8th Annual Global LambdaGrid Workshop
1-2 October 2008
Bell Harbor Conference Center, Seattle, United States

Technical and Control Working Group – Opening Session

Approximately 60 persons attended the meeting chaired by Erik-Jan Bos (SURFnet) & Gigi Karmous-Edwards (MCNC). The secretary was Peter Szegedi (TERENA).

1. Welcome and Apologies

The first joint Technical and Control Plane Working Group meeting was opened by Erik-Jan Bos (SURFnet). He announced that the two former working groups had been officially merged upon the decision of both groups during the winter meetings in Honolulu. The main motivation was the common interest of the participants of the groups in the activities of the other.

The new mailing list arrangements were also announced. For practical reasons, the proposal was to re-use the existing technical mailing list <tech@glif.is> and simply close the control plane mailing list <controlplane@glif.is> (all of the control plane list members will be added to the technical list).

Action 20081001-1: TERENA to migrate the control plane mailing list to the technical mailing list.

2. Minutes of last meeting and approval of agenda

There were no comments on the minutes of the last meetings of either the Technical or Control Plane Working Groups. The minutes could therefore be taken to be approved.

The meeting approved the proposed agenda.

3. GOLE updates

AARNet

Ivan Phillips (AARNet) introduced the network footprint of AARNet, briefly covering the DWDM, TDM and IP layer architecture (see http://www.glif.is/meetings/2008/tech/phillips-aarnet.pdf). It was pointed out that the fibre build business is booming and the number of available lightpaths is increasing in Australia. AARNet is looking to deploy dynamic optical control plane solutions (e.g. DCN/DRAGON), 10G IP access and national VPLS network in the near future.

CANARIE

Eric Bernier, the new CTO of CANARIE, mentioned the personnel changes at CANARIE, that included a new CEO (Guy Bujold) and Chief Engineer (Thomas Tam). The CANARIE network was briefly introduced using Nortel's ROADMs (see http://www.glif.is/meetings/

2008/tech/bernier-canarie.pdf). Eric emphasised the opportunity of wavelength swap businesses in Canada. At the TDM level, the old Cisco ONSs were replaced with four new Nortel OMEs with better management capabilities. The HPDMnet and perfSONAR demonstrations as well as a new link between Seattle and KRLight were mentioned.

CzechLight

Stanislav Sima (CESNET) provided an update on CzechLight (see http://www.glif.is/meetings/2008/tech/sima-czechlight.pdf), with new lambdas available throughout the Czech Republic and with the metropolitan DWDM in Prague. They have 2 x 10G lambdas to NetherLight and StarLight. Beside the lambdas there are VC connections to the CESNET2 network, universities and other experimental facilities.

Stanislav was asked about CzechLight's plans within respect to dynamic lightpaths. He answered that only a few lambdas exist in the network, so has not yet become an issue and everything is manually configured at present.

KRLight

Dongkyun Kim (KISTI) provided an update on KRLight (see http://www.glif.is/meetings/2008/tech/kim-krlight.pdf). This has a total of nine connections to three other GOLEs, three NRENs, and three other institutes. They had started a migration from Cisco ONS to Nortel OME6500. Three new lambdas were also mentioned; the Norway-Korea medical lightpath, the USA-Korea VisualCasting lightpath and the perfSONAR demonstration lightpath. A brand new lightpath will be established in October 2008 through Japan-Korea-Taiwan-USA. KRLignt has started to share network resource information with GOLE partners.

MAN LAN

Christian Todorov (Internet2) provided an update on MAN LAN (see http://www.glif.is/meetings/2008/tech/todorov-manlan.pdf). The Nortel HDXc had been upgraded to 5.0 code, there was a new connection Ankabut (UAE), and they had implemented perfSONAR for circuit monitoring. In future, it is planned that GÉANT will remove their NYC router and connect their OC192 link through a 10GE on the MAN LAN switch, with IP peerings occuring on their router in Amsterdam. In addition, MAN LAN will investigate adding dynamic or user controlled capability to the network. An important step is to fully enable perfSONAR for monitoring dynamic circuits established through MAN LAN.

NetherLight

Wouter Huisman (SURFnet) introduced NetherLight's SDH/Ethernet hybrid node architecture (see http://www.glif.is/meetings/2008/tech/huisman-netherlight.pdf). The details of the global connectivity tender were briefly discussed with optical 40 Gb/s connections between Amsterdam–MAN LAN and Amsterdam–StarLight replacing the Global Crossing links. NLR, NSF/IRNC and NORDUnet was also participating in this tender. Five on-going projects were also discussed, as well as a Service Template to easily track services/projects by the SURFnet NOC.

TaiwanLight

Eugene Yeh (NCHC) introduced the fault-tolerant TWAREN network that had two STM-16 connections to the USA (see http://www.glif.is/meetings/2008/tech/yeh-taiwanlight.pdf). TaiwanLight is the international part of TWAREN. Switches are used to connect lightpaths

and MPLS-VPN service is provided on top of these. One of the latest achievements was the Taiwan-Czech Republic link for live medical video streaming that was demonstrated during the CESNET'08 conference. The OptIPuter video streaming project will come in March next year, and they also plan to switch from an SDH based platform to pure DWDM transport architecture.

Pacific Wave

Geoff Lakeman (University of Washington) showed the three distributed sites of the PacificWave network in the USA. The Nortel HDXc upgrades and the new Internet2 DCN connections were also announced. In the future they plan to have a NLR upgrade on the west coast.

Geoff was asked about the type of the DCN connections. He answered that the connection at the Los Angeles PoP is a combination of SDH (L1) and Ethernet (L2). In Seattle, only L2 connection is available. All the connections require manual configuration of the boxes at present.

LHCNet

Artur Barczyk (CERN) introduced the network topology of the LHCNet with path diversity among 4 PoPs (CERN, StarLight, MANLAN, NetherLight). The LHCnet introduced semi-automated (scripted) provisioning in the Ciena domain and also deployed the DCN software suite. They recently joined the DICE group and they have plans to increase the network capacities. The current circuits are only semi-dynamic, modified by management system. The automated mash protection capability of the network is based on the in-built features of the Ciena platform.

4. perfSONAR Circuit Monitoring

Jeff Boote and Thomas Tam (CANARIE) provided a brief overview of PerfSONAR that would feature in the evening demonstration session (see http://www.glif.is/meetings/2008/tech/tam-perfsonar.pdf). The main aim of PerfSONAR is the monitoring of circuit status.

Jeff introduced the distributed status architecture of PerfSONAR, where archives are implemented, and the web interface pulls data from the archives at each domain controller. Thomas showed the PerfSONAR lightpath monitoring details and showed what happened if a fault was injected into a lightpath.

Task Forces Sessions

The meeting then broke-out into separate sessions for each of the task forces.

GNI Specifications Task Force

Evangelos Chaniotakis (ESnet), as leader of task force, welcomed the participants. It was the first meeting after the initial preparation in Hawaii, and he provided an overview of the aims of the task force (see http://www.glif.is/meetings/2008/taskforces/chaniotakis-intro.pdf). The main objectives were short-term interoperability between systems, and to relate experiences to the OGF NSI Working Group. The group should also develop an API and framework for

interoperability, and the first version of the framework has just been committed to SVN (see http://gusi.inocybe.ca).

There were also several APIs available (e.g. for AutoBAHN, Phosphorus and NSI,), but because of different implementation approaches, there was no way to select one as a unified API. This was why the task force proposed to develop a common API called the GLIF Unified User Interface (GUSI), which could work with a distributed system, but did not assume anything about the data plane, AAI, or topology exchange issues.

For future work, an iterative approach was suggested, starting with development of the framework and resource managers, testing if that worked, and then learning from these experiences. An interoperability demo at the GLIF winter meeting would be good, as well as a semi-formalised API specification, and a 'lessons learned' document.

Action 20081001-2: GNI Specifications Task Force to prepare API interoperability demo for the GLIF winter meeting.

Gigi Karmous-Edwards (MCNC) summarised that the initial goal of the task force was to ensure interoperability between current and future APIs (see http://www.glif.is/meetings/2008/taskforces/bos-gusi.pdf). The community had made some progress but work was not as fast as it could be, mainly because of its voluntary basis. The proposal was established a GLIF-funded open source development project which would be led by the Technical Working Group Co-Chairs under the control of a Steering Committee which would consist of the current API developers. This would be discussed further during the final session of the Technical and Control Plane Working Group.

Afrodite Sevasti (GRNET) then discussed the AutoBAHN tool which had been developed by the GN2 community for circuit and lightpath provisioning (see http://www.glif.is/meetings/2008/taskforces/sevasti-autobahn.pdf). AutoBAHN is a research activity for engineering, automating and streamlining the inter-domain setup of guaranteed capacity (Gb/s) end-to-end paths across multiple heterogeneous domains. The tool has been developed by the GN2 community for circuit and lightpath provisioning.

The main challenge was to find a common inter-domain control solution for the various intra-domain tools in a heterogeneous (technical and administrative) environment. That is why the tool had two main parts: a generic inter-domain part, and a technology specific intra-domain part. The southband interface is now available for the European NRENs to implement and adopt. The AuthoBAHN is IDC-compatible and will become a production level service in GN3.

John Vollbrecht (Internet2) provided an outline of the IDC protocol and infrastructure (see http://www.glif.is/meetings/2008/taskforces/vollbrecht-idc.pdf). The DCN Software Suite developed by Internet2, ESnet, ISI-East and Mid Atlantic Crossroads was now up to version 0.3.1, and supported GMPLS for intra-domain, and IDC for inter-domain control. Version 0.4 would support VLAN translation at network edges (if the hardware supported it), improves internal state handling, and includes a notification broker to interface between IDCs and other services.The IDC is being further extended with an additional authentication and authorisation infrastructure, and a SONET to Ethernet adaptation between domains, and the protocol needs to be standardised at OGF in the near future.

NetherLight and MAN LAN said they were considering creating Dynamic GOLEs using GUSI or IDC. John expressed this seemed a good goal for GLIF, but the future depended on developing global interoperation standards.

Michel Savoie (CRC Canada) discussed the Phosphorus project and provided an overview of the Harmony system (see http://www.glif.is/meetings/2008/taskforces/savoie-harmony.pdf). Phosphorus addresses some of the key technical challenges in enabling on-demand e2e network services across multiple, heterogeneous domains. It has been demonstrated on a testbed involving European NRENs, GÉANT2, Cross Border Dark Fibres and GLIF.

Harmony (produced by Work Package 1 of Phosphorus) was an inter/multi-domain path provisioning architecture/system where both users and Grid applications can book paths and network resources over heterogeneous domains in advance. It has a distributed, peer-to-peer architecture with integrated AAI capabilities, and G2MPLS support. The plan is that the Harmony tool has translators to G2MPLS, Internet2 IDC and AutoBAHN, and possibly also to GUSI.

Although not part of the original agenda, Freek Dijkstra (University of Amsterdam) gave a short talk on the achievements of the OGF Network Markup Language. The main goal is to define network topologies, and there were several parallel activities. The NDL was developed by UvA to provide generic descriptions with as few technology dependencies as possible; the cNIS was developed by DANTE for IP and Layer 2 topologies only; whilst the NMWG schema was developed by OGF and used by PerfSONAR. There had been a lot of progress, and standardisation was really needed in this area.

SLA Task Force

This meeting was cancelled, although the introductory slides are available on the GLIF website (see http://www.glif.is/meetings/2008/taskforces/todorov-sla.pdf).

Global Identifier Task Force

It was decided to combine this session with that of the GNI Specification Task Force.

GNI Specifications/Security & Authorisation/Global Identifier Task Forces

Erik-Jan opened the session and announced the separate parallel sessions originally planned had been merged upon the request of those task forces. It therefore meant major changes to the agenda.

Evangelos Chaniotakis (ESnet) summarise the overall architecture of the GUSI. He explained the message exchanges between GUSI instances, as well as the message types. As previous discussed, it had been agreed that further discussions were needed on this issue, and a semi-formalised API specification, and a ‘lessons learned’ document should be drafted.

Action 20081001-3: Evangelos Chaniotakis to start drafting API specification and ‘lessons learned’ document.

Mathieu Lemay (Inocybe Technologies) gave a presentation via videoconference about the Service Implementation System in GUSI (see http://www.glif.is/meetings/2008/taskforces/

lemay-gusi.pdf). The GUSI goals can be summarised as providing an open source implementation of interfaces like the Network Service Interface (NSI) or Grid Network Interface (GNI); having a common code base to accelerate developments; avoid having each team implement the same service and interface with different security mechanisms; accelerate new services and adoption of standards in the community; provide a graphical interface to GLIF services.

Mathieu then summarised the GUSI architecture, technologies and security features (see http://gusi.inocybe.ca/). The project is being developed using a variant of the OpenUP software development process. The project is still in the elaboration phase, so any feedback on the architecture, API and schemas or additional ideas are welcomed. The construction phase should begin shortly and the goal is to have a running prototype by the next GLIF winter meeting.

Action 20081001-4: Technical Working Group to provide feedback on the GUSI architecture.

Mathieu was asked about the GUSI services. The GUSI has a standalone service set (not just a subset of the integrated services) but allows the selection ofintegrated services (e.g. Harmony, AutoBAHN). It was also mentioned that the translation service is stateless in GUSI. At present, it was just a configuration tool without a graphical interface, and in practical terms a recommended piece of open source code rather than an interface. Agreement is still needed on the security aspects, because different security mechanisms are used in the integrated services and there is no guarantee for that. Contribution to the source code writing was welcomed (Java and UML based). GUSI is not an obligation just for convenience, if you implement GNI/NSI there is no need for GUSI.

ACTION 20081001-5: Technical Working Group to contribute to the GUSI source code writing

Ronald van der Pol (SARA) presented the Global Identifier work (see http://www.glif.is/meetings/2008/tech/vdpol-globalid.pdf). The main challenge was that the operators have to identify the lightpaths, and the current practice of using local domain identifiers in trouble tickets was not adequate (note that global IDs do not replace local IDs though). The preferred global identifier format contains a global part and a local part, which means that no centralised registry is needed. The maximum length of the identifier should be agreed and a well-defined character set should be used. There are various naming schemes: the DANTE naming scheme is used in the LHCOPN and DEISA networks; the Internet2 naming scheme is used in Internet2 (the identifiers are called Global Resource Identifiers); whilst uniform Resource Names (URNs) are defined in RFC 2141.

A discussion about the semantics of global identifiers has been started. The Internet2 and URN like naming scheme seem to get the most support. During provisioning, we need to start following the Sourcing Organisation model.

Finally, Tom Lehman (USC/ISI) discussed URN identifiers. A basic principle was that identifiers should not include attributes, because small changes in these would mean a change of identifier. Two proposals were discussed: urn:glif and urn:ogf. On one hand, GLIF is not a standardisation body and did not maintain the top level domain. On the other hand OGF was a standardisation body, but longer names would need to be applied because other organisations were also using urn:ogf.

The group tended to favour urn:ogf, but the discussion will continue via the mailing list. It was agreed the Global Identifer Task Force should agree on what to use before the GLIF winter meeting.

Action 20081001-6: Global Identifier Task Force to agree on URN policy before the GLIF winter meeting.

Technical and Control Working Group – Closing Session

5. OGF-NSI Working Group update

Guy Roberts (DANTE) gave a talk about the OGF Network Service Interface (see http://www.glif.is/meetings/2008/tech/roberts-ogfnsi.pdf). There is increasing interest in allowing Grids to request dynamic circuits automatically, but standard mechanisms are needed for describing and publishing network descriptions, performance characteristics, and to enable dynamic control.

Two meetings was organised at OGF23: a DICE-initiated BoF to standardise network-network control plane; and a OGF GHPN-initiated BoF to standardise Grid middleware requests to networks for circuits. It had been agreed to merge the two groups into a new working group called NSI, the structure and deliverables of which were presented. It was pointed out that GLIF should provide input to the OGF, as GLIF was not a standardisation body.

6. MAX: VLAN scheme proposal

Matt Siniscal (Mid Atlantic Crossroads) presented a proposal for a VLAN scheme (see http://www.glif.is/meetings/2008/tech/siniscal-max.pdf). This started with a short MAX update describing the Ethernet layer topology. A cheap Virtual Label Switch Router implementation was also presented, as well as a successful demonstration of a 500 Mb/s connection being set up.

The proposed VLAN scheme called for a static VLAN range among GLIF members. The main motivation behind this was that dynamic VLAN provisioning is not available on all edge switches, and it therefore requires VLAN translation or mutually agreed VLANs for all switches within the circuit path. As an example, one block has already been agreed upon by Awave, Starlight and Netherlight (VLANs 451-500), although up to 50 additional VLANs may be need in future. If a switch in the circuit path already has a VLAN in use, the group provisioning can skip to the next one.

The sooner VLAN ranges are agreed upon; the less conflict between tags will arise. However, some concern was expressed by the group as it is difficult to identify continuous ranges (practical problem) and it causes restrictions from a GOLE's perspective.

Action 20081001-7: Erik-Jan Bos and Gigi Karmous-Edwards to more closely evaluate the VLAN scheme proposal.

7. PBB (802.1ah) in NetherLight

Wouter Huisman (SURFnet) discussed their experiences of the PBB (802.1ah) deployment at NetherLight (see http://www.glif.is/meetings/2008/tech/huisman-802.1ah.pdf). The

NetherLight network architecture is based on SDH/Ethernet lightpaths using a Nortel platform, but the benefits of the Ethernet transport layer are now clear.

There are currently two Ethernet based services: tagged VLANs and untagged VLANs. In the case of tagged service, the limitations of 802.1Q are well known, with the alternatives being 802.1ad (QinQ) and 802.1ah (Provider Backbone Bridging).

PBB is more scalable and provides clear traffic separation. Its implementation is possible in parallel with the QinQ which has solved the basic interoperability issues. However, there might be some interoperability problems with PBB since only proprietary implementations are available. Foundry and Nortel have already implemented PBB, Cisco will implement in the future and Juniper does not plan it yet.

The group agreed there was increasing interest in PBB. For example, all JANET TDM circuits will be replaced with EoMPLS links in the near future.

8. GLIF Evolution and Roadmap

Erik-Jan Bos (SURFnet) and Gigi Karmous-Edwards (MCNC) summarised the grand technical vision of GLIF (see http://www.glif.is/meetings/2008/tech/roadmap.pdf).

Erik-Jan said the trends showed the fibre market to be increasing, with technology evolving towards 100 Gb/s and beyond. Not so long ago, EoSDH was the preferred solution, but now EoMPLS seemed to be the most promising transport option. The control of dynamic lightpaths will also be highly important in future.

Gigi pointed out that GLIF was now in the Phase 2, and needed to accelerate the work. The task forces were currently voluntary, but one option might be for partners to make firm commitments of money and/or effort, which would allow open source development projects to be commenced.

The group discussed this proposal, but the main question was who would manage such technical projects because more coordination was needed than in the case of a task force. The IETF had found it difficult to move money from one organisation (providing money) to another (providing man-months), and some people were not convinced that was a good way forward for GLIF. However, the counter argument was that the goal of GLIF was to establish partnerships to provide a global facility, and software was just the next step in undertaking this. Nevertheless, it was agreed the discussion should be continued on the mailing list.

Action 20081001-8: Technical Working Group to continue the discussion about the open source development project proposal on the mailing list.

Erik-Jan moved to discussing the task forces. He proposed that the GNI, Global Identifier and Security Task Forces should continue. In addition, whilst the perfSONAR Task Force was originally established to set-up a demonstration in Seattle, there were a number of unresolved issues for which it would be useful to keep running.

It was agreed the SLS Task Force should be dissolved, but it was proposed to establish a replacement task force on dynamic GLIF services. This would be discussed further on the mailing list.

Action 20081001-9: Erik-Jan Bos and Gigi Karmous-Edwards to further discuss task force on Dynamic GLIF Services on the mailing list.

6. Date and venue of next meeting

Erik-Jan said there were two options for the next meeting: on 4-5 February 2009 in College Station, Texas, USA (in conjunction with the Internet2/ESCC Joint Techs Workshop), or on 2-6 March 2009 in Catania, Italy (in conjunction with OGF25).

The advantages of co-locating with the Internet2 meeting was that a provisional offer had already been received and no registration fee would be necessary. By contrast, the proposal to co-locate with OGF had been made at very short notice, there was no offer from OGF yet, and registration fees would be quite high. In addition, the dates of OGF25 directly clashed with those of the APAN 27 meetings.

The group felt that meeting in Catania would offer a good opportunity for Europeans to attend GLIF Working Group meetings, and for OGF people to participate in the discussions. The GLIF Secretariat was therefore asked to negotiate with the OGF Programme Committee to see whether a GLIF Working Group meeting could be scheduled during OGF25.

Action 20081001-10: GLIF Secretariat to speak with OGF Programme Committee about possibility of scheduling of a GLIF Working Group meeting during OGF25.

Open Actions

20081001-1 TERENA to migrate the control plane mailing list to the technical mailing list.

20081001-2 GNI Specifications Task Force to prepare API interoperability demo for the GLIF winter meeting.

20081001-3 Evangelos Chaniotakis to start drafting API specification and 'lessons learned' document.

20081001-4 Technical Working Group to provide feedback on the GUSI architecture.

20081001-5 Technical Working Group to contribute to the GUSI source code writing

20081001-5 Technical Working Group to contribute to the GUSI source code writing

20081001-6 Global Identifier Task Force to agree on URN policy before the GLIF winter meeting.

20081001-7 Erik-Jan Bos and Gigi Karmous-Edwards to more closely evaluate the VLAN scheme proposal.

20081001-8 Technical Working Group to continue the discussion about the open source development project proposal on the mailing list.

20081001-9 Erik-Jan Bos and Gigi Karmous-Edwards to further discuss task force on Dynamic GLIF Services on the mailing list.

20081001-10 GLIF Secretariat to speak with OGF Programme Committee about possibility of scheduling of a GLIF Working Group meeting during OGF25.