GLIF 12th Annual Global LambdaGrid Workshop

DEMONSTRATIONS, Electronic Visualization Laboratory (EVL)
October 11, 2012

Demonstration Descriptions

CA01: New Techniques for Uncompressed 4K Video Transmission – Ciena Corporation

CZ01: Use of High-Speed Networking in Film Restoration Process – CESNET

CZ02: Remote Collaboration Over 8K Visualization Using UltraGrid and SAGE – CESNET

NE01: Collaborative Analysis of Climate Models using Remote Scientific Visualization – SARA

NE02: Multipathing with MPTCP and OpenFlow – SARA and CERN

NO01: Plugfest 2 Windy City: NSI v2.0 First Look + GLIF Automated GOLE Pilot Project – NORDUnet

PO01: High-Resolution Advanced Visualization over Long-Distance Optical Networks – PSNC

US01: NexCAVE to CAVE2 – Calit2/UCSD

US02: The Open Science Data Cloud – University of Chicago / Computation Institute

US03: Highly Efficient LHC Data Transfer over WAN among 40Gbps Disk Servers using FDT – Caltech

US04: Slice Around the World – Northwestern University/iCAIR

US05: 100Gbps High-End Computer Networking for Petascale Science – Northwestern Univ/iCAIR

US06: InstaGENI Distributed Dynamically Programmable Environment – Northwestern Univ/iCAIR

US07: International OpenFlow Experimental Network Testbed – Northwestern University/iCAIR

US08: High Performance Digital Media Network – Northwestern University/iCAIR

US09: EVL CAVE2 Demonstrations

GLIF 2012 gratefully acknowledges the support of the following sponsors:

CA01: New Techniques for Uncompressed 4K Video Transmission

Contact: Rodney G. Wilson, Ciena Corporation, rwilson@ciena.com

Collaborators: Ciena, CANARIE, *Canada;* Northwestern/iCAIR, University of Illinois at Chicago/EVL, *U.S.*

A temporary highly reliable 100Gbps optical network connection from Ciena Research Labs in Ottawa, via Canada's CANARIE network, through StarLight to EVL, enables a demonstration of 4K uncompressed video streaming that originates from Ciena Labs and appears on a 4K video monitor at EVL. Extreme high-definition, multimedia transmission enables the broadcast industry to produce content such as films and TV series and is being tested for use in a research environment. The demonstration features Ciena's video MOTR multimedia networking card on the 5200 Advanced Services Platform.

CZ01: Use of High-Speed Networking in Film Restoration Process

Contact: Michal Krsek, CESNET, michal.krsek@cesnet.cz

Collaborators: CESNET, Czech Technical University, Masaryk University, *Czech Republic;* CineGrid, *International Research & Education Community*

www.cinegrid.org

Film restoration is an important process in cultural heritage preservation. Currently, all films are stored in a printed format that requires special storage conditions (temperature, humidity). Film materials degrade even if properly stored and analog replication does not preserve quality. In the film restoration process, film gets scanned into digital form and all restoration work is done digitally using digital post-production methods. The work is done by system operators, but results need to be critically screened by the original film authors to get the original, intended result. Authors (particularly of older films) are hard to reach (they are busy or elderly) and sometimes they do not want to travel. Our demonstration shows the principle of critical viewing of restored material and interactive work between an operator (in Prague) and an original author (in Chicago).

CZ02: Remote Collaboration Over 8K Visualization Using UltraGrid and SAGE

Contact: Petr Holub, CESNET, petr.holub@cesnet.cz

Collaborators: CESNET, Masaryk University, *Czech Republic;* Northwestern/iCAIR, University of Illinois at Chicago/EVL, *U.S.*

http://ultragrid.sitola.cz
http://www.sagecommons.org

The project demonstrates multi-site interaction of users with low-latency streaming of an 8K pre-rendered planetarium visualization to SAGE screens located at two geographically separated locations. The streaming process employs on-the-fly compression and decompression (JPEG and DXT) of the content using high-performance GPU-accelerated processing in the UltraGrid platform which is then rendered to SAGE screens, to demonstrate capabilities of commodity hardware on high-speed networks.

NE01: Collaborative Analysis of Climate Models using Remote Scientific Visualization

Contact: Paul Wielinga, SARA, wielinga@sara.nl

Collaborators: SARA, IMAU Utrecht University, Netherlands eScience Center (NLeSC), SURFnet, *The Netherlands;* Los Alamos National Laboratory, *U.S.; South Africa*

Climate modeling is becoming a multi-disciplinary, internationally spread science. A remote collaboration and visualization environment – as demonstrated at GLIF – allows leading climatologists worldwide to discuss breakthrough simulation results as soon as they become available, which is of seminal importance for enhancing the speed and quality of scientific discovery in the field. During the demo, climatologists located on three continents (North America, Africa, Europe), will discuss results of simulations and modeling parameters by using a newly developed visualization tool for climate data, combined with existing video-conferencing technologies and SAGE middleware to share the visual data.

NE02: Multipathing with MPTCP and OpenFlow

Contacts: Ronald van der Pol, SARA, rvdp@sara.nl; Artur Barczyk, Caltech, Artur.Barczyk@cern.ch

Collaborators: SARA, *The Netherlands;* Caltech, *U.S.;* CERN, *Switzerland*

http://nrg.sara.nl/GLIF-2012

Datasets in e-science are increasing rapidly in size. To transfer these extreme datasets, we need to make efficient use of all available network capacity. This means using multiple paths where available. In this demonstration, Multipath TCP (MPTCP, RFC 6182) is used in combination with an OpenFlow-based network. Multipath TCP has the distinct advantage over other multipath solutions in that it distributes the load based on the available bandwidth of each path, rather than evenly distributing the load across each path. This makes it particularly suitable when the links have different capacities. The resulting multipath flows will be monitored during the demo using traffic statistics and Ethernet OAM continuity checks. We use large data transfers (e.g., migrating VMs and copying datasets) and show the effect of interrupting certain links on the overall throughput. The actual load on the various links will be shown on a website.

NO01: Plugfest 2 Windy City: NSI v2.0 First Look + GLIF Automated GOLE Pilot Project

Contacts: Jerry Sobieski, NORDUnet, jerry@nordu.net

GLIF Automated GOLE Pilot Project Collaborators: NORDUnet, *Denmark;* GÉANT, *European R&E Community;* Poznan Supercomputing and Networking Center, PIONIER, *Poland;* CESNET, *Czech Republic;* NetherLight, SURFnet, University of Amsterdam, *The Netherlands;* AIST, KDDI Labs, JGN-X, *Japan;* KISTI, KRLight, *Korea;* GLORIAD, StarLight, USLHCnet, WIX, MAX, ACE, ESnet, *U.S.*

http://www.glif.is/working-groups/tech/GLIF%20NSI%20Implementation%20Task%20Force%20v2.pdf

This demonstration provides an alpha test of the newly minted OGF NSI version 2.0 protocols utilizing an expanded Global Automated GOLE testbed facility. The Network Services Interface (NSI) is a standard

for global automated provisioning of performance-guaranteed network services. Such automated provisioning of connection services forms the basis for future network virtualization, globally distributed and highly dynamic networked applications, and is an integral part of the emerging Open Lightpath Exchange architectures for projects such as LHCONE, radio astronomy, and advanced global video/telepresence services.

As a reprise to the successful NSI "Plugfest Rio" demonstrations at GLIF 2011 in Rio de Janeiro, the Plugfest Windy City demonstration shows new features of NSI v2.0 as well as interoperability among numerous NSI implementations. The demonstrations include a functional interoperability matrix that the NSI developers exercise. And, some of the subtle but very important new features of NSI v2 are highlighted and explained with informal presentations and demos – features such as enhanced topology management, sophisticated path planning, ease of use, versioning, authorization, and many more.

PO01: High-Resolution Advanced Visualization over Long-Distance Optical Networks

Contact: Michał Kierzynka, Poznan Supercomputing and Networking Center (PSNC), michal.kierzynka@man.poznan.pl

Collaborators: PSNC, PIONIER, *Poland;* University of Illinois at Chicago/EVL, *U.S.;* NSI-enabled infrastructure in PIONIER and other NRENs along a path from Poland, to the GLIF venue in Chicago

This demonstration shows users interacting with a visualization application using dedicated hardware. Rendering is performed remotely on hardware running on the PSNC premises, while the display is at the GLIF venue. One can interact with displayed pictures – steering signals are sent from EVL through StarLight to Poznan using the same dedicated channel. The demonstration shows the difference between using a dedicated circuit setup and a regular IP connection without any QoS features available. The demonstration also shows how easily a dedicated circuit can be created on demand (by the application) in a short timeframe using an NSI-enabled global infrastructure (i.e., the GLIF infrastructure, NSI-enabled global testbed, and the PIONIER network). The network is dynamically configured when needed, and the circuit is commissioned and decommissioned at various times during the demo, showing its influence on the interactive visualization.

US01: NexCAVE to CAVE2™

Contact: Tom DeFanti, Calit2/UCSD, tdefanti@ucsd.edu

Collaborators: Calit2/UCSD, University of Illinois at Chicago/EVL, *U.S.*

http://www.calit2.net/newsroom/rss.php?id=1862
http://www.evl.uic.edu/core.php?mod=4&type=4&indi=824

Advanced networking is used to share and collaborate with sets of quarter-gigapixel images in virtual reality. Also, new results of the virtual-reality CAVEcam work are shown, highlighting 360-degree photographic panoramas and LiDAR scans interactively shared in the CAVE2 at EVL and the NexCAVE and Vroom at Calit2.

US02: The Open Science Data Cloud

Contact: Robert Grossman, University of Chicago/Computation Institute, robert.grossman@uchicago.edu

Collaborators: University of Chicago/Computation Institute, Johns Hopkins University/Department of Astronomy, Northwestern University/iCAIR, NASA Goddard Space Flight Center, *U.S.*

www.opencloudconsortium.org

The Open Science Data Cloud (OSDC) is a distributed science cloud managed and operated by the not-for-profit Open Cloud Consortium that provides cloud-based infrastructure to scientific collaborations interested in managing, analyzing, sharing and archiving medium to large scientific datasets. For GLIF, we are demonstrating OSDC applications from the biological and earth sciences.

US03: Highly Efficient LHC Data Transfer over WAN among 40Gbps Disk Servers using FDT

Contact: Harvey Newman, Caltech, newman@hep.caltech.edu; Artur Barcyzk, Caltech, artur.barczyk@cern.ch; Azher Mughal, Caltech, azher@hep.caltech.edu

Collaborators: Caltech, *U.S.;* Ciena Corporation, *Canada;* SURFnet, *The Netherlands*

Caltech's High Energy Physics (HEP) team utilizes a 40Gbps testbed between two data centers to demonstrate highly efficient transfers of LHC data between these sites. Each end site hosts storage servers with 40GE Ethernet NICs and PCIe Gen3 motherboards. The demonstration uses a 40Gbps path between Geneva and Amsterdam.

After a successful SC'11 demonstration, we continue our work on data-transfer systems and have seen very good, consistent results – particularly in terms of single-node performance, transfer stability and system robustness. According to researchers, these efforts are helping to establish new ways to transport the increasingly large quantities of data that traverse continents and oceans via global networks of optical fibers. These new methods are needed for the next generation of network technology, which allows transfer rates of true 40 and 100Gbps.

US04: Slice Around the World: Highly Distributed Computational Based On Global Clouds Closely Integrated With Highly Programmable Networks

Contact: Joe Mambretti, Northwestern University/iCAIR, j-mambretti@northwestern.edu

Collaborators: Renaissance Computing Institute (RENCI), Princeton University, Duke University, Northwestern University/iCAIR, Hewlett Packard Research Labs, University of Utah, *U.S.;* NICTA, *Australia;* ANSP, Centro de Pesquisa e Desenvolvomento de Telabras, Rede Nacional de Ensino e Pesquisa (RNP); *Brazil;* Communications Research Center, *Canada;* Chinese Academy of Sciences/Computer Network Information Center, *China;* Technical University of Kaiserslautern/G-Lab, *Germany;* University of Tokyo, National Institute of Information and Communications Technology, *Japan;* Electronic and Telecommunications Research Institute, Korea Institute of Science and Technology Information (KISTI), *Korea;* University of Amsterdam, SARA, *The Netherlands;* Applied Research

Center for Computer Network at Skolkovo, *Russia;* National Kao Hsiung University of Applied Science, Kaohsiung, NCHC, National Cheng-Kung University, *Taiwan;* University of Essex, *U.K.*

www.geni.net

The "Slice Around the World" initiative demonstrates the powerful potential of designing and implementing worldwide environments of global computational clouds closely integrated with highly programmable networks. The initiative was established by network research centers and research labs that participate in next-generation network efforts to develop large-scale distributed experimental network environments – such as the U.S. National Science Foundation's Global Environment for Network Innovations (GENI), the EU Future Internet Research Environment (FIRE), the Japanese New Generation Internet, the Korean Future Internet initiative, the German Future Internet Lab (G-Lab), and others. These environments are being developed *by researchers, for researchers.* An important goal is to have persistent global environments, developed and managed by the research community, in support of experimental research.

The GLIF demonstrations focus on applications that create ad-hoc high-performance networks integrated with cloud computing for scientific visualization, specifically related to nanotechnology, which requires remotely steered visualizations of remotely rendered, real-time distributed data. "Slice Around the World" demonstrates several techniques. Finite Difference Time Domain (FDTD) is one of most commonly used computational electrodynamics modeling techniques for many research and industry simulations, such as LSI design electro verification. Under current HPC workflow techniques, researchers submit jobs, retrieve results, visualize those results and then resubmit the job with modifications, additional information, data, etc. Today this is a tedious, manual slow process, in part because of the limitations of today's networks. The GLIF demonstration shows how using dynamically programmable networks closely integrated with computational and storage clouds provide capabilities to create interactive simulation/visualization instruments that significantly improve this traditional process.

US05: 100Gbps High-End Computer Networking for Petascale Science

Contact: Joe Mambretti, Northwestern University/iCAIR, j-mambretti@northwestern.edu

Collaborators: Northwestern University/iCAIR, NASA Goddard Space Flight Center (GSFC), MAX (Mid-Atlantic Crossroads), *U.S.*

www.hec.nasa.gov

A research consortium demonstrates an advanced optical-fiber-based network testbed, established to develop 100Gbps services and technologies for next-generation data-intensive "petascale science." The testbed has been implemented as a contiguous facility from the Washington DC area to the StarLight International/National Communications Exchange Facility in Chicago. This consortium demonstrates high-performance 100Gbps services and networking technology that are being developed for petascale science. These activities include experiments, evaluations and demonstrations of high-performance 100Gbps IPv4/IPv6 disk-to-disk file copying, memory-to-memory transfer, ultra-large data transfer, and access tools and methods across LANs and WANs. These activities demonstrate innovations in network architecture, services, tools, methods and technologies required to effectively use 100Gbps channels for large-scale, data-intensive science across multiple LAN and WAN environments, including using devices from multiple vendors that integrate both available components and emerging alpha/beta components that have not yet been commercialized.

US06: InstaGENI Distributed Dynamically Programmable Environment

Contact: Joe Mambretti, Northwestern University/iCAIR, j-mambretti@northwestern.edu

Collaborators: Northwestern University/iCAIR, Hewlett-Packard Labs, Open Networking Institute, Princeton University, University of Utah, *U.S.;* Technical University of Kaiserslautern/ G-Lab, *Germany;* NCHC, *Taiwan;* SARA, *The Netherlands;* University of Tokyo, *Japan*

www.geni.net

InstaGENI is a distributed cloud based on programmable networks designed for GENI (Global Environment for Network Innovations), a U.S. National Science Foundation project funded through the GENI Project Office. GENI is creating a national experimental network research environment with international extensions, in part through the International GENI (iGENI) initiative. The InstaGENI architecture closely integrates light weight cluster design with Software Defined Networking, Hardware-as-a-Service, and Containers-as-a-Service, remote monitoring and management, and high-performance inter-site networking. The initial InstaGENI deployment instantiates multiple sites across the U.S., interconnected through a specialized GENI backbone network deployed over national, regional and campus research and education networks, with international network extensions to sites across the world.

US07: International OpenFlow Experimental Network Testbed

Contact: Joe Mambretti, Northwestern University/iCAIR, j-mambretti@northwestern.edu

Collaborators: Northwestern University/iCAIR, *U.S.;* NCHC, National Kaohsiung University of Applied Sciences (KUAS), National Cheng-Kung University (NCKU), *Taiwan;* SARA, SURFnet, *The Netherlands;* Caltech, *U.S.*

www.icair.org

The International Center for Advanced Internet Research (iCAIR) at Northwestern University and its international research collaborators demonstrate new how large-scale international OpenFlow-based techniques can create innovative, ad hoc, specialized networks, including those that for extremely high-volume, high-resolution, scientific visualization for large-scale data at multiple remote sites.

US08: High Performance Digital Media Network

Contact: Joe Mambretti, Northwestern University/iCAIR, j-mambretti@northwestern.edu

Collaborators: Northwestern University/iCAIR, StarLight, Calit2/UCSD, University of Illinois at Chicago/EVL, University of Illinois at Urbana-Champaign/National Center for Supercomputing Applications (NCSA), *U.S.;* Communication Research Centre Canada, Inocybe, Synchromedia, Ciena, CANARIE, *Canada;* KISTI, *Korea;* SARA, University of Amsterdam, NetherLight, SURFnet, *The Netherlands;* University of Essex, *U.K.;* Poznan Supercomputing and Networking Center (PSNC), *Poland;* I2CAT, *Spain;* National Center for High-Performance Computing (NCHC), *Taiwan;* Technische Universität Carolo-Wilhelmina zu Braunschweig/Institute of Computer and Network Engineering, The

Braunschweig University of Art, *Germany;* Korea Institute of Science and Technology Information (KISTI), *Korea*

www.hpdmnet.net

The High Performance Digital Media Network (HPDMnet) and its related activities comprise an experimental network research initiative that is designing, developing and implementing the world's first international high-performance service specifically for high-quality, large-scale digital media, including support for extremely high-volume media streams. This initiative represents a complete departure from existing services, which are based on legacy services and technologies that cannot meet emerging needs for high-quality, reliable services. By using advanced concepts, architecture, and technology, HPDMnet is providing a foundation for future digital media services, as well as for other data-intensive applications. HPDMnet was established as a cooperative partnership by several major network research centers. Because this initiative is a research project, it is continually progressing through ongoing iterations. Major milestones are indicated in part by large-scale demonstrations at national and international forums.

US09: EVL CAVE2™ Demonstrations

Contact: Maxine Brown, University of Illinois at Chicago/Electronic Visualization Laboratory (EVL), *U.S.*

Collaborators: University of Illinois at Chicago/EVL, Argonne National Laboratory, *U.S.*

http://www.evl.uic.edu/core.php?mod=4&type=4&indi=824

CAVE2™, the next-generation large-scale virtual-reality environment, is a hybrid system that combines the benefits of both scalable-resolution display walls and virtual-reality systems to create a seamless 2D/3D environment that supports both information-rich analysis and virtual-reality simulation exploration at a resolution matching human visual acuity.

CAVE2™ is approximately 24 feet in diameter and 8 feet tall, and consists of 72 near-seamless passive-stereo off-axis-optimized 3D LCD panels, a 36-node high-performance computer cluster, a 20-speaker surround audio system, a 10-camera optical tracking system and a 100Gbps connection to the outside world. CAVE2 provides users with a 320-degree panoramic environment for displaying information at 37 Megapixels in 3D or 74 Megapixels in 2D with a horizontal visual acuity of 20/20 – almost 10 times the 3D resolution of the original CAVE™.