ARISTOTLE
CLOUD FEDERATION

CC*DNI DIBBs: Data Analysis and Management Building Blocks for Multi-Campus Cyberinfrastructure through Cloud Federation

Monthly Report 10/2015

Report 1 of 18

Submitted by David Lifka (PI)
lifka@cornell.edu

This is the first required monthly report of the Aristotle Cloud Federation team. We plan to report on plans and activities for each area of the project Work Breakdown Structure (WBS).

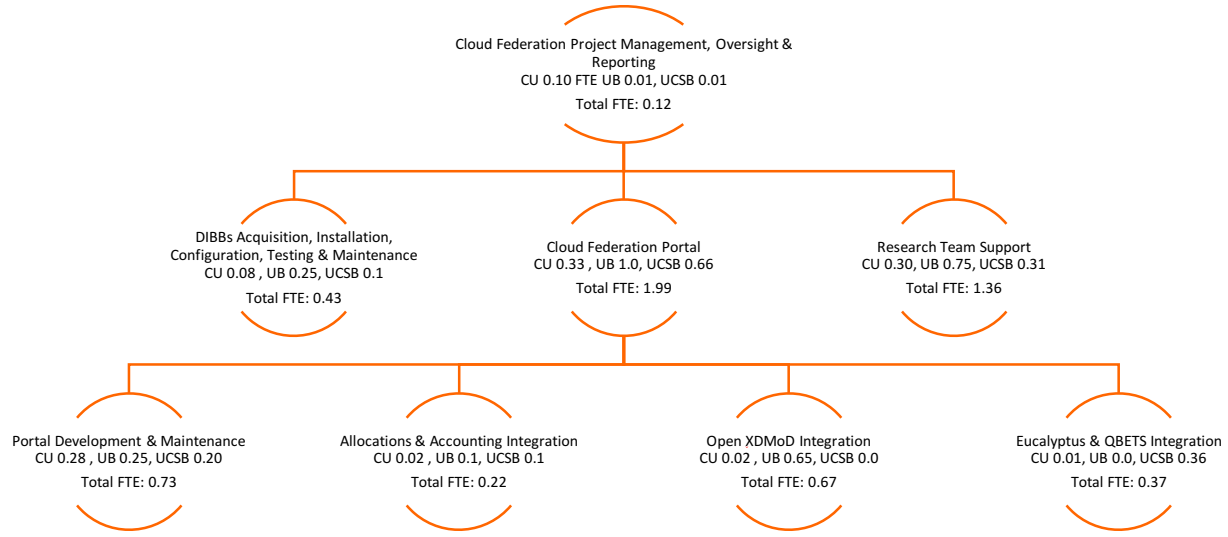

Cloud Federation Project Management, Oversight & Reporting Report

1. We have established WBS Leadership for the Aristotle Team. For the time being, Rich Wolski will serve as UCSB Lead on all Aristotle Cloud Federation activities until subcontracts are in place and he can hire additional support staff. Recruiting efforts are underway. Wolski will remain the UCSB lead for the Project Management WBS group.
    - **Cloud Federation Project Management, Oversight & Reporting**
        - Team Lead Lifka, UB Lead Tom Furlani, UCSB Lead Wolski
    - **DIBBs Acquisition, Installation, Configuration, Testing & Maintenance**
        - Team Lead Resa Reynolds, UB Lead Guercio Salvatore & Andrew Bruno, UCSB Wolski/TBD
    - **Cloud Federation Portal**
        - Team Lead Susan Mehringer
            - **Portal Development & Maintenance**
                - Lead Mehringer
            - **Allocations & Accounting Integration**
                - Lead Lucia Walle
            - **Open XDMoD Integration**
                - UB Lead Steven Gallo
            - **Eucalyptus and QBETS Integration**
                - UCSB Lead Wolski/TBD
    - **Research Team Support**
        - Team Lead Adam Brazier, UB Lead Varun Chandola, UCSB Lead Wolski/TBD

2. Lifka has completed necessary subcontracts paperwork with the Cornell Office of Sponsored Programs. We have confirmation that they have everything they need to execute subcontracts for UB & UCSB.

3. We have established every other Tuesday from 12-1pm Eastern Time as our all-hands team call starting from October 27th. We have had two team meetings thus far. Topics of discussion thus far included:
    - Review of Project Plan expectations and initial action items for all WBS areas and associated leaders.
    - Lifka will circulate monthly reports to the team for review before submitting to NSF Program Director Amy Walton. This first report will be submitted by 10/29/2015.
    - Reynolds setup a team mailing list on Google (Completed 10/27).
    - Lifka to report back to the team on PEP & Cooperative Agreement status after first project status update call 10/30/2015 with Amy Walton.
    - Lifka requested status updates for this report for all WBS areas.

4. We have established our External Advisory Committee-:

Name	Affiliation	Project	Email
John Towns	NCSA	XSEDE	jtowns@illinois.edu
Craig Stewart	IU	Jetstream	stewart@indiana.edu
Jamie Kinney	AWS	SciCo	jkinney@amazon.com
Rick Wagner	SDSC	Comet	rpwagner@sdsc.edu
Ian Foster	UC	Globus	foster@cs.uchicago.edu
Dan Nurmi	HP	Eucalyptus	nurmi@hp.com
Ben Rosen	Dell	Big Data & Cloud	Ben_Rosen@dell.com
Steve Johnson	WCMC	NIH CTSC	johnsos@med.cornell.edu

5. Lifka presented information about the Aristotle Cloud Federation at the Fall CASC meeting on October 14th in Arlington, Virginia. Several institutions expressed immediate interest in joining the federation. Lifka explained that is certainly a goal and we will report back to CASC when we are ready.

DIBBs Acquisition, Installation, Configuration, Testing & Maintenance Report

1. **Federation Resource Status Updates**
   - **CU Red Cloud Status**

     Red Cloud has been in production at Cornell (and available to the national community under a cost recovery model) for close to 5 years. Recently CAC established a second geographical location at Weill Cornell Medical College (WCMC) in Manhattan, NY for fail over and additional capacity. Red Cloud details and user documentation are available at www.cac.cornell.edu/redcloud. Much of this documentation, including "how to burst to AWS," will be useful to the Aristotle Cloud Federation. On 10/27/2015 CAC upgraded Red Cloud at WCMC to Eucalyptus 4.2. The upgrade went very smoothly and CAC is now planning to upgrade the Red Cloud resources in Ithaca.
   - **UB Cloud Status**

     Procurement (5/14-7/1) (UB funds not DIBBs):
     The PO for UB cloud hardware and software went out on 5/14. Various components were delivered throughout the month of June. Everything from the original PO was received by July 1st. During pre-installation, discussions with Eucalyptus engineers determined that UB/CCR needed to order an additional server to be used as a dedicated Cloud Controller and the servers that were ordered for the Cluster Controllers needed more memory. CCR also had to make additional network connections in order to support the configuration that was required. This additional equipment was received mid-July.

     Hardware Installation (7/1-7-15):
     After all the additional equipment arrived, everything was racked and cabled up. The servers were provisioned utilizing Foreman and Puppet. UB/CCR IT staff created Puppet configurations for each server type based on the specifications provided by

Eucalyptus engineers and tested the installations several times. New servers can be provisioned in a matter of minutes with minimal effort. At this point, everything was ready for the installation of the Eucalyptus software stack.

Cloud Installation (8/18-8/20):
Two engineers from HP were on site at CCR for 3 days to perform the cloud software installation and configuration. After day two, the cloud was mostly functional minus a few features that needed to be further investigated by Eucalyptus. Working with the Eucalyptus engineers, CCR staff, led by Guercio and Bruno, were able to solve all the problems that came up during the installation. Most of these issues were due the complexities of UB's installation, which features a split-cloud setup consisting of a private and public cloud that are independent of each other but share the same cloud controller. This is a non standard setup and required significant customizations on the network site to keep the two separate.

Training (9/22-9/24):
One of the HP engineers that performed the installation came back to CCR to perform administrative training. Each of CCR's 5 systems administrators completed the course material. After this point, the UB cloud was fully functional.

Current Status (10/16):
At the current moment, the cloud is up and running and CCR can create instances and this has been tested. CCR admin staff are using a cloud instance to perform the system administration/programming test for potential candidates. The UB cloud is not fully integrated with the rest of CCR's services infrastructure yet and therefore users are not yet able to self provision their own instances. It is anticipated that this capability will be available during the next reporting period.

Next Steps:
CCR staff, led by Guercio and Bruno, are currently working on integrating the cloud infrastructure with their existing core services. Specifically, they are trying to get the cloud to pull in the user account information from the existing central authentication system so that they can manage the Eucalyptus accounts in the same place that the rest of the CCR accounts are managed. Also in progress is creating the usage policies/terms and governance for the cloud so that CCR users understand proper usage. CCR is exploring the use of charge back model similar to what Cornell has in place for Red Cloud. This will help alleviate many of the concerns with possible system abuse and misuse. CCR is also working through the branding element of this service offering. Similarly to how Cornell markets Red Cloud, CCR would like to present this service to the University community to fill in some of the gaps that currently exist in the University's IT Service Catalogue.

- **UCSB Cloud Status**

  The Engineering Computing Infrastructure (ECI) cloud has been operational and in production for approximately 2 years at UCSB. It consists of:
  -- 22 nodes with 32GB memory and 4 Intel E3-1230 3.3 GHz cores each.
  -- 8 nodes with 64GB memory and 6 Intel E5-2620 2.1 GHz cores each.
  -- 100GB Seagate 600 Pro SSD per node and 200GB STA disk per node.
  -- Dell EqualLogic PS4100E SAN with 13TB of storage.
  The system is maintained by the UCSB College of Engineering's technical staff. The current system is running Eucalyptus 4.1.2, but because it is

used for both educational and research purposes, it will not be upgraded to Eucalyptus 4.2 until the quarter break in December.

2. **Eucalyptus 4.2** is now available. Cornell has installed it on Red Cloud in Manhattan. The fix we have been waiting for a security issue appears to resolve the problem. Our next step is to install 4.2 in Ithaca and test the new "federated key" capability. Assuming this works, it means all Aristotle Cloud Federation users will only need one cloud key to launch instances at any of the partner sites. Our intention is for the Aristotle Cloud Federation sites to go live with Eucalyptus 4.2, with upgrades and testing, before the end of 2015.

3. **CloudLaunch** is a new capability being developed by Cornell. CloudLaunch consists of extensions to the SLURM scheduler which allow researchers to use cloud resources with traditional HPC methods. Each site will offer a "login node" where researchers can log in, compile codes, check output, and submit, cancel and check job status. When a job is submitted to a SLURM queue on this login node, the requested number of virtual "node instances" will be spun up and the job run on them as if it were a dedicated HPC cluster. When the jobs complete, SLURM will shut the virtual node instances down making the cloud resources available for other types of work. Note: job queues will have different node instances associated with them (e.g., number of cores, RAM/core). Recent testing is progressing nicely and we expect several of our proposed science use cases will want to take advantage of the CloudLaunch capability.

   Once testing is complete, we will investigate how to have queues that point to federation resources as well as AWS and, hopefully, NSF cloud resources eventually for users that have allocations.

   Our goal is to make this code available to the community and, hopefully, gain community support. Lifka has already discussed this with John Towns and Rick Wagner (Aristotle External Advisors). We believe creating a group of folks interested in contributions of cloud-related tools, virtual machines, and cloudy-ready, pre-packaged software will be of great value to the NSF community and likely something that makes sense for XSEDE to host.

Cloud Federation Portal Report

1. **InCommon** access to the Aristotle Cloud Federation portal. We have successfully demonstrated that we can use InCommon to password protect user access to sensitive project data and to their Eucalyptus keys.

2. **Allocations & Accounting** is being planned in a phased approach. Initially, we will leverage the CAC Red Cloud Accounting and Allocation software we've been running in production for close to five years. We will help the partner sites implement usage accounting that reports back to Cornell and measures usage of all federation resources. At the same time, we are actively developing a longer term plan for breaking out the allocations, accounting and all critical federation components into redistributable building blocks that others can use to deploy their own cloud federation or to join the Aristotle Cloud Federation. Our goal (part of the project plan) is to have this plan completed, agreed upon and reviewed by a subset of the External Advisory Committee before the end of the calendar year.

Research Team Support

1. Reynolds, Steven Lee, and Mehringer will setup RT help ticket queues and provide appropriate partner access and ticket redirection in order to provide timely, collaborative support to the federation users.

2. In order for the Aristotle Cloud Federation to succeed, it must support the 7 proposed science use cases effectively and, ultimately, provide options/paths for faster "time to science." To this end, we have already engaged several of the leads from the 7 science teams to start gathering their goals, requirements, and support requirements. We have limited staff effort in this award so our focus is to provide initial hand holding as necessary and use those experiences to create and improve user documentation and training materials so that future users can be as self-sufficient as possible.

   Our next focus is to get project plans in place with each science team and help them understand what their initial allocations are likely to be. Initial discussions with the 7 science teams are already underway. It is clear that all science teams are anxious to get started. Patrick Reed, Sarah Pryor, Angela Douglas, and Varun Chandola have all expressed a desire to get their research teams trained as soon as possible. We are considering providing limited exploratory allocations to Red Cloud to get them started while we await the first installment of the DIBBs infrastructure that will be allocated to the science teams.

   - **Use Case 1:** A Cloud-Based Framework for Visualization and Analysis of Big Geospatial Data
   - **Use Case 2:** Global Market Efficiency Impact
     - Varun Chandola and Cristian Tiu (UB) have a kickoff meeting scheduled for October 30th to initiate plans for porting of UB science cases on the Aristotle Cloud Federation. Brazier also has a call scheduled with Chandola (UB) to help get the UB science teams up and running.
   - **Use Case 3:** High Fidelity Modeling and Analytics for Improved Understanding of Climate-Relevant Aerosol Properties
   - **Use Case 4:** Transient Detection in Radio Astronomy Search Data
     - Jim Cordes is interested in NANOGrav research as well as searching the PALFA data set for transient events. We will focus initially on the PALFA work, but happily support NANOGrav and other projects once it is is working successfully and assuming we have adequate resources to allocate to him for the additional projects
   - **Use Case 5:** Water Resource Management Using OpenMORDM
     - The Reed team's biggest usage has been 526,000 cores. They are a likely candidate to work with AWS. We will work to make the connections with Reed's team and the AWS SciCo Team. We believe this would be an exciting capability, demonstrating value on an important science problem: drought management. We think AWS may consider supporting this use case because their research product is in part for municipalities that could in principle pay for the necessary computing for the product in operations mode.
     - Reed's collaborators at Penn State expressed interest in joining the federation with a new computer they are acquiring. There are several technical and logistical issues to discuss but, as we had hoped, we are already seeing interest from other institutions in participating in the Aristotle Cloud Federation!
   - **Use Case 6:** Mapping Transcriptome Data to Metabolic Models of Gut Microbiota
     - Brandon Barker has started requirement discussions with Angela Douglas.
   - **Use Case 7:** Multi-Sourced Data Analytics to Improve Food Production