

GENI Quarterly Report for DiCloud project

April 1, 2010 - June 30, 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 third quarter of the DiCloud project includes the following major accomplishments:

- Investigation of Amazon CloudWatch to monitor EC2, S3 and EBS resources.
- Development of a proxy aggregate manager for Amazon resources.
- Preparation of a demo for GEC8 integrating data from Vise and Dome projects in an experiment using Amazon resources for storage and computation.
- Dissemination activities including presentations in Europe and a research paper at LANMAN 2010.

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

I.A. Milestones Achieved

We achieved the following milestones in the 3rd quarter as specified in our original Statement-of-Work.

- **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.

I.B. Milestones in Progress

- **July 20, 2010:** (deliverable S2.g) Demo 2 at GEC8: Researchers use an initial version of the proxy aggregate manager to interact with EC2 control interfaces (e.g., to reboot servers or get console output). Note that these EC2 control interfaces will not include interfaces to EBS (see Demo 3 in Year 2). We will demonstrate a sample experiment using an initial version of our proxy aggregate manager, that uses meteorological detection algorithms tested on historical weather data.
- **August 13, 2010:** (deliverable S2.h) Release initial proxy aggregate manager. Finish and release an initial version of the proxy aggregate manager to monitor fine-grained usage

metrics, such as the number of I/Os and the aggregate network bandwidth. Expose Amazon's API for EC2 to users through the proxy.

- **September 15, 2010:** (deliverable S2.i) Extend ViSE's web portal to include functions for users to lease EC2, S3, and EBS resources. Augment ViSE Trac website with documentation for installing and using software artifacts in conjunction with Orca.

II. Deliverables Made

Deliverable S2.f has been produced on May 15, 2010 to explain how CloudWatch can be used to monitor Amazon resource usage.

III. Description of Work Performed During Last Quarter

The primary work during the quarter has been the implementation of an Orca proxy aggregate manager and the preparation of a demonstration at GEC8. We also presented Cloud Computing and GENI in Europe as well as published a research paper at LANMAN 2010.

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)

David Irwin, Prashant Shenoy, Emmanuel Cecchet, and Michael Zink – Resource Management in Data-Intensive Clouds: Opportunities and Challenges – Proceedings of the 17th IEEE Workshop on Local and Metropolitan Area Networks (LANMAN 2010), May 5-7, 2010, Long Branch, New Jersey, USA.

III.C. Outreach Activities

Emmanuel Cecchet gave a lecture on Cloud Computing at the seminar for INRIA (French Research Institute in Computer Science and Control) research engineers. The approach used in the GENI DiCloud project to monitor and control Amazon AWS resources in the Orca framework was discussed. The European community is interested in the experiment conducted in GENI.

III.D. Collaborations

We collaborated with other Cluster D projects during the quarter. First, we provided feedback and participated to discussions on the mailing about the new Bella release. Second, we worked with the Dome and Vise project to prepare the GEC8 demo.