3rd Annual Global Lambda Grid Workshop

Distributed November 8, 2003

National Library, Reykjavik, Iceland

August 27, 2003

Table of Contents

1. Attendees

Erik-Jan Bos	SURFnet	bos@surfnet.nl
Heather Boyles	Internet2	heather@internet2.edu
Maxine Brown	UIC/ StarLight	maxine@uic.edu
Vinton G. Cerf	MCI	vinton.g.cerf@mci.com
Peter Clarke	U College London/ UKLight	clarke@hep.ucl.ac.uk
Steve Corbató	Internet2	corbato@internet2.edu
Cees de Laat (organizer)	University of Amsterdam	delaat@science.uva.nl
Tom DeFanti	UIC/ StarLight	tom@uic.edu
David Foster	CERN/ DataTAG	david.foster@cern.ch
Tom Greene	US NSF	tjg@mit.edu
Jan Gruntorad	CESnet/ CzechLight	jan.gruntorad@cesnet.cz
René Hatem	CANARIE	rene.hatem@canarie.ca
Saethór L. Jónsson	NORDUnet/ NorthernLight	slj@hi.is
Akira Kato	WIDE	kato@wide.ad.jp
Joe Mambretti	Northwestern U/ StarLight	j-mambretti@northwestern.edu
Olivier H. Martin	CERN/ DataTAG	Olivier.Martin@cern.ch
Osamu Nakamura	WIDE	osamu@wide.ad.jp
Kees Neggers (chair)	SURFnet	Kees.Neggers@SURFnet.nl

Harvey Newman	Caltech/ DataTAG	newman@HEP.CALTECH.EDU
Marius Olafsson	NORDUnet/ NorthernLight	marius@hi.is
Dennis Paus	SURFnet	dennis.paus@surfnet.nl
Malcolm Read	JISC/ UKLight	M.Read@jisc.ac.uk
David Richardson	U Washington/ NLR	drr@washington.edu
Don Riley	IEEAF/ USAwaves	drriley@umd.edu
Markus Sadeniemi	NORDUnet/ NorthernLight	Markus.Sadeniemi@csc.fi
David Salmon	UKERNA/ UKLight	D.Salmon@ukerna.ac.uk
Larry Smarr	UCSD/Cal-(IT)^2	lsmarr@ucsd.edu
Bill St. Arnaud	CANARIE	Bill.St.Arnaud@canarie.ca
Karel Vietsch	TERENA	vietsch@terena.nl
Peter Villemoes (host)	NORDUnet/ NorthernLight	Peter.Villemoes@nordu.net
Steve Wallace	Indiana U	ssw@indiana.edu
Linda Winkler	ANL/ TeraGrid/ StarLight	winkler@mcs.anl.gov
Alicia Wise	JISC/ UKLight	a.wise@jisc.ac.uk

Invited, but could not attend:

Tomonori Aoyama	U Tokyo/ JGN	aoyama@mlab.t.u-tokyo.ac.jp
Shoichiro Asano	SuperSINET	asano@nii.ac.jp
Dai Davies	DANTE/ GÉANT	Dai.Davies@dante.org.uk
Bob Day	UKERNA/ UKLight	R.Day@ukerna.ac.uk
Jón Ingi Einarsson	NORDUnet/ NorthernLight	jie@rhnet.is
Bengt Gördén	NORDUnet/ NorthernLight	bengan@sunet.se
Ron Johnson	U Washington/ PacificWave/ NLR	ronj@cac.washington.edu
Xing Li	CERnet	xing@cernet.edu.cn
Michael McRobbie	Internet2/ TransPAC	mcrobbie@ovpit.ucs.indiana.edu
Jun Murai	WIDE	junsec@wide.ad.jp
Tony Rimovsky	NCSA/ TeraGrid	tony@ncsa.uiuc.edu
Ken-ichi Sato	NTT	kenichi@exa.onlab.ntt.co.jp
Rick Stevens	ANL/ TeraGrid	stevens@mcs.anl.gov
Douglas Van Houweling	Internet2	DVH@internet2.edu

2. Introductions

2.A. Agenda

In 2001, TERENA hosted the first, Lambda Grid meeting (invition only), followed by an open Lambda Grid Workshop; the first lambda was on order. In 2002, iGrid 2002 was an open event, followed by the second, invitation-only Lambda Grid meeting. In 2003, we are holding the third closed meeting of the academic community with truly global attendance.

We have lots of opportunities, but lots of challenges – which is today's agenda. We are splitting into three groups:

- ***Technical Issues (TEC)*** – Erik-Jan Bos
- ***Governance and Growth (GOV)*** – Kees Neggers
- ***Research and Applications (RAP)*** – Cees de Laat

2.B. Cees de Laat Presentation

See Appendix A for a copy of de Laat's Presentation.
<DeLAAT.talk-03-08-27-lamdaWS.ppt>

Overall goal: How does this group take its next steps towards an international Lambda Grid?

Proposed problem statement for the *Research and Applications (RAP)* group:

- Demonstrators for SC2003 and other conferences
- Layer 1 versus 2 model: what services does the user community want?

Proposed problem statement for *Technical Issues (TEC)* group:
- Connectivity requirements, equipment (WAN PHY, switches)
- Functionality, services

Proposed problem statements for the *Governance and Growth (GOV)* group:
- Goals for next year, in terms of Lambdas, connections, and application support
- Governance: SLA's, SLS's, cross-domain lambda policies

2.C. Discussions

Does the TEC group know enough about the requirements of the RAP group in order to meet separately? Alternatively, what can the TEC group produce in two years? How far can it stretch the technology? This defines a "practical" limit. We discussed combining RAP and TEC, but there was concern that the group would be too large. The three groups will need to interact.

Has there been a discussion of having a broadband packet interface we can steer, or having a multiple-interface architecture? We need a consistent system picture, from the traditional to the experimental. Hypothetically, if the "light path" were extended to the workstation, it would be dual port. There are two models: at the edge of the network and in the system.

The networking and applications people need to work together to build the system. Can we aggregate the various testbeds (TransLight, NetherLight, StarLight) where these people work? We need to allow for parallel experiments. The TEC group needs to secure basic building blocks (end-to-end lambdas); what's done on this can be undefined. For this effort to succeed and interest vendors, it just can't be switched lambdas – it *must* come in contact with production networks, research networks, etc.

What do we want to come out of today's meeting? What is everyone's level of expectation? From the TEC group, maybe they should just look at all the ways of providing end-to-end light paths.

3. Morning Breakout Reports

See Appendix B for detailed Breakout Notes and Reports.

3.A. Governance and Growth (GOV)

Kees Neggers, Chair
Joe Mambretti, Scribe and Presenter
Participants: Vint Cerf, Steve Corbató (morning), David Foster, Saethór Jónsson, Joe Mambretti, Kees Neggers, Malcolm Read, Don Riley, Markus Sadeniemi, Peter Villemoes

3.B. Research and Applications (RAP)

Cees de Laat, Chair and Presenter
Maxine Brown, Scribe
Participants: Cees de Laat, Alicia Wise, Heather Boyles, Maxine Brown, Harvey Newman, Tom Greene, Larry Smarr, Peter Clarke

3.C. Technical Issues (TEC)

Chair and presenter: Erik-Jan Bos
Scribe: Dennis Paus
Participants: Erik-Jan Bos, Steve Corbató (afternoon), Tom DeFanti, Jan Gruntorad, René Hatem, Akira Kato, Olivier Martin, Osamu Nakamura, Marius Olafsson, Dennis Paus, David Richardson, David Salmon, Bill St. Arnaud, Karel Vietsch, Steve Wallace, Linda Winkler

4. Morning Group Discussion

RAP proposed something larger than a Lambda Grid – a *World Optical Networking Laboratory for the Sciences*.

In addition to development activities, the TEC group had more practical concerns, such as: UK will connect; NY's MAN LAN connection will come up; Seattle/Japan link; CA*net4 integration across North America; NLR; USAWaves – will allow us to interconnect a global Lambda Grid testbed. There will be a lot of activity in Europe in two years (2005), with EU 6th Generation Programme and GÉANT II.

If we take the RAP process – for example, medical imaging – which requires 2-3 sites in the USA be lambda-tized for connectivity to the UK, Amsterdam and Tokyo, then it will put pressure on NLR. NLR is only building to cities, not researchers. Building out to specific facilities is important. Clarke also advocates HPC applications.

Are there plans to extend the Lambda Grid into Japan in the next 12 months? SuperSINET provides access to scientists, and WIDE is extending its backbone.

From an engineering view, what are we hooking lambdas into? We don't just need computing clusters; we also need clusters optimized for data and for visualization. There are several models: the TeraGrid, the Newman model of a small cluster with 10Gb NICs, optically switched systems, etc. Smarr sees a commonality; we need a lab's eye view. We need to get those labs in our experiment so they can learn to spec the right systems. (Note: HPC community in Manchester, UK, does not have COTS clusters.)

Within a few months, we should identify application candidates: people/labs, and identify any last-mile issues.

Working groups need to follow up.

- ***TEC –*** What links are coming and how will we connect them? Include discussions with NLR and USAWaves on coupling.
- ***RAP –*** Identify people we want to contact, and invite them to investigate opportunities with us.
- ***GOV –*** Mambretti's notes have a list of issues.

NLR (Steve Corbató): Internet2 will be allocated one of the four initial NLR lambdas, and it is Internet2's intent to put it in the Lambda Grid space for 5 years. Corbató wants to tie NLR together with the Europeans, the Asians and the Canadians. With Erik-Jan Bos and René Hatem, Corbató will put a TDM multiplexer in NY as the first step.

We should involve application people who are already playing in the Lambda Grid space. Be opportunistic. The OptIPuter has Mark Ellisman (medical imaging) and John Orcutt (Earth Sciences).

EU Data Grid presentation at NORDUnet 2003 had a nice timeline of when various applications will come "on board" with respect to needing large bandwidth.

Clarke asks, that if we find applications, develop an experiment, come up with milestones – does this group agree to run these demonstrations on TransLight? Neggers says that TransLight is committed; we are cooperating now for SC 2003. The milestones need to be agreed upon by the applications people and the networking people, so we get usable and useful results. Other showcase events are Telecom World 2003, Geneva (joint sessions with Internet2 Member Meeting) in October and World Summit Information Society (WSIS) in Geneva in December. One problem with conferences is getting the appropriate infrastructure at a conference.

UK is proposing *proof-of-benefit experiments* to help their researchers get a competitive advantage and attract others in their field. They do not want to do demos that aren't persistent. Moreover, the scientists should give a presentation at the next Lambda Grid/e-Science conference on the benefits achieved.

GOV issues:

- There is a worldwide commitment to make the Lambda Grid evolve and grow. For governance, what constitutes a "member" and what does it take to "bring something to the table"?
- TransLight has such information in its *Operating Principles and Procedures* document; how does this group relate to the TransLight group?
- Newman again proposed that we adapt the name *World Optical Networking Laboratory for the Sciences.* Another term is Global Lambda Grid (GLG). Also, Gigabit Optical Development (GOD). (Note: Whatever label, we need to very well define the label and have a common message.) *(Newman suggests avoiding the terms "grid" and "light," as the former sounds limiting and the latter sounds elitist.)*

Smarr notes that creating an entity requires a lot of political overhead, too. Perhaps, as an interim step for the coming year so we can get some science done, we use TransLight – as the core of an emerging, living laboratory. Perhaps a working group of TransLight?

TEC group:

- Provision lightpaths
- Look at links coming up
- Look at how we will provision
- Look at NY
- Look at web page requirements

RAP group:

- Define the process

5. Afternoon Breakout Reports

See Appendix C for detailed Breakout Notes and Reports.

5.A. Governance and Growth (GOV)

Kees Neggers, Chair
Joe Mambretti, Scribe and Presenter

Name of initiative: **Global Lambda Integrated Facility (GLIF).**

There will be a website that can link to all participating groups (TransLight, StarLight, NetherLight, etc.) This is a cooperative activity; there shouldn't be rigid legal and/or hierarchical structures. For the coming year, we think we can survive with a few volunteers to encourage growth and development.

5.B. Research and Applications (RAP)

Cees de Laat, Chair and Presenter
Maxine Brown, Scribe

Discipline	Spokesperson	Who's doing contact
Particle physics – CMS	David Stickland, others	Peter Clarke, Harvey Newman
Particle physics – Fermilab/NIKHEF		Peter Clarke, Cees de Laat
Particle physics – ATLAS (worldwide)		Bill St. Arnaud
VLBI	Steve Parsley, Harvey Butcher, Richard Hughes-Jones, Alan Whitney (Haystack)	Peter Clarke, Cees de Laat
Medical	Mark Ellisman	Larry Smarr, Maxine Brown
Earth Science	John Orcutt	Larry Smarr
HPC – RealityGrid/TeraGrid	Peter Coveny, Rick Stevens	Peter Clarke
HPC – Remote steering	Ed Seidel	Peter Clarke

5.C. Technical Issues (TEC)

Chair and presenter: Erik-Jan Bos
Scribe: Dennis Paus

Erik-Jan explained a connectivity diagram on the whiteboard. (See Appendix C.)

6. Afternoon Group Discussion

Persistent – how long is that? It should be as long as needed to get experiments done. Rather than say how long that is, it means that circuits don't end with an event. We need to caution, however, that we cannot lock up circuits for one application. People need to share.

A *vision* statement needs to be crafted. Eventually, a mission statement should also be crafted. [Note: After the meeting, Malcolm Read wrote and submitted the following vision statement for our consideration.]

> **Modern [*Advanced?*] research requires new forms of network connectivity. High bandwidth (typically of 1Gbps and more) needs to be dedicated to applications for fixed periods (typically one hour to one day [*or whatever*]). This bandwidth provision (often referred to as "lambdas" from the Greek symbol for wavelength) must be supported by a range of applications to provide secure and controlled access to online resources; these applications are an integral part of the network.**
>
> **The Global Lambda Integrated Facility (GLIF) joins together the world's most advanced lambda networks to provide a flexible infrastructure to meet the increasingly demanding needs of researchers. GLIF will, through international collaborations and partnerships of network providers and the research community, form the basis of a new Internet support infrastructure. GLIF will initially support the challenging applications of e-Science [*GRID? cyber-infrastructure*?], but will be designed to meet the broader needs of research and education in years to come.**
>
> **A major goal of GLIF will be to integrate existing developments in networking and information management [*define "middleware" here?*] and to define the necessary technical and operational standards to achieve this objective.**

Mailing lists will be created so we can keep in touch on the various topics. We should have ongoing discussion threads – email is less efficient than today's face-to-face meeting, but vital.

Maybe the various working groups (not a full meeting) should meet face-to-face before another year goes by? The TEC Group needs a lot more time and should schedule a meeting in the near future. Groups should organize around other conferences (GGF, APAN, etc.).

Middleware should be built into our thinking from the beginning. Malcolm Read likes the word "Facility" in the GLIF title, as it's more than a Network that we're building.

Direction for the next few years is to build up this Facility into a System that can be used to accomplish real work.

Today we developed a paradigm for how to proceed with applications, which is very powerful. GLIF is a *Laboratory* (an experimental learning experience) and for the sciences and engineering (and the broader research community).

7. Next Meeting

There will be another Lambda Workshop next year in September, hosted by UKLight in London. It will be held in conjunction with an e-Science meeting.

8. Appendix A: Introduction/ Cees de Laat Presentation

<DeLAAT.talk-03-08-27-lamdaWS.ppt>

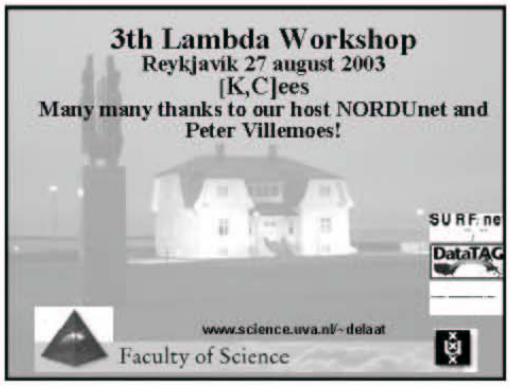

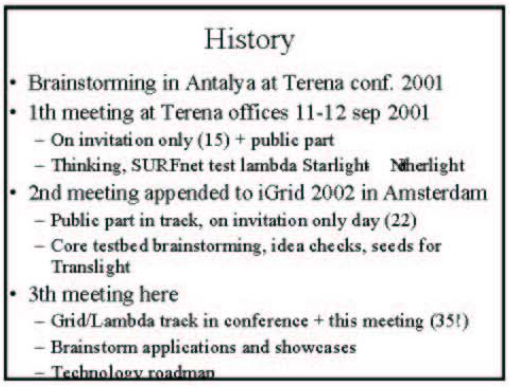

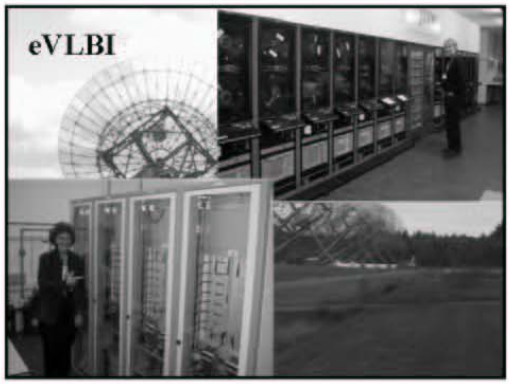

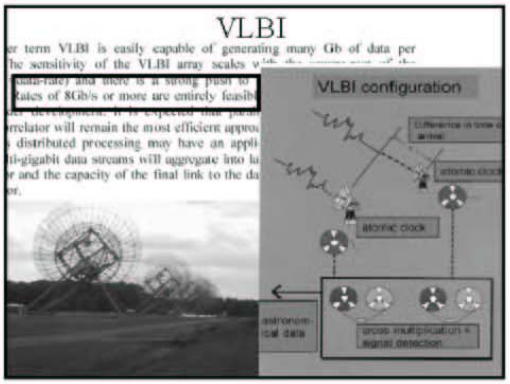

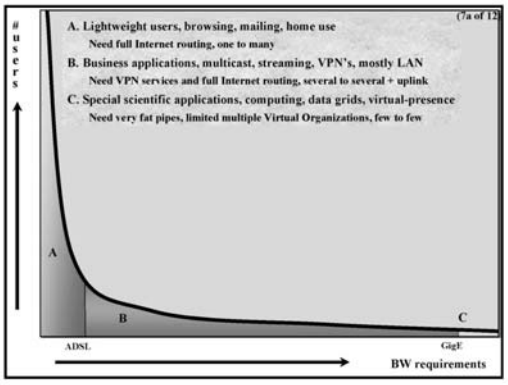

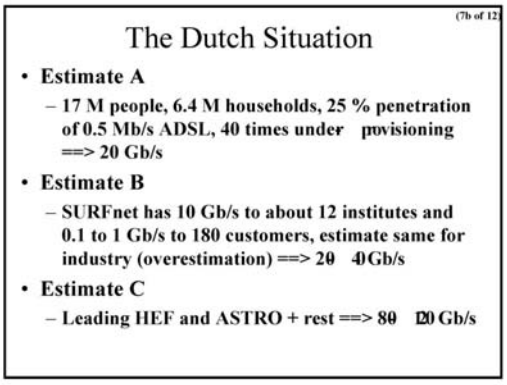

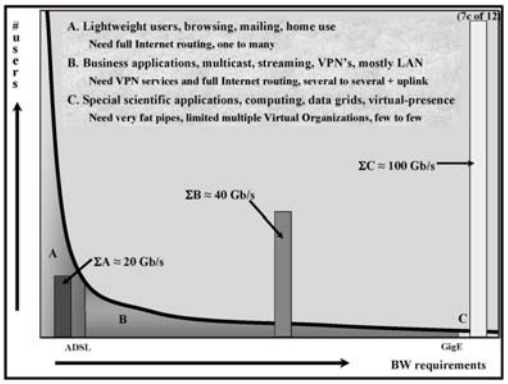

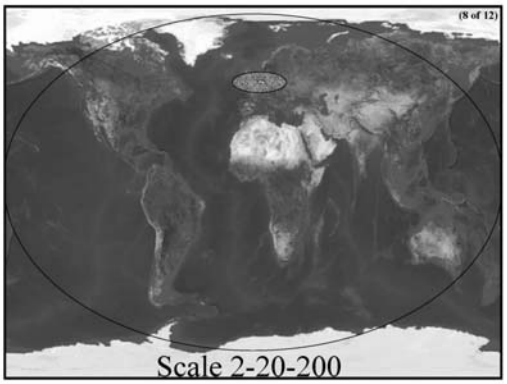

Services

SCALE / CLASS	2 Metro	20 National/ regional	200 World
A	Switching/ routing	Routing	ROUTERS
B	VPN's, Switching (G)MPLS	VPN's Switching Routing	ROUTERS Switches
C $\#\lambda(rtt) \approx \frac{200 e^{(t-200)}}{rtt}$	dark fiber Optical switching	Lambda switching	Sub lambdas, ethernet sdh

UVA/EVL's 64*64 Optical Switch @ NetherLight in SURFnet POP @ SARA Costs 1/100th of a similar throughput router but with specific services!

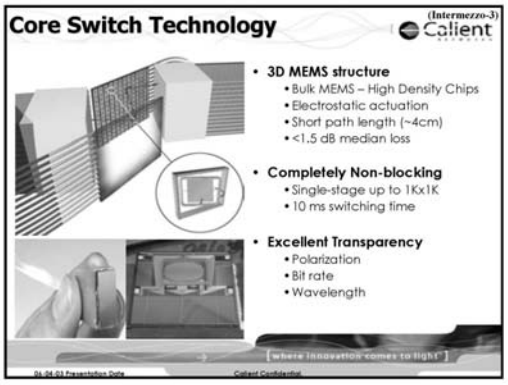

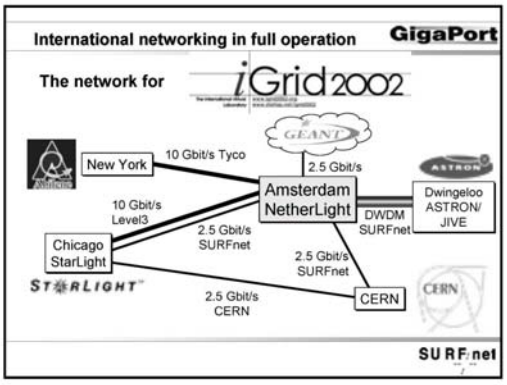

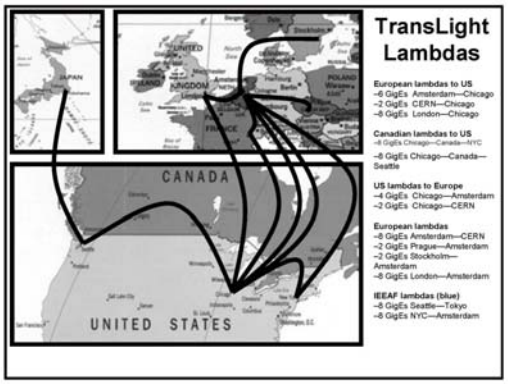

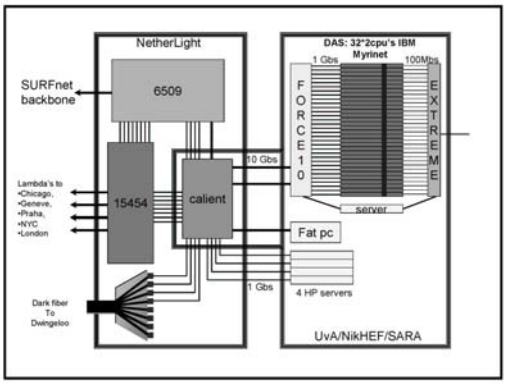

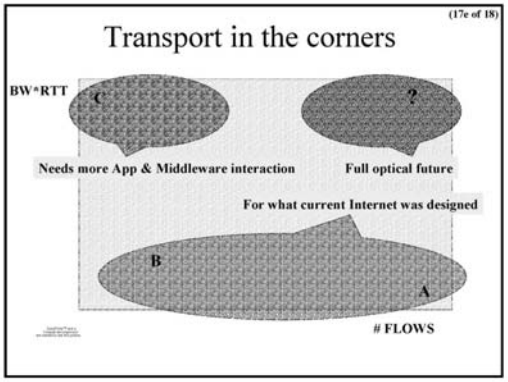

Agenda

chair Kees Neggers
- **09h00 introduction by chair**
- **09h15 agenda bashing**
- **09h20 introduction of participants**
- **10h00 format of subgroup discussions , three subgroups:**
  (RAP) research and applications, chaired by Cees de Laat
  (TEC) technical issues, chaired by Erik-Jan Bos
  (GOV) governance and growth issues, chaired by Kees Neggers
- **10h30 break**
- **10h45 formation of subgroups, work on the problem statements (next slide)**
- **12h30 break for lunch**
- **13h30 reconvene, report from RAP, GOV and TEC**
- **15h00 testbeds and deployments on 1, 2, 5 year scale**
- **15h30 break**
- **15h45 expected mid and long term developments, scaling up grid**
- **16h30 identify technical work to be done and establish working groups for that**
- **17h15 discussion about future of this kind of meeting**
- **17h30 end**

Problem statements & expectations

* overall goal: how to make next step towards an international Lambda Grid

* proposed problem statement for RAP subgroup:
  - Demonstrators for SC2003 and other conferences
  - layer 1 vs 2 model, services wanted by user community
* proposed problem statement for TEC subgroup:
  - connectivity requirements, equipment (wanphy, switches)
  - functionality, services
* proposed problem statements for GOV subgroup:
  - goals for next year in terms of Lambda's, connections, application support
  - governance: SLA's, SLS's, cross domain Lambda policies GOV

9. Appendix B: Morning Breakouts

9.A. Governance and Growth (GOV)

<MAMBRETTI.Global_Lambda_Grid.ppt>

Kees Neggers, Chair
Joe Mambretti, Scribe and Presenter
Participants: Vint Cerf, Steve Corbató (morning), David Foster, Saethór Jónsson, Joe Mambretti, Kees Neggers, Malcolm Read, Don Riley, Markus Sadeniemi, Peter Villemoes

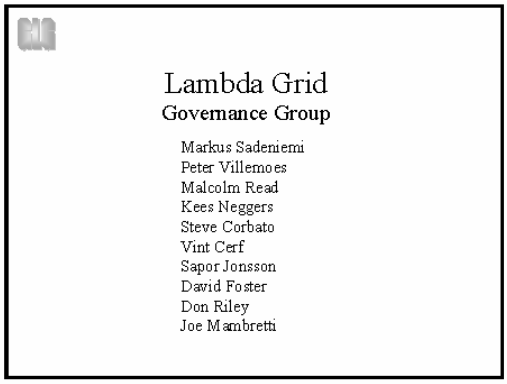

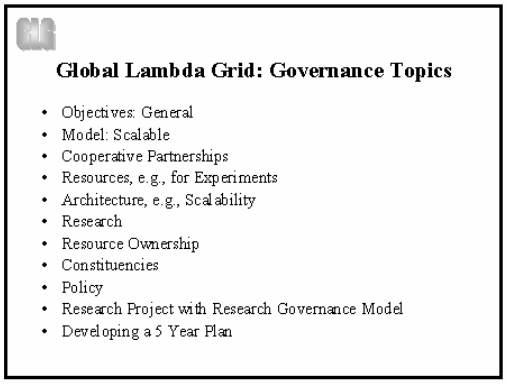

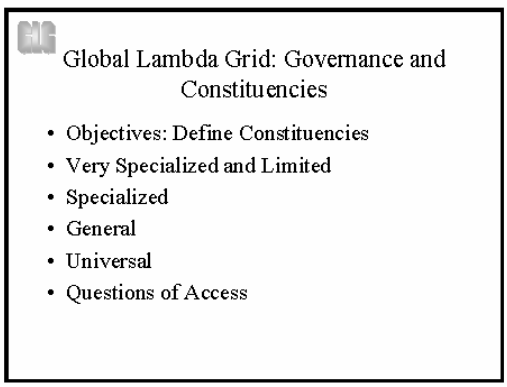

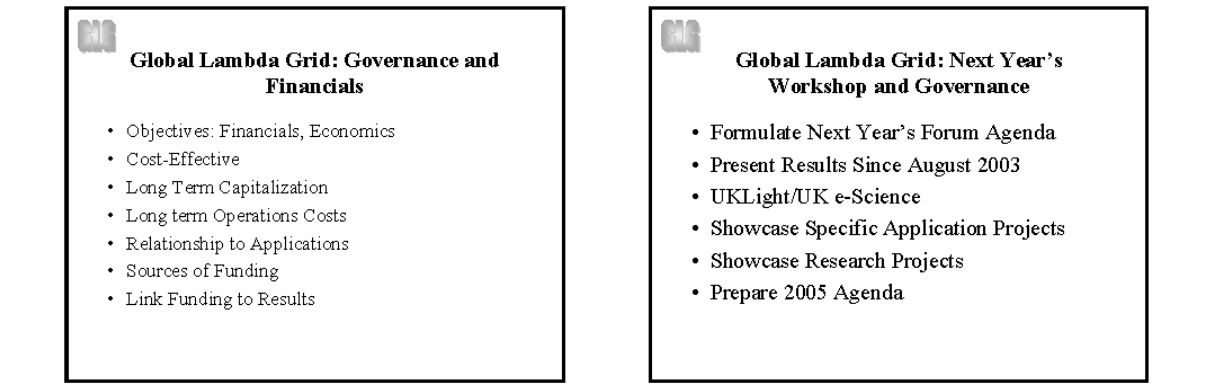

9.B. Research and Applications (RAP)

Cees de Laat, Chair and Presenter
Maxine Brown, Scribe
Participants: Cees de Laat, Alicia Wise, Heather Boyles, Maxine Brown, Harvey Newman, Tom Greene, Larry Smarr, Peter Clarke

Discussions

Two approaches:

- *Research on Lambdas:* Simulation, Control Plane, Transport Protocols, etc. Network research: Schema, search schemes, semantic tools, creating lambdas, protocols, circuit switching, etc.
- *Lambda-Enabled Research; Application Expectations (requirements):* amount of bandwidth? (We can only give minimums, such as 1GigE.)

How do we connect to the user group? Find good users?

How much exposure to lambda details does an application want/need?

Do applications require connections that are point-to-point, host-to-host, server-to-host, site-to-site, etc?

From grid efforts, application people don't want to see many details, unless there are strategic consequences. They really don't use toolkits. *Grid-enabled Analysis Environments* (Newman) and *Quanta* (Leigh) advocate intelligent systems that require low-level access to make choices, informing the user when issues arise; e.g., layered complexity encapsulation and transparency.

- Little complexity, large user group.
- Lots of complexity, small user group.

The integration of networks and grid technologies (creation of a Lambda Grid) toward grid applications needs to be discussed.

In the coming year, Clarke wants to have important people see useful application results from using lambdas. In the UK, user communities are: particle physicists, astronomy, high-performance computing community (computational steering) and medical imaging. Initially, one person per group needs to get involved.

- Success story for particle physics: simulations requiring Monte Carlo production farms (distributed computing)
- Radio astronomy
- HPC: (Clarke needs to poll UK people, concerned with the major HPC centers, joining TeraGrid, do protein folding, chemical systems, etc.)
- Medical imaging (on-demand connections and large demand for bandwidth for a short time).

Work on a simple XML service request.

Query estimators.

Need demonstrators to show how you apply intelligence to decision tree. There are those decisions that deal with networks (can it run on production links or wavelengths?).

We need to explore application *processes,* not the applications themselves.

Engage with a few user communities and set up testbeds. What do application communities want to accomplish in a year?

Have early adaptors do demonstrations (SC, iGrid) to show broader communities in order to get them engaged in R&D efforts.

Two goals:

- Masses can use GigE connections to do 10GB transfers
- For high-end users who want to use TBs of data, use 10GigE paths over long distances

Greene notes that we use the terms *lambda* and *grid* – so we need to be sensitive when working with grid communities. What are we going to do with the grid vector as well as the lambda vector (Foster and Kesselman don't agree with the term Lambda Grid).

Smarr made the analogy of getting users to adapt new technologies – he's seen this before, from getting early adopters 15 years ago for supercomputers, to supernetworks today. We need to pick the *alpha* person who is in charge of a long-term interactive activity and get them intellectually engaged so they become advocates. At supercomputer centers, there are the hardware and systems software types – those are the types of people at this meeting (NRN people; group #1 below) and application people (group #3). At supercomputer centers, there are also integrators (computers, mass stores, visualization, etc). Those people are missing here. Three groups; users cannot do their work until the first 2 groups are engaged.

- Group #1: Lambda interconnectivity and interoperability
- Group #2: "Tuning" systems for labs, users (the Human User Interface)
- Group #3: Users, asking scientific questions that cannot be asked

The supercomputer community made the transition from a small number of users to lots of users by building the supercomputer centers – maintenance, integration, training, etc. A few places are emerging: StarLight, NetherLight, CERN. Can we pick 4-5 user communities and bring them to these environments to do their work?

Smarr advocates forming an organization to crystallize what we are saying…to put a name on it (e.g., World Optical Networking Laboratory for the Sciences)…this makes it clear to the scientific community that we are serious.

Clarke wants to see a policy statement come out of this meeting to put more energy into the application communities and doing demos in the course of the year. Radio astronomy people are already doing it. Particle physics is difficult; Newman and Clarke can do things about it. (Need an LHC experiment, particularly CMS.) In the UK, the HPC community is very interested; but, needs to be organized. In the UK, Clarke knows e-health people, but this needs to be organized. (Smarr is an advisor to the NIH, so he can help.) Smarr also advocates the Earth Sciences; he is the only IT person on the NASA Council.) *(Having these insiders – Newman, Clarke, Smarr – is good. Need to work with other insiders for the other scientific communities.)*

Boyle notes that we're liaisons between the technologists and the users. We need descriptions of what facilities we have and what we can offer applications people.

How do we do this correctly? How do we get scientists to use many lambdas? How do we avoid a disruption between packet switching and lambda usage? Brown pointed out that in past years, for demonstrations, UIC/EVL has been funded to deploy its students ("smart bombs") to work with the application scientists.

We need to approach communities and find out their requirements. Newman cautions against requirement studies, as applications people say they want it now, want it to work, and don't want their software affected. Newman says our goal should be to launch keystone applications. There are technical milestones we should achieve.

To develop a process for the next year: (a) name a global facility (North America/Europe/Asia), (b) define experiments to be done on that facility, and (c) achieve milestones developed by the applications community. We don't have to choose the milestones; we have to define the process. From San Diego/ Chicago/ Netherlands/ CERN/

Tokyo – define a World Optical Networking Laboratory for the Sciences; partner (join and co-evolve) the facility/technology groups with existing application groups; work with facilities to define milestones. This "Laboratory" isn't displacing services, but may help define future services. In these leading problems, there are technology drivers and application drivers. We need the convergence of technology-driven applications and application-driven technologies.

After three iGrids, we have a trusted community of network engineers, computer scientists and computational scientists.

TransLight is a subset (core) of the proposed World Optical Networking Laboratory for the Sciences.

Take an opportunistic approach to finding existing groups (that come with their own funding) to make use of this Laboratory; such as BIRN, EarthScope, OptIPuter, etc.

Smarr feels that NLR, SURFnet, etc. will want to contribute resources, to demo applications that can have commercial opportunities.

Conclusions

Process – Need milestones and roadmaps (specific experiments); Need to specify services in a common way (SLS's: Service Level Specifications).

Research efforts should be ongoing (need to go to an end-to-end "systems" approach to using lambdas; lambda switching by itself doesn't solve anything).

Laboratory, services, efforts (be opportunistic).

Go after applications!

In Summary…

Research: Research is going well in several areas, but we need a "systems" approach that includes applications.

Continuum: Given this core of lambdas (StarLight, NetherLight, etc.), how much should they extend into the user space? There are various classes of applications, from those that don't want to understand the networks to those that do. There is a continuum of how far the networking should extend into the application space.

Process: How do we get to demonstrate applications tied into this work? Researchers do work bottom-up, but we prefer doing it top-down, to have other application groups endorse the need for lambdas.

Select thought leaders of key applications. Need to be opportunistic, and rely on current relationships, and build on them during the coming year.

What sort of world-scale laboratory do we have? We have TransLight. Grow and evolve, and find applications that can benefit from this infrastructure. Define milestones.

In our world, there are 3 concentric circles: NRNs and network engineering at the core, and applications on the outer circle. We are missing the in-between piece: the service group. We must bridge that gap. We must deliver back to this group 3-5 applications that can, in one year's time, set milestones/requirements and showcase progress.

9.C. Technical Issues (TEC)

Chair and presenter: Erik-Jan Bos
Scribe: Dennis Paus
Participants: Erik-Jan Bos, Steve Corbató (afternoon), Tom DeFanti, Jan Gruntorad, René Hatem, Akira Kato, Olivier Martin, Osamu Nakamura, Marius Olafsson, Dennis Paus, David Richardson, David Salmon, Bill St. Arnaud, Karel Vietsch, Steve Wallace, Linda Winkler

Erik-Jan opened the meeting; the scope is *technology push!* Look ahead towards the bleeding/leading edge of what is possible and what we should accomplish between now and the next Global Lambda Grid Workshop. After the RAP breakout, the requirements of the applications will be known. Since the scope of the RAP working group is four years out, and the TEC working group won't go much further than 1-2 years, there's no conflict.

The agenda items to be covered by the working group, in addition to Cees de Laat's proposed agenda, are:

1. Status update.
2. Connectivity requirements, equipment, building blocks: We know what we use today. Are we satisfied with it? What additionally do we need? What are the building blocks regarding types of equipment?
3. Functionality and services: What do we want at the edges? What do we want (and foresee) connecting to, today and in the future?
4. Capacity: 1GigE and 10GigE. Are we able to deliver something in between, keeping more users on the 10Gb lambda?
5. Measurement and management (provisioning): What is required from infrastructure operators? What features are missing to do so?

Status update

David Richardson (National Lambda Rail, USA): NLR intends to build an infrastructure that supplies multiple waves – a DWDM system of up to 40 wavelengths, to be installed throughout the US. NLR will install both 10GigE switches (layer 2) but will not be limited to that for switching lambdas. Connectivity will be provided mostly to currently connected Abilene sites. NLR expects to support applications like HDTV, etc.

David Salmon (UKLight/UKERNA, London): UKLight will realize two international connections at 10Gb each, one to Chicago and one to Amsterdam, providing GigE light paths to end users. Also, a local switch will be installed that will be earmarked for testing with local researchers, not for production. UKLight/UKERNA also got funding (6.5-million pounds) to extend its national backbone to reach inside the country by adding wavelengths to current connections. UKLight/UKERNA is also providing dark fiber to national researchers, targeting 5-6 connected APS facilities. UKLight/UKERNA is looking for metro-like equipment for SDH with Ethernet encapsulation. They have a few Cisco boxes installed, but procurement is open to all equipment vendors. Spread of equipment will require a control plane on the Ethernet level to get light paths provisioned. Timelines: January/February international circuits ready. UK national circuits will come sooner, depending on equipment delivery. Circuits are already in place, equipment needed to get GigEs from it.

Akira Kato (WIDE Project, Japan): Tyco/IEEAF is providing an unprotected OC-192 and a protected OC-12 between Tokyo and Portland. NLR will provide transit between Portland and Seattle in the January/February 2004 timeframe. The WIDE Project is looking at how the OC-192 can best be subdivided into smaller light paths for researchers. In Tokyo, a lot of researchers already have access to light paths and the WIDE Project is now having negotiations with circuit providers to bring connectivity further into Japan. No confirmed date of when things will be operational.

Jan Gruntorad (CzechLight/CESnet, Prague, Czech Republic): An OC-48c came up earlier this year, but funding wise, it is difficult to extend the circuit further into the Czech Republic. CzechLight will be extended to the Grid community via 300km dark fiber, which is to be acquired. Currently, CESnet is in negotiation with Cisco for pre-GARDEN (6NET money) equipment. Furthermore, CzechLight/CESnet is looking at DWDM equipment for extending the Lambda Grid to Poland and Austria; however, nothing is certain yet.

Tom DeFanti (StarLight, Chicago, Illinois, USA): The TransPAC group (US/APAN) will dedicate 1GigE between Chicago and Tokyo (part of an OC-48c) for TransLight. At the StarLight facility, a Force10 will be installed, 6x10GigE and 24xGigE, for international networking. Also, a 128x128 Calient optical switch will be installed.

René Hatem (CANARIE, Canada): Canada has two OC-192s from Vancouver to Halifax and Vancouver to Seattle. CANARIE is also in Chicago and will be at the New York MAN LAN in September. Not all of Canada's regional research networks are able to extend light paths to the end users yet; some areas are more difficult than others. CANARIE is waiting with contracts for waves until the Tokyo lambda materializes. In September, CANARIE will have a meeting with its light-path provisioning software team, with perhaps an end-of-year working prototype.

Olivier H. Martin (CERN, Geneva): Currently there's an OC-48c circuit between CERN and Chicago. In September, it will be upgraded to 10Gb. CERN did not find a layer 1 product that could do sub-10Gb connections; i.e., no 10GigE LAN on long-distance SONET multiplexers. So, they bought two Juniper M20s to interface the conversions on layer3 and with layer2 Juniper CCC tunnels.

Erik-Jan Bos (NetherLight/SURFnet, Amsterdam): NetherLight is now at two locations: SARA and COLT (GÉANT's PoP in The Netherlands). The Euro-Link 10Gb from Chicago terminates at COLT, where 5Gb peers with GÉANT and 5Gb continues to NetherLight. SARA and COLT will soon be interconnected with 10Gb. The tender for SURFnet6 (2004-2007) is out; SURFnet6 ultimately wants to connect all SURFnet customers to light

paths (next to a port with layer3 functionality).

Discussion (Agenda items 2-5)

René: Interested in philosophy behind giving an end-user control over light paths. Do we want this at all? Or do we just want optical bypass? There are different techniques to switch light paths (layer1 and layer2). What's the difference, from both a network provider and an application's point of view?

Tom: At the end of today we should have a few goals to be reached this year.

David: We should further explore Vint Cerf's question about having multiple interfaces on hosts.

Steve W: Would like to make kind of internet draft with light path thinking.

Bill: Global Grid Forum (GGF) is doing something like that; he will forward it (but the essence will be slightly different.) Therefore, the TEC group should try to write something over the coming year.

René: We are thinking about using a martini tunnel (standardized Juniper CCC) for creating a layer2 GigE light path through the routed cloud of a commercial provider. Internet2 is thinking about using this technology through the Abilene backbone. POS is sometimes able to support larger frames than Ethernet framing over POS.

Steve would like to demonstrate martini light paths at SC2003; however, we need to have an application that requires (and demonstrates) it.

Question: Will there be different types of light paths possible in the Global Lambda Grid?
Answer: Perhaps, but the characteristics of each light-path type needs to be very clear in order to concatenate them and ensure that the end user gets what he/she expects. There will probably be more framing types in the future: Merinet, FibreChannel, SANs, etc. At least one type (Type A) will be defined now, based on Ethernet framing.
Conclusion: For SC2003, we will only look at Light Path Type A which uses Ethernet framing!

Problems result when connecting multiple Type A light paths (in different technological domains, but in one administrative domain!) behind each other:

- Support of jumbo frames?
- Tagged/untagged Ethernet?
- Guaranteed bandwidth?

Tom: Type A, today, is 1GigE and 10GigE. Can we have 1GigE or 10GigE interfaces and provide 3.2 or 5.1Gb? What is the time granularity of light paths; i.e., how dynamic will they be? We should not be handing out circuits forever.

Erik-Jan: What can we do to offer >1GigE paths to end users? CANARIE did some tests with Force10 10GigE WAN PHY and 10GigE LAN PHY connected to an ONS 15454. They could only push 90% of the LAN PHY speed, while always (regardless of how fat the pipe was) occupying the full OC-192 on the WAN side.

Conclusion: The light-path demarcation interface (e.g., 10GigE) should not be coupled to light-path capacity (e.g., 5.1 Gbit/s).

Either the application should avoid reordering (give them four light paths from 1GigE) or the light-path provider should take care (give them one interface and tunnel it over 4 light paths). **Definition:** One light path should always deliver packets in order. What about providing four light paths that do not travel the same route? Will applications have problems with that? What about martini tunnels that are asymmetrically routed?

Tom: We will not have a total control plane next year. How, then, are we going to process requests for setup of light paths? Users should not ask for too much time in order to give everybody access. Automate scheduling by using application forms and calendar with occupied light paths? What information is needed, such as:

- Endpoint A, Endpoint B (define *endpoint!* Address, room, GPS coordinates, ports?)
- Required demarcation interface
- Required bandwidth
- Support of jumbo frames
- Desired start date
- End date of the circuit

- Etc.

René: Would like to know if Abilene/NLR will do user-enabled light paths in the far future?
David: NLR will, in the foreseeable future, do static provisioning. End users should make cross connections dynamic.
Erik-Jan: How to deal with transit through NLR or Abilene? This is particularly difficult because regional networks also need to be crossed to reach the end users.

Summary

Three goals for next year:

- Work on Internet-style informational draft: put in different usage models (Light paths between routers, between end users), make definitions more accurate, also useful for applications people to find more information about the nature of light paths
- Define interconnection of various light path technologies
- Work on setup and tear-down of light paths; i.e., scheduling, automated provisioning, etc.

To do:

- Start informational Internet draft on light-path definition
- Research and document light-path interconnection (BCP)
- Set up a mailing list for the TEC Working Group with today's participants
- Set up a web page; from there, enable access to a lambda reservation tool

Presentation to Global Lambda Grid participants

The above discussions were summarized in PPTs, shown below.

<BOS.GLG TEC Reykjavik.ppt>

GLG TEC

Reykjavik, Iceland
27 August 2003

Participants

- StarLight/Euro-Link/…
- CA*net 4
- WIDE
- UCAID/IU/Abilene
- UW/NLR
- CERN
- CESnet
- UKERNA/UKlight
- SURFnet/NetherLight

Main Topic

- Look ahead towards the bleeding/leading edge on what is possible and how we should accomplish this
- Between now and the next Workshop

Topics

- Definitions
- Updates from participants
- Connectivity requirements
- Equipment (WAN PHY, switches)
- Building blocks
- Functionality, services - 1G en 10GE but also 6.34 Gbit/s (example)
- Management and measurements (provisioning)

Discussions

- Lightpaths type A (ethernet framed):
  - 1G
  - 10G
  - In between 1G and 10G, ways to do this
- Provisioning (set-up and tear down):
  - resources agenda
  - scheduling form with all details
  - User empowered?! Yes, don't know, perhaps…

Three goals

- document “lightpath” (including examples)
- define interconnection for various lightpath technologies
- scheduling/automatic provisioning

To Do

- Start draft on lightpath definition
- Document lightpath interconnection (BCP)
- Set up mailing list for GLG TEC
- Set up web page

10. Appendix C: Afternoon Breakouts

10.A. Governance and Growth (GOV)

Refer to PPTs in Section 9.A, which contains information from the afternoon breakout.
<MAMBRETTI.Global_Lambda_Grid.ppt>

Kees Neggers, Chair
Joe Mambretti, Scribe and Presenter

Name of initiative: **Global Lambda Integrated Facility (GLIF).**

There will be a website that can link to all participating groups (TransLight, StarLight, NetherLight, etc.)

This is a cooperative activity; there shouldn't be rigid legal and/or hierarchical structures. For the coming year, we think we can survive with a few volunteers to encourage growth and development.

10.B. Research and Applications (RAP)

Cees de Laat, Chair and Presenter
Maxine Brown, Scribe

We want persistent applications.

Particle physics: For CMS, Harvey Newman suggests first getting buy-in from the CMS software person and the computing person (David Stickland) by talking about proof-of-concept and cost benefits. Explain that this is a larger concept than one application, as configurable circuits are here to stay. Then, meet with the "spokesperson" and develop milestones. Other CERN groups would ultimately buy in. We need to be careful that the costs of diverting resources don't overwhelm them in a 12-month cycle. Competitiveness is also leverage we can use. Clarke will talk to Stickland. Clarke is also interested in Fermilab and the D0 experiment, and they work with NIKHEF.

While discipline discussions are appropriate now, geographical/regional discussions (such as extending to Berlin) may also be of interest downstream.

VLBI: UK (Richard Hughes-Jones)/JIVE (Steve Parsley) applications are taking place. Clarke/Hughes-Jones and de Laat will pursue. Can Japan be involved?

Medical: Smarr suggests that Mark Ellisman, UCSD, focus on NIH-funded BIRN (Biomedical Informatics Research Network) efforts. Smarr and Brown will talk with Ellisman, and talk with Tony Hey to find a contact in the UK.

Earth Sciences: John Orcutt, UCSD Scripps Institution of Oceanography, is involved with the NSF-funded EarthScope. NEPTUNE is a 3000-mile fiber-optic system off the Canadian/Northwest US group. NSF is about to put together a large Ocean Science project; Smarr could look at the CERN, VLBI and BIRN models, as he is on an NSF GEO committee to propose a model for this Ocean Science project. Smarr will talk with Orcutt to discuss his participation, and also involve the UK and/or Europe.

High Performance Computing community: (a) TeraGrid/Argonne National Laboratory application and/or (b) get a user, like Ed Seidel – remote steering/remote visualization. Clarke has a specific angle; a group in Manchester that does "RealityGrid" (Peter Coveny, Ron Pero) is talking with Rick Stevens next week. It sounds like we have two ideas in this area. Clarke and de Laat can follow up with Seidel and the "RealityGrid" group.

Discipline	Spokesperson	Contact
Particle physics – CMS	David Stickland, others	Peter Clarke, Harvey Newman
Particle physics – Fermilab/NIKHEF		Peter Clarke, Cees de Laat
Particle physics – ATLAS (worldwide)		Bill St. Arnaud
VLBI	Steve Parsley, Harvey Butcher, Richard Hughes-Jones, Alan Whitney (Haystack)	Peter Clarke, Cees de Laat
Medical	Mark Ellisman	Larry Smarr, Maxine Brown
Earth Science	John Orcutt	Larry Smarr

HPC – RealityGrid/TeraGrid	Peter Coveny, Rick Stevens	Peter Clarke
HPC – Remote steering	Ed Seidel	Peter Clarke

10.C. Technical Issues (TEC)

Chair and presenter: Erik-Jan Bos
Scribe: Dennis Paus

The Working Group discussed ways to engineer the upcoming links into the current lambda network. This involves:

- Tyco/Tokyo ↔ Portland 10Gb link
- Tyco/New York ↔ Amsterdam 10Gb link
- NLR
- Chicago ↔ Tokyo TransPAC GigE link
- UKLight
- CzechLight

Discussion

Steve C: Abilene will put an ONS 15454 in New York next to the MAN LAN switch. It will be connected to the CANARIE ONS (New York), SURFnet ONS (Amsterdam) and Abilene router at 10Gb POS. Erik-Jan suggested a scenario with a Lucent LambdaUnite in Amsterdam and New York. This looks rather promising and is worth further exploration. Other TDM multiplexer vendors might have interesting products: WhiteRock, Ciena, etc.

David will find out if Juniper supports multiple OC-48c virtual ports on one OC-192c interface and report back to the group.

Erik-Jan will find out more specifics about LambdaUnite from Lucent (10GigE LAN PHY roadmap, 10GigE LAN PHY occupying less than OC-192 bandwidth, use in WAN environment, use of concatenated circuits).

It is unclear how SuperSINET would like to connect in New York. This need to be checked with them.

For NLR, Steve C. would like to use the Internet2 wavelength circuits in the TDM space, working with Cisco to further explore ONS 15454 and beyond. Use of wavelength circuits for transit and/or national projects is still a blank piece of paper. For sake of clarity: in Q1of 2004, NLR could be seen like CANARIE – a 10Gb cloud connecting large cities on OC-192, 10GigE LAN PHY and switched GigE. Seattle and Chicago for sure. New York is being pushed but not certain yet.

See picture below.

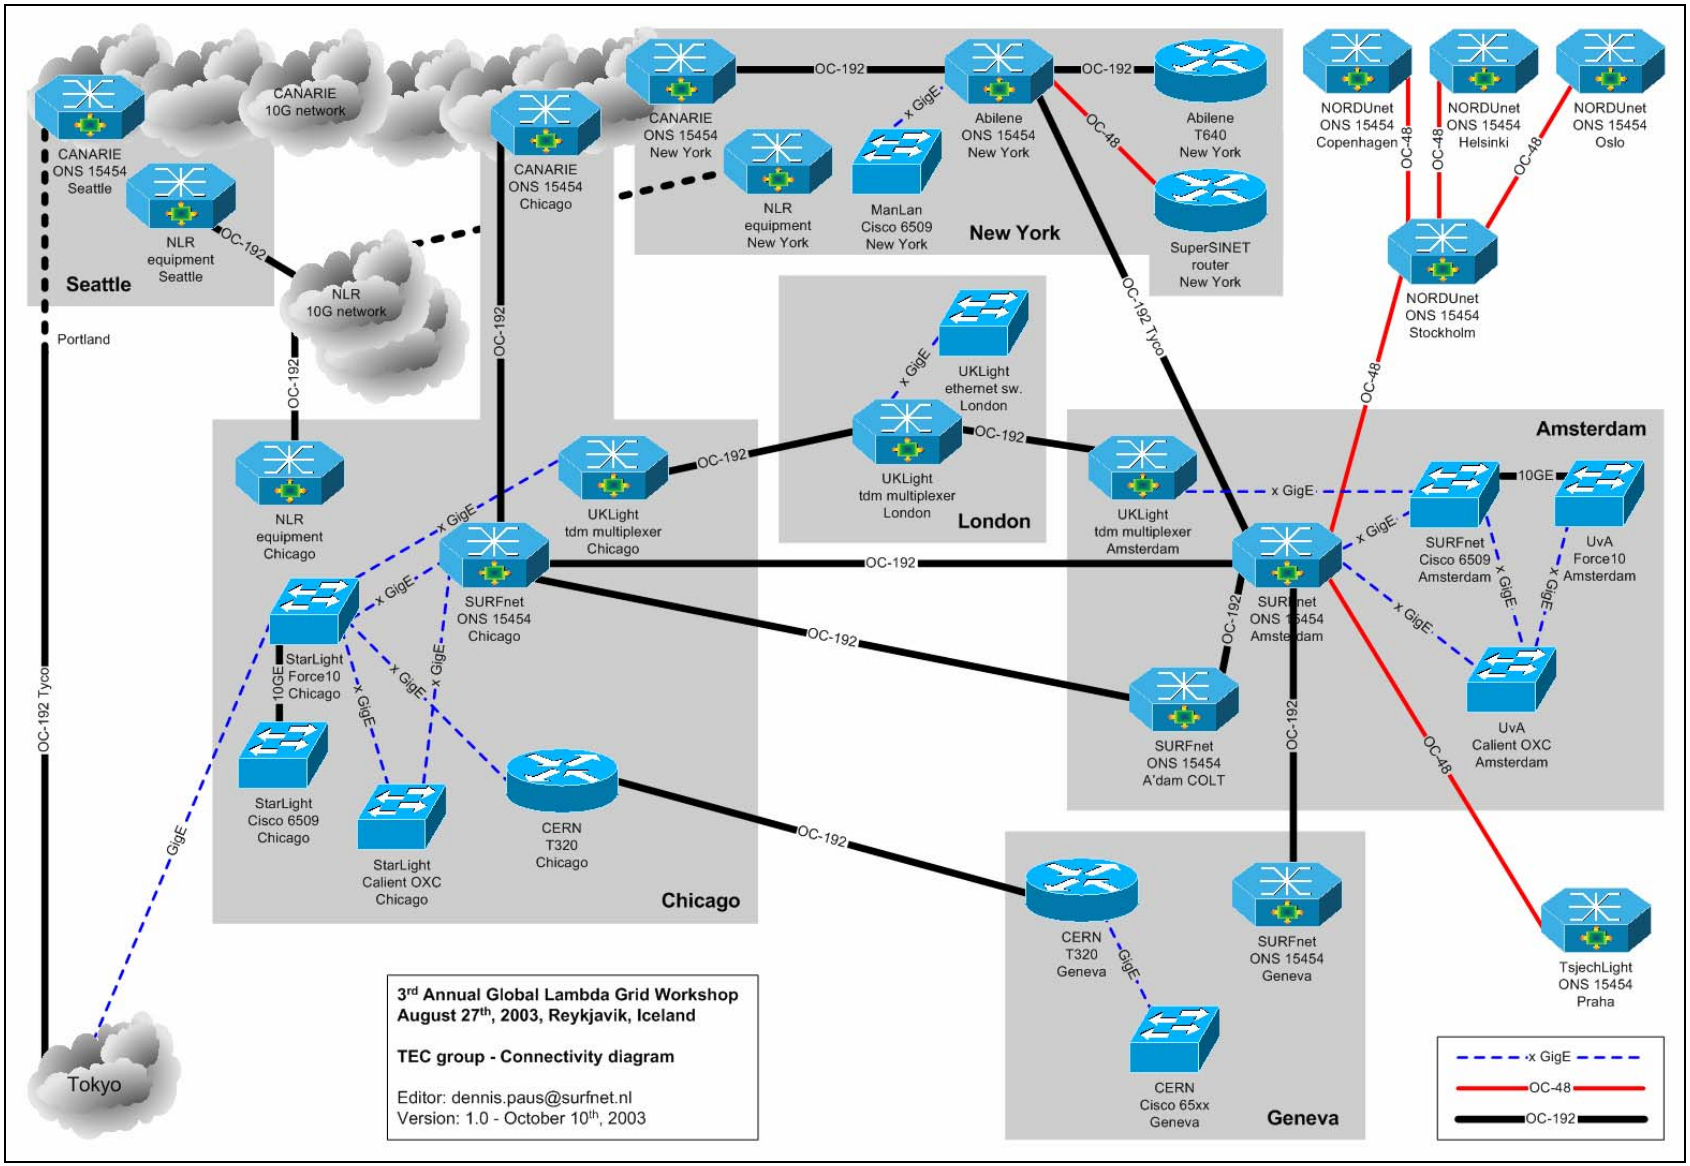
CANARIE
10G network
CANARIE
ONS 15454
Seattle
NLR
equipment
Seattle
OC-192
Seattle
NLR
10G network
Portland
OC-192
OC-192 Tyco
GigE
Tokyo
CANARIE
ONS 15454
Chicago
OC-192
CANARIE
ONS 15454
New York
OC-192
x GigE
Abilene
ONS 15454
New York
OC-192
Abilene
T640
New York
OC-48
NLR
equipment
New York
ManLan
Cisco 6509
New York
New York
SuperSINET
router
New York
OC-192 Tyco
NORDUnet
ONS 15454
Copenhagen
NORDUnet
ONS 15454
Helsinki
NORDUnet
ONS 15454
Oslo
OC-48
NORDUnet
ONS 15454
Stockholm
OC-48
UKLight
ethernet sw.
London
x GigE
UKLight
tdm multiplexer
London
OC-192
London
NLR
equipment
Chicago
x GigE
UKLight
tdm multiplexer
Chicago
OC-192
SURFnet
ONS 15454
Chicago
x GigE
StarLight
Force10
Chicago
10GE
StarLight
Cisco 6509
Chicago
x GigE
StarLight
Calient OXC
Chicago
CERN
T320
Chicago
OC-192
Chicago
UKLight
tdm multiplexer
Amsterdam
x GigE
Amsterdam
SURFnet
Cisco 6509
Amsterdam
10GE
UvA
Force10
Amsterdam
x GigE
SURFnet
ONS 15454
Amsterdam
OC-192
SURFnet
ONS 15454
A'dam COLT
UvA
Calient OXC
Amsterdam
OC-48
OC-192
CERN
T320
Geneva
GigE
SURFnet
ONS 15454
Geneva
CERN
Cisco 65xx
Geneva
Geneva
TsjechLight
ONS 15454
Praha
3rd Annual Global Lambda Grid Workshop
August 27th, 2003, Reykjavik, Iceland
TEC group - Connectivity diagram
Editor: dennis.paus@surfnet.nl
Version: 1.0 - October 10th, 2003
x GigE
OC-48
OC-192