

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 6th TF-NGN meeting
10th and 11th of January 2002
CERN, Geneva, Switzerland.

Valentino Cavalli, Kevin Meynell,
Issue 2, 25 February 2002

ATTENDEES

Name	Organisation	Country
-----	-----	-----
Michael Allenby	UKERNA	United Kingdom
Wim Barbaix	Alcatel	Belgium
Alain Bidaud	Crihan/RENATER	France
Artur Binczewski	PSNC POL-34/155	Poland
Mauro Campanella	GARR-INFN	Italy
Valentino Cavalli (Secr)	TERENA	-
Tim Chown	Univ. of Southampton	United Kingdom
Axel Clauberg	Cisco Systems	Germany
Daniel Davids	CERN	Switzerland
Tiziana Ferrari	INFN-CNAF Bologna	Italy
Alexander Gall	SWITCH	Switzerland
Avgust Jauk	ARNES	Slovenia
Joop Joosten	CERN	Switzerland
Dimitrios Kalogeras	GRNET	Greece
Tom Kosnar	CESNET	Czech Republic
Sabine Kühn	DANTE	-
Olav Kvittem	Uninett	Norway
Simon Leinen	SWITCH	Switzerland
Ladislav Lhotka	CESNET	Czech Republic
Athanassios Liakopoulos	GRNET	Greece
Riccardo de Luca	Telscom	Switzerland
Guido Luyts	Alcatel	Belgium
Olivier Martin	CERN	Switzerland
Kevin Meynell	TERENA	-
Marcin Michalak	ULB-STC	Belgium
Jari Miettinen	CSC/FUNET	Finland
János Mohácsi	DANTE	-
Paolo Moroni	CERN	Switzerland
Christian Müller-Böhm	JOIN/Univ. of Münster	Germany
Ruy de Oliveira	University of Berne	Switzerland
Arne Oslebo	UNINETT	Norway
Michal Przybylski	PSNC	Poland
Sylvain Ravot	CERN	Switzerland
Jürgen Rauschenbach	DFN	Germany
Rudolf Roth	Fraunhofer FOKUS	Germany
Luis Rodrigues	ITU	-
Roberto Sabatino (Chair)	DANTE	-
Pavel Satrapa	CESNET	Czech Republic
Yves Schaaf	Restena	Luxembourg
Afrodite Sevasti	GRNET	Greece
Nicolas Simar	DANTE	-

Wim Sjouw	Univ. of Amsterdam	The Netherlands
Trond Skjesol	Uninett	Norway
Vladimir Smotlacha	CESNET	Czech Republic
Robert Stoy	DFN	Germany
Bernard Tuy	RENATER	France
Jean-Marc Uzé	Juniper Networks	France
Stig Venaas	Uninett	Norway
Steven Williams	UKERNA/UW Swansea	United Kingdom
Wilfried Woeber	Univie AConet	Austria

APOLOGIES

Octavio Medina	EST Bretagne	France
Miguel Angel Sotos	RedIRIS	Spain
Victor Reijs	HEAnet	Ireland
Franz Widhofner	University Linz/ACOnet	Austria

Online presentations

<http://www.terena.nl/task-forces/tf-ngn/presentations6.html>

1. Introduction and GÉANT update

Roberto Sabatino welcomed the participants, reviewed the list of open actions and presented the latest GÉANT updates. He said that the network is up and running since the 1st of December but there are some open issues regarding Athens and Lisbon for which connectivity had not been implemented yet and for the time being required some reuse of the TEN-155 circuits; the GÉANT routers in Poland were not in place yet. Roberto said that DANTE was in the middle of a tender process for connecting Israel, Cyprus and Malta at a capacity that realistically would range from 34 to 155 Mbps. Circuits had been delivered for connectivity to North America Research Networks. Connectivity to Germany was under test, in the United Kingdom it was put in place but they had not started testing yet. He then mentioned the initiative to build a worldwide regional network called GTREN, a proof of concept in the current perspective, involving some commitment from UCAID to buy capacity for the transatlantic connection that would be further extended to Seattle and Tokyo.

Roberto talked about provisions for commodity IP traffic. He said that the US PoP was still in 25, Broadway. DANTE had been running a tender that was won by KPNQ and Global Crossing. Roberto also showed the interconnection points and capacities they are implementing for a planned RFS of mid-March. He added that global IP connectivity was available in Europe in each of the 10 Gbps PoPs, and that the arrangement with INFONET would be phased out. He then showed the indicative IP subscription rates requested by the NRENs. All traffic will be marked and rate will be limited at NREN access rate. The configuration of internal setup for multicast was finished. There were plans about making extensive usage of multicast monitoring tools but the goal was also to monitor traffic volumes. Support for Premium IP was planned. A test case H.323 was started in collaboration with the SEQUIN project. The GÉANT routers are configured with policy scheduling, it was now the matter to install counters to count packets that have DSCP different from 0 (Best Effort). Currently 95 % of packets coming from NRENs are BE. Roberto mentioned that a EUMEDCONNECT phase 1 study of the possibilities and options in the Mediterranean area was being worked out and concluded his presentation by reminding the audience that according to the updated GÉANT roadmap two new deliverables were due: a capacity management study deliverable (D9.8, due fall 2002) and a study on Less than Best

Effort (D9.9). More deliverables to be produced soon were an appendix to D9.5, about MBS replacement, D9.6 on IPv6 and D9.7 about traffic measurement.

2. Premium IP update, H.323 test experience

Mauro Campanella reported briefly on test activities going on between Italy, Germany, Greece, Switzerland and GÉANT on EF and H.323. He said that some basic testing had already taken place in the last days of December 2001 and the group was planning to finalise the configuration of equipment within three weeks of the Geneva meeting. However they still needed tools to monitor packet-loss, delay sensitivity and other QoS parameters. Debugging videoconference itself is not a goal, the actual video conference will only start after they have tested the QoS parameters. They plan to do some very limited test with network congestion because the experiment is using production lines. Finally Mauro announced a forthcoming SEQUIN workshop that would be held on the 1st of February in Amsterdam targeted to users who might be interested in using the Premium IP service.

3. Less than Best Effort

At the TF-NGN meeting in Athens a lot of interest was expressed in investigating into LBE (Less than Best Effort), based on the developments of the Internet2 Scavenger project. Tim Chown chaired the session to initially coordinate the investigation and mentioned the document produced by Ben Teitelbaum about Internet2 stepping away from Premium IP and moving to network over-provisioning to solve the QoS problem. Tim remarked the need to understand LBE rationale and identify scenarios to tag packets as LBE: some applications that need it are GRID/E-science bulk data transfer, student hall of residence and others. An implementation method could be to use a common DSCP with Internet2 but the whole idea would need more investigation. Tim suggested some ways to go forward including identifying trials and deployment areas, again look at Internet2, find the infrastructure to conduct the trial. The ideal testbed would be GÉANT and/or some NREN but one has to identify which particular ones. It would also be possible to create a new TF-NGN activity group. He recalled that a deliverable is due as mentioned by Roberto in the roadmap update, but it is still open about the content, and can also be a report about the activities of Internet2. Tim concluded with the impression that LBE would be very useful to be tested on international link.

Simon Leinen said that he is member of the Internet2 relevant lists and recalled that a document came out recently about deployment recommendations to make use of the service, he had the impression that it has been used already for control of bandwidth using QoS appliances.

Tiziana mentioned the IP+ service on GÉANT and asked if the LBE study would be focused on implementing it. Roberto clarified that LBE is a new item whereas IP+ has been set aside for now, based on the analysis made by Simon Leinen within SEQUIN.

Mauro expressed some concerns about LBE, particularly in relation to implementing congestion detection. Tiziana replied that it can be done with Waited Fair Queue (WFQ) and Waited Round Robin (WRR) basically because LBE has a very simple mechanisms.

Jean Marc supported Tiziana in saying that in the Juniper boxes LBE (Scavenger) is today already implemented and under test on Transpac link between Tokio and Chicago. (see "JUNOS QBSS Configuration Notes" at <http://www.shepfarm.com/juniper/re/JUNOS-QBSS.html>). He added that there are 4 queues per port on the router, but 2 are sufficient to support Best Effort + LBE services, the one takes 1% of

the traffic marked as LBE traffic and the other uses the remaining 99% for BE, if there is no congestion BE can take up to 100% of the available bandwidth.

Lada asked how to stimulate users to use the service and a related question about how to charge for the Premium IP service. Mauro replied that this is up to the GÉANT policy committee to decide and it does not just depend on the amount of bandwidth used but on the quality of the service.

To make a decision about how to follow up on the LBE activity Tim asked which NREN would be available to run such service in principle and who would be able to find the resources to contribute to do some testing. There were only a few hands raised, many NRENs seemed not to have time and resources to work on that, but Tiziana said that interest in the GRID community is high because their tests congest the network and this is a problem on production networks. She also remarked that it is very simple to implement it, much more than Premium IP. Tim said that UKERNA might be able to do some tests, but he needs to go back to them. Tiziana agreed to work up an initial work plan for test and implementation of the LBE service.

ACTION 6.1 Tiziana to work up an initial work plan for test and implementation of the LBE service

4. Multicasting

Janos gave a presentation on the GÉANT multicast service.

Dimitrios asked whether the multicast service would use different routing announcements than the normal unicast service. Janos replied they planned to use the same routing rules for multicast routes but different routing table for easier debugging and convenience.

Dimitrios also asked whether announcements would be allowed from different ranges. Janos explained in his presentation that SSM range is filtered thus source specific multicast is not allowed to use MSDP.

Artur gave a presentation on the C multicast beacon that was being developed by Roman Lapacz. The intention was to deploy this in all GÉANT PoPs, although some features still needed to be implemented. In particular, a configurable packet size, message frequency, and receiver-only option.

5. IPv6

Tim said the group needed to discuss how to progress the GÉANT IPv6 Test Programme alongside 6NET. It also needed to identify the status of IPv6 deployment at NREN and university sites, identify any non-Cisco equipment that could be trialled, and identify testbeds outside of 6NET.

Tim then updated the group on the status of GTPv6. Deliverable 9.3 had been approved and was now online. The addendum needed some more work, but the first draft should be available by the end of January. This should also declare future deliverables.

Janos followed with a presentation on the proposed GÉANT topology and 6NET overlay. He said this would be discussed further at the 6NET Work Package 1 meeting later that day (11 January).

Bernard gave an update on the status of the Juniper M5 router in Paris. Six NRENs were currently connected to this and the link to ETRI in Korea was also operational. The intention was that all GÉANT routes would be announced by this router.

Tim raised the issue of the Telebit router. This has been in place since the QUANTUM Test Programme and Ericsson were not participating in the GÉANT Test Programme. However, as they had not asked for the router to be returned, it was agreed it should be used in the GTPv6 testbed for now.

The addressing plan was discussed, although there were still connectivity issues with Belgium and Finland. It was agreed that a 2001:0660 prefix should be used for peering, whilst 3FFE should be avoided if possible. It was also agreed that participants could announce their own routes, but that 2001:0660 should be aggregated as /35 and 3FFE as /28.

Routing protocols were also discussed after Joop provided an overview of BGP4+. It was agreed that both iBGP4+ and eBGP4+ should be tested, along with IS-IS and tunneling. Access controls were also considered, but these were not currently supported by Juniper.

Bernard gave a presentation on the M6Bone which was being used to run videoconferences. There were still some platform-specific issues with PIM Sparse Mode and some of the tools (RATv6 in particular), but a successful all-day seminar had recently been held between a number of sites.

A work programme was agreed for TF-NGN which attempted to avoid too much overlap with 6NET activities.

The main priorities were interoperability testing, application testing, evaluation of some transitional tools, and possibly some consideration of addressing and registry issues. Lower priority work was to investigate multicasting and possibly to undertake some firewall and IPSec testing. It was decided that DNS, multihoming, monitoring and wireless issues should not be considered as these were being undertaken in the 6NET project.

Tim said there had been some interest in creating a separate mobility group within TF-NGN to complement the mobility activities in the 6NET and 6WINIT projects. He had discussed this with Valentino, but a decision yet to be made.

6. Over-provisioned Networks

Wim Sjouw summarized the experiment he already reported in Athens. He had presented results on loading an empty 2.5 Gbps network and measure if traffic can get through without packet loss and jitter. The idea behind the experiment was that if you have a large pipe of 2.5 to 10 Gbps and more, before doing any QoS manipulation you need to know if QoS techniques are really necessary. He found out that the bottleneck is in the routers at the entrance and at the end in that they affect the performance of the pipe. Fast routers are needed at the ends. At the moment the most sophisticated router cross-connects have the same construction as the less fast ones: every input port is connected to an output port with 30 microseconds throughput i.e. 50% of the cross-connect and this is the main cause of delay in a router. If on the same input channel packets arrive for different output channels they are split into different buffers and this accelerates the process. A connection is over provisioned as long as there is no buffering for one output channel, i.e. no packet is queued behind each other. Over provisioning requires minimum delay in the routers. If there is any buffering there is no over provisioning and then you have to

prioritise. Michal Przybylski said that at PSNC they have measured routing time and that normal routers have 50 microseconds routing time. Steve Williams remarked that when you have aggregate you have to queue in some way and asked which minimum delay would still allow you to say that the network is over provisioned. He suggested that this kind of effect might be introduced by the GigE interface on the PC. Wim confirmed that without the two routers in the path the pattern was different. He experimented with Cisco 12016 routers with engine2 cards. The PCs were not directly connected to the router, but to the Cabletron 6800 switch. Mauro noticed that the effect could be explained by the switch queue: the reason should have been a buffer in the metrics of the switch. Wim, agreed, but restated his main point, to identify a feasible way of measuring over provisioning. His method was to send simple series of ping and look at the RTT and he wanted to try it with longer delay on a lambda to Chicago, as well as to replace the Cabletron switch with a faster one. He did not consider Jitter nor one-way-delay, because he does not have a clock, but he was still willing to look at other parameters. There was a discussion about some experiences in this respect and Steve said he had started looking at buffer depth. Mauro suggested considering no (or almost) packet-loss as another metric and pointed out the need to choose a percentage limit. A threshold about network load like the common sense 70% figure still being used today from the telecommunication industry might be a good indication for network over provisioning.

A more general discussion about the whole activity started. Mauro said that one of the goals was to measure the GÉANT network when it was empty, but it could not be made for time limitations. Wim had collected baseline results for tests on the unloaded network, and would like to try the same experiments now that the network is live. Tiziana, supported him in saying that the main interesting thing was to measure the live network and characterize its behaviour. Dimitrios remarked also that in the past months it did not make sense to experiment on old equipment, but now it was the right time to start them. As many other attendees Olav thought that ping was not enough and should be combined with some other parameter, including more passive-measurement stuff. About measurement tools, it was suggested that one must have some probes occasionally, but not active measurement continuously running in order to avoid congesting the production network. Generally a need was felt to define more clearly the parameters to be tested and the related kind of measurement. Mauro suggested that packet-loss and queue depth should be an indication of over provisioning. Afrodite Savasti urged to make a decision about how much frequently the measurement is going to be done.

Wim raised the issue of coordinating the continuation of the activity and asked who was willing to collaborate. He also mentioned a leadership issue since it appeared that Dimitrios Kalogeras had taken over from Trifon Chiotis. Nobody committed to do the work but there was consensus about continuation and targeting at measure one-way delay across GÉANT. Dimitrios promised to coordinate the development of an activity plan.

ACTION 6.2 Dimitrios to coordinate the development of an activity plan to continue and extend the OvPN performance test.

Vladimir Smotlacha said he has tools to measure the accuracy of NTP which have been tested on a MAN in Prague but need time synchronisation in order to be used on a wider network like GÉANT, however they do not need GPS antennas. His proposal was to install NTP on Linux workstations, get the signal from GPS antennas and compare it with the NTP clock to check the quality of time-synch. If the resulting time tolerance is around one microsecond then the tool can effectively be used for the measurement. The GPS antennas can be installed at the NREN node, they do not need to be on GÉANT. Tiziana expressed some concern because she ran similar tests in TF-TANT and NTP proved to be unuseful, but Vladimir remarked that there have been lots of improvements in the latest version (4) of NTP and promised to prepare a proposal about Accuracy of the Network Time Synchronization.

ACTION 6.3 Vladimir to prepare a proposal about Accuracy of the Network Time Synchronization
- DONE

A discussion about the availability of RIPE NCC measurement boxes followed. Tiziana and Dimitrios have installed them and found that they are very appropriate for the over provisioning tests. Dimitrios agreed to poll the TF-NGN email distribution list to find out how many RIPE boxes are installed in NREN POPs and then define how they can be used in the OvPN performance measurement.

ACTION 6.4 Dimitrios to poll the TF-NGN email distribution list to find out how many RIPE boxes are installed in NREN POPs and then define how they can be used in the OvPN performance measurement.

7. Flowma

Simon gave an update on the IETF IPFIX Working Group (<http://ipfix.doit.wisc.edu/>) which was working on NetFlow standardisation.

The Working Group had been established in September 2001 following a successful BoF at IETF 51 in London. The first meeting was held at IETF 52 in Salt Lake City in December, and primarily discussed the UDP problems with NetFlow.

Wilfried asked why reliability was an issue when NetFlow was only sampling data anyway. Simon replied it was important to have a reliable transport mechanism and it was suggested that TCP should always be used.

Tom then gave an update on the measurement system he was developing that could also handle IPv6.

Olav asked when the code would be available. Tom replied that he had not yet finished work on it, but it should be available by July.

8. CCC on GEANT

Roberto Sabatino addressed the issue of Managed Bandwidth Service (MBS) replacement and proposed a solution based on MPLS/CCC. The service would be complementary to the Premium IP service, but so far it had been requested only to support the IPv6 test bed. It should address both an end-to-end and edge-to-edge solution, both must be compatible but the edge-to-edge solution can be very different. The service should provide granularity to support from 2 to 622 Mbps. The proposal from DANTE consisted in using MPLS/CCC, an edge-to-edge solution based on ATM PVCs or frame relay or GigE V-LANs. DANTE would support it because it is ready and tested, and they know how it works. The requirement for NRENS would be to have separate access to this service. Roberto proposed a scenario where ATM switches are placed next to the GÉANT router in the GÉANT POP connected to an ATM switch at the NREN PoP. On the NREN site it would be done either by connecting a native ATM PVC or, if they do not support ATM, at the same way as GÉANT, by using CCC. There was some concern about the proposed solution, which would be proprietary, and some discussion with Dimitrios about the ATM switch. Roberto was open to discuss whether ATM or GigE is the best option and then adopt that one.

IPv6 in TF-NGN and 6net are candidates for the service alpha testing. However it was not clear how many NRENs would be interested in doing that. There was a suggestion that Tim could possibly be interested, but since he was attending another session somebody had to contact him off-line. Steve had some concern about deploying such a complex architecture if it is was only meant for IPv6, in relation to the fact that IPv6 does not need MBS but just needs transport. Roberto replied that it is also needed to support Atrium. Olivier Martin said that CERN might have another request for MBS but it could also be point-to-point to Japan. Dimitrios said he would like to test if the proposed solution would work with ATOM on the one side and CCC on the other side.

ACTION 6.5 Dimitrios, Roberto, to work out scenarios for CCC on GÉANT

Roberto asked if there was interest in carrying out some more generic work on MPLS testing, there was some discussion ending with some expressions of interest in working on L2VPN.

9. Delay+jitter

Tiziana Ferrari has extensively tested QoS on Cisco Catalyst 6506 with GigE interface. The tests basically concerned policing and scheduling, considering that congesting a GigE locally is not realistic. Tiziana expressed a wish to have other Catalysts available and carry out some more tests on the WAN. The input queuing had one strict priority queue and four single-threshold WRR queues, the output queuing consisted of one strict priority queue and two double-threshold WRR queues. The IOS image was version 12.1(8b)EX2. Tiziana said that policing works according to token bucket and depends on the following parameters:

- CIR Committed Information Rate, the minim sustainable rate
- Norm_burst_size token bucket size related to the CIR rate
- PIR peak rate
- Maximum_burst_size token bucket size to LE traffic rate peak to PIR
- Conform Action

Policing can be of 2 types: aggregate policer and per-flow policer

Interfaces can be in one of the following states:

- Untrusted
- Trust IP precedence
- Trust DSCP
- Trust CoS (layer 2)

The purpose of the test was to verify the correctness of the policing mechanism. Tiziana ran different cases: test A with variable normal_burst_size and fixed CIR at 50 Mbps; test B, with different CIR; test C, with CIR fixed at 50 Mbps and variable PIR; and test D, with fixed PIR at 71 Mbps and variable Maximum_burst_size. Details about the tests are available at <http://www.cnaf.infn.it/~ferrari/tfngn/cat6500/police/>

Tiziana said there was an open case with Cisco about WRR so she wanted to test it on GSR routers. They have GSR in France and Poland, but PlaGe is not available at the moment, and those in Poland are operational so cannot be used. Tiziana would be ready to continue the tests but it seems that there seemed to be no platform available to perform them. She remarked that Engine 2 would be sufficient for the test and asked the audience if there was some option. It turned out that it would be possible to obtain a window for Tiziana on a GSR at Alcatel in Antwerp to carry out the test remotely, however on one router only she would be able to perform only some basic tests, to complete the test schedule she would need at least three routers. Dimitrios said that he has some boxes available in Greece and would follow up with Tiziana.

Besides the Catalyst tests Tiziana reported some notes on the C12000 queuing architecture for MDRR testing. She described the architecture of the 12000 GSR line card, and discussed advantages and disadvantages. There was some discussion about the formula to calculate the MDRR queue quantum, Tiziana promised to check the documentation and found out the correct one. As a conclusion of the session she had an additional presentation about ECN testing, a simple derogation from WRED that need to be tuned before being tested. Olav had expressed an interested in participating in the test but he was planning to do that after he got an IOS image with IPv6 support and Uninett do not have it available, he was willing to investigate if he could use a Linux box.

10. ATRIUM

Wim Barbaix updated the audience about the current status of the ATRIUM project. Three 7770 routers have been deployed and one more will be connected via OC-3 at BELNET that will later be connected to GÉANT either directly or via RENATER. The project is carrying out some interoperability tests also with Juniper and Cisco routers. The equipment in Antwerp is available to the group for testing. E-LSP, Diffserv QoS tests are ongoing and WFQ is available. Alcatel is also working on Traffic Engineering tools with facilities to display the network routers, interface and capacity as well as primary and backup LSPs. Extensions of the tests are planned to Poland and Italy. Plans also include tests with 10 Gbps on interconnections in June 2002 and ATM-IP mediation tests using MPLS-TE OPTIP (the Alcatel optical IP programme). So far the ongoing tests were completely separated but they will probably be brought into Atrium. Eventually there was a plan to upgrade the 7770 software by the end of 2002 to demonstrate IPv6 functionality.

Wim was willing to organize a workshop in Antwerp about the 7770 with an extra focus on the IP Diffserv features. It would be a two-days event with hands-on initially limited to 12 people. In April the full Diffserv QoS functionality (strict priority queuing) would be available. Dimitrios said he might be interested in that but asked to look into some documentation, including internal architecture. Tiziana found it useful to have training only if the lab could be used after that. Wim said that the Atrium lab is available to the TF-NGN group, and the tests can be done both locally and remotely (but only once they are connected to BELNET)

Guido Luyts presented OPTIP, the Alcatel optical strategy. The current core network is service specific and multi-layer but in the future it will become a production service. IP-MPLS is going directly into the optical layer and GMPLS is used to connect between routers and optical cross light. The benefits consist in a standard way of handling connection, not centrally based anymore and fast and end-to-end provision of bandwidth (although not multi-domain). There were plans of using the UNI interface that would enable overlay and hybrid operation. What Alcatel has implemented so far consists in the ODSI interoperability tests, which is also part of the UNI interoperability test. They are implementing Core Node 1, Network Dynamic through Management, and hope to release the product by end of 1st quarter 2002. This would be the same platform as the Alcatel ALMA management tool but it is not integrated with it.

11. Lambda testbed proposal

Valentino Cavalli reported about the offer from FLAG Telecom to provide an international Lambda testbed for TF-NGN. FLAG Telecom is a TERENA associated member since October 2001 and has transcontinental optical links from Europe to Asia, the Middle East and the US. Their European network serves London, Paris, Amsterdam, Frankfurt, Duesseldorf and Brussels, with extensions to Geneva and Milan planned in 2002. A brainstorming meeting about the proposal was held in London on 6 December 2001 between representatives from FLAG Telecom, Victor Reijs from

HEAnet and John Dyer from TERENA. Valentino mentioned that Victor raised a very interactive discussion about the goals of TF-NGN Optical Networking activity and presented some ideas about how to get experience on the European scale. There was consensus about going ahead with the initiative but there was a need to discuss the technical options in the community and TERENA offered to organize a meeting in Amsterdam on the 12th of February 2002 inviting experts to discuss the proposal and define in more details the goals of the tests and other issues like, for instance, how to handle POP locations and equipment, define principles to handle costs to connect and so on. Valentino announced that TERENA had established the mailing list lambda-testbed@terena.nl, invited people to subscribe to it and to attend the meeting.

There was some interest but also some expression of concern about the economical viability, particularly in relation to the availability of test equipment and the cost of local loop to connect NREN PoPs to the FLAG network. The group felt that many technical details were missing and suggested to start a discussion online before the meeting in order to focus the goals and motivate people to participate actively in the initiative.

12. PIONIER

The previous sessions ran longer than planned therefore Artur Binczewski had only a few minutes time for his presentation about the Polish optical network and limited himself to announce that Pionier will be operational this year. He showed details about the planned topology being deployed in three stages now and to be completed at the end of 2002.

13. date of next meeting

The 7th TF-NGN meeting will be held at the University of Southampton, United Kingdom on 18-19 April 2002.

14. Actions from previous meetings

3.4 Tijany to report on the usage of the new generation RIPE measurement boxes.
- Ongoing

4.2 Simon and Victor to work on specification of the monitoring system for premium IP service. They ought to report to TF-NGN
- Ongoing

4.6 GARR and GRnet to start performing the baseline tests.
- Closed

4.9 Wim Barbaix to collect proposals from TF-NGN for further interconnections and experiments with ATRIUM.

1 proposal about QoS testing was received. More proposals not coming yet.
- Ongoing

4.13 Task Force members to send Mauro requirements for experiments on the Garr-G testbed. No request was received. The test bed is still available until November 2001
- Closed

5.1 Robert Stoy to provide detailed work plan for the QoS & Multicast activity before the 6th TF-NGN meeting.

?

5.2 Tiziana to coordinate discussion about loan of Equipment.

Exchanges going on with Juniper Networks for testing equipment after Cisco

- Ongoing

5.3 the TF-NGN group to discuss how to carry over the MPLS activity.

At the meeting Roberto will report on potential usage of MPLS techniques in GEANT

- Ongoing

5.4 Lada to invite Sven Ubik and his colleagues in Prague to get involved in the measurement activity.

- Done

15. Open actions

3.4 Tijany to report on the usage of the new generation RIPE measurement boxes.

- Ongoing

4.2 Simon and Victor to work on specification of the monitoring system for premium IP service.

They ought to report to TF-NGN

- Ongoing

4.9 Wim Barbaix to collect proposals from TF-NGN for further interconnections and experiments with ATRIUM.

1 proposal about QoS testing was received. More proposals not coming yet.

- Ongoing

5.1 Robert Stoy to provide detailed work plan for the QoS & Multicast activity before the 6th TF-NGN meeting.

?

5.2 Tiziana to coordinate discussion about loan of Equipment.

Exchanges going on with Juniper Networks for testing equipment after Cisco

- Ongoing

5.3 the TF-NGN group to discuss how to carry over the MPLS activity.

At the meeting Roberto will report on potential usage of MPLS techniques in GEANT

- Ongoing

6.1 Tiziana to work up an initial work plan for test and implementation of the LBE service

6.2 Dimitrios to coordinate the development of an activity plan to continue and extend the OvPN performance test.

6.3 Vladimir to prepare a proposal about Accuracy of the Network Time Synchronization

- DONE

6.4 Dimitrios to poll the TF-NGN email distribution list to find out how many RIPE boxes are installed in NREN POPs and then define how they can be used in the OvPN performance measurement.

6.5 Dimitrios, Roberto, to work out scenarios for CCC on GÉANT