

TERENA TASK FORCE ON NEXT GENERATION NETWORKING

Report on the 17th TF-NGN meeting

14-15 April 2005

SWITCH, Zürich, Switzerland

Issue 1, Kevin Meynell

Attendees

| <u>Name</u> | <u>Organisation</u> | <u>Country</u> |
|---------------------------|---------------------------|-----------------|
| Martin Büchli | DANTE | - |
| Emilie Camisard | RENATER | France |
| Mauro Campanella | Consortium GARR | Italy |
| Valentino Cavalli | TERENA | - |
| Tim Chown | University of Southampton | United Kingdom |
| Thomas Dipasquale | Nortel | Canada |
| Jerome Durand | RENATER | France |
| Michael Enrico (Chair) | DANTE | - |
| Alexander Gall | SWITCH | Switzerland |
| Marcin Garstka | PSNC | Poland |
| Jon Kåre Hellan | UNINETT | Norway |
| Gabor Ivanszky | NIIF/HUNGARNET | Hungary |
| Avgust Jauk | ARNES | Slovenia |
| Dimitrios Kalogeras | GRNET | Greece |
| Antonin Kral | CESNET | Czech Republic |
| Otto Kreiter | DANTE | - |
| Radek Krzyania | PSNC | Poland |
| Felix Kugler | SWITCH | Switzerland |
| Olav Kvittem | Uninett | Norway |
| Yolanda Lamilla | Cisco Systems | Spain |
| Simon Leinen | SWITCH | Switzerland |
| Athanassios Liakopoulos | GRNET | Greece |
| Kevin Meynell (Secretary) | TERENA | - |
| János Mohácsi | NIIF/HUNGARNET | Hungary |
| Cristian Morariu | University of Zürich | Switzerland |
| Victor Olifer | UKERNA | United Kingdom |
| Dennis Paus | SURFnet | The Netherlands |
| Jan Radil | CESNET | Czech Republic |
| Govinda Rajan | Bell Labs, Lucent | United States |
| Victor Reijs | HEAnet | Ireland |
| Esther Robles | RedIRIS | Spain |
| Yves Schaaf | RESTENA | Luxembourg |
| Matthew Schmitz | Cisco Systems | |
| Laura Serrano | RedIRIS | Spain |
| Trond Skjesol | Uninett | Norway |
| Bernard Tuy | RENATER | France |
| Sven Ubik | CESNET | Czech Republic |

| | | |
|-----------------|------------------|----------------|
| Jean-Marc Uze | Juniper Networks | France |
| Stig Venaas | Uninett | Norway |
| Josef Vojtech | CESNET | Czech Republic |
| Chris Welti | SWITCH | Switzerland |
| Steve Williams | UKERNA | United Kingdom |
| Wilfried Woeber | ACOnet-CERT | Austria |

Apologies

| <u>Name</u> | <u>Organisation</u> | <u>Country</u> |
|---------------------|---------------------|----------------|
| Carlos Friças | FCCN | Portugal |
| Jürgen Rauschenbach | DFN | Germany |

Meeting Presentations

All presentations from the meeting are available online at:
<http://www.terena.nl/tech/task-forces/tf-ngn/presentations17.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 Update, Maarten Büchli

Maarten gave a presentation on the current status of GÉANT2 (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/buechli-geant2.pdf>). Preferred and backup suppliers had been selected for most routes, and they were currently in the final round of negotiations. Laboratory tests had also been run on the equipment of three different vendors, and a decision was expected by the end of April. In both cases, it was expected that contracts would be signed by early-May.

The proposed topology was largely be determined by the offers received, although it was hoped to have two diverse routes out of each country where possible. Other considerations were minimising concentration of capacity and avoiding unnecessarily long routes. Unfortunately, the topology could not yet be publicly disclosed, although NRENS had been consulted via a workshop held in Amsterdam.

A couple of different types of GN2 PoP had been identified - those where dark fibre and managed wavelengths were available, and those where managed wavelengths only were available. This would determine what equipment needed to be installed, and what set-ups would be required.

The most important consideration for GÉANT2 would be to support the services and access types requested. This would require Ethernet flows to be identified by VCGs internally, and a focus on SDH in the backbone. It would also be necessary to consider

protection and restoration mechanisms, network management and monitoring, and multi-domain control plane developments. Five scenarios were therefore outlined.

Simon said that asymmetric link failures were possible with OADM and asked whether one-way routing would still work. Michael thought that the link would automatically shutdown in both directions in case of failure in one direction (at least where POS interfaces were being used), but this should perhaps be investigated.

2. Introduction to MPLS VPN & VPLS work area, János Mohácsi

Janos outlined the various VPN types (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/mohacsi-vpn.pdf>). These could be classified as Virtual Leased Lines (VLLs)/Virtual Private Wire Services (VPWSs), Virtual Private Routed Networks (VPRNs), Virtual Private LAN Services (VPLSs) and Virtual Private Dial Networks (VPDNs).

VPWSs are IP or MPLS tunnels forming a point-to-point link to emulate a physical leased line or dedicated connection, and requiring a tunnelling mechanism. Examples include GRE tunnels, L2TP, MPLS and IPsec.

VPRNs emulate a multi-site routed network over IP, and consist of a mesh of tunnels between ISP routers and stub links connecting CPE to ISP routers. This arrangement simplifies configuration of CPE routers because the ISP edge routers appear to be neighbours. Unfortunately, the burden of tunnel maintenance then falls on ISPs and each VPRN only supports a single network layer protocol.

VPLSs emulate a LAN segment over IP in similar fashion to VPRNs, except that tunnels extend all the way to CPE routers and ISP edge routers provide link-layer bridging. CPE nodes can then act either as bridges or routers and are protocol transparent.

VPDNs allows remote users to connect on demand into sites through ad-hoc tunnels. However, strong binding between a particular user and a central site requires authentication and security. Examples include L2TP and PPP.

The IETF Working Group on Layer 2 VPNs had produced a couple of draft standards on the service requirements for Layer 2 Provider Provisioned VPNs (draft-ietf-l2vpn-requirements-04.txt) and a framework for Layer 2 VPNs (draft-ietf-l2vpn-framework-05.txt) which were currently being reviewed by the IESG. Several other drafts relating to VPNs had also been produced.

3. Intra- & Inter-domain VPLS: configuration and deployment experiences, Laura Serrano

Laura gave a presentation on VPLS (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/serrano-vpls.pdf>) which is way of providing Ethernet or MPLS VPNs between sites located at different geographical points. It is useful for advanced applications and distributed computing over Layer 3 networks, even between different domains.

With VPLS, the provider networks work as a single LAN, and dynamically map MAC addresses which are then sent to other provider edge devices. There are currently two draft standards for this supported by Cisco and Juniper.

RedIRIS was currently testing VPLS between three sites in Madrid, Seville and Valencia, two of which were in the RedIRIS domain and one in the GESGA domain. MP-BGP was set-up between the PE routers and MPLS/LSPs extended between them so that sites in different ASs appeared to be on the same LAN. This provides more security and flexibility as only trusted hosts are included, and the customer is able to manage their network without provider control. It is also easy to add further sites as it is only necessary to establish one IBGP session between the RR and each new PE device. However, it should be noted that the PE devices require a Tunnel Physical Services Card, and the same technology must be used at both ends of the VPLS connection.

The next stage is to install a complete Grid testbed using nodes connected by VPLS, and they are currently looking for other interested domains. They also aim to install a remote boot server to boot remote nodes using a VPLS connection, and to establish management and monitoring facilities.

4. LOBSTER Project, Sven Ubik

Sven gave a presentation on the LOBSTER project (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/ubik-lobster.pdf>). This was a nine-partner Specific Support Action under the IST programme which aims to deploy a passive monitoring architecture with support for distributed monitoring and configurable anonymisation.

The project would build on the work of the SCAMPI project (<http://www.ist-scampi.org>) by utilising the COMBO hardware and MAPI software developed by that project, and extending this to support distributed monitoring. It would also improve on the initial intrusion and denial-of-service detection work, and a number of other useful monitoring applications. It was hoped that worms and viruses in particular could be detected before all machines were affected, but this would require a distributed infrastructure of intelligent sensors. The project was therefore looking for sites interested in cooperating in this activity.

Maarten asked what interfaces were supported by the COMBO card. Sven replied the currently available add-on boards supported 4 x 1 GE copper, 2 x 10 GE copper, and 2 x 10 GE optical. A 2 x OC-48 board was also under development.

? asked what type of PC was necessary to run the COMBO cards. Sven replied that pretty much any medium specification PC would be suitable, but it was currently advisable to use an industrial chassis due to heat issues. This would be fixed with the next generation of cards.

Tim asked why sites wanted to create virtual LANs. Janos replied they wanted to run the likes of NetBIOS which wasn't routable.

Olav said that not everyone was using Juniper and asked when a common protocol would be available. Janos replied there had been discussions with Cisco and Siemens and it seemed they were working on something similar. Jean-Marc added that was only necessary for two of the PEs to have BGP-based MPLS.

5. Multicast Performance (Moderator: Steve Williams)

5.1 Multicast Weathermap, Dimitirios Kalogeras

Dimitios gave a presentation about the Multicast Weathermap application being developed by GRNET (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/kalogeras-weathermap.pdf>). The problem with multicast is that it is easy to deploy but difficult to manage, and there are few tools available with good interfaces. There was therefore a requirement for a multicast tool with a simple intuitive interface, which was extensible, and which could work with existing standard components.

The Multicast Weathermap was designed with a web-based interface to display quantitative and qualitative information about multicast traffic. It presents highly processed and filtered data generated by the standard MIBs for IP multicast routing, PIM, IGMP and various other protocols. It is able to display multicast distribution paths, overall multicast traffic, fine-grained measurement of group-specific traffic, IGMP memberships, and various per node information through plug-ins. A modular architecture is employed which includes a data collector which takes raw information from managed nodes, a data processor which extracts meaningful information, a presentation mechanism for visually displaying the information, and finally a storage mechanism.

The Multicast Weathermap is suitable for the majority of multicast networks because it is independent of the data collecting instrumentation, and separates data from presentation. The visual interface was based on code from the original weathermap and could probably be used for other projects as well.

5.2 Requirements for discovery of dynamic SSM sources, Stig Venaas

Stig outlined the requirements for SSM applications (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/venaas-ssm.pdf>). These needed to know source addresses and for multi-party applications like conferencing this could be difficult. It may not be known ahead of time who would participate, and participants may come and go during a session. Some sort of application RP functionality was therefore necessary, which should be standardised to avoid proprietary and non-compatible solutions. It may also be decoupled from specific applications, thus being able to serve a variety of tasks.

Any solution must offer discovery of dynamic SSM sources during sessions, with a comparable level of performance to current ASM architecture. It should not introduce additional requirements on the network; it must work in the SSM address space; and should be manageable and secure. It should also allow co-existence with other source discovery mechanisms, and gradual deployment must be possible.

The specific host and signalling requirements were outlined, but some input was still needed about application requirements. In particular, what was considered acceptable delay from when a new source became available until others could receive it, were there any examples of dynamic sources other than conferencing, and were there any applications with more dynamicity?

5.3 dbeacon: Distribution Beacon Implementation, Stig Venaas

Stig gave a presentation on dbeacon. This was a distributed IPv4/IPv6 monitoring tool that gathered various statistics about the state of multicast connectivity. Each beacon joins a common multicast group and sends/receives various information to/from it. Some statistics such as TTL, loss, delay and jitter are computed, but other information is sent explicitly,

Unlike the original DAST/NLANR implementation, dbeacon is written in C++ rather than Java which makes it less portable but faster. It also allows information to be dumped to an XML file for storage and/or later processing. In addition, it implements a SSM beaconing feature and a new specific protocol instead of RTP/RTCP. A large-scale beaconing mechanism is currently being studied, with the aim of coping with 1000+ beacons.

Bernard asked about the status of the IPv4-IPv6 multicast gateways. Stig replied there was still one employed at RENATER.

Dimitrios asked whether there was any standardised solution for multicast monitoring. Stig replied the IETF was working on this, but it would take at least another year to agree on a standard.

5.4 SSMping, Stig Venaas

Stig gave a presentation on ssm ping which is a tool for testing multicast connectivity (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/venaas-ssmping.pdf>). Its behaviour is bit like normal ping, although a server must run ssm pingd before client can ping it. A server responds with both unicast and multicast replies which allows the client to check that it receives SSM.

The tool supports both IPv4 and IPv6 and runs on Linux, FreeBSD and possibly other BSD variants. Solaris will also be supported shortly. The aim is to extend it to IGMP/MLD, although there is currently not much interest from the IETF in this.

More information is available at <http://www.venaas.no/multicast/ssmping/>

6. IPv6 (Moderator: Tim Chown)

6.1 SWITCH IPv6 Experiences, Alexander Gall

Alexander gave an overview of the SWITCH IPv6 network which had been running as dual-stack since May 2004 (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/gall-ipv6.pdf>). There was currently a sustained IPv6 data rate of around 50 Mbps which was mostly NNTP traffic. Onwards connections existed to TeliaSonera via an IPv6/GRE tunnel from Zürich, GBLX via an IPv6 tunnel from CERN, and GÉANT via a native connection from CERN.

The Cisco 6500 and 7200s used for the core and border routers were generally very stable although they had experienced some small glitches often related to link-local multicast. The EFT “Rockies2” software being used offered a number of interesting features including SSM, embedded RP, SSM mapping and IPv6-enabled QoS, although no IPv6-enabled SAFI yet. This beta software would soon be available as IOS 12.2(18)SXE.

The Cisco 3750s being used in their main office had proved very stable although they did not yet offer multicast or interface ACLs. These features were promised soon.

6.2 6PE over CsC – The SEEREN case, Athanassios Liakopoulos

Thanassis gave a presentation on 6PE over CsC as used in the SEEREN network (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/liakopoulos-6pe.pdf>). 6PE is a method that allows an MPLS-based ISP to offer IPv6 interconnection services to customers without upgrading the entire network to support IPv6 protocols. It is similar to the MPLS VPN model, and IPv6 functionality only needs to be enabled on the edge routers. It is perhaps most suitable where only a small number of customers require IPv6, or where an ISP wants to avoid upgrading the core network and/or run Layer 2/3 tunnels. An overview of how to set-up and configure 6PE was then provided.

Within SEEREN, much of the connectivity is provided with CsC MPLS VPNs. Carrier-supporting-Carrier is designed for ISPs that are VPN customers of larger MPLS-based ISPs. The 6PE functionality is then installed on the customer edge (CsC-CE) rather than the provider edge (CsC-PE), and 6PE peers therefore belong to a different administrative domain. Operational experiences suggest that the technique is similar to MPLS VPNs in terms of technical implementation and complexity, although other solutions should be considered for large-scale IPv6 services.

Laura asked why Layer 3 VPNs were needed. Thanassis replied that you could choose to use Layer 2 VPNs, but you would then lose certain functionality such as QoS.

6.3 IPv6 Intrusion Detection, Tim Chown

Tim gave a presentation on the requirements of intrusion detection systems for use with IPv6 (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/chown-ids.pdf>). These were already widely available for IPv4, but IPv6 intrusion detection systems were still in their infancy and there were a number of issues to consider. These included new header formats, new header options, and the widespread use of tunnelling which could give rise to new types of attack.

Tim asked whether any NRENs were interested in investigating these issues. He had written an IPv6 packet generator and was modifying Snort as the first step.

Matthew suggested Tim contact Steve Wallace at Cisco who was doing similar work.

6.4 MRD6 – an IPv6 Multicast Router for Linux, Stig Venaas

Stig gave a presentation on the MRD6 project, which was developing an open-source multicast router for Linux (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/venaas-mrd6.zip>). This started in August 2004 at IT-Aveiro as an internal development project, but was used in the IST-funded Daidalos project.

MRD6 is strictly a multicast router and only supports IPv6 in order to push the new technology. It currently features full MLDv1 and v2 with proxy support; PIM-(S)SM with BSR, Static-RP and Embedded RP; MRIB capability; and M-BGP support. It can support both native and tunnel interfaces, and so can be used as a tunnel broker. It is very modular so new interface handlers can easily be added for local management, new protocols or even custom access and accounting modules.

MRD6 is currently being used in the M6Bone and is used to stream a variety of content. Additional features such as Anycast RP and proper kernel multicast forwarding are being developed, and full documentation should be available shortly. In the longer-term it is hoped to develop a more robust version of M-BGP.

More information is available at <http://artemis.av.it.pt/~mrd6/>

Tim asked who was running the M6Bone nowadays. Bernard replied that RENATER was.

6.5 GÉANT IPv6 Multicast Testing, Stig Venaas

Stig gave a short presentation about the proposed IPv6 multicast tests on GÉANT (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/venaas-multicasting.pdf>). Around eight NRENs were now able to get IPv6 multicast through GÉANT, some natively using production links, and some via tunnels to the nearest PoP. This needed to be tested to ensure that it works as expected and required cooperation between the different NRENs. It was proposed to use Dbeacon, Java Beacon, ssm ping to verify SSM connectivity, and GÉANT Looking Glass amongst other utilities. However, the main problem was providing multicast content as not much was available. They were therefore looking for additional NRENs and end-sites to get involved.

6.6 Migration Broker service deployment in RENATER, Bernard Tuy

Bernard gave a presentation on the migration broker service being deployed in RENATER (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/tuy-migrationbroker.pdf>). This is a server that provides IPv6 connectivity to anyone who is allowed to request it, by creating dynamic tunnels based on TSP (Tunnel Setup Protocol). Users subscribe to the service via a web interface.

The server is a rackable PC running NetBSD and the HexOS software written by Hexago. End-users need to install a TSP client that is available for Windows, MacOS, Linux and FreeBSD, and which establishes a tunnel to the migration broker. This allows IPv6 connectivity to be provided to sites that are not connected to RENATER's IPv6 service, mobile users, and for temporary events. Other available features are DNS updates, NAT traversal, Digest-MD5 authentication and SNMP monitoring. They are also hoping to provide SSL and multicast support in the near future.

6.7 6DISS Project, Tim Chown

Tim gave a presentation about the 6DISS project (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/chown-6diss.pdf>). This was a 30-month IST project to disseminate information about IPv6 in the Balkan, Mediterranean, Caribbean, Central Asian, Sub-Saharan Africa, Southern African and South/Central American regions. It would also collaborate with selected organisations in India and China.

The project was organised around a series of workshops (one per region) that would draw upon the documentation and training material produced by 6NET and other projects. The project would also feature hands-on IPv6 training for future trainers in Brussels, as well as ‘Tiger Teams’ that would provide advice on IPv6 issues. The aim of the project was not only to raise awareness of IPv6, but to build contacts in other regions for future collaboration.

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

Marcin reported on the evaluation of the Cisco 4948 10 GE switch that had been undertaken by PSNC (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/garstka-evaluation.pdf>). This featured two 10 GE optical ports and forty-eight 10/100/1000 copper ports. It had been tested as Layer 2 switch, and the forwarding, autosensing and autonegotiation on the 10/100/1000 ports worked without problems, although they were only tested with other Cisco devices. Traffic filtering based on MAC addresses was less successful as the filters did not work if an Ethernet frame contained an IP packet. However, traffic filtering based on IP address did work.

Bandwidth limiting on physical port worked okay, but not with VLAN ports. It was possible to limit traffic per VLAN, but then traffic will be limited on all ports in the VLAN. MAC address limiting also working okay on the 10/100/1000 ports, but not on the 10 GE ports. Equally, frame counters worked on all physical ports, but not on the VLAN ports. The spanning tree and Layer 3 routing tests were successful, with the exception that it was not possible to maintain a full BGP table.

In general, the switch worked well although there were still some minor gaps in functionality.

Wilfried thought that bandwidth limiting should apply across all ports in a VLAN otherwise traffic allocations could be circumvented. Marcin replied that per-port limiting was necessary as some of their customers had several links.

Dimitrios asked whether filtering had been tested on the trunk ports. Marcin replied the limits applied to the whole port, not just for the VLAN.

Bernard asked whether any of the management and/or monitoring functions had been evaluated. Marcin replied that had not been done through lack of time.

8. Transport Protocols Work Area, Radoslaw Krzywania

Radoslaw reported that an initial range of enhancements had been selected and the PERT knowledge base had been suggested as the place for state-of-the-art output (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/krzywania-transport.pdf>).

The plan was to finish the specification of applications in the first half of May, then define the test procedures for evaluating the protocols.

In addition, a new IST project proposal (FASTNESS) would be submitted in late-September. This aimed to collect and compare existing transport protocol solutions, and develop innovation solutions based on application and user requirements.

Michael asked for volunteers to help with this activity. Radoslaw said he would circulate the details of the mailing list.

Action 050414-1: Radoslaw Krzywania to send details of Transport Protocols mailing list to the TF-NGN mailing list.

9. IP Routing Work Area, Mauro Campanella

Mauro outlined the proposed list of topics on which to work (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/campanella-routing.pdf>). These included investigating the status of routing on NRENs; peering load balancing; routing on dual-stack routers; traditional IGP behaviour; extensions to IGP protocols; terabit routing; QoS routing; path computation; testing of virtual routers; dynamic circuit set-up and tear-down; dynamic reconfiguration of routing and services; and VPN communication. In addition, it was proposed to develop a worldwide multi-layer weathermap.

This work area would collaborate with the GN2-JRAs, CAIDA, the NSF-funded DRAGON project, the MUPBED IST project, the LASER project, and the IETF working groups on rtgwg and ccamp. It would also coordinate with the Regional Internet Registries.

Janos asked whether the work area included Layer 2 VPN issues. Mauro thought that it should stick to Layer 3 issues, but perhaps this could be discussed further.

10. GN2 Research and Service Activity Updates

10.1 JRA1 Update, Nicolas Simar

Nicolas gave an overview of the JRA1 activities (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/simar-jra1.pdf>). These aim to develop monitoring tools that can retrieve information from several networks using a pre-defined message format. Another aim is to develop a visualisation tool to showcase the concept. There will be close collaboration with the Internet2 piPEs activities as well as ESnet.

The requirements and metrics (OWD, OWPL, IPDV, RTT, available bandwidth and interface errors) had already been defined, and an initial set of tools (DFN IPPM, ipperf/BWCTL, CNM and nemo) had been chosen in conjunction with SA3. However, they still need to choose a NetFlow and packet capture tool.

The next stage was to develop a prototype by July, as the overall system architecture needed to be defined by August.

10.2 SA3 Update, Simon Leinen

Simon reported on the SA3 activities (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/leinen-pace.pdf>). These are comprised of a network performance monitoring infrastructure, a performance and enhancement response team (PERT) and QoS provisioning. The SA1 and JRA1 activities agreed to jointly select a set of preferred tools, which SA3 would then deploy operationally and JRA1 would use for development and testing. The selected tools were DFN-IPPM for delay measurements, iperf/bwctl with a SUNET-developed front-end for achievable bandwidth measurements, whilst the SCAMPI system was being evaluated for passive monitoring. Four DFN-IPPM devices and 2 bwctl servers were already deployed in GÉANT, but additional equipment would not be ordered until the GN2 PoPs were confirmed.

The pilot phase of PERT had concluded and the ‘production’ phase started on 1 April. There were number of currently active issues, including throughout problems with some VLBI sites, and asymmetric delays between ESTEC and ESRIN, but the work could be followed via the PERT diary at <http://diary.pert.switch.ch/>.

A number of pilot tools including the PERT tracker, knowledge base and diary had already been installed by DANTE, whilst the ticket system and website developed by PSNC would be available shortly. The knowledge base at <http://kb.pert.switch.ch/> was now starting to become useful, although more input would be useful.

The QoS provisioning activities were building on the work of the SEQUIN project and the Premium IP allocation policy had just been completed. The policy aimed to distinguish between early phases when service guarantees are limited, and later phases when enhancements meant more accurate SLAs could be offered. The next phase was to develop the provisioning system to ensure that the total amount of Premium IP across any given network link did not exceed a pre-defined figure (typically 30%). The software would be offered to all networks that participate in the Premium IP service, but other systems may be used provided they respected the defined interfaces. Development of the provisioning system was slightly behind schedule, and was now expected in early-September.

Olav asked whether static or signalled QoS would be provided. Simon replied that set-up was not sub-millisecond, but it should be in the order of hours.

10.3 JRA3 Update, Maarten Büchli

Maarten reported on the JRA3 activity which is focused on the research and development of end-to-end, connection-oriented bandwidth-on-demand (BoD) services (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/buechli-jra3.pdf>).

Work Item 2 had already commenced and the BoD User and Application Survey had just been finalised. The initial review of BoD-related technologies was also near to completion. A number of users and applications had already been identified including the European VLBI network, DEISA, the EGEE and SEEGRID projects, the MUPBED testbed, the GN2 testbed and the LOFAR astronomical array. A questionnaire had been sent to these activities in order to obtain feedback on their requirements which included connection to high-end computing facilities at speeds of between 10 Mbps and 10 Gbps, and standardised interfaces for resource reservation and service monitoring.

With respect to Work Item 3, the framework and architecture discussions were still ongoing and a face-to-face meeting had been held on the 13th of April. This included defining the scope of the work, producing a glossary of terminology, identification of the necessary modules and services, and the development of a functional specification.

11. Optical Networking Work Area, Victor Reijs

Victor outlined the general principles for the work area (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/reijs-optical.pdf>). These included end-to-end tests for long-haul links, development of nomenclature, liaising with providers and manufacturers, information exchange between partners, and testing of equipment. The medium-term aim was to further develop the optical exchange concept, utilise pure optical switching and routing, and develop open WDM standards.

11.1 CzechLight & CzechLight Amplifiers, Josef Vojtech

Josef gave an overview of the CzechLight network and amplifiers (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/vojtech-czechlight.pdf>). This currently had connections to GLIF, the GN2 testbed and to various institutions within the Czech Republic.

CzechLight decided to build its own amplifiers from commercially available elements as there was a lack of suitable equipment for NRENs. Commercial optical amplifiers needed modifications which were expensive and lacked the possibility of further development. Three types of EDFA module were currently being used, including a 2iv1 module suitable for OSA inline with CD compensation, an inline module, and a hi-power booster.

The amplifiers are housed in an industrial PC with redundant power supplies and flash disk to reduce vibration. The management system is based on Linux and offers access via

SSH and SNMP. Critical warnings are sent via e-mail, with a Net-SNMP package under development.

The next stage is to develop semiconductor amplifiers with RZ/NRZ conversion; raman fibre amplifiers for ultra-long spans, and true broadband (100 nm) amplifiers for MANs.

Av gust asked which lasers could be used in the amplifiers. Josef replied they used Class 3B, although they could go up to Class 4 if automatic shutdown protection was used.

11.2 LambdaTunnel: a pilot optical service, Felix Kugler

Felix gave a presentation on the pilot lambda service being established by SWITCH (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/kugler-lambdatunnel.pdf>). Lambda tunnels are a fraction of the CWDM spectrum on a lit fibre that can be assigned to customers. These customers can then set-up and run their own communication links using equipment they are familiar with. Lambda tunnels can be added or dropped anywhere along existing fibre links and are not restricted to existing PoPs.

Since 2002, SWITCH had employed thirteen bi-directional links using POCs, and several regional networks had adopted the technology as well. SWITCH uses a standard colour set for bi-directional GE links which is 1530/1550 nm and within the EDFA range, thus allowing amplified links up to 200 km. When adding lambda tunnels on a fibre, it is necessary to use different colour pairs and calculate with 4dB loss per link added. It is also important to carefully choose the colours if links are long.

Lambdas with the lowest attenuation are generally preferred, but this is dependent on the fibre type which is not always known. Based on various experiences, it was determined that older G.652A fibres perform better in the upper CWDM channel range, with 1590 and 1610 nm being the next best choice. Additional channels can be added by using lambdas below 1530 nm, but attenuation might increase depending on the fibres used.

SWITCH planned to start a two-year pilot project to get practical experience of this type of infrastructure sharing, and there was already a lambda tunnel being trialled between Lausanne and Brig. The trunk fibre operator was responsible for attenuation, chromatic and polarisation dispersion within defined specifications, and providing working OADMs as agreed. Customers on the other hand, were responsible for the design of their link, providing the PoP equipment, and proper operation of their connection. One problem though is that customers have limited flexibility to debug fibre paths, and that fault localisation needs to be a joint effort with the fibre operator.

Victor asked whether optical power was a problem. Felix replied the filters were very effective, although they were just about on the distance limit of what was possible.

Dennis asked why SWITCH undertook their own measurements. Felix replied it was because they had a mixture of dark and supplier-provided fibres and needed to know exactly where the problems lay.

11.3 Introduction to Intelligent Control Plane Architectures Work Area, Otto Kreiter

Otto outlined the proposed topics for this work area (<http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/kreiter-controlplane.pdf>).

MPLS DiffServ-TE makes MPLS-TE aware of classes of service and the fault-tolerance properties of MPLS. This still requires admission and policing control, but it can limit the proportion of traffic from a particular class on a link, maintain relative proportions of traffic on a link, and guarantee bandwidth services. It can be used to develop QoS services for end customers based on signalling rather than provisioning. Possible test topics were alternative solutions for premium services, and interoperability between different vendor implementations.

Dynamic provisioning of intra-domain GMPLS LSPs aims to reduce the number of switching layers, utilise label switching and re-use the MPLS-TE protocol mechanism. It was proposed to start with single-layer tests to trial basic configurations, then to move onto multi-layer tests using more advanced functions (e.g. FA, link bundling, loose paths, LSP re-grooming). This would allow GMPLS management to be compared to legacy management.

The work area would also look at recovery techniques for GMPLS, focusing on data and transport plane failures, protection and restoration at the TDM, LSC and FSC layers, and mesh and ring-like topologies. Proposed test topics included end-to-end, LSP segment and span recovery, and consideration of robustness, resource optimisation and multi-layer recovery.

With respect to the inter-domain aspect of GMPLS, three distinct options had been identified for signalling TE LSP across multiple domains: LSP stitching, contiguous LSP, and LSP nesting. LSP stitching was deployed in GÉANT for MPLS, but more experience was required for other layers. Contiguous LSP was also deployed to some extent in GÉANT, but more control on transit domains was required. LSP nesting had not been tested at all yet.

GMPLS was a superset of MPLS, although some aspects were different. Unfortunately, migration techniques had not yet been defined by the IETF, so this was something to consider in future.

The tests should start as soon as possible as two of the topics were also proposed as possible tests for JRA4. However, the involvement of vendors was needed in order to get support for the necessary equipment.

12. Dissemination of Flow-Specification Rules, Jean-Marc Uze

Jean-Marc gave a presentation on flow-specification rules (see <http://www.terena.nl/tech/task-forces/tf-ngn/presentations/tf-ngn17/uze-flowspec.pdf>). At the present time, filtering engines are configured with static entries in contrast to routing which is usually dynamic. There was therefore a need to coordinate flow filtering within and between domains, or more specifically to extend routing information with flow specification. This would be particularly useful with respect to distributed denial-of-service attacks, as it would allow the network to quickly identify and block them.

One solution is to use BGP to propagate filtering rules independently of unicast routing information, although validated against it. Filtering actions could be a combination of accept, discard, rate-limit sample and re-direct etc., and could be controlled by unicast routing advertisements. BGP offers the advantage of being an incremental addition to already deployed mechanisms, and takes care of the distribution complexity.

There was currently an IETF draft (draft-marques-idr-flow-spec-02) outlining this technique, and a mailing list that could be subscribed to at <http://www.cqr.org/mailman/listinfo/flow-spec/>.

13. Next meetings

The next TF-NGN meetings will be held as follows:

- 28-29 July 2005 in Paris, France (prior to IETF 63)
- 20-21 October 2005 in Athens, Greece

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. | Outstanding |
| 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. | 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 050113-04 | Radoslaw to prepare a state-of-the-art document on transport protocols before the next TF-NGN meeting. | Outstanding |
| ACTION 050113-05 | Jean-Marc to send Valentino slides on JUNOS IPv6 support to be made available on the TF-NGN website. | Completed |

| | | |
|------------------|---|-----------|
| ACTION 050113-06 | Yolanda to find out schedules of MIB implementations for IPv6. | Completed |
| ACTION 050113-07 | Tim/Stig to solicit ideas and lead discussion about a test plan concerning the GÉANT multicast IPv6 service, to be ready by end of March. | Completed |
| ACTION 050113-08 | Michael to discuss with Janos about the possibly to extend the scope of the VPLS work item. | Completed |
| ACTION 050414-01 | Radoslaw Krzywania to send details of Transport Protocols mailing list to the TF-NGN mailing list. | NEW |