



11th Annual Global LambdaGrid Workshop
13-14 September 2011
Rio de Janeiro, Brazil

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

Approximately 50 persons attended the GLIF Technical Working Group session at the 11th Annual Global LambdaGrid Workshop, co-chaired by Erik-Jan Bos (NORDUnet) and Gigi Karmous-Edwards (NCSU). Erik-Jan Bos temporarily substituted Lars Fischer, CTO of NORDUnet, who was on paternity leave. The secretary was Peter Szegedi (TERENA).

1. Welcome, apologies, and agenda

The Technical Working Group session was opened by Erik-Jan Bos. A brief introduction to GLIF, and particularly to the active task forces of the working group, was given (<http://www.glif.is/meetings/2011/tech/130911-intro-ejb.pdf>). After the overview of the meeting agenda, Erik-Jan also announced the next coming GLIF Tech WG Meeting that will be held on 25-26 January 2011, in Baton Rouge, Louisiana, USA. The meeting will be co-located with the Winter 2012 ESCC/Internet2 Joint Techs to be hosted by the Louisiana Optical Network Initiative (LONI) on the campus of the Louisiana State University.

2. GOLE and resource updates

The technical programme started with flash presentations about the major updates from GOLEs and other resources.

- *StarLight*

Joe Mambretti (Northwestern University) gave an update on the current StarLight infrastructure and the StarLight International Consortium and he also mentioned several projects on-going. The TransLight/StarLight team is pursuing application experiments on international networks with 6 different activity areas. These areas are:

- High-Performance Digital Media Network (HPDMnet),
- iGENI: the GENI-funded international GENI project,
- SAGE: connecting people and their data at high-res,
- CineGrid: it's all about visual communications,
- GreenLight International: less watts/terabyte, and
- Science Cloud Communication Services Network (SCCSnet): the impending disruption.

Joe mentioned that it should be easy to contribute to these initiatives from the US, joining from other world regions is a bit more difficult. The trans-oceanic resources must be in place first. The GLIF community can take a role here.

(<http://www.glif.is/meetings/2011/tech/130911-StarLight-jm.pdf>)

- *Pacific Wave*

Dave Reese (CENIC) gave an update on Pacific Wave. There is sufficient capacity in the network to handle the current traffic although additional 10G capacity can be made available as needed. New 2x10G capacity from Hawaii is landing at Los Osos, CA extending to Los Angeles and Seattle (under University of Hawaii's NSF ARI grant). There is an on-going work on lighting the West Coast fibre to support 100G.

Pacific Wave is now integrated with Western Region Network (WRN) for added access to Hawaii, Colorado (FRGP), and New Mexico (UNM).
(<http://www.glif.is/meetings/2011/tech/130911-PacificWave-dr.pdf>)

- *NetherLight and LHCONE*

Gerben van Malenstein (SURFnet) mentioned NetherLight's new 1G Taj-GLORIAD link to Egypt, the 100G clear channel transmission to CERN, Geneva, and the n*10G upgrade to NORDUnet. NG-Ethernet is available as pre-production service that means that PBB-TE port type can be selected by the users. There is a new NetherLight website with references to projects such as LHCONE. A SURFnet paper on Open Exchanges is also published on the web. The new SURFnet7 equipment vendor is Ciena and Ciena have announced to work on implementing NSI into their production network management system.

In the future all NetherLight switches will be replaced by the Ciena 5410 Ethernet platform and extended measurement and management information will be provided. NetherLight also plans to connect commercial service suppliers such as storage providers and green data centres.

(<http://www.glif.is/meetings/2011/tech/130911-NetherLight-gvm.pdf>)

- *AMPATH*

Julio Ibarra (Florida International University) talked about the enhancements at AMPATH and AMLIGHT. AMLIGHT is linking the research and education networks in the Americas with international exchange points such as SouthernLight, AMPATH, SoX, MAX, and MANLAN. The Caribbean region is not on the AMLIGHT map at the moment but there are activates there too. RedCLARA provides the leadership in Latin America. A couple of countries have not yet been connected but the plan is to have them all.

The hybrid networking services implemented at AMPATH are based on OSCARS and capable of dynamic circuit provisioning. OpenFlow is being deployed in collaboration with Internet2, RNP, and NDDI as well as Open Science Data Cloud (OSDC) in collaboration with University of Chicago. One of the future plans is to develop AtlanticWave's 5-year plan, including plan for additional capacity, support for hybrid network services and performance monitoring.

(<http://www.glif.is/meetings/2011/tech/130911-ampath-ji.pdf>)

- *TWAREN*

Li-Chi Ku (NCHC) gave an update on TAIWANLight. The domestic network TWAREN is resiliently connected to US with 2*2.5G pacific crossing lightpaths. A single 622 Mbps connection was expanded to Europe in 2006. OpenFlow test bed has been deployed on top of TWAREN that is extended to iGENI@iCAIR and HPDMnet. The Taiwan Integrated Research Network is participating in the iGENI initiative.

(<http://www.glif.is/meetings/2011/tech/130911-TAIWANLight-lck.pdf>)

- *ENSTINET and GLORIAD*

Majid M. Al Sadek (ENSTINET) introduced the Egyptian NREN and its international connections for the first time. The current network is based on GbE connections and it is 10G-ready. The on-going research activities include the activation of a LHCnode at the Academy of Scientific Research & Technology. The collaboration with African

countries (Sudan, Kenya, etc.) is facilitated by the recently established Egyptian GOLE. The Egyptian S&T portal has recently been launched for further reference. (<http://www.glif.is/meetings/2011/tech/130911-egypt-mmas.pdf>)

3. Demonstrations

During the demo session on the first day, several demonstrations were organised that included

- the interoperability of seven different systems implementing OGF's Network Services Interface (NSI) protocol;
- a CPqD 4K visualisation tiled display wall featuring both collaboration software for sharing multiple windows of information and high-performance audio-visual media transmitted from the University of Essex in the UK using Intopix equipment at both ends;
- the microscopic observation of high bandwidth streaming performed by Keio University;
- and the dynamic set-up and tear-down of lightpaths between user institutions of RNP in Brazil and Internet2 in the US, demonstrating the interoperation of their respective Cipó (pre-production) and ION (production) on-demand dynamic circuit services.
- There was also a closed circuit demonstration of 100 Gb/s DWDM transmission technologies done by CPqD, along with a high-performance packet inspection system.

(<http://www.glif.is/meetings/2011/demos/>)

4. GLIF Open Lightpath Exchanges Policy Group

Bill St. Arnaud gave an introduction to the topic of Governance-Policy Framework for Open LightPath Exchanges (GOLEs) and Connecting Networks and moderated an open discussion afterwards.

The white paper that has previously been published with the initial considerations on open exchanges is available at

http://www.glif.is/publications/papers/20110519BStA_Open_Exchanges.pdf

In summary, it has been agreed upon so far that an Open Exchange is “policy free” and it is non-blocking in terms of bandwidth between connectors. Interconnection between OLEs is independent of ownership of OLEs. Governance does not mean central management or control, in this framework it refers to issues of policy of a federation of resources. The governance models for interconnecting pipes should be a self-defined autonomous architecture (like the Internet eXchange model).

(<http://www.glif.is/meetings/2011/tech/130911-Governance-Policy-bsa.pdf>)

Discussion:

The discussion started at clarifying the term “policy free”. It means that there are no restrictions on technology, bandwidth, traffic type, content, etc. at the lightpath exchange itself but there might be recommended best practices for interconnection (i.e., eduroam type governance). Restrictive policy is defined by the connectors only. By definition, it means that commercials can connect to an OLE too. It's important for connecting commercial storage providers, for instance. The circuits between OLEs are different story. The circuit owner may enforce policy on those links but the decision of the circuit owner is separate from the OLE.

There was a question about the definition of a distributed OLE. A distributed OLE consists of OLE nodes and connections in between. The whole distributed OLE structure must be policy

free and non-blocking by definition that imposes some requirements on the technology of the internal connections. Pacific Wave noted that they don't want to compete with commercials providing policy free connections between OLE nodes. This raised the issue that at some places around the world OLEs should not allow commercials to connect to other commercials. Joe Mambretti noted that the policy free commercial to commercial interconnection is apparently illegal by the US law. It cannot be done policy free!

A clear definition of commercial traffic is needed to understand this issue better. OLE is not just a switch but a service. Some ports can be part of a policy free OLE some ports not, although it is some kind of policy after all. But the policy of the links must be deferent than the overall OLE policy. Kees Neggers commented that different policy can be set up by the OLE operator but the operator needs to be clear what policy is that. OLEs need to advertise their policy. A protocol is needed to publish policy and make sure that the policy is understood in technical means. The Distributed Topology Exchange Task Force (DTOX) has to work on that how policy information can be exchanged in a machine readable format. The University of Amsterdam is working on policy and trust models and offered to contribute to DTOX. Bill St. Arnaud summarised that we need an agreement on a "policy engine" hopefully during the next coming GLIF Tech WG meeting in Louisiana.

There was another discussion about the governance models for interconnecting pipes. Different policy models (such as the old telephony approach or the dynamic LCHONE approach) can live in parallel (maybe overlap) for the moment, we do not have to choose now. The question was that if we should discuss this further in a Policy Task Force. Erik-Jan Bos felt that we need to discuss about it more frequently than GLIF meetings happen. Instead of creating a new task force, we have to work on a problem statement document first.

Gigi Karmous-Edwards summarised the outcome of the discussions as follows:

- 1) Regarding the definition of OLE; an update is needed to the existing white paper to be presented at the GLIF Tech Winter meeting in Louisiana.
- 2) Regarding the link policy issue; a problem statement document has to be written first. A new task force may be created later on.

Bill St. Arnaud, Erik-Jan Bos (NORDUnet), and Inder Monga (ESnet) agreed to work on the documents mentioned above.

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

The second day was chaired by Gigi Karmous-Edwards (NCSU). She gave an update on the currently running and officially closing task forces of the Technical Working Group:

- The Dynamic GOLE Services Task Force led by Jerry Sobiesky (NORDUnet) and the Distributed Topology Exchange Task Force led by Jeroen van der Ham (UvA) are progressing. The Campus Networking Task Force led by Ronald van der Pol (SARA) has given its final report at this GLIF meeting in Rio.
- The Generic Network Interface Specifications Task Force led by Evangelos Chaniotakis (ESnet) was concluded last time in Hong Kong and now is officially closed. The Next Generation GOLE Architecture Task Force led by Éric Bernier (former CANARIE) is also considered to be officially closed. The Resource Allocation Task Force led by Gigi Karmous-Edwards (NCSU) is dormant; suggested to be closed and put the efforts in potential new task forces.

Gigi also announced that she is stepping down as co-chair of the Technical Working Group. Karmous-Edwards led the original Control Plane Working Group until its merger with the Technical Working Group in 2008, when she became a co-chair of that group. GLIF extended its thanks for her work over the years, and plans to confirm her replacement before the next Technical Working Group meeting on 25-26 January 2012 in Baton Rouge, Louisiana, USA.

5. Campus Networking Task Force closing

Ronald van der Pol (SARA) presented the final result of the campus networking survey. (<http://www.glif.is/meetings/2011/tech/140911-campus-rvdp.pdf>)

Kwangjong Cho (KISTI) presented some lightpath issues between NREN and campus networks in South Korea. As he stated, still the most IT staff (including networking staff) of campus networks are not familiar with lightpath technology. Sometimes the researcher knows more than the IT department. Usually in South Korea commercials are connected to the campus border routers, but do not provide lightpath services because of the additional expenses. There is no real competition with the NREN in that respect. (<http://www.glif.is/meetings/2011/tech/140911-LightpathIssues-kc.pdf>)

Ronald elaborated on the future of the task force. The task force charter was to “determine the needs and requirements of campus networkers, produce information on how to setup and use lightpaths, and encourage and support tests and demonstration.” Some documents on demonstration set ups (done by KISTI and SURFnet) have been collected and made available on the GLIF Wiki at http://wiki.glif.is/index.php/Campus_Networking. No real best current practices have been collected so far, but the task force did all the possible facilitations within the GLIF community and can be concluded. The Campus Networking Task Force was officially closed with the remark that the Wiki page is still available for collecting best practice documents. (<http://www.glif.is/meetings/2011/tech/140911-campusfuture-rvdp.pdf>)

Peter Szegedi (TERENA), the organiser of the former TERENA End-to-End Provisioning Workshops, noted that there is a plan for 2012 to organise a “global” Lightpath Services Workshop for Campuses. This mainly campus oriented event should facilitate and boost the uptake of lightpaths and dynamic circuit services at campuses. Such an activity must be built on liaisons with all global peer groups (i.e. GLIF, TERENA, Internet2, APAN) possible. Erik-Jan supported the idea and encouraged TERENA to try and organise such an event in coordination with the global peers. The GLIF community would definitely welcome and support this initiative.

6. Distributed Topology Exchange Task Force update

Jeroen van der Ham (UvA) gave an update on the latest progress of the DTOX task force. For the Automated GOLE demonstration the control plane topology was collected and there is no information on the data plane topology. The LHCONE use case demonstrates that labels are becoming necessary ingredients for endpoint determination although label swapping is not yet universal. It is planned to demonstrate fully automated provisioning through NSI with accurate topology exchange and optionally with end point discovery function by September next year.

Jeroen showed some theoretical results on the chance of successful path finding against the topology aggregation method. The more knowledge on the internal topology is collected the

better chance to have a successful path finding. Jeroen clearly advocated that some knowledge on the internal topology is needed.

(<http://www.glif.is/meetings/2011/tech/140911-DTOX-jvdh.pdf>)

DTOX task force is planning to issue a survey about the basic principles of path finding (e.g., what information should not be imposed). Volunteers to agree on the survey questions were Erik-Jan Bos and Jerry Sobieski (NORDUnet).

7. Dynamic GOLE Services Task Force update

Jerry Sobieski (NORDUnet) talked about the Automated GOLE Pilot Project and the NSI Interoperability Plugfest.

The Automated GOLE project has demonstrated a basic ability to establish inter-domain connections across the GOLE infrastructure using automated systems such as Fenius. The NSI framework has leveraged the experience of Fenius (and other similar efforts) to define a more comprehensive and long term protocol. Since the Hong Kong meeting, the task force has been focused on the deployment of NSI as the common provisioning protocol across the Automated GOLE fabric. NSI is now an OGF standard, and offers a much more comprehensive scalable distributed architecture for automated inter-domain provisioning. A collaborative effort among the NSI Working Group, the DTOX Task Force, and the Automated GOLE project defined a very simple NSI compliant inter-domain ontology for use within the demonstration performed during the meeting. The current focus is on the service state presentation/display tool. The project participants are now looking to SC2011 (November, in Seattle) to demonstrate NSI-CS across real GOLE physical infrastructure.

The NSI Plugfest is a demonstration of progress towards deployable software incorporating the NSI CS protocol and framework. There were 7 NSI-CS v1.0 implementations participating in the live demonstration: OpenNSA (NORUnet), AutoBAHN (GÉANT), OpenDRAC (SURFnet), G-LAMBDA (AIST), G-LAMBDA (KDDI Labs), OSCARS (ESnet), and DynamicKL (KISTI). With the collaborative efforts of the NSI Working Group, the GLIF DTOX Task Force, and the Automated GOLE project, it was possible to take existing technology and adapt it to fit NSI Topology model for the Plugfest. Do note that the Plugfest topology approach was a hack done for the demo. It lacks a number of important features before we can view it as a long term answer. More about the achievements can be found in the presentation.

(<http://www.glif.is/meetings/2011/tech/140911-AutoGOLE-js.pdf>)

8. Deployment of NSI 1.0 in the network

Inder Monga (ESnet) gave an update on the OGF NSI Working Group activity. He briefly summarised the three important concepts in NSI:

- Recursive Framework of “requestor” and “provider” agents that scales over multiple Network Service Agents (NSAs)
- Abstraction of multi-layer physical topology into “service topology”
- “Composable” Services Model

The current standard still has some gaps what concerns the Security Profile, Service Definition, Topology, and NSI Connection-service Gap Analysis. GLIF activity is not about the standardisation but the mapping of use cases to the NSI protocol and the protocol integration with network operations. The Authentication and Authorization framework as well as the NSI Service Verification are issues for further investigation.

(<http://www.glif.is/meetings/2011/tech/140911-NSI-im.pdf>)

9. JIVE's activities

Paul Boven (JIVE) talked about NEXPreS project that gives an e-VLBI use case for NSI implementation. NEXPreS is a three-year EU funded project aimed at further developing e-VLBI services of the European VLBI Network (EVN), with the additional goal of incorporating e-VLBI into every astronomical observation conducted by the EVN. JIVE is ready to build a “requester” or a full NSI client in line with the standard. They are still waiting for the final NSI version although the on-demand path finding feature is not necessary for their application. The only requirement is that the path must be tested in advance and they need path stability and redundancy (maybe on multiple 10G paths).

(<http://www.glif.is/meetings/2011/tech/140911-JIVE-pb.pdf>)

10. Future steps

Gigi Karmous-Edwards (NCSU) concluded the GLIF Technical Working Group meeting and summarised the agreements. Basically, we are hoping to form three new important task forces as agreed to below:

- A new task force has been proposed on the topic of *End-to-End Service Verification*. Steve Wolff (Internet2) and Jerry Sobieski (NORDUnet) volunteered to co-chair the task force. Both Steve and Jerry will send an email to the list articulating the End-to-End Service Verification problem space and the proposed effort of the task force. This email will be sent to the group shortly. If all goes well, the first meeting of this task force will be held in the next winter meeting.
- Another new task force has been proposed on the topic of *Defining GLIF Architecture*. Bill St. Arnaud, Erik-Jan Bos (NORDUnet), and Inder Monga (ESnet) volunteered to co-chair the task force. They will first send out an email to the technical working group list to articulate the exact problem space and the proposed effort of this task force. This email will be sent out shortly. This task force will be critical for moving forward, since it aspires to capture the efforts of the last 10 years of GLIF and more so forge a direction for the next 10 years. If all goes well, the first task force meeting will be held in the next winter technical meeting.
- Inder Monga (ESnet) proposed a task force on the topic of *NSI Implementation*. Inder agreed to write up and circulate a text with the proposal to be discussed on the next coming GLIF Winter meeting.

The GLIF Technical Working Group session was closed on time.

Summary of task forces

Closed	Running	Just proposed
<p><i>Campus Networking</i> Chaired by Ronald van der Pol (SARA)</p>	<p><i>Dynamic GOLE Services</i> Chaired by Jerry Sobieski (NORDUnet)</p>	<p><i>End-to-End Service Verification</i> Co-chaired by Steve Wolff (Internet2) and Jerry Sobieski (NORDUnet)</p>
<p><i>Generic Network Interface Specifications</i> Chaired by Evangelos Chaniotakis (ESnet)</p>	<p><i>Distributed Topology Exchange</i> Chaired by Jeroen van der Ham (UvA)</p>	<p><i>Defining GLIF Architecture</i> Co-chaired by Bill St. Arnaud, Erik-Jan Bos (NORDUnet), and Inder Monga (ESnet)</p>
<p><i>Next Generation GOLE Architecture</i> Chaired by Éric Bernier (CANARIE)</p>		<p><i>NSI Implementation</i> Chaired by Inder Monga (ESnet)</p>
<p><i>Resource Allocation</i> Chaired by Gigi Karmous-Edwards (NCSU)</p>		