iGENI Quarterly Report

GENI Project #1719

For the Period January 1, 2011 Through March 31, 2011

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

I. 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 already begun conducting experiments that utilize these 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, c) to make GENI available to more research communities, and d) to provide a platform for demonstrating its capabilities for supporting experiments. In Q2, the iGENI initiative undertook multiple efforts to expand GENI international infrastructure. For example, in January the iGENI initiative organized a meeting at the offices of the Computer Network Information Center, Chinese Academy of Sciences (CNIC/CAS) in Beijing, China, which was focused on establishing cooperative projects centered on iGENI/GENI activities. A related meeting was held in January with SingAREN in Singapore to establish a series of iGENI/GENI projects, and to plan infrastructure implementations, including those based on the newly established SingLight GLIF open communications exchange, which was inaugurated on January 31, 2011, in partnership with GLORIAD. Also, during Q2, beginning January 1, 2011, 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 among StarLight/iCAIR and and number of universities in Taiwan, Korea, Brazil, and Germany. iGENI planned and orchestrated several demonstrations at GEC 10 in San Juan, Puerto Rico March 15-17 2011, including several TransCloud demonstrations, in partnership with the TransCloud consortium, and a large scale demonstration of dynamic federated cloud provisioning with dynamic federated network provisioning. Also, planning was undertaken for demonstrations at GEC 11, which will be hosted by iCAIR/Northwestern University in Chicago in July 2011. The iGENI community has continued planning the implementation of additional sites in multiple countries, including Australia, Brazil, Canada, China, Germany, Japan, Korea, Singapore, Taiwan, Spain, New Zealand, Poland, the UK, and other countries. Also, the iGENI project participated in, and presented, at a GENI workshop hosted on February 18 at the University of Houston. iGENI participants also discussed GENI and iGENI at several other workshops, including the ON*VECTOR Photonic Workshop at UCSD (March), the DOE Office of Science Terabit Networks for Extreme Scale Science in Rockville Maryland (February), and the multi-federal agency Future Heterogeneous Networks at the NASA Ames Research Center in Mountain View (March).

A. Milestones Achieved

The basic infrastructure architecture and design has been developed for the iGENI distributed, federated environment, and these concepts have been presented and discussed at various forums, including at GEC forums -- most recently at GEC 10 conference in Puerto Rico in March 2011. The ORCA clearinghouse that has been implemented at iCAIR was updated in this quarter. 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 Communications Exchange Facility. Also, private dedicated fiber has been implemented between iCAIR/StarLight to the NLR core node at a large co-location space at 111 North Canal in Chicago. During this quarter, a second fiber pair was implemented between these sites. Optical fiber testing was undertaken for that new connection. The preliminary design was originally developed for a network to interconnect all GENI Cluster-D sites – a GENI Cluster D Network (GCDnet) 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, Ohio State, HP Research Lab, and sites in Korea, Taiwan, and Brazil. During this quarter, this network was extended to Kaiserslautern University in Germany, from that university to NetherLight in Amsterdam to the StarLight Exchange in Chicago across the GLIF infrastructure. (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.) This network will soon be extended to several additional Cluster-D sites, including the University of Houston and Columbia University. 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 national C-Wave lightpaths. Preliminary concepts and options are being discussed 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.

B. Deliverables Made

The initial design of the iGENI infrastructure has been developed and implemented, and this infrastructure is currently being extended nationally, and internationally. The majority of current activities are focused on a) planning for additional extensions nationally and internationally, providing additional resources for those extensions, providing control frameworks for those extensions, planning for researcher use of those resources, and demonstrating the capabilities of the exiting platform. Q2 extensions included the implementation of the new path to Kaiserslautern University. Also, plans are being developed to provide for sites in China and Singapore as subsequent processes to the meetings in January. The iGENI initiative organized a meeting with the Computer Network Information Center, Chinese Academy of Sciences (CNIC/CAS)in Beijing in January to establishing cooperative projects centered on iGENI/GENI activities, in partnership with the GLORIAD initiative, which will provide part of the infrastructure that will support those activities. A subsequent, related meeting was held in January with SingAREN in Singapore to establish the processes that will support a related series of iGENI/GENI projects, and to plan

infrastructure implementations, including those based on the newly established SingLight GLIF open communications exchange, which was inaugurated on January 31, 2011, also in partnership with GLORIAD. At both meetings, the potential for further extensions to other sites in Asia was also discussed. The SingLight open exchange will be a core connection facility for multiple Asian R&E networks. 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 implementation has been demonstrated as useful to support multiple major demonstrations. All future extensions are based on the current model. In addition, options for various tunneling techniques are being explored. 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. In addition, recently iGENI has been working with PlanetLab, ProtoGENI, and GpENI, to allow for additional connections and framework integration.

II. Description of Work Performed During 2nd Quarter

II..a. Activities and Findings

Q2 activities were focused on planning for, testing, and providing for additional resources, extending prototypes based on core infrastructure architectural concepts, testing and evaluating the current implementations, planning for demonstrations, staging demonstrations, and presenting the iGENI and GENI environments at various forums. 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 and to other control frameworks. 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. iCAIR supported the recent Fenius demonstration at the GLIF Technical Working Group Meeting in Honk Kong February 24-25, 2011. 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-RPC 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 Communications 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 platform was showcased during the individual project demonstrations at GEC 10 in Puerto Rico, on Tuesday March 15th. 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. (Ref : Figure 1 below).

Another GEC 10 demonstration (plenary demonstration) showcased TransCloud, which illustrated the potential for creating a highly scalable distributed computing environment comprised of federated dynamically programmable clouds integrated with dynamic network provisioning.. The TransCloud demonstrations illustrate 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, Northwestern, and Kaiserslautern) 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. Four clouds were federated and interconnected via iGENI infrastructure, HP OpenCirrus, iCAIR's OpenCloud, a research cloud at UCSD, and a cloud at Kaiserslautern University and they were used to stream digital media from repositories and live streams. (Ref: Figures 2-8)

II.b. Project Participants

The iGENI initiative consists of multiple organizational partners, including International Center for Advanced Internet Research (iCAIR), Northwestern University, Electronic Visualization Laboratory,

University of Illinois at Chicago, California Institute for Telecommunications and Information Technology (Calit2), University of California, San Diego, the computer science department, UCSD, the StarLight consortium, the Metropolitan Research and Education Network, RENCI, Duke University, North Carolina State University, MCNC, the University of Massachusetts at Amherst, Ohio State University, Wane State University, the University of Houston, Columbia University, the University of Oklahoma, University of Utah, Princeton University, University of Victoria, HP Research Labs, PlanetWorks, Kaiserslautern University, DFN, Cisco, the NLR, Merit, MAX, NOX, the GLORIAD consortium, the Global Lambda Integrated Facility (GLIF) Consortium, SURFnet, NetherLight, SARA, University of Amsterdam, CANARIE, CAnet, Communications Research Center, the Computer Network Information Center, Chinese Academy of Sciences, SingAREN, KREONET, KISTI, GIST, TWAREN, National Center for High Performance Computing, Taiwan, KAUS, NCKU, Ampath, CRP, RNP, and others..

2.c. Publications and Presentations

iGENI was presented at a meeting at the Chinese Academy of Sciences in Beijing and at a SingAREN hosted forum in Singapore on January 31, 2011. The iGENI project was integrated into a presentation at the ON*VECTOR Photonics workshop at UCSD in March, 2011. An update of the iGENI project was presented at GEC 10. 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 March 2011. The iGENI initiative has begun planning for demonstrations in the LAC/CAIR booth at the SC11 international supercomputing conference in November 2011 in Seattle.

II.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 has also been presented at multiple community forums, e.g., those related to state wide, metro, and regional R&E networking..

II.e Collaborations

The primary activities among the initiative partners noted in II.B have been focused on continuing to a) design and implement GENI infrastructure, b) design and conduct demonstrations for GEC 10, including GCDnet and TransCloud 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 10 with the ORCA framework developers, d) plan for future infrastructure implementations and collaborative activities and e) plan for future demonstrations including at GEC 11, a technology workshop in Kaiserslautern in August, the GLIF workshop in Rio in Brazil in September, and SC11 in Seattle in November.

II.f Other Contributions

The iGENI initiative is currently exploring options for integrating GENI with other major national and international testneds.

Figure 1: GCDnet

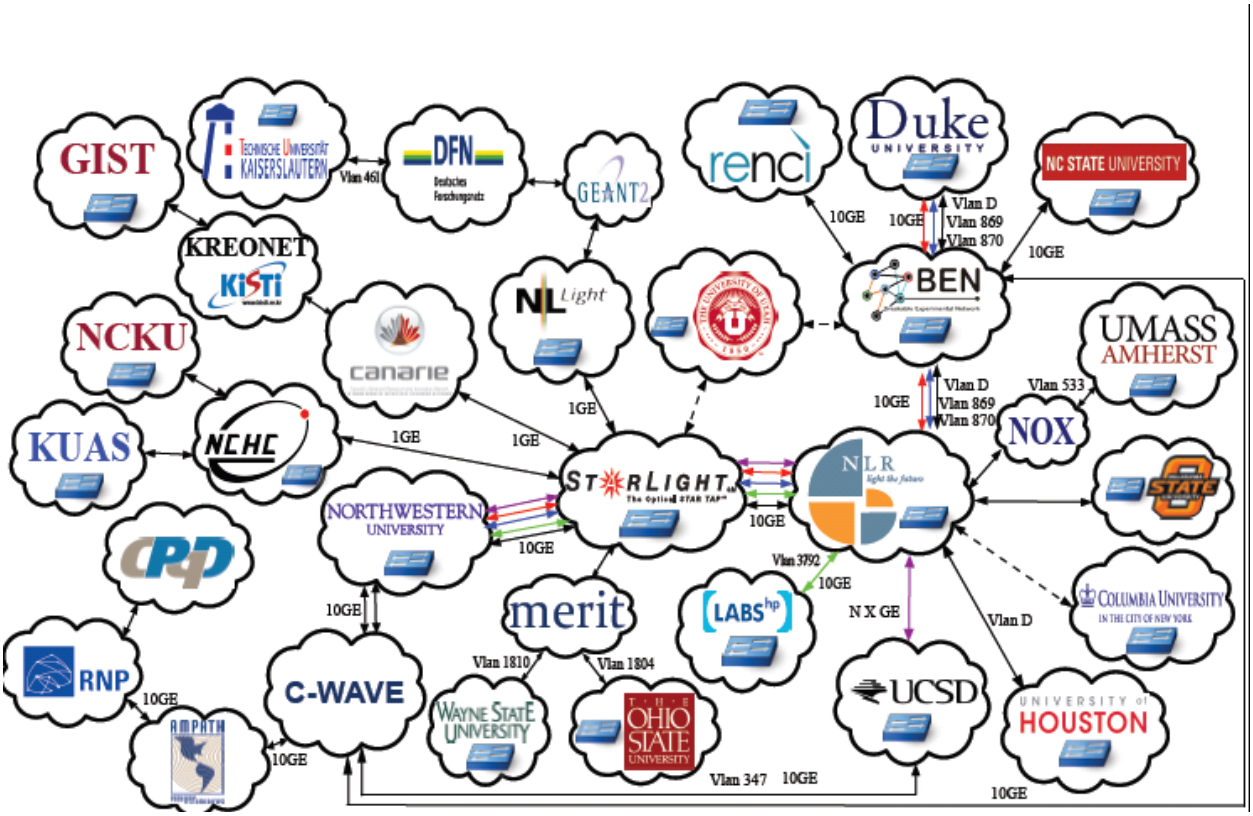

Figure 2: TransCloud Overview

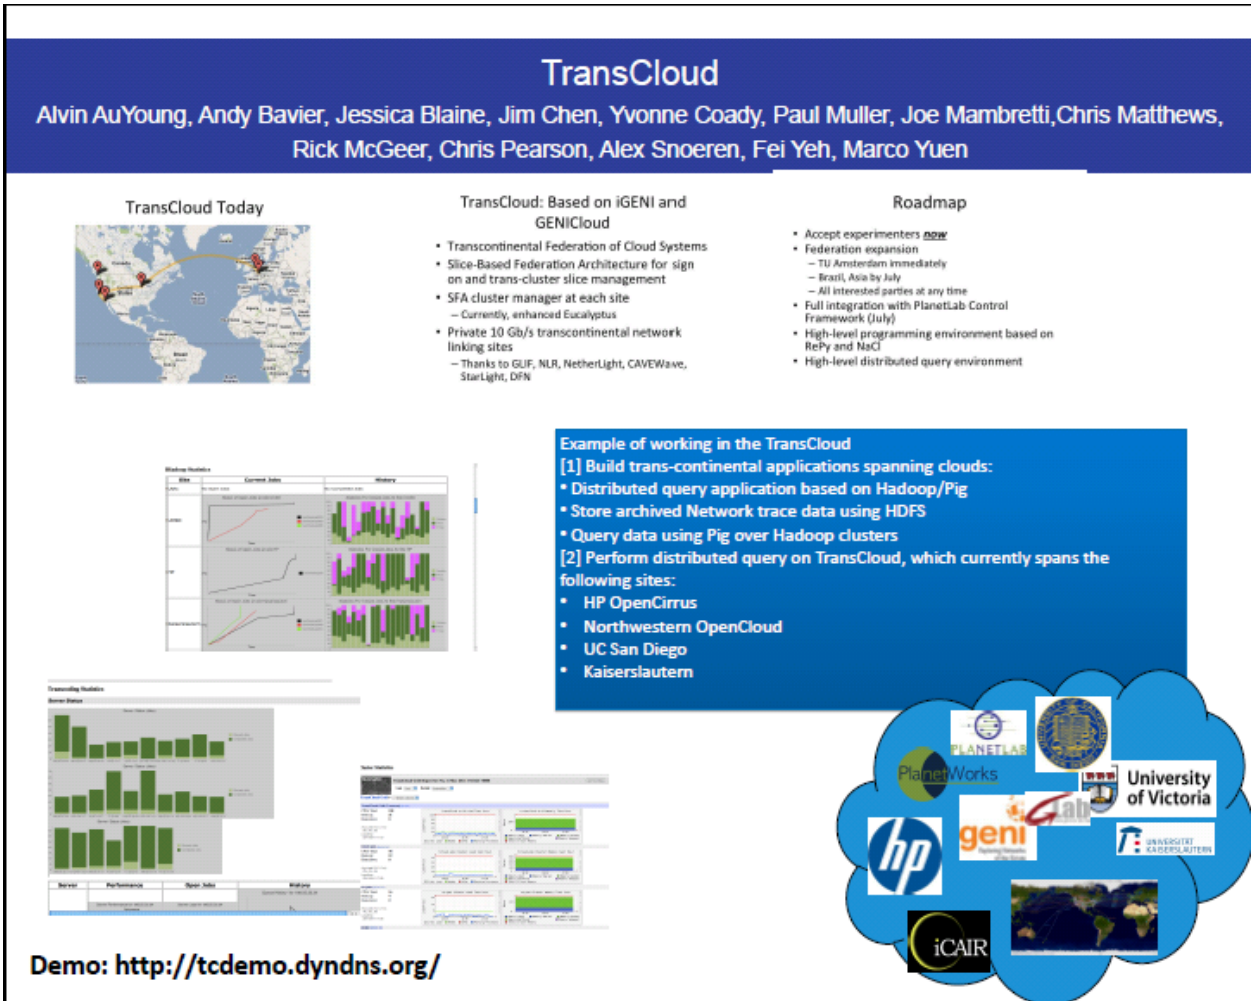

Figure 3: **TransCloud Integration**

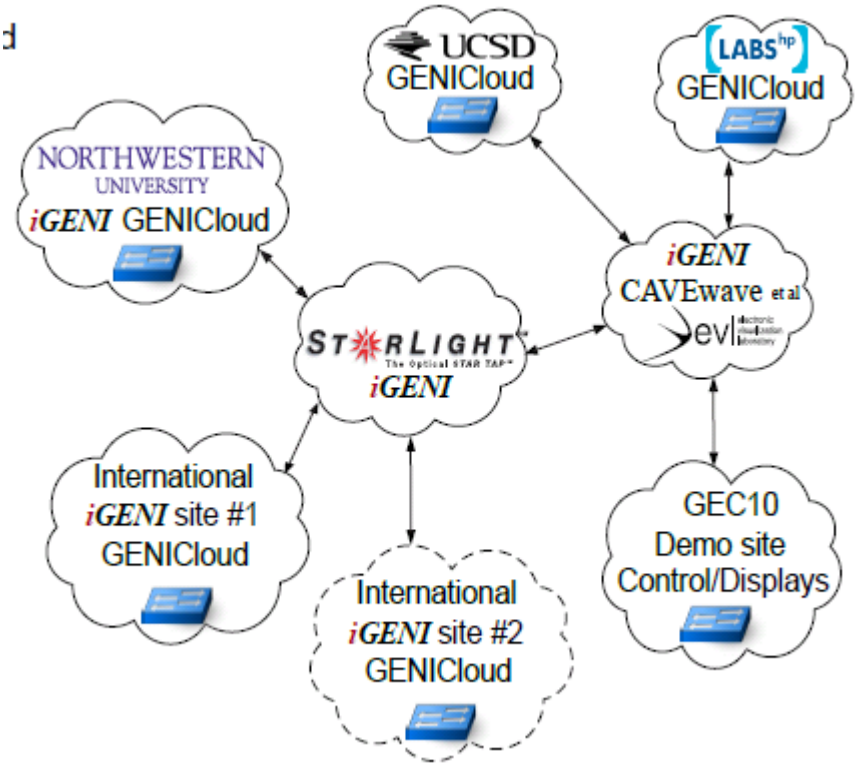

Figure 4: GEC 10 March 2011 San Juan Puerto Rico

Figure 5: GEC 10 March 2011 San Juan Puerto Rico

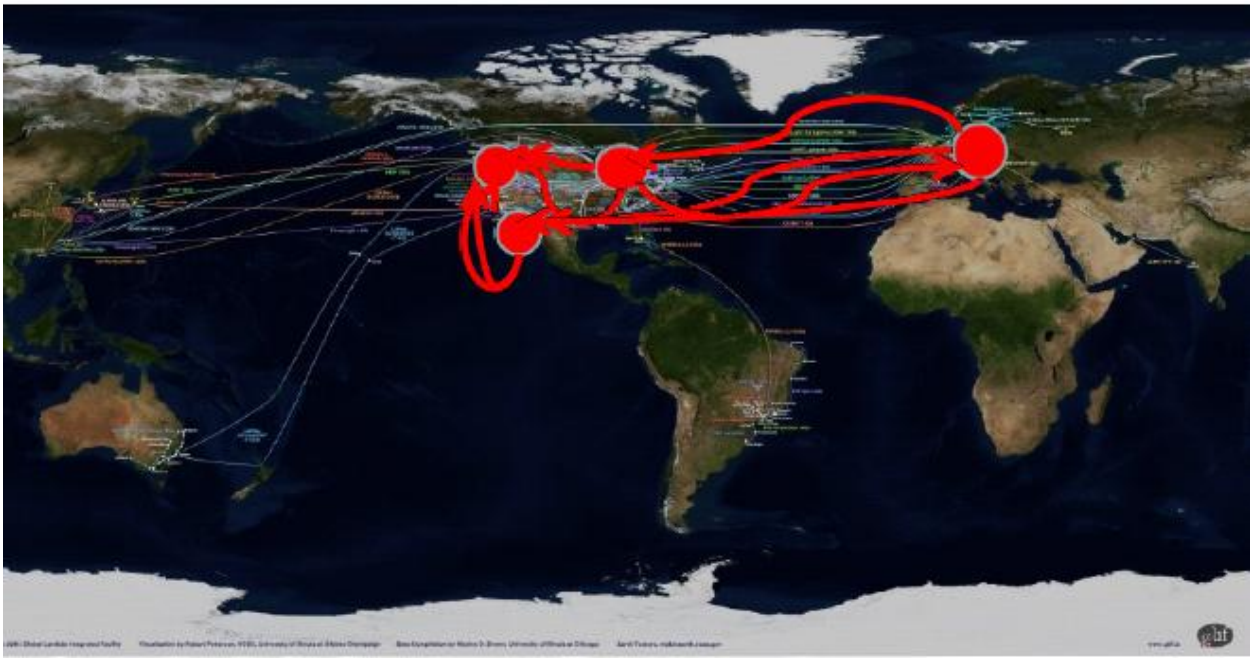

Figure 6: GEC 10 March 2011 San Juan Puerto Rico

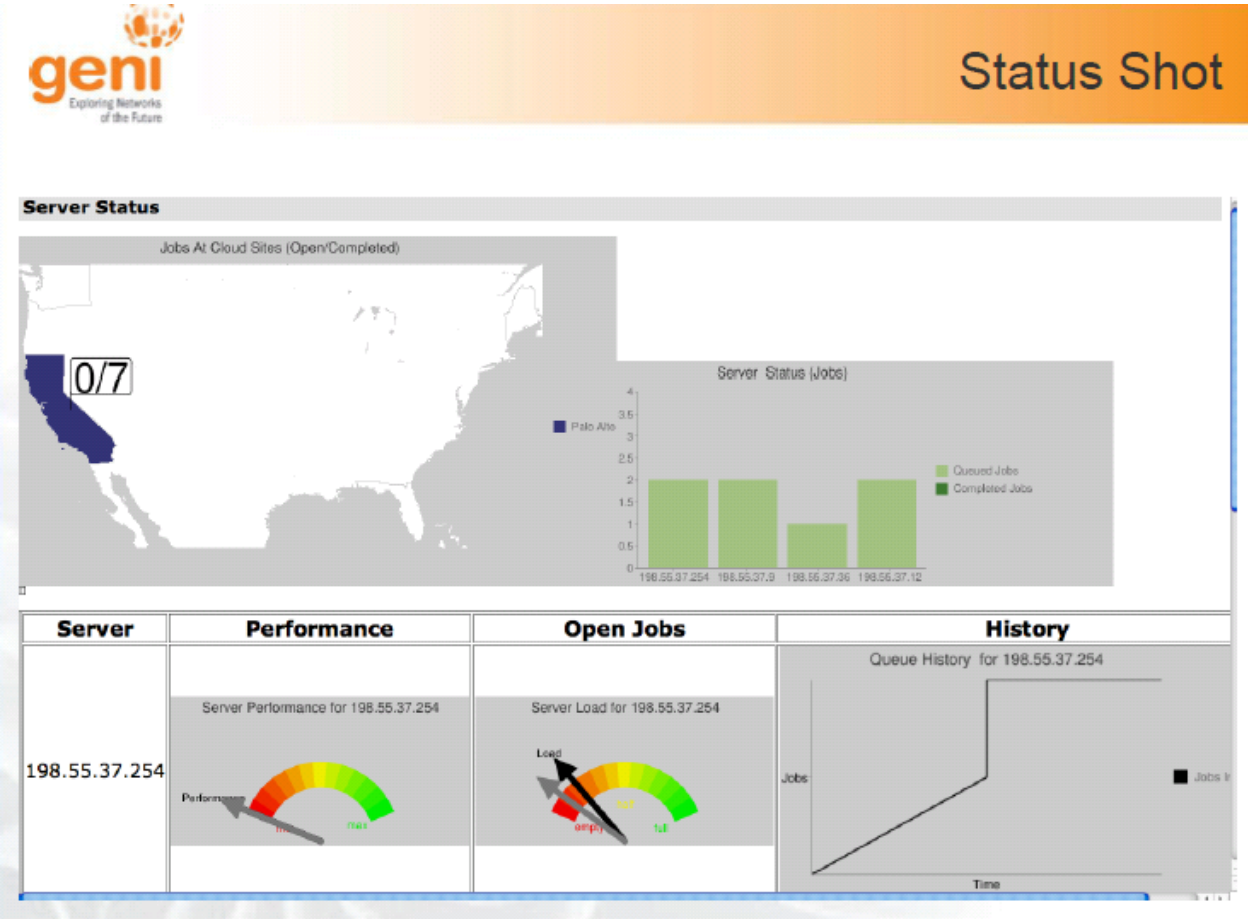

Figure 7: GEC 10 March 2011 San Juan Puerto Rico

Figure 8: GEC 10 March 2011 San Juan Puerto Rico