# Use case analysis GLIF Architecture Task Force –DRAFT

September 25, 2012

Prepared by:

Bill St. Arnaud bill.st.arnaud@gmail.com

Erik-Jan Bos bos@nordu.net

Inder Monga imonga@es.net

The purpose of this document is to specify the potential use cases for the GLIF architecture task force to enable end to end lightpath connectivity for a researcher to access remote databases, computers and instruments. Although great success has been achieved in establishing lightpaths across multiple independent managed networks and GOLES achieving connectivity across campus networks to the researcher's desktop remains an elusive goal. Recently technologies such as Science DeMilitarized Zones (DMZs) and campus Software Defined Networks (SDN) promise to alleviate some of the campus lightpath challenges. But the interconnection and interoperability of DMZs, SDNs and other campus network architectures with global interconnected lightpaths as a seamless architectural vision remains an unrealized objective. Not only is a seamless physical interconnection required, but the specification of all the user interface, management, measurement, operational and control aspects of the architecture must be detailed as well.

Compounding the problem of defining an end-to-end lightpath architecture is the increasing need for researchers to interconnect lightpaths to commercial databases, clouds and computational resources. In some cases the end-to-end solution may not even touch the campus network. Instead a researcher may wish to connect to the output of a remote instrument directly to a commercial cloud. Building end to end solutions in the academic/research world with its commitment to openness and collaboration is one thing, but this can be quite a bit more problematic in the commercial world with its concerns about competition, privacy, security etc.

The ultimate vision of the GLIF architecture task force is that a researcher, or an application can compose or create an end-to-end lighpath solution across a campus, multiple GOLES and networks using a simple interface such as SURFconext, Globus OnLine or COmanage. All the necessary management, measurement and control tools would also be incorporated in such an interface. To simplify the complexity of interconnecting many independent lightpaths across multiple networks, many services may be consolidated into a much smaller number of abstracted services which can also be an advertised service as part of an end-to-end solution. Clearly any likely architecture will likely therefore be a recursive Service Oriented Architecture (SOA). Several SOA network platforms have been explored by the research community including OpenNaaS, Mantychore, OSCARS, UCLP etc. and some, like OSCARS and Argia are have been actively deployed in production R&E networks. Some even include support for lightpaths and multi-domain path setup protocols e.g. Argia, OSCARS, OpenNaaS and UCLP. The challenge for the GLIF architecture committee is to identify strengths and weaknesses of these platforms and the missing elements to facilitate a true end-to-end architecture.

With any type of end-to-end switched architecture the role and process of initiating and terminating parties must be carefully addressed. How does a researcher at one campus initiate an end-to-end lightpath to a research or database another campus if they have no authority or credentials to setup a lightpath at the destination campus? At the network to network level authenticating and accepting lighpath requests across a GOLE or intervening network , although not trivial, is relatively easy in comparison. Can a researcher delegate authority to allow external parties to setup a lighpath across the campus network? Or should a researcher only be authorized to "meet in the middle" at a GOLE or similar facility i.e. all lightpaths terminate at GOLEs , one set from the designated originator and another set from the designated recipient? The ability or non-ability to receive external lightpaths will be a major determining factor on the GLIF end-to-end architecture.

To help clarify the requirements for the GLIF end-to-end architecture it would be useful to document the range of possible use cases that would have to be addressed by such an architecture. The following list is a summary of the various possible use cases with a more detailed analysis of each case given separately:

(a) True lightpath connectivity across campus with direct interconnect to global GLIF services;

(b) SDN network on campus (most likely OpenFlow) with configured Ethernet VLANs as lightpaths to interconnect GLIF lightpaths at campus interface;

(c) Campus DMZ outside of campus network interconnected to GLIF facilities on the outward facing connection and IP connection facing inward;

(d) Campus IP network with VPNs (MPLS) or VLANs to interconnect to GLIF facilities at campus egress;

(e) Terminating end-to-end lightpath on a commercial interface: e.g. cloud; and

(e) Establishing lightpath connection on a remote instrument network to a commercial cloud or database

Each of these use cases will be explored more fully in the following sections:

# True lightpath connectivity across campus with direct interconnect to global GLIF services

A small number of university and research campuses have local area optical networks with dynamic switching of optical lightpaths across the campus as well as direct connections to GLIF network facilities. Most of these optical networks are operated completely independent of the campus IP network and are responsible for their own security and global connectivity. In some cases servers connected to the optical campus network are firewalled from the campus IP network.

Some of these networks support NSI (or its predecessor protocols ). As such setting up end-to-end lighpaths is rather trivial compared to the other use cases. Questions that remain outstanding are:

(a) How to identify connected devices and ports – NDL ontologies?

(b) Can campus connectivity be delegated to third parties who wish to access campus resources via an externally originating lightpath?

(c) How to incorporate end-to-end measurement and management services such as PERFSonar?

# SDN network on campus (most likely OpenFlow) with configured Ethernet VLANs as lightpaths to interconnect GLIF lightpaths at campus interface

A growing number of universities, research campuses and large data centers are deploying various SDN networks, mostly variations of OpenFlow. SDN or OpenFlow allows the network manager to easily and quickly configure dedicated VLAN Ethernet networks to various researchers and users on campus. With OpenFlow these devices can be centrally managed and configured – which is often very appealing to a campus network manager.

For the most part ingress and egress to the campus is at the IP layer through a campus border router. Considerable research is going on to map OpenFlow VLANs to MPLS and GMPLS VPNs. A few examples include proof of concept with OSCARS demonstrated at SC11 (http://sc11.supercomputing.org/schedule/event_detail.php?evid=rsand110) and work at Internet2 with NDDI.

Common mechanisms including policies need to be developed to map OpenFlow VLANs to NSI optical paths. This would probably major focus of activity for the GLIF Architecture task force.

Outstanding issues are as follows:

(a) Mechanisms to automate the mapping of OpenFlow VLANs to optical lightpaths via NSI at campus border switch

(b) Does OpenFlow support guaranteed bandwidth VLANs to act as surrogate lightpaths?

(c) Can campus connectivity be delegated to third parties who wish to access campus resources via an externally originating lightpath?

(d) How to incorporate end-to-end measurement and management services such as PerfSonar?

(e) Can a researcher, independent of a campus network supervisor, encapsulate, abstract and advertise their OpenFlow VLAN mapping to NSI to third parties, so that external parties can initiate end-to-end lightpath connection?

# Campus DMZ outside of campus network interconnected to GLIF facilities on the outward facing connection and IP connection facing inward;

To get around many of the bandwidth and connectivity limitations of campus networks, ESnet in particular has been promoting the concept of DMZs. With a DMZ a campus researcher can upload or download large data files to a server outside of the campus firewall. The DMZ is directly connected to GLIF optical infrastructure. For the most part the DMZ is considering the terminating device and the rest of the campus network including researcher's services remain

hidden from external users. DMZ also come configured with PerFSonar and other network management devices which makes measurement easier.

As researcher's progressively move to using commercial clouds for storage and computation the DMZ may in fact become an intermediate stop point for a data flow between an instrument and a commercial cloud facility. The interconnection to the campus network becomes less relevant and would let researchers do large data analysis from their local coffee shop. In that case the ability to set up lightpaths from the DMZ or the originating instrument itself to a commercial cloud becomes important.

Outstanding issues:

(a) ScienceDMZ needs to be enhanced with a security blueprint that supports the different security requirements of most campuses without affecting performance

(b) Various Science DMZ design patterns that work integrate seamless with GLIF and OpenFlow/SDN, NSI and performance management

(c) Can researcher delegate third party lightpath access to local DMZ and PerfSonar?

# Campus IP network with VPNs (MPLS) or VLANs to interconnect to GLIF facilities at campus egress;

This is the most common interconnection, other than using general IP for interconnecting researchers with GLIF infrastructure. In many cases, campus configuration problems bedevil the setup of end-to-end lightpaths which has resulted in the deployment of Science DMZs.

Considerable work has been done in NSI, IDCP and other lighpath switched protocols to map optical lightpaths to MPLS tunnels.

Outstanding issues include:

(a) Can campus connectivity be delegated to third parties who wish to access campus resources via an externally originating lightpath?

(b) How to incorporate end-to-end measurement and management services such as PerfSonar?

(c) Can a researcher, independent of a campus network supervisor, encapsulate, abstract and advertise their OpenFlow VLAN mapping to NSI to third parties, so that external parties can initiate end-to-end lightpath connection?

# Terminating end-to-end lightpath on a commercial interface: e.g. cloud

As mentioned previously there is growing demand by researchers to use commercial clouds and databases for the uploading downloading of large datasets, as well as directly forward data from instruments.

In most environments the connection to a commercial cloud provide such as Google, Amazon, Azure, GreenQloud, etc is owned and controlled by a NREN. Connectivity is provided at the IP layer through standard IP addressing and naming. However, the need for a researcher to have a direct connection independent of the IP service layer is growing. This will introduce a host of problems of how to terminate individual lightpaths through perhaps a single 10G pipe to a cloud service provider? Most commercial cloud providers have not yet scaled up to handle this type of large IO data flows (although they do handle teabits of IP flows).

As most commercial cloud providers are not yet ready to accept lightpaths, it is likely that the NREN will have to offer a proxy service to terminate and manage lightpath requests to a commercial cloud – in effect operating a "reverse" DMZ on behalf of the commercial cloud operator. The ability, therefore to terminate originating lightpaths from third parties will be an essential feature.

Considerable more research has to be done for this use case. SURFnet in partnership with GreenQloud in Iceland is probably the most advanced in this field. Their experience will be a useful in helping other NRENs and researchers use lightpaths to transmit and receive data from cloud providers.

# Establishing lightpath connection on a remote instrument network to a commercial cloud or database

This use case is the ultimate example of third party delegation of lighpaths – where an independent researcher may, for example, want to setup a lightpath from CERN to a commercial cloud provider such as Amazon. None of the lightpaths may terminate or come even close to touching the researcher's own campus network.

Clearly addressing issues of third party delegate as part of the GLIF architecture will be essential before this use case can be addressed.

***