GLIF Technical & Control Plane Working Groups
19-20 January 2008
University of Hawaii, Honolulu, United States

Attendees

Name	Organisation	Country
Celeste Anderson	USC/Pacific Wave	United States
Artur Barcyk	Caltech	United States
Erik-Jan Bos (Tech WG Chair)	SURFnet	The Netherlands
Eric Boyd	Internet2	United States
Heather Boyles	Internet2	United States
Matt Crawford	Fermilab	United States
Lars Fischer	NORDUnet	-
Licia Florio (CP WG Secretary)	TERENA	-
David Foster	CERN	-
Clifford Frost	UC Berkeley	United States
John Graham	Global NOC	United States
Gigi Karmous-Edwards (CP WG Chair)	MCNC	United States
Akira Kato	WIDE	Japan
Dongkyun Kim	KISTI	South Korea
JongWon Kim	KISTI	South Korea
JongUk Kong	KISTI	South Korea
Li-Chi Ku	NCHC	Taiwan
Tomohiro Kudoh	AIST	Japan
Geoff Lakeman	Pacific Northwest GigaPoP	United States
Hui-Lan Lee	NCHC	United States
Tom Lehman	USC/ISI	United States
Mathieu Lemay	Inocybe Technologues	Canada
Te-Lung Liu	NCHC	Taiwan
Iara Machado	RNP	Brazil
Kevin Meynell (Tech WG Secretary)	TERENA	-
John Moore	MCNC	United States
Dan Nae	Caltech	United States
Kees Neggers	SURFnet	The Netherlands
Harvey Newman	Caltech	United States
Bram Peeters	SURFnet	The Netherlands
Ivan Phillips	AARNet	Australia
Jan Radil	CESNET	Czech Republic
Don Riley	Univ. Maryland/IEEAF/SURA	United States
Don Robertson	AARNet	Australia
Matt Siniscal	MAX	United States
Michael Stanton	RNP	Brazil
Thomas Tam	CANARIE	Canada
Christian Todorov	Internet2	United States
Andree Toonk	BCNET	Canada
Ronald van der Pol	SARA	The Netherlands
Alan Verlo	UIC/StarLight	United States
Josef Vojtech	CESNET	Czech Republic
John Vollbrecht	Internet2	United States
James Williams	Univ. Indiana	United States
Linda Winkler	StarLight	United States

1. Joint Session of Technical and Control Working Groups

1.1 Introduction

Erik-Jan welcomed everyone to the meeting and introduced Gigi as the co-chair of the session. He outlined the objectives of the Technical and Control Plane Working Groups; the former focusing on provisioning global lightpaths today, with the latter looking at how to provision them in future. However, there were increasingly activities of common interest, which was why it had been agreed to hold a joint session.

Thanks were extended to Internet2 and the University of Hawaii for hosting the GLIF Working Group meetings.

1.2 Actions from last meeting

20070917-1 Catalin Meirosu to create new 'resources' section on the GLIF Wiki to allow AARNet and other GLIF lambda contributors to list their resource information.
- Done.

20070917-2 GOLE operators without an account on the GLIF wiki to contact GLIF Secretariat for login details.
- Done.

20070917-3 Internet2, NLR, and Pacific Wave to send information about how to request bandwidth for SC'07.
- Done.

20070917-4 René Hatem to update mailing list about the new working areas on lightpath visualisation and management.
- Superseded.

1.3 DCN GOLE: Visions and Challenges

John V presented some ideas about how to establish automated exchange points in dynamic circuit networks (see http://www.glif.is/meetings/2008/winter/joint-session/vollbrecht-dcn.pdf). This had been demonstrated between Internet2 and SURFnet, utilising MAN LAN and NetherLight, but had raised some operational issues that needed to be addressed.

The main issue was how to establish policies for each GOLE. For example, what priority should intra-domain scheduling take over inter-domain connections, how is contention resolved where projects may not always be using their circuits, and how can certain classes of user be prioritised? There are a number of possible solutions, but these need to be discussed within GLIF.

1.4 Experiences from LHC OPN

David gave a presentation on the Large Hadron Collider Optical Private Network (see http://www.glif.is/meetings/2008/winter/joint-session/foster-lhcopn.pdf). This aims to connect the LHC at CERN (Tier 0) with the twelve Tier 1 sites that will process the collected data and manage the experiment requests. The infrastructure started as non-redundant star topology, but the requirements have increasingly leveraged the procurement of cross-border fibre that has enabled resilience to be added.

Whilst the LHC OPN will essentially remain a private closed network, there is scope for the GLIF to become involved with connecting Tier 2 to Tier 1 sites, and even with respect to adding additional resilience between Tier 1 sites. In addition, the GLIF could work to establish connectivity to those countries involved in the LHC that were still missing from the GLIF map.

Kees asked whether Tier 2 sites would have permanent or dynamic connections. David replied it would be dependent on the nature of the site as some Tier 2 sites were bigger than the Tier 1s.

1.5 VLANs and Lightpaths

Alan gave a presentation about the network requirements of the CineGrid demonstration that had been given during the 7th Annual Global Lambda Workshop in Prague (see http://www.glif.is/meetings/2008/winter/tech/verlo-vlans.pdf). CineGrid was an activity to promote the production and transmission of high-quality digital media over optical networks, and had undertaken a number of demonstrations during the 2007.

In Prague, CineGrid demonstrated two applications: a 4K JPEG2000 stream from several sites to Charles University, and colour correction/post-production in real-time from Ryerson University. This required a 10 Gbps circuit from each site, over which were provisioned IP-based VLANs using private addressing.

A number of issues had been encountered such as non-contiguous lightpaths, Layer 1 and 2 transitions along the path, terminating equipment that did not always support GFP/VCAT/LCAS, and configuration differences between domains. In addition, not all applications required a full 10 Gbps, and some only used the bandwidth sporadically.

2. Technical Working Group Meeting

2.1 GOLE Updates

AARNet

Ivan reported on the current status of AARNet's domestic dark fibre footprint, and their international links to other research and education networks (see http://www.glif.is/meetings/2008/winter/tech/phillips-aarnet.pdf). They had a loop of dark fibre between Brisbane, Sydney, Canberra, Melbourne and Adelaide over which two channels were provisioned (with the possibility to increase this to 32) for production IP traffic and SDH-based lightpaths. They also had dark fibre between Adelaide and Perth, but currently lacked funding to light this. There was an international STM-64 to CENIC via Hawaii that could be used for lightpath provisioning, and this had previously been used for OptiPuter and EXPReS e-VLBI experiments.

The plans for 2008 were to add additional STM-64 channels on all domestic legs for supporting radio astronomy and OptiPuter projects. However, they were also looking for other applications that could utilise this infrastructure.

John asked whether AARNet had an IRU on its fibres. Ivan replied this was the case.

CANARIE

Thomas gave an overview of the CANARIE infrastructure (see http://www.glif.is/meetings/2008/winter/tech/tam-canarie.pdf). CANARIE-operated wavelengths existed in the more populated regions in the east and west of Canada, with carrier-provisioned wavelengths used for the geographic expanse in-between, as well as from Montreal to the Maritime Provinces. They also had a wavelength-swap agreement with NLR that provided them with a 10 Gbps wavelength from Chicago to Seattle, in return for four 10 Gbps wavelengths from Boston to New York, and a 10 Gbps wavelength from Boston to Chicago. These wavelengths were currently provisioned using SONET.

CANARIE had been involved in the High Performance Digital Media network testbed (HPDMnet) that started in April 2007. This aimed to explore new Layer 1 and 2 capabilities to support a large-scale HPDM service, and a number of optical multicast experiments had already been demonstrated. They were also using the TL1 toolkit to develop a lightpath status page similar to that provided by NetherLight, and were extending UCLP to support their GOLEs.

CzechLight

Jan reported that a new Gigabit Ethernet circuit had been added from CzechLight to VINI, via NetherLight and StarLight (see http://www.glif.is/meetings/2008/winter/tech/radil-czechlight.pdf). This meant there were now two VINI nodes in the Czech Republic. They also had a connection to Brookhaven National Laboratory via NetherLight, MANLAN and ESnet. In addition, an NDL description of CzechLight was available at http://www.ces.net/network/cesnet-gole.rdf.

Erik-Jan pointed-out that CESNET wasn't represented on the CzechLight diagram. He said it was useful to know about all connectivity from a GOLE, and that CESNET2 should be represented as a cloud on the diagram.

ACTION 20080119-1: Jan Radil to add CESNET2 to the CzechLight diagram.

KRLight

Dongkyun gave an update on the status of KRLight (see http://www.glif.is/meetings/2008/winter/tech/kim-krlight.pdf) that currently had links to StarLight (United States), PacificWave (United States), CERNET (China), HKOEP (Hong Kong), the University of Hong Kong, and KREONet2 (South Korea). The existing Cisco ONS15600 and 15454s would be replaced by Nortel OME6500s later in the year.

KRLight was also involved in a pilot project to provide live HD medical transmissions between Yonsei Hospital in Seoul, and the National Technical University in Trondheim.

MAN LAN

Christian gave an update on the status of MAN LAN (see http://www.glif.is/meetings/2008/winter/tech/todorov-manlan.pdf). The equipment had previously been subsidised by Internet2, but they now needed to move to a self-sustaining cost model. As a result, the connection fees would be stepped-up in 2008 and 2009, although no increase was anticipated for 2010. The HDXc code would be upgraded during the first quarter of 2008, and there were plans to deploy DRAC and perfSONAR on an experimental basis during the second quarter.

Erik-Jan asked whether there would any additional fees for VLANs. Christian replied the fees only applied to physical connections.

Erik-Jan also asked what the fees would be for 40 Gbps connections. Christian said these had yet to be determined.

Christian added that the introduction of a full cost recovery model meant they needed to start thinking about having a proper service level agreement. He thought it might be useful to start defining a common SLA definition for all GOLEs, even if there might have to be some differences to take into account the nature of the involved organisations.

After some discussion, it was agreed that a small task force should be formed to start drafting some ideas for this. The volunteers were Walter van Dijk (Leader), Geoff Lakeman, Christian Todorov and Jim Williams

ACTION 20080119-2: SLA Task Force to draft some common SLA definitions.

Mid-Atlantic Crossroads

Matt provided an overview of the MAX topology (see http://www.glif.is/meetings/2008/winter/tech/siniscal-max.pdf) that was comprised of seven switches interconnected by a mixture of 10 and 40 Gbps wavelengths running SONET and Ethernet circuits. Three switches would be added in Baltimore, Boston and New York in April 2008, and GMPLS-enabled VLSR switches would be installed throughout the DRAGON group by June 2008. RNP and CLARA would also use the MAX infrastructure to connect to CERN via StarLight and NetherLight.

John asked what had happened to their Movaz ROADMs. Matt replied that since Adva that taken over Movaz they had learnt they would be discontinued, so had decided to replace them. However, they planned to add GMPLS code to their Adva ROADMs.

NetherLight

Ronald reported that the cross-border fibre from NetherLight to NOX via Hamburg was now operational (see http://www.glif.is/meetings/2008/winter/tech/vanderpol-netherlight.pdf). A 10 Gbps link to Taipei (Taiwan), and a 1 Gbps link to Mumbai (India) had also recently been added. The IRNC link now terminated on the NetherLight ERS8600 rather than the UvA Force10 switch which removed one domain from the path.

The Nortel OME6500s in Amsterdam and Geneva had been upgraded to release REL0400Z to support the 10 Gbps EPL cards at full line rate. The HDXc would also shortly be upgraded to 4.0.2 which would allow all eight ports on the 10 GE card to be used (rather than the four as now).

A new monitoring system was now available that can provide an overview of lightpath status, as well as timeslot usage on the interface of a network device via a web page. A total of 31 lightpaths were currently running through NetherLight.

NORDUnet Optical Exchange

Lars reported that NOX had become a distributed GOLE with connection points in Copenhagen, Helsinki, Oslo and Stockholm (see http://www.glif.is/meetings/2008/winter/tech/fischer-nox.pdf). These were interconnected with dark fibre, and onwards to NetherLight via Hamburg (currently 5 x 10 Gbps). There were also five 10 Gbps connections to GÉANT, a 10 Gbps connection to MoscowLight, and a 2.5 Gbps connection to RUNnet (Russia).

The NORDUnet infrastructure was scheduled to be upgraded with Wavelength Selectable Switches which should reduce the need for OEO switching. In addition, a new OC-48/192 link to Iceland was being planned, and further cross-border dark fibre links to Germany, Poland and the Baltic states were being investigated.

In Sweden, the OptoSUNET network now offered a full DWDM service, whilst in Norway, Uninett offered a 1/10 Gbps national lambda service. In Denmark, the fibre tender had been completed, but the DWDM equipment tender was still ongoing. In Finland, the tenders had still to be finalised, although some lambda services were available. Unfortunately, there were not currently any lambda services available in Iceland.

Erik-Jan Bos said that it would be useful to know about the topologies of any national and regional networks connected to GOLEs, so that one could determine whether lightpaths could be established between particular locations. He asked whether GOLEs could publish this information on the GLIF Wiki.

ACTION 20080119-3: GOLEs to publish topologies of national and regional networks connected to them.

PacificWave

Geoff provided an overview of the current PacificWave facilities (see http://www.glif.is/meetings/2008/winter/tech/lakeman-pacwave.pdf). This had three nodes in Seattle, Sunnyvale and Los Angeles with onward connections to AARNet, CANARIE, HKOEP, KRLight, T-LEX and StarLight. It also provided connectivity to CENIC (the Californian REN) and indirectly to Internet2 and NLR via the Pacific Northwest GigaPoP.

StarLight

Linda gave a short update on StarLight. This currently had connections to CANARIE, CERN, CzechLight, MAN LAN, MAX, NetherLight, the Pacific Northwest GigaPoP, T-LEX, and UKLight.

T-LEX

Akira said there was not much to report about T-LEX, except the ASCC had been replaced by the ASCG (see http://www.glif.is/meetings/2008/winter/tech/kato-t-lex.pdf).

They had also been involved in a 4K uncompressed video streaming demonstration (6 Gbps) between Kyoto and Stockholm via Chicago, Prague and Amsterdam that had provided very successful, despite a problem with their Nortel switch.

TaiwanLight

Li-Chi provided an overview of the TWAREN infrastructure (see http://www.glif.is/meetings/2008/winter/tech/ku-taiwanlight.pdf) that comprises four core nodes and twelve GigaPoPs interconnected by 10 Gbps links. TaiwanLight operates the two 2.5 Gbps international links to Los Angeles and Palo Alto.

Live ophthalmic surgery was demonstrated between the NCKU College of Medicine and the Czech Central Military Hospital in October 2007.

Erik-Jan mentioned there was a Layer 1 link between TaiwanLight and NetherLight that should probably be represented on their diagram.

ACTION 20080119-4: Li-Chi to add link to NetherLight on the TaiwanLight diagram.

Other resources

Erik-Jan explained that a number of European NRENs were currently procuring dark fibre to their neighbours that could be utilised in a wider pan-European network (such as GÉANT2). This was categorised into A, B and C types, with type A being utilised by DANTE for GÉANT2 using the owner NRENs as suppliers.

Josef then showed a map produced by Stanislav Sima that displayed all the known cross-border fibre in Europe.

2.2 Management of Dynamic Lightpaths

Ronald gave a presentation on some of the management issues related to dynamic lightpaths (see http://www.glif.is/meetings/2008/winter/tech/vanderpol-management.pdf). Methods needed to be found to handle alarm conditions that occurred when provisioning lightpaths, such as when a circuit is not available end-to-end (unequipped alarm) or when there is an outage (link-down alarm). In addition, a number of control plane systems (e.g. UCLP and DRAC) required identifiers to be generated for each lightpath. As dynamic lightpaths are established across different domains, methods of generating globally unique identifiers need to be agreed.

Once lightpaths are established, there also needs to be ways of monitoring them. How should this be undertaken, and what status and configuration data should be shared between domains? Finally, there needed to be agreement on points of contact for each lightpath in case of planned work or outages. At the present time, trouble tickets were globally broadcast globally, but this approach was ultimately not scalable.

After some discussion on these issues, it was agreed that a small task force should be formed to start working on global lightpath identification. The volunteers were Ronald van der Pol (Leader), Lars Fischer, Tom Lehman and Thomas Tam.

ACTION 20080119-5: Global Identifier Task Force to draft some ideas for generating unique lightpath names.

It was also thought there should be a demonstration of perfSONAR during the 8th Global Lambda Workshop in Seattle, as a possible technique for monitoring lightpaths. Thomas Tam (Leader), Jeff Boote, Lars Fischer, Dongkyun Kim and Ronald van der Pol agreed to work on this.

ACTION 20080119-6: perfSONAR Task Force to organise demonstration of perfSONAR at 8th Global Lambda Workshop.

Finally, it was suggested that the SLA Task Force could also consider ticketing issues as well.

ACTION 20080119-7: SLA Task Force to draft some common SLA definitions.

3. Control Plane Working Group Meeting

The meeting was opened with a review of actions from last meeting. One of main outcomes of the last Control Plane meeting was the agreement to share existing interfaces from several global projects to define the specifications for a GNI API. It is important to note that the main focus is interoperability of the existing control plane software (GMPLS, UCLP and so on). The importance of defining real operational scenarios in order to tailor the API design was stressed.

3.1 GNI Functionalities

The design for a possible API architecture, as agreed by the group during last meeting, was reviewed. This architecture would also support Grid heterogeneous resources. A weak point in the architecture is due to the fact that the communications is asynchronous, which makes the system quite fragile. In the model presented, the owner of the resources writes their own policies.

A discussion followed where it was agreed that there is need for a policy broker, which is currently missing.

The group agreed that coordination of network and storage is not needed during the first phase but will be addressed after GNI is developed. It was also pointed out that the Grid software does not implement storage interfaces properly. Everybody agreed that for the time being it would be better to focus on network interfaces only. The architecture should be able in the future to incorporate other resources.

It was pointed out that the current design of the architecture lacks high-level topology functionalities and that this should be added.

It was further agreed to define the GNI in a way that it can support both an hierarchical and a chain model. Most of the participants felt that this could be carried out by the GNI API and will work on making this a reality.

The group suggested asking industry research leaders to come and talk during the next Control Plane meeting about their existing solutions

ACTION 20080119-8: Gigi Karmous-Edwards to gather all comments on the GNI architecture and prepare a new version for the next GLIF meeting.

ACTION 20080119-9: Gigi Karmous-Edwards to invite commercial companies (e.g. IBM, Microsoft and Cisco) to the next meeting to present their work on control plane issues.

3.2 DICE IDE

John V presented the DICE architecture and also reported on the demo at SC'07 (see http://www.glif.is/meetings/2008/winter/controlplane/vollbrecht-idc.pdf). DICE, the inter-domain circuit network, allows dynamic circuits to be created across multiple circuit networks. Internet2, ESnet and Dante have worked and continue working on the DICE control plane. DICE has also recently started a collaboration with the University of Amsterdam and Nortel.

John explained the differences between the Inter-Domain Controller (IDC) and the Domain Controller (DC) in DICE. The inter-domain functionalities are provided by the IDC, whereas the DC only works within its domain. Each IDC can be seen as a representation of a particular network. The IDC is also responsible to request services from its DC. The DC is only responsible to control its network, whatever technologies are used. Interoperability at the inter-domain level is ensured by IDC-to-IDC communications.

The IDC might be implemented in a way that it would know all the domains, but it might also be implemented in order to only know how to get to the next hop. In the current implementation the IDC only knows how to get to the next hop. The main issue is to define the way the various IDCs talk amongst each other.

The talk focused mainly on the status of the IDC implementation and it was said that:

- RSVB is not implemented to setup circuits as not all vendors use this technology, but the plan is to support it;
- topology exchange does not use dynamic exchange at the moment;
- topology and path computation is still done by configuring paths between endpoints, which implies that the topology needs to know how to get to the end-destination.

Tom and John discussed also about the IDC future. In this respect two options are being explored: one option would be to continue the IDC development, whereas the second option would be to approach standard bodies (i.e. IETF) and pursue the standardisation path. This triggered a discussion on how best to proceed.

The conclusion was to work together as a community, to utilize existing APIs from DICE, G-lambda, UCLP, etc.. and formulate a GNI to start a dialogue with the IETF.

ACTION 20080119-10: Gigi Karmous-Edwards and John Vollbrecht to call a few IETF Area Directors to see whether there is any interest in IDC standardisation.

3.3 Open GOLE Concept

Inder Monga presented the Nortel IDC implementation and the way the communication with another IDC takes place (see). It was proposed to standardise the interface between the IDC and the domain controller.

A discussion about policies took place, mainly concerning whether the GOLE should be policy free or not. The conclusion was that a GOLE could be truly policy free only if policies are in place in the various domains. In the current implementation GOLEs are connected to each other and the policy determines what it is allowed on an inter GOLE link.

Inder also said that the role of the GOLE could be different, a GOLE could for instance aggregate information about another domain. Especially for the GLIF community he suggested to work to define DCN GOLE.

ACTION 20080119-11: Gigi Karmous-Edwards to discuss the role of GOLEs together with the Technical Working Group.

3.3 Inter-Domain Controller API

Evangelos gave a presentation on the IDC API (currently in pre-production), designed by a partnership composed of DICE, Phosphorus, TeraPaths(BNL), LambdaStation(FNAL) and Nortel (see http://www.glif.is/meetings/2008/winter/controlplane/chaniotakis-idc.pdf). The API allows for path computation, scheduling and provisioning; the topology exchange is IDC-to-IDC only.

The API has a web interface and supports SOAP/WSDL. The data types used for the API are compliant with the specifications defined by the Network Management Working Group (NMWG) OGF working group; this is a plus for the design but in the long term if the specifications changes it will require changes in the API implementation.

Current open issues with the API are:

- reservation process quite invisible for the end users in particular in case of problem; users get very limited feedback on what they do.
- data plane quite slow.
- topology exchange: some domains only want to expose an abstracted view of their local topology, and only to trusted peers; this might lead to problems handling topology exchange, maintenance and outages. Evangelos pointed out that the topology exchange mechanism used for the API is quite comparable to the BGP functionalities and therefore it would be wise to leverage with available BGP implementations.
- reservation debugging not really possible right now

Evangelos also presented a proposal to address the current API issues. Lots of discussion followed to understand how other to use the API in other projects/environments.

ACTION 20080119-12: Evangelos Chaniotakis to upload current WSDL of DICE to GLIF Wiki.

3.4 GNS-WSI

Tomohiro gave an update on the progresses of the GNS-WSI (Grid Network Service / Web Services Interface), which allows reserving bandwidth between end points (see http://www.glif.is/meetings/2008/winter/controlplane/kudoh-gni.pdf). This interface was used for EnLIGHTened project and it is currently used for G-lambda project.

GNS-WSI functionalities include basic operations (reserve, cancel and modify), 2-phase commit protocol, per-request hierarchical architecture (this can support chain model) and authentication policy. Concerning the authentication policy Tomohiro said that a component can either delegate the authentication certificate it receives from a client (using GSI), or use its own certificate.

Because information service is required to find resources, get availability and attributes, the GNS-WSI also support policy-based information service. Each reservation is identified by a reservation ID, a URI that represents the path.

During the discussion that followed it was noted that the same functional blocks are found in the various projects, even if they might have different names, whereas a common interface to connect the various efforts is still missing.

As result of the discussion it was agreed to only concentrate on an interface that communicates to one domain. In the case of a GNI, it was agreed to look at the communication IDC-2-IDC.

ACTION 20080119-13: Tomohiro Kudoh to upload current WSDL of GNS-WSI API to GLIF Wiki.

3.5 Any other business

It was agreed to create the following task forces to work on specific issues:

GNI Specifications Task Force
This group will write the GNI specification, using existing interfaces to capture the minimum set of calls and parameters. They will share existing WSDL by having all the key groups who successfully demonstrated an inter-domain API upload their version of the WSDL to the GLIF WIKI. This includes DICE effort, G-lambda's GNS-WSI, and Phosphorous API, and any others that have already been successfully demonstrated. The GNI API should at a minimum include requests from clients to a domain for a path or partial path set-up and teardown. The group will be led by Evangelos Chaniotakis (ESNet) and will also include Tom Lehman (USC/ISI), Mathieu Lemay (CRC), Tomohiro Kudoh (AIST), Inder Monga (Nortel), Bram Peeters (SURFnet), and John Vollbrecht (Internet2).

ACTION 20080119-14: Mathieu Lemay to upload current WSDL of UCLP to GLIF Wiki.

Security and Authorization Task Force
This group will investigate security and authorization with respect to multi-domain lightpaths. The group will be led by Mathieu Lemay.

Topology Information Task Force
This group will investigate how to undertake topology information exchange. Its activities still needed to be further detailed.

Failure Modes Task Force
This group will investigate failure modes. Participants will be Mathieu Lemay and Bram Peeters.

4. Future of the Working Groups

Erik-Jan outlined the differences between the Technical and Control Plane Working Groups. The Technical Working Group had largely focused on current operational issues, whilst the Control Plane was more focused on future development. However, it had become clear there were increasing overlaps between the groups, particularly as optical networks started to move towards more dynamic provisioning.

It was therefore proposed that the two working groups should be merged, with small sub-groups being formed if specific issues needed to be worked-on. In addition, break-out sessions could still be organised during meetings for those issues that were not of interest to everyone (such as GOLE updates).

Those present agreed that the working groups be merged, subject to ratification by the GLIF Governance Working Group. It was suggested that Erik-Jan Bos and Gigi Karmous-Edwards remain as co-chairs of the new group, a name for which would be proposed in due course.

ACTION 20080219-15: Erik-Jan Bos and Gigi Karmous-Edwards to propose that Technical and Control Plane Working Groups be merged, and to suggest a name for the new group.

5. Date of next meeting

The 8th Annual Global Lambda Workshop would be held on 1-2 October 2008 in Seattle, USA. This would include a meeting of the newly-merged Technical and Control Plane Working Group.

Given that there would be more than eight months until that meeting, it was agreed than an interim teleconference should probably be held in June.

Open Actions

20080119-1 Jan Radil to add CESNET2 to the CzechLight diagram.

20080119-2 SLA Task Force to draft some common SLA definitions.

20080119-3 GOLEs to publish topologies of national and regional networks connected to them.

20080119-4 Li-Chi to add link to NetherLight on the TaiwanLight diagram.

20080119-5 Global Identifier Task Force to draft some ideas for generating unique lightpath names.

20080119-6 perfSONAR Task Force to organise demonstration of perfSONAR at 8th Global Lambda Workshop.

20080119-7 SLA Task Force to draft some common SLA definitions.

20080119-8 Gigi Karmous-Edwards to gather all comments on the GNI architecture and prepare a new version for the next GLIF meeting.

20080119-9 Gigi Karmous-Edwards to invite commercial companies (e.g. IBM, Microsoft and Cisco) to the next meeting to present their work on control plane issues.

20080119-10 Gigi Karmous-Edwards and John Vollbrecht to call a few IETF Area Directors to see whether there is any interest in IDC standardisation.

20080119-11 Gigi Karmous-Edwards to discuss the role of GOLEs together with the Technical Working Group.

20080119-12 Evangelos Chaniotakis to upload current WSDL of DICE to GLIF Wiki.

20080119-13 Tomohiro Kudoh to upload current WSDL of GNS-WSI API to GLIF Wiki.

20080119-14 Mathieu Lemay to upload current WSDL of UCLP to GLIF Wiki.

20080219-15 Erik-Jan Bos and Gigi Karmous-Edwards to propose that Technical and Control Plane Working Groups be merged, and to suggest a name for the new group.