ExoGENI and GIMI: GENI Racks and Their Measurement and Instrumentation Tools

Michael Zink
Electrical and Computer Engineering Department
University of Massachusetts Amherst
Amherst, MA, USA
zink@ecs.umass.edu

Jeff Chase
Department of Computer Science
Duke University
Durham, NC, USA
chase@cs.duke.edu

Ilia Baldine
RENCI/UNC Chapel Hill
Chapel Hill, NC, USA
ibaldin@renci.org

Max Ott
Networks Research Group
NICTA
Sydney, Australia
max.ott@nicta.com.au

I. INTRODUCTION

In this paper, we give an overview on two current GENI projects. First, we describe ExoGENI, a new GENI testbed that links GENI to two advances in virtual infrastructure services outside of GENI: open cloud computing (OpenStack and Eucalyptus) and dynamic circuit fabrics. ExoGENI orchestrates a federation of independent cloud sites and circuit providers through their native IaaS interfaces, and links them to other GENI tools and resources. The ExoGENI deployment consists of cloud site “racks" on host campuses within the US, linked with national research networks and other circuit networks through programmable exchange points. The ExoGENI sites and control software are enabled for software-defined networking using OpenFlow. ExoGENI offers a powerful unified hosting platform for deeply networked, multi-domain, multi-site cloud applications.

Second, we describe the GENI Measurement and Instrumentation project GIMI. This project is developing and deploying the GIMI instrumentation and measurement framework, capable of supporting the needs of both GENI experimenters and GENI infrastructure operators. It uses the ORBIT Measurement Framework and Library (OMF/OML) and integrated Rule Oriented Data System (iRODS) as its basis. It provides libraries to instrument resources, to filter and process measurement flows, and to consume measurement flows. It uses the iRODS data grid for archiving and further processing and will include access control based on accepted GENI policy and authorization mechanisms.

Finally, we will demonstrate how GIMI can be used to instrument an experiment running on an ExoGENI slice.

II. EXOGENI

ExoGENI is a new testbed at the intersection of networking and cloud computing, funded through NSF's GENI project. 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.

ExoGENI is based on an extended Infrastructure-as-a-Service (IaaS) cloud model with orchestrated provisioning across sites. Each site is a private IaaS cloud using a standard cloud software stack to manage a pool of servers. The sites federate by delegating 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 hybrid cloud.

ExoGENI combines this structure with a high degree of control over networking functions: traditional vlan-based and OpenFlow networking within each site, multi-homed cloud servers that can act as virtual routers, site connectivity to national circuit backbone fabrics through host campus networks, and linkages to international circuits through programmable exchange points. The project aims to enhance US research cyberinfrastructure capabilities in four inter-related ways:

- The missing link. ExoGENI interconnects clouds to dynamic circuit fabrics, enabling a range of networked cloud applications and services, including data-intensive interaction, distributed data sharing, geo-replication, alternative packet networks, and location-aware services.
- On-ramps to advanced network fabrics. ExoGENI shows how to use campus clouds to bridge from campus networks to national transport network fabrics. ExoGENI cloud sites can act as virtual colocation centers that offer on-demand cloud services adjacent to fabric access points. Sites at fabric intersection points can also act as virtual network exchanges to bridge “air gaps” between network fabrics stemming from lack of direct connectivity or incompatible circuit interfaces.
- Cloud peering and data mobility. ExoGENI enhances the potential for peering and sharing of private clouds. It offers a means to bring data and computation together by migrating datasets to compute sites or placing computation close to data at rest.
- Networking as a service. ExoGENI brings flexible network configuration to cloud computing. It also enables

experimental deployments of new packet networking models over a flexible link substrate. Built-to-order virtual networks can implement routing overlays using IP or other packet-layer protocols. Testbed users may deploy custom node operating systems with alternative networking stacks into their nodes, and use OpenFlow data paths and/or virtual routers to implement new network services at the cloud edge and at network intersection points.

III. GIMI

The fist goal of GIMI is to provide easy-to-use I&M services for experimenters, who are deploying slices in GENI on selected types of servers, VMs, racks, and WiMAX base stations, interconnected by various types of network paths.

Its second goal is to provide comprehensive infrastructure measurement services for infrastructure operators, who are deploying measurement slices on selected types of servers, VMs and racks, interconnected by a various types of network paths. In addition, slices established by other infrastructure operators, or by experimenters, can be authorized to gather data from the measurement slices.

This project will build and operate two persistent services: the GENI Measurement Data Archive service (i.e., iRODS) and the GENI Experimenter Portal Service. GIMI builds on the Orbit Measurement Framework and Library (OMF/OML), which support testbed control, measurement, and management.

The initial GIMI prototype was developed with the goal to support experiments executed on top of the ExoGENI infrastructure. In the following, we give an overview on the major GIMI components and then describe an example experiment performed on and ExoGENI slice.

A. OMF/OML

1) OMF - OMF provides a set of tools that allow an experimenter to describe and instrument an experiment, execute it and collect its results. It has substantial scripting capabilities, which allow experimenters to perform automated, large-scale experiments.

2) OML - OML is a measurement library that allows application writers to define customizable measurement points inside applications. Experimenters can then direct measurement streams from these measurement points to storage in a (remote) measurement database. OML consists of a client library, an OML server that is responsible for collecting the measurement data, and an OML proxy server that can buffer measurement data locally and sent them to the server later (e.g., to support mobile experiments).

3) GIMI Support - To simplify the use of OMF and OML in GIMI, we have created an ExoGENI virtual machine image that is publicly available to the GENI experimenter community. Booting this image in an ExoGENI rack provides an experimenter with the tools to perform and OML/OMF-based experiment. This image also includes and iRODS client (see below) to enable the storage of measurement results in the GIMI Measurement Data Archive.

B. Integrated Rule Oriented Data Sytsem (iRODS)

iRODS is a data grid software system that allows the sharing of data across space. It was designed to support Data Grids, Digital Libraries, Persistent Archives, and Real-time Data Systems. We selected iRODS as the basic data storage service for the Measurement Data Archive for the following reasons: (1) iRODS allows federation of data grids; thus several iRODS servers can be *federated* into one data system, (2) iRODS provides a metadata catalog (iCAT) which tracks the state of the data stored in the system, (3) metadata can be easily added to the raw data; an important feature when it comes to the storage of measurement data, finally (4) iRODS offers a large set of client software that allow users easy access to their data using different operating systems.

We have created an initial Measurement Data archive for the GIMI project, consisting of three data servers at RENCI, NICTA, and UMass Amherst. In addition, an iCAT sever for these data system is running at RENCI. For measurements that are executed with GIMI the OML server automatically stores the measurement results in the iRODS home directory of the user performing the experiment. In addition, rudimentary metadata information is also automatically created and stored along with the measurements.

C. Internet Remote Emulation Experiment Laboratory (IREEL)

IREEL was originally designed to conduct experiments with real Internet applications and protocols in the context of network course for university students. We are currently in the process of expanding IREEL into what we describe a “Lab Wiki”. This extended version of IREEL will serve as the measurement portal for GENI experimenters through which they can define and execute experiments. In addition, a digital version of a lab book will allow experimenters to document and annotate their experiments. Currently, the GIMI project provides a persistent IREEL server that allows experimenters to execute experiments on previously allocated ExoGENI resources.

IV. DEMONSTRATION

We will use a measurement experiment executed with GIMI tools on an ExoGENI slice to demonstrate the capabilities of both projects. For this experiment a simple triangular topology is set up. On top of this topology two separate Iperf measurements are performed. One measures the throughput between two directly connected nodes, while the other one measures the throughput for the same pair of nodes, which are connected through a software router. We will demonstrate the setup of the ExoGENI slice and the execution of the experiment based on GIMI tools.

V. ACKNOWLEDGEMENTS

The work presented in this paper is supported by the NSF GENI grants “Network-agile multi-provisioned infrastructure for GENI and ExoGENI” and “GIMI: Large-scale GENI Instrumentation and Measurement Infrastructure.”