

GENI Quarterly Report for DiCloud project

January 1, 2010 - March 31, 2010

Michael Zink, Prashant Shenoy, Jim Kurose, David Irwin and Emmanuel Cecchet
{zink, shenoy, kurose, irwin, cecchet}@cs.umass.edu

University of Massachusetts, Amherst
140 Governors Drive
Amherst, MA 01003-9264

I. Major accomplishments

The first quarter of the DiCloud project includes the following major accomplishments:

- Development of Orca handlers and brokers for Amazon EC2 resources.
- A successful intra-cluster demonstration at GEC7 that show-cased ViSE sensor data streaming from a web server running on Amazon EC2 over a dynamically stitched VLAN through NLR and BEN to servers at RENCi using the Orca control framework. The demonstration was presented at a plenary session of GEC7. To complete this demonstration, we worked with the GPO in Boston, Joe Mambretti at the Starlight facility, Iliia Baldine at RENCi, and Jeff Chase at Duke University.
- Continued our collaboration with the University of Puerto Rico, Mayagüez (UPRM), as part of our outreach plan. In mid-January Michael Zink and David Irwin organized and led a 3-day seminar at UPRM on GENI, virtualization, cloud computing and sensor networks.
- Testing and prototyping of various options to implement Orca handlers for EC2 resources.

The rest of this document describes in detail the major accomplishments above.

I.A. Milestones Achieved

We achieved the following milestones in the 2nd quarter as specified in our original Statement-of-Work.

- **January 29, 2010.** Develop 3 Orca handlers to allocate and revoke resources from Amazon's Elastic Compute Cloud (EC2), Simple Storage Service (S3), and Elastic Block Store (EBS) cloud services. Explore the feasibility of integrating third-party handlers from either Eucalyptus or OpenNebula into GENI/Orca. Note that the handlers interact with Amazon's API perform allocation/authorization functions, but do not expose the Cloud API to slice controllers.

We have developed three Orca handlers for Amazon's Elastic Compute Cloud (EC2), Simple Storage Service (S3), and Elastic Block Store (EBS) cloud services. The handlers directly invoke the native EC2 tools (s3cmd for S3). While Eucalyptus and OpenNebula offered interesting alternatives to the EC2 native APIs they don't offer the monitoring capabilities that are required in the future proxy to estimate the cost of resource usage.

- **February 10, 2010.** Develop a first-come first-served (FCFS) clearinghouse (broker) policy that tracks the amount of resource time incurred by each cloud user. Note that this

policy does not track fine-grained usage costs, such as the number of I/Os (for EBS) or the aggregate network traffic (for EC2/S3). The proxy will serve this function.

We have re-used existing Orca brokers accounting node resources as the current handlers cannot report resource usage. We have explained how to define a given budget into EC2 resources that can be managed by existing brokers in the Orca implementation.

- **March 16, 2010.** Demo 1: Archiving of sensor data on a storage service. We will demonstrate an early version of our handlers and our simple FCFS Clearinghouse policy, which will allow a user to lease servers from EC2 and storage volumes from EBS (as an aggregate). We will highlight the addition of storage by uploading sensor data (from CASA's Oklahoma testbed) to EBS through an EC2 server. The functions will be exposed by extending ViSE's web portal.

We performed a demo at GEC7 where RENCI dynamically stood up a web server after stitching together the larger network that queried archived radar data from a web server running on an EC2 node on Amazon, which ViSE bridged onto the VLAN using OpenVPN through `visetestbed.cs.umass.edu` (actually its private VLAN IP). The radar data was displayed as a Google Maps animation after Orca finished stitching the links from BEN, NLR, Starlight, to UMass.

I.B. Milestones in Progress

- **May 15, 2010.** (deliverable S2f) Report on the feasibility of using Amazon's new CloudWatch service to monitor fine-grained usage costs. Amazon's CloudWatch was not available at the time of the proposal; in the initial proposal we planned to not give users root access on servers, and, instead, proposed a per-server root monitoring daemon. CloudWatch may remove this need.

We have been looking at multiple options to monitor EC2 resource usage and its cost. While a proxy can easily determine the running time of VMs, estimating network and disk usage accurately is much harder. Amazon CloudWatch seems to be a good solution for real time monitoring and cost can be adjusted at the end of a lease by extracting the real cost from Amazon billing service. We are testing different proxy strategies to get the most accurate resource utilization information.

II. Deliverables Made

Deliverable S2.c has been produced on January 28, 2010 along with the Orca handlers code for Amazon EC2 cloud resources. Deliverable S2.d has been provided on February 11, 2010 to explain how EC2 resources can be used with existing brokers. Deliverable S2.e has been demonstrated at GEC7.

III. Description of Work Performed During Last Quarter

The primary work during the quarter has been the implementation of Orca handlers and the demonstration code at GEC7. We also have the materials for the GENI Seminar we conducted in Puerto Rico in mid-January.

III.A. Project Participants

The primary PI is Michael Zink. Co-PIs are Prashant Shenoy, and Jim Kurose. Research Staff is David Irwin and Emmanuel Cecchet.

III.B. Publications (individual and organizational)

Bo An, Victor Lesser, David Irwin, and Michael Zink - *Automated Negotiation with Decommitment for Dynamic Resource Allocation in Cloud Computing* – Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS), Toronto, Canada, May 2010.

III.C. Outreach Activities

As part of our involvement in the GENI project, we scheduled a seminar at the University of Puerto Rico, Mayaguez in early January. The primary purpose of the seminar was to teach students at UPRM about emerging technologies in virtualization, cloud computing, wireless communication, networking, and sensing that make it possible to multiplex experimental testbeds, such as those being incorporated into the GENI prototype. Over the workshop's two days Michael Zink and David Irwin gave lectures and tutorials on virtualization, the GENI project, wireless communication, research efforts at UMass-Amherst and cloud computing.

III.D. Collaborations

We collaborated with other Cluster D projects significantly during the quarter. First, we provided documentation on the Cluster D mailing list about how to port controllers to the new Bella release. Second, we worked with the GPO, RENCI, Duke, and Northwestern to arrange for the VLAN control in the GEC7 demonstration.