

Minutes of the 21st TF-NGN meeting, Ljubljana, 4-5 July 2006

issue 1.3, Catalin Meirosu (TERENA)

Michael Enrico - GÉANT2 Update

Michael started his talk with a brief historical view on the evolution of the bandwidth growth for trans-European research and education networks. The bandwidth available on the network, calculated based on that of the fastest connection in the network, increased exponentially. The current network, GÉANT2, is based on dark fibre and even though the bandwidth of a single connection remains unchanged from the previous GÉANT network, multiple such connections could be provisioned over a single fibre. The network has 25 POPs, serves 31 NRENs and comprises 11600km of dark fibre. Wavelengths or lower-speed circuits were provisioned in places where dark fibre was unavailable or not efficient cost wise. The network is built using Juniper routers (T640, M160 and M40) and Alcatel 1626LM and 1678MCC boxes. The network has a hybrid architecture and is thus able to offer both IP connectivity (using three classes of traffic plus multicast) and point-to-point connections at lower layers. The point-to-point services support emulated circuits (CCC implementation) or native connectivity (GE, 10 GE) either through encapsulation on GFP or directly over a managed wavelength. The total transatlantic connectivity for research and education network in June amounted to 11 circuits, each with a capacity of 10 Gbit/s. Four of these circuits were operated as being part of the GÉANT2 network. The capabilities of the GÉANT2 network are being updated through research and service activities in the framework of the GN2 project: improved performance monitoring in JRA1 and SA3, automated PIP provisioning in SA3, network security best practices in JRA2.

Victor Reijs – UCLP 1.5 for HEAnet

Victor presented the results of the work carried out for adapting the UCLP v1.5 environment to the HEAnet network. UCLP (User Controlled Lightpath Provisioning) is a tool developed within the CANARIE's Advanced Research program to address the problem of configuring and managing network partitions at the physical / data link layer. The UCLP-HEAnet variant, implemented by i2cat in Spain, supports Cisco 760x and Catalyst 3750 devices (via a Cisco E-DI server) and a MPLS edge pseudo-device, in addition to the devices supported by the standard v1.5 implementation. The modular implementation allows for additional devices to be added by implementing a new Switch Communication Services software component. The goal of this implementation was to control and support the Ethernet Relay Service and Ethernet Wire Service that are offered to HEAnet's customers. A Java-based GUI is provided for establishing, managing and querying these connections. i2cat built a testbed and technology demonstrator at their premises, using 7 switches equipped with 1G and 10G connections. HEAnet will evaluate the UCLP-HEAnet implementation in view of a potential deployment in their network.

Jean-Marc observed that the architecture looks to be centralised hence there is no peering between the federations. Victor agreed. David asked about the long term vision and whether this tool would remain confined to research networks. Victor answered that UCLP is one solution for providing a particular type of resources to people in a standardised way.

Isabel Gandia - The Scientific Ring: the Catalan RREN Network Infrastructure

Isabel presented the Centre de Supercomputacio de Catalunya (CESCA), a public consortium formed by a partnership between the Catalan government, universities and funding bodies. It manages the Anella Cientifica, a network that connects directly 56 institutions (including research departments of private companies). The current version of the network is based on Gigabit Ethernet deployed over WDM links and has 69 points of access. The core of the network has a full mesh topology while the access is organised in rings. Some of the connections use dark fibre in the metropolitan area. The Anella Cientifica has thus a fully redundant topology. The external connectivity is assured by RedIRIS and also by commercial ISPs. This is because of the fact that some of the connected institutions run business traffic and the policy of RedIRIS would not allow this traffic to pass. Example of use cases included the Universidad Autonoma de Barcelona, for which special links were deployed to support its collaboration in CERN's LHC project and in the Open Opera initiative. CESCA is also home of the Catalan Internet Exchange (CATNIX). In the near future, CESCA plans to connect 3000 schools in the region using IPv6 technology.

Alexander Gall: GENI – A new breed of testbed for network innovation

Alexander talked about GENI, the Global Environment for Network Innovations. This is an initiative supported by the National Science Foundation in the US. A Project Execution Plan document, issued in January 2006 by the participants in the project, identified fundamental deficiencies in the current Internet and proposed to address these problems from a clean-slate approach. Alexander considered that GENI intends to take the idea of the PlanetLab testbed to a next level. This would be done by providing a high-performance network substrate at the physical layer, managed by a framework that would allow incremental and clean-slate experiments to take place simultaneously within perfectly isolated slices of the network. Also, the management framework could allow separated slices to interact when the experiments require it. It is envisaged that the National Lambda Rail assures the connectivity for GENI within the US. A typical GENI PoP would interconnect sensor networks, wireless networks and LANs to a high-performance backbone through a programmable router. Cisco started discussions with the NREN community on a project that would build a GENI equivalent in Europe (codenamed ENGINE – Experimental Next Generation Internet in Europe). This proposal was in a very early stage of development.

Marcin Gartska - Juniper Networks TX Matrix testing (including STM-256 card and DWDM PIC)

Marcin presented the results of the tests that were run on a Juniper TX Matrix router on 20-21 June at the Juniper lab in Amsterdam in the framework of activity 9.7 of the ToR. Several members of TF-NGN planned and participated in the trials that were actually carried on by support people from Juniper. The comprehensive series of tests included switching between wavelengths on the DWDM PIC (a 10 GE LAN PHY capable of using one of 45 wavelength from the standard ITU grid), measuring the performance of the forwarding plane of the STM-256 card for IPv4 and IPv6 packets, reliability tests and convergence of routing protocols (through provoking failures by removing the engines from the chassis). The convergence time for routing protocols was also measured using a maximum number of 500000 prefixes in the routing table for BGP and iBGP and 25000 prefixes for ISIS. The use of Bidirectional Forwarding Detection was found to improve the ISIS convergence of the test scenario by roughly two orders of magnitude (from 22s to 0.3s). The QoS features were evaluated by configuring three classes of service (as currently deployed in the GÉANT network). The team concluded that the tests proved that the TX Matrix is really a single router in multiple chassis.

Stig Venaas – IPv6 session

Stig started his talk by announcing that 6Bone was shutdown on the 6th of June, and therefore IPv6 addresses in the 6bone range should not longer be used or advertised. Embedded RP is now deployed in GÉANT, and Stig asked everybody to use it whenever possible. Administrative scopes were defined in order to limit IPv4 multicast traffic, and Stig called for defining similar scopes for IPv6. He gave the example of the BBC that uses a country-wide scope for their IPv6 multicast. Victor thought that the scope in IPv6 is limited and perhaps the associated field (4 bits) should have been larger. Stig announced the availability of an open source DoCoMo implementation for the Secure Neighbour Discovery protocol, a new version of ssm ping and a looking glass for ssm ping. He continued by showing slides made by Carlos Friacas of FCCN, debating whether IPv6 still is a topic for NGN. Carlos argued that because of the scarce deployments (mainly in R&D networks, even though certain commercial offerings exist) and several open issues being addressed by workgroups in the IETF, IPv6 should still be considered as a topic for NGN. However, Carlos also argued that successful deployment in research networks and the end of lifetime for 6Bone marked a turning point with respect to IPv6 being a topic for NGN. Victor considered IPv6 services a regular production services today. Michael suggested that Stig takes the debate on the tf-ngn mailing list.

David Boyle - Some challenges for Optical Networking

David started his presentation by highlighting the challenges seen by 40 Gbps transmission system deployments. The operators see a clear need for carrying any service

(Ethernet, SONET, SAN) at any speed (up to 40 Gbps now) at line rate. Systems operating at speeds of 40 Gbps over the installed base of optical fibres have to face challenges related to the chromatic dispersion, polarisation mode dispersion, optical signal to noise ratio and the width of the available spectrum. Single stream transmissions at 100 Gbps are being researched, but unlikely to be available commercially in the near future. The major shift in traffic type from TDM to IP will make Ethernet a major option for transport network interfaces. Ethernet services will be offered by the transport network in a peer-to-peer transparent optical layer or at wavelength and sub-wavelength level. This means that no single technology will dominate (CWDM, DWDM and ROADM will all be used. Optical planning and management systems need to deliver end-to-end circuits over various network topologies. Meriton became a corporate member of Internet2 and is involved in the HOPI project. Currently, 8 wavelengths carrying 10 GE LAN PHY were deployed on Meriton equipment.

Felix Strohmeier: The MOME Meta-Database for Monitoring and Measurement Tools and Traces

Felix started with a brief introduction of the MOME project, funded by the EU FP6 IST project in the “Broadband for All” strategic research direction. MOME formally ended in March 2006 and provided a meta-database to assist researchers in finding traffic traces and network measurement tools that were most appropriate for their research. A centralised directory was developed to integrate documentation on the measurement environment and scenario, statistic information about the content of the traffic traces and links to local repositories that actually store the data. In this way, the researchers that collected the data in the first place would retain total control over their data. The database can accommodate and classify accordingly packet-level traces, flow traces, routing data, HTTP traces, etc. A module that allows to automatically analyse packet traces (in libpcap and DAG formats) was developed. The results of the analysis may be presented as general statistics or as bitrate plots. The MOME Tools database classified the large majority of tools (123 for now) existing in the network measurement and monitoring area. Even though the project ended in March 2006, the website is still maintained by TERENA and the databases are actively supported by Salzburg Research, Fraunhofer FOKUS and the Warsaw University of Technology.

Victor asked whether the project collaborates with JRA1. Felix answered that Nicolas Simar was contacted but there was no follow-up. Victor and Michael suggested to Felix to renew the contacts with JRA1. Catalin thought that an export module in a ns-2 or nistnet usable format would be useful for researchers. Felix agreed that this might be an improvement to be considered in the framework of a new project.

Victor Reijts – BLUEnet (L2 provisioning) Tool

Victor described BLUEnet, a tool that was developed in cooperation by HEAnet and GRnet to automate the provisioning of point-to-point L2 circuits. The manual process for

provisioning such a link is time consuming and prone to errors. BLUEnet is based on the ANS_tool developed by GRnet and modified to provision MPLS L2 circuits. A new workflow and front end interface were also developed for the HEAnet version. BLUEnet performs auto-discovery and automatically maps new devices added to the network. The configuration and monitoring setups are ensured to be consistent. The system provided only one reference point for looking up circuit information. HEAnet plans to use BLUEnet as their operational tool for provisioning production point-to-point L2 circuits. The envisaged use of UCLP is more of a research tool, also in view of the concept of resource virtualisation that is represented by UCLP. HEAnet plans to integrate both BLUEnet and UCLP v1.5 in the Bandwidth on Demand activity in GN2 (JRA3).

Av gust asked how difficult would be for someone to replicate the BLUEnet installation and what was the licence for the code. Victor answered that HEAnet developed good documentation in English and he would be ready to give away the source code for use in an NREN environment. However, he would need to have the agreement of GRnet as part of the code was developed by them.