Global Lambda Integrated Facility Technical Working Group
17 – 18 September 2007
Prague, Czech Republic

The session was chaired by Erik-Jan Bos (SURFnet) and René Hatem (CANARIE). The scribe was Catalin Meirosu (TERENA).

## 1. GLIF Introduction and review of the action items

Erik-Jan Bos presented the working group chairs and the secretary to the audience. The Tech working group has an email list at (tech_at_glif.is). The website of the working group, located at http://www.glif.is/tech, contains more details on the activity.
The goals of the working group are:

- to design and implement a global hybrid networking infrastructure
- identify equipment being used in Research and Education Networks, requirements for new connections
- identify functions and services to be provided.

## 2. Resource updates GOLE

A recurring and much appreciated item on the agenda of the Tech WG is the GOLE resource updates, during which the resource operators present the developments and advances within their domain.

*CzechLight*
Jan ?? mentioned that CzechLight is located in Central Europe and currently had two international 10G lambdas, i.e. to NetherLight (Amsterdam) and to StarLight (Chicago). The list of resources was the same as in previous GLIF meeting in Minnesota. Now they can offer multiple 10 Gbps connections to Poland, Slovakia and Austria.

*HKLight*
Jun Li ?? announced that last year the Hong Kong to Seattle link was upgraded to 2.5 Gbps. The GOLE had the following circuits installed: 10Gbps to KISTI, 2.5Gbps to CSTNET and GLORIAD, multiple 1 Gbps circuits. For the future, they were looking onto the possibility to add more connectivity to South-East Asia (India and Pakistan in particular). An OME6500 platform was installed in 2007.

The BJlight (Beijing light) activities were based on the CNGI (China Next Generation Internet) project. Currently it consisted of a 2.5 Gbps connection to HKlight, expected to be deployed within 3 months. Afterwards, the circuit may be upgraded depending on the applications using it. The International Cosmic Ray Pre-Processing Centre was cited as an example of application to be using this connection.

*KRLight*

Dongkyun Kim announced that new multi-Gigabit connections to CERNET and University of Hong-Kong will be installed from September 2007. He presented the network diagram for the HD video experiments performed between Chicago and Gwangiu in September 2007. Dongkyun announced that NDL descriptions of the KRLight circuits will be made available on the KRLight webpage soon. In the near future (September – October 2007) an experimental HD video transmission will take place between the Yonsei hospital in Seoul and the university hospital in Trondheim. Once the schedule of the event will be finalised, support from the GOLEs will be needed for the provisioning of the lightpath.

*MANLAN*
Christian Todorov mentioned the ongoing discussions with NLR to add another 10 Gbps connection. These are yet to be finalised. He was still working on some of the testing and support related to interoperability issues for the Ciena to Alcatel equipment connections. He expected that all the outstanding issues should be solved shortly. CANARIE reconfigured their network in New York to retire their ONS 15454. Now the MANLAN ONS 15454 terminates CANARIE provided LightPaths for 3 active connections – two for TWAREN, one for GLORIAD.

Dan Nae asked about the ownership of the HDXc. Christian answered that the device was owned and operated by Internet2. The optical equipment in the MANLAN diagram on the GOLE wiki was separated by discontinued lines on the graph just to draw attention to the network layer functionality separation. The NOC of MANLAN was operated by Indiana University. The reason for not having a Ciena box was that the MANLAN existed prior to the current Internet2 network, and at the time a number of exchange points were using Nortel equipment (for example: NetherLight). The way the GOLE was currently setup allows for different capabilities. They could do sub-lambda switching on the OME6500 and lambda-level switching on the HDXc. Recently the OME was upgraded with a 10 Gbps interface allowing for a direct 10GE connection.

*MoscowLight*
Name of Russian participant ??? announced the installation of a new POP in Moscow in the telephone exchange station. MoscowLight has thus two sites in Moscow. A new 10 Gbps link from Moscow to Amsterdam via Stockholm was being installed.

*NetherLight*
Erik-Jan Bos presented the NetherLight update. The current footprint in Amsterdam had an HDXc, an OME6500 and a new ERS8600 that supported conversion from LAN PHY to WAN PHY. Active equipment (an OME6500) was also installed at CERN.

In terms of new connections, a 10 Gbps lambda was installed to Academia Sinica Grid Centre (ASGC) Taiwan, as well as five 1 GE lightpaths for GEANT2 for the Phosphorus and EXPReS (to Torun, Medicina, 2x Jodrell Bank) projects.

Three Cross Border Fibre (CBF) projects have been completed or were expected to be completed in the near future. The CBF to Munster was in operation, and carrying traffic

for the LHC OPN. The CBF to Hamburg was awaiting final patches with NORDUnet. The CBF to Aachen was expected to be available in October 2007. A new 10Gbps connection to CERN will be installed on a diverse route for fibre resiliency. The RIPN connection to MoscowLight is to be upgraded from 2.5 Gbps to10 Gbps.

A few connections were decommissioned since the last GLIF Working Groups meeting. The lambda to UKlight was replaced by GEs over GÉANT+. Also, the IEAAF 10Gbps lightpath to MANLAN is expected to be shutdown soon.

*NOX*
Lars Fischer announced NOX as the new name for the NorthernLight GOLE. Following an internal re-branding, the NorthernLight name will be kept for the network while the GOLE will be known as NOX. The distributed nature of the GOLE was not affected by the re-branding: POPs will be maintained in Oslo, Stockholm and Helsinki. New focus areas for work to be approached in the near future consist of advanced control plane issues, provisioning, reservation, as well as local applications and outreach.

Connections status: the Moscow link at OC192 speeds will be up late September. All the required equipment has been installed, but the system in Stockholm had yet to be commissioned. The plans for the future include a new OC48 connection to Reykjavik.

*Pacific Wave*
Dave McGaugh said that the connectivity map between PNW and CANARIE was changed to reflect the possibility of having multiple 10 Gbps connections possible onto the CANARIE network. More connectivity options were now available out of Sunnyvale. Dave announced planned connectivity for the Internet2 *Dynamic Circuit Network (DCN)* into the Los Angeles ONS. The NLR FrameNet connectivity will be provisioned through CENIC. Currently there was no connectivity between the 7600 and the HDXC. One of the challenges they were facing was being able to provide greater than 1Gbps circuits over the 10Gbps lightpaths for uncompressed video applications.

*StarLight*
Linda Winkler stated that the two changes in the StarLight GOLE map reflected the changes in the CANARIE ROADM network, and the CzechLight connection that was added this year.

*T-LEX*
Akira Kato said that the configuration presented on the GOLE wiki is up to date. However, several new circuits were in the planning phase. One such connection would allow attaching Academia Sinica Grid Centre (ASGC) in parallel to the Academia Sinica Computing Centre (ASCC). Also, Layer2 services were enabled with Sinet, while the discussions on whether giving access to Layer1 connectivity were still ongoing.

*US LHCnet*

Dan Nae announced the intention of a new GOLE to participate in the GLIF. This would be represented by the CERN Tier0 to US Tier1 HEP network. Currently, they operate three transatlantic links, over 3 different submarine cables, as well as continental links. Ciena CoreDirector devices were recently deployed in all end points. The connection to Amsterdam was installed during summer 2007. They will replace the Force10 connection to HDXc with a Ciena. A process of re-tendering for the long-haul infrastructure was under way. He would expect to obtain an additional transatlantic 10 Gbps connection and the associated transcontinental links by the end of this year. Some of the transatlantic connections will be kept on the Force10 switch, but the main links were now connected to CoreDirectors. As the major part of the infrastructure was already in place, Dan concluded that they are ready to be more involved in GLIF.

*AARNet*
Alex Reid noted that Australia does not have a GOLE but this might perhaps change in case lightpath connectivity towards New Zealand would run through Australia. Lambdas were available for researchers. Therefore they would like to have an entry for them in the wiki to show the available lambdas.

Erik-Jan announced that Maxine Brown was working on a new version of the GLIF map, and asked the participants to please update her on the newest circuits.

*Action point:* Catalin Meirosu to find a way for AARNet and other GLIF lambda contributors to present their resource information on the GLIF wiki. An additional entry on the right hand side of the webpage will be created underneath the “GOLE and Lambdas” menu.

Heather Boyles asked whether the lambdas discussed were intercontinental or also networks that provide internal lambdas for their constituency could be listed there. Erik-Jan confirmed that only lambdas between GOLEs or between a GOLE and a hybrid network should be listed.

Dave pointed out that many of the hybrid networks that were able to provide lambdas should be present on the GOLE maps as clouds.
René Hatem explained that a cloud meant that if someone saw the same cloud on different diagrams they could get to different places from the same cloud, but they were spared of the details of internal connectivity throughout the cloud (in this example, the CANARIE network in Canada). Erik-Jan said that when further details on the internal cloud connectivity were needed, the map of the network could be consulted on their own website (in this case, the CANARIE website) to find out what the connectivity was. René noted that people should remember that clouds on the GOLE maps still represent lightpath connectivity, not Layer 3.

Erik-Jan asked people that operate resources to please contact Catalin for an account on the GLIF wiki. They could then make sure their information was updated and then the Chairs and secretary will look on how this might be perhaps better presented graphically.

*Action point:* GOLE operators that do not have account on the GLIF wiki or do not remember their password to contact Catalin or the GLIF secretariat for new accounts or password resets.

**3. SC'07 network details and contacts**

Linda Winkler announced that the connectivity will be provided by NLR pretty much in the same way as last year. She had not heard what researchers asked for in terms of access. Internet2 was expected to bring five 10 Gbps lambdas from different directions, up to a total of ten lambdas. Some portion was on that standard Internet2 L3 access, the rest was carved for ESNET and other projects that worked with Internet2. The network in Reno was being built during the week of the GLIF Workshop.

*Action point:* People from Internet2, NLR, and Pacific Wave to send contact details for what to do when people want to request bandwidth for SC'07, and Erik-Jan to relay this information to the Tech WG mailing list.

**4. Establishing of new working groups**

René Hatem reminded the participants that the GLIF Working Group meeting was not a conference, but a working group session. Therefore it would be nice to work on projects during the actual meeting. He was aware of the fact that the participants were very busy. In René's opinion, the best way to proceed would be to work on activities which are homebase priorities anyway. One area where he thought there was particularly interested was Lightpath or wavelength utilization visualisation. He asked the participants whether there would be interest in forming an ad-hoc working group so the volunteers could advance their own capabilities or share efforts and maybe accelerate the evolution of this area. GLIF Tech would serve as a forum for monitoring and discussing progress on a 6-month basis. But this would not be a GLIF deliverable as such, so it could be rather informal.

The other idea for a new working group had to do with managing lightpaths. René reported that CANARIE was using an Excel spreadsheet to do this. This was another problem that CANARIE engineers would be interested in working on collaboratively with the GLIF-Tech community. René noted that there were other areas, such as VLAN management, which might also benefit from collaborative work

Regarding the visualisation of lightpaths, Ronald van der Pol, Lars Fischer, Christian Todorov, Dongkyun Kim, and Linda Winkler indicated interest.

Christian noted that in the case of schedulable resources, a static snapshot would not tell very much about the actual allocation of resources at a given moment in time. Therefore,

a visualisation tool would need to have some notion of time. Otherwise there was a chance of losing the sense of what was really available. René thought this was an excellent point and suggested that the working group would have to discuss this and that different participants might opt for different solutions according to their homebase priorities or capabilities.

*Action point*: René to send a note to the GLIF mailing list about the new working areas on lightpath visualisation and management.

**Phosphorus and DICE IDC: Multi-domain Project approaches - Inder Monga, (Nortel) and Bram Peeters (SURFnet)**

The talk was focused on providing a brief overview of the Phosphorus and DICE Inter-Domain Controllers (IDC) and raise points for further discussion. Phosphorus aims at creating a transparent network service for the middleware. They were working to get a prototype ready during the first year of the project by integrating existing heterogeneous domain controllers that used technologies like UCLP, Viola, and DRAC. Many areas would be open to discussions in the GLIF environment. These included architectural models (centralised vs. distributed vs. some combination of both), the scope of topology exchanges, the way path computation could be performed (at the first domain, or on a hop-by-hop basis), signalling and path management. Phosphorus implemented a centralised, tree-based architectural system, that was viewed as better adapted to co-scheduling multiple types of resources in typical Grid computing applications. Bram and Inder would like to share research experiences and agree on topology exchange, reserve messages and a path compute model. The aim would be to converge on an IDC implementation for a longer term “live”research testbed.

ACTION: ???

**The PHOSPHORUS project supported infrastructure – Artur Binczewski, PSNC**

The Phosphorus project was presented as a European and Global alliance of partners to develop advanced solution of application-level middleware and underlying management and control plane technologies. Artur mentioned that Phosphorus was partly funded by the European Commission. The project aimed to

- demonstrate on demand service delivery across multi-domain/multi-vendor research network test-beds on a European and Worldwide scale
- develop integration between application middleware and transport networks, based on three planes: service plane, network service provisioning plane, control plane
- study resource management and job scheduling algorithms incorporating network-awareness, constraint based routing and advance reservation techniques and develop a simulation environment

Artur reported that a set of applications were identified as early adopters of the Phosphorus infrastructure and the project was working closely with the application teams.

Some example applications included the following: WISDOM (Wide In Silica Docking On Malaria) was a compute-intensive application; KoDaVis focused on distributed visualisation techniques using Unicore middleware; TOPS (Technology for Optical Pixel Streaming) involved constant streaming of ultra high resolution data sets.

ACTION: ????

**Experiences with Fault Resolution Process – Ronald van der Pol (SARA)**

The ordering and fault resolution of lightpaths have been a topic for discussions in the last few GLIF Working Group on Technical Issues. Since June 2006, a document available on the group website contains best practices and lists common issues in this area. Ronald reported that in his experience current practice in ordering lightpaths did not follow the recommendations in this document. Often, there is no clear sourcing organisation and the entire process is driven through a long email thread. This results in unclear clear consumer – provider relationships, with no sub-tasking for fault resolution. Especially during the circuit setup phase, the network elements generated many alarms which contributed to making debugging difficult. In order to address these problems, Ronald suggested to make ordering and fault resolution independent processes. Also, a more efficient information flow between the GOLEs needs to be created. Better monitoring systems to allow for precise fault localisation would facilitate the resolution process. Also debug tools could help by providing facilities for loopbacks, dumping MAC address tables and providing Ethernet traffic statistics. Erik-Jan added that resource operators are encouraged to try the processes listed, and to point potential requesters to the processes.

ACTION: ???

**Multi Domain Coordination in GÉANT2 – Hans Döbbeling (DANTE)**

The GÉANT2 network is a pan-European interconnecting 31 management domains. Multi-domain operations and interactions were thus part of the daily activities. The perfSONAR framework is used for network monitoring across the domains. The End-to-End Coordination Unit (E2ECU) was created in order to monitor and report on circuit services, provide ownership and coordinated problem resolution, synchronize set-up and removal of services, and manage a unified naming for services provided. The E2ECU coordinates these issues directly with the NOCs involved. A Common Network Information System (cNISv2) was being developed to coordinate concepts of operations and adapt NREN business processes to accommodate multi-domain operations use cases.

ACTION: ???