

GENI Quarterly Report for DiCloud project

October 1, 2009 - December 31, 2009

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:

- In conjunction with the ViSE and DOME projects, DiCloud has been working with BBN, NLR, and our Cluster D peers to enable VLAN connections. In the last quarter we successfully tested a VLAN connection to BBN's offices in Cambridge
- Participation to GEC6, presentations and discussions to manage EC2 resources within the Orca framework.
- Study of options and cost implications for GENI network connectivity with Amazon EC2 cloud resources.
- 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 1st quarter as specified in our original Statement-of-Work.

- **November 16, 2009.** Collaborate with GPO and other projects in Cluster D, to complete a Spiral 2 plan for the setup of VLANs between aggregates, to be carried by the Internet 2 (or NLR) backbone network between the aggregates.

During the quarter we were able to successfully test a static VLAN connection from our offices in the Computer Science Department at UMass-Amherst to BBN's offices in Cambridge, Massachusetts. We are currently working with Kathy Benninger at NLR and the GPO to determine the best way to deliver our traffic to NLR, since UMass-Amherst is not an active member of NLR. The current plan is to apply to NLR to obtain a separate port at the Northern Crossroads (NOX) in Boston for UMass GENI traffic. As a backup, we will use the BBN port at the NOX to allow our traffic to get on NLR. We are also working with our Cluster D peers to enable cross-site VLANs using Northwestern's Starlite project.

- **December 23, 2009.** Report on the options and cost implications for network connectivity with Amazon's cloud, including the use of publicly addressable IPs, Amazon's recently announced Virtual Private Cloud Service, and VLANs now and in the future. Pick initial approach for use in year 1.

Deliverable S2.b contains a description of the current options available to access Amazon EC2 cloud resources with their associated cost. By default, EC2 servers are provided with a public IP address for generic access. Recently, a new service called Amazon Virtual Private Cloud (VPC) has been added to provide layer 3 VLAN capabilities. This service is only available in beta for specific software that is only available in 2 product brands. As no layer 2 solution is currently available, we recommend the usage of a free software solution such as OpenVPN to run inside the virtual machines hosted on EC2. We expect the Amazon offering to expand quickly over the next year at which stage we will be able to offer a revised recommendation.

I.B. Milestones in Progress

- **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 been prototyping different approaches to implement Orca handlers including Ant task or direct invocations of the EC2 tools. While Eucalyptus and OpenNebula offer interesting alternatives to the EC2 native APIs they don't offer the monitoring capabilities that will be required in the future to estimate the cost of resource usage. We are leaning towards a solution based on the EC2 native tools.

- **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 been thinking about various ways of exposing EC2 resources to the users. While EC2 instances can be registered and managed as a pool of servers, disk resources (S3 or EBS) can be created dynamically, are likely to have a longer life/lease than one experiment and can potentially be shared among multiple users. There is also a cost for persistent data even if it is not used and it will be necessary to reclaim these resources after a certain expiration date.

II. Deliverables Made

Deliverable S2.b has been produced on December 22, 2009. The document describes the current network connectivity solutions available to access Amazon EC2 cloud resources. As no layer 2 solution is currently available, our recommendation for the first year is to use a free software solution such as OpenVPN running inside the Amazon EC2 instances.

III. Description of Work Performed During Last Quarter

The primary work during the quarter, including our Activities and Findings, centered on achieving the milestones described above and making progress toward our upcoming milestones. We also attended GEC 6 where we presented our approach and discussed the specifics of EC2 resources within the Orca framework.

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)

We did not submit any external publications besides the deliverable due during this quarter.

III.C. Outreach Activities

As part of the ViSE project, we are teaching seminars and tutorials on GENI at the University of Puerto Rico Mayaguez from January 11th to January 15th, 2010. As part of this lecture series and the DiCloud project, we are also including seminars on cloud resources from Amazon, including EC2, S3, and EBS tutorials.

III.D. Collaborations

We are actively working with the GPO to connect our VLANs to NLR. In the last quarter, we tested our VLAN from our offices to the BBN's offices in Cambridge. Additionally, we are also working with both RENCi and Northwestern on cross-site VLANs. Finally, we are having ongoing conversations with CASA researchers on their experience running experiments in end-to-end sensor networks and that we plan to keep them informed about the activities and the progress in DiCloud.