iGENI Quarterly Report

GENI Project #1719

For the Period Oct 1, 2010 through December 31, 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 have begun conducting experiments that utilize resources based on multiple aggregates at multiple sites. The iGENI infrastructure is being defined in collaboration with the GPO and other GENI projects a) to expand the variety of resources, especially controllable transport services, available to GENI researchers, b) to add additional capabilities to that infrastructure, and c) to make GENI available to more research communities. During Q1, beginning October 1, 2010, iGENI continued to participate in the design and implementation of aggregate interconnections among multiple Cluster-D participant sites. iGENI also enhanced the capabilities of several iGENI international testbeds that have been established between StarLight/iCAIR and several universities in Taiwan, Korea, and Brazil. Earlier, 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, with GIST as a partner. iGENI planned and orchestrated several demonstrations at GEC 9 in Washington DC, including several GENICloud demonstrations, based on the TransCloud initiative, and a large scale demonstration of edge and core dynamic testbed provisioning. Also, planning was undertaken for demonstrations at GEC 10 in Puerto Rico planned for March 2011. In this quarter, the iGENI community engaged in GENI and iGENI planning additional discussions with networking research groups from multiple countries, including Australia, Brazil, Canada, China, Germany, Japan, Korea, Singapore, Taiwan, Spain, New Zealand, Poland, 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 the annual meeting of the Global Lambda Integrated Facility consortium (GLIF) in September at CERN and at GEC 10 at the conference venue in Washington DC in November 2010. The ORCA clearinghouse that has been implemented at iCAIR was again extended to include 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. During this quarter, planning was undertaken to allow for

implementation of a second fiber pair. In addition, optical testing for connections and db loss was undertaken for that 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) has now been completed and implemented at among multiple 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 discuss related to international path implementations, to Canada, Asia, South America, and Europe. These activities are being assisted through funding from a NSF award for the TransLight/StarLight proposal under the International Research Network Connections (IRNC) program. This program is providing some support for iGENI international activities.

1. Milestones Achieved

iGENI milestones are described on the GENI wiki.

Previously, the initial design of the iGENI infrastructure has been developed, reviewed, and implemented as a preliminary prototype. The majority of current activities are focused on a) planning for extensions nationally and internationally, providing resources for those extensions, providing control frameworks for those extensions, and planning for researcher use of those resources.. 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. The initial implementation was useful as a demonstration model for establishing similar connections to other Cluster D sites. The basic model developed is now being extended. All future extensions are based on that model. In addition, various tunneling techniques are being explored. 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. During this period the ORCA instantiation was upgraded several times.

2. Description of Work Performed During 4nd Quarter

2.a. Activities and Findings

Q1 activities were focused on planning for, testing, and providing for additional resources, extending prototypes based on core infrastructure architectural concepts, and evaluating the current implementations. The iGENI initiative has developed processes and procedures for integrating core resources with an ORCA based control plane framework, including L2/L1 paths. iGENI continues to plan additional resource extensions, including those related to cloud computing. The current 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. These investigations include consideration of and experimentation with other control frameworks and APIs to those frameworks. For example, iCAIR is actively supporting the GLIF Fenius API experimental and demonstration activities. Plans have been developed to enable core L1/L2 resources to be identified using standard L1/l2 resource addressing. Techniques are being investigated for developing identification methods for experimental L1/L2 core resources allowing for a level of abstraction that can be integrated into an XML-based resource description language. 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. 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.

As noted, 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 has been operational on a server in one of the iCAIR research labs for over a year, 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 9 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 9, based on dynamic and static L1/l2 paths among multiple Cluster-D sites. (Ref : Figure 1 below). This set of capabilities supported the plenary demonstration presented at GEC 9 by the University of Massachusetts at Amherst

Another GEC 9 demonstration showcased 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: TransCloud -- 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. 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 iGENI infratstructure, HP OpenCirrus, iCAIR's OpenCloud, and a cloud at UCSD and used to stream digital media from repositories and live streams (organized by iGENI). Figures 4-6 shows images from GEC 9 and SC10.

2.b. Project Participant Activity

The primary activities in Q1 have been focused on continuing to a) design and implement GENI infrastructure, b) design and conduct demonstrations for GEC 9, c) participate in 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) plan for future infrastructure implementations and collaborative activities and e) plan for future demonstrations

2.c. Publications and Presentations

The iGENI project was presented at the GLIF 10th Annual Global LambdaGrid Workshop, 12-14 October 2010 in Geneva, Switzerland an event hosted by CERN. An update of the iGENI project was presented at GEC 9. GLIF 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 project was also described during a meeting at CERN in October 2010 of the GLIF North American GOLE participants. The iGENI project was also presented at meetings with many groups of visitors, including international visitors, at iCAIR. 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 iCAIR in December 2010, The iGENI initiative organized demonstrations and presented a poster in the LAC/CAIR booth at the SC10 international supercomputing conference in November 2010 in New Orleans.

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. Plans were undertaken to organize meetings on iGENI and GENI in China and in Singapore in 2011.

Figure 1:

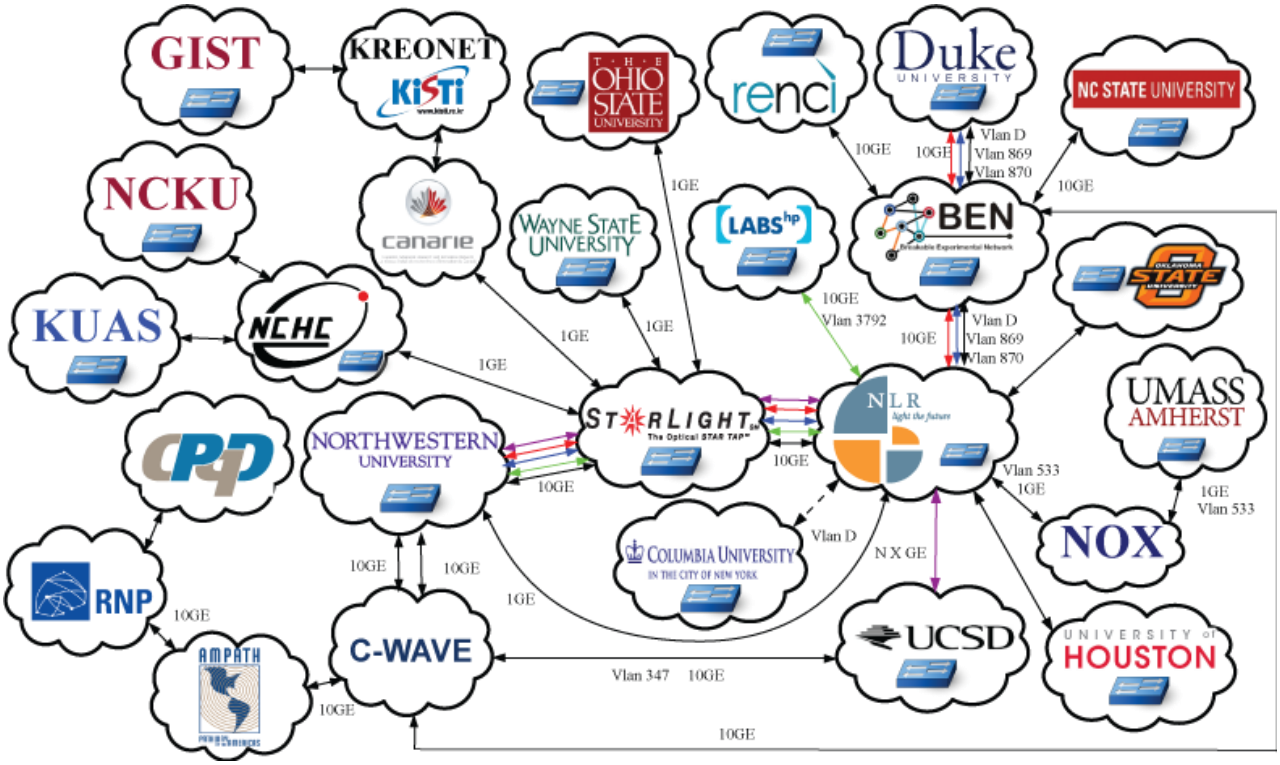
GIST
KREONET
KISTI
THE OHIO STATE UNIVERSITY
renci
Duke UNIVERSITY
NC STATE UNIVERSITY
NCKU
canarie
WAYNE STATE UNIVERSITY
LABS hp
BEN
KUAS
NCHC
STARLIGHT
The Optical STAR TAP
NLR
light the future
UMASS AMHERST
CPD
NORTHWESTERN UNIVERSITY
COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK
NOX
RNP
AMPATH
C-WAVE
UCSD
UNIVERSITY of HOUSTON
10GE
1GE
Vlan D
Vlan 869
Vlan 870
Vlan 3792
Vlan 533
N X GE
Vlan 347

Figure 2:

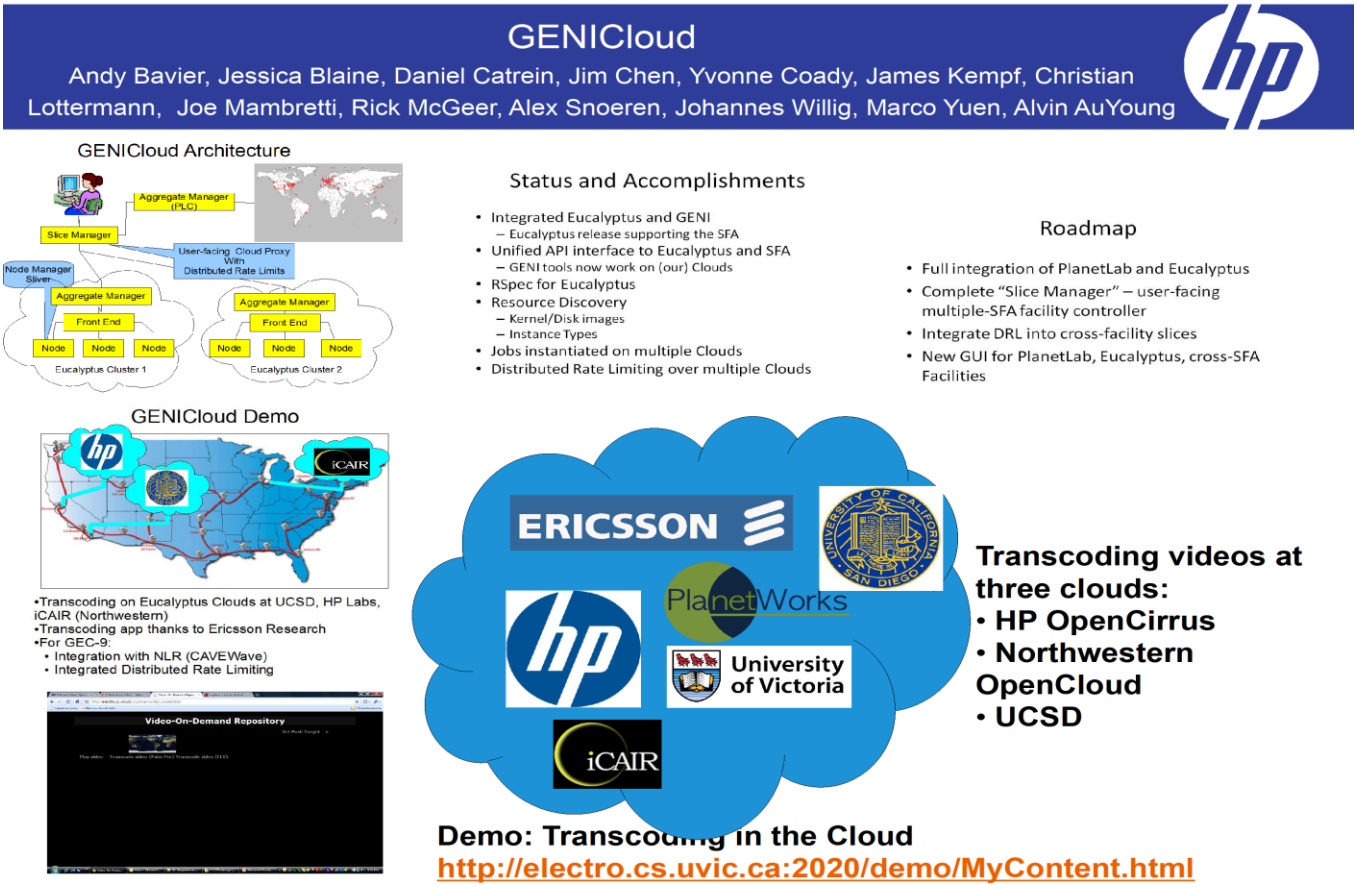
GENICloud
Andy Bavier, Jessica Blaine, Daniel Catrein, Jim Chen, Yvonne Coady, James Kempf, Christian Lottermann, Joe Mambretti, Rick McGeer, Alex Snoeren, Johannes Willig, Marco Yuen, Alvin AuYoung
GENICloud Architecture
Aggregate Manager (PLC)
Slice Manager
User-facing Cloud Proxy With Distributed Rate Limits
Node Manager Sliver
Aggregate Manager
Front End
Node
Eucalyptus Cluster 1
Eucalyptus Cluster 2
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
iCAIR
•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
Video-On-Demand Repository
ERICSSON
PlanetWorks
University of Victoria
iCAIR
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:

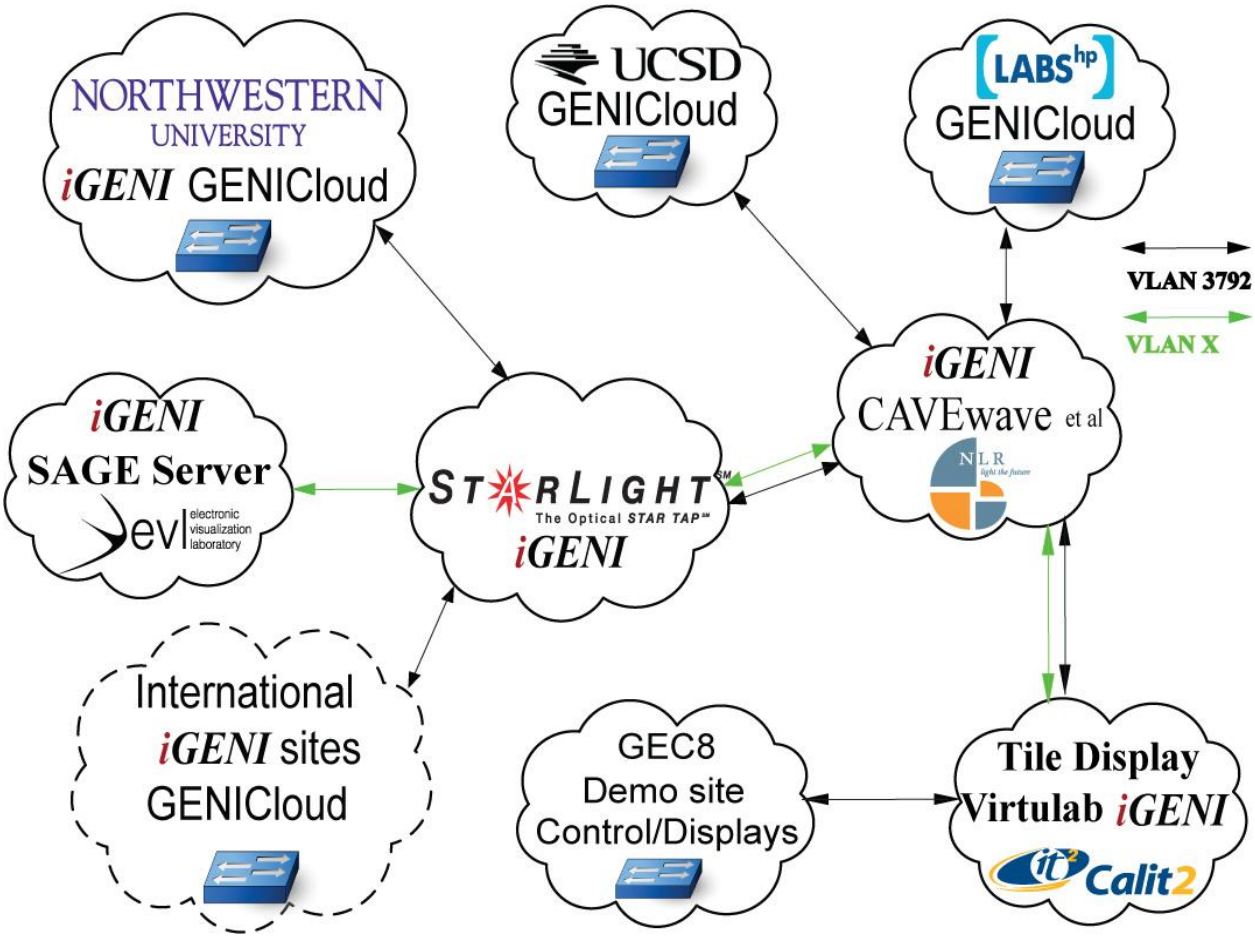
NORTHWESTERN
UNIVERSITY
iGENI GENICloud
UCSD
GENICloud
LABS hp
GENICloud
VLAN 3792
VLAN X
iGENI
CAVEwave et al
NLR
iGENI
SAGE Server
evl electronic visualization laboratory
STARLIGHT
The Optical STAR TAP
iGENI
International
iGENI sites
GENICloud
GEC8
Demo site
Control/Displays
Tile Display
Virtulab iGENI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

Figure 4:

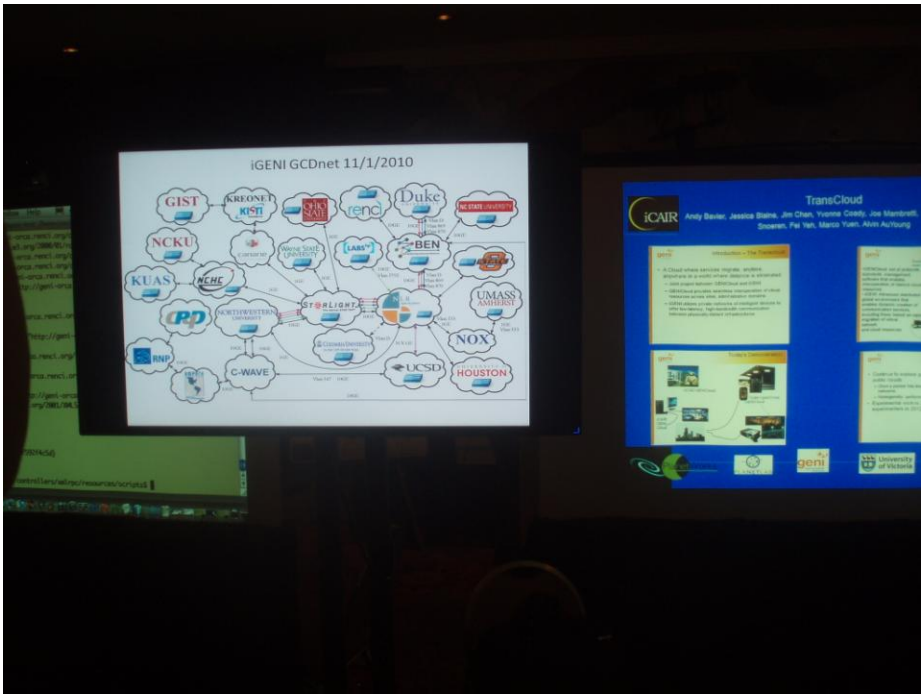

GEC 9 November 2010 Wash DC

Figure 5

SC10 November 2010 New Orleans, La

Figure 6

SC10 November 2010 New Orleans, La