

GLIF North America (GLIF-NA) Community Meeting
Tuesday, October 12, 2010 (4:00 – 7:00 pm)
Room 3-003, Building 13
CERN, Geneva, Switzerland

DRAFT Minutes May 26, 2011

Attendees

Heidi Alvarez	AMPATH/AmLight/FIU	heidi@fiu.edu
Artur Barczyk	Caltech/ USLHCNET	artur.barczyk@cern.ch
Eric Boyd	Internet2	eboyd@internet2.edu
Jacqueline Brown	Pacific Wave	jborwn@ms.uw.edu
Maxine Brown	UIC/StarLight	maxine@uic.edu
Evangelos Chaniotakis	ESnet	haniotak@es.net
Greg Cole	GLORIAD	gcole@gloriad.org
Steve Cotter	ESnet	steve@es.net
Julio Ibarra	AMPATH/AmLight/FIU	julio@fiu.edu
Joe Mambretti	iCAIR/NU/StarLight/MREN	j-mambretti@northwestern.edu
Grant Miller	NCO/NITRD	millier@nitrd.gov
Dave Reese	CENIC	dave@cenic.org
Jerry Sobieski	NORDUnet	jerry@nordu.net
Bill St. Arnaud	(self)	bill.st.arnaud@gmail.com
Michael Stanton	RNP	michael@rnp.br
Thomas Tam	CANARIE	thomas.tam@canarie.ca
Linda Winkler	ANL/StarLight	winkler@mcs.anl.gov
David Yeung	Ciena	dyeung@ciena.com

A. Meeting Objectives

This meeting provided an opportunity for North American GLIF participants to meet, discuss mutual opportunities and identify key topics of the next five years. Key findings from the 2009 GLIF-NA Daejeon minutes were discussed, notably ongoing requirements for providing very large network capacity and for flexibility in managing services and capabilities.

B. Major Trends In North American International Connection Requirements, and Communities Served

1. Overview of major 100Gb projects

CANARIE, in collaboration with Ciena, created the very first international 100Gb network, which was implemented as a testbed from StarLight in Chicago to Ciena’s research lab in Ottawa (1200 kilometers) with no regenerators or amplifiers. It performed beyond spec.

There are other 100Gb projects: ESnet’s Advanced Network Initiative (ANI), a national network research 100Gb backbone testbed (which will connect to several GLIF GOLEs); and, several individual projects, such as NCSA Blue Waters, CHI-MAN in Chicago, and the proposed 100Gb bankbone upgrades to NLR and Internet. NASA also has a prototype 100Gb project (currently instantiated as 80Gb between Goddard Space Flight Center in Maryland and StarLight), with a planned 100Gb demonstration showcase for SC10

in New Orleans in November. StarLight/MREN was recently funded by the National Science Foundation (NSF) to bring up a 100Gb exchange and will be planning this in the next 12 months.

2. Projections for resources required *beyond* bandwidth 2011-2016

Large-scale bandwidth provisioning is more important than bandwidth alone; i.e., dynamic provisioning (granulated control) of streams. As part of a GLIF Working Group project, several GLIF members organized a large-scale demonstration of international dynamic provisioning among multiple GOLEs for the GLIF Workshop later this week. One key component of the demonstration was Fenius, an API developed by ESnet (Evangelos Chaniotakis), which enables individuals to access dynamic provisioning capabilities at GOLEs using a common interface so that they do not have to deal directly with individual GOLE control frameworks. This is a follow-up to the demo they showed last year at GLIF. Plans are also being made to showcase this capability at SC10.

3. Areas of major potential new opportunity

One opportunity is the migration of capabilities now being demonstrated (e.g., dynamic provisioning) from research projects to real, persistent, production services at the GOLEs.

4. Potential options for addressing requirements

How does one migrate demonstration projects and prototypes (e.g., dynamic provisioning) to real services? It's not necessary to support implementations at all GOLEs at once; it may be best to target a few GOLEs initially to implement new services. It may take awhile to develop a common service definition among all the GOLEs at once; however a subset of GOLEs should be able to agree on initial prototype services, even if they are not totally finalized, and proceed; i.e., they can create a *de facto standard*.

5. Expected emerging major challenges to be addressed

Funding for such new services remains an ongoing challenge. Another challenge is socio-political among GOLEs; that is, even at the policy level, one must adopt the Fenius model with common interfaces (an umbrella API) to individual policies. There should be a generalized definition of services; details should be done in the control frames of the individual GOLEs.

C. Summary Reviews of Key Initiatives Related to North American R&E International Networking, and US/Canadian Government Support of International Networking

1. NSF International Research Network Connections Program (IRNC)

a. GLORIAD

GLORIAD-US has new location at University of Tennessee, additional new staff, and offices in DC. The GLORIAD team is doing a lot of measurement, monitoring, etc., as well as involved in major infrastructure developments:

- (1) Korea-China-US circuit upgraded
- (2) China-US circuit upgraded
- (3) US-Russia circuit upgraded
- (4) New US-Nordic (IceLink) circuits operational (and Nordic-Russia dark fiber operational)
- (5) First Egypt/Africa-US R&E trial link operational in March (new 1Gb circuit operational by November)

- (6) New circuit from Mumbai-Singapore-Hong Kong-Seattle ready to be provisioned
- (7) New Canada-US (trans-North America) circuits operational
- (8) GLORIAD-CSTNET peering with ESnet in Seattle

b. AmLight

Michael Stanton reported that the AmLight infrastructure is a progression of WHREN-LILA, with the additional partner LEARN in Texas. West coast link's goal is to get to Ensenada. RedCLARA links and country links are increasing in South America.

c. TransLight/PacificWave

Jacqueline Brown reported that she believes there are no recent new developments. Dave Reese said there is growth (added more 20Gb links). David Lassner received an NSF ARI award to connect major islands (in a ring) and then on to Seattle and Los Angeles.

d. TransPAC3 and ACE

- (1) Jacqueline Brown reported on behalf of Jim Williams. TransPAC3/ACE is creating a "round the globe" circle with partners APAN/GÉANT2 TEIN3, NICT-Japan, NII-Japan, CERNET-China, Data Logistics Toolkit, and others.
- (2) TransPAC3/ACE is interested in continuing and strengthening its work on end-to-end measurement and security.
- (3) A TransPAC3/ACE topology map can be found at <http://internationalnetworking.indiana.edu>

There are many other measurement activities happening, and a suggestion was made to provide an overview of all these projects. GLIF-NA can only identify the projects and suggest ways that measurement activities can share concepts and data.

While ACE mentions London and Frankfurt as endpoints, these locations are not firm yet. Williams plans to use OSCARS over ACE across the Atlantic.

e. TransLight/StarLight

The TransLight/StarLight team (Calit2, UIC, NU) continues to build cooperative partnerships and to develop services and technologies on international networks:

- (1) HPDMnet
- (2) iGENI
- (3) SAGE
- (4) CineGrid
- (5) GreenLight International
- (6) Science Cloud Communications Services Network (SCCSnet)

f. IRNC projects IRIS and DyGIR

IRIS and DyGIR – Eric Boyd, Internet2 is PI and Ana Preston Hunsinger is co-PI. The goal is to enable successful IRNC ProNet awardees. IRIS (perfSONAR) is a software framework that provides infrastructure to identify, diagnose and eventually correct network performance problems for paths traversing IRNC links. DyGIR (OSCARS) is a software framework to enable dynamic circuit services for a pair of IRNC links.

g. IRNC projects 6Watch and CAIDA

CAIDA provided a presentation on its IRNC activity. This initiative will undertake sustainable data handling and analysis methodologies for IRNC networks. They will help make operational network data available to the research community.

6Watch provided a presentation on its IRNC funded activity. This effort will help the IPv6 community with measurement, monitoring and diagnosis.

2. DOE ESnet international networking initiatives

Some ESnet efforts are focused on supporting dynamic circuits and on making provisions for measurements. ESnet is involved in DICE activities with CANARIE, Internet2 and GÉANT. The Advanced Networking Initiative (ANI) testbed is currently a table-top testbed and will eventually (in 2011) be implemented on a large scale, beginning in Rhode Island, subsequently extended to New York and Chicago, and then across the nation. The ANI testbed will be open to all researchers and will support national and international users.

ESnet sent out a transport RFP in June and received responses in July, which are now under evaluation. ESnet will negotiate with finalists in the next few weeks and make selections at the end of the month. It may take 6 – 8 months to build the network testbed.

ESnet is also evaluating 100Gb routers in the ESnet lab and has identified bugs that vendors are investigating. Once ESnet concludes the transport RFP, it will issue a router RFP (in November/December timeframe).

3. DOE USLHCnet

LHC is transmitting real (not simulation) data. USLHCnet is seeing very large flows in support of experiments. A new funding cycle starts March 2011, so USLHCnet is waiting for DOE feedback.

Some transatlantic carriers are talking about deploying 40Gb over the Atlantic (but they need to develop a business case), and they are talking about deploying 100Gb within the next five years. Carriers need to assess needs vs. what's available (and affordable). [Note: Joe Mambretti mentioned that in previous discussions, GLIF-NA talked about buying capacity as a consortium. Steve Cotter says the issues aren't technical; the issue is that providers are being "old school". Bill St. Arnaud says consortiums aren't popular anymore; individuals can buy.]

Experiments that USLHCnet is considering:

- (1) Going from a "push" model to "pull"; a change that makes service quality and service guarantees more important.
- (2) Moving to a mesh topology (now a hierarchy of Tier0, Tier1, Tier2, etc); e.g., Tier2-to-Tier2 movement. Some sites would require dedicated infrastructure (exchange points) rather than general networking. (David Foster will comment during his presentation at the GLIF meeting.)
- (3) Remote data access – during execution, reading files remotely. This service would create zillions of tiny flows.

Tier2 equipment is old and needs to be refreshed. What if Tier2 sites decide to compute with clouds? The group discussed how clouds are currently not built for large-scale science data; they are designed for smaller enterprise data and consumers. There are science cloud testbeds that are developing new techniques for addressing these issues – that understand large capacity, dynamic provisioning, etc. – but

they are not yet ready for prime time (for example, Bob Grossman has implemented the national Open Science Cloud testbed). Eventually this approach can provide a model for large-scale science projects like the LHC. DOE is also experimenting with science clouds with its Magellan project. Supercomputers have a different allocation model than clouds – the supercomputers are oversubscribed while clouds need to provide cycles on demand. *Cloud computing is an issue GLIF-NA should discuss in the future.*

4. Internet2 international networking perspective

Internet2 R&D is focused on end-to-end performance, and it is heavily involved in OGF standards and protocols (NSI, NML, Network Measurement, NMC/monitoring) going forward for interoperability. Internet2 is actively engaged with GLIF (perfSONAR, GOLE). Internet2 is working with GÉANT2 and ESnet to come up with a common definition of a dynamic networking service. What is a common service definition that makes sense but is not so big that it is overwhelming? Internet2 is committed to supporting the success of the IRNC ProNet awardees. Internet2 has work going in the Federation space, as well as work going on with GENI and the Openflow space.

Dave Lambert, Internet2 CEO, is very interested in the international space and how Internet2 can work with international groups to provide better service. [Joe Mambretti remarked that Internet2 has been very successful in helping create new NRNs globally, in areas of the world where the concept was not very familiar and Internet2 assured us that this work would continue.]

5. NLR international networking perspective

Although no NLR representative was present, NLR provided a subsequent statement: The NRL continues to provide communication services support at all major US GLIF GOLES. Also, the NLR provides support for a 10Gb trans-Atlantic circuit between New York and the NetherLight exchange in Amsterdam.

6. CANARIE international networking perspective

CANARIE implemented a successful 100Gb wave testbed from Ottawa to Chicago, without any amplification or regeneration between sites. *This was the first 100Gb link at a GLIF GOLE, and the first international demonstration of a 100Gb path using emerging technologies.* CANARIE upgraded the Chicago/StarLight GOLE equipment, which is now based on a Ciena framework, and it will soon upgrade its Seattle/Pacific Wave equipment. There are no current plans to upgrade New York.

The Green Star project is up and running – has three sites currently all powered by renewable energy. It is adding more sites. It needs high-speed uncongested bandwidth, which is perfect for GLIF-related activities. If a site goes down, it transfers virtualization to another site, depending on the power at that site. The FCC is interested in this concept. [BTOP funded programs will soon be implementing networks that can benefit from these concepts.] *Green networking is another issue GLIF-NA should pursue.*

7. NCO Joint Engineering Taskforce (JET)

JET is interested in perfSONAR, IPv6, E2E enabling of international applications, dynamic optical networking, and cloud computing.

D. Potential for Leveraging New International Trans-Oceanic Fiber Builds

Discussed above. *Trans-oceanic fiber builds should be an issue GLIF-NA should pursue.*

E. Projected Future Plans (2-3-5 Years) for North American GLIF GOLE Exchange Points

1. Pacific Wave/PNWGP

Dave Reese gave Pacific Wave report above.

2. MAN LAN

There were no new items to discuss. MAN LAN will continue to monitor capacity requirements. Jim Williams has not confirmed US cities yet for ACE. Intends to run an IDC in MAN LAN (not part of Internet2). MAN LAN wants to have a perfSONAR installation. MAN LAN will replace its HDX in 6 months and it will use the Core Director that was in Chicago.

3. AMPATH

AMPATH will support 2 x 10Gb to South America as part of ProNet. AMPATH is working with various Caribbean nations, the West Indies, and Barbados for connectivity. AMPATH is involved with Atlantic Wave.

Atlantic Wave has an MOU with its exchange points that expires in 2011. A need exists to discuss Atlantic Wave's evolution, what it means, capacity upgrade, services, etc. [100Gb equipment will come on the market in 2011, so now is a good time to start planning for a 100Gb backbone.]

4. StarLight

StarLight is supporting several demonstrations at the GLIF Workshop, including Fenius and HPDMnet (which supports multi-domain dynamic provisioning, including lightpath provisioning). One HPDMnet demonstration will show high-resolution CERN particle collisions. As discussed previously, the StarLight-related IRNC initiatives will provide a number of new advanced programmable services. Also, StarLight/MREN received an ARI grant to upgrade to multi-100Gb capacity over the next 12 months. Multiple new StarLight activities are focused on GENI community activities, and more activities are planned in the future because the next round of funding for Spiral 3 will fund infrastructure. [GENI is an instrument for network research; it is not research itself.]

5. MAX

No representative attended. MAX will upgrade to 100Gb infrastructure.

F. Potential to Implement a New Services Model Across International Facilities, Especially for Large-Scale Research Projects

This issue will be covered in presentations and demonstrations over the next few days. GLIF will provide leadership in distributing services across fabrics.

G. GLIF Dynamic Provisioning Architecture and Services, Fenius Trials

A GLIF Working Group will demonstrate end-to-end dynamic provisioning across multiple GOLES this week at GLIF, using the Fenius interface developed by Evangelos Chaniotakis. Fenius works well and should be considered a potential model for GLIF GOLEs. What has worked especially well is the process of organizing the developers and networking people to allow for this collaboration; the software, too, was

a major contribution. GLIF provided the infrastructure and helped focus everyone to produce the resulting international demo. The ultimate goal is to create standards to make this architecture interoperable across multiple facilities. [GOLEs can implement Fenius today as a prototype, but it is not yet production ready. The Fenius group should develop a method to allow for progress toward a production implementation.]

H. International Research Testbeds (US's iGENI/GENI, Europe's FIRE, Japan's New Generation Network (NWGN), etc.)

Besides GENI, there are other testbeds in development globally. These testbeds would like to obtain dedicated bandwidth in the GLIF. We need to carve out channels within the GLIF infrastructure to support these activities. *This is another issue that GLIF-NA should pursue.*

I. Potential for a Distributed GLIF-NA GOLE NOC and Related Processes

This issue is related to measurements. We have an opportunity to monitor a distributed GLIF space and work more directly with the measurement people. We can leverage the work of architecting a distributed NOC. (USLHCnet does that between Caltech and CERN, and GLORIAD/Korea does this.) Another way to address these challenges consists of automating as many processes as possible.

J. Implications of BTOP Awards for International Networking

\$7.2B has been allocated by the US NTIA agency, which is part of the Commerce Department under the Broadband Technologies Opportunity Program, to advance high-performance communications in regions throughout the nation, although very little has been spent to date. Much of this funding will be used to implement fiber-based community owned and managed communication services. How will this relate to international networking? Some recipients have international applications. For example, medical communities want to have high-performance access to international audiences. We should keep our eye on the developments being funded under these grants.

K. Wrap-Up

The community is looking forward to the GLIF meeting in Rio next year on September 13-14, hosted by the GLIF participants from Brazil. The next GLIF-NA community meeting will be held during that week.

Heidi Alvarez, Julio Ibarra and Michael Stanton will make suggestions to rename GLIF-NA to best represent GLIF North and South America.