

## TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Report on the 20th TF-NGN meeting

13 January 2006

New Hall College, Cambridge, UK

Issue 1.2, Valentino Cavalli

### Attendees

<u>Name</u>	<u>Organisation</u>	<u>Country</u>
Lada Altmannova	CESNET	Czech Republic
Martin Attenson	Glimmerglass	UK
Bartek Belter	PSNC	Poland
Mauro Campanella	Consortium GARR	Italy
Valentino Cavalli (p.p. Secretary)	TERENA	-
John Chevers	DANTE	-
Tim Chown	University of Southampton	UK
Hans Döbbeling	DANTE	-
Jerome Durand	RENATER	France
Michael Enrico (Chair)	DANTE	-
Rob Evans	JANET	UK
Ernst Heiri	SWITCH	Switzerland
András Jákó	NIIF/HUNGARNET	Hungary
Avgust Jauk	ARNES	Slovenia
Dimitrios Kalogeras	GRNET	Greece
Radek Krzywania	PSNC	Poland
Olav Kvittem	UNINETT	Norway
Felix Kugler	SWITCH	Switzerland
Simon Leinen	SWITCH	Switzerland
Athanassios Liakopoulos	GRNET	Greece
Simon Muyal	RENATER	France
Wojcieck Naurot	PSNC	Poland
Victor Olifer	UKERNA	UK
Jan Radil	CESNET	Czech Republic
Victor Reijs	HEAnet	Ireland
Esther Robles	RedIRIS	Spain
Laura Serrano	RedIRIS	Spain
Robert Stoy	DFN	Germany
Stanislav Sima	CESNET	Czech Republic
Derek Stagg	Marconi	UK
Miguel Angel Sotos	RedIRIS	Spain
Robert Stoy	DFN	Germany
John Taylor	Glimmerglass	UK
Sven Ubik	CESNET	Czech Republic
Stig Venaas	Uninett	Norway
Joseph Vojtech	CESNET	Czech Republic

## **Meeting Presentations**

All presentations from the meeting are available online at <http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn20/presentations.html>.

This report records the main discussion items and actions arising during the meeting. Readers should refer to the presentations for detailed information.

### **1. IPv6**

#### **1.1 IPv6 Deployment on GÉANT, Marian Garcia Vidondo**

Marian presented a comparison between the volume of IPv6 traffic to and from the GÉANT network in 2004 and 2005. There was a significant growth in 2005, almost doubling the volume by April. The volume decreased in the following months.

Marian showed a chart with the IPv6 utilization volume per NREN in December 2005, half of them have less than 2.5Gbit/s/month traffic.

Approximately 50 multicast IPv6 routes were received in GÉANT2.

The GÉANT2 rollout was ongoing. It was not yet a stable environment for introducing such new features as embedded RP, and tests and monitoring of the new network would need to be carried out first. A JunOS upgrade was due in February, and DANTE would take this opportunity to test new features. As a result, the current estimation for introduction of Embedded RP was in April 2006.

At the time of the last APM meeting only NORDUnet and RENATER had implemented Embedded RP. In the meantime UKERNA had implemented it, but not on the production network. DANTE will carry out the test and, if needed, exchange experiences with those NRENs that have done it already.

#### **1.2 IPv6 Multicast Deployment Status on m6bone and IPv6 Multicast Address Allocation Methods, Jerome Durand**

There were no updates about the m6bone. Approximately two connections were requested each month on a regular basis, but for the rest there was not much activity.

Jerome focused the first part of his talk on the following tools for IPv6 multicast monitoring:

- DBeacon has a distributed architecture, is more robust than multicast Beacon, and has a good development/support team backing it. Further development is needed in the area of documentation, alarms and unicast reporting.
- QoS Metrics probes can analyse streams of a multicast conference and rate the quality based on ITU-defined standards.
- MRD6 is an IPv6 multicast router for Linux which implements PIM-SMv2 and MLDv2.

The second part of Jerome's presentation was on Embedded RP. He presented a plan for support of Embedded RP in GÉANT2 and NREN networks. The idea was to have a single global RP in RENATER, so that NRENs with no support for Embedded RP could statically configure RENATER Embedded RP. He provided details of the schedule, and discussed the issues associated with his proposal. The main issues were the choice of address - how applications find the address to be used. Jerome had attempted to investigate the issues in some IETF draft proposals, but got very little support. However, this should not prevent NRENs from making use of Embedded RP.

Jerome mentioned training events organised by RENATER on multicast in February, and on IPv6 in April.

### **1.3 Multicast RSS Model and Planning Embedded RP Application tests, Stig Venaas**

Stig focused the first part of his talk on the presentation of the SSMPing tool for testing multicast connections. SSMPing had been slightly updated and was available for Debian and other platforms.

RSS feeds are based on the use of HTTP, with clients regularly polling a server for updates. Stig proposed an implementation of RSS using multicast, in which clients join a multicast group and a server pushes out the updates. The main advantage is that the clients do not need to wait for the server to poll at certain intervals, and would get the information feed immediately.

Stig explained the architecture and the usage of his prototype. He invited the meeting attendees to download and test the software, but warned them that the current implementation is rather simple and there were no reliability guarantees.

### **1.4 6Bone Addressing Phase-Out and IPv6 Training Material, Tim Chown**

Tim reminded the audience that the 6Bone has been problematic, and that the 6Bone address space (3ffe::/16) was being phased out. He asked whether this was already happening in NRENs, if they needed this address space, or whether there were any other concerns with the phase-out. Nobody in the audience was using 6Bone addresses, and neither was it known whether other NRENs were using them.

Tim suggested that NRENs start pooling IPv6 training material. He mentioned existing material from UKERNA and the 6DISS project. 4-5 NRENs in the audience said they had such material, and Tim suggested sending a pointer to the UKERNA material in order to start the discussion.

## **2. Brief GÉANT2 Rollout Update, Michael Enrico**

The strategy for the GÉANT2 Rollout was to build a new network, by using new and existing resources in order to keep communications between research organizations unaffected at all times.

Michael provided a very brief overview of the rollout and some details about the network implementation – an overview of a typical GÉANT2 PoP, as well as new routes. The last slide of his presentation provided a snapshot of the network implementation at the time of the meeting.

## **3. TEIN2: A Regional Research Network in Asia, Michael Enrico**

The European Commission has been funding a number of regional research networks, including SEEREN, ALICE, EUMEDCONNECT, and TEIN2 (the Trans-Eurasia Information Network). The Virtual Silk Highway network had been funded by NATO.

Michael summarised the TEIN2 project stages and partnership. Beneficiary countries are: China (CERNET), Indonesia (ITB), Malaysia (MDC), Philippines (ASTI), Thailand (ThaiREN), Vietnam (MOST). The Commission has been funding 80% of the costs for the beneficiary countries.

The TEIN2 network was in the process of being built, based on the topology agreed in Tokyo on July 2005 (as depicted on slide 6). There were three TEIN2 PoP: one in Singapore was operational at the time of the meeting, whereas the ones in Beijing and Hong Kong were under construction. The TEIN2 network NOC is going to be operated by the CERNET NOC. The routing policy was being defined at the time of the meeting, but would regulate transit to Abilene via TransPAC2 through Japan. OWD on the direct link between Singapore and Frankfurt was 130msec .

Although APAN pre-existed the TEIN2 project, the former had not really deployed a regional network, but just created a mix of bilateral agreements between existing networks. Michael remarked that with TEIN2 the first substantial inter-Asian research network entering into service.

## **4. Dark Matter Estimation and MPLS VPNs ++ , Dimitrios Kalogeras**

Dimitrios explained that “Dark Matter” was addresses which were allocated but not assigned, and constituted a large part of the Internet. These addresses can be exploited by Internet scanners and worms, and can be used to detect early signs of attacks or worm contamination.

There are two possible ways of discovering them:

- Querying of the routing table and subtraction from the allocated size
- Monitoring the ARP table or the DHCP server of a campus network for requests.

This knowledge could be used to collect and aggregate the dark matter addresses, monitor the network traffic and forward it to recipients. However, there would be some issues with this, so Dimitrios proposed to use MPLS VPNs ++ as a tool to detect them.

There was positive feedback from the audience about this, as some NRENs had already had similar thoughts. It was also suggested to develop the ideas in a proposal to add to JRA2, but in general it was thought that this would not be work for a JRA because tools are already available. NRENs could just go ahead and implement the suggested proposal.

## **5. User Support and VLBI, John Chevers**

John introduced the general concept of underlying user support in GÉANT2, and more specifically support for VLBI. This was one of several demanding pan-European projects supported by GÉANT2. John explained the user-support mechanism for GÉANT2, which includes user questionnaires, evaluation, planning, designing and actually supporting the users involved in the projects.

John reminded the audience about the VLBI project and how it worked. He also presented the VLBI network in Europe and reiterated how this overcame the problems with transferring many tapes from all over Europe to the correlator in Dwingeloo.

The GÉANT2 user support for eVLBI includes: the EVN-NREN forum for the communities to meet and address planning issues, solving problems and discussing results; monitoring tools like the GÉANT Weathermap and Taxametro; counters on GÉANT links; diagnostic tools (BWCTL); the PERT and also specialised tools such as UDPmon.

John listed the progress of the eVLBI in 2004 and 2005, and discussed the issues related to bulk TCP transfer. He mentioned possible options to solve the problem: the usage of dedicated network links (lightpaths) and using Logistical Session Layer (LSL) over the existing IP network. He then explained LSL in more detail, which was currently being tested on GÉANT2 under SA3.

John concluded with the view about the future extension of eVLBI into EXPReS, which would support 16 telescopes in Europe, Asia, Australia, Latin America and North America.

## **6. GMPLS in GÉANT2, Otto Kreiter**

Otto reported on the support for GMPLS in the Alcatel 1678 Metro Core Connect (MCC) platform, and provided details about the technology it uses for routing and signalling.

Support for interfaces in the Alcatel code was as following:

- NNI in operation
- UNI 1.0 user network interface for SDH only
- UNI 2.0 not yet in production, Ethernet user network interface support was only available in experimental code
- E-NNI only in experimental code.

Otto showed that the network can support multiple path initiators and provision different types of connections. He illustrated the types of protection supported, and then described how a GÉANT2 hybrid node would be managed. The legacy element would be controlled by the SDH NMS, and the new network element by the GMPLS control plane (this holds only for the testbed and should not interfere with the production network).

The status of support for multi-region multi-layer network MRN/MLN was as following:

- GMPLS control plane is available for SDH only
- GMPLS control plane for Ethernet is being discussed at the IETF – Alcatel is very active there.

Otto provided an example how to set up a path involving multi-domain/multi-region operation. Different control planes are needed now because they do not interact with each other. A unified control plane was not available yet.

Otto concluded by mentioning a GMPLS test session that was scheduled to take place at the end of January at Alcatel. This aimed to better understand the GMPLS operational aspects and how the hybrid node should be used; the importance of MRN/MLN future developments for GÉANT2; and the helpfulness of a stand-alone GMPLS testbed.

Some final remarks were made about ASON OIF UNI and IETF UNI implementations. These do not practically interact.

## **7. UCLPv2, Sergi Figuerola**

Sergi introduced the i2CAT foundation, then listed the drivers to user-controlled networks and explained the UCLP and Service Oriented Architecture (SOA). The latter enables network to be integrated with other web service applications.

Sergi gave detailed examples of how UCLPv2 works, how to set up a path, configure nodes, and provision nodes etc...

The audience asked some questions about policy restrictions, which are currently dealt with by means of access lists – are they managed, and in that case how is it possible to handle them? There was some consensus that these issues would be solved now by means of certificates (AAI).

## **8. Integrating Layer 1 Switching into Advanced Optical Networks, John Taylor**

John introduced Glimmerglass and what the company did. The basic goal of their equipment is to add dynamicity at the physical layer (L1), normally a static or human controlled environment. He gave examples of the technology based on customers' projects and applications. These include:

- control plane developments at CERN/Caltech Mona Lisa software
- Optiputer
- HOPI node providing interface between NLR and Abilene

In terms of network protection, John illustrated the case of their largest customer, the Amsterdam Internet Exchange: 237 customers, 382 ports, 127Gbit/s peak traffic, 10% monthly growth.

Glimmerglass equipment does not have its own control plane. Their boxes have GUI and command line interfaces, but these have simple functionality and there was no plan to support a distributed control plane in the near future.

## **9. MUPBED Update, Mauro Campanella**

Mauro gave a brief update on the MUPBED project. He started by explaining its rationale, which builds on a number of testbeds already available in Telco premises. These are interconnected by MUPBED in order to integrate and validate “telecommunication operator’s concepts” within research networking infrastructure, particularly with respect to on-demand circuit switching techniques like ASON/GMPLS. Interoperability and understanding of ASON GMPLS interaction is important for GÉANT2 and NRENs.

Mauro focused on Data Plane, Control Plane, Management Plane; MUPBED wants applications to interact and be able to control all that. He showed the project testbed and explained in detail the set-up of the node in Berlin. The project has been testing UNI-C 2.0 Ethernet and E-NNI which have been successfully demonstrated in some recent experiments (see slide 16 and next).

A demonstration on the testbed was carried out by sending 400 Mbps uncompressed video flows for more than two days. This worked well without any need to provision QoS.

## 10. Routing Integrity with BoD Circuits, Dave Wilson

Dave discussed the implications of establishing many bandwidth-on-demand connections to arbitrary locations on routing. Campus networks have routing policies that consist of lists of assigned IP prefixes, and information on how those prefixes are routed. These are based on assumptions about a hierarchical structure.

Dave saw no problems with creating arbitrary connections between NRENs, but there would be issues doing this between institutions and between users inside institutions there, because their networks were not built with complex routing policies in mind. There were three possible solutions, none of which were fully free from problems.

A discussion followed where some people agreed on the issues, whilst others did not see any inherent problems with customer addressing as it depended on how the customers used BoD. In this case the issue need to be tackled in a different way, by teaching the users how to make better use of the BoD service. The service had to come with a user manual telling users how to behave, and outlining the known issues etc...

## 11. Next Meetings

A host is being sought for the following meeting in May or June 2006.

### Summary of Actions

ACTION 050113-01	Victor to provide a draft of the term definition document before the next TF-NGN meeting.	Outstanding
ACTION 050113-03	Marcin to provide a list of features to be tested and the test methodology in work item 9.7.	Outstanding
ACTION 050728-01	Stig Venaas to draft test plan for running Embedded-RP on GN1 testbed.	Outstanding
ACTION 040728-02	Kevin Meynell to contact GRNET about arranging alternative dates for the next meeting.	Done
ACTION 050729-03	Michael Enrico and Kevin Meynell to discuss organisation of workshop on dynamic routing configuration.	Outstanding