

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Draft Minutes of the 7th TF-NGN meeting 18th and 19th of April 2002 University of Southampton, United Kingdom.

Valentino Cavalli, Kevin Meynell,
Issue 1, 21 May 2002

Attendees

Name	Organisation	Country
-----	-----	-----
Michael Allenby	UKERNA	United Kingdom
Alain Bidaud	Crihan/RENATER	France
Mauro Campanella	GARR-INFN	Italy
Valentino Cavalli (Secr)	TERENA	-
Tim Chown	Univ. of Southampton	United Kingdom
Bruno Ciscato	Cisco Systems	Italy
Axel Clauberg	Cisco Systems	Germany
Tiziana Ferrari	INFN-CNAF Bologna	Italy
Avgust Jauk	ARNES	Slovenia
Joop Joosten	CERN	Switzerland
Dimitrios Kalogeras	GRNET	Greece
Sabine Kühn	DANTE	-
Olav Kvittem	Uninett	Norway
Ladislav Lhotka	CESNET	Czech Republic
Kevin Meynell	TERENA	-
Marcin Michalak	ULB-STC	Belgium
János Mohácsi	DANTE	-
Christian Müller-Böhm	JOIN/Univ. of Münster	Germany
Julio Orozco	ENST Bretagne	France
Antonio Pinizzotto	IIT-CNR	Italy
Agnès Pouélé	DANTE	-
Michal Przybylski	PSNC	Poland
Jürgen Rauschenbach	DFN	Germany
Esther Robles	RedIRIS	Spain
Rudolf Roth	Fraunhofer FOKUS	Germany
Roberto Sabatino (Chair)	DANTE	-
Reeta Salonen	FUNET	Finland
Trond Skjesol	Uninett	Norway
Vladimir Smotlacha	CESNET	Czech Republic
Robert Stoy	DFN	Germany
Jean-Marc Uzé	Juniper Networks	France
Stig Venaas	Uninett	Norway
Steven Williams	UKERNA/UW Swansea	United Kingdom
Octavio Medina	ENST Bretagne	France
Miguel Angel Sotos	RedIRIS	Spain
Victor Reijs	HEAnet	Ireland
Chris Ward	SPiRENT	United Kingdom

Alan Way
Franz Widhofner

SPIRENT
University Linz/ACOnet

United Kingdom
Austria

Apologies

Yves Schaaf
Wim Sjouw

Restena
Univ. of Amsterdam

Luxembourg
The Netherlands

Online presentations

<http://www.terena.nl/task-forces/tf-ngn/presentations7.html>

1. Introduction and GÉANT update

Roberto Sabatino welcomed the participants, reviewed the list of open actions and presented the latest GÉANT updates since January 2002. He said that the STM-4 to Lisbon was in place and the GÉANT router in Poland was fully connected. The link to Romania had been upgraded. Bulgaria was not connected yet and it was not known when this would become possible. Tenders for connectivity to Israel, Malta and Croatia had been finalised, DANTE had just started a tender process to upgrade Ireland capacity to GÉANT.

Test of connectivity to North American research networks were completed and the operational phase started. Eight connections to GÉANT for Commodity IP traffic had been set up (4 from KPNQ, 4 Global Crossing), seven of them were ready whereas the one to Vienna was still under testing.

By contract GÉANT supports EC projects and user groups. Contacts are well ahead to support ATRIUM, DataTAG and four more projects at the moment. ATRIUM plans to connect the Alcatel 7770 in Poznan to the rest of the ATRIUM test-bed in France and Belgium. Roberto said that DANTE was working on a proposal to provide capacity for the test-bed on the basis of two possible alternatives, the one being to extract it out of NRENS by using DLCI on their production interfaces and the other to provide additional GÉANT access capacity. The second option would technically much simpler but administratively more complex.

2. Update on results of sequin

SEQUIN held a successful workshop on 1st February 2002 with 35 attendees. Several projects have expressed their requirements and two of them already started collaboration. AQUILA has sites in Vienna and Warsaw connected with Premium IP at 2 Mb/s, MOICANE will follow in April with sites in Greece, Italy, Portugal and maybe Romania.

The conclusion of SEQUIN showed how to implement Premium IP with real users, and defined a monitoring architecture part of which is already implemented (the taksometro, at the moment providing measurements on bandwidth and packet loss). As next steps NRENS will have to conform to the specifications to support Premium IP. The results will be presented to the NREN PC meeting on 14 May 2002.

3. Less than Best Effort + Alternative Best Effort

Tiziana presented results of her investigation into LBE and ABE, she recalled that a GÉANT deliverable on LBE is due in September this year. Tiziana maintains web pages for the group activity at <http://www.cnaf.infn.it/~ferrari/tfngn/lbe/>.

LBE is more penalized than BE and the performance of BE should be better under congestion than in a flat service model. LBE is suited best for applications tolerant of large packet loss, high delay and jitter. It is potentially useful to protect production traffic from test bulk traffic (like traffic originating from the DataGRID project), but is not suited for high-performance TCP-based applications. Additionally, Mauro remarked that LBE can waste resources because packets can travel 90% of the path and get dropped before they reach their final destination whereas with Premium IP packets are dropped as soon as possible.

Tiziana proposed a work programme for the LBE activity in TF-NGN based on two deployment scenarios:

- 1) a customer-based approach where marking is done at the end-systems and scheduling is located at the customer network/NREN borders;
- 2) a customer and NREN/GÉANT-based approach with defined LBE domains at NRENS and GÉANT and scheduling at all possible congestion points.

She suggested to do service specification first, then start a phase 1 with baseline testing of the implementation and the feasibility of LBE followed by a phase 2 with performance analysis based on scenario 1 (DataGRID test-bed sites possible) and a phase 3 with implementation of scenario 2, possibly on production networks. Interest to do some work in the proposal was expressed by people from the UK, Slovenia, Spain and Norway. Tiziana expected to be able to lead the activity until Summer 2002 and agreed to share the overall coordination with Tim. Tim suggested to involve some people from the USA and said he would liaise with them be at the Internet2 Member Meeting in May. If feasible TERENA could organize a BoF on the topic at TNC-2002 (<http://www.terena.nl/tnc2002/>).

Tiziana explained the principles of ABE (Alternative Best Effort) which uses the Diffserv blocks but differs in providing delay guarantee. BE Packets are marked as green (low end-to-end delay) or blue (higher overall throughput). The service is useful for limited classes of applications, and therefore should be seen as complementary to Premium IP. Anyway, it would be very beneficial to applications like IP-telephony which are not very robust with respect to packet loss. What makes ABE interesting is that it addresses a different market from existing Diffserv proposals, moreover it does not need policing. It was observed that ABE is an interesting research area, but there is no real deployment or even good implementation of ABE schedulers available at the moment. A lot of useful information about ABE is provided on the web at <http://www.abeservice.com/>

4. Multicasting

Janos gave a presentation on the GÉANT multicast service.

Seventeen of the twenty-six NRNs had already been connected to the multicast service and the remaining six would be added shortly. These were EENet, HEANET, LITNET, RENATER, RESTENA and RoEduNet. Multicast Beacon Server was being used to check whether multicast traffic was being forwarded inside GÉANT and not via external routes.

There had been several problems with the original Beacon implementation that was written in Java, so it had been rewritten in PERL and was now up to Version 0.8. This still worked with the older Java-based clients, although the client was also being rewritten in PERL. The software could

be found at <http://dast.nlanr.net/Projects/Beacon/> although it also required the Apache web server to be installed.

Monitoring of the multicast traffic was enabling a service level specification to be developed. For example, up to 17 Mbps of multicast traffic had been recorded from Abilene.

Michal also gave a presentation on the PSNC effort to add functionality to Multicast Beacon. This included a history feature, statistical graphs on the web-based interface, a message trigger module and authentication via SSL. The latter feature would allow secure access via the web to the server. This software could be found at <http://noc.man.poznan.pl/>.

Robert then gave an update on the DFN/G-WIN multicast service. They had taken steps to improve their best-effort service and had fixed a bug on their Cisco GSRs that caused multicast packet loss when an asynchronous interface was enabled. Multicast forwarding had also improved with the upgrade to IOS Version 12.0.21S. They were also monitoring performance using Multicast Beacon.

5. QoS Measurement

Valentino briefed the participants about the current status of the Internet2 End-to-End Performance Initiative. He attended a workshop in Tempe AZ in January 2002 and provided a report available on the TERENA web site at <http://www.terena.nl/news/reports/2002/report01.pdf>.

The latest development was a draft document provided by Russ Hobby about a performance measurement infrastructure <http://e2epi.internet2.edu/pma01.shtml>. The End-to-End Performance Initiative web site contains many more pieces of useful information <http://e2epi.internet2.edu/>.

In a discussion about the activities coordinated by the Internet2 group Victor recalled the idea of organizing a network of Performance Response Teams (PRT) alongside the scheme of the Security Incident Response Team organization idea. Valentino suggested him to contact people at Internet2 to find out more updated information about that. It was remarked that an organization infrastructure for performance monitoring/debugging needs to be implemented in GÉANT.

Victor reported on the workshop on scalable QoS held in Annapolis. The presentation examined various reasons why no real QoS services are widely deployed, features of QoS provisioning and QoS metrics. A number of open issues were discussed, like the need of having accurate performance measurement data before starting to implement QoS solutions.

Valdimir gave an extensive presentation of Accuracy Measurement of NTP client (http://www.terena.nl/task-forces/tf-ngn/presentations/tf-ngn7/owd_sync.ppt). Participants in the activity are , HEANet, GARR, FhG/FOKUS and CESNET, in collaboration with RIPE TTM.

The CESNET tools can be used for one-way delay measurement, using two boxes synchronised to NTP server(s), one-way delay is measured as the difference between timestamps. Vladimir explained three different setups: one using two NTP servers, a second set up using only one NTP server and (where the second box is synchronised by the network) a third one, structurally more complex, but very general and capable of being used both at NREN POPs and institutions.

6. Scalable efficient personalised and automated QoS provisioning

Dimitrios presented a proposal for personalized and semi-automated QoS provisioning based on user/application requirements. Slides are provided on the TF-NGN web site, a demonstration is available at <http://linux.noc.ntua.gr/qos>.

One of the goals of the project is to provide (almost) automated administration, allocation of resources would be done by the administrator, implementation is coming from user request. Dimitrios showed two approaches, the one LAN-based (RSVP) and the other WAN-based (Diffserv).

Problems that need to be addressed regard authentication, personalisation, and signaling. Dimitrios proposed a trust model based on the example of gold traffic between NTUA and UoP passing through GRNET. LDAP is useful to store profiles, however it is not easy to store policies with LDAP, and coping with personal information through many domains does not scale, so in order to provide the service some manual operation is needed. Modelling is provided on the base of profiles and policies.

7. Current status and perspectives of AF testing

Octavio Medina introduced Julio Orozco who will continue the work started by him in this domain. Octavio presented the status of the experiments carried out so far and the major results achieved.

In the future he would like to test and simulate something closer to real Diffserv networks with multiple domains. AF offers priority-based drop in case of congestion, but current transport protocols do not know how to take advantage of that, this is an example of what can be done to improve the PHB for some specific applications. On a longer perspective one can also imagine a service proposal for CBR sources.

There was a discussion about the decreased interest for interest for IP+ and in that context AF was felt as a possible alternative to LBE. Testing could be done about how can AF improve ECN and any other possible use of it. Tiziana would be interested in collaborating in the IP+ and ECN part.

8. MDDR Testing

Tiziana showed the key results of a number of tests of MDRR scheduling on the CISCO GSR router carried out during the past three months as documented at <http://www.cnaf.infn.it/~ferrari/tnfgn/mdrr/no/>.

The first test measured the performance of Engine 2 and Engine 3 interfaces for different levels of traffic load and congestion and showed that the CPU load is very stable with Engine 3 and does not increase as much as with Engine 2 cards.

Test 2 verified the capability of MDRR to protect traffic in two queues (precedence 1 and background precedence 0 traffic) with or without using WRED, and showed that it can be ineffective whenever buffer memories become unavailable in a given memory pool due to congestion. However, if there is no contention of buffer memory between different classes MDRR effectively protects precedence-1 traffic from loss.

Test 3 analysed MDRR traffic isolation in presence of five classes with and without WRED. Benefits are seen with WRED but some packet loss was still experienced because the configuration was too aggressive. Changing configuration parameters showed that the drop was due to WRED and not to contention. The test was also carried out with another extreme configuration by relaxing the minimum and maximum values of WRED parameters.

The fourth test focused on DRR and low-latency (strict priority) queuing delay and IP delay variation.

9. IPv6

Alan Way of Spirent plc gave a presentation on their Adtech AX/4000 broadband analysers that supported IPv6. Spirent had recently acquired both Adtech and Smartbits who were well-known for their network analysers. A demonstration of the system was also provided.

It was hoped that DANTE could be loaned some equipment in order to start generating some realistic traffic. At the present time, there were insufficient applications to generate high volumes of traffic which meant the network could not be seriously stress-tested.

Janos then gave an update on the 6NET project. The connections had been ordered, but had been slightly delayed until late-May. An interconnection to GTPv6 via the Juniper M5 in Paris was also planned. In addition, the access policy for non-6NET participants was currently under discussion.

Bruno reported on progress at the 6NET staging laboratory in Brussels. All the routers had been delivered and configured on-site, which enabled a complete implementation of the network to be tested. The routers were now being shipped out and recipients were asked to open the packaging and check the contents as soon as possible. The recipients were asked to contact Bruno if anything was missing.

Three new 6NET partners had also been approved by the European Commission. These were CESNET (Czech Republic), HUNGARNET (Hungary) and PSNC (Poland). It was likely that ETRI (South Korea) would join at a later stage.

Tim asked how the new partners would be connected. Bruno replied they would initially have tunneled connections, but there was a proposal for them to transfer some budget to DANTE to arrange native connectivity.

Bruno also asked whether partners could add location records (latitude and longitude) to their DNS servers. This would allow an accurate network map to be plotted.

Michal then gave a presentation on the IPv6 network in Poland. They were using the Zebra software router running on Linux to support upwards of thirty connections. The next phase was to deploy a production-quality network using Cisco 7500 routers.

Finally, Ladislav also gave an update on the CESNET IPv6 network (<http://www.cesnet.cz/english/project/ipv6/>).

10. Optical Networking

Victor presented the ASTON project, a collaborative proposal aiming at testing technologies towards optical networking. The proposal was presented to the TERENA TTC, with the purpose of evaluation only because some commitments were still missing from a few NRENs, equipment manufacturers and FLAG Telecom. The TTC feedback was very positive and there was hope that all pending issues would be solved on time for project submission in June 2002.

The topology of the test-bed had been proposed by looking at the activities and the availability of PoPs on the FLAG Telecom network, however it was not yet finalized and could have been different in principle. The proposed activities were Bandwidth on Demand, Network Management, Transport at 40Gb/s and beyond, 10Gb/s Ethernet connectivity over long distances and preparation of a follow up optical test-bed proposal to the EC 6th Framework Programme.

Victor suggested to start some activities within TF-NGN without waiting for commitment from the providers. It was agreed that Mauro and Michal would work on the 40 Gb/s, CESNET and PSNC would work on long-distance 10GE. Moreover, TF-NGN would steer a brainstorming about BoD on the FLAG-testbed mailing list. One could start discussing technology (G-MPLS, OBG, L2-VPN, etc.), interfaces: (web-G-MPLS) and equipment requirements.

Michal made a brief presentation about testing Alcatel 1696 Metro Span. PSNC tested configuration options using craft terminals. The presentation with details of initial tests and results is available at <http://qos.man.poznan.pl/files/optical.ppt>. People from the TF-NGN group were invited to collaborate, even remotely for some limited time.

11. Date of next meeting

The 8th TF-NGN meeting will be held at Fraunhofer FOKUS in Berlin, Germany on 1-2 July 2002.

12. Actions from previous meetings

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

- CLOSED

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

- ONGOING

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

Proposals were received.

- CLOSED

5.1 Robert Stoy to provide detailed work plan for the QoS & Multicast activity before the 6th TF-NGN meeting. Robert presented initial plans at TF-NGN 7 in Southampton

- CLOSED

5.2 Tiziana to coordinate discussion about loan of Equipment.

Discussion with Juniper Networks was taking place at TF-NGN 7 in Southampton

- ONGOING

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

Discussion was taking place at TF-NGN 7 in Southampton

- CLOSED

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

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

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

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.
- ONGOING

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

13. Open actions

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

5.2 Tiziana to coordinate discussion about loan of Equipment.
Discussion with Juniper Networks was taking place at TF-NGN 7 in Southampton
- ONGOING

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.
- ONGOING