



9th Annual Global LambdaGrid Workshop
27-28 October 2009
Daejeon Convention Centre, Daejeon, South Korea

Technical Working Group session – Day 1 (in the afternoon)

Approximately 60 persons attended the GLIF Technical Working Group session co-chaired by Erik-Jan Bos (SURFnet) and Gigi Karmous-Edwards (NCSU). The secretary was Peter Szegedi (TERENA).

1. Welcome, Apologies, and Agenda

The Technical Working Group session was opened by Erik-Jan Bos. A brief overview about GLIF, and particularly about the task forces of the working group, was given. After the overview of the meeting agenda Erik-Jan called for the GOLE/lambda and hybrid optical network resource updates.

<Slides: <http://www.glif.is/meetings/2009/tech/intro.pdf>>

2. GOLE and resource updates

- *KRLight*

Dongkyun Kim introduced KREONET2 (Korea's National Science & Research Network), KRLight, and GLORIAD (Global Ring Network for Advanced Applications), the first 10 Gbps international hybrid network in Korea with lightpath connections to USA (Seattle) and China (Hong Kong). KRLight, the GOLE of Korea, has 9 connections as of today and lightpath provisioning capability with the highest data transfer rate ever. The Korean test bed is connected to GENI, HPDMnet, and Phosphorus for further collaborations. The GLIF Global Identifier Naming Scheme will be adapted earlier next year and the dynamic GOLE developments are on going.

<Slides: <http://www.glif.is/meetings/2009/tech/kim-krlight.pdf>>

- *TaiwanLight*

Te-Lung Liu presented the domestic network TWAREN (TaiWan Advanced Research and Education Network) and its international part the TaiwanLight. The connections have been expanded to Europe using a 622 Mbps lightpath between MAN LAN (New York City) and NetherLight (Amsterdam). The longest lightpath ever with a single hop from Taiwan to Czech Republic has been created and a live medical streaming application has been tested over it. Recently, TaiwanLight is participating in HPDMnet. The future developments include the migration of TWAREN core nodes to RODAM technology and the implementation of voice quality measurement and hybrid alarm correlation systems on perfSONAR.

<Slides: <http://www.glif.is/meetings/2009/tech/liu-taiwanlight.pdf>>

- *HKOEP Hong Kong*

Haina Tang gave a talk about China's Science & Technology Network and its impressive usage statistic and traffic analysis results. The Hong Kong Open Exchange has 9 lightpath

connections even to commercials like Google. HKOEP is involved in e-VLBI and various e-Science projects worldwide.

<Slides: <http://www.glif.is/meetings/2009/tech/tang-hkoep.pdf>>

- *MoscowLight*

Natalia Bulashova introduced the MoscowLight network and its two PoPs in Moscow. Many local institutions are connected by fibre optics and the GOLE has international lightpath connections to NorthernLight, NetherLight and StarLight. RIPN (Russian Institute for Public Networks), the funding organization of MoscowLight, is in partnership with various regions in Russia. Especially in the Moscow region the High energy Physics (HEP) applications use the network heavily.

<Slides: <http://www.glif.is/meetings/2009/tech/bulashova-moscowlight.pdf>>

- *AARnet*

Chris Hancock, on behalf of Ivan Philips, presented about AARnet, Australia's NREN, including the domestic footprint as well as the optical and the IP network topologies and equipment. During the past year various lightpath related activities were performed such as: the first 4k Super High Definition video stream delivered via AARNet3 to the SC'08 venue (TX, Huston), the EN4R (Experimental Networks for Researchers) support during International Year of Astronomy, the 10 Gbps-ready eVLBI correlation study at Curtin University which highlighted the road towards SKA (Square Kilometre Array) readiness, and the transmission of the total solar eclipse (which was visible in Japan, but not in Australia) from Keio University in Japan to Questacon, to an audience of high school science students. The fibre built business is still booming in Australia and AARnet is deploying regional DWDM networks on top of it. They are also investigating 40G and 100G networking for MPLS backbone (L3VPN and L2VPN/VPLS) and IP network capacity upgrades to support 10G IP customers.

<Slides: <http://www.glif.is/meetings/2009/tech/hancock-aarnet.pdf>>

- *SouthernLight*

Michael Stanton introduced SouthernLight, the distributed GOLE co-administered by ANSP and RNP. The GOLE interconnects the three Brazilian participating networks and now is connected to the AMPATH GOLE in Miami by means of 2x10G links, installed between July and October 2009, replacing the previous 2.5G link. The first 10G link was inaugurated in July 2009 with a lightpath demo involving the digital cinema community with the transmission of compressed 4K digital media (400 Mbps) and uncompressed HD videoconferencing (900 Mbps) between the FILE 4KT in São Paulo, UCSD (USA) and Keio University (Japan). Regular e-science usage expected in support of HEP, e-VLBI and Dark Energy Survey communities. There are on going collaborations with the GENI and FIRE communities for Brazilian inclusion in future network research test beds. Recently, RNP is carrying out a comparative study of alternative schemes for dynamic lightpath provisioning, with the aim of introduction of a pilot service in 2010. The RedClara2 (Latin American regional network) is acquiring cross-border fibre supporting 10G links interconnecting Brazil to Chile and Argentina by 2010. It is expected that these links will also provide lightpath connections from SouthernLight to these countries.

<Slides: <http://www.glif.is/meetings/2009/tech/stanton-southernlight.pdf>>

- *CANARIE*

Eric Bernier showed the CANARIE network footprint, SONET over ROADM infrastructure, and Seattle and Chicago GOLEs. Via the StarLight and PacificWave GOLEs CANARIE has many international lightpath connections. They are actively participating in HPDMnet demonstrations. Last year additional IP connections were set up to reduce the latency in the IP backbone. Recently they are upgrading the IP infrastructure. Five new small size Juniper MX480 routers were purchased with L2/3 switching capability and 240G full duplex backplane. In addition, new Nortel OME6500 20G L2SS cards were deployed (Layer 2 Service Switch) and new OME6500 10GE wavelength cards were installed in the core network.

<Slides: <http://www.glif.is/meetings/2009/tech/bernier-canarie.pdf>>

- *StarLight*

Joe Mambretti, on behalf of Alan Verlo and all the StarLight engineers, introduced StarLight infrastructure with many lambda connections and collaborators. Among the collaborators, the CineGrid project, the National Centre for Data Mining expanding TeraFlow and OpenCloud test beds (twice as big as Google has in its data centre), the HPDMnet, and the OptIPuter were mentioned. HP Labs recently connected to StarLight via NLR-CENIC-CAVEwave testing novel protocols. The iGENI (International Global Environment for Network Innovation) award was formally announced during the GLIF plenary.

<Slides: <http://www.glif.is/meetings/2009/tech/mambretti-starlight.pdf>>

- *MAN LAN*

Eric Boyd announced that MAN LAN is now being managed by Dale Finkelson. In terms of IP connectivity Qatar Foundation changed provider and now is connected at 1G Ethernet, MIT is connected at 10G, and NLR has at least one new 10G connection to Europe. In terms of bandwidth MAN LAN is planning to migrate to support 100 Gbps, as funding becomes available. Internet2 has introduced ION service (i.e., open source software includes OSCARS and DRAGON) across backbone, enabling automatic provisioning of circuits. Internet2 ION service and ESnet Science Data Network are capable of peering with other IDCs at MAN LAN.

MAN LAN supports E2EMON monitoring of static circuits using perfSONAR. It is planned to support automatic monitoring of dynamic circuits, as is currently supported for ION service. Internet2 and its partners will support open source software for implementing perfSONAR and IDCs on dynamic GOLEs. Internet2 will contribute to and follow lead of GLIF and OGF in defining advanced cyberinfrastructure protocols.

<Slides: <http://www.glif.is/meetings/2009/tech/boyd-manlan.pdf>>

- *PacificWave*

Dave Reese described PacificWave, a distributed GOLE over three PoPs (Seattle, WA, Sunnyvale/San Jose, CA, Los Angeles, CA), jointly operated by PNWGP and CENIC. Recently there was an upgrade to 2x10G Ethernet interconnections among the PoPs. They are working on deploying GLIF naming standards within PacificWave naming/numbering scheme. The architecture consists of L2 switches will support demonstrations during SC'09.

<Slides: <http://www.glif.is/meetings/2009/tech/reese-pacwave.pdf>>

- *AtlanticWave*

Julio Ibarra introduced AtlanticWave, a distributed exchange point peering service along the Atlantic rim. It consists of MANLAN in New York City, AMPATH in Miami, MAX Gigapop and NGIX-East in Washington DC, and SoX Gigapop in Atlanta. SouthernLight in Sao Paulo is in progress to join the AtlanticWave collaboration.

AMPATH is an international exchange point serving network-enabled U.S., Latin America, and Caribbean science research and education communities. AMPATH provides multiple L2/L3 paths to U.S. and international research and education backbone networks. There is a production Layer 2 10G Ethernet transport service via AtlanticWave to U.S. national and international backbone networks and an experimental Layer 2 10G Ethernet transport service via Cisco Research Wave (C-Wave). PerfSONAR test bed is installed at AMPATH. They are participating in SC'09 Bandwidth Challenge Preparations and GENI-Brazil development test bed project.

<Slides: <http://www.glif.is/meetings/2009/tech/ibarra-atlanticwave.pdf>>

- *US LHCnet*

Artur Barczyk gave an overview on US LHCnet topology and architecture. US LHCnet is a mission oriented network with 4 PoPs providing trans-Atlantic network infrastructure to support DOE and NSF HEP programs, focused on LHC. Recently, the providers of some circuits have been changed. There was no major impact on the production traffic. The Ciena CoreDirectors at MANLAN, StarLight, and CERN have also been upgraded. The reliability of the network was enhanced by adding resiliency features through path diversity and mesh protection. Dynamic circuit switching capability between all PoPs is being implemented using DCN SS. This is interoperable with the Internet2 ION and ESnet OSCARS. The long term objective is to connect dynamic circuit networks on both sides of the Atlantic.

<Slides: <http://www.glif.is/meetings/2009/tech/barczyk-uslhcnnet.pdf>>

- *NorthernLight*

Lars Fischer presented NorthernLight, now also scheduled to connect the North Atlantic. The emerging importance of Greenland and Iceland was mentioned in the context of placing green data centres there. The optical core network is based on DWDM technology with wavelength selector switches, reconfigurable and tunable 10G /40G (100G in the future) waves up to 96 channels. A multi-domain 40G alien wave transport terminated by Nortel end-systems and transported by Alcatel-Lucent boxes was successfully demonstrated. The transport infrastructure can be used by parallel virtual IP networks supported by MANTICORE. In the future the cross border fibre efforts are planned to continue together with the 100G trials and dynamic GOLE deployments.

<Slides: <http://www.glif.is/meetings/2009/tech/fischer-northernlight.pdf>>

- *JANET Lightpath*

Rob Evans described the JANET lightpath service. In the past the service was based on SDH technology with STM-64 circuits between backbone PoPs. It was reliable but expensive solution with high operational overhead. That is why the core SDH infrastructure has been replaced by Juniper MX-960 boxes using Ethernet over MPLS technology and 10G Ethernet LAN PHY connections. The current provisioning is manual, restricted to point-to-point connections only, and not applying QoS. But, it was installed five months ago and working well.

<Slides: <http://www.glif.is/meetings/2009/tech/evans-janet.pdf>>

- *NetherLight*

Wouter Huisman showed the map of NetherLight GOLE connectivity. The network is capable to provide both SDH and Ethernet VLAN based services and even to combine the two. The interconnection between the SDH and Ethernet data planes can be done by the newly installed 20G L2SS cards. The cards support 2x10G and 8x1G Ethernet ports with SDH GFP mapping towards the backplane. The PBB-TE was introduced in September 2009. It gives additional flexibility, security and scalability to the network. In the future the replacement of HDXs with Nortel OME6500 Double Deckers and new cross border fibre system installation to CERN are planned.

<Slides: <http://www.glif.is/meetings/2009/tech/huisman-netherlight.pdf>>

At the end of the GOLE and resource update session Erik-Jan acknowledged the reasonable upgrades since the last meeting in Seattle. It is always useful to share the latest information about equipment upgrades/replacements, technologies and networks, he said. The resource owners were kindly asked to update the GLIF Wiki (<http://wiki.glif.is>) according to the latest developments. Please contact the GLIF Secretariat for credentials if you do not have them yet, or lost yours.

3. GNI API Task Force (1/3)

Evangelos Chaniotakis (ESnet) presented the GNI API Task Force objectives and the latest achievements since the winter meeting in Catania, Italy. The overview of the activity was dated back to the Honolulu meeting where it was decided to bring the various APIs under the same umbrella, known as Generic Network Interface (GNI). The aim is to develop a software framework to make sure that any kind of API (Application Programming Interface) translation can seamlessly be done. To achieve this, one common interface is needed and all API calls must then be translated in there. Originally, it was not the objective of the task force to develop a standard interface that is ready for production.

Recently KISTI and NCSU joined to the task force development team. New external and internal interfaces have been drafted and the GUSI framework (introduced during the Catania meeting) has been completely rewritten and renamed to Fenius. In Fenius there are translators implemented for IDC, dynamicKL, Harmony, and G-lambda. Currently the documentation is lacking but the code is open source and can be found on Google Code (<http://code.google.com/p/fenius>).

The future plans are to collaborate, grow, develop, and deploy. The collaboration with OGF is essential to provide feedback to NSI (Network Service Interface) working group, to provide a framework rapidly testing potential NSI interfaces, and to use NML (Network Mark-up Language) for topology exchange. It is planned to invite AutoBAHN and others (e.g., DRAC) into the task force to grow, as well as to align Fenius and Harmony. The future development plans include the improvement of internal and external interfaces, the additional security layer, and the additional topology exchange feature add-in to Fenius. The long term plan is to deploy the improved Fenius instances on real networks and GOLEs.

The current Fenius version is not ready for production but good enough for demonstrations. There is no security feature implemented yet. For short term solution a simple username / password protection is preferred with the potential improvement to include IdPs in the future. The two main questions in the context of production use are; who wants to run Fenius on top

of their own domain and who will support it. However, before it turns to production we need to proof of the concept by volunteers. Finally, Evangelos thanked the hard work to each and every developer and briefly described the Fenius demo set up.

<Slides: <http://www.glif.is/meetings/2009/tech/chaniotakis-gni.pdf>>

4. perfSONAR Task Force

Thomas Tam (CANARIE) gave an overview on the lightpaths monitoring challenges, introduced the distributed status architecture of perfSONAR and explained the perfSONAR link status monitoring demonstration set up. End-to-end lightpaths often cross multiple domains that cause many issues (i.e., various administrators, operating hours and time zone differences, etc.) in case of unexpected failures. The perfSONAR link status service can help the NOCs to manage cross-domain lightpaths effectively. The demonstration lightpath spans six network domains (KRLight/KREONet, CANARIE, Internet2, Netherlight, NorthernLight, and UNINETT) and a single E2EMon web client collects all the link statuses.

The latest perfSONAR enhancements are the extension of the XML schema and the URL referral exposed through the web service interface with the operation and administration states. This URL link could be used by other NOC tools such as ticketing system, performance monitoring counters, existing MRTG graphs, etc.

Thomas mentioned the current perfSONAR deployment challenges. Monitoring metrics may vary from domain to domain and network operations require time to integrate their existing tools to perfSONAR service. However, there are some on going enhancement of usability. The monitoring schema will be expanded for new functionalities and a new web GUI needs to be developed. The plan is to involve a few more network operators in the perfSONAR Task Force and ask everyone to join and express interest in the tool.

<Slides: <http://www.glif.is/meetings/2009/tech/tam-perfsonar.pdf>>

5. Demonstrations

After the Technical Working Group afternoon session the demonstrations went in parallel repeated six times and the six groups of attendees had chance to rotate among the presentation booths accordingly.

The list of demonstrations was as follows:

- HPMDnet - *iCAIR, KISTI & HPDMNet*
- perfSONAR - *CANARIE*
- IDC: MAN LAN - NetherLight - *MAN LAN & NetherLight*
- Fast Data Transfer for LHC - *CERN*
- GNI API Fenius interoperability framework - *KISTI, G-lambda, ESnet, & HPDMnet*
- Medical Video Streaming - *NTNU*

It can be noted that the Fenius interoperability demonstration was the first working demonstration ever of the new common interface developed by the GNI API Task Force. The task force will document the lessons learned during the demonstration and share them with GLIF community as well as the OGF NSI Working Group.

All demonstrations went very well, with lots of interaction between the demonstrators and the large number of attendees.

Technical Working Group session – Day 2 (in the morning)

The second day was chaired by Gigi Karmous-Edwards (NCSU) and Erik-Jan Bos (SURFnet). Gigi introduced the speakers and the planned agenda.

6. Dynamic GOLE Services Task Force

Cees de Laat (UvA), on behalf of John Vollbrecht, presented an automated GOLE pilot project. The pilot is an operational project to bring automated provisioning to some GLIF exchange points. A small subset of GOLE resources is expected to be available for a minimum of a year by each participant of the project. The pilot includes analyzing and deploying software for automated switching and finding applications that can use the pilot resources.

Three main project phases have been proposed (to be revised by project team):

- *Phase 1:* Evaluate pilot requirements, resources at exchange points, policy requirements, potential users, and connectivity to and between GOLEs.
- *Phase 2:* Implement hardware and software at participating sites, connect to participating networks and debug, and finally run demos.
- *Phase 3:* Evaluate and report on pilot, determine next steps.

This automated GOLE pilot will be limited only for Layer 2 (if successful, follow-on with other layers). The initial aim is to protect the existing static exchange resources from automated software trials. The pilot will have an interface to allow connections between GOLE ports and will connect to automated networks at both the transport and control level allowing end-to-end connections that include GOLEs and multiple networks. Cees showed the ideal GOLE architecture and the proposed L2 GOLE pilot architecture. As a main principle, the pilot protects operational switch by moving switching to isolated location (i.e., a separated Ethernet switch will be used for switching / VLAN re-tagging). The role of GOLEs in an automated connection infrastructure was explained. The pilot will determine what is required for software to be “pilot compatible” (Fenius is a potential candidate for the pilot).

The next steps are to identify participants, chairs, and set up a kick-off call. Gigi asked for the interest level among the meeting participants. SURFnet, NORDUnet, Internet2, ESnet, and many others expressed their interest. The task force will be chaired by John Vollbrecht (Internet2) but a co-chair is also needed (open for volunteers).

<Slides: <http://www.glif.is/meetings/2009/tech/delaat-dynamic-gole.pdf>>

7. GNI API Task Force (2/3)

Evangelos Chaniotakis (ESnet) continued the discussion on Fenius from the previous day. Three main issues were discussed: security, topology exchange and pathfinding.

Regarding the security features of Fenius it is needed to define what the bare minimum requirement is and what desired on top of that. Security depends on whether the peering trust model or the centralised trust model is applied for now. In long term, the peering trust model is preferred. In principle, all the meeting participants agreed that a secure Fenius is needed. In short term, a very basic approach can be followed using SSL for message transport and simple HTTP authentication with username / password. This will work because there is only a small and well-known set of agents that perform Fenius request currently. In the future the task force will consider alternate schemes for the authentication, identity, and authorization

problems depending on the first experiences and as required by the community. It highly depends on the maturity of IdMs at various domains.

Regarding the topology exchange issue basically it was decided that Fenius will not do topology exchange at all (it is outside the API's scope), but a view of the topology is needed so that Fenius clients can do multi-domain pathfinding. PerfSONAR has a feasible framework to leverage on its topology service and it is flexible enough to accept various schemas. GNA API Task Force will recommend that Fenius-enabled sites use the PerfSONAR topology service and topology schemas to expose their topology to the world.

In principle, Fenius is not smart and cannot do pathfinding. It was agreed that the pathfinding problem is not strictly inside the Fenius scope however it was mentioned that a simple agent that can do basic pathfinding and perform requests will be very useful. It can also be used as a starting point for interested parties to build more sophisticated agents. In long term probably a broker is needed on top of Fenius translator to do the pathfinding.

In the context of Fenius a reference implantation of intra-domain pathfinder can be built up. The future development plan is to take a PerfSONAR topology component, write a (very simple) inter-domain pathfinder, and optionally add on top the existing Web GUI.

<Slides: <http://www.glif.is/meetings/2009/tech/chaniotakis-gni.pdf>>

8. GNI API Task Force (3/3) – on open source DRAC

Rodney Wilson (Nortel) gave a follow on presentation on the Nortel's open source DRAC announcement done during the plenary. In order to fully realize the potential of DRAC and to meet the needs of the research and education networking collaborators Nortel announced their intention to make DRAC available for the GLIF Community under an open source license.

DRAC is for the applications to see the network as manageable resource. Rodney gave a history overview. In 2004, Nortel did set up a proof of concept trans-continental, multi-domain link controlled by DRAC. In 2005, an application controlled dynamic path was demonstrated. In 2006, they demonstrated the on-the-fly migration of an operating system and running applications seamlessly from one continent to another. DRAC was the main engine that allows it. In 2008 and 2009 DRAC development was robust. Nortel realised that if DRAC is open source GLIF can lift it up high.

Bram Peeters (SURFnet) presented DRAC from SURFnet's perspective. DRAC was introduced to SURFnet network in 2004 as a result of the SURFnet6 tender. The DRAC-based dynamic network service was launched at the end of 2008. Currently there is no multi-layer, multi-technology and multi-vendor functionality in DRAC but it is good enough for basic intra-domain production service and some code is underway for interdomain. Dynamic Services are a core part of the future developments of SURFnet.

Bram introduced the basic service design, tools, administration interfaces and the Web GUI. The current service is capable to provision single or protected paths (scheduled, repeatable, first fit timeslot allocation) between transparent L2 Gigabit Ethernet (GE) ports or between GE and STM-64 ports. There is a clear administrative separation from other services. The dynamic service is implemented on Nortel OME6500 boxes. In the near future, thanks to the new L2SS cards, VLAN to GE and VLAN to STM-64 provisioning will also be available. The basic principle is to have no limitation on vendor and equipment in SURFnet.

The open source DRAC is a big advantage to the GLIF community, but not just the code needs to be open source but a whole collaboration platform is needed. SURFnet and Nortel called for an open source community to develop DRAC. GPLv3 licence will be used. DRAC is one of the candidates for the south-band interface of Fenius to support multi-domain capability. The development roadmap can include the multi-vendor support, Carrier Ethernet (PBB-TE, T-MPLS) support, planned works management, incidents/service management, and group management following a federated approach.

<Slides: <http://www.glif.is/meetings/2009/tech/peeters-drac.pdf>>

The DRAC development server is available here: <http://drac.surfnet.nl/> (username: GLIF2009 password: Glif2009!)

9. Global Identifiers Task Force – status update

Ronald van der Pol (SARA) gave a status report on Global Identifiers Task Force. The task force was established in 2008. In April, 2009 the proposed naming scheme was adopted. The actual status is as follows:

- MAN LAN, T-LEX and DANTE/GÉANT3 are still in the initial discussion/orientation phase.
- NorthernLight, TawianLight, SouthernLight, MoscowLight and PacificWave are already in the planning phase (decided to use).
- At CANARIE and KRLight there is an on going deployment scheduled in new monitoring system.
- JANET and NetherLight have fully deployed the scheme in monitoring and ticketing systems.
- No information about the others.

The meeting participants agreed that the task force has achieved its goals. It will be kept alive until the Technical Working Group winter for another reporting period, and then it will be closed.

<Slides: <http://www.glif.is/meetings/2009/tech/Ronald-gid-status.pdf>>

10. Discussions

The discussion session was led by Erik-Jan Bos (SURFnet) and Gigi Karmous-Edwards (NCSU). Three emerging topics were discussed in more details; end site challenges, next generation GOLE architectures, and green GLIF proposal.

- *Access to Hybrid Networks - End Site Challenges*

Ronald van der Pol (SARA) gave a presentation about the end site challenges. The traditional IP based end site connection practice was compared to the novel hybrid L1/L2 connection model where the administrative boundary is not so clear, the routers are by-passed thus BGP and ACLs cannot be controlled. To illustrate this issue, SARA routing policy was introduced. Many open questions were raised by Ronald regarding how to protect against misconfiguration of links, Ethernet loops and broadcast storms, what about firewalling, IP addressing, performance challenges, and so on.

It was agreed that there is a lack of technical knowledge and necessary software tools at the end sites. The challenges should be collected and categorised on a wiki page. Peter Szegedi (TERENA) mentioned the End-to-End Workshop series organized by TERENA to its broader community as a similar activity dealing with end site challenges. It was not clear yet whether GLIF should create a new task force on this issue. Further discussions are needed and the

outcome will be presented at the Technical Working Group winter meeting in Salt Lake City, UT, USA.

<Slides: <http://www.glif.is/meetings/2009/tech/Ronald-challenges.pdf>>

- *GOLE - HDXc Refresh Next Generation Architecture*

Eric Bernier (CANARIE) made a proposal on a new GLIF task force under the Technical Working Group. One of the motivations is that the manufacture of the Nortel HDXc optical cross-connects, used at many places in the GLIF network, discontinued after 30 June, 2009 and the end of support is on the horizon. A replacement strategy is needed and this opportunity can be used to define new features and architecture. The proposal is to create a new task force with the objective to define the next generation GOLE architecture. Large number of GOLE operators presented at the meeting agreed to create this new task force. The Next Generation GOLE Architecture Task Force will be led by Eric Bernier.

<Slides: <http://www.glif.is/meetings/2009/tech/bernier-hdxc.pdf>>

- *GLIF Technical Working Group meeting summary*

Gigi Karmous-Edwards (NCSU) summarised the current task forces and the proposed new activities in the Technical Working Group.

Among the operational task forces the Global Identifiers Task Force completed its objectives. It remains open until the winter meeting in Salt Lake City to report back with the latest status of the scheme uptake. The perfSONAR Task Force is still active. Some progress is expected in the area of inter-domain lightpath management. The Dynamic GOLE Services Task Force agreed to develop a Layer 2 operational dynamic GOLE experiment.

The only one research task force currently running is the Generic Network Interface (GNI-API) Specifications Task Force. A successful demonstration of a simple adaptor was performed during the meeting. However, security needs a simple temporary solution soon and the multi-domain pathfinding should also be solved. The topology exchange requires a simple mechanism based on the perfSONAR framework. The open source DRAC gives a good opportunity to develop these missing features for Fenius.

Among the new activities three topics were mentioned by Gigi. The emerging end site challenges were highlighted by Ronald. The meeting participants agreed on not to create a new task force yet, but try to clarify the issues by the winter meeting in Salt Lake City. The new Next Generation GOLE Architecture Task Force proposed and led by Eric has been accepted by the participants. Finally, the Green ITS panel discussion was announced by Gigi as a potential new task force activity, if appropriate.

<Slides: <http://www.glif.is/meetings/2009/tech/conclusions.pdf>>

As the closing of the Technical Working Group session Erik-Jan asked again for the GLIF Wiki update with the latest GOLE and resource information. It was announced that the next coming GLIF Technical Working Group winter meeting will be held in Salt Lake City, UT, USA, on 3-4 February, 2010, in conjunction with the ESCC/Internet2 Joint Techs meeting.