TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Report on the 18th TF-NGN meeting
28-29 July 2005
RENATER, Paris, France

Issue 2, Kevin Meynell

Attendees

Name	Organisation	Country
Eddie Beier	T-Systems	Germany
Emilie Camisard	RENATER	France
Mauro Campanella	Consortium GARR	Italy
Andrea di Donato	UKERNA	United Kingdom
Jerome Durand	RENATER	France
Larry Dunn	Cisco Systems	-
Michael Enrico (Chair)	DANTE	-
Rob Evans	ULCC	United Kingdom
Adam Filby	Force10 Networks	-
Alexander Gall	SWITCH	Switzerland
Marcin Garstka	PSNC	Poland
Vangelis Haniotakis	University of Crete	Greece
Adrás Jákó	Hungarnet	Hungary
Avgust Jauk	ARNES	Slovenia
Radek Krzywania	PSNC	Poland
Loukik Kudarimoti	DANTE	-
Olav Kvittem	Uninett	Norway
Simon Leinen	SWITCH	Switzerland
Kevin Meynell (Secretary)	TERENA	-
Simon Muyal	RENATER	France
Jordi Palet	Consulintel	Spain
Esther Robles	RedIRIS	Spain
Victor Reijs	HEAnet	Ireland
Nicolas Simar	DANTE	-
Trond Skjesol	Uninett	Norway
Andreas Åkre Solberg	Uninett	Norway
Bernard Tuy	RENATER	France
Sven Ubik	CESNET	Czech Republic
Jean-Marc Uzé	Juniper Networks	France
Jeroen Valcke	BELNET	Belgium
Stig Venaas	Uninett	Norway
Chris Welti	SWITCH	Switzerland

Apologies

Name	Organisation	Country
Tim Chown	University of Southampton	United Kingdom

Dimitrios Kalogeras	GRNET	Greece
Felix Kugler	SWITCH	Switzerland
János Mohácsi	NIIF/HUNGARNET	Hungary
Dennis Paus	SURFnet	The Netherlands
Dirks Schroetter	Cisco Systems	-
Wilfried Woeber	ACOnet-CERT	Austria

Meeting Presentations

All presentations from the meeting are available online at: http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/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. GÉANT2 Network Infrastructure, Michael Enrico

Michael gave a presentation about the GÉANT2 services (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/enrico-geant2update.pdf). The procurement process was close to completion, with connectivity and equipment providers having been selected, and contracts being finalised.

The GÉANT network had 21 PoPs utilising Juniper M160 and M40 routers. There were 16 x 10 Gbps and 13 x 2.5 Gbps lambdas, as well as several lower speed links serving 29 NRENs (at speeds of up to 10 Gbps). There was also a 12 Gbps connection to US research and education networks via the New York PoP, plus other connections to the Mediterranean, and Latin American regions.

By contrast, GÉANT2 would have 25 PoPs utilising Juniper T640, M160 and M40 routers serving 30 NRENs. There was a total of 11,600 kms of fibre, which included 50 x 10 Gbps owner-operated lambdas, 9 x 10 Gbps leased lambdas, and 8 x 2.5 Gbps leased lambdas. The US connectivity would also be upgraded to 4 x 10 Gbps (via the New York PoP) and an additional connection to South-East Asia added.

GÉANT2 would have a hybrid infrastructure offering both IP services and lambdas. Both Alcatel 1626 Light Manager and Alcatel 1678 Metro Core Connect switches would be utilised, as these were able to support both SONET/SDH and GE over the distances required.

Marcin asked when the packet switching matrix would be available for VLANs. Michael replied this would be available by November 2005.

Marcin also asked what would be used to mark the VLANs. Michael replied VC4 would be utilised.

Eddie asked whether spanning tree would be available to those using an Ethernet switch fabric. Michael was uncertain about this.

2. GÉANT2 Services, Michael Enrico

Michael gave a presentation about the forthcoming GÉANT2 services (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/enrico-geant2services.pdf). GÉANT2 would continue to support the usual IP services such as Best Effort, LBE, PIP and Multicast (IPv4 and IPv6). These would be enhanced by new features as required (e.g. Embedded RP for IPv6 Multicast), improved performance monitoring, automated PIP provisioning, and improved network security. Emulated L2 P2P services would continue to be supported, and there would be an additional portfolio of Layer 1 and 2 P2P services. This would include high-capacity (up to 10 Gbps) lightpath services largely based on Ethernet framing, although there would also be experiments with translational services such as native GE to GFP VC-4-7v over SDH. In addition, support for various protection modes will be tested.

Marcin asked how many ports would each user be entitled to. Michael replied that it would depend on the type of subscription, the details of which still needed to be worked-out.

3. GN2 Research and Service Activity Updates

3.1 JRA1 Update, Nicolas Simar

Nicolas gave an update on the status of the JRA1 activity (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/simar-jra1update.pdf). The first pieces of equipment (IPPM-AB) had already been purchased, and the Axis and Tomcat-based web services were available. The interfaces followed the GGF Network Measurement Working Group definitions.

The OWD-IPPM and RIPE TTM work was ongoing, and the selection of NetFlow tools was being finalized. The selection of passive monitoring and visualization applications was in still in the early stages.

Phase I was due to be implemented by the end of December 2005, and included path-based visualisation, IPPM integration, BWCTL integration, CNM integration, as well as base services. Phase II was due by the end of May 2006, and would include provision of looking glass services, IPPM multicast, and NEMO integration. Phase III was due by the end of October 2006 and included NetFlow integration, authorization and authentication, and improved visualisation and end-user tools.

More information was available on the JRA1 website at http://www.geant2.net/server/show/nav.00d00a001. Information about DFN IPPM could

be found at http://www-win.rrze.uni-erlangen.de/ippm/, and about DFN CNM at http://www.cnm.dfn.de/. The Uninett NEMO website was at http://software.uninett.no/.

3.2 JRA3 Update, Michael Enrico

Michael gave an update on JRA3 (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/enrico-jra3.pdf). Over the past nine months, JRA3 had produced the requirements documents, and started work on the service specifications. These were approaching completion, and implementation and testing would start shortly. The service definition called for a point-to-point connection oriented service, which might employ SDH with GFP, MPLS L2VPN, as well as native Ethernet. Advance reservation would be possible.

The updated workplan specified the interfaces (UNI, NNI and IDM to DM), and the inter-domain manager process. The inter-domain manager coordinated the handover point between domains, managed sessions, scheduled reservations, and determined the routing (next domain and egress port). This would use a credit-based accounting scheme and would utilise existing network management systems and protocols.

They were also considering holding a joint meeting with Internet2, ESnet and CANARIE to exchange information and explore the possibility of wider collaboration. This might also include other bandwidth-on-demand activities such as CHEETAH, DRAGON, MUPBED, NOBEL and VIOLA.

Victor pointed-out that several activities in JRA3 had been over-specified and were not actually needed. Michael replied that most of these activities had not yet been undertaken, so they could still be reviewed.

3.3 JRA4 Update, Michael Enrico

Michael gave an update on JRA4 (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/enrico-jra4.pdf). This activity aimed to build and operate an enhanced testbed facility for use by other JRAs, GN2 partners and FP6 projects. It would undertake a variety of technology testing activities, including cross-border testing of dark fibre. The experience gained from these pilot activities would then be applied to the operational aspects of GN2.

Unfortunately, although GN2 user requirements had been collected, the construction of the testbed had been delayed as it was necessary to finish building the GÉANT2 production network first.

4. High Performance Ethernet for Grid & Cluster Applications, Adam Filby

Adam gave a presentation about using high-speed Ethernet for high-performance applications (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/filby-hpe.pdf). Computing clusters were traditionally high-performance computers that were separated with WANs providing only a fraction of the internal cluster bandwidth. However, this was increasingly inadequate given the amounts of data being generated by some applications. CERN alone would be generating a minimum of 8 Petabytes per year which would require 70,000 CPUs to process. The new computing grids therefore needed high bandwidth connectivity over long ranges.

The problem is that existing POS technology relies on expensive underlying SDH/ATM technology, in order to offer service differentiation and resilience. As Ethernet is the standard for interconnecting LANs, it makes sense to move it into the MAN and WAN environment.

A number of standards have been developed to facilitate this, such as 802.1ae that provides high-speed Ethernet with a seamless WAN interconnect. 802.1ad provides resilience and link aggregation, whilst 802.1q and 802.1p provide service separation and differentiation. Management functions can be provided with SNMP and RFC 3176.

In order to use 10 GE as a WAN technology, the platforms providing it need to offer carrier-grade resilience and stability. This includes redundant power supplies and cooling, redundant route processors, and redundant switch fabrics. There should be complete separation of control and forwarding, and no forwarding packets should traverse the control plane. The use of XFP interface technology also allows port densities to be increased, with a corresponding reduction in the price per port.

The Force E1200 offers a fourteen slot chassis, supporting 56.25 Gbps per slot and an overall switching capacity of 1.6875 Tbps. It can provide up 672 ports of line-rate GE, or 56 ports of line-rate 10 GE. The FTOS modular operating system is powered by 3 CPUs which each handle different functionality. The control processor handles configuration and management functions; the first router processor handles IP protocols, routing and access control; whilst the second router processor handles Ethernet and other low-level protocols.

The smaller Force10 E600s have been used to provide wide-area Ethernet over an OC-192c lightpath circuit between SURFnet and CANARIE. This demonstrated throughput of up to 9.24 Gbps with traffic generators, up to 6 Gbps using UDP, and 5.5 Gbps using TCP.

The next consideration was line speeds beyond 10 GE. The next logical step was 40 GE in order to leverage OC-768/STM-64 standardisation efforts, but the projected cost of OC-768 ports (>USD 1 million each) was not a significant enough improvement to warrant development and adoption costs. The development costs for 100 GE were likely to be similar to those for 40 GE, so it made sense to jump directly to the higher

bandwidth. It was projected that 100 GE would initially come-in around USD 10K per port.

The requirements for 100 GE were a backplane supporting 334 Gbps per slot, and optics that could support 100 Gbps over a single fibre (either as 4 x 25 Gbps or 10 x 10 Gbps). This required high-speed silicon that could run at 25 Gbps, which in itself created power supply issues as each slot would require 425W. Nevertheless, the adoption of SerDes technology (high-speed serial interface circuitry) in the next generation of ASICs, would improve speeds and significantly reduce power consumption.

Force10 took an active role in the standardisation process and were working towards both 40 and 100 Gbps technologies. They were committed to supporting whatever standards were adopted.

Mauro asked whether the equipment was really suitable for wide area networking as the 20 MB buffering per port was unlikely to be sufficient. Adam replied they could add more memory or implement congestion deployment mechanisms, and he asked how much memory should be added. Mauro suggested this could only be determined through trial and error because it depended on traffic patterns.

Marcin asked whether VPLS was supported on the 10 GE interfaces. Adam replied this was not yet the case.

Marcin also asked whether Force10 would be interested in making some of its equipment available for TF-NGN to test. Adam replied they were indeed interested in third-party testing, and had already loaned some equipment to the HOPI project in the US.

5. Transport Protocols, Radoslaw Krzywania

Radoslaw gave an update on the transport protocols activity (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/krzywania-transprotocols.pdf). Descriptions of the Quench, XCP, HS-TCP and H-TCP protocols had been added to the GN2 Wiki pages at http://pace.geant2.net, whilst descriptions for the Westwood TCP, Fast TCP and Scalable TCP protocols were currently being prepared.

The next stage was to gather performance benchmarks of existing protocols, in order to compare the relative performance of the new protocols. These results would be published on the web pages as they became available.

The testing of the new protocols would be undertaken on two testbeds. The CESNET-SWITCH testbed was already available, whilst the CESNET-PSNC testbed would be ready shortly.

PSNC was also gaining experience with the Network Simulator tool (available from http://www.isi.edu/nsnam/ns/). This was a C++ based tool that could be used to simulate

network environments using different parameters, and modules were available for various TCP implementations.

Simon added that version 2.6.13 of the Linux kernel had been restructured so that each socket could select which type of TCP to use.

6. IP Routing Work Area, Mauro Campanella

Mauro gave an update on the IP Routing Work Area (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/campanella-iprouting.pdf). There were not yet any terms of references for this activity, but the suggested research items were BGP peering load balancing, routing security, dynamic circuit set-up and tear-down, and routing and inter-domain communication. In addition, some consideration should be given to analysis of traffic patterns.

The tests could be run on the LHC (Large Hadron Collider) computing grid, as well as the cross-border fibres between Italy and Switzerland. However, volunteers were needed to start running these. It was proposed to hold a videoconference in September to discuss the work, and then meet during the next TF-NGN meeting.

Jean-Marc asked for clarification as to what routing security entailed. Mauro replied that it was about securing the routing protocols themselves.

7. IPsphere, Jean-Marc Uzé

Jean-Marc gave a presentation about the IPsphere Forum (http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/uze-ipsphere.pdf). This is an international industry-wide association of vendors, telcos and networking companies to promote the deployment of public IP infrastructures that can offer the performance and security of private networks. It was started by Juniper as the Infranet, but has moved to a self-funded model comprising more than 40 companies, including Cisco and Alcatel.

The goals are for the Internet to be a service of the overall IP infrastructure, rather itself being the infrastructure. In other words, IP service is presented as a virtual environment alongside other environments. In addition, it should be possible for different network operators to interconnect their services according to clearly defined business relationships, whilst at the same time not allowing users to access the underlying infrastructure.

IPsphere divides a service into elements, which include access, transport/connection, content/processing, and user access. Any given service may have any number of elements, and services are created by combining these. A provider participating in the IPsphere will contribute at least one element for at least one service, but will likely contribute many elements for many services. The creation of a pan-provider service

involves a web service-based exchange of messages between the administrative owner who coordinates the service on behalf of the customer, and the providers who contribute elements to the service. This process takes place at the Service Signalling Stratum level of the IPsphere Reference Architecture.

IPsphere builds upon, and conforms to IETF, ITU-T and MFA standards. The forum currently has six working groups in the areas of reference architecture, business dimensions, use case, standards synergies, showcasing, and publication. The next meeting will be on 12-13 September 2005 in Europe.

The IPsphere is a model for putting network services into a business context by linking service creation with service ordering and fulfillment. It is based on web service principles for easy exchange of information and ease of managing higher-layer services that require identity management and reliable communications. However, it is not a strategy to create services on the network, provide QoS, or manage resources at the physical level. Neither is it an alternative to the Internet – rather it is an alternative to the Internet model applied to non-Internet services.

Finally, Jean-Marc asked whether anyone from TF-NGN was willing to attend the next IPsphere meeting. Any volunteers should contact him.

8. Hands-on evaluation of new routers and switches, Marcin Garstka

Marcin gave an update on the Cisco CRS-1 that PSNC was currently evaluating (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/garstka-newrouters.pdf). This was the new carrier-class router and was available with two types of chassis: one with eight slots and a total capacity of 640 Gbps, and one with sixteen slots and a total capacity of 1.2 Tbps. There was also a multi-shelf configuration for up 72 line cards with a total capacity of 92 Tbps.

Interfaces were available for OC-768c/STM-256c POS (the first 40 Gbps routing interface), OC-192c/STM-64 POS/DTP, OC-48c/STM-16 POS/DTP, and 10 GE. Lower-speed interfaces were expected in future.

9. Optical Networking Session (Moderator: Victor Reijs)

9.1 Dark fibres on RENATER-4, Emilie Camisard

Emilie Camisard gave a presentation about the deployment of RENATER-4 (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/camisard-renater4.pdf). In June 2004, calls for tender were issued for equipment, lines, maintenance and network operations. Cisco were selected to supply the routers (GSR 12400) and switches (Catalyst 4500 and 6500), whilst Alcatel would supply the WDM equipment (1696 Metrospan). For the links between the metropolitan nodes (except in the Paris area), Cegetel and

France Telecom would provide leased lambdas, whilst Cegetal, neuf and Level3 would provide dark fibre.

The routing and switching equipment, as well as the leased lambda links were now being deployed. The dark fibres were due for delivery between September and December, with the WDM equipment arriving during the autumn.

The RENATER nodes would provide signal regeneration and lambda add/drop where projects required this, whereas shelters would provide signal re-amplification (except in a few places where signal regeneration was needed). An 8 lambda configuration would be used which required less regeneration points than if 16 lambdas were used.

They were interested to see how well transceivers from different manufacturers would interoperate, as well as the signal compliance between the RENATER, GÉANT2 and other regional networks. It would also provide them with experience of tunable lasers and 100 Gbps link technology.

Victor asked why the Cisco equipment was not connected together, as this would allow the dark fibre to be tested and perhaps allow the leased fibre to be relinquished. Emilie replied it was largely a matter of cost.

9.2 UKLight infrastructure update and the Optical Control Plane activity, Andrea di Donato

Andrea gave a presentation about UKLight (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/donato-uklight.pdf). UKLight was a facility to provide national and international 'bandwidth channels' to research and e-science projects. It also deployed transmission equipment to provide point-to-point links, which would eventually be separate from the JANET production network. This would allow projects to undertake potentially disruptive work without impact on the production IP service, as well as the ability to test new equipment and protocols.

UKLight operated its own transmission equipment in conjunction with SDH switches and multiplexors. There was full support for Ethernet GFP/VCAT/LCAS to ensure bandwidth optimization, service flexibility and interoperability across multi-vendor networks. =

The equipment used for the core network was the Ciena Core Director which is a high-density device with STM-64 and GE interface options. The switch matrix offers VC3 granularity (about 50 Mbps), and the control plane supports O-UNI, I-NNI, E-NNI and GMPLS.

Ciena Metro Directors were used as the edge devices, and offer STM-64, GE, 10/100 Ethernet, 801.1q VLANs and sub-155 Mbps SDH interfaces. They also have a VC3 switch matrix, although no control plane.

The Ciena CN3600 is used in the place of some Metro Directors that are not carrying production traffic, This offers similar facilities, but with the additional of digital cross-connect capabilities.

Unfortunately, this means that UKLight is essentially a Ciena network, which potentially causes problems with regional networks and end-sites. There are interoperability problems between UNI and GMPLS, and whilst UCLP could be tried, there are still some concerns regarding the maturity of such software. This means that the regional/end-site packet networks need to be bypassed with WDM links when possible, and MPLS DiffServ TE should be enabled. These issues still need to be resolved by the standard bodies and vendors.

10. IPv6 (Moderator: Stig Venaas)

Prior to the start of this session, Jordi distributed copies of the 'Distributed IPv6 Broadband' book. Further copies were available from him (in boxes of 50), provided the recipients were willing to cover the shipping costs.

Bernard asked whether there was any chance of Embedded-RP being implemented on GÉANT2. Michael replied the routers would need to be upgraded to enable this feature, but there simply was not time during the current migration.

Sven did not think it should be too problematic if the correct software versions were installed, as both the NORDUnet and Abilene networks were already running Embedded-RP.

Bernard suggested using the GN1 testbed which was still available. Michael thought this should not be a problem provided that a test plan was submitted. As RENATER, SWITCH, RedIRIS and the University of Southampton expressed interest in this, Stig agreed to draft the plan.

Action 050728-01: Stig Venaas to draft test plan for running Embedded-RP on GN1 testbed.

10.1 Auto-Transition, Jordi Palet

Jordi gave a presentation on auto-transitioning mechanisms for IPv6 (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/palet-autotransition.pdf). The aim was to ensure that any device can obtain IPv6 connectivity, even if such a device is connected to an IPv4-only network. This would also facilitate IPv6 deployment in a seamless fashion.

There are already several methods of IPv6 autoconfiguration, as well as transition mechanisms for obtaining IPv6 connectivity over IPv4 networks. However, these all

require the initial availability of native IPv6, rather than providing true plug-and-plug capabilities.

It is proposed to develop support in the IPv6 stacks of end devices, so that users can obtain IPv6 connectivity without needing to know anything about it. Native IPv6 would be preferred, but other mechanisms (e.g. TS with proto-41, TS with UDP, ISATAP, STEP, 6to4 and Teredo) could be selected if they offer better performance. Such mechanisms would also need to utilise tunnelling methods that could traverse NATs, proxies and firewalls.

Devices supporting the auto-transition mechanism would need to know where to find the tunnel end-points (TEP) that provide IPv6 connectivity, in case native IPv6 was not available. This could be facilitated through the use of specific resource records (SVR or CNAME) in the DNS, or through anycast addresses. The process could be improved by describing transition mechanism policies using PBNs, which would allow interaction with routing, Qos and security policies.

Some of the preliminary research in this area is documented in a paper available from [URL required].

10.2 IPv6 Distributed Security, Jordi Palet

Jordi gave a presentation about aspects of IPv6 security (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/palet-ipv6security.pdf). One of the reasons for the slow adoption of IPv6, was that network administrators did not always understand the security implementations. For example, IPv6 introduced an end-to-end paradigm, and many of the devices would be nomadic.

The conventional network-based security model assumes that threats are external and that everyone on the same LAN is trusted. It also assumes there are no backdoors though WLAN, ADSL and dial-up, and that hosts will not need to be directly accessed from outside. Such an approach does simplify matters as it minimises the number of points of configuration. Unfortunately though, it creates a single point of failure and does not address internal threats. It also usually means that secure end-to-end connections are not possible, and makes virtual networks (e.g. grids) difficult to operate.

By contrast, a host-based security model assumes that threats can come from anywhere, and that each host can be uniquely and securely identified. Security policies would be applied to each host, although this increases complexity with respect to managing and identifying hosts.

However, the proposed approach is to centrally define policies using a policy specification language, and have them distributed via a policy exchange protocol. This would protect against internal attacks, whilst maintaining central control and not

depending on where the host was connected. It would not facilitate the end-to-end connectivity model, but also the expected upsurge in the number of mobile hosts.

11. GRNET Advanced Network Services Tool, Vangelis Haniotakis

Vangelis gave a presentation on the network services tool that had been developed for GRNET (see http://www.terena.nl/tech/task-forces/tf-ngn/tf-ngn18/haniotakis-anstool.pdf). The aim was to develop a simple, extensible tool that could obtain administrative information about the GRNET network and store this in a database. Information could then be viewed via a web interface, and complex configuration tasks undertaken.

The topology database maintained information about member institutions and networks (AS numbers, domains etc..), physical points of presence, provider routers, interfaces, edge switches, and various network services (QoS, VPNs, etc..). This was updated by a custom perl script which examined router configurations, and created objects as necessary.

The ANStool was based on PHP4 and MySQL, and utilised a common framework for the basic functionality, to which service modules could be added. The tool allowed clients to submit requests for services and view service status, whereas network managers could review requests, and ask for a recommended router configuration. It would also shortly be possible to upload such configurations to routers. At the present time, the ANStool supported the following services: MPLS L2 VPNs, MPLS L3 VPNs, IP Premium, Managed Bandwidth, and Network Dimensioning.

A demonstration of the ANStool was available at http://edet.ucnet.uoc.gr/demo/html/. It was possible to submit a few example requests, view the outstanding requests, and see the generated configurations.

The next steps were to improve the provisioning scheme for network services; in particular with respect to QoS and MPLS VPNs. Support would also be added for L2TP services, and the tool would be upgraded to GN2 standards. In addition, the use of XML for the topology database would be investigated as it was currently difficult to add new functionality.

12. Next meetings

The next meeting had provisionally been arranged for 10-11 October 2005 in Athens, but it was pointed-out that this would clash with RIPE 51. Kevin said that he would contact the hosts GRNET to see if alternative dates could be found. [It was subsequently agreed to hold the meeting on 3-4 November 2005.]

Action 040728-02: Kevin Meynell to contact GRNET about arranging alternative dates for the next meeting.

It was also mentioned that a host was sought for the subsequent meeting, around late-January/early-February time.

13. Any other business

Further to Action 040929-03, there was a suggestion to organise a workshop on dynamic routing configuration, either at the next TF-NGN meeting or as a standalone event at the TERENA offices (subject to numbers). Michael and Kevin would discuss how best to facilitate this.

Action 050729-03: Michael Enrico and Kevin Meynell to discuss organisation of workshop on dynamic routing configuration.

Summary of Actions

ACTION 040929-03	Jean-Marc Uzé, Michael Enrico to organise a discussion on dynamic configuration, customer-empowerment at next TF-NGN meeting.	Superseded
ACTION 050113-01	Victor to provide a draft of the term definition document before the next TF-NGN meeting.	Outstanding
ACTION 050113-02	Victor to send a call for participation in optical activities on the TF-NGN mailing list.	Done
ACTION 050113-03	Marcin to provide a list of features to be tested and the test methodology in work item 9.7.	Outstanding
ACTION 050113-04	Radoslaw to prepare a state-of-the-art document on transport protocols before the next TF-NGN meeting.	Superseded
ACTION 050414-01	Radoslaw Krzywania to send details of Transport Protocols mailing list to the TF-NGN mailing list.	Superseded
ACTION 050728-01	Stig Venaas to draft test plan for running Embedded-RP on GN1 testbed.	NEW
ACTION 040728-02	Kevin Meynell to contact GRNET about arranging alternative dates for the next meeting.	NEW
ACTION 050729-03	Michael Enrico and Kevin Meynell to discuss organisation of workshop on dynamic routing configuration.	NEW