DiCloud Project Status Report
June 30, 2010 - November 5, 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:

- Completion of Orca handlers for EC2, S3 and EBS resources.
- Development of a proxy aggregate manager for Amazon resources.
- Development of the DiCloud web portal to manage AWS resources and make resources available to users.
- Demos at GEC8 and GEC9 integrating data from Vise and CASA projects in an experiment using Amazon resources for storage and computation.

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

I.A. Milestones Achieved

We achieved the following milestones in the 4th quarter of Spiral 2 and for the beginning of Spiral 3.

- **July 20, 2010:** (deliverable S2.g) Demo 2 at GEC8: We demonstrated a simple experiment using an initial version of our proxy aggregate manager, that uses meteorological detection algorithms tested on historical weather data. Weather radar data was uploaded and stored to S3, then downloaded to a compute server on EC2 where it was processed to generate weather map stored back into S3. Images were pulled out of S3 to be shown live on the demo floor.
- **August 13, 2010:** (deliverable S2.h) Release of our initial proxy aggregate manager. This proxy aggregate manager monitors fine-grained usage metrics, such as the number of I/Os and the aggregate network bandwidth using Amazon CloudWatch. An S3 proxy is also available for put/get operations in S3 buckets.
- **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.
- **September 30, 2010:** (deliverable S2.j) Initial set of Amazon resources (EC2 and S3) made available to users through DiCloud portal.
- **September 30, 2010:** (deliverable S2.k and S2.l) Emmanuel Cecchet designed POC to GENI response team and GENI security team.
- **November 5, 2010:** GEC 9 demo showing EC2 and S3 resource allocation, Amazon Web Services accounting capabilities, and Web portal reporting. The demo showcased execution of radar workflows and forecasting algorithms, developed by CASA scientists, on GENI and cloud networks that also include computing and sensing resources reserved on-demand.

- Documentation and Code Release (due 11/18/10, release early):

    - Code release including the AWS accounting service, updated Orca handlers for EC2, S3 and EBS, as well as the DiCloud Web Portal. User guide documentation as well as software design and administration documentation provided.

    - Expose through the web portal the capability to lease EBS volumes and EC2 servers independently, and bind them to EC2 servers as needed. See user guide documentation on how to use.

I.B. Milestones in Progress

- Consult with Gush Project PI Jeannie Albrecht on a plan for Gush integration given the current state of Gush, ViSE, and Orca as of Fall 2010.

II. Deliverables Made

Deliverable S2.g, S2.h, S2.i, S2.j, S2.k and S2.l have all been made available through the DiCloud wiki.
The software releases including AWS accounting service, Orca handlers, DiCloud Web Portal, DiCloud server and console software are all available on the DiCloud wiki for download.
The user guide, design and administrator documentations are also available on the DiCloud wiki for download.

III. Description of Work Performed During Last Quarter

III.A. Activities and Findings

The primary work during the quarter has been the completion of Orca handlers and proxy aggregate manager to manage EC2, S3 and EBS resources. These resources are now available to users through the DiCloud Web portal and were also finalized this quarter. We demonstrated the different functionalities at GEC8 and GEC9 with demos involving weather data from ViSE and CASA projects and resources from EC2 and S3.

III.B. 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.C. Publications (individual and organizational)

- David Irwin, Navin Sharma, Prashant Shenoy, and Michael Zink – *Towards a Virtualized Sensing Environment* – Proceedings of the Sixth International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities --Testbed Practices Track (TridentCom), Berlin, Germany, May 2010.
- Ilia Baldine, Yufeng Xin, Anirban Mandal, Chris Heermann, Jeff Chase, Varun Marupadi, Aydan Yumerefendi, and David Irwin – Networked Cloud Orchestration: A GENI Perspective – IEEE Workshop on Management of Emerging Networks and Services (MENS), Miami, Florida, December, 2010

III.D. Outreach Activities

No additional outreach activities were carried out this quarter.

III.E. Collaborations

We collaborated with other Cluster D projects during the quarter. First, we provided feedback and participated to discussions on the mailing. Second, we worked with the CASA, Dome and ViSE projects to prepare the GEC8 and GEC9 demo.