



GLIF Technical Working Group
20th Meeting, 3-4 October 2013
Nanyang Technological University

This meeting was held as a session during the 13th Annual Global LambdaGrid Workshop.

1. Welcome

Lars Fischer welcomed everyone to the meeting and outlined the aims of the meeting. He reiterated that GLIF was not really a network or a facility, but was about sharing ideas and developing contacts in order to reach consensus on architectures research networking.

2. Approval of minutes

There were no comments on the minutes of the last meeting, and were therefore taken to be approved.

3. GOLE updates

3.1 KRLight

Hyeak Ro Lee provided an update on KRLight (see <http://www.glif.is/meetings/2013/tech/lee-krlight.pdf>). In 2013, they had established a new 10 Gb/s link between Daejeon and Chicago, as well as upgraded their backup link between Daejeon and Hong Kong to 2.5 Gb/s. KREONET had also been upgraded to a dark fibre infrastructure that currently provided multiple connections of between 10 and 100 Gb/s, on which NSIv2.0 and OSCARS had been deployed to offer dynamic circuit provisioning.

3.2 NetherLight

Gerben van Malenstein provided an update on NetherLight (see <http://www.glif.is/meetings/2013/tech/vanmalenstein-netherlight.pdf>). 2 x 100 Gb/s to NORDUnet via London had been in service August 2012, a dark cross-border fibre link had been established between Maastricht (The Netherlands) and Hasselt (Belgium), and a second STM-4 to Egypt had been live since January 2013. The Trans-Atlantic 100 Gb/s link had been demonstrated at TNC 2013, whilst a second 100 Gb/s link to CERN would go live soon.

NetherLight was primarily based at the Amsterdam Science Park, but a Ciena 5150 (to be upgraded to a Ciena 5160 in 2014) had also been added to second location in Amsterdam to offer redundancy.

3.3 StarLight

Joe Mambretti reported that StarLight had established several initiatives aimed at supporting dynamic services at 100 Gb/s and higher (see <http://www.glif.is/meetings/2013/tech/mambretti-starlight.pdf>). StarWave was a new multi-100 Gb/s exchange facility within the StarLight communications facility, and was being used to support several SC13 demonstrations. It was also a key element of the ESnet 100 Gb/s networking initiative, as well as the CA*net/Ciena/StarLight/iCAIR 100 Gb/s testbed.

3.4 SingLight

Francis Lee provided an overview of SingLight (see <http://www.glif.is/meetings/2013/tech/lee-singlight.pdf>). This was a relatively new GOLE and had connections of 1 Gb/s to NICT (Japan), GLORIAD and Internet2, as well as a 622 Mb/s link to AARNet, and a 90 Mb/s link to the TEIN4 network.

The SingAREN-Lightwave Internet Exchange (SLIX) was a 5-year project involving SingAREN, A*STAR, NTU and NUS that was co-funded by NRF. This aimed to upgrade the local infrastructure to a leased fibre core that would support multiple 100 Gb/s lightpaths to member institutes.

3.5 MAN LAN / WIX

Dale Finkelson provided an update on MAN LAN and WIX (see <http://www.glif.is/meetings/2013/tech/finkelson-manlan.pdf>). Both MAN LAN and WIX currently had 100 Gb/s connections to Internet2 via AL2S and ESnet, with WIX also having a 100 Gb/s connection to MAX. The ANA 100 Gb/s circuit terminated at MAN LAN.

Both static and dynamic circuits were provisioned with NSI and OSCARS, although OpenFlow was not currently supported. perfSONAR boxes were also available for testing circuits in and out of MAN LAN and WIX.

4. AutoGOLE Task Force

Gerben van Malenstein reported on the work of the AutoGOLE Task Force (see <http://www.glif.is/meetings/2013/tech/vanmalenstein-autogole.pdf>). The goals for 2013 had been to implement NSI-CSv2 and topology exchange, and there would be a demonstration of this later in the day.

A number of lessons had been learned in that topologies were difficult to update and verify, and there were often mismatches between the data and control planes. Implementing the protocol in the real world showed that an authentication and authorisation infrastructure for end users was still required, as well that real-time resource availability for efficient path finding.

The next steps were to bring NSI-CS v2 into production and to implement other NSI framework components such as monitoring and discovery services. There were also plans to introduce an E-LAN switching capability and add support for OpenFlow and SDN-T.

Lars Fischer welcomed the development of an important milestone, but felt the task force now needed to focus on the elements that were needed to bring the AutoGOLE system into production. It should consider what was ready and what still needed to be improved to allow real usage.

5. Distributed Topology Exchange Task Force

Jeroen van der Ham presented the final report of the Distributed Topology Task Force (see <http://www.glif.is/meetings/2013/tech/vdham-dtox.pdf>). This work had started back in 2005 and had introduced some innovative concepts that had culminated in the NML Base Schema

currently going through the OGF standardisation process. This could be further extended to other computing resources to describe complete infrastructures, but that was something for the future.

Cees de Laat then announced that Jeroen would be moving to take-up a new position at the University of Amsterdam after eight years involvement for GLIF. He moved for a motion of thanks which was unanimously agreed by the participants on behalf of the GLIF community.

6. NSI Implementation Task Force

John MacAuley reported on the work of the NSI Implementation Task Force (see <http://www.glif.is/meetings/2013/tech/monga-nsi.pdf>).

They had held a 2-day meeting prior to the workshop to figure out some of the main issues, and had agreed to use digital certificates and SSL authentication between NSA servers. This required manual provisioning of certificates on client and server sides, which mean the current authorization model was based on transient trust. However, this approach would be increasingly difficult to scale for a larger number of NSI clients, so the need for more client friendly authentication mechanisms such as HTTP BASIC and OAuth2 had been identified.

Topology distribution took a simple REST over HTTP(S) approach, with extensibility to add further features. Nevertheless, there was a large operational overhead associated with maintaining accurate topology files which meant more automated topology discovery was necessary.

It been decided to change the service definition by removing service specifics from NSI CS v2 and adding three basic building blocks for Point-to-Point, Ethernet Transport, and Ethernet VLAN Transport services. Additional services could then be defined as needed.

7. Performance Verification Task Force

Jerry Sobieski provided an update on the Performance Verification Task Force (see <http://www.glif.is/meetings/2013/tech/sobieski-pv.pdf>). Most of the activity had focused on attempting to define the scope of what should be verified, and how link performance could be independently verified without perturbing the flows. In addition, the group was considering where monitoring points should be located, and how measurements should be correlated.

8. Testbeds as a Service: A view into a new GN3plus Service

Jerry Sobieski gave a presentation on the GN3plus Service Activity 2 (see <http://www.glif.is/meetings/2013/tech/sobieski-sa2.pdf>). This aimed to deliver two production-quality services that included a dynamic packet testbed over the GN3plus network, and a photonic testbed between a limited set of sites in Europe. It would also leverage other GN3plus services such as bandwidth-on-demand provisioned by AutoBAHN, TaaS, and authentication and authorisation infrastructures.

The plan was to involve several NRENs in a multi-domain interoperability trial by mid-2014, with a view to reaching a consensus before the start of GN4.

9. OpenNaaS

Sergi Figuerola gave a presentation on OpenNaaS (see <http://www.glif.is/meetings/2013/tech/figuerola-opennaas.pdf>). This provided an open source framework for virtualising and abstracting network resources, managing physical and virtual networks, support for slicing and aggregation, and deploying dynamic infrastructures. It allowed building of different types of network service, with HEAnet aiming to deploy the first pilot infrastructure by mid-November 2013.

The OpenNaaS founders were i2CAT, DeIC, HEAnet, Juniper Networks, NORDUnet, Telefonica and the University of Bristol. The aim was to develop an open-source consortium to support and sustain further development of the framework, allowing different stakeholders to benefit from customised solutions.

10. Network Virtualisation

Eric Boyd gave a presentation about enabling innovation through use of network virtualisation and integration of computing and storage resources (see <http://www.glif.is/meetings/2013/tech/boyd-virtualization.pdf>). The Internet2 Advanced Layer2 Service Deployment provides a 100 Gb/s platform for introducing virtualised networking by decoupling the control plane from the data plane, enabling multiple virtual control planes on a common physical network. The question was whether a common GLIF approach was required to this, and how the computing and storage community could be integrated with it.

Internet2 in partnership with Indiana University was been developing and deploying an architecture to provide network multi-tenancy at Layer 2 and 3, and to enforce non-overlapping Layer 2 tag-based flowspace. They had initially looked at FlowVisor, but this did not support VLAN Tag range-based policy, and with 28 switches with 10 ports and 4096 policies per port, it was unable to load a sufficient number of rules in an acceptable time. It was therefore decided to develop FlowSpace firewall to keep an OpenFlow application within a defined slice. This was built on FloodLight in collaboration with the GlobalNOC Software Engineering.

11. Global Network Architecture Developments

Erik-Jan Bos reported on the Global Network Architecture activity (see <http://www.glif.is/meetings/2013/tech/bos-gna.pdf>) that had been started by Internet2, CERNET and Janet under the auspices of the CEO Forum to organise strategic dialogue on coordinating intercontinental research and education links. Two groups had been established in May 2013 to discuss general aims and architecture, with a view producing an initial discussion paper in December 2013.

12. Connecting Researchers with Clouds

It was felt important to consider how GLIF could connect big research users to cloud services, and to this end the meeting heard from the Helix Nebula (see [http://www.glif.is/meetings/2013\[KM2\]/tech/jenkins-helixnebula.pdf](http://www.glif.is/meetings/2013[KM2]/tech/jenkins-helixnebula.pdf)) and EGI FedCloud (see <http://www.glif.is/meetings/2013/tech/wallom-fedcloud.pdf>) initiatives, as well as from Boudewijn Lelieveldt on the computing and networking challenges in population imaging studies

13. Date of next meeting

Thanks were extended to SingAREN and Nanyang Technological University for hosting the meeting.

It is proposed to hold the next meeting on 19-20 March 2014 in Atlanta, United States. This would be in conjunction with the 19th GENI Engineering Conference (GEC 19), and it was hoped could be confirmed in due course.