



10th Annual Global LambdaGrid Workshop
13-14 October 2010
CERN, Geneva, Switzerland

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

Approximately 50 persons attended the GLIF Technical Working Group session at 10th Annual Global LambdaGrid Workshop, co-chaired by Erik-Jan Bos (SURFnet) and Gigi Karmous-Edwards (NCSU). The secretary was Peter Szegegi (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/2010/tech/bos-welcome.pdf>). After the overview of the meeting agenda, Erik-Jan also announced the next coming GLIF Tech WG Meeting that will be held on 24-25 February 2011, in Hong Kong, China. The meeting will be co-located with APRICOT-APAN.Asia/2011 Conference (21-25 February 2011, <http://www.apricot-apan.asia/>)

2. GOLE and resource updates

The GLIF Technical Working Group (<http://www.glif.is/working-groups/tech/>) opened with updates on current GLIF Open Lightpath Exchanges (GOLEs), the potential new GOLEs in Egypt and Singapore. This was followed by discussions on the ongoing work to develop dynamic lightpath services inside GOLEs, as well the practical applications for this.

- *StarLight*

Joe Mambretti (Northwestern University) gave an update on the current StarLight infrastructure and the StarLight International Consortium on behalf of Linda Winkler and Alan Verlo (<http://www.terena.org/activities/tf-storage/ws10/>).

He introduced the IRNC:ProNet: TransLight/StarLight project funded by NSF. Under this project the Translight/Starlight team will be pursuing application experiments on international networks with 6 different activity areas. These areas are:

- Greenlight International - TL/SL will investigate monitoring for energy consumption in selected international-reach experimental networks.
- SCCSnet - TL/SL will support the Science Cloud Communication Services Network by assisting projects that are developing high-performance communication services tailored for computational clouds used by data-intensive scientific applications.
- CineGrid - TL/SL will help enable the production, use and exchange of very-high-quality digital media, including scientific collaborations and visualizations, over photonic networks.
- High-Performance Digital Media Network (HPDMnet) - TL/SL will work with international partners to develop new services based on high-performance optical transport (i.e., optical multicast) of large-scale data streams, including high-resolution digital media streams, as well as new capabilities for high-volume scientific data.
- iGENI (International GENI) - TL/SL will connect existing StarLight resources with current GENI backbone transport resources.

- SAGE - TL/S will help transition SAGE from a transformative research prototype to a hardened technology that provides production-quality, community-driven open services for visualization and collaboration utilizing shared national and international cyberinfrastructure.

Finally, Joe talked about the StarLight 100 Gbps/Tbps initiatives including CA*net/Ciena/StarLight/iCAIR 100 Gbps test bed implemented in September 2010 and the new NSF award to create a new 100 Gbps exchange at the StarLight facility for data intensive science.

- *HKOEP Hong Kong*

Jiangning Chen (Chinese Academy of Science) gave an update on the HKOEP links, cross-connections, equipment and their involvement in the GLORIAD-Taj program (<http://www.glif.is/meetings/2010/tech/chen-hkoep.pdf>).

She mentioned some important projects such as the network management and measurement activity within GLORIAD, the integrated services (i.e., security, massive data storage, video conference, collaborative environment, accelerate transmission) for ITER, the lightpath provisioning for eVLBI, and the High Performance Transfer Protocol testing for Future Internet Research.

- *Egyptian GOLE*

Greg Cole (GLORIAD) presented a unique public/private approach to establishing the first African GLIF Open Lightpath Exchange in Egypt (slides are not available).

As a first step, GLORIAD and GLIF community brought up new network (Cairo-Chicago) in < 5 days (with Tata and Telecom Egypt). Commitment made at AMCOST to launch major R&E networking initiative in Africa.

Greg also gave a brief update on the Egypt National Research & Education Network (ENSTINET). Egypt geographical location helps it so much to bridge US/Europe with both Africa and Arab Region. Egypt is looking to be African Network Hub.

- *NetherLight*

Gerben van Malenstein (SURFnet) gave an overview on the NetherLight connectivity (<http://www.glif.is/meetings/2010/tech/malenstein-netherlight.pdf>).

He announced the new 100G ready Amsterdam -Geneva DWDM system in production, spanning 1650 km and 22 hops. In 2010, NetherLight established a 40G production link to CERN, replaced the Nortel HDXc platform to Ciena OME6500 Double Decker, implemented the GOLE SLS at NetherLight, and launched the Automated GOLE pilot. The Automated GOLE pilot in NetherLight is based on Open DRAC software with Fenius interface and perfSONAR monitoring on the Force10 platform.

The plans for 2011 include a 40G trans-Atlantic transmission and a 100G clear channel transmission between Amsterdam and Geneva. Next Generation Ethernet will be available as pre-production service as well as Dynamic GOLE in NetherLight.

3. Update on Internet2's ION service

Eric Boyd (Internet2) talked about the Internet2's ION service

(<http://www.glif.is/meetings/2010/tech/boyd-ion.pdf>).

He explained that how Internet2 is planning to extend the IDC-based Dynamic Circuit Networking service. ION is a production service since 2009 managed by the Internet2 NOC

on the Internet2 network backbone. Internet2 ION service transitioned from dedicated network built atop Ciena CoreDirectors to Juniper MX960s (complete as of July 6th, 2010). Internet2, in consultation with partners, plans to evolve working IDC-based service instances over time in keeping with emerging consensus on standards and best practices within the OGF and GLIF. Internet2 is also engaged in a common service definition of a dynamic circuit networking service with DICE partners.

Internet2 would like to continue to extend dynamic circuit networking services (specifically service definitions based on the IDC protocol) towards US regional networks, US campuses, and US-funded international links.

Eric also gave an overview on the NSF funded projects IRIS, DyGIR, and DYNES. DYNES project's aim is to extend Internet2's ION service into regional networks and campuses, based on OSCARS implementation of IDC protocol (developed in partnership with ESnet) in order to support large, long-distance scientific data flows in the LHC, other leading programs in data intensive science (such as LIGO, Virtual Observatory, and other large scale sky surveys), and the broader scientific community. The aggregated goal of the three projects is to ensure that the integrated end-to-end performance monitoring and automatic provisioning of circuits will be "normal part of networking" by 2013.

4. GridARS: Resource Management Framework for Multi-domain Cloud

Tomohiro Kudoh (AIST) talked about GridARS, the Resource Management Framework for Multi-domain Cloud (<http://www.glif.is/meetings/2010/tech/tomohiro-gridars.pdf>).

GridARS is a unified provisioning system of network and compute resources (following the IaaS concept). Its features include the dynamic scheduling and automatic set up of IP addresses, VLAN (host and network), SSH and file system. The monitoring of virtual infrastructure is the essential part of the framework. It provides requester with monitoring information of the provisioned infrastructure (network and CPU) only. The monitoring system interoperates with the provisioning system and applies policy based filtering. At the end of his talk, Tomohiro invited everyone to visit the GridARS demonstration later on.

5. Overview on OGF NSI-WG

Inder Monga (ESnet), co-chair of OGF NSI-WG, gave an overview on Network Services Framework (<http://www.glif.is/meetings/2010/tech/monga-nsi.pdf>).

Inder mentioned that GLIF supported the case to create NSI-WG. Fenius, developed by GLIF GNI-API Task Force participants, can be an intermediate step to move multiple implementations towards the same API. It is important to ensure that the Fenius effort stays in sync with NSI WG and all the GLIF requirements are properly represented for NSI v1 protocol specification.

The most important thing is to put resources in GLIF to help define and create a prototype implementation of the NSI protocol. Focus on learning, documenting and transitioning Fenius to NSI. The next coming big challenge would be a demonstration with NSI protocol at the 11th Annual GLIF 2011.

6. perfSONAR Task Force wrap up and demo introduction

Thomas Tam (CANARIE) wrapped up the perfSONAR Task Force at GLIF (<http://www.glif.is/meetings/2010/tech/tam-perfsonar.pdf>).

The perfSONAR Task Force demonstrated the benefits and potential of using perfSONAR to monitor lightpaths and GOLEs, and GOLE operators agreed on keeping the perfSONAR instances running even though the task force was wrapped up. Its leader, Thomas Tam (CANARIE), was thanked for his continued efforts.

7. GNI API Task Force

Evangelos Chaniotakis (ESnet) presented the GNI API Task Force objectives and the latest achievements since the winter meeting in Salt Lake City (<http://www.glif.is/meetings/2010/tech/chaniotakis-gniapi.pdf>).

Evangelos recapped that several GOLEs already exist and collaborate in the GLIF forum, but until recently, interconnecting has been a manual process initiated by the network administrators at each GOLE. Because of the lack of standards, any automation in the process was only accessible through a proprietary interface. This lack of interoperability has hindered the development and use of virtual circuit services that cross more than a few GOLEs at a time.

The GNI API Task Force's objective for the Geneva meeting was to demonstrate that we can indeed have automated, dynamic GOLEs that can provision virtual circuits with no manual intervention initiated by the end-user through the Fenius common interface. The demonstration was a complete success.

Since Fenius and automated GOLE applications clearly demonstrated the value of interoperability, the next steps will be to help define and develop an open-source implementation of NSI, a standard protocol that will establish native interoperability between the various provisioning software systems such as OSCARS, AutoBAHN, G-Lambda, Open-DRAC, etc.

After the talk there was a discussion related to the future of Fenius:

The basic question was how to continue? Everybody agreed that what have been done so far is powerful. John Vollbrecht - The NSI direction is something that we should force; the two activities should work together. Michal Krsek - The documentation for Fenius is not yet available. Moreover, Fenius is a centralised super-agent that is not scalable and someone still needs to compute the route. Evangelos Chaniotakis - It was not intended so far to be scalable but we don't need to worry about it now only in longer term. Erik-Jan Bos - Open standards are the key. Cees de Laat - UvA is looking at scalability of path finding. Although, over a given scale it is nothing else than an ISDN network at lambda speed. We don't have to think about things that have already been solved. Jerry Sobiewski - We should focus on the general functionality and not on the specific Fenius software. Not just look at Fenius but the issues that we have to address, namely: robustness, reliability, who supports the tools, etc. Erik-Jan Bos - The recommendations should come from OGF NSI and be adopted in Fenius, after that we can think about the standardisation process. Lars Fischer - GLIF is not for standardization. Working code and rough consensus is just fine. The working Fenius code must be translated to NSI language, but there are so many open issues. Gigi Karmous-Edwards - If the Fenius API, which has been implemented as the native DRAC implementation, conforms to the newly developed NSI API then that could be enough for an early standard. We need a successful implementation to move towards a standard. Inder Monga - NSI and Fenius have common principles but GLIF should focus on new network services. Make sure that we innovate! John Vollbrecht - Fenius should become the first early implementation of NSI. Great challenge to make Fenius as an NSI v1.0 working code by the GLIF Tech Winter meeting. Evangelos Chaniotakis - A bit more interoperability towards NSI can be achieved soon.

8. Dynamic GOLE Services Task Force progress and demo preparation

John Vollbrecht (Internet2) gave an update on the Dynamic GOLE Services Task Force's progress (<http://www.glif.is/meetings/2010/tech/vollbrecht-dynamicgole.pdf>).

It was agreed in Salt Lake City that the task force has to create a pilot. This pilot is to be a long running service, up for at least a year after the initial demonstration in Geneva. The pilot creates an “Inter-GOLE Infrastructure” that allows users to schedule a connection between Infrastructure edges. GOLEs participating are NetherLight, MANLAN, StarLight, JGN2, CernLight, NorthernLight, CzechLight, UvA and PSNC (others may join).

John gave an introduction to the demonstration. The demo showed the very recently completed implementation of Automated GOLE pilot, and the capability of using common software interface to control switching in multiple GOLEs. All applications in this demo were on GOLE provider controlled resources.

The pilot implements a common interface for creating connections (NSI WG at OGF is developing a standard for this interface). John suggested that GLIF should agree to implement standard as it becomes available. He is going to report the state of the current GLIF work to NSI at the next OGF meeting.

Finally, John mentioned that he is retiring by the end of November 2010, so the group will need a new coordinator. Lars Fischer said that this task force is really important for GLIF. Jerry Sobieski (NORDUnet) volunteered to lead the Dynamic GOLE Services Task Force in the future. Gigi Karmous-Edwards thanked John Vollbrecht for the hard work.

9. ClearStream demo introduction

Case de Laat (UvA) closed the programme with an introduction to the ClearStream demonstration (<http://www.glif.is/meetings/2010/tech/cdl-demo.pdf>).

ClearStream demonstration was about an end-to-end ultra fast transmission over a wide area 40 Gbps lambda from UvA, Amsterdam to CERN, Geneva. The main goals of the successful demo were to demonstrate single stream single wave performance end to end operating without any obstruction, to break the 10 Gbps barrier, as well as to test a real application on this infrastructure.

The first day was closed with the demonstration session.

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 agenda of the day.

10. Distributed Topology Exchange Task Force discussions

Jeroen van der Ham (UvA) summarised the progress of the Distributed Topology Exchange Task Force (<http://www.glif.is/meetings/2010/tech/vanderham-dtox.pdf>).

Network Description is needed in order to provide an overview of resources, to make path discovery easier, and to do simple problem detection. It has been demonstrated that the automatic inter-domain provisioning is now a reality that makes NDL more important.

The discussion after the talk focused on the principles of the topology exchange issue: *Jeroen van der Ham - Nobody answered to the request to share topology information. We need information to work with. Freek Dijkstra and Jerry Sobieski - We should define the problem first; why we need topology exchange. Actually, a real topology is not needed to implement and test a prototype. Lars Fischer - Finding the best solution on the wrong problem is just a waste of time. Evangelos Chaniotakis - During the Fenius demo sophisticated path finding was not needed. Cees de Laat - Automated generation of NDL files are already solved by UvA. It is a question if GLIF is interested in investing efforts on that.*

Gigi Karmous-Edwards - Can GOLEs agree on using NDL? Evangelos Chaniotakis - Overall policy is missing. Erik-Jan Bos - Who will use the information that is extracted from the UvA tool? It is a chicken and egg problem. These innovations are not used by e.g., OSCARS and DRAC currently! But this work should be continued in order to have this component ready whenever it will be needed. Gigi Karmous-Edwards - Some GOLEs can start using NDL and agree on the labels. More discussions are needed on this.

11. The GN3 Project AutoBAHN Stitching Prototype Development

Andrew Mackarel (HEAnet) and Soren Krum (NORDUnet) talked about the AutoBAHN Stitching Prototype Development (<http://www.glif.is/meetings/2010/tech/mackarel-krum-gn3.pdf>).

Within the lifetime of the GN2 project the stitching framework was defined, responsible for specifying and agreeing technical attributes for inter-domain circuits' definitions. Although this framework is a key point for inter-domain links management by automated system, only a minor part of it was implemented in AutoBAHN at the end of GN2 project. GN3 is about to extend the stitching framework with new logic and attributes, and also to ensure that AutoBAHN implements proper mechanisms for automated links configuration. The GN3 Stitching Framework, data model and stitching engine were discussed in details. The technology stitching engine must be able to handle any number of parameters and to properly evaluate all attributes of the parameters in the path.

A short discussion followed the presentation:

Evangelos Chaniotakis - The deployment of the basic logic is a valuable exercise. Perhaps the business logic in GLIF is not the same. How to address that? Andrew Mackarel - The input parameters have to be analyzed first. Gigi Karmous-Edwards - The basic question is how we know where the resources are. Regarding the architecture, - Cees de Laat - pushing information would not scale well and will fail if people forget to push it. Keep the information where they are, no central place is needed. Distributed and linked system is required. Gigi Karmous-Edwards - The point is to find the resources. Cees de Laat - The European Commission funded NOVI and GEYSERS projects are doing the same. The idea can be extended to GLIF. Erik-Jan Bos - This topic is alien for networkers. The overall aim is to work on models to have one interoperable universal system providing seamless experience for the user.

12. Campus Networking Task Force discussions

Ronald van der Pol (SARA) gave an update on the Campus Networking Task Force (<http://www.glif.is/meetings/2010/tech/vanderpol-campusnet.pdf>).

He reported that there have been little discussions and feedback on the mailing list. Ronald asked if there is interest in continuing this task force. He reminded that the aims of the task force are to investigate use cases, write Best Current Practices documents, and build a community.

The discussion after the talk emphasised the importance of this task force:

Ronald van der Pol - SARA is doing strait staff at on-site and willing to share that knowledge. Evangelos Chaniotakis - Campuses are not just providing the on-site connections nowadays. Valuable people who know each other are there. Lars Fischer - Solving the campus problem is the biggest challenge. We need to educate campus people! Erik-Jan Bos - Campus people will never come to GLIF to consulate with them. Although, we have some experiences with end-to-end connections such as the 40G demo at CERN or the Rio de Janeiro – Eindhoven link setup. We have to share our experiences on that. Peter Szegedi - TERENA has been

running end-to-end workshop series (between 2008-2010) that is addressing the same issue and already built up a community around the workshop series with more than 150 people on the e2e mailing list. It would be interesting to channel in those people to GLIF. Ronald van der Pol proposed to put up a questionnaire about end-site issues that may facilitate the further discussions.

13. Next Generation GOLE Architecture Task Force update

Eric Bernier (CANARIE) summarised the achievements of the Next Generation GOLE Architecture Task Force (<http://www.glif.is/meetings/2010/tech/bernier-nextgen.pdf>). The main objective of the Next Generation GOLE Taskforce is to develop the future requirements and reference architecture for GOLE. Create a blueprint that facilitates equipment replacement at GOLE. The blueprint helps framing what are the functions that are fundamental for a GOLE and which are optional.

Eric talked about the extension to the CANARIE services towards IPv6, Settlement Free Peering and the new ROADM and fibre footprint. North America first 100G demo was done between Ottawa and Chicago facilitated by the new Ciena 100G Adaptive Engine. It was acknowledged that a tremendous job has been done at CANARIE. In order to share all of these practical experiences a White Paper will be compiled and presented during the next GLIF Tech Winter meeting in Hong Kong that will wrap up this task force.

14. Open DRAC update

Gerben van Malenstein (SURFnet) talked about the latest OpenDRAC developments (<http://www.glif.is/meetings/2010/tech/macaulay-openrdac.pdf>).

It was reported that the basic Ethernet/VLAN support of Force10 switches, the Fenius 2.0 inter domain component, the extended static topology and OpenDRAC WSDL 3.0 have recently been completed. The OpenDRAC layer-2 control of NetherLight and CERNLight was demonstrated during GLIF Automated GOLE demo. Eventually, OpenDRAC becomes a true multi-layer controller.

The plans for 2011 include the OGF NSI reference implementation, the NOC integration for DRAC alarms/events, resources, and topology, the federated authentication, authorization and policies, the implementation of Next Generation Ethernet-based on technology outcome of SURFnet tender, a research solution to the last-mile problem of campus edge to researcher/instrument, and a new support team to address technical issues and assist in setting up dynamic end-to-end services.

Gerben mentioned that OpenDRAC can be used by GOLEs where there is no such a tool in place. GEYSERS project may provide the solution on how virtual resources (and non-hardware resources such as repositories) can be included in this dynamic environment.

15. GEYSERS: Generalised Architecture for Dynamic Infrastructure Services

Pascale Vicat-Blanc (Iyatiss) gave an introduction to the GEYSERS project (<http://www.glif.is/meetings/2010/tech/vicatblanc-geysers.pdf>).

The Internet as it stands today doesn't necessarily provide the unlimited, on-demand scalability that cloud computing providers like to promise. Dynamic networks connecting user premises are needed to improve capacity, reliability and reduce costs. GEYSERS's (Generalized Architecture for Dynamic Infrastructure Services) goal is to bring optical networks to the cloud. This goal will be reached by a novel architecture and tools for the composition of virtual infrastructures, an enhanced network control plane, a definition of new business roles that impact on legacy infrastructure providers, network & IT operators and

application providers and an international test bed for prototyping and validating the model and tools.

It was felt that GLIF should carefully listen to non-ICT communities dealing with huge and ever growing amount of data. Interoperability of multi-domain networks is the key; a single solution for all domains doesn't work. Participants agreed that GEYSERS project is important for GLIF.

16. GLIF: What's next?

At the end of the day Erik-Jan Bos (SURFnet) summarised the Technical Working Group's conclusions (<http://www.glif.is/meetings/2010/tech/ejb-nextsteps.pdf>).

He quoted Kees Neggers saying "providing users with seamless access to a global e-Infrastructure Ecosystem" is the key. GLIF should reach out to other communities, like the Grid community (computing and storage) and scientific domains.

Summarising the task force discussions Erik-Jan said:

- *Everyone should contribute to the Distributed Topology Exchange Task Force as Jeroen van der Ham (UvA) requested.*
- *The lessons have been learnt from the Campus Networking Task Force. Ronald van der Pol (SARA) will put together a survey in order to collect more information from end-site and campus network operators.*
- *The Research Allocation Task Force should work together with projects such as GEYSERS.*
- *Next Generation GOLE Architecture Task Force will wrap up with a white paper presented on the next coming GLIF Tech Winter meeting.*
- *perfSONAR Task Force has been wrapped up with the Geneva demo but will be available for other task forces in the future.*
- *The leadership of the Dynamic GOLE Services Task Force has been taken over by Jerry Sobieski (NORDUnet). John Vollbrecht was thanked for his hard work.*
- *GNI API Task Force should work in close cooperation with OGF NSI WG. The overall objective is to move to a standard as fast as possible.*

Finally, Erik-Jan reminded again that the next coming GLIF Tech WG Meeting that will be held on 24-25 February 2011, in Hong Kong, China, co-located with APRICOT-APAN.Asia/2011. The meeting registration is already available here (<http://www.glif.is/meetings/2011/winter/>).