Managing Dynamic Networked Cloud Infrastructure for Data Driven Scientific Workflows Using Proactive Introspection

Anirban Mandal, Paul Ruth, Ilya Baldin,
Yufeng Xin, Claris Castillo,
Fan Jiang, Charles Schmitt
RENCI - UNC Chapel Hill
{ibaldin,pruth,anirban,yxin,claris,dcvan,cschmitt}@renci.org

Gideon Juve, Mats Rynge,
Ewa Deelman
ISI - USC
{gideon,rynge,deelman}@isi.edu

Jeff Chase
Duke University
chase@cs.duke.edu

***Abstract*—This demonstration will showcase a novel, dynamically adaptable cloud infrastructure driven by the demand of a data-driven scientific workflow. It will use resources from ExoGENI - a Networked Infrastructure-as-a-Service (NIaaS) testbed funded through NSF's Global Environment for Network Innovation (GENI) project. The demo will connect compute and data resources in the RENCI SC14 booth to a large dynamically provisioned 'slice' spanning multiple ExoGENI cloud sites that are interconnected using dynamically provisioned connections from Internet2 and ESnet. The slice will be used to execute a scientific workflow driven from a computer in the RENCI SC14 booth connected to the slice via SCinet. We propose to demonstrate the features of "ShadowQ" , an entity that predicts future resource needs of a workflow, and runs alongside the Pegasus workflow management system. This workflow introspection feature will be used to adapt the slice to the demands of the workflow as it executes, by adjusting the amount of resources used.**

I. Overview

This demonstration will showcase a novel, dynamically adaptable cloud infrastructure driven by the demands of a scientific data-driven workflow. It will use resources from ExoGENI - a Networked Infrastructure-as-a-Service (NIaaS) testbed funded through NSF's Global Environment for Network Innovation (GENI) project. GENI is a federation of testbeds created to support experimentation in distributed and networking systems. ExoGENI links computational, network, and storage resources from a growing number of institutional cloud sites using dynamically-provisioned layer-2 network connections via Starlight, Internet2, ESnet, and BEN (http://ben.renci.org).

ExoGENI (Figure 1) is based on an extended networked Infrastructure-as-a-Service (IaaS) cloud model with orchestrated provisioning of computation, network, and storage resources from independently operated cloud sites. It resides at the intersection of networking and cloud computing. Independent resource administrators delegate certain functions for identity management, authorization, and resource management to common coordinator services. This structure enables a network of private clouds to operate as a single federated hybrid community cloud. The ExoGENI testbed is designed to support research and innovation in networking, operating systems, distributed systems, future Internet architectures, and deeply networked, data-intensive cloud computing. It is envisioned as a novel distributed cloud infrastructure that bridges advanced distributed systems experimentation with computational and data-driven sciences, high-performance and high-throughput computing.

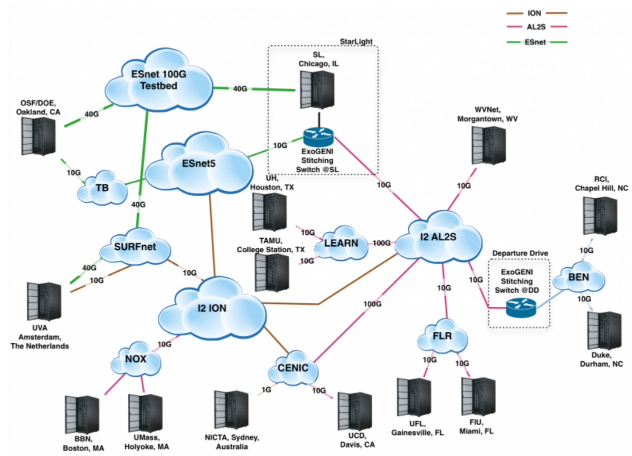

Fig. 1. Snapshot of current state of ExoGENI

The proposed demonstration will connect compute and data resources in the RENCI SC14 booth to a large dynamically provisioned 'slice' spanning multiple ExoGENI cloud sites that is interconnected using dynamically provisioned connections over I2, ESnet, and BEN. SCinet resources will be required to connect the resources in the booth to the slice. The slice will be used to execute a Pegasus [2] scientific workflow initiated from a computer in the booth. Data for the workflow execution will also be contained at the booth and will flow into the slice.

This demonstration will also showcase how the provisioning of the underlying networked cloud resources is automatically controlled by the needs of the workflow. We have developed "ShadowQ" , an entity that predicts future resource needs of a workflow, and runs alongside the Pegasus workflow management system. This workflow introspection feature will be used to adapt the slice to the demands of the workflow as it executes, by adjusting the amount of resources used.

The ExoGENI team has extensive experience with provisioning and executing scientific workflows on networked cloud infrastructures. Figure 2 shows an example ExoGENI slice request and manifest for workflow resources.

II. Technology Description

This demonstration will focus on a few of the recent ExoGENI and Pegasus software features:

- **ShadowQ**: An entity that predicts future resource needs of a workflow through proactive introspection, and runs alongside Pegasus.
- **Application to NIaaS**: Support software layer to facilitate ease of use of networked cloud systems, which gleans requirements from ShadowQ and automatically constructs appropriate topology requests.
- **On-ramps, Storage Slivering**: Advanced features like bandwidth-provisioned network links between ExoGENI resources and physical resource controlled by third party and dynamically provisioned iSCSI storage volumes.
- **Resource Adaptation**: Dynamic modification of Infrastructure based on predictions from ShadowQ.

Specifically, the demonstration will use a data-driven Pegasus scientific workflow running on a dedicated physical machine in the RENCI booth on the SC14 floor. The workflow deals with high-throughput genome sequencing. A large slice of ExoGENI resources including dozens of virtual machines and iSCSI storage volumes will be requested from ExoGENI sites and an on-ramp to the workflow machine in the RENCI booth. Workflow data will flow from the machine in the booth into the slice using the SCinet on-ramp to Internet2. The resulting output of the workflow will be sent back from the slice to the machine in the booth. Resource provisioning and adaptation will be controlled by ShadowQ and Application to NIaaS software. We will also demonstrate user tools and workflow packaging technologies, which will leverage the Application to NIaaS software to instantiate slices for workflows on ExoGENI.

III. Relevance to the Community

Advanced networks are an essential element of data-driven science enabled by next-generation cyberinfrastructure environments. Recent advances in enabling on-demand network circuits in backbones, advances in Software Defined Networking (SDN) like OpenFlow and programmable edge technologies like OpenStack create a unique opportunity to enable complex scientific applications to run on dynamic, specially tailored to their needs, infrastructure that includes compute, storage and network resources.

Workflows, especially data-driven workflows and workflow ensembles are becoming a centerpiece of modern computational science. However, scientists lack the tools that integrate the operation of workflow-driven science applications on top of dynamic infrastructures that link campus, institutional and national resources into connected arrangements targeted at solving a specific problem. The ExoGENI testbed [1] enables computational scientific workflows that currently use existing static high-throughput and high-performance computing facilities to dynamically provision custom network topologies in order to move data to computation or vice versa.

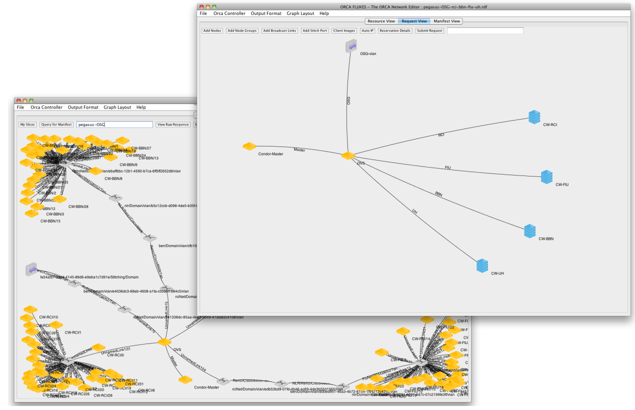

Fig. 2. Example Slice Request (top right) and Manifest (bottom left) for a workflow

IV. Description of Resources

The demonstration will use compute, storage, and network resources from multiple ExoGENI sites along with network resources from existing ExoGENI partner transit fabrics (ESnet, Internet2, BEN, Starlight). Figure 1 shows many of the existing ExoGENI partners. These resources are currently available through ExoGENI and no additional configuration is necessary.

In addition to existing resources, SCinet resources will be used to extend the ExoGENI infrastructure to the RENCI booth on the SC14 exhibition floor. This additional functionality will require a bandwidth provisioned 10Gbps link to the RENCI both and one or more VLANs across SCinet between a the RENCI booth and a national network transit provider (preferably Internet2).

Acknowledgment

This work was supported by the NSF CC-NIE ADAMANT project (NSF ACI 1245926), DoE DROPS project (ASCR DE-SC0005286), DoE SciDAC SUPER (DE-FG02-11ER26050 / DE-SC0006925) project, NSF SDCI Missing Link project (NSF ACI 1032573), NSF ACI 1148515: SI2-SSI project, and the NSF GENI project (GENI Award #1872).

References

[1] I. Baldine, Y. Xin, A. Mandal, P. Ruth, A. Yumerefendi, and J. Chase. ExoGENI: A Multi-Domain Infrastructure-as-a-Service Testbed. *8th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM 2012)*, June 2012.

[2] E. Deelman, G. Singh, M.-H. Su, J. Blythe, Y. Gil, C. Kesselman, G. Mehta, K. Vahi, G. B. Berriman, J. Good, A. Laity, J. C. Jacob, and D. S. Katz. Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Sci. Program.*, 13(3):219–237, 2005.