Visions of a Future Internet: The ExoGENI Example

Building networks of the future based on the Network Infrastructure as a Service (NIaaS) model

A RENCI/National Consortium for Data Science WHITE PAPER

renci ncds

CONTACT INFORMATION
Ilya Baldin, PhD
Telephone: 919.445.9685
Email: ibaldin@renci.org

Summary

GENI (Global Environment for Network Innovations, www.GENI.net) is a virtual laboratory that explores future Internet architectures and provides opportunities for understanding and innovation in the field of global networking. GENI, funded by the National Science Foundation (NSF), provides experimenters with large-scale virtual network infrastructure for their experiments, which are custom built from a resource pool offered by multiple service providers.

ExoGENI is a new GENI test bed that links the NSF GENI project to two advances in the arena of virtual infrastructure services, which are outside of GENI:

1. Open cloud computing – ExoGENI participates in the OpenStack (www.openstack.org) community, a global collaboration among technologists and developers that is producing an open source cloud computing platform for public and private clouds.
2. Dynamic circuit fabrics - computer networking technology that allows users to specify network services in real time.

ExoGENI enables networked cloud sites, within GENI and at external sites, to provide resources and applications which support distributed experimental and production activities, including data-intensive computing. ExoGENI also provides a high degree of control over networking functions, allowing a user to create specified virtual network topologies.

Networked Infrastructure as a Service Model

Networked Infrastructure as a Service (NIaaS) is a model that provides access to network resources in a virtual environment. It allows for network provisioning functions alongside compute and storage resource provisioning. By providing users with Application Program Interfaces (API) and resource descriptions, a user can specify the details of a network without going into the details of its implementation.

NIaaS has the ability to "stitch" resources together – this includes the ability to attach cloud resources to network links, concatenate network links from multiple providers, and allow for proper order of resource provisioning. A NIaaS model also provides the capability to plan and schedule resources in a dynamic fashion based on user needs.

The benefits of using a NIaaS model include:

1. **Reduction in cost of ownership**. With the availability of virtual networks, businesses and research laboratories no longer have to invest in hardware and other infrastructure that are expensive and become outdated quickly.
2. **Reduction of operational expenses.** With networks available through a "pay per use" model, businesses and research organizations will pay only for actual usage instead of paying for network operation even when it is not in use.
3. **Dynamic scaling.** Dynamic allocation of network services as needed provides flexibility and eliminates the need to maintain high bandwidth network connections in perpetuity.

At a Glance

- The Networked Infrastructure as a Service (NIaaS) model allows dynamic provisioning of virtual network resources. This leads to the growth of scalable and cost-effective networks.
- ExoGENI explores the future of the Internet and global networking based on the model of Network Infrastructure as a Service.
- So far 11 ExoGENI racks have been deployed at RENCI, GPO (Boston, MA), NICTA (Sydney, Australia), Duke University (Durham, NC), University of North Carolina (Chapel Hill, NC), Florida International University (Miami, FL), University of Houston (Houston, TX), North Carolina State University (Cisco Research Triangle Park Campus, NC), University of Florida (Gainesville, FL), National Energy Research Scientific Computing Center, Oakland Scientific Facility (Oakland, CA), and University of Amsterdam (Amsterdam, The Netherlands).

The Team

Written by **Ilya Baldin**, Director, Network Research & Infrastructure, RENCI; with contributions from **Paul Ruth**, RENCI Senior Distributed Systems Researcher; **Yufeng Xin**, RENCI Senior Network Research Engineer; **Anirban Mandal**, RENCI Senior Research Scientist; and **Jeff Chase,** Department of Computer Science, Duke University. Molecular systems applications contributed by **Jeffrey Tilson**, RENCI Senior Research Scientist; and **Steven Cox**, RENCI Cyberinfrastructure Engagement Lead. Additional writing support by **Smitha K. Prasad.**

Vol. 1, No. 4 in the RENCI White Paper Series, December 2013. Created in collabortion with the **National Consortium for Data Science** (www.data2discovery.org)

ExoGENI as a NIaaS Model

In ExoGENI, every piece of infrastructure is an interconnected service. An ExoGENI test bed consists of cloud site racks comprised of servers, switches and storage. Each rack is a small networked cloud that runs software allowing sites to operate with each other. ExoGENI racks are separate aggregates but have the ability to act as a single aggregate due to transparent stitching of resources.

An ExoGENI user can request Virtual Machines (VM) of varying size or bare-metal nodes (physical processing machine), specify the type of network topology linking the nodes or VMs, both within a cloud and between clouds, and specify the storage required. Each of these resources, offered as a service, is termed a "sliver." The user request that is served up by ExoGENI is termed a "slice."

An ExoGENI slice may contain a network topology with programmable nodes and links, thus acting as a virtual distributed environment. Software running within a slice manages the network using Internet Protocol (IP) or alternative protocols at the discretion of the slice owner.

An ExoGENI slice can span multiple sites and link to other GENI resources or external resources. A slice user can operate within the slice in isolation from other users and their slices. The creator of a slice can also control access to the specific slice, which interacts with the outside world via controlled interfaces. At the end of the operation, the slice creator releases the slice (and all associated resources) back into the aggregate resource pool.

ExoGENI enhances research in cyberinfrastructure capabilities by:

- Connecting cloud services to different dynamic circuit fabrics. Having the capability to establish network links dynamically allows ExoGENI to utilize a network of cloud based applications and services.
- Serving as a bridge from campus networks to national research networks.
- Offering data mobility. ExoGENI brings together data and computation either by migrating data sets to compute sites, or enabling computing to be done close to the data location. ExoGENI users operate in isolation; hence there are significant reductions in the amount of time that it takes to transfer data from one location to another. Similarly, the ability to instantiate a VM on demand enables computation to take place at the data site.
- Providing flexibility of network configuration. ExoGENI allows experimenters to work with built-to-order virtual networks.
- Incorporates resources outside of GENI into a federation. ExoGENI is able to connect to various storage resources and provides network resources using, for example, OpenStack or Eucalyptus (open source software for providing Infrastructure as a service).
- Offering control of resource utilization. Based on an 'opt –in' mode, resource providers are supplied with finely tuned resource owner controls that control how much of their resources are available to external users and for what period of time.

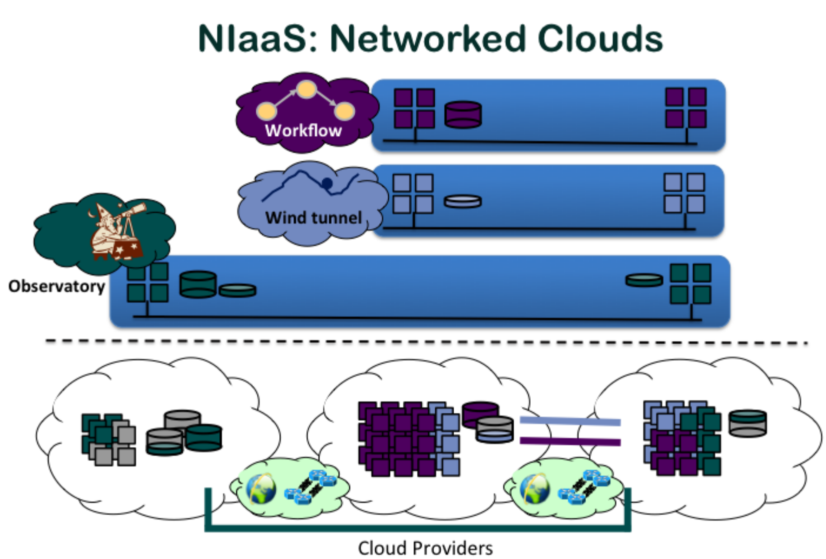

Figure 1: Network Infrastructure as a Service (NIaas) model. Multiple experimenters are shown working in isolated slices using the same infrastructure. Sophisticated orchestration capabilities facilitate the sharing network/cloud infrastructure among users. *(Graphic courtesy of Ilya Baldin, RENCI).*

ExoGENI Key Components

Each ExoGENI site is a private IaaS cloud using a standard cloud stack to manage a pool of servers with a layer of federation and orchestration control software on top. In line with the IaaS model, ExoGENI allows users to specify a virtual topology consisting of virtual machines, programmable switches and virtual network links based on Ethernet standards. The deployment is based on an evolving set of technologies including point-to-point Ethernet circuits, OpenFlow-enabled hybrid Ethernet switches, and standard cloud computing software such as OpenStack and xCAT. ExoGENI aims to leverage these technologies by:

1. Separating infrastructure control from orchestration. Each provider communicates with other services in the federation by means of APIs.
2. Using off-the-shelf software and services to control and release virtual resources to cloud sites and circuit services.
3. Being compatible with standard back-end infrastructure. This enables independent resource providers to join the federation.
4. Enabling participating providers to be autonomous. Each provider can approve or deny requests to allocate virtual resources according to their policies.

ExoGENI Key Features

ExoGENI, a funded project by the GENI Project Office, is a joint effort between the Renaissance Computing Institute (RENCI) at the University of North Carolina at Chapel Hill and Duke University (Baldin, et al. 2012) in collaboration with IBM, which provides the hardware platforms for ExoGENI deployments.

The ExoGENI test bed uses the ORCA control framework for resource orchestration. ExoGENI is a result of a long-term collaboration between RENCI and Duke across several related projects, including GENI-ORCA, Networked Clouds, and CC-NIE ADAMANT.

ExoGENI is based on the Networked Infrastructure as a service (NIaaS) model. ExoGENI can be thought of as:

- An independent collection of “racks” from which experimenters can pick the resources themselves or,
- A single aggregate with significant power to automatically locate and provision available resources in different racks and stitch them together.

The first model is consistent with the GENI approach of experimenters locating the resources that they need and manually connecting them together. In the second model, the ExoGENI control software identifies and allocates resources based on the user’s request.

ExoGENI Control Software

The ExoGENI Control framework, ORCA (Open Resource Control Architecture) enables federated orchestration of off-the-shelf private clouds and dynamic network circuit providers. Each rack contains an ORCA Aggregate Manager that delegates resources and coordinates the allocation of slices. The ORCA Aggregate Manager uses a local Broker Agent to allocate resources within one rack and a global Broker Agent to coordinate allocation for slices spanning more than one rack.

The Breakable Experimental Network (BEN), a multi-layer optical test bed built by the state of North Carolina and managed by RENCI, is one of several circuit providers for ExoGENI. The ORCA framework contains plugins that plan requested paths through the BEN network by issuing queries to models that describe the network. Since the initial phase of implementing the BEN circuit service, newer ORCA plugins have been developed to interface with other external circuit services.

ExoGENI also utilizes national research circuit providers, including Internet2, NLR FrameNet and ESnet.

ExoGENI Use Case: High Throughput Computational Screening of Molecular Systems

The global energy consumption rate is conservatively projected to increase from 13.5 Terawatts (2001) to 27 TW in 2050 (Lewis & Nocera 2006). In the U.S., this energy consumption is overwhelmingly petroleum-based (Eisenberg & Nocera 2005) even though the sun plays a unique role in energy and could, in principle, easily supply our global energy needs (Lewis 2007). However, significant technological hurdles exist in using solar energy, including its capture, storage, and dispatch of the energy to end users. One important long-term solution is the direct conversion of solar energy into useful chemical fuels: solar fuels.

Energy storage is a critical bottleneck that must be eliminated in order for solar energy to be usable as a significant energy supply. As such, researchers in the solar fuels community are seeking a solution to the energy storage problem through as-yet-unavailable high-performance, dye-sensitized photoelectrosynthesis cells (DSPEC) based loosely on the Grätzel cell (O'Regan & Gratzel 1991; Gratzel 2003). A depiction of a DSPEC device is shown in Figure 2.

Each of the indicated six steps in the DSPEC figure are substantial chemical problems being analyzed by teams of researchers. Step 2, which involves electrode injection, is an area where scientists are researching the use of Perovskites as materials of interest for use as electrodes. The research seeks to build resilient infrastructure that supports performing comprehensive sweeps on modified Perovskites using computational chemistry techniques. This yields a short set of candidate systems that can be used for further testing.

The scientific calculations involved in the research of potential electrode materials are an example of a workflow currently processed via ExoGENI. ExoGENI provides a usable, coupled, multi-scale framework based on workflow technologies and elastic resource provisioning to carry out scientific computations and resulting workflows. Since such workflows are complex and unpredictable, ExoGENI provides an ideal environment as it eliminates the need to identify in advance the scale of resources required to execute the workflow.

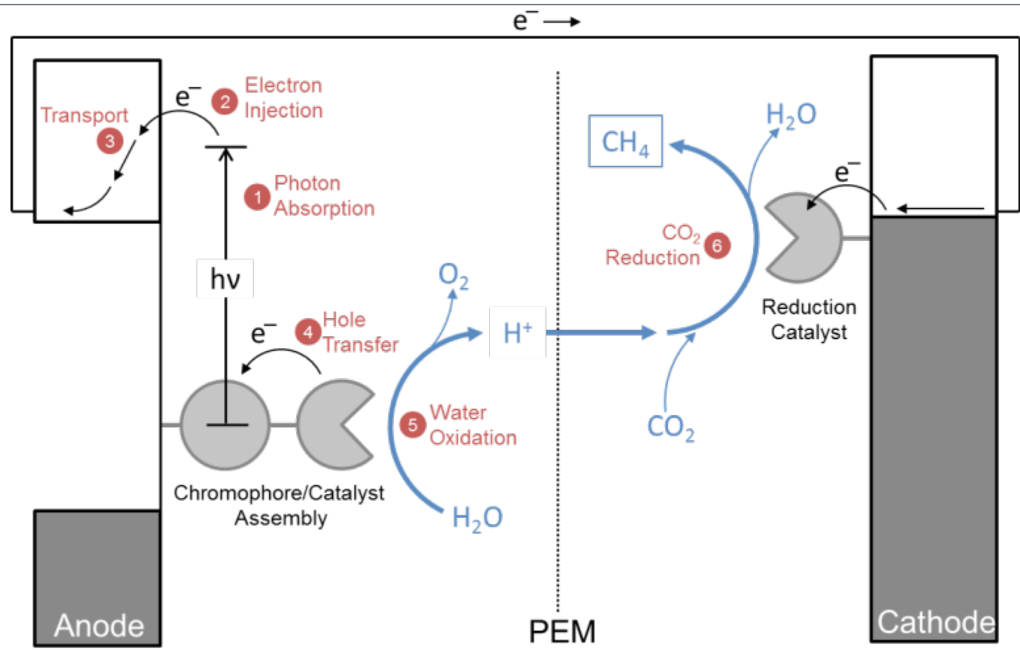

Figure 2: Depiction of a dye-sensitized photoelectrosynthesis cell (DSPEC). Photon absorption (1) by the Chromophore excites an electron to a state above the electrode Fermi level. The electron is injected (2). The remaining hole is transferred to the catalyst (4) which yields differentially stabilized Ru(IV) species. Water oxidation is catalyzed at (5). The reduction step (6) produces CH4 and the catalyst is restored. *(Graphic courtesy of the Energy Frontier Research Center, UNC-Chapel Hill. This research is funded by the U.S. Department of Energy Office of Science award DE-FC02-06ER25764. SciDAC-e: Enhancing Productivity of Materials Discovery Computations for Solar Fuels).*

The Big Picture

Distributed data intensive research is currently carried out with data being centralized and computation built around the data as a permanent function. This makes the research infrastructure very expensive to maintain, and hard to share and adapt. ExoGENI supports data intensive collaborations by creating long-lived slices in which data can reside and be distributed, augmented by short-term compute slices that perform computations on the data or on-ramps (bridges between a physical unit and a virtual machine) that allow the processed data to come in or leave the slice.

Future areas of ExoGENI research include:

- Integration of on-ramps with OpenFlow "campus" networks. This would allow ExoGENI to be deployed on OpenFlow-enabled campuses and augment their existing infrastructures.
- Inter-slice stitching, which would allow multiple slices to work together with data policies and authentication between the slices.
- Storage "slice-as-a-service" would allow data retrieval from the location closest to the user, data caching, and intelligent data management for the needs of a virtual organization or a single experiment.
- Improving performance benchmarks, which would enable networking researches to better measure the performance on network slices and NIaaS
- Slice topology templating. With topology templates in place, a user could simply select the type of topology required from a menu without having to specify the details of the topology.
- The GENI Shakedown project, which compares the performance of an application when it is run on a physical server compared to when it is run on GENI/ExoGENI, and specifies the areas of improvement.
- Integration of ExoGENI with scientific workflow tools. Integration of ExoGENI with scientific workflow tools would allow better resource utilization especially when the workflow is dynamic and difficult to predict in advance.

The Upshot

ExoGENI has pioneered the first opportunity to evaluate a federated IaaS model at scale with the ability of network provisioning.

It is a prototype for larger evolving federations that encourage participation from, and collaboration between, independent cloud solution providers. ExoGENI lays the foundation for an expandable, reconfigurable, sustainable networked "intercloud." It also enables deployment of innovative distributive services and resources as needed, thus eliminating the need for the user to maintain these resources at all times. ExoGENI provides the vision of a "Future Internet" that conforms to the needs of the user, providing bandwidth and resources as needed.

About RENCI

RENCI, an institute of UNC Chapel Hill, develops and deploys advanced technologies to enable research discoveries and practical innovations. RENCI partners with scientists, policy makers, and industry to engage and solve the problems that affect North Carolina, the U.S., and the world. RENCI is a collaboration involving UNC Chapel Hill, Duke University and North Carolina State University. For more information, see www.renci.org.

About NCDS

The National Consortium for Data Science (NCDS) is a public-private partnership that offers a foundation for advancing data science research, educating the next generation of data scientists, and translating data innovations into economic opportunity. The NCDS formed in 2012 in the Research Triangle area of North Carolina. For more information, see www.data2discovery.org

How to reference this paper:
Baldin, I., Ruth, P., Xin, Y., Mandal, A., Chase, J., Tilson, J., & Prasad. S. (2013): Visions of a Future Internet: The ExoGENI Example. RENCI/NCDS, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. http://dx.doi.org/10.7921/G0CC0XM1

References

Baldin, I., Xin, Y., Mandal, A., Ruth, P., Heerman, C., & Chase, J (2012). ExoGENI: A Multi-Domain Infrastructure-as-a-Service Testbed. Proceedings of the 8th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom)

Eisenberg, R. & Nocera, D.G. (2005). Preface: Overview of the Forum on Solar and Renewable Energy. Inorg. Chem, 44, 6799-801.

Global Environment for Network Innovations (GENI) http://www.geni.net [Accessed October 23, 2013]

Grätzel, M. (2003). Dye-sensitized solar cells, Journal of Photochemistry and Photobiology. C: Photochemistry Reviews, 4, 145-53.

Irwin, D., Chase, J.S., Grit, L., Yumerefendi, A., Becker,D., & Yocum, K.G. (2006). Sharing Networked Resources with Brokered Leases. Proceedings of the USENIX Technical Conference.

Lewis, N.S. & Nocera, D.G. (2006). Powering the planet: Chemical challenges in solar energy utilization. Proceedings of the National Academy of Sciences, 103, 15729-35.

Lewis, N.S. (2007). Toward Cost-Effective Solar Energy Use, Science, 315, 798-801.

O'Regan, B. & Gratzel,M. (1991). A low-cost, high-efficiency solar cell based on dye-sensitized colloidal TiO2 films. Nature, 353, 737-40.

ncds | THE NATIONAL CONSORTIUM for DATA SCIENCE

renci

RESEARCH \ ENGAGEMENT \ INNOVATION

ExoGENI

ExoGENI is a joint effort between RENCI/UNC-Chapel Hill and Duke University in collaboration with IBM, which provides the hardware platforms for ExoGENI deployments. ExoGENI is funded by the GENI Project Office. Learn more at www.exogeni.net