There are quite a few issues at the moment, including:

- New PIM Hello option with respect to link locals.
- Varying option numbers for ICMPv6 types for MLDv2 (need to upgrade all Linux workstations), and with hosts not falling back to MLDv1.
- Embedded RP is now at RFC status, RP to be used is chosen by the client (it is embedded in the multicast address). Do we want the client to choose? How does the application learn the multicast address (end users cannot manually create the address)? Some kind of multicast address assignment is needed.
- There is a new BSR draft - do we need this if we have embedded RP? The new spec should be as backwardly compatible as possible, ideally. Many routers will not have scoped BSR, so what will the old BSR routers use?

- It is not possible to redistribute static multicast routes in MBGP on Cisco, though OK on JUNOS, and is a problem when non-congruent topology is used, so it is difficult to connect new test sites.
- Juniper embedded RP and new BSR is still needed, as is IPv6-in-IPv6 tunnelling (Jean-Marc will check this). Not clear yet whether new BSR will be needed due to embedded RP.
- The Perl beacon is not yet working properly with IPv6, while the Java beacon seems to have a memory leak. There is an interoperability (crash!) problem with different beacon client versions, but this may now be fixed already.

The m6bone network is still growing, including Taiwan.

RedIRIS is streaming operas with IPv6 Multicast at DVD quality (MPEG2). RedIRIS will announce more details soon (Nov/Dec'04).

János reported that SSM has IPR attached, through Apple? Are vendors refraining from implementing due to this IPR issue?

Jean-Marc asked if anyone was interested in Automatic Multicast Tunnelling (AMT) The operator can avoid multiple unicast streams, and break a chicken and egg problem. There is FreeBSD code available for this; Jean-Marc will circulate this. There is an IETF draft on AMT, Juniper is considering implementing it.

9.4. IPv6 peer-to-peer, János Mohácsi

Most peer-to-peer applications hard code IPv4 addresses. BitTorrent and GNUnet do not, and thus may be good candidates for IPv6 porting. A few issues involved are:

- What is the IPv6 advantage? Is it avoiding NAT traversal?
- What about new peer-to-peer architectures one the assumption that IPv6 is there?
- Can peer-to-peer IPv6 systems make NREN network usage more efficient?

BitTorrent is written in Python. Scalability may be an issue (because of the tracker behaviour). The multitracker may assist with this.

There is an IPv6 patch written for BitTorrent 3.3. that enables use of wildcard IPv4/IPv6 addresses. Version 3.4 has changed the tracker protocol though, with IP+port being a 6 byte binary! BitTornado is a shadow client that has also implemented the v3.4 changes.

ABC BitTorrent client (Windows) that supports IPv6 on Linux, partially (no IPv6 in Windows Python binaries yet, but v2.4 of Python should do). One needs to consider IPv4-IPv6 peer interactions. The developer claims DNS is unreliable, hence wishes to use addresses only. IPv6 is seen as an unnecessary complexity, maybe it will be added in late 2005.

GNUNET is promising, but hard to compile and use. It can transfer over channels including SMTP!

Other peer-to-peer systems with IPv6 are Multipeer and ThreeDegrees. We should talk to the Internet2 Lionshare people.

9.5. Tools to assist deployment

If native IPv6 isn't available, most NRENs will offer tunnels to sites wishing to deploy it. We know how to connect sites. The bigger problem is finding cheap CPEs to allow the "masses" to connect, e.g. xDSL. What can TF-NGN do to help there? A few issues/ideas are that:

- Global IP addresses and ease of configuration should be an advantage.
- Universities already have the address space, student hostels may use NAT.
- There must be boxes being made in Asia with IPv6? We should investigate.

9.6. Summary of actions/activities required for TF-NGN IPv6

1) How do users in non multicast capable IPv6 NRENs get multicast IPv6 access? And who are these NRENs?

- DANTE can offer one tunnel per NREN (when embedded RP is ready)
- Phase 1 - Run tunnels to Renater's Juniper M5 while GÉANT gets ready
- Phase 2 - Move tunnels to GÉANT when ready, with Renater being the community RP, though with embedded RP we don't care where it is.

2) Plan tests for embedded RP. Should a core router run an RP given it may be a DoS target? Probably run the RP at the edge?

3) Agree community RP and plan for GÉANT phased deployment.

- 6NET RP is in Renater, so use Renater?

4) Define SSM tests, and have SSM architecture discussion - is SSMSDP good or bad? Will we need IPv6 ASM? Track Apple's IPR claims on SSM.

5) Need tests show how to monitor multicast IPv6 (required in JRA1).

6) Think about IPv6 security issues for JRA2. Include security of transition mechanisms, and recommendations from these.

7) Define how the new GN2 testbed (JRA4) can be made available for TF-NGN tests.

8) Feed IPv6 issues on mobility and ubiquity into JRA5. Summarise what is done in available testbeds already.

9) Multicast and IPv6 with QoS, e.g. for IPv6 AccessGrid.

10) Premium IPv6 should be straight forward to define and use. Does IPv6 Flow Label help at all here? How do you combine Premium IPv4 and IPv6?

11) Take part in RedIRIS multicast IPv6 opera streaming tests.

12) Recommendations for IPv6 Multicast deployment, including scoped BSR, embedded RP, AMT.

13) Promote IPv6 to universities to get real usage of IPv6. Identify IPv6 applications to encourage university adoption, e.g. new peer-to-peer applications (messaging like Jabber or file sharing like BitTorrent) especially where IPv6 architecture is leveraged. Investigate Lionshare?

14) Determine which transition aids should be deployed in support of IPv6 deployment by campuses.

15) Identify gaps in standards and tools to monitor and manage IPv6 networks (currently happening in WP6 of 6NET).

16) Produce short guides on how to configure basic services for campuses, including SMTP enabling.

ACTION 040929-4: Tim to share long-list of possible new work items via the tf-ngn mailing list and encourage the group to nail down to some feasible IPv6 activity.

10. Future of NRENs: Migration from Pure IP to Hybrid IP+Optical Networks, Chris McGugan (Cisco Systems) and Olivier Jerphagnon (Calient Networks)

Chris and Olivier gave a joint presentation on the basic technologies, protocols and architectures required for evolving NRENs from pure IP to hybrid IP+Optical networks.

Drivers for such an evolution are identified in dynamic and high-bandwidth demanding applications and services, performance requirements, scalability issues and economic factors. To match emerging needs NRENs should be able to provide production and experimental services in a flexible way from layer 3 to layer 1 and be able to allocate large capacity on demand.

The presentation went through the details of the available technologies and the network architectures enabling the provision of network services on a hybrid infrastructure and explained in detail the status of GMPLS/ASON developments towards the provision of a Unified network Control Plane.

Chris and Olivier provided examples of NRENs who have been deploying hybrid IP+Optical networks, including SuperSINET (Japan), NLR (USA),TWAREN (Taiwan), Optiputer (USA, The Netherlands) and Japan Gigabit Network JGN II.

Chris and Olivier identified the following areas of joint research interests with NRENs: Bandwidth on Demand, Policy and Bandwidth Broker, AAA Integration, protection services, multicast services, extending capabilities of optical transmission, multi-domain provisioning and hybrid router architectures. In addition, middleware developments are needed to support applications and make full use of resources at layer 1, 2 and 3. They said NRENs have a critical role in driving such developments.

Cisco Systems and Calient Networks are working with other institutions on a new optical router architecture that can scale up to 100Tbit/s called LASOR, label switched optical router. The project is funded by DARPA and started six months ago. Key technologies are new routing protocols, architecture, all-optical and fast wavelength converters, etc.

Mauro asked if it possible for the TF-NGN community to collaborate in the project. It was also remarked that similar developments in these areas have started in Europe in GN2 JRA3 as well as in the MUPPETT and the VIOLA projects. Olivier said Calient was interested in getting new contacts and encouraged people to talk to him offline.

11. GN2 common components and architectures, Mauro Campanella, Afrodite Sevasti

Mauro explained his view on the network architecture needed to support the evolution of NREN requirements. These requirements include support for services needing high network Availability, with resiliency at various layers, plus the capability to allocate Bandwidth on Demand and to provide accounting. Additional complexity is due to scaling in many directions at the same time. However, NRENs want to preserve key principles of the Internet, like end-to-end, simplicity, transparency etc.

The idea proposed by Mauro consists of a layered architecture based on a robust information system. Mauro provided an overview of the proposed network logical architecture and the underlying Information System, which would store functional descriptions and semantic content of network objects.

Mauro provided more details on the proposed architecture and proposed an example based on GN2. He said all new services in GN2 require the same technical components to provide different functionality: measurement in JRA1, provisioning of Premium IP (PIP) in SA3, BoD provisioning in JRA3. There is a need for a common harmonised system.

Afrodite explained that the reason for holding this session in the meeting was that with the start of the GN2 project many related issues had started to be discussed on various specific mailing lists and it was important to have a common feedback from TF-NGN.

Afrodite built on the example given by Mauro and started her presentation by observating that JRA1, SA3 and JRA3 need a common information system and a unified communication model for multi-domain provisioning. These may additional be useful also in the ubiquity (JRA5) and security (JRA2) activities. Such information system and communication model could be shared by various services, all of which entail path discovery, AA requirements, brokerage, policies etc. Afrodite provided some specific examples based on the operation of the BoD provisioning process and the various elements involved.

The conclusion was a proposal to adopt a common architecture for all relevant services, adopt an interoperable information system schema, identify dependencies between different GN2 activities and set up an informal 'architecture team' monitoring the specifications proposed by individual JRAs.

A positive feedback was expressed by Victor about the benefit of coordinating different GN2 services. Jean-Marc said the goal of providing a common information system did not seem to be too ambitious. He though the Internet needs something like this and vendors would look at it with lot of interest and would like to work along with the community in this specific project.

People in the audience remarked that very similar problems have been experienced already in the respective NRENs. Afrodite and Mauro added that GRNET and GARR had indeed already developed a database addressing the issues discussed and this database could be replicated (if not even reused) at the GN2 level.

Some detailed discussion on the network logical architecture proposed by Mauro followed the two presentations. Mauro asked the participants if there was any missing functionality. People from the audience expressed some concerns about the need for a Main Broker as well as about the suitability of a hierarchically layered structure. Wilfried suggested that the information repository should not be centralised and follow a DNS-style model instead. More generally, it was remarked that the proposed architecture should not just provide an ad-hoc solution to GN2 issues, but be discussed in a broader forum, possibly including vendors.

12. Next meeting

The 16th TF-NGN meeting would be held on 13-14 January 2005 in Brussels, Belgium

13. Summary of OPEN ACTIONS

ACTION 040929-01	Michael Enrico to draft a text in the new Terms of Reference about producing How-To documentation.	NEW
ACTION 040929-02	Valentino Cavalli to make sure that new TF-NGN work items are finalised via the task force email list.	NEW
ACTION 040929-03	Jean-Marc Uzé, Michael Enrico to organise a discussion on dynamic configuration, customer-empowerment at next TF-NGN meeting.	NEW
ACTION 040929-4	Tim to share long-list of possible new work items via the tf-ngn mailing list and encourage the group to nail down to some feasible IPv6 activity.	DONE