iGENI Quarterly Report

GENI Project #1719

For the Period July 1, 2010 through September 30, 2010

Joe Mambretti, International Center for Advanced Internet Research (iCAIR), Northwestern University, j-mambretti@northwestern.edu
Maxine Brown, Electronic Visualization Laboratory, University of Illinois at Chicago, maxine@uic.edu
Thomas A. DeFanti, California Institute for Telecommunications and Information Technology (Calit2), University of California, San Diego, tdefanti@ucsd.edu

1. Major Accomplishments

This project is defining, designing, and implementing the International GENI (iGENI), a distributed network research infrastructure, which is being integrated with current and emerging GENI resources and which will be operated for GENI researchers, who will conduct experiments that involve multiple aggregates (at multiple sites). The iGENI infrastructure is being defined in collaboration with the GPO and other GENI projects to expand the variety of controllable transport services available to GENI researchers, to add additional capabilities to that infrastructure, and to make GENI available to more research communities. During Q4, iGENI participated in the design and implementation of aggregate interconnections among multiple Cluster-D participant sites. iGENI also enhanced the capabilities of two iGENI international testbeds that it established between the StarLight/iCAIR and several universities in Taiwan. In addition, iGENI and RENCI implemented an iGENI testbed international testbed among the BEN testbed in North Carolina, the StarLight international communications exchange facility and several universities in South Korea. iGENI participated in several demonstrations at GEC 8, including the GENICloud demonstration, TransCloud, and led a large scale demonstration of edge and core dynamic testbed provisioning. Also, planning was undertaken for demonstrations at GEC 9 in Washington DC. The iGENI community has had GENI and iGENI planning discussions with networking research groups from Australia, Brazil, Canada, China, Egypt, Germany, India, Japan, Korea, Singapore, Taiwan, Spain, New Zealand, Sweden, Switzerland, Poland, Saudi Arabia, the UK, and others.

Current Capabilities

Preliminary infrastructure architecture and design concepts were developed for the iGENI US infrastructure and presented and discussed at various forums, including at GEC forums, most recently at GEC 8 at the conference venue in La Jolla, July 20-22, 2010 and the Quilt GENI Workshop, at the same location, July 22-23, 2010. The ORCA clearinghouse that has been implemented at iCAIR was extended to additional resources. This core facility, which has been implemented within one of the iCAIR network research labs, is connected by dedicated optical fiber to high performance switches at the StarLight International/National Exchange Facility. Also, private dedicated fiber has been implemented between iCAIR/StarLight to the NLR core node at 111 North Canal in Chicago. Planning was undertaken for a second fiber pair. A preliminary design that was originally developed for a network to interconnect all GENI Cluster-D sites – a GENI Cluster D Network (GCDnet) was completed and implemented at several sites, including RENCI/BEN in North Carolina,, Northwestern University, the StarLight facility, BBN Research Lab, the University of Massachusetts at Amherst (through the NOX in Boston), Wayne State,

and Ohio State. The iGENI community has established planning processes directed at providing additional connections from existing resources at the StarLight national and international communications exchange with current GENI backbone transport resources, with an initial path based on NLR Layer 2/Ethernet VLANs) using 10 Gbps NLR FrameNet and C-Wave lightpaths. Preliminary concepts and options are being explored for international path implementations, to Canada, Asia, and Europe. These activities have been assisted with the recent announcement of the NSF award for the TransLight/StarLight proposal under the International Research Network Connections (IRNC) program. This program is planning to provide support for iGENI international activities.

1. Milestones Achieved

iGENI milestones are described on the GENI wiki.

The initial design of the iGENI infrastructure has been developed and reviewed, and the initial prototype was implemented and then extended. The majority of current activities relate to planning for extensions nationally and internationally. iGENI has been integrated as an aggregate with the ORCA control framework in Cluster D, with persistent and dynamic L1/L2 paths among multiple Cluster-D sites using GCDnet. This first implementation served as a demonstration model for establishing similar connections to other Cluster D sites. The ORCA clearinghouse at RENCI is being used for GCDnet provisioning. The ORCA GENI Cluster D implementation includes one Broker, multiple Service Managers, and multiple Site/Domain Authorities. iGENI has been integrated with ORCA, through an initial lab implementation at iCAIR.

2. Description of Work Performed During 4nd Quarter

2.a. Activities and Findings

Q4 activities were focused on extending prototypes based on core infrastructure architectural concepts and implementation plans. The iGENI initiative has developed processes and procedures for integrating core resources with an ORCA based control plane framework, including L2/L1 paths. iGENI also moved forward with initial prototype implementations. These implementations have allowed for resources to selectively advertise their external interfaces, including vLANs, enabling interconnects among dedicated GENI resources, initially among Cluster-D sites (to be followed later, among resources provided by regional networks, nationals R&E networks, internationals R&E networks, non-profit R&D organizations, corporate R&D organizations, and other sites, facilities and institutions). Investigations are also being conducted to determine options for supporting multiple types of L1/L2 paths, including vLANs, tunneling services, e2e lightpaths, standard optical L2 framing, and others. Plans are also being developed to enable core L1/L2 resources to be identified using standard L1/l2 resource addressing while experimental L1/L2 core resources will be identified by using a method that allows for a level of abstraction that will be integrated into an XML-based resource description language. The ORCA control framework currently is being used for the network resource allocation. Within the iGENI infrastructure, calls are mapped onto an addressable L1/L2 path infrastructure, using static, semi-dynamic and dynamic infrastructures. Edge resources use a private addressing scheme. (The Cluster-D lead organizations sponsored a network resource specification workshop at GEC8.) This design anticipates that the core resource infrastructure framework and the experimental research infrastructure will be operated by

distributed operational NOC processes. Core infrastructure will be addressed by a management plane based on common L3 secure channels in addition to the control plane framework.

The ORCA control framework has been integrated with the iGENI infrastructure. iGENI Consortium has implemented the Open Resource Control Architecture (ORCA) control framework at the StarLight international exchange facility. An instantiation of ORCA was installed on a server in one of the iCAIR research labs and it has been integrated with facilities equipment. A second implementation integrates iCAIR and StarLight facilities with the ORCA clearinghouse at RENCI. This implementation is integrated with switches and servers at a core node in the StarLight facility. iGENI is now integrated as an aggregate with that implementation of the ORCA control framework in Cluster D, with L1/L2 paths among StarLight, RENCI/BEN, and other Cluster-D sites. This initial implementation is serving as a model for establishing connections to other sites. Through ORCA, available resources in iGENI can be discovered; services can be setup and managed; and, individual traffic streams will be controlled and managed. This project has implemented interfaces to ORCA that allow dynamic control of network services involving iGENI, associated transport resources and GENI aggregates. It is possible to setup services using prepackaged or customized configurations and topologies.

The current prototype was demonstrated at the GEC 8 workshop. The iGENI dynamic network provisioning demonstrations showcased capabilities for large scale (national and international) multiple domain dynamic provisioning, including L1/L2 path involving multiple sites, using specialized signaling and implementation techniques. In partnership with RENCI (Renaissance Computing Institute), Duke University, the University of Massachusetts, Wayne State, Ohio State, and other D-Cluster participants iGENI supported a demonstration of dynamic vLAN provisioning at GEC 8, based on dynamic and static L1/l2 paths among multiple Cluster-D sites. (Ref : Figure 1 below).

Another GEC 8 demonstration showcased GENICloud, TransCloud, which illustrated the potential for creating a highly scalable distributed computing environment integrated with dynamic network provisioning (See Figures 2, 3, and 4 below). The TransCloud demonstrations show the potential for creating powerful new capabilities and services based on distributed environments by integrating multiple clouds (established at highly distributed sites: GENICloud -- HP OpenCirrus, UCSD, and Northwestern) with the dynamic network provisioning envisioned by GENI. The TransCloud demonstration shows that separate infrastructures do not have to be implemented for different types of delivery platforms. The single TransCloud environment can provide streams available to mobile devices, computers, tablets, tile displays and any other edge device. This demonstration was a prototype for one that is being prepared for a GEC 9 Alpha demonstration. To demonstrate the utility of these capabilities, application was selected (transcoding) to emphasize the advantages of using the prototype environment in contrast to legacy approaches, which require different infrastructure for each edge platform (e.g., mobiles, tablets, computers, tile displays etc.) This multi-organization TransCloud demonstration showcased a capability for using dynamic large scale cloud and network infrastructure for highly distributed specialized capabilities among multiple sites connected by the iGENI network, including digital media transcoding and streaming to multiple edge platforms, supported by scaleable cloud computing and network provisioning. Three clouds were interconnected via GENI, HP OpenCirrus, iCAIR's OpenCloud, and a cloud at UCSD and used to stream digital media from repositories and live streams (organized by iGENI). Figure 4 show an image of live streaming from the StarLight facility to a tile display at GEC 8.

2.b. Project Participant Activity

The primary activities in Q4 have been focused on a) designing and implementing GENI infrastructure, b) designing and conducting demonstrations for GEC 8, c) R&D meetings with GENI Cluster D partners, national research networking organizations, and international research network organizations, as well as conference calls and meetings at GEC 8 with the ORCA framework developers, d) planning for future infrastructure implementations and e) planning for future demonstrations

2.c. Publications and Presentations

The iGENI project was presented at a GLORIAD workshop and at a Future Internet conference in Seoul, South Korea in June. The iGENI project was also presented at GEC 8, at the Quilt GENI Workshop, and at meetings with several groups of international visitors at iCAIR, including as part of a GLORIAD project meeting at iCAIR, which included representatives from the new R&E international communications exchange being implemented in Cairo, Egypt. Another iGENI presentation was given at an international meeting organized by iCAIR for representatives from KISTI, a research center in Daejeon, Korea. The current activities of the iGENI project were presented at the quarterly meeting of the Executive Committee of the Metropolitan Research and Education Network (MREN) at Argonne National Laboratory in August 2010, and are scheduled to be presented at the GLIF 10th Annual Global LambdaGrid Workshop, 12-14 October 2010 in Geneva, Switzerland an event hosted by CERN. GLIF (Global Lambda Integrated Facility) participants include National Research and Education Networks (NRENs), consortia and institutions that are creating a globally distributed infrastructure testbed facility based on optical-fiber lightpaths and are involved in multiple, innovative communication services and technology projects. An overview of GENI and iGENI was presented at the ON*VECTOR Photonics Workshop at UCSD in early 2010. The iGENI project was presented at a GLORIAD workshop and at a Future Internet conference in Seoul, South Korea in June 2010. The iGENI project was also presented at GEC 8, at the Quilt GENI Workshop, and at meetings with several groups of international visitors at iCAIR, including as part of a GLORIAD project meeting at iCAIR, which included representatives from the new R&E international communications exchange being implemented in Cairo, Egypt. Another iGENI presentation was given at an international meeting organized by iCAIR for representatives from KISTI, a research center in Daejeon, Korea. The iGENI project was presented at the quarterly meeting of the Executive Committee of the Metropolitan Research and Education Network (MREN) at Argonne National Laboratory in August 2010. Another presentation was scheduled for the GLIF 10th Annual Global LambdaGrid Workshop, 12-14 October 2010 in Geneva, Switzerland an event hosted by CERN. GLIF (Global Lambda Integrated Facility) participants include National Research and Education Networks (NRENs), consortia and institutions that are creating a globally distributed infrastructure testbed facility based on optical-fiber lightpaths and that are involved in multiple, innovative communication services and technology projects. The iGENI community participated in planning meetings for large scale dynamic network provisioning demonstrations at the SC10 international supercomputing conference in November 2010 in New Orleans, and at the international 2010 Global LambdaGrid (GLIF) Workshop at CERN, near Geneva, Switzerland.

2.d. Outreach Activities

The iGENI community has had GENI and iGENI planning discussions with networking research groups from Australia, Brazil, Canada, China, Egypt, Germany, India, Japan, Korea, Singapore, Taiwan, Spain,

New Zealand, Sweden, Switzerland, Poland, Saudi Arabia, the UK, and others. iGENI was discussed at two Science Cloud Workshops in Chicago in July 2010 and at the Open Gid Forum in Chicago in July 2010.

Figure 1:

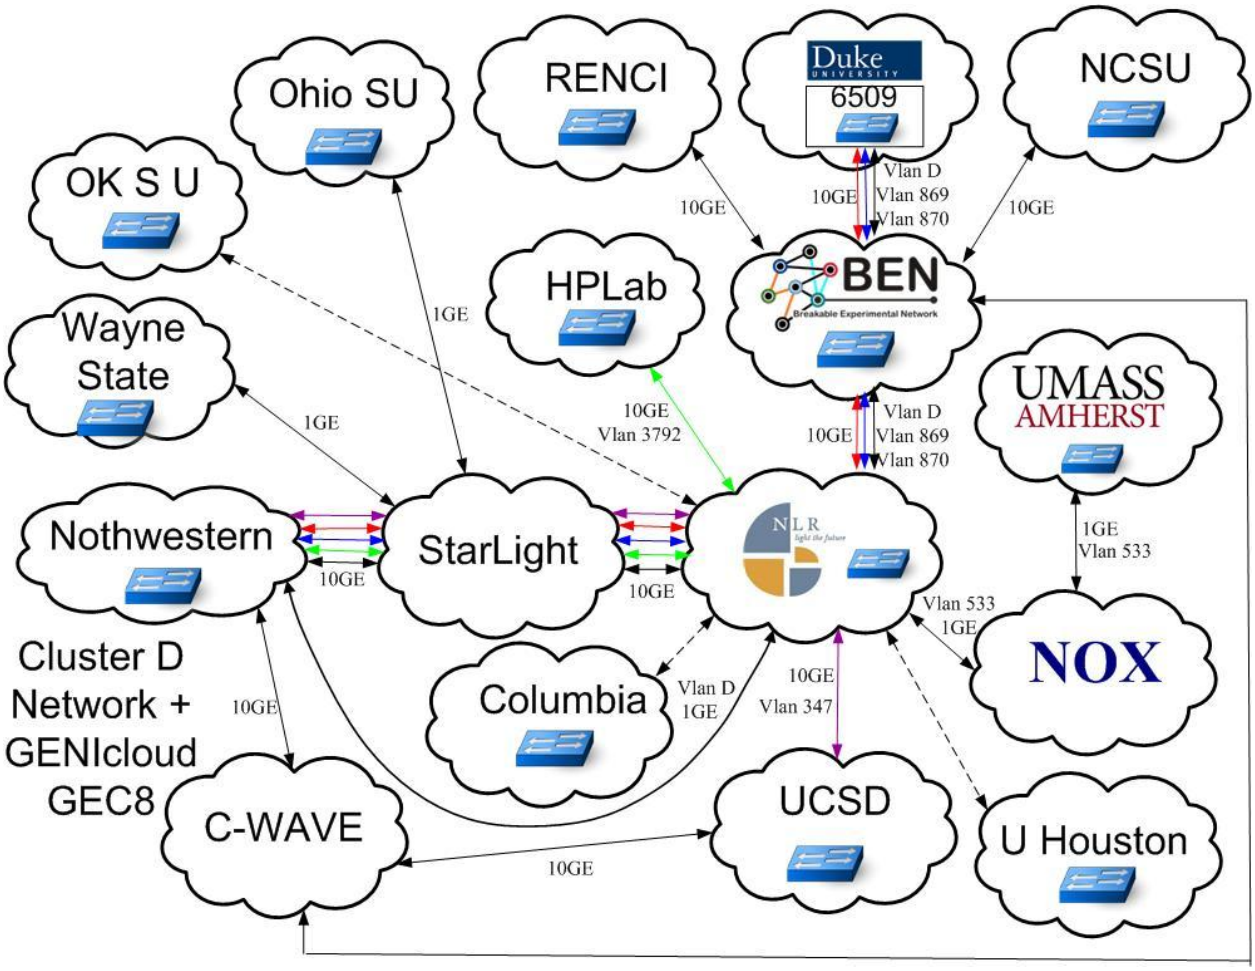

Figure 2:

GENICloud

Andy Bavier, Jessica Blaine, Daniel Catrein, Jim Chen, Yvonne Coady, James Kempf, Christian Lottermann, Joe Mambretti, Rick McGeer, Alex Snoeren, Johannes Willig, Marco Yuen, Alvin AuYoung

GENICloud Architecture

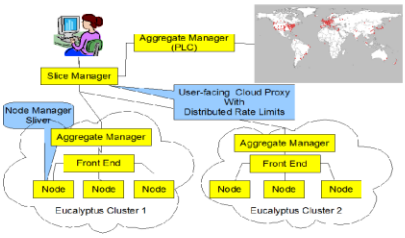

Status and Accomplishments

- Integrated Eucalyptus and GENI
  - Eucalyptus release supporting the SFA
- Unified API interface to Eucalyptus and SFA
  - GENI tools now work on (our) Clouds
- RSpec for Eucalyptus
- Resource Discovery
  - Kernel/Disk images
  - Instance Types
- Jobs instantiated on multiple Clouds
- Distributed Rate Limiting over multiple Clouds

Roadmap

- Full integration of PlanetLab and Eucalyptus
- Complete "Slice Manager" – user-facing multiple-SFA facility controller
- Integrate DRL into cross-facility slices
- New GUI for PlanetLab, Eucalyptus, cross-SFA Facilities

GENICloud Demo

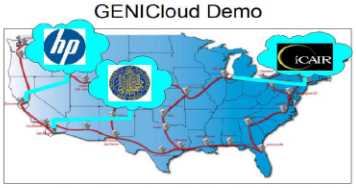

- Transcoding on Eucalyptus Clouds at UCSD, HP Labs, iCAIR (Northwestern)
- Transcoding app thanks to Ericsson Research
- For GEC-9:
  - Integration with NLR (CAVEWave)
  - Integrated Distributed Rate Limiting

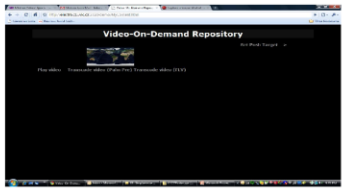

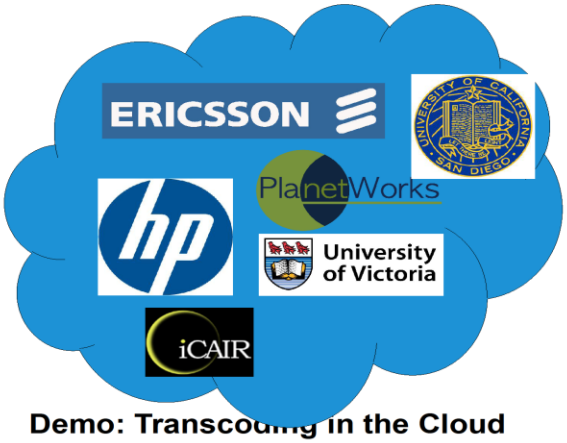

Transcoding videos at three clouds:
- **HP OpenCirrus**
- **Northwestern OpenCloud**
- **UCSD**

Demo: Transcoding in the Cloud
http://electro.cs.uvic.ca:2020/demo/MyContent.html

Figure 3:

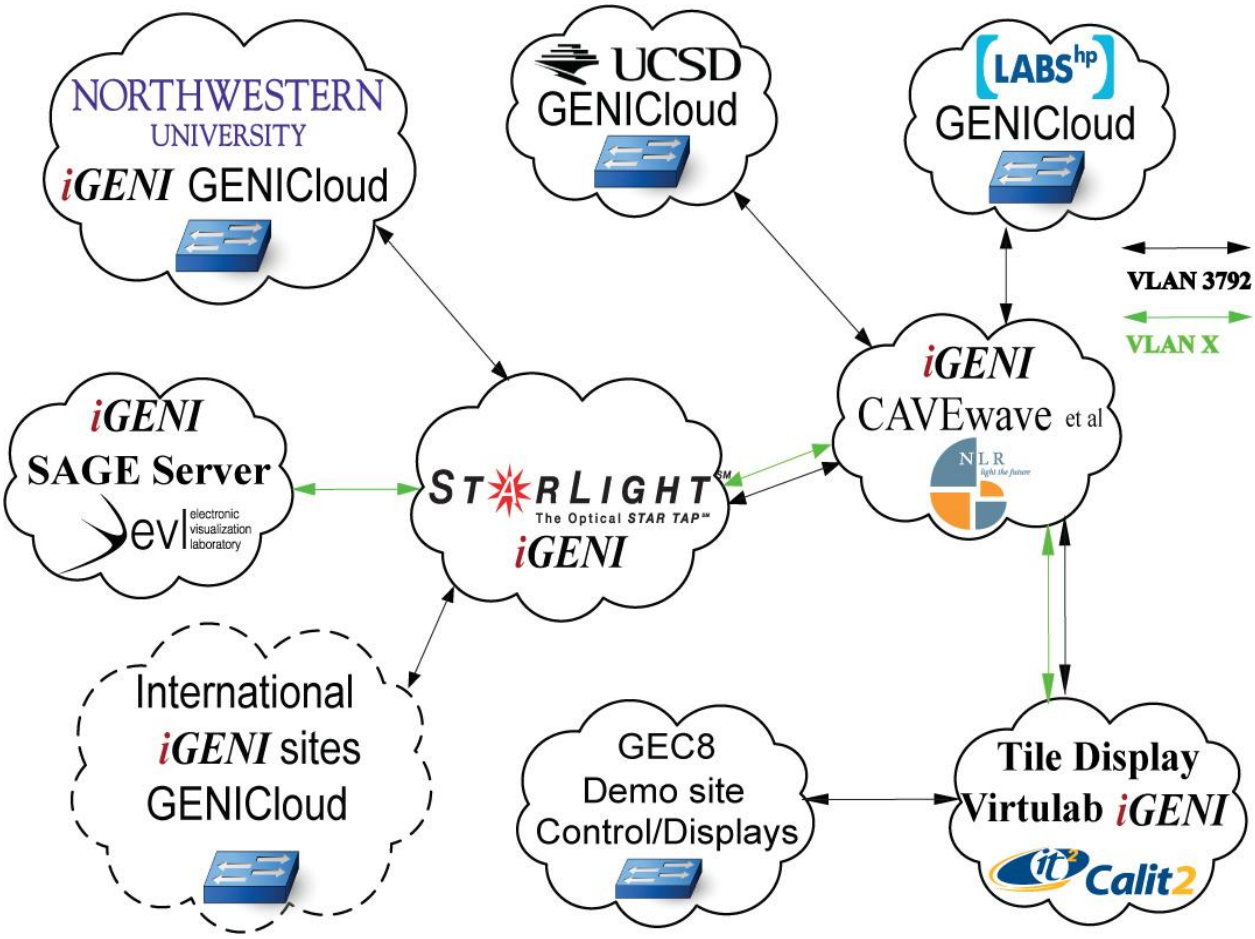
NORTHWESTERN
UNIVERSITY
iGENI GENICloud
UCSD
GENICloud
LABS hp
GENICloud
VLAN 3792
VLAN X
iGENI
CAVEwave et al
NLR
iGENI
SAGE Server
evl electronic visualization laboratory
STARLIGHT
The Optical STAR TAP
iGENI
International
iGENI sites
GENICloud
GEC8
Demo site
Control/Displays
Tile Display
Virtulab iGENI
Calit2

Figure 4: