IRNC ProNet: TransLight/StarLight

NSF Cooperative Agreement OCI-0962997
www.startap.net/translight

ANNUAL REPORT July 1, 2010 – June 30, 2011
Submitted June 20, 2011

Thomas A. DeFanti, Tajana Šimunić Rosing
University of California, San Diego
tdefanti@ucsd.edu

Maxine D. Brown
University of Illinois at Chicago

Joe Mambretti
Northwestern University

Table of Contents

1. Participants

1.A. Primary Personnel

Participant's Name(s)	Project Role(s)	>160 Hours/Yr
Thomas A. DeFanti	Principal Investigator (PI)	Yes
Tajana Šimunić Rosing	Co-Principal Investigator	Yes
Maxine Brown	Co-Principal Investigator (and PI, UIC subaward)	Yes
Joe Mambretti	Senior Personnel (and PI, NU subaward)	Yes

1.B. Other Senior Personnel (Excluding PI and Co-PI)

Participant's Name(s)	Project Role(s)	>160 Hours/Yr
Alan Verlo	Senior Personnel	Yes
Fei Yeh	Professional staff	Yes
Benjamin Ellis (1)	REU	Yes
Thirith Hout (1)	REU	Yes
Eric Michael Dorman (1)	REU	Yes

(1) UCSD REUs under the direction of Tajana Rosing for AY 2010-2011

1.C. Other Organizations That Have Been Involved as Partners

Metropolitan Research and Education Network
The Metropolitan Research and Education Network (MREN) <www.mren.org> is an advanced network that provides service to Research & Education communities in seven states the upper Midwest. MREN manages a metro-area optical networking facility located at StarLight. The MREN facility exclusively provides service and infrastructure support for large-scale data-intensive Research & Education activities; MREN does not provide general Internet services.

StarLight Consortium
The StarLight Consortium manages the StarLight International/National Optical Communications Exchange Facility <www.startap.net/starlight>. Consortium members consist of the International Center for Advanced Internet Research at Northwestern University, the Electronic Visualization Laboratory at the University of Illinois at Chicago, the California Institute for Telecommunications and Information Technology (Calit2) at University of California, San Diego, and the Mathematics and Computer Science Division at Argonne National Laboratory, in partnership with Canada's CANARIE and the Netherlands' SURFnet.

1.D. Other Collaborators or Contacts

CineGrid
CineGrid <www.cinegrid.org> is an international, interdisciplinary community that focuses on the research, development, and demonstration of networked collaborative tools to enable the production, use, preservation, and exchange of very-high-quality digital media over photonic networks.

Global Lambda Integrated Facility (GLIF)
GLIF <www.glif.is> is an international virtual organization of NRENs, consortia and institutions that promotes lambda networking. GLIF provides lambdas internationally as an integrated facility to support data-intensive application and middleware development. It brings together premier networking engineers to develop an international infrastructure by identifying equipment, connection requirements, and necessary engineering functions and services.

2. Activities and Findings

2.A. Research Activities

The goal of the IRNC ProNet: TransLight/StarLight initiative is ***to continue to expand upon experimental networking technologies through the development of several international communication services and data-intensive applications, leveraging existing collaborations to make significant science impacts.***

More specifically, the TransLight/StarLight team is involved in several efforts:

- *e-science application and middleware experiments across geographically distributed sites,* notably: CineGrid, GreenLight International, High-Performance Digital Media Network (HPDMnet), iGENI (International GENI), SAGE, and Science Cloud Communication Services Network (SCCSnet).
- *Cooperative community partnerships:* TransLight/StarLight leadership has established international partnerships, communication channels, forums, and processes to ensure ongoing successful interactions among all its constituents. Team members attend and participate in JET meetings, Joint Techs meetings, SCInet planning meetings, the annual SC conference, Internet2 Member Meetings, NLR All Hands Meetings, GLIF meetings, and meetings with representatives of science communities, international networking partners, equipment manufacturers, and IRNC partners.
- *Ongoing GLIF partnership:* TransLight/StarLight continues to: (a) participate with GLIF groups who are advancing dynamic provisioning techniques and technologies across single and multiple domains, (b) participate in the GLIF Technical, Governance and Research & Applications working groups, (c) participate in a joint GLIF/Open Grid Forum working group on open architectural standards for dynamic provisioning, (d) implement prototype services, (e) support science domain projects that use these services, and (f) update the GLIF map.
- *New services:* TransLight/StarLight assists new international services as they pertain to the development of CineGrid, GreenLight International, HPDMnet, iGENI, SAGE, SCCSnet and GLIF, and with other services as they emerge. Also, TransLight/StarLight continues to work with industry to utilize advanced networking services and install/test prototype equipment borrowed or bought with other funding.

TransLight/StarLight relies on links from multiple sources, including those provided by IRNC:ProNet awardees, international NRNs who bring their own 10Gbps links to the US, major consortiums such as GLIF, Federal agency networks, and others. TransLight/StarLight remains involved in activities related to the coordination and optimal implementation of IRNC circuits, in partnership with other members of the IRNC community.

Specific efforts and accomplishments related to the abovementioned efforts can be found in Section 2.B: Research Findings.

While some of the above efforts have other R&D funding, IRNC support leverages this investment and makes it possible to spend time expanding collaborations globally and assisting with documentation, tech support, education and outreach.

Some of the activities described here overlap with the IRNC TransLight/StarLight OCI-0441094 Final Report, which only had ~$5,000 in funds remaining. That funding, combined with this IRNC award, have provided adequate support to ensure successful outcomes.

2.B. Research Findings

2.B.1. e-science Application and Middleware Experiments

Below is a list of advanced application and middleware experiments that TransLight/StarLight has participated in and/or has provided network engineering support.

2.B.1.1. CineGrid

This map, produced by Keio University (Japan) in February 2011, shows the active CineGrid participants.

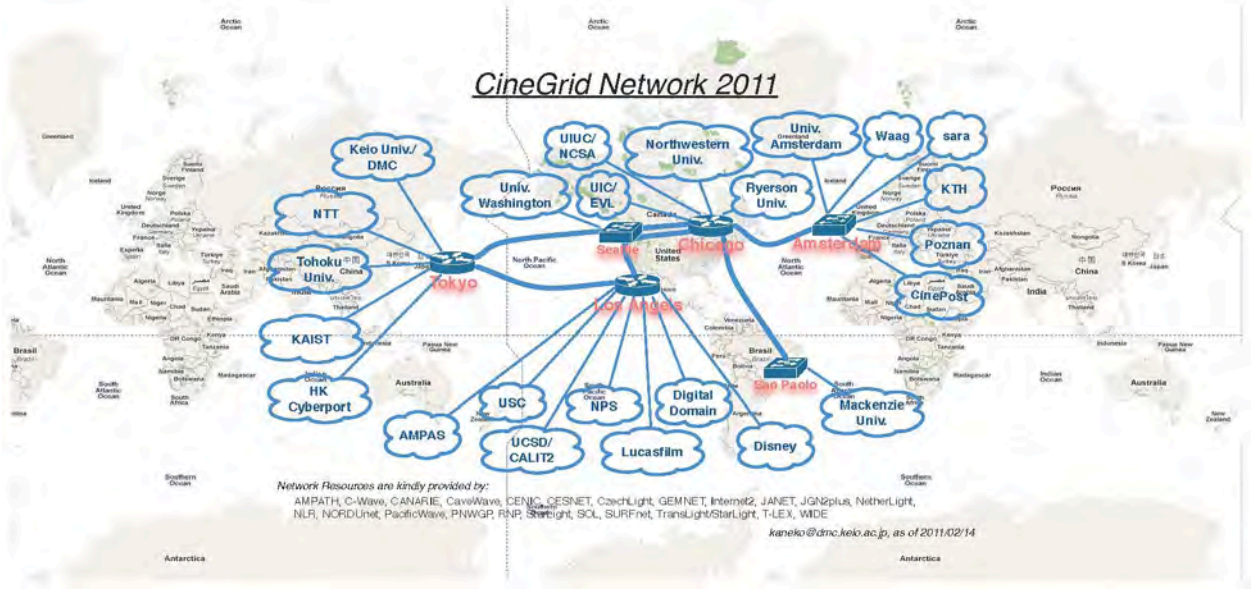

CineGrid Experiments 2010-2011
CineGrid @ TERENA Networking Conference 2011 https://tnc2011.terena.org/ *Collaborators:* • CineGrid; Calit2; US • CESNET; Czech Republic The TERENA Networking Conference 2011 took place in Prague, May 16-19. CESNET did a CineGrid demo, streaming JPEG 2000 compressed 4K movies from San Diego (Calit2) to the conference location in Prague.
CineGrid and iRODS http://www.cinegrid.org/index.php?option=com_content&view=article&id=73&Itemid=32 *Collaborators:* • CineGrid; Calit2; UIC/EVL; US • University of Amsterdam; Netherlands • Keio University; Japan The CineGrid Exchange (CX) is a distributed digital media repository designed to support CineGrid member-driven testbeds for digital media asset management, distribution and preservation applications. CX consists of digital media (visual and aural) of varying resolution, subject matter and format made accessible to CineGrid

members via iRODS (developed by Reagan Moore of RENCI) and secure high-speed networks (provided by GLIF).

iRODS had not been working correctly and remote CX nodes (in Amsterdam and Japan) were not updated properly. On May 10, 2011, Calit2 resolved difficulties upgrading its servers to the new version of iRODS, and was making progress converting CX assets in JPEG 2000 from the old NTT version to the new NTT version, which will be replicated at all CX nodes using updated iRODS rules.

CineGrid: 5th Annual CineGrid Workshop 2010

http://www.cinegrid.org/index.php?option=com_content&view=article&id=143:cinegrid-takes-digital-cinema-into-next-dimension&catid=1&Itemid=11

Collaborators:

- CineGrid; Calit2; US

December 12-15, 2010, Calit2 hosted the 5th Annual CineGrid Workshop. Tom DeFanti served as a member of the Program Committee. EVL participated: Luc Renambot gave the presentation "SAGE OptIPortable with iVisto for Collaboration," Jason Leigh gave the presentation "Big Picture Collaboration Using SAGE OptIPortals," Alan Verlo spoke about "Networking for Collaboration," and Dana Plepys spoke about the "CineGrid Exchange." In addition, Renambot conducted several tele-demos between Calit2 and EVL using SAGE software.

TransLight/StarLight funds cover some of the annual CineGrid workshop participant fees. It should be noted that the media and entertainment industries are actively involved in CineGrid to better understand the role of robust, stable large-bandwidth networks in applications from content creation and remote collaboration, and are beginning to adopt IRNC-style networks and high-end distributed telepresence technologies for their companies and their global suppliers. After CineGrid 2010, in a personal correspondence, one such company manager said, "The need for 'smart networks' that can help expedite increasingly large file-based, media dependent and capable applications can serve the needs of many industries. The infrastructure for this capability will absolutely be a competitive, creative and enabling technology for our needs and I am sure will have lasting and important benefits across many sectors of public and private interests."

CineGrid: CineGrid@TIFF 2010

http://www.cinegrid.org/index.php?option=com_content&view=article&id=142:tokyo-film-festival&catid=1:news&Itemid=11

Collaborators:

- CineGrid; Calit2; UIC/EVL; Disney Studios; Pacific Interface; US
- NTT Network Innovation Laboratories; Keio University; Japan

CineGrid members, including Tom DeFanti, organized and participated in the CineGrid@TIFF 2010 (Tokyo International Film Festival) one-day Workshop, October 28, 2010. Speakers conducted tutorials and a panel discussion explaining how information and computer technology (ICT) is changing digital media post-production, distribution, exchange and preservation. DeFanti gave the presentation "Collaborative Visualization for Science, Education and Entertainment." The event also showcased HD/4K live streaming.

CineGrid: CineGrid@Disney 2010 (Includes SAGE demonstration)

www.cenic.org/publications/cenictoday/20110201CT.html#5

Collaborators:

- CineGrid; Calit2; UIC/EVL; Disney Studios; Pacific Interface; US
- NTT Network Innovation Laboratories; Keio University; Japan

On October 13, 2010, the CineGrid@Disney demonstration was held at the Frank G. Wells Theatre at Disney Studios in Burbank, CA, for an internal audience of 100. Remote collaboration for movie making and movie selling is a "grand challenge" from CineGrid's perspective. Disney's Digital Studio team has defined, and continues to refine, the Disney user's requirements: a remote collaboration working environment that presents picture and sound at the highest calibrated quality, that enables easy creative interaction among colleagues, and that fosters the real-time intimacy that comes from seeing and hearing the same things at the same time. Historically, this has only been possible by bringing everybody involved into the same purpose-built room.

CineGrid@Disney was by far the most complex and ambitious CineGrid project to date. It represented the first prototype of a next-generation, multi-function, networked, remote collaboration facility for globally distributed cinema post-production and marketing. It brought together several different creative workflows, linking multiple remote locations, into a single room using very-high-quality media running over high-speed networks for interactive, real-time, remote collaboration.

Specific use cases demonstrated included: 4K/60p telepresence virtual conference room; critical viewing of digitally restored archival film elements at 4K and 2K resolutions streaming from a remote server; Digital Intermediate (DI) color grading; critical viewing of 3D high-definition stereoscopic visual effects; collaborative audio editing and mixing; and, use of EVL's SAGE software and multi-panel display walls for collaborative review of multimedia marketing materials. While most use cases involved point-to-point bilateral collaboration, several required triangular collaboration configurations using multicasting. Below is the network configuration for the SAGE collaboration workflow demo. Note that EVL provided network engineering support at StarLight for all the abovementioned demos.

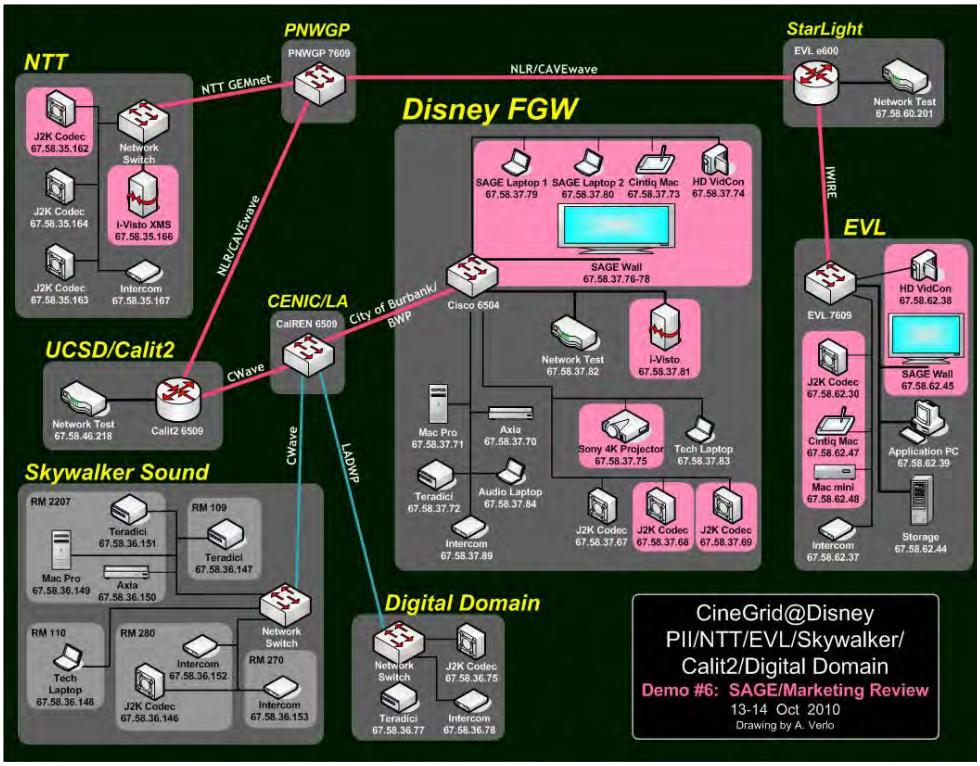

SAGE enabled remote collaboration between Disney in Burbank and EVL in Chicago. Calit2 shipped a tiled display wall to Burbank for the event. This wall connected over a high-speed network to EVL's Cyber-Commons. Participants shared and discussed marketing materials for the new "TRON Legacy" movie.

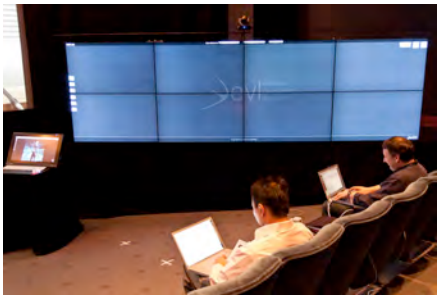

Some of those participating in the event were (in photo below, left to right: Leon Silverman, General Manager of Digital Studio Operations, Walt Disney Studios; Mark Harrah, Executive Director, Post Production, Walt Disney Studios; Jeff

Schectman, Senior Vice President of Creative Film Services, Walt Disney Studios; Jason Leigh, Director, Electronic Visualization Laboratory, University of Illinois at Chicago; Andy Fowler, Vice President and Senior Visual Effects Producer, Walt Disney Studios; Sara Duran-Singer, Senior Vice President of Worldwide Post Production, Walt Disney Studios; Laurin Herr, President, Pacific Interface Inc.; Phil Benson; Executive, Skywalker Sound, a Lucasfilm Company. (Photo courtesy of Pacific Interface, Inc.)

In addition to CineGrid member institutions that participated, several CineGrid networking organizations also contributed: CENIC, JGN2, GEMnet, PNWGP, and StarLight. (NTT used its GEMNET transoceanic network to stream compressed 4K and 2K using NTT's JPEG 2000 codecs; to stream uncompressed HD video via NTT's iVISTO server/gateway; and, to stream 3-way HD teleconferencing, including echo cancellation, via NTT's RealTalk.)

The CineGrid@Disney demo received the CENIC (Corporation for Education Network Initiatives in California) 2011 Innovations in Networking Award, in the Experimental/Developmental Applications category. The award was presented at the 2011 CENIC Annual Conference, held March 7-9, 2011 at UC Irvine.

CineGrid: DELF 2010

http://www.cinegrid.org/index.php?option=com_content&view=article&id=134:cyberport-becomes-first-cinegrid-member-in-greater-china-region&catid=1:news&Itemid=11
http://www.cinegrid.org/index.php?option=com_content&view=article&id=130:cyberport-and-cinegrid-present-hong-kongs-first-4k-live-streaming&catid=1:news&Itemid=11
http://www.cinegrid.org/index.php?option=com_content&view=article&id=127:laurin-herr-delivers-keynote-at-delf-210&catid=1:news&Itemid=11

Collaborators:

- CineGrid international consortium
- Cyberport; Hong Kong

On March 23, 2010, CineGrid participated in the Digital Entertainment Leadership Forum (DELF 2010) in Hong Kong, organized by the Hong Kong Cyberport Management Company Limited (Cyberport). Cyberport subsequently became the first CineGrid member in the Greater China Region.

On July 19, 2010, Cyberport did its first CineGrid 4K live streaming demo at Cyberport's Speaker Series in Hong Kong. The event, themed "CineGrid: Streaming 4K Motion Pictures from Around the World," attracted over 130 industry professionals.

2.B.1.2. GreenLight International

TransLight/StarLight leadership is quantifying the energy costs of data computation and movement for large-scale scientific applications using selected experimental networks.

GreenLight International Experiments 2010-2011

Switch Throughput and Power Measurements

http://greenlight.calit2.net/
http://seelab.ucsd.edu/

Collaborators:

- Project GreenLight, Calit2/UCSD; US

Project GreenLight measured the maximum throughput and energy cost of moving large amount of data, both within the UCSD GreenLight datacenter container and across campus (container and Calit2). The team tested both 1Gbps and 10Gbps networks, and found that power consumption per Ethernet switch is constant, regardless of how much traffic goes through it, while the maximum throughput ranges from 60-80% of rated capacity depending on the switch type. The team also measured the energy cost and throughput through optical switches, and found that the power consumption also does not vary with throughput, but throughput capabilities are much larger than what is possible in Ethernet switches.

The capability of Network Interface Card (NIC) on the server limited the maximum throughput that could be measured in the case of the optical switch. The disadvantage of optical switches is that setting up a circuit takes approximately 15ms, which is unacceptable for many transient flows, but may be leveraged by long-term, high-bandwidth flows.

In summary, networking equipment within a campus and within a campus datacenter is not energy proportional today. However, our project with Cisco indicates that it will be possible to scale down line-card power with throughput going forward. Thus, adaptive energy measurement and traffic scheduling/routing will become key for future large-scale systems.

For the complete report of this study, see Section 7.

GreenLight / GreenStar Collaboration

www.calit2.net/newsroom/release.php?id=1748
http://greenlight.calit2.net/
http://www.greenstarnetwork.com

Collaborators:

- Project GreenLight, Calit2/UCSD; US
- Project GreenStar Network; Canada

Starting October 2010, Project GreenLight was expanded internationally to provide DC powered compute nodes to Canada's Project GreenStar Network. This will support a "follow the sun/wind" processing paradigm in efforts to minimize ICT equipment Greenhouse gas (GHG) emissions by the dynamic [re]allocation of resources and [re]location of processing. In short, the network is used to communicate renewable-energy-powered ICT resource availability on an international scale and support the dynamic migration of computational tasks to minimize ICT related GHG emissions. Collaborative activities include the addition

of GreenStar nodes, the instrumentation of intermediate networking devices, the development of a power-managed hierarchical network infrastructure and the sharing of real-time energy consumption data.

The rapid growth in data-intensive scientific research has fueled an explosion in computing facilities and demand for electricity to power them. Energy usage per compute server rack is growing from approximately 2 Kilowatts (KW) per rack in 2000 to an estimated 30 KW per rack in 2010. Every dollar spent on power for IT equipment requires that another dollar be spent on cooling – equivalent to doubling the cost of the hardware itself over three years. As a result, cooling and power issues, to include "cleanliness" of power, are now major factors in system design. If we are going to continue to benefit from advances in computing, we need to better understand power and cooling. Scientists from all domains will choose more efficient systems as they invest in new cyberinfrastructure. GreenLight and GreenStar will hopefully give them the data they need to make environmentally sound decisions. Some scientific computing jobs need more powerful processors, some can do with less memory, some can use specialized processors – important requirements to understand so the optimally configured cluster can be chosen and scheduled through virtualization techniques each time they are needed. Availability of renewable-energy-powered ICT resources should also be taken into consideration.

GreenLight Minority-Serving Institutions Cyberinfrastructure Coalition (MSI-CIEC) Workshop

http://greenlight.calit2.net/outreach/

Collaborators:

- GreenLight Minority-Serving Institutions Cyberinfrastructure Coalition (MSI-CIEC); UIC/EVL; Calit2; US

June 9-10, 2010, Calit2 and SDSC organized the GreenLight Minority-Serving Institutions Cyberinfrastructure Coalition (MSI-CIEC) Workshop at UCSD. It focused on the Calit2 NSF-funded GreenLight project on the first day, and on OptIPortals the second day. Luc Renambot and Maxine Brown of EVL gave a remote SAGE presentation and demo to attendees in Calit2 Virtulab.

2.B.1.3. High-Performance Digital Media Network (HPDMnet)

This diagram illustrates HPDMnet topologies used for 2010-2011 demonstrations and experiments.

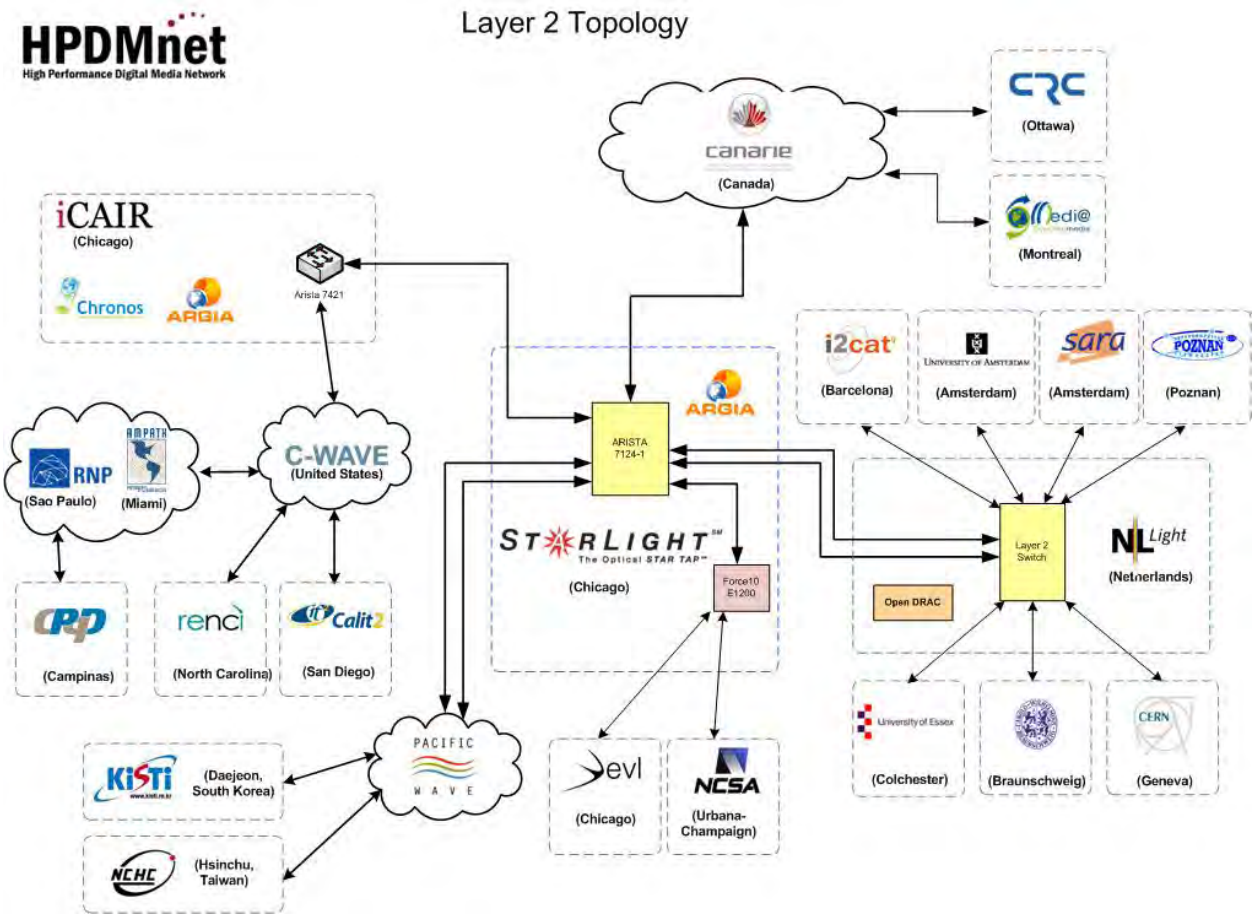

HPDMnet Experiments 2010-2011

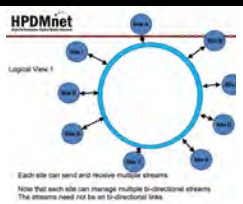

HPDMnet: HPDMnet @ GLIF 2010 and SC 2010

www.hpdmnet.org

Collaborators:

- CANARIE; Ciena; Communications Research Centre (CRC); Inocybe Technologies Inc; Synchromedia; University of Québec; Canada
- International Center for Advanced Internet Research (iCAIR), Northwestern University; UIC/EVL; Calit2/UCSD; National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign; StarLight Consortium; Renaissance Computing Institute (RENCI); North Carolina State University; US
- I2CAT; Barcelona
- Institute of Computer and Network Engineering, Technische Universitä Carolo-Wilhelmina zu Braunschweig; Braunschweig University of Art; Germany
- Korea Institute of Science and Technology Information (KISTI); Korea
- National Center for High-Performance Computing (NCHC); Taiwan
- SARA; University Van Amsterdam; NetherLight; SURFnet; The Netherlands

- Poznan Supercomputing and Networking Center; Poland
- Photonic Network Laboratory of High Performance Networking Group, University of Essex; JANET; UKLight; UK
- CPqD, National Education and Research Network (RNP), Brazil
- Global Lambda Integrated Facility (GLIF) International

Over the last six months, HPDMnet participants identified key objectives, developed a strategic plan, and organized several working groups. The primary objectives are to a) provide a persistent international network infrastructure for experiments and demonstrations, b) enable ongoing implementation of persistent dynamic connection services in an operational infrastructure, c) provide for the implementation of OpenFlow-enabled nodes in the HPDMnet testbed and explore the potential integration with similar infrastructures (e.g. iGENI-US, Ofelia-EU, etc.), d) develop and validate new innovations related to media streaming technologies in the HPDMnet infrastructure, including multicast L1/L2 streaming, services for various modalities, including HD, HD-3D, 4K, 4K-3D and 8K streaming as well as interactive HD, HD-3D and 4K-3D services, e) design and develop content-aware networks supported by dynamic and programmable networks, and f) provide for the technological validation of ultra-high-resolution video-conferencing facilities among multiple sites in the HPDMnet infrastructure.

The following chart describes the interactions among these working groups.

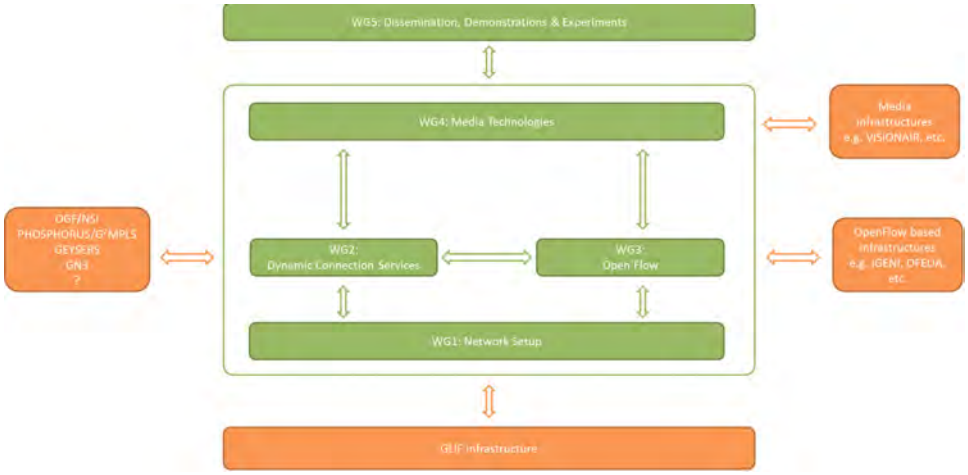

October 12-14, 2010, at the GLIF Workshop at CERN, HPDMnet Consortium members showcased dynamically provisioned inter-domain international services for high-performance digital media and other data-intensive applications. More specifically, in one demonstration, collaborators demonstrated the integration of Fenius and Argia.

In another demonstration, the Poznan Supercomputing and Networking Center in Poland used HPDMnet to showcase live streaming of 4K digital media from Poznan to CERN.

In another demonstration, iCAIR and CERN's networking and visualization groups showcased the "CERN Fireworks Demo" which displayed ultra-high-resolution visualizations of LHC experiment particle collisions (shown below).

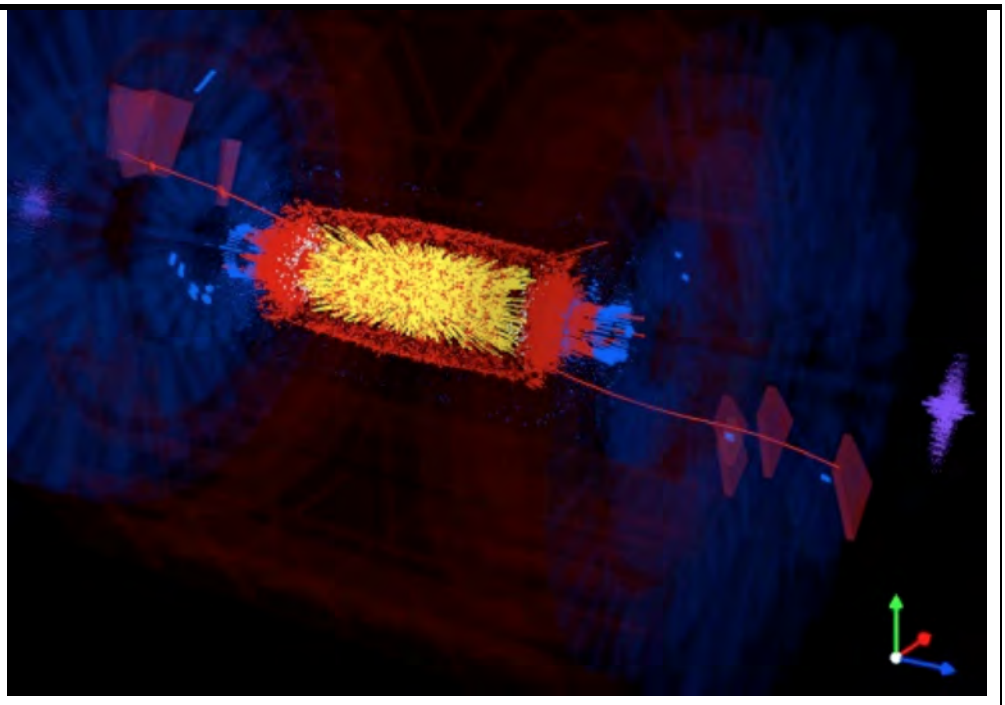

November 14-19 2010, at the SC10 conference in New Orleans, the HPDMnet consortium demonstrated several international digital media communication services based on an advanced architecture, network middleware, and lightpath-based channels on optical fiber. These demonstrations showed multiple high-resolution digital media streams transported between North America, Europe, and Asia over lightpaths, which could be scheduled or configured on demand. These demonstrations used the HPDMnet testbed, C-Wave, SCInet, the NLR and GLIF. The SC10 demonstrations showcased diverse content from Chicago, Ottawa, Taiwan, Daejeon, Poznan (Poland), and South Korea.

2.B.1.4. iGENI (International GENI)

Northwestern University (Joe Mambretti, principal investigator), Calit2 (Tom DeFanti) and EVL (Maxine Brown and Alan Verlo) are principals of "iGENI," the international Global Environment for Network Innovations (GENI), which is an NSF-funded project and part of the overall NSF GENI program.

The iGENI project optimizes private networks of intelligent devices capable of dynamically provisioned, low-latency, high-performance communications among multiple physically distributed infrastructures and federated domains.

iGENI Experiments 2010-2011

GENI Engineering Conference (GEC8) (with SAGE)

http://groups.geni.net/geni/wiki/Gec8Agenda

On July 20-22, 2010, Joe Mambretti participated in GEC8, held at Calit2/ UCSD. Mambretti gave three presentations: (a) "Partnerships for Advanced Global Networking: GENI, iGENI, and the International Research Community;" (b) to the Quilt workshop on iGENI from a regional network perspective (MREN): "iGENI Project Overview;" and, (c) to the Cluster-D group on incorporating ORCA and Resource Advertisements into iGENI.

Mambretti's team also did demos. One aspect of the demo focused on dynamic network provisioning integrated with cloud computing to support transcoding for digital media. Using SAGE, digital media content was streamed from Chicago to Calit2's SAGE wall.

This iGENI demo is a prototype of a demo being prepared with partners from HP Labs and Ericsson for a GEC9 VIP event in Washington DC in November 2010. The full demo will involve: (a) digital media transcoding, (b) for multiple edge platforms (e.g., smart phones, tablets, computers, tiled displays, standard displays), (c) supported by scalable cloud clusters, (d) based on a fabric of dynamically provisioned network L1/L2 paths. Due to logistic (and other) reasons, transcoding for GEC8 was limited to small edge platforms and not high-definition streams; however, L1/L2 network provisioning to support high-end digital streams was shown at GEC8 to convey the overall concept. Transcoding for HD will be shown at GEC9.

SAGE will be the client used for the full demonstration at GEC9, when the transcoding capability will be extended to a wider range of platforms. The goal is to create a cloud-based platform that can transcode to random devices (versus having separate services for each device).

Northwestern University established Eucalyptus clouds on clusters at HP Labs, Northwestern, and UCSD's CS department, interconnected with vLANs. These three interconnected clouds have been implemented as scalable compute resources to support transcoding.

GENI Engineering Conference (GEC9)

http://groups.geni.net/geni/wiki/Gec9Agenda

On November 2-4, 2010, Joe Mambretti participated in the GEC9 in Washington DC. His group staged several complex iGENI demonstrations and gave several presentations.

In addition to being part of a special showcase at this GEC venue, Mambretti showcased the iGENI projects that his team has been developing over the last year. The majority of these focused on interconnections among multiple US sites; however, the team has also successfully established two international iGENI testbeds: (1) one among iCAIR/StarLight, NCHC in Taiwan and several universities in Taiwan, and (2) one connecting several universities in Korea, StarLight, iCAIR, RENCI and Duke University. He is also working to establish a third with Brazil.

GENI Engineering Conference (GEC10)

http://groups.geni.net/geni/wiki/Gec10Agenda

On March 15-17, 2011, Joe Mambretti participated in the GEC10 in Puerto Rico. His team demonstrated a prototype platform using the ORCA control architecture to discover resources in iGENI; i.e., it can set up and manage network services and control/manage individual traffic streams.

Mambretti and HP did a demo and poster session on TransCloud, a transcontinental federation of cloud systems, which is a joint project between iGENI and GENICloud. GENICloud provides seamless inter-operability of cloud resources across n-sites, n-administrative domains.

TransCloud uses a Slice-based Federation Architecture (SFA) for sign-on and trans-cluster slice management. A SFA cluster manager resides at each site, currently based on Eucalyptus. A 10Gbps transcontinental network links several sites – HP OpenCirrus, Northwestern OpenCloud, UCSD, Universitat Kaiserslautern, University of Victoria – thanks to PlanetLab, GLIF, NLR, NetherLight, CAVEwave, StarLight and DFN. Rick McGeer of HP said the attendees found it "compelling," noting it took 19 people at eight different organizations working on this project to create the underlying system. More sites will be added in the future.

The *TransCloud Transcoding* demo showed how TransCloud can provide a capability to use one high-performance distributed environment for transcoding to multiple platforms, such as computers, tablets, tiled display walls, and mobile phones. Traditionally, digital media has used different infrastructure for different edge delivery platforms.

2.B.1.5. SAGE

This map, which appears on the SAGE website, shows that there are 65 SAGE users worldwide <http://www.sagecommons.org/index.php?option=com_content&view=article&id=50>.

SAGE Experiments 2010-2011

SAGE: IEEE e-science 2010

www.escience2010.org/

Collaborators:

- Calit2; UIC/EVL; US
- University of Queensland; AARNet; Australia

December 7-10, 2010, the IEEE e-science 2010 conference was held at Queensland University of Technology (QUT) in Brisbane, Australia. University of Queensland did SAGE demos on OptIPortals at QUT that were Visualcasted to AARNet in Sydney. They ran Paraview in MPI-mode across AARNet's OptIPortal cluster.

University of Queensland also conducted *Tiled Display Wall* workshops, designed for both new users as well as researchers already operating tiled display walls. The current and future applications of display wall technology were discussed, including the role of SAGEBridge (for Visualcasting), and opportunities for participants to collaborate and contribute to a growing community. Jurgen Schultze, Calit2/UCSD, spoke at the workshop and was also one of the conference's keynote speakers.

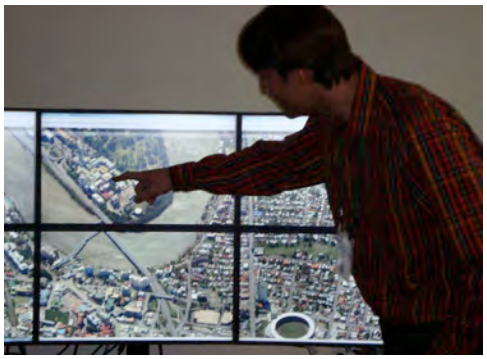

In the photo on the left, Schultze indicates where he was on 7 Dec 2010 – at QUT in Graden's Point on the Brisbane river. The image was taken from a UAV and shown on a 3x2 OptIPortal. The photo on the right shows Professor Bernard Pailthorpe, at the conference site, being Visualcasted.

Note: EVL provides technical support to University of Queensland SAGE users.

Note: The AARNet November 2010 newsletter reported that the use of OptIPortal visualization walls among Australian institutions has grown to a level where interaction beyond the usual point-to-point has become feasible, and has opened up opportunities for multi-site collaborations, which drive demand for stable middleware to manage such wall-to-multiwall interactions. EVL's SAGE is gaining broad acceptance in serving this purpose. To further the use of SAGE for collaboration opportunities among the Australian research community, the University of Queensland and AARNet deployed a SAGEBridge trial service with 10Gbps backbone connectivity. SAGEBridge enables interactions among multiple walls by replicating visual streams from one wall to other participating sites with a variety of bandwidth capabilities – a method dubbed "Visualcasting" or just "VisCasting". Various VisCasting sessions were conducted among SAGE-enabled OptIPortals, particularly among AARNet (Sydney), University of Queensland and

NICT in Japan, generating up to 7Gbps across the network, and over 10Gbps in the lab! See <http://www.aarnet.edu.au/library/AARNews_Issue19.pdf>.

SAGE: 10th Annual ON*VECTOR International Photonics Workshop

*Collaborators: ON*VECTOR (Optical Networked Virtual Environments for Collaborative Trans-Oceanic Research) is a joint project of:*

- NTT Network Innovation Laboratories; Keio University's Institute for Digital Media and Content (DMC); the University of Tokyo's Morikawa Laboratory; Japan
- UIC/EVL; Calit2/UCSD; Pacific Interface Inc (PII); US

This year's meeting took place March 1-4, 2011 at Calit2/UCSD. EVL's Luc Renambot gave the presentation "SAGE/iVisto Integration Experiment: SAGE Application Perspective" and worked with NTT on a real-time demonstration. iVisto is an NTT product for video streaming.

Note: This demo was also done for the NTT-sponsored Digital Media Analysis, Search and Management (DMASM 2011) Workshop held earlier in the week.

SAGE: SARA 40Gbps Streaming Demo at GLIF 2010

https://noc.sara.nl/nrg/
http://www.isgtw.org/feature/feature-ultra-fast-networks-final-frontier

Collaborators:

- UIC/EVL; US
- SARA; The Netherlands
- CERN; Switzerland

October 13-14, 2010. For the GLIF 2010 Workshop, SARA's Network Research Group streamed ultra-high-resolution scientific visualization animations over a SURFnet 40Gbps lightpath between SARA and CERN.

EVL put considerable time into working with SARA colleagues to trouble shoot issues encountered with SAGE; however, these bottlenecks could not be resolved and ultimately SAGE was not used. This demo is included here because it provided a rich proving ground to analyze SAGE capabilities, as explained below.

Tijs de Kler of the SARA Visualization Group wrote a thoughtful analysis of why SAGE failed. Essentially, the research problem they wanted to solve was doing 40Gbps from a single machine and disk; i.e., there was significant I/O from disk to memory to network, all from one machine.

The way SAGE streams data is to first read it into a buffer, and then hand it off to networking threads. Here, the data is divided into blocks, and streamed to receiving nodes with some header attached to identify what block it is. Streaming is done dynamically, with SAGE determining what blocks go to what receiving nodes and, if necessary, to multiple nodes if the image falls across a border. In SARA's case, SAGE streaming didn't scale well. It couldn't get past the 16Gbps mark, and was mainly running out of CPU. However, *netperf* showed that higher speeds (up to 39.6Gbps) could be reached with a fraction of the CPU used.

Additional benchmarking with a simple application, doing nothing more that reading data from a file and sending it over a UDP socket, showed similar bottlenecks. But, if either disk reading or network streaming was taken out, it did fine. The conclusion was that moving data from disk to user space memory, and back to the network was the problem, requiring too much of the CPU. SARA then came up with a solution that would do it more efficiently. While their solution is unwieldy and not general purpose (like SAGE), it works for this demo.

Note: SARA repeated the demonstration at SC10, but streamed locally within the Dutch Research booth.

SAGE: QUESTnet 2010 (Queensland Education, Science and Technology Network)

https://www.questnet.edu.au/display/qnc2010/Home

Collaborators:

- University of Queensland; AARNet; Australia

QUESTnet (Queensland Education, Science and Technology Network) is the Queensland regional component of the Australian Academic and Research Network

AARNet. The QUESTnet 2010 conference was held July 6-9, 2010 at Griffith University, Australia. Christopher Willing of University of Queensland and Bernard Meade of University of Melbourne gave the presentation "OptIPortal Collaborations: Requirements and Progress Examples" and did SAGE demonstrations to explain collaboration. The demos showed how two OptIPortals (AARNet divided theirs into two 5x3 and 3x2 displays) could run the same application and how changes to one could be simultaneously displayed on both; i.e., not just a broadcast from one controlling site, but rather controlled from either site.

SAGE: Knowledge Capital of Japan

http://www.kmo-jp.com/en/trial/2010/salon.html

Collaborators:

- Calit2; UIC/EVL; US
- NICT; Osaka University; Japan

On June 9-10, 2010, Japan's National Institute of Information and Communications Technology (NICT) organized a major event in Osaka to promote it as the Knowledge Capital of Japan. One demonstration featured SAGE Visualcasting. Participants included NICT in Tokyo, Osaka University's VisLab, NCHC and NTUT in Taiwan, University of Queensland in Australia, and Calit2 in the US. EVL remotely operated the SAGE Wall at Calit2.

SAGE: World Congress on Information Technology (WCIT 2010)

http://www.wcit2010.org/

Collaborators:

- UIC/EVL; US
- SARA; The Netherlands

The WCIT 2010 was held in the RAI in Amsterdam, May 25-27, 2010. The theme of WCIT 2010 was Challenges of Change; nine tracks focused on the role of ICT with regard to social issues, such as climate change, energy saving, and water management. EVL collaborator SARA participated in the Track Energy. In the Energy pavilion, SARA gave demonstrations on a tiled display panel.

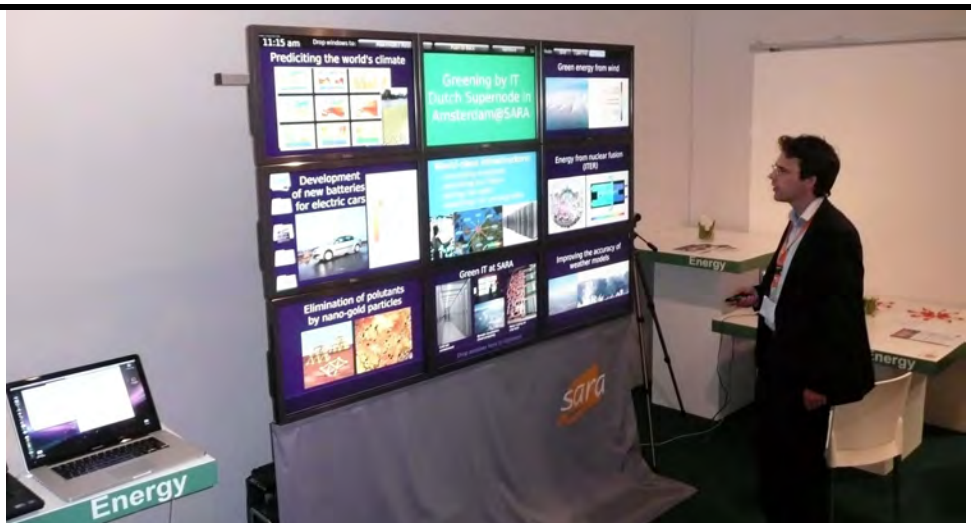

SARA used the latest SAGE user interface modules, which had not yet been thoroughly tested on different display systems. UIC/EVL provided a lot of support for this demonstration.

SAGE: PRAGMA 18

Collaborators:

- Calit2; UIC/EVL; US
- National Institute of Adv. Industrial Science and Technology (AIST); Japan

The PRAGMA 18 workshop was held at Calit2/UCSD, March 3-4, 2010. Calit2 and EVL organized several demonstrations to showcase networked visualization, including a SAGE demo where imagery was streamed from a storage server located at StarLight to the SAGE tiled display wall at Calit2. In particular, EVL streamed an animation of Mt. Hood, Oregon, created by PRAGMA member and SAGE user AIST of Japan. EVL and Calit2 were also connected via Polycom (output directed to the tiled display wall) so the two groups could talk to one another.

2.B.1.6. Science Cloud Communication Services Network (SCCSnet)

SCCSnet Experiments 2010-2011

Science Cloud Communication Services Network

Collaborators

- International Center for Advanced Internet Research (iCAIR), Northwestern University; Open Cloud Consortium; Open Science Data Cloud Consortium; TransCloud Consortium; StarLight Consortium; Metropolitan Research and Education Network (MREN); US
- iGENI Consortium; Global Lambda Integrated Facility (GLIF); (international)
- DFN; Germany

Current cloud architecture, technology and infrastructure are oriented to consumer and enterprise cloud applications on best-effort networks; they cannot support large-scale data-intensive science applications. SCCSnet is investigating, designing, and implementing prototype services, architecture, and technology to integrate dynamic networking and dynamic cloud provisioning for large-scale data-intensive science, as well as developing new communication services. This initiative leverages several major cloud research project partnerships, including the Open Cloud Consortium, the Open Science Data Cloud Consortium, and the TransCloud Consortium.

One initiative is the design of communication services based on advanced techniques for virtualized overlay networks. This design will be implemented on a national/international testbed and demonstrated at SC11 in Seattle in November. As part of this effort, several experimental control-plane techniques are being investigated, including those being developed on international network testbeds that extend from StarLight to Germany, Korea and Taiwan.

During the GENI Engineering Conference (GEC) 10 conference in San Juan, Puerto Rico, March 15-17, 2011, the TransCloud Consortium showed a prototype of its cloud networking techniques during the plenary session (see Section 2.B.1.4: iGENI). In particular, they demoed the integration of programmable, dynamic networking with programmable clouds. Multiple clouds (established at distributed sites: TransCloud – HP OpenCirrus, UCSD, Northwestern, and Kaiserslautern University in Germany) were interconnected with dynamic network provisioning to create new capabilities and services. This demo showed that separate infrastructures do not have to be implemented for different types of edge-delivery platforms. The single TransCloud environment provides streams to any other edge device. Whereas legacy approaches require different infrastructure for each edge platform, this transcoding application illustrates the potential to create a scalable, distributed computing environment comprised of federated dynamically programmable clouds integrated with dynamic network provisioning.

Reference: Robert L. Grossman, Yunhong Gu, Joe Mambretti, Michal Sabala, Alex Szalay, Kevin White, “An Overview of the Open Science Data Cloud,” Proceedings, Science Cloud Workshop, June 21, 2010, Chicago Illinois.

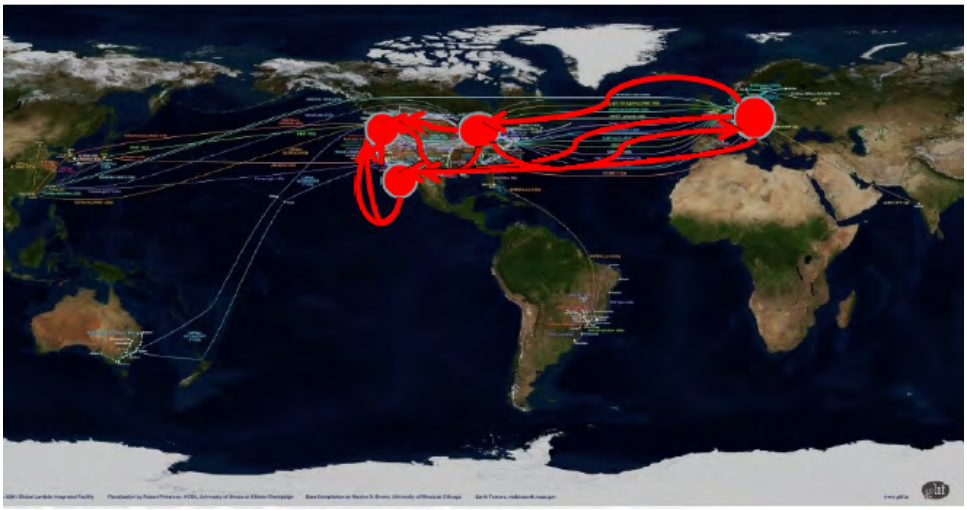

2.B.1.7. Additional Experiments Aided by TransLight/StarLight

Additional Experiments 2010-2011

HDTV and Medical Joint Session (Live Demo of 3D Robot Surgery from the Czech Republic) @ APAN 2011

http://www.apan.net/meetings/HongKong2011/Session/Medical.php
http://www.apan.net/meetings/HongKong2011/

Collaborators:

- Kyushu University; Japan
- GIST; Korea
- Masaryk Hospital in Ústí nad Labem (MHUL), CESNET; Czech Republic

This live demo, which took place at the APAN 2011 meeting in Hong Kong, February 21-25, 2011, was for the high-level education for surgeons. MHUL in the Czech Republic connected to the APAN venue in Hong Kong, where Dr. Jan Schraml at MHUL performed a live surgery of a patient with carcinoma of the prostate to demonstrate the advantages of mini-invasive robotic-assisted radical prostatectomy in preserving neurovascular bundles and bladder neck.

The networks involved, illustrated above, were CESNET (Prague to Amsterdam), TransLight/ StarLight (Amsterdam to Chicago), CANARIE (Chicago to Seattle), KREONet2 (Seattle to Korea), and then to Hong Kong. StarLight network engineer Alan Verlo helped configure this network. Below is a TransLight/StarLight MRTG graph for February 23, 2011 (showing 4Gbps usage).

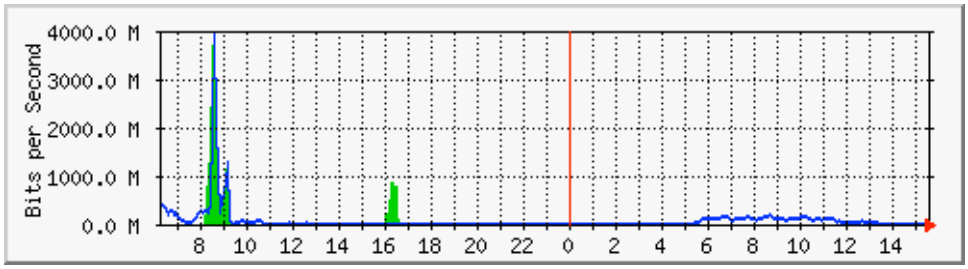

Below are statistics captured by the Czech and Korean NOCs:

TenGigE0/0/0/4 -- CR1.apricot-apan.asia

System: CR1.apricot-apan.asia in
Maintainer: 2011-infra@googlegroups.com
Description: TenGigE0/0/0/4 APAN: KRLight / CESNET 3D Video [to Czech]
ifType: ethernetCsmacd (6)
ifName: TenGigE0/0/0/4
Max Speed: 10.0 Gbits/s
Ip: 67.58.60.148 ()

The statistics were last updated **Wednesday, 23 February 2011 at 8:35**,
at which time **'CR1.apricot-apan.asia'** had been up for **2 days, 16:39:55**.

`Daily' Graph (5 Minute Average)

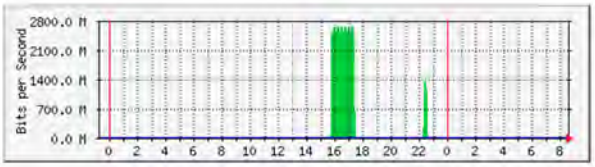

	Max	Average	Current
In	2659.2 Mb/s (26.6%)	128.4 Mb/s (1.3%)	256.0 b/s (0.0%)
Out	24.0 b/s (0.0%)	0.0 b/s (0.0%)	0.0 b/s (0.0%)

`Weekly' Graph (30 Minute Average)

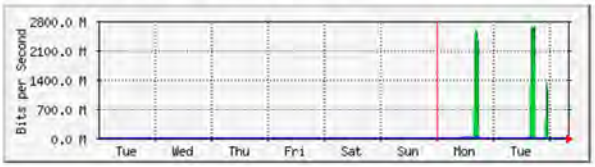

	Max	Average	Current
In	2663.8 Mb/s (26.6%)	115.8 Mb/s (1.2%)	256.0 b/s (0.0%)
Out	78.6 kb/s (0.0%)	240.0 b/s (0.0%)	0.0 b/s (0.0%)

GREEN ### Incoming Traffic in Bits per Second
BLUE ### Outgoing Traffic in Bits per Second
DARK GREEN ### Maximal 5 Minute Incoming Traffic
MAGENTA ### Maximal 5 Minute Outgoing Traffic

MRTG MULTI ROUTER TRAFFIC GRAPHER
2.14.5
Tobias Oetiker <tobi@oetiker.ch>
and Dave Rand <dlr@bungi.com>

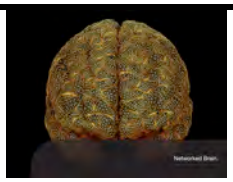

CBRAIN: Canadian Brain Imaging Research Network

http://www.cbrain.mcgill.ca

Collaborators:

- Montreal Neurological Institute, McGill University; Rotman Research Institute, Baycrest Hospital; Robarts Institute, University of Western Ontario; L'Unité de Neuroimagerie Fonctionelle de l'Université de Montréal; Neuroscience Department, University of British Columbia; CANARIE; Canada
- National Center for Microscopy and Imaging Research (NCMIR), UCSD; US

The CBRAIN Project is a pan-Canadian network of five leading brain imaging research centers in Canada, with an extension to UCSD/NCMIR (Mark Ellisman's lab). Specifically, McGill university has partnered with NCMIR to collaboratively view 3D visualizations of ultra-high-resolution brain scans on tiled display walls. Initially McGill installed a 4x3 (27Mpixel) tiled display wall with a 12-node visualization cluster, using about 7.5Gbps bandwidth. At StarLight, network engineers interconnected CANARIE and Cisco's C-Wave for this purpose.

McGill built a multi-resolution 3D viewer that allows them to remotely browse through massive models over the network. Their current high-resolution brain scan totals 1TB of data on disk. They designed the software to require bandwidth as a function of the number of pixels displayed on the visualization window, independent of maximum resolution and full model size. (They do not know of other existing software systems capable of browsing this amount of data interactively in real time.) The purpose of working with NCMIR is to perform a 2-point collaborative visualization (McGill-UCSD) of the highest resolution full brain scan in existence. The UCSD setup will use NCMIR's expertise in tiled display walls (walls with up to 70 full HD monitors), thus requiring large dedicated bandwidth to access the brain model (1TB+ total) located at McGill University.

Ciena and CANARIE Demonstrate 100Gbps Network Delivering Innovative and Powerful Tele-Medicine and Brain Research Application

http://www.fiercewireless.com/press-releases/canarie-and-ciena-demonstrate-100g-network-delivering-innovative-and-powerful-tele-me
http://www.lightwaveonline.com/networking/news/CANARIE-Ciena-demonstrate-100-Gbps-network-for-tele-medicine-and-brain-research-105175949.html

Collaborators:

- CANARIE; Ciena; Canada
- StarLight; NASA Goddard Space Flight Center; International Center for Advanced Internet Research (iCAIR), Northwestern University; US

Throughout October, 2010, Ciena, a networking equipment manufacturer, and CANARIE, Canada's Advanced Research and Innovation Network, conducted successful 100Gbps demonstrations over CANARIE's network to highlight their ability to deliver applications across a wide range of industrial, educational and research sectors. The demonstrations showcased two CANARIE-funded research projects – the Health Service Virtual Organization (HSVO) and the Canadian Brain Imaging Research Network (CBRAIN) – that use networks to accelerate medical discoveries and innovations in brain research, and help eliminate the challenges of geography in delivering medical education to remote and rural areas. The demonstrations were part of an advanced technology event for Ciena customers.

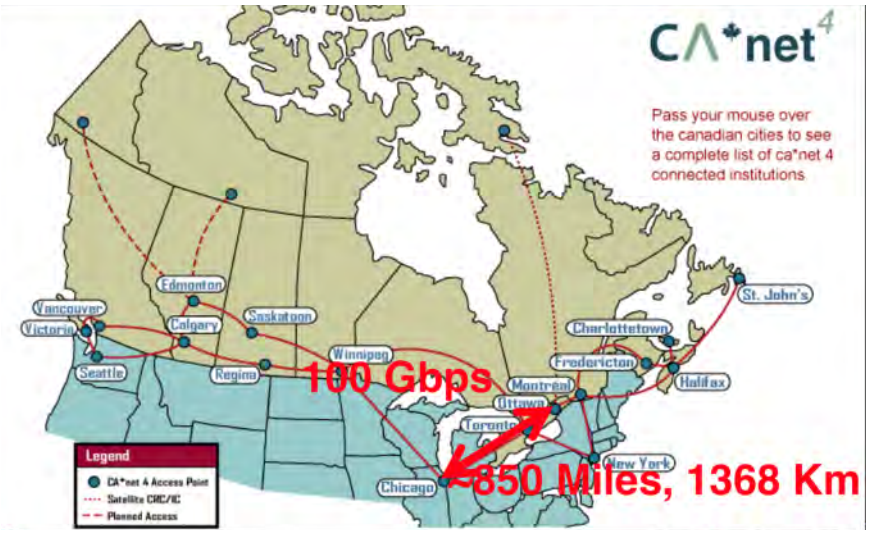

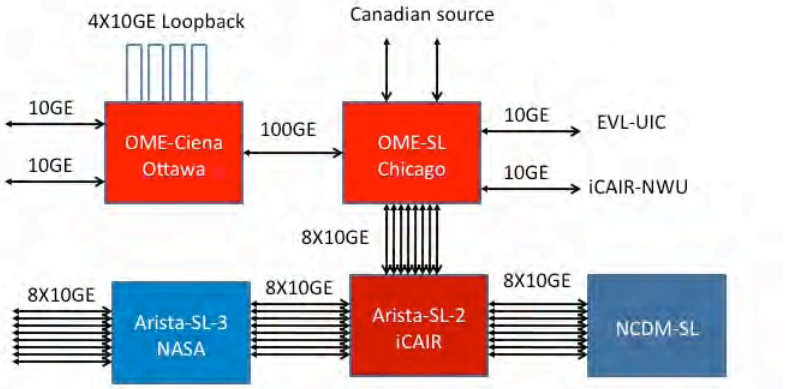

CANARIE's network was easily upgradeable to 100Gbps by simply adding 100Gbps interfaces to its installed Ciena ActivFlex 6500 equipment without changing any network infrastructure or completely re-engineering the network. Using the ActivFlex 6500 and ActivSpan Common Photonic Layer, the demos originated from Chicago's StarLight facility and included uncompressed HD video data streams from the NASA Goddard Space Flight Center, the International Center for Advanced Internet Research (iCAIR) at Northwestern University, and other sources, along with the CANARIE applications. Using Ciena's 40/100Gbps Adaptive Optical Engine technology – which includes coherent optics, electronic dispersion compensation, and directionless and colorless ROADM functionality – the 100Gbps wavelength traveled 1,300 km to Ciena's Ottawa labs for displaying the video content.

FENIUS @ GLIF 2010 and SC10

http://www.glif.is/publications/press/20101129.html
http://www.glif.is/meetings/2010/tech/vollbrecht-dynamicgole.pdf
http://code.google.com/p/fenius/
http://www.ultralight.org/~azher/idc/Glif2010/

Collaborators:

- *Participating GOLEs:* CERNLight (Switzerland), CzechLight (Czech Republic), JGNLight (Japan), MAN LAN (US), NetherLight (the Netherlands), NorthernLight (Sweden), PSNCLight (Poland) and StarLight (US)
- *Other participating networks and institutions:* AIST (Japan), CESNET (Czech Republic), Internet2 (US), KDDI (Japan), USLHCnet (US), University of Amsterdam (Netherlands), and University of Essex (UK)

TransLight/StarLight previously reported on this in its OCI-0441094 annual report for the period February 1, 2010 – January 31, 2011. It is mentioned again here, as subsequent activities reference this GLIF initiative.

GLIF successfully demonstrated a pilot implementation of an automated lightpath system *(the Automated GOLE Pilot Project)* during SC10 in New Orleans on 15-18 November 2010. This expanded on the system that was first demonstrated at GLIF's 10th Annual Global LambdaGrid Workshop in Geneva on 13 October 2010.

While GLIF has been instrumental in facilitating high-performance applications worldwide, it has been necessary to manually set up lightpaths at each intervening GLIF Open Lightpath Exchange (GOLE). This not only requires significant administrative effort, but it can take several hours to configure each lightpath. There are significant advantages to being able to automatically establish lightpaths.

One obstacle to automating this process is that GOLEs utilize a variety of equipment with different control mechanisms (e.g., Argia/UCLP, DCN/OSCARS, DRAC, G-Lambda), not to mention they fall under different administrative domains. To facilitate dynamic set up, it is necessary to have a standard way of requesting and configuring lightpaths across all GOLEs. It is also necessary to advertise resource availability at each GOLE so that a comprehensive topology can be built up for reachability purposes.

The Automated GOLE Pilot builds on Fenius software developed by ESnet, which is

a common API for setting up lightpaths, but translates requests to the different underlying control mechanisms at the GOLEs. This allows lightpaths to be established on demand or reserved in advance for specific periods, along with dedicated capacity and performance characteristics.

Eight GOLEs participated in the demonstrations, along with three additional sites, including the SC10 venue. Lightpaths were successively brought up and torn down between each location in conjunction with the perfSONAR PinGER service for monitoring and display, providing proof-of-concept and a platform for further development.

GLIF Automated GOLE Logical Topology

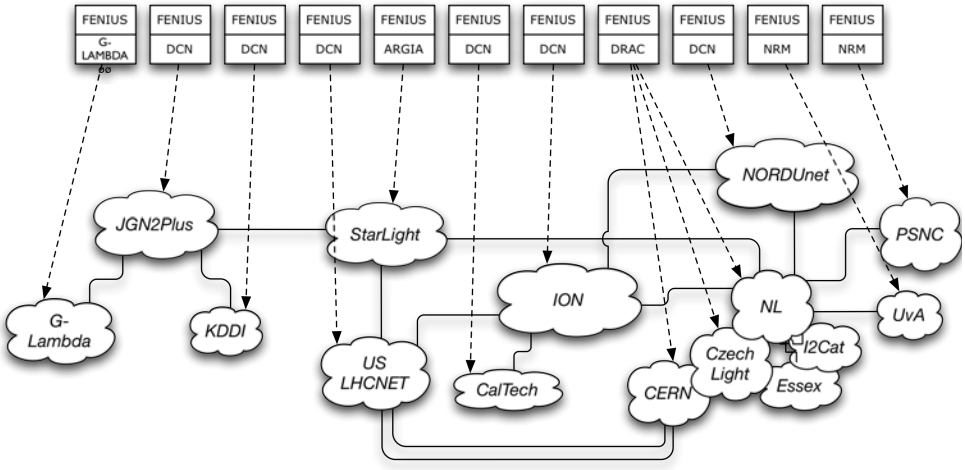

The demonstrations represented the initial phase of the pilot project, and there are now two further objectives. The first objective is to work with the Open Grid Forum (OGF) Network Services Interface Working Group to develop a standardized API that can be implemented by GLIF. Complementary to this is the further development of topology distribution and IP layer negotiation protocols, with a view to making these elements more robust. The second objective is to start involving demanding users and application developers who are interested in the capacity and performance offered by dynamic lightpaths, and who would be willing to work with GOLE operators by providing feedback and helping develop best practices.

Below are snapshots, taken by Caltech, for some of the USLHCNet-provided circuits used at GLIF 2010 that were dynamically created from ESnet Fenius (through USLHCNet IDC). *Note that StarLight is part of the overall configuration. When a request is made for a connection between, for example, StarLight and NetherLight, that connection is shown on the displays when it becomes active.*

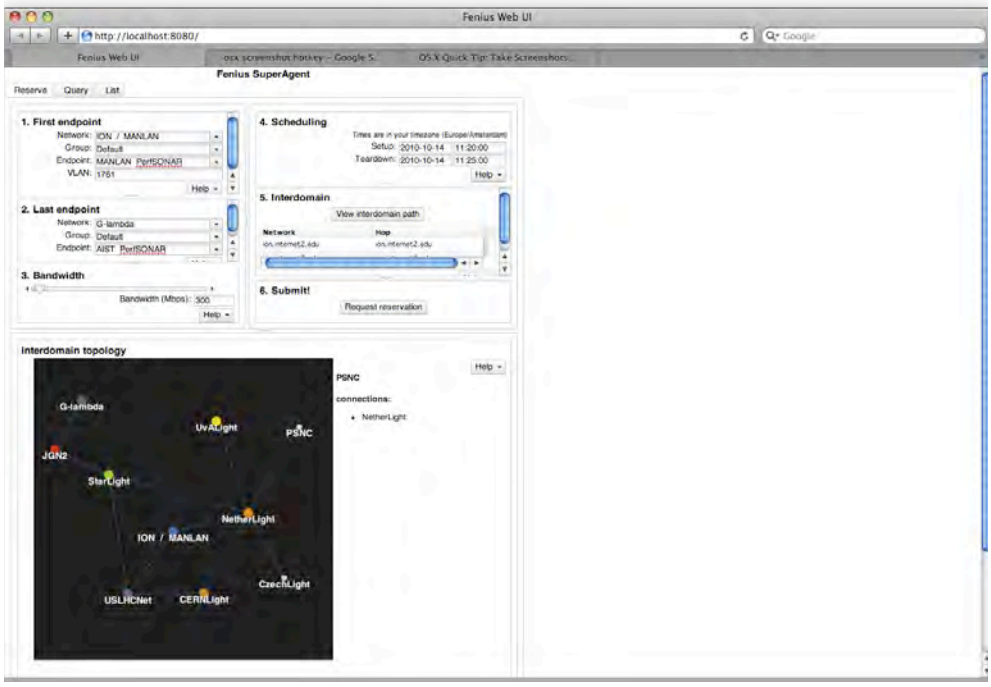

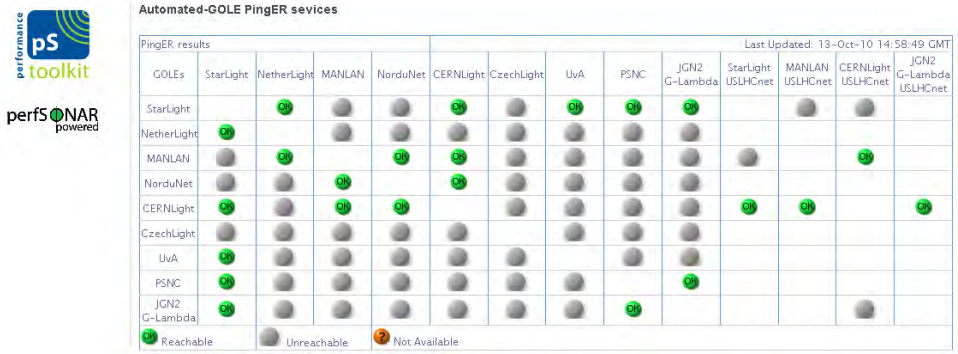

FENIUS @ APRICOT-APAN 2011

http://www.glif.is/meetings/2011/winter/

Collaborators:

- *Participating GOLEs:* CERNLight (Switzerland), CzechLight (Czech Republic), JGNLight (Japan), MAN LAN (US), NetherLight (the Netherlands), NorthernLight (Sweden), PSNCLight (Poland) and StarLight (US)
- *Other participating networks and institutions:* AIST (Japan), CESNET (Czech Republic), Internet2 (US), KDDI (Japan), USLHCnet (US), University of Amsterdam (Netherlands), and University of Essex (UK)

Global Lambda Integrated Facility (GLIF) Technical and Control Plane Working Group held a meeting at the APRICOT-APAN 2011 conference in Hong Kong, February 21-25, 2011. At that time, they successfully repeated the FENIUS Automated GOLE Pilot Project demonstration done at the GLIF 2010 LambdaGrid Workshop and at SC10.

Alan Verlo, StarLight network engineer, worked on this project.

2.B.2. Cooperative Community Partnerships

2.B.2.1. Workshop Organizing and Support

Tom DeFanti and Maxine Brown have been involved with the following organizations and conferences throughout the past year, whose goals are to find and encourage application and middleware development.

- **12th Annual LambdaGrid Workshop,** 2012 – GLIF has asked Maxine Brown and Joe Mambretti about the possibility of hosting the annual GLIF Workshop in Chicago in October 2012, and co-locating with OGF, IEEE e-Science, and Microsoft eScience Workshop. Maxine and Joe have spoken to OGF and IEEE e-Science organizers and are investigating whether this is feasible.
- **11th Annual LambdaGrid Workshop,** September 13-14, 2011, in Rio de Janeiro, hosted by RNP Brazil and sponsored by GLIF. Maxine Brown is a member of the Program Committee.
- **CineGrid@TIFF 2011 (Tokyo International Film Festival)**, October 2011 – hosting a one-day CineGrid workshop at TIFF 2011 is currently being discussed. Tom DeFanti is one of the organizers.
- **10th Annual ON*VECTOR International Photonics Workshop,** March 1-4, 2011, Calit2/UCSD. Tom DeFanti and Maxine Brown were members of the Program Committee. DeFanti, Brown and Mambretti participated.
- **10th Annual LambdaGrid Workshop,** October 13-14, 2010, at CERN, Geneva, Switzerland, sponsored by GLIF. Maxine Brown served as a member of the Program Committee. Tom DeFanti, Tajana Rosing and Joe Mambretti participated.
- **CineGrid@TIFF 2010 (Tokyo International Film Festival)**, October 28, 2010, was a one-day CineGrid workshop at TIFF. Tom DeFanti was one of the organizers, as well as a participant.
- **5th Annual CineGrid International Workshop 2010,** sponsored by CineGrid and hosted by Calit2 at UCSD, December 12-15, 2010. Tom DeFanti was a co-organizer and member of the Program Committee.

2.B.2.2. Meetings and Events

In addition to the above-mentioned events, TransLight/StarLight principals participated in the following meetings and conferences, promoting IRNC efforts.

June 24, 2011. Tom DeFanti participated in the NSF SBIR/STTR CISE Fundamental Research panel, held at NSF in Arlington, VA.

June 21, 2011. Donga-Ilbo, a major newspaper in Korea, wishes to write an article on the use of information technology (especially computing) in the media industry. A delegation will visit Calit2 (Tom DeFanti) and other CineGrid sites to learn firsthand about the latest technologies.

June 14-27, 2011. Daniel Acevedo and Steve Cutchin (KAUST Visualization Laboratory, Saudi Arabia) are spending two weeks visiting Tom DeFanti at Calit2/UCSD (a KAUST partner) to work on joint projects, start preparing research agendas for next year, and finalize plans for the KAUST booth at ACM SIGGRAPH in August.

June 6-7, 2011. Calit2 hosted an international CineGrid Exchange meeting. Participants included Laurin Herr and Natalie Van Osdol (Pacific Interface), Lisa Padilla (communicate.io), Fred Davis (CineGrid), Dana Plepys (UIC/EVL), Jeff Weekley (Naval Postgraduate School), Michal Krsek (CESNET, Czech Republic), and Nathan Brock, Qian Liu, Yuma Matsui, Joe Keefe and Tom DeFanti (Calit2).

May 28, 2011. Amit Apte of the Centre For Applicable Mathematics at the Tata Institute of Fundamental Research Centre is working with Greg Cole on BXLight (a Taj-GLORIAD GOLE in Bangalore, India). He was in the US on other business and visited Chicago to meet with Joe Mambretti and Maxine Brown about StarLight.

May 17, 2011. Alan Verlo and Joe Mambretti participated in a JET meeting.

May 16-29, 2011. Jason Leigh and Luc Renambot of EVL visited KAUST (King Abdullah University for Science and Technology) in Saudi Arabia. EVL receives funding from KAUST to develop and deploy some of its visualization technologies. KAUST most interested in SAGE as a tool for distributed, collaborative research over high-speed networks with funded groups in the US and elsewhere.

May 16, 2011. Seok-Myun Kwon and Jin Kim of KISTI visited EVL for the day to learn more about SAGE user interface development, as KISTI wants to standardize on SAGE and tiled display walls for distributed, collaborative research and education.

April 19, 2011. Alan Verlo and Joe Mambretti participated in a JET meeting.

April 5-May 3, 2011. Tom DeFanti visited KAUST (King Abdullah University for Science and Technology) in Saudi Arabia, as well as archaeological sites in Cairo and Luxor, Egypt. Calit2 receives KAUST funding to develop and deploy visualization technologies. In addition, DeFanti was instrumental in their acquiring a 10Gbps circuit from the campus to NetherLight for transatlantic bandwidth connectivity. DeFanti is working on capturing cultural heritage images, to be displayed in 3D on KAUST's virtual-reality systems, but these same very-high-resolution images can also be shared among KAUST, Calit2 and EVL over optical networks. During his stay, DeFanti and his team visited cultural heritage sites in Madain Saleh (Saudi Arabia) and Luxor (Egypt).

April 1, 2011. Tom DeFanti met with Dave Lambert, Internet2 CEO. (Lambert also met with Larry Smarr on April 4.) Subsequently, in May, Ben Fineman and Mike LaHaye of Internet2 – along with Erik Hofer of University of Michigan – contacted EVL. Internet2 is interested in learning more about using SAGE for multi-site collaboration. (Discussions have started and Internet2 plans to visit EVL June 16, 2011.)

March 31, 2011. Maxine Brown participated via VTC with Larry Smarr at Calit2/UCSD, who was visited by Chilean representatives: Harald Beyer (PhD in economics from UCLA, and director of the CEP think tank in Chile); Eugenio Guzmán (Dean of the School of Government of the Universidad del Desarrollo in Chile); Alvaro Fischer (President, Chairman of the Board, Fundacion Chile). Maxine discussed international networking, and provided Calit2 with a draft of the new GLIF map of South America to display on its OptIPortal.

March 15, 2011. Alan Verlo participated in a JET meeting.

March 1-4, 2011. 10th Annual ON*VECTOR International Photonics Workshop. Tom DeFanti and Maxine Brown were members of the Program Committee. DeFanti gave the presentation "GreenLight Project: Progress and Plans" and chaired a session on Designing and Securing Cloud Computing Systems. He also organized demos. Joe Mambretti gave the presentation "Future GOLEs: StarWave" and chaired the session on 40-100Gbps Technology Research. Alan Verlo attended.

March 2011. Tom DeFanti was in Australia and Singapore, and gave a number of presentations.

- March 9, 2011. Presented "Eliminating the Bottlenecks in Visualization of Large Data and Research Collaboration," at Monash University as part of the opening of the MASSIVE

visualization center.

- March 17, 2011. Presented "Eliminating the Bottlenecks in Visualization of Large Data and Research Collaboration," at the CSIRO conference on Computational and Simulation Science in Melbourne in the eResearch Technologies session.
- March 19, 2011. Invited speaker. Presented "GreenLight Energy Usage Displays" at the Third Workshop on Cloud-Mobile Convergence for Virtual Reality (CMCVR 2011), held in conjunction with IEEE VR 2011 in Singapore <www.cmcvr.org/CMCVR11.html>.
- March 20, 2011. Invited speaker. Presented "The Future of the CAVE" at the JST/CREST Workshop held in conjunction with IEEE VR 2011 in Singapore
- March 23, 2011. Invited speaker. Presented "Beyond Digital Cinema: Globally Shared Visualization and Virtual Reality" at the School of Art, Design, and Media at Nanyang Technological University in Singapore.
- March 25, 2011. Presented "The Future of the CAVE" at University of Queensland.

February 8, 2011. Kevin Thompson, NSF, was at UCSD and visited with Tom DeFanti, Tajana Rosing and Jerry Sheehan at Calit2. They discussed network measurement (with respect to GreenLight), international networking (with respect to GreenStar), and advanced visualization applications.

January 31, 2011. Alan Verlo participated in a JET meeting held at Joint Techs.

January 30 – February 3, 2011. Alan Verlo attended the Winter 2011 ESCC/Internet2 Joint Techs meeting in Clemson, South Carolina.

January 28-31, 2011. Joe Mambretti attended iGENI and GLORIAD meetings with SingAREN in Singapore. The group inaugurated the new SingLight GLORIAD/GLIF GOLE on January 31.

January 27, 2011. The Chinese Academy of Sciences (CAS) Computer Network Information Center (CNIC) and Joe Mambretti of Northwestern/iCAIR (International Center for Advanced Internet Research at Northwestern University) held a meeting on Internet innovation and experimentation cooperation in Beijing. CNIC and iCAIR signed Memorandum of Understanding (MOU) during this meeting to enhance cooperation in future Internet related activities including research, development, and networking. Dai Bowei, vice director of CNIC and director of CSTNet (China Science and Technology Network), Joe Mambretti, director of iCAIR, director of Metropolitan Research and Education Network (MREN) and PI of iGENI, signed the MOU on behalf of CNIC and iCAIR respectively. Emily Yaung Ashworth, director of the NSF China Office, Chen Mingqi, director of Division of Informatization Affairs of General Office, and Zhang Shizhuan, vice director of Division of American and Oceanian Affairs in Bureau of International Cooperation, also witnessed this ceremony, together with other experts from CNIC. CNIC with CSTNet and iCAIR will undertake cooperative partnership activities in three areas: to design and implement processes that optimize research and development of the future Internet through programs based on new architecture, technology, and experimental testbeds; undertake the development and implementation of experimental research testbeds and cooperate on developing the International Global Environment for Network Innovations (iGENI); and, design and implement processes that will enable communities to researchers to use experimental testbeds. During the meeting, both sides shared their experiences and achievements on network construction, innovative research and e-science applications. Mambretti gave the presentation "Creating 21st Century Communication Services and Networks: Architecture, Technology, and Facilities." For further information, see <http://english.cnic.cas.cn/ns/es/201101/t20110130_64913.html> and <http://www.cnic.cn/xw/rdxx/201101/t20110130_3069456.html>.

January 14, 2011. Tom DeFanti participated in the Australian American Leadership Dialogue (AALD) workshop, held at Calit2, and spoke in a panel on wired broadband futures and GLIF international networking, particularly with respect to Australia's AARNet. AALD is a group of government, business and academic leaders from Australia who have worked at the highest levels with their US counterparts for 20 years. Luc Renambot (EVL) demonstrated SAGE between two OptIPortals in Chicago/EVL and San Diego/Calit2.

January 10, 2011. Calit2 had 30 visitors from the Harvard Business School. During site tours, Tom DeFanti did a collaborative SAGE demonstration between Calit2 and EVL.

January 6, 2011. Larry Smarr, Calit2, was a Distinguished Lecturer for the Hawaii International Conference on System Sciences (HICSS-44) in Kauai, HI. He gave the presentation "Building a Global Collaboration System for Data-Intensive Discovery" and talked about SAGE, OptIPortals and international collaboration. His PPTs are at <http://lsmarr.calit2.net/presentations?slideshow=6482766>.

December 13, 2010. Instituto Tecnologico de Costa Rica (ITCR) became a PRAGMA member and contacted other PRAGMA members about using OptIPortals to collaborate on metagenomics. While UIC/EVL does not do metagenomics, we offered some alternatives. Ideas are still pending.

December 6-17, 2010. Tom DeFanti had visitors from Petrobras, a Brazilian oil conglomerate, and their Brazilian university partners, who are charged with building a networked visualization center. They visited Calit2 to learn what to do. They also attended the CineGrid Workshop. Luc Renambot met with them December 9-10 to talk about SAGE and how it enables video streaming over international high-performance networks. Alan Verlo met with them to talk about StarLight, GLIF and our collaborations with RNP in Brazil. Tom DeFanti covered such topics as building NexCAVEs, OptIPortables, SAGE, and the CineGrid Exchange. The visitors were: Luciano Pereira Dos Reis, Technical Consultant, Reservoir Geological Engineering, Petroleo Brasilia, S.A., Mario Pimenta, responsible for the IT issues related to the Visualization Center at Petrobras, and from Pontifícia Universidade Católica (PUC), Alberto Raposo, the manager of the VR team, and Thiago Bastos Tecgraf/PUC-Rio (Pontifícia Universidade Católica).

October 13-14, 2010. Tom DeFanti and Tajana Rosing (UCSD), Maxine Brown (EVL) and Joe Mambretti (NU) participated in the GLIF 10th Annual Global LambdaGrid Workshop, held at CERN in Geneva, Switzerland. Brown is co-chair of the GLIF Research & Applications Working Group, and as part of her session, Tom DeFanti gave the presentation "TransLight/StarLight: The Lightpath is not the Goal, the GOLE is the Goal." SAGE was one of several international application and middleware experiments mentioned in his presentation. In the same session, Mambretti gave two presentations: "HPDMnet @ GLIF 2010" and "iGENI." Mambretti also did HPDMnet demos and assisted Poznan Supercomputer Center to showcase their 4K initiative (see Section 2.B.1.3: HPDMnet).

August 11-12, 2010. Jason Leigh (EVL) visited Seoul National University and KISTI in Korea. At both locations, he gave the presentation "Global Cyber-Commons: Supporting Global Collaborative Research, Development, and Education in Cyber-Infrastructure-Enhanced Environments." *In the next three years, the Korean government wants to construct and operate ~40 persistent Cyber-Commons-like collaboration nodes for education and research, all running SAGE. Currently, OptIPortal-like facilities in Korea consist of one at KISTI and six at national universities, used for remote education.* Jason met with Jysoo

Lee, head of the KISTI Supercomputer Center, to discuss system architecture, software, and content.

August 9, 2010. David Lassner, IRNC TransLight/Pacific Wave principal investigator, gave an IRNC update at the APAN meeting in Hanoi. He requested Powerpoint slides from IRNC projects to present.

2.B.2.3. Education, Outreach and Broader Participation

EVL and Calit2 do a number of tours for high-school students and undergraduate students to excite them about going to college and to encourage them to pursue careers in science and/or engineering. Other tours include campus VIPs, Board of Trustee members, press, etc. Most tours consist of an overview presentation of collaborative research, including IRNC/GLIF activities, followed by hands-on demonstrations of advanced, networked visualization technologies. EVL and Calit2 participated in the following broader outreach activities over the past year:

South Metro Career Center

Casa Familiar

May 13, 2011. Calit2 built and deployed two OptIPortables running SAGE to community centers in the San Diego area. They previously brought in staff and students who work/visit the center to teach them SAGE. In the past, Calit2 has constructed or helped other universities build their own tiled display walls, because the systems make it easier for users to access, visualize and analyze data remotely. Now Calit2 is deploying OptIPortables closer to home – helping underserved communities in San Diego to interact with one another and with UCSD, without leaving their own neighborhoods. The technology is expected to be particularly useful because underserved communities in San Diego are often so geographically distant from academic and commercial centers (and from one another) that it can impede education, job growth and social reform. One is deployed to South Metro Career Center (SMCC) in San Diego's Mount Hope community, and another to Casa Familiar in San Ysidro. The hope is that both organizations will network with one another via their OptIPortables and to Calit2 to enable communities to share experiences, content and programs with one another and be tightly coupled with UCSD students. More information is available at <http://www.calit2.net/newsroom/article.php?id=1808>.

April 13-15, 2011. Sukumar and Joe Keefe of Calit2 did SAGE network demonstrations using SAGE at the 3rd annual site visit of the NSF ERC Center for Integrated Access Networks (CIAN) at University of Arizona. Maxine Brown and Luc Renambot provided presentation materials.

April 13, 2011. EVL hosted teacher Melvin Slater and 30 students from Austin Polytech High, who visited the UIC campus under the auspices of the University of Illinois Affiliate Project Lead The Way (PLTW) program.

April 6, 2011. EVL hosted a local troupe of ~30 Girl Scouts ("Brownies," ages 7-9). They primarily got hands-on demonstrations of the technologies, and thought it the best field trip, ever! Some now want to be professors when they grow up.

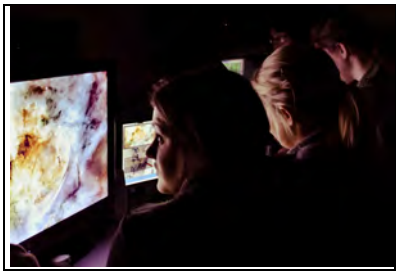

March 1, 2011. EVL hosted 20 high-school girls attending the UIC WISE GEM-SET pre-college outreach program (Women in Science and Engineering Program's Mentoring for Success, a grant funded by the U.S. Department of Education Women's Educational Equity Act WEEA Program).

February 21 2011. EVL had an open house for UIC Engineering Week, and also hosted a Cub Scout troupe from Indiana.

October 27, 2010. Maxine Brown of EVL hosted a teacher and 18 students from Chicago's Austin Polytechnic High School as part of the University of Illinois Affiliate Project Lead The Way (PLTW) program. The students wrote a review of their trip to UIC for their school paper, noting, "We found out that the lab [EVL] was founded in 1973 and Jason Leigh is the current director. The lab supports about 18 graduate students and 6 undergraduate students. The lab researches advanced display systems, visualization systems, high-speed networking and more. This gave students an example of what a particular day at college would be like…"

October 12, 2010. EVL hosted journalism students from Northwestern University's Medill School. Their professor, Donna Leff, takes her students to visit a variety of places, with the goal of having the students write articles about subjects they find interesting.

September 24, 2010. EVL hosted 2nd graders from a local school (a UIC employee's child and her friends). They enjoyed EVL's interactive environments.

July 22, 2010. A group of teachers and college-eligible African American male high school juniors visited UIC to learn more about its programs. As several students expressed interest in studying computer science, the group visited EVL and got demos and tours.

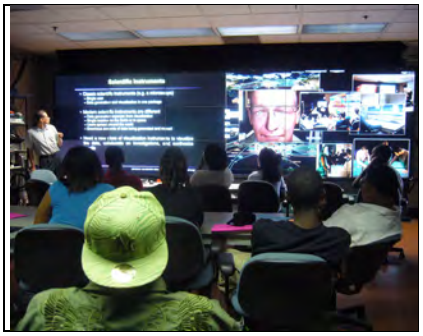

July 9, 2010. Chicago State University Minority Engineering program staff and students, who are future UIC Engineering transfer student candidates, visited UIC and took tours of the campus. One stop was EVL.

2.B.3. GLIF Partnership

2.B.3.1. GLIF GOLE Governance

Earlier this year, GLIF asked Bill St. Arnaud to investigate GOLE (GLIF Open Lightpath Exchanges) Governance. As members of the StarLight management consortium, Joe Mambretti and Tom DeFanti have been involved in several discussions. Notably, St. Arnaud is exploring how inter-GOLE lightpaths are managed; for the most part, GOLEs at each end establish their own access policies. In the future, we may need new Internet architectures built around GOLEs – creating a pool of "federated" inter-GOLE lightpaths with common policies.

More significantly, new major eScience programs like LHCONE are designed specifically around the concept of open lightpath exchange points, which is likely to be adopted by other big science programs. The concept of federated networks interconnecting at open lighpath exchange points is also likely to be the major architectural underpinning of future cyberinfrastructure, as well as R&E and community network architectures in the US and North America.

We need to move away from traditional hierarchical networks to GOLEs enabling the direct peering of networks, institutions and/or researchers. GOLEs enable a policy free interconnection with no bandwidth constraints or blocking between the connecting parties and therefore research is not constrained by policy or bandwidth issues as on traditional networks. GOLEs are also important for the R&E networks in the new research-intensive nations like Brazil, China, South Africa, Korea, etc. In the future, networks and researchers of these nations will be far less dependent on having to transit intermediary networks by peering directly with networks or institutions in designated host nations.

For example, Internet 2 has made a proposal to the NSF to build a number of distributed optical exchange points in the US to support the needs initially of high-energy physics, but other big science applications as well. Given these recent developments, a number of organizations involved with cyberinfrastructure and optical networking have agreed that it would be useful to start discussing "policy/governance" issues with respect to these developments.

In this case, governance does not imply any type of central control or management, but to work with the GLIF, network operators, users and funding agencies to set in place common agreements, definitions, funding models and policies as to what defines an open exchange and how resources designated for common usage should be allocated.

The CANARIE UCLP (User Controlled LightPaths) and more recently Internet 2's OS3E service are intended to allow institutions, or even researchers to establish their own private networks for specific virtual organizations or communities of interest who can interconnect at these GOLEs.

GOLEs will challenge traditional business models for R&E networking, but the first priority must be to enable the needs of the researchers themselves in the exponential growth of data-driven science.

2.B.3.2. GLIF 10th Annual Global LambdaGrid Workshop

Tom DeFanti, Tajana Rosing, Maxine Brown and Joe Mambretti attended the GLIF 10th Annual Global LambdaGrid Workshop, 13-14 October 2010 at CERN, Geneva, Switzerland. Brown was a member of the Program Committee. Brown also serves as co-chair of the GLIF Research & Applications (RAP) Working Group, and as part of the RAP session during the meeting, Tom DeFanti gave the presentation "TransLight/StarLight: The Lightpath is not the Goal, the GOLE is the Goal." GreenLight and SAGE were some of the international application and middleware experiments mentioned in his presentation. In the same session, Mambretti gave two presentations: "HPDMnet @ GLIF 2010" and "iGENI." Mambretti also did HPDMnet demos and assisted Poznan Supercomputer Center to showcase their 4K initiative (see Section 2.B.1.3).

2.B.3.3. GLIF Map

In May 2011, GLIF published new world maps showing GLIF GOLEs and circuits. Maxine Brown compiled the data for this map over the past year, which was produced by Robert Patterson of the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC). More information about the map can be found in the press release <http://www.glif.is/publications/press/20110525.html>. The world map, as well as closeup regional maps, can be downloaded from the GLIF website <http://www.glif.is/publications/maps/>.

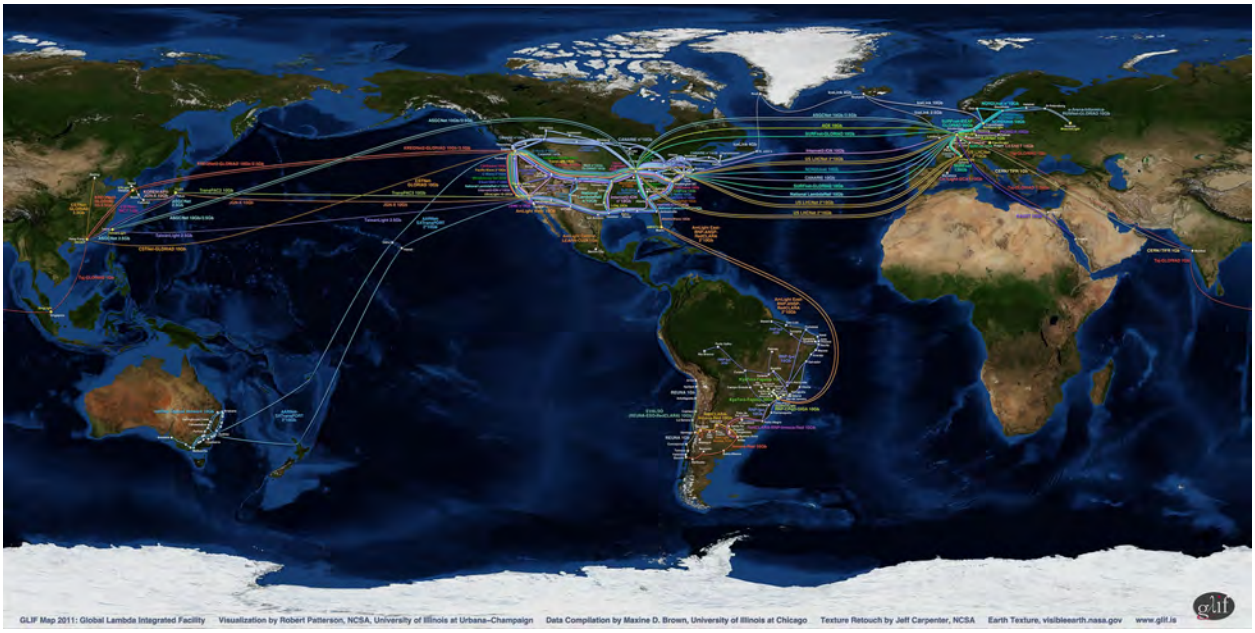

2.B.4. New Network Services: Network Engineering and Consulting

In addition to participating in the many experiments described in Sections above, the TransLight/ StarLight team also worked with international partners to facilitate network engineering at StarLight.

- Alan Verlo facilitated the peering of US LHCNet and Internet2-ION at StarLight
- The Universidad Nacional Autonoma de México (UNAM) is connected to StarLight through NLR (1Gb). In May 2011, they began implementing a redundant connection to StarLight through Cogent (also 1Gb). These connections are used to bring LHC data to UNAM.
- In May 2011, Greg Cole e-introduced Tom DeFanti, Joe Mambretti and Maxine Brown to Taj-GLORIAD Indian partners Leena Wadia, Spenta Wadia and Amit Apte. Cole is working with them to redirect the current Taj-GLORIAD circuit in Mumbai to Bangalore. There is interest in creating a StarLight-like facility in Bangalore. We provided the Indians with several journal articles on

StarLight, TransLight and OptIPuter. (Subsequently, Amit Apte visited Chicago on May 28, 2011, and met with Joe Mambretti and Maxine Brown.)

- In May 2011, Greg Cole e-introduced Tom DeFanti, Joe Mambretti and Maxine Brown to Taj-GLORIAD Egyptian partner Ola Laurence. Prior to this, in October 2010, Joe Mambretti and Gigi Karmous-Edwards consulted with Greg Cole on GLORIAD-Egypt 1Gbps connectivity to NetherLight, as well as Telecom Egypt's interest in developing a regional GOLE in Cairo.
- As of April 2011, StarLight supports the Tier 2 Atlas Project Network, devoted to CERN's Atlas LHC experiments, primarily focused at University of Michigan and Michigan State University. Here is the network topology.

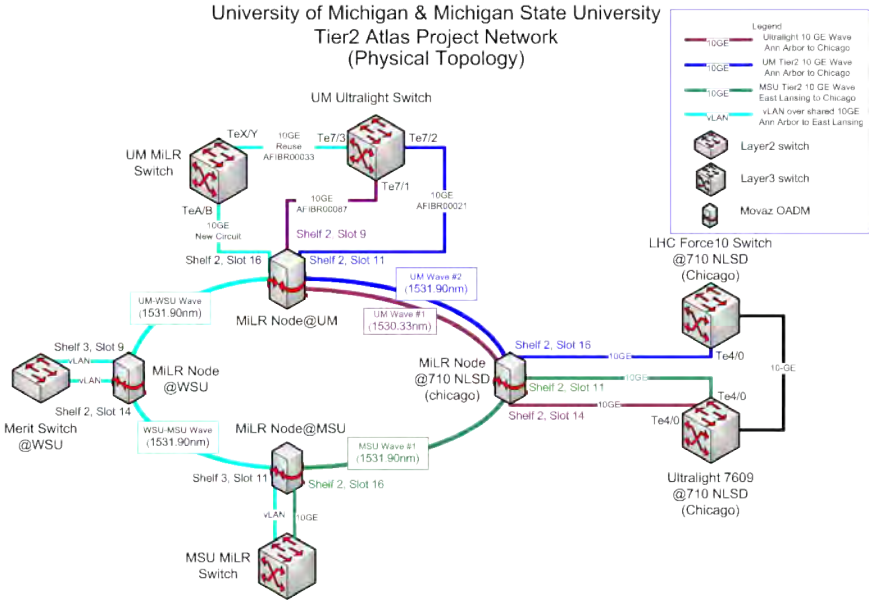

- Starting April 2011, StarLight has been working with the NSF Ocean Observatories Initiative (OOI) to create OOInet. While not global *yet*, this national network has potential.

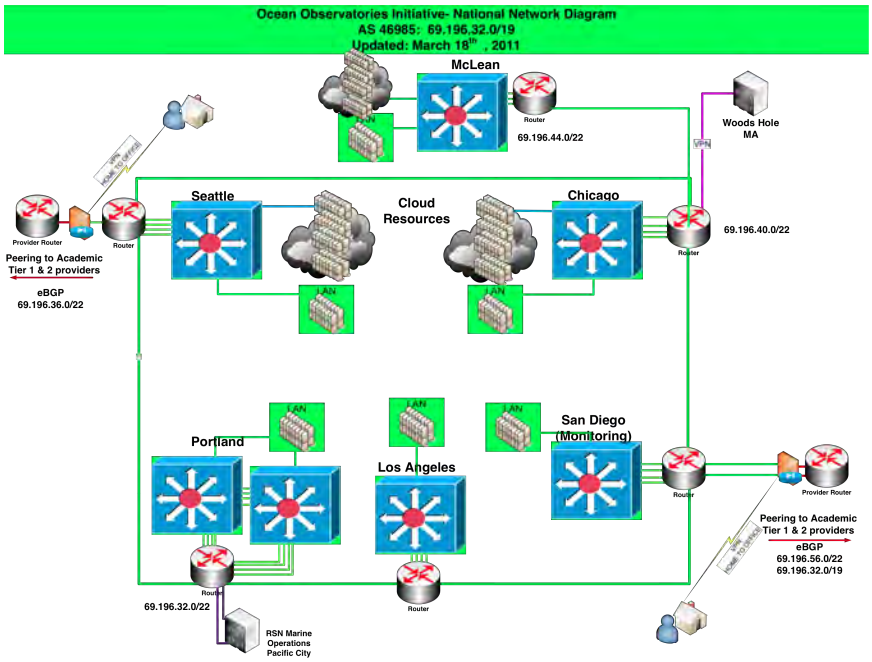

- In April 2011, David Foster of CERN sent Mambretti a nice note saying "We have just had some success with LHCONE and as it's running through StarLight, I thought I should mention it to you." LHCONE (LHC Open Network Environment) is a new LHC data-distribution network model to ensure better access to the most important datasets by the worldwide HEP (High Energy Physics) community and hence improve data analysis. Specifically, LHCONE flattens the Tier1-Tier2-Tier3 hierarchy MONARC model so that any site may connect with any other; revisits the implementation of strategic data placement; and, reinforces the use of dataset caching – rather like caching web pages. Initially implemented with existing infrastructure, LHCONE has started exchanging traffic between CERN and Caltech; the route is: CERN-AMS (SURFnet-CERN fiber 10G) - NY (USLHCNET 5G) - CHI (USLHCNET 5G) - Caltech (Shared 10G). Traffic monitoring of the transatlantic segment is at <http://repository.uslhcnet.org/display> and statistics for June 2, 2011 are below.

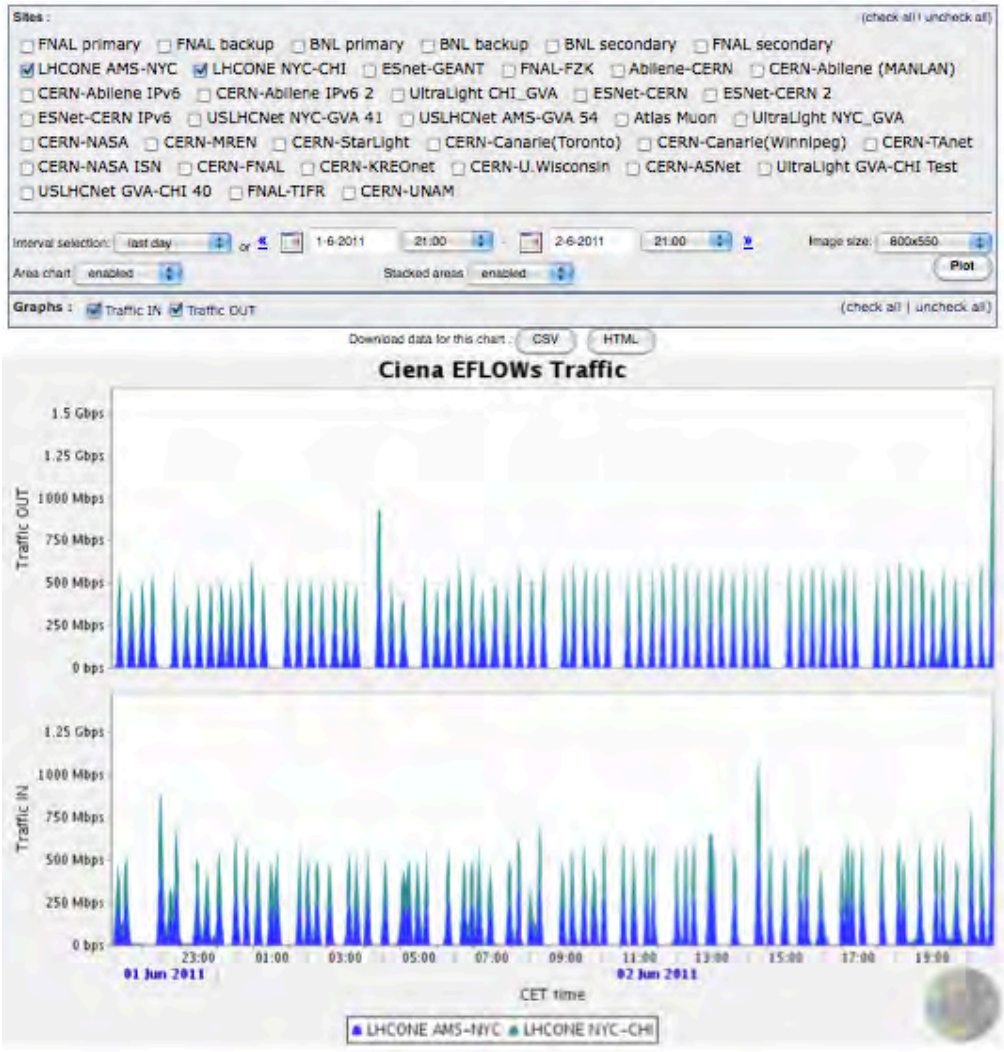

- Joe Mambretti consulted with Louis Fox. University of Washington/PNWGP received an NSF ARI award to create Northern Wave, a distributed facility similar to Pacific Wave that connects from Seattle eastward. Funding isn't adequate to purchase circuits all the way to Chicago. The current plan is to get to Minneapolis, where Northern Wave has an agreement with BOREAS to get to StarLight.
- In September 2010, Joe Mambretti consulted with Greg Cole about an IRNC GLORIAD Chicago-to-Stockholm 10Gbps circuit for US/Russia connectivity.

2.C. Research Training and Development

For over two decades, TransLight/StarLight senior leadership has developed and nurtured national/international partnerships, communication channels, forums, and processes to ensure ongoing successful interactions among its constituents. Our success has been in the development of teams, tools, hardware, system software, and human interface models on an accelerated schedule to enable multi-site collaborations for complex problem solving.

TransLight/StarLight and its global partners have elected to work together, on a persistent basis, to further the goals of international e-science collaboration. Partners consist of researchers, middleware developers, network engineers and international NRENs who provide all the layers of global cyberinfrastructure – to develop "grass-roots" experiments using "best-of-breed" processes in developing next-generation, shared, open-source optical networking services. Together, this group facilitates greater advances in global networking than single-investigators can afford to do on their own.

While TransLight/StarLight leadership consists primarily of computer scientists, it always puts the needs of computational scientists first, and works with application drivers to create useful and usable infrastructure, technologies, tools and techniques to enhance e-science.

TransLight/StarLight does research training via web documentation, journal articles, presentations and demonstrations at major networking conferences and workshops (e.g., GLIF, SC), scientific conferences (e.g., AAAS), and NREN user meetings (e.g., Internet2, NLR), as well as distribution of PowerPoint presentations and other instructional material.

2.D. Outreach Activities

For the computer science, computational science and networking communities, TransLight/StarLight's application and middleware efforts directly enhance infrastructure for research and education. Project results are published on the web, in journal articles, in PowerPoint presentations and in other teaching materials provided to collaborators. Project results are also demonstrated at major conferences and workshops. TransLight/StarLight principals have also gotten successful news media coverage (television, print magazines), providing avenues to broadly disseminate their research and enhance public knowledge.

In addition, Calit2, UIC/EVL and Northwestern host numerous international visitors and work closely with their respective campus recruitment and community minority outreach programs to organize lab tours for K-12 teachers and students, to excite them to attend college and to pursue advanced degrees.

Calit2 and UIC faculty also teach undergraduate and graduate computer science classes, and integrate research activities into the curriculum. At UIC/EVL, one such course on Video Game Design and Development attracts non-computer-science majors as well as computer science students. The course is simultaneously taught both at UIC and Louisiana State University using video-teleconferencing systems and high-speed networks. Students are given the semester to create a video game – they form small teams with one another *and* with half the students from Chicago and half from Louisiana. And, they must implement these games on advanced visualization devices whose development is funded by NSF MRI awards. Opportunities such as these help train the next-generation global workforce.

3. Publications and Products

3.A. Journals/Papers

Joe Mambretti, Mathieu Lemay, Scott Campbell, Hervé Guy, Thomas Tam, Eric Bernier, Bobby Ho, Michel Savoie, Cees de Laat, Ronald van der Pol, Jim Chen, Fei Yeh, Sergi Figuerola, Pau Minoves, Dimitra Simeonidou, Eduard Escalona, Norberto Amaya Gonzalez, Admela Jukan, Wolfgang Bziuk, Dongkyun Kim, KwangJong Cho, Hui-Lan Lee, Te-Lung Liu, "High Performance Digital Media Network (HPDMnet): An Advanced International Research Initiative and Global Experimental Testbed," Future Generation Computer Systems, Elsevier, Vol. 25, 2011, doi:10.1016/j.future.2010.12.012

3.B. Books/Publications

Gordon Cook, "Out of Many One – Resurrected: How Human Ingenuity, DIY Technology, and Global R&E Networks Are Remaking the World," Cook Report, February - April 2011, 431 pp., <http://www.cookreport.com/index.php?option=com_content&view=article&id=261:feb-apr-2011&catid=38:current-issues&Itemid=73>

Joe Mambretti, Tom DeFanti, Maxine Brown, "StarLight: Next-Generation Communication Services, Exchanges, and Global Facilities" (chapter), Advances in Computers, Vol. 80, Marvin V. Zelkowitz (editor), Elsevier, 2010, pp 191-207, doi: 10.1016/S0065-2458(10)80005-1

3.C. Internet Dissemination

www.startap.net/translight

3.D. Other Specific Products

Other than the information reported here, we have not developed any other specific product of significance.

4. Contributions

4.A. Contributions within Discipline

TransLight/StarLight advances knowledge and understanding within the high-performance computing and networking community. TransLight/StarLight experiments and experiences help lead and motivate services for 21st century e-science and computer science innovation by working with global cyberinfrastructure and e-science virtual organizations that use and/or develop next-generation grid services.

4.B. Contributions to Other Disciplines

All science is not well served by one protocol at one network layer. Through its aggressive use of networks to conduct end-to-end experiments, TransLight/StarLight is helping discover new methods and technologies to customize services and capabilities for individual science disciplines.

4.C. Contributions to Human Resource Development

TransLight/StarLight disseminates results broadly to enhance scientific and technological understanding among computer scientists, computational scientists, network engineers, faculty, staff and students. Its leadership is committed to establishing global collaborations that benefit US academic institutions, industry, and government laboratories. Through the development of e-science application and middleware experiments, community partnerships, GLIF partnerships and new services, TransLight/StarLight is stimulating and supporting the development and dissemination of next-generation instrumentation, multi-user facilities, and other shared research and education platforms.

4.D. Contributions to Resources for Research and Education

TransLight/StarLight has international partnerships, communication channels, forums, and processes to ensure ongoing successful dissemination of information and results to its constituents. Its leadership continually works with domain scientists to better understand application requirements and the need for customized services; has organized workshops and events to accelerate network adoption and advancements; has produced special issues of international journals to disseminate success stories; has participated in other NSF-funded efforts such as Blue Waters, TeraGrid and PRAGMA; has an active role in GLIF and is instrumental in producing updated GLIF maps; helps facilitate demonstrations at major international workshops and conferences, such as SC, AAAS, Internet2 and GLIF; and, incorporates its research into undergraduate and graduate computer science class curricula.

4.E. Contributions Beyond Science and Engineering

Helping pioneer new technologies, tools and techniques among virtual organizations of computer scientists and computational scientists is helping accelerate solutions to issues of global importance, such as climate change, energy, homeland security, disaster response, etc. In addition, TransLight/StarLight works with network equipment manufacturers and telecommunication providers to create and showcase new markets for wavelength-based network services and products. Our users expect us to grow in capacity and sophistication, and we look forward to the technical challenges ahead.

5. Conference Proceedings

None.

6. Special Requirements

6.A. Objectives and Scope

A brief summary of the work to be performed during the next year of support if changed from the original proposal.
Our scope of work has not changed.

6.B. Special Reporting Requirements

Do special terms and conditions of your award require you to report any specific information that you have not yet reported?
No.

6.C. Animals, Biohazards, Human Subjects

Has there been any significant change in animal care and use, biohazards, or use of human subjects from what was originally approved (or approved later)?
No.

7. Summary of Switch Throughput and Power Measurements

This is the complete report of the study mentioned in Section 2.B.1.2.

Summary of switch throughput and power measurements

Eric Dorman with Prof. Rosing
Greenlight Project, UCSD

Introduction

This document describes the power and throughput measurements recorded from experiments over 1G and 10G packet switches and an optical switch, both in the Greenlight box and in Calit2.

Switches being measured

So far I've measured the 1G and 10G switches in the Greenlight black box, and the Glimmerglass optical switch in Calit2. There are multiple 1G switches in the black box. For 1G measurements, I've been using the 1G rack-level switches on racks 2 and 4. There is only one 10G switch in the box, and I've been using that one for the 10G measurements.

Here are the switch specs:

- 1G Greenlight switch**:** X350-48ts Extreme
  http://www.extremenetworks.com/libraries/products/DSSumX350_1399.pdf
- 10G Greenlight switch: Fujitsu XG2000
  http://www.fujitsu.com/us/services/computing/peripherals/ethernet/xg2000.html
- Optical switch: Glimmerglass MEMS gg308-c

Greenlight host accounts

Rack	OS	Host details	GridFTP install	Used for
2	Linux	KOI Supermicro headnode with Intel Xeon E5430 CPU and 32GB memory	v5.2	1G measurements
4	Linux	KOI Supermicro headnode with Intel E5430 CPU and 32GB memory	v5.2	1G measurements
5	Solaris	Sun Fire X4540 storage server, two Quad-Core AMD Opteron 2356 CPU, and 32GB memory	v4.2	10G measurements
7	Solaris	Sun Fire X4540 storage server, two Quad-Core AMD Opteron 2356 CPU, and 32GB memory	v4.2	10G measurements

Figure 1: Greenlight host accounts.

Throughput measurements

All throughput measurements over the packet switches were taken using GridFTP. We are using two different versions of it (for OS compatibility reasons): v5.2 on racks 2 and 4 (used for 1G transfer) and v4.2 on racks 5 and 7 (used for 10G transfer). We are using the "ftp" protocol for the transfers, not "sshftp". We perform non-disk transfers (because disk speeds become a bottleneck in higher-bandwidth connections). The format of the GridFTP commands is thus as follows:

On client side, for non-disk transfer:

```
$ globus-url-copy -p <parallelism> -tcp-bs <tcp buffer>
file:///dev/zero ftp://<destinationip>/dev/null
```

On server side:

```
$ globus-gridftp-server -control-interface <server ip> -aa
-p 5001
```

We calculate tcp-bs according to the following formula:

*-tcp-bs = bandwidth in Megabits per second (Mbs) * RTT in milliseconds (ms) * 1000 / 8*

We determine optimal parallelism experimentally.

Power measurements

All the switches we have been measuring are connected to Avocent PDUs, which we are polling via SNMP. Our power measurement methodology has been as follows: first, we take baseline power readings over the past day, week, month and year. This gives us a baseline power draw. Then we do a sustained (5 minute) transfer over the switch. We measure power levels during this transfer and compare them to our baseline.

Note that because our switches have varying numbers of ports in use, baseline power tends to vary substantially between switches. In these measurements, we are more interested in power variation that baseline. To get an approximate value of what the power consumption of a switch would be with all of its ports in use, the reported power value should be scaled by the number of unused ports (note that this is a very rough approximation). Below is a table containing the port usage for every switch at the time these measurements were taken.

Switch	**Active ports**	**Total Ports**
Top rack-level switch for Rack 2 (1G)	30	48
Top rack-level switch for Rack 2 (1G)	39	48
10G switch on Rack 5	~15	20
Glimmerglass optical switch	47	48

Detailed experiment results

Here we detail the results of four experiments, presenting both power and throughput measurements:

1) Transfer over 1G switches between two racks in the Greenlight blackbox
2) Transfer over a 10G switch between two racks in the Greenlight blackbox
3) Transfer from Calit2 to Greenlight blackbox
4) Transfer between two hosts over Glimmerglass MEMS optical switch

Experiment 1: Transfer over 1G switches between two racks in the Greenlight blackbox

Configuration	Transfer from head node on Rack 2 to head node Rack 4 in the Greenlight blackbox. These racks are geographically isolated, but very close. There are two 1G rack-level switches on each rack which are being measured.
Bandwidth	1Gbs
RTT	1.2 ms
TCP-BS	153600 bytes
Parallelism	12
Throughput	112.2 MB/sec

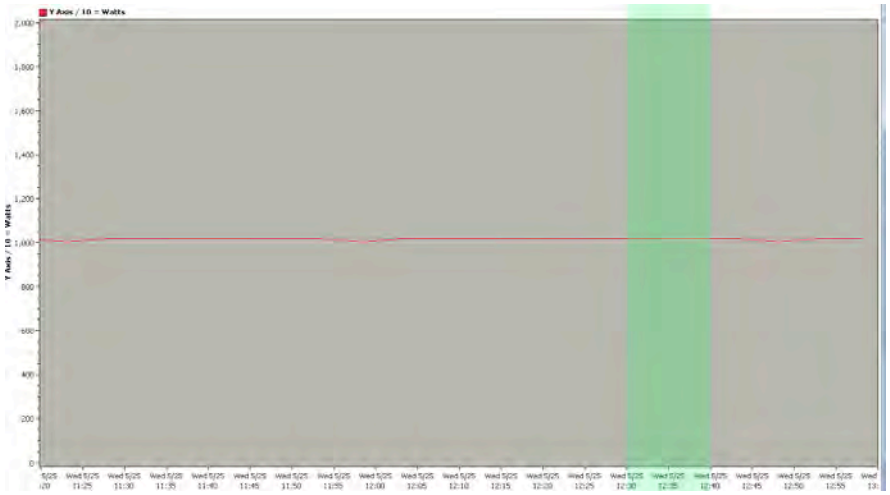

Figure 1: The power level (in volt-amps) of the top Rack 2 switch over the period of an hour. A ten-minute GridFTP data transfer utilizing the entire bandwidth of the switch occurs from 12:30 to 12:40, and is shown in green.

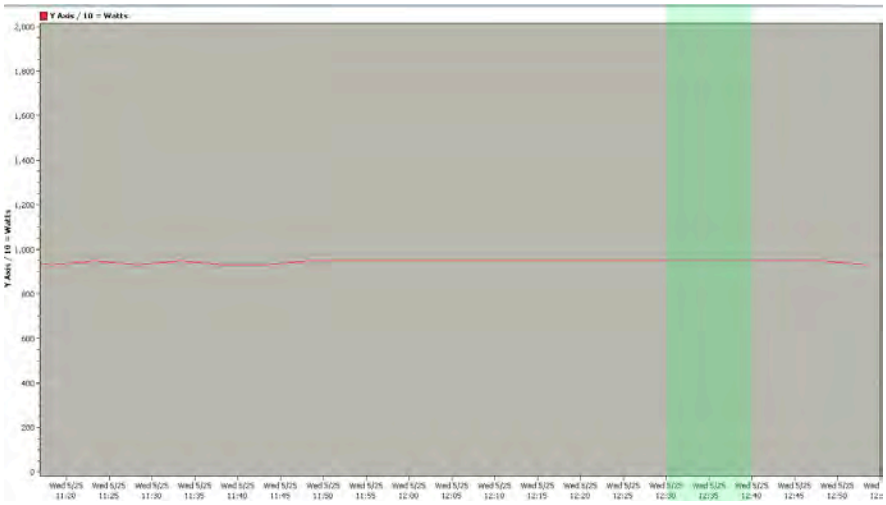

Figure 2: The power level (in volt-amps) of the top Rack 4 switch over the period of an hour. A ten-minute GridFTP data transfer utilizing the entire bandwidth of the switch occurs from 12:30 to 12:40, and is shown in green.

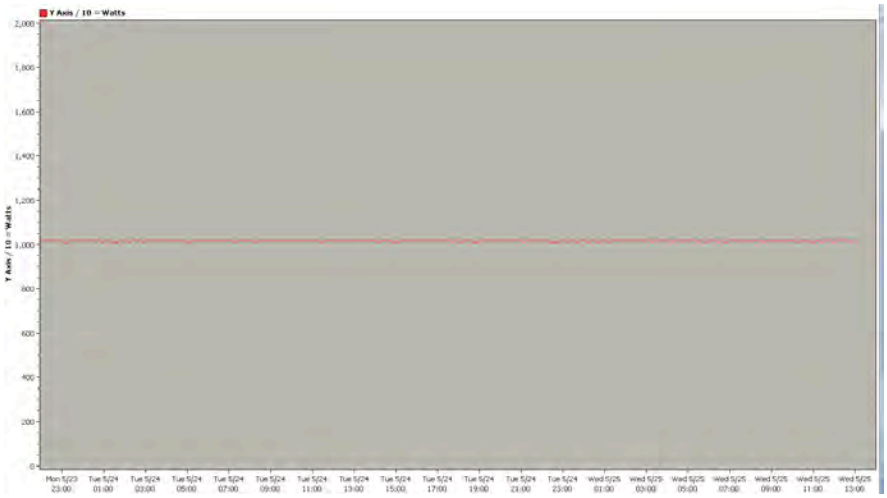

Figure 3: The power level (in volt-amps) of the top Rack 2 switch over the period of a day.

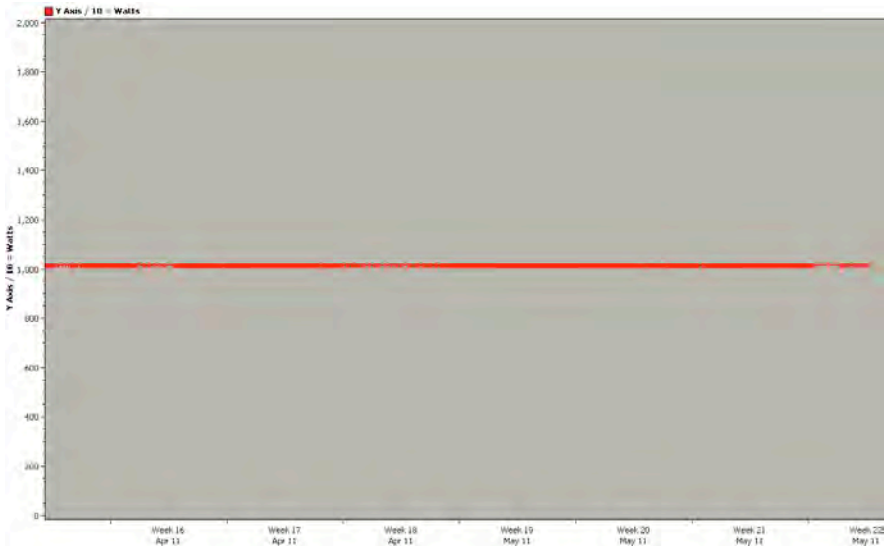

Figure 4: The power level (in volt-amps) of the top Rack 2 switch over the period of a month.

Experiment 2: Transfer over 10G switch between two racks in the Greenlight blackbox

Configuration	Transfer from a node on Rack 5 to a head node on Rack 7 in the Greenlight blackbox. Both nodes have connectivity to a 10G switch on Rack 5. These racks are geographically isolated, but very close. The switch being measured is the 10G switch.
Bandwidth	10Gbs
RTT	0.241 ms
TCP-BS	308480 bytes
Parallelism	16
Throughput	897.1 MB/sec

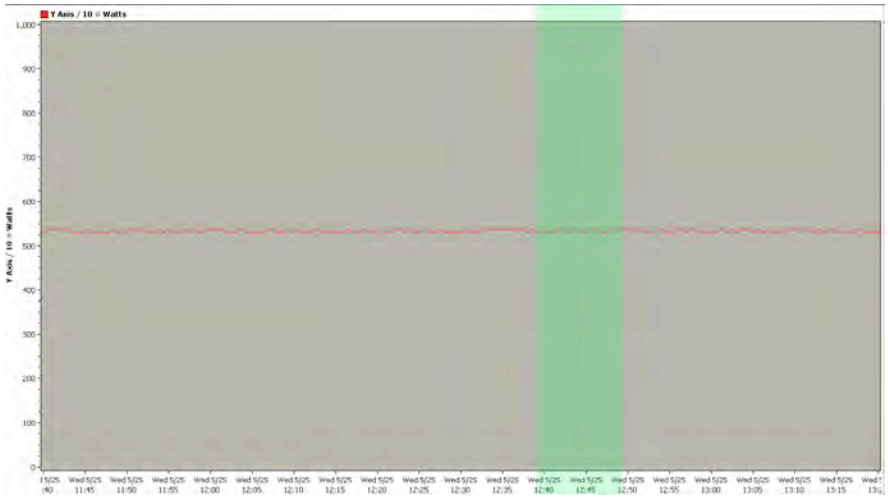

Figure 5: The power level (in volt-amps) of the 10G Rack 5 switch over the period of an hour. A ten-minute GridFTP data transfer utilizing the entire bandwidth of the switch occurs from 12:38 to 12:48, and is shown in green.

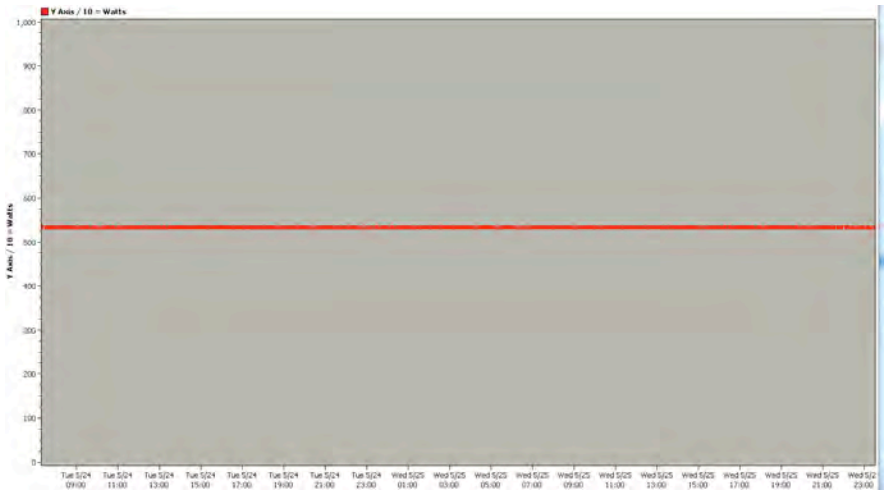

Figure 6: The power level (in volt-amps) of the 10G Rack 5 switch over the period of a day.

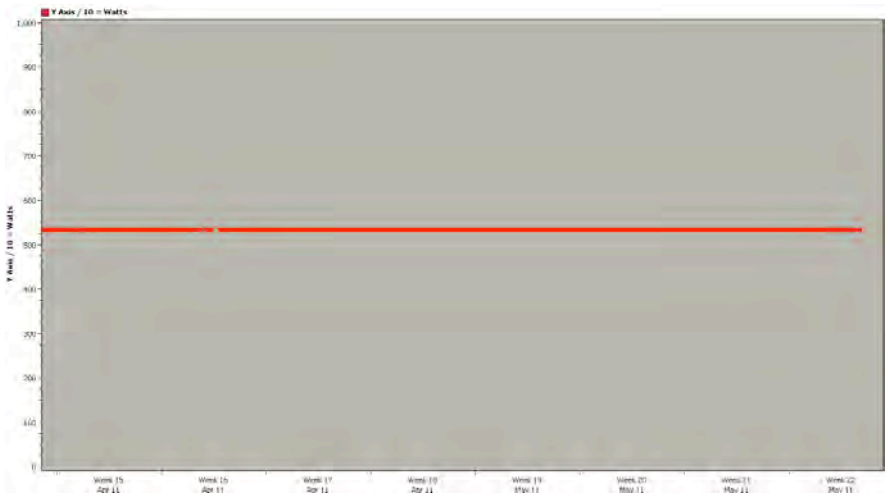

Figure 7: The power level (in volt-amps) of the 10G Rack 5 switch over the period of a month.

Experiment 3: Transfer over 10G switch from Calit2 to Greenlight blackbox.

Configuration	Transfer from a node on Rack 5 to a head node on Rack 7 in the Greenlight blackbox. Both nodes have connectivity to a 10G switch on Rack 5. These racks are geographically isolated, but very close. The switch being measured is the 10G switch.
Bandwidth	10Gbs
RTT	1.95 ms
TCP-BS	1MB
Parallelism	24
Throughput	481.4 MB/sec

The power results for this transfer are the same as they were for Experiment 2.

Experiment 4: Transfer over optical switch.

Configuration	Transfer between two hosts, both connected to a Glimmerglass MEMS optical switch and a 10G packet switch. Experiments were conducted by sending traffic through the optical switch in one path, then sending through in another, forcing the Glimmerglass to rotate its mirrors to allow for the new flow. Traffic was generated using IPERF, not GridFTP, due to testbed limitations.
Bandwidth	Limited by 10Gbs NIC
RTT	NA
TCP-BS	NA
Parallelism	NA
Throughput	1.17 GB/sec (temporarily drops substantially when switching between flows)

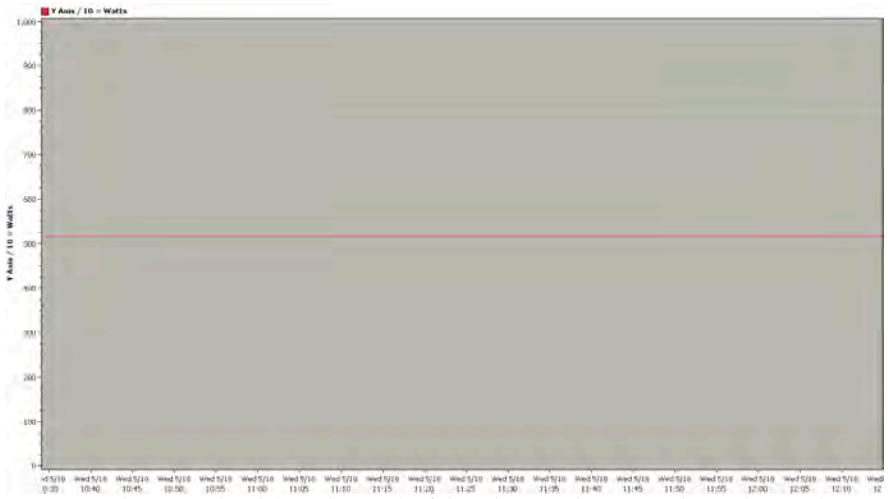

Figure 7: The power level (in volt-amps) of the Glimmerglass optical switch over the period of an hour, during which several data transfers using IPERF were conducted.

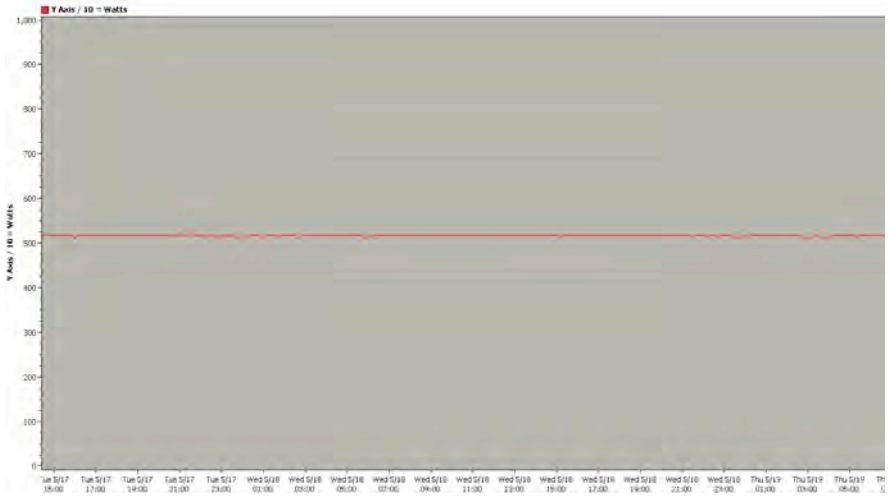

Figure 9: The power level (in volt-amps) of the Glimmerglass optical switch over the period of one day.

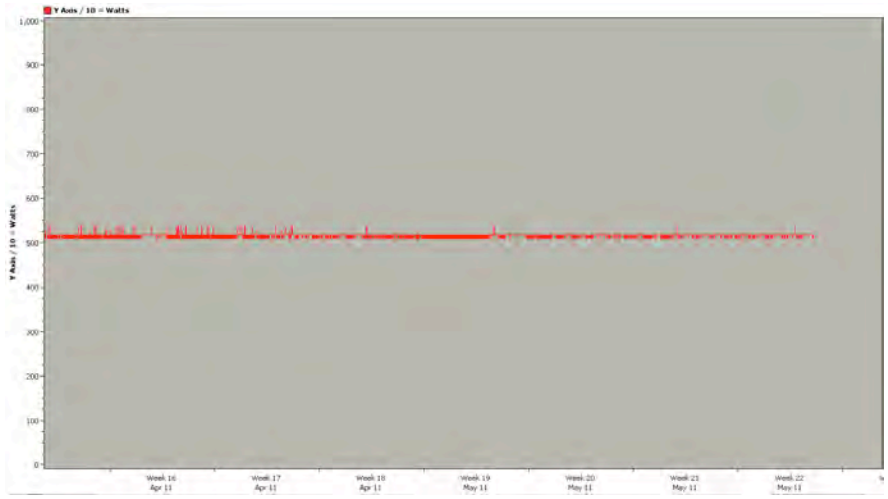

Figure 10: The power level (in volt-amps) of the Glimmerglass optical switch over the period of one month.