

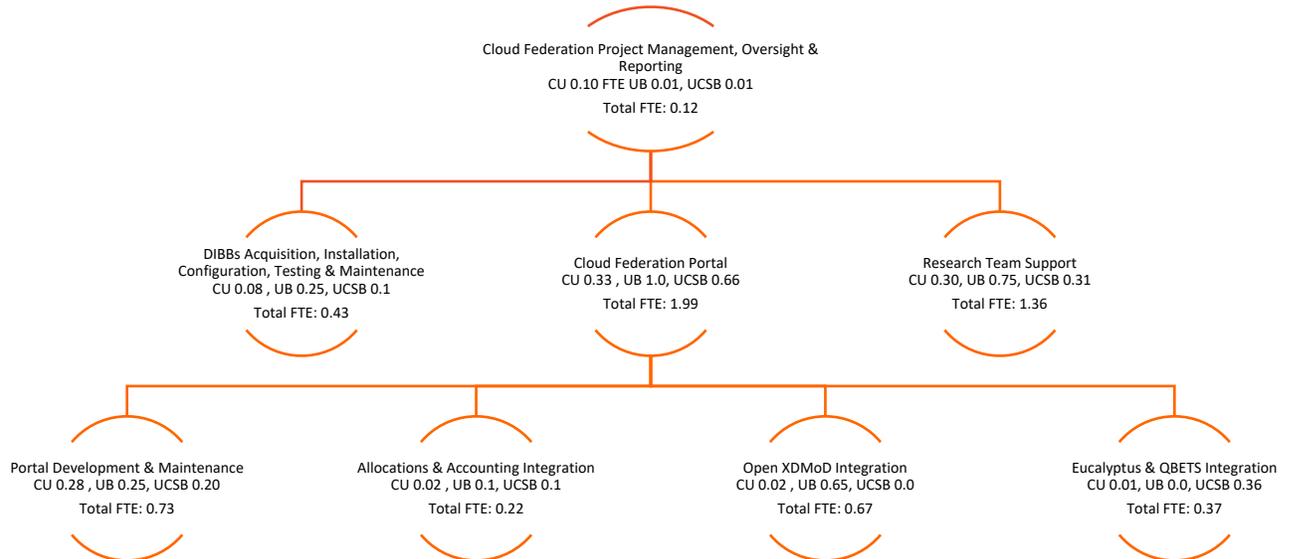
## CC\*DNI DIBBs: Data Analysis and Management Building Blocks for Multi-Campus Cyberinfrastructure through Cloud Federation

### Program Year 4: Quarterly Report 2

3/28/2019

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

This is the Program Year 4: Quarterly Report 2 of the Aristotle Cloud Federation team. We report on plans and activities for each area of the project Work Breakdown Structure (WBS).



## Contents

<b>1.0 Cloud Federation Project Management, Oversight &amp; Reporting Report .....</b>	<b>3</b>
1.1 Subcontracts .....	3
1.2 Project Change Request.....	3
1.3 Project Execution Plan.....	3
1.4 PI/Partner Meetings.....	3
1.5 Project Status Calls.....	3
<b>2.0 DIBBs Acquisition, Installation, Configuration, Testing &amp; Maintenance Report.....</b>	<b>5</b>
2.1 Hardware Acquisition.....	5
2.2 Installation, Configuration, and Testing.....	5
2.3 Federated Identity Management.....	5
2.4 Cloud Status by Site.....	5
2.5 Tools.....	6
<b>3.0 Cloud Federation Portal Report.....</b>	<b>6</b>
3.1 Software Requirements & Portal Platform .....	8
3.2 Integrating Open XDMoD and DrAFTS into the Portal .....	8
3.3 Allocations & Accounting .....	10
<b>4.0 Research Team Support .....</b>	<b>11</b>
4.1 Science Use Case Team Updates .....	11
Use Case 1: A Cloud-Based Framework for Visualization & Analysis of Big Geospatial Data	11
Use Case 2: Global Market Efficiency Impact .....	11
Use Case 3: High Fidelity Modeling and Analytics for Improved Understanding of Climate .	11
Use Case 4: Transient Detection in Radio Astronomy Search Data .....	12
Use Case 5: Water Resource Management Using OpenMORDM.....	12
Use Case 6: Mapping Transcriptome Data to Metabolic Models of Gut Microbiota.....	13
Use Case 7: Multi-Sourced Data Analytics to Improve Food Production & Security .....	13
<b>5.0 Community Outreach and Education.....</b>	<b>13</b>
5.1 Community Outreach .....	13
5.2 Education.....	14

## 1.0 Cloud Federation Project Management, Oversight & Reporting Report

### 1.1 Subcontracts

All subcontracts are in place. Nothing new to report.

### 1.2 Project Change Request

No new project change requests were made this quarter.

### 1.3 Project Execution Plan

The Project Execution Plan (PEP) was approved by NSF on 12/18/2015. We are operating as planned and continuously updating our PEP on a monthly basis.

### 1.4 PI/Partner Meetings

Knepper continued strategic discussions with RightScale and met with Aristotle staff who will be further assessing the cloud management platform. A technical meeting between RightScale staff and Aristotle infrastructure and science leaders was held during the first week of March.

### 1.5 Project Status Calls

1/15/2019 status call:

- Multi-instance MPI is moving forward in the radio astronomy use case.
- John Towns (XSEDE) and Rich Knepper (Cornell Aristotle) are discussing a partnership led by Rich to help U.S. campuses set up OpenStack clouds onsite. Participating campuses would become potential users of the Aristotle Cloud Federation stack. Dartmouth would be the first university to test the documentation; Aristotle staff would make improvements to the documentation based on feedback from Dartmouth. Aristotle team got Dartmouth into the Aristotle GitHub repository where all the puppet scripts, etc. are.
- UCSB is setting up their OpenStack configuration and moving projects over and installing federated XDMoD at the same time. They are also ordering more Ceph equipment to make their Ceph storage larger.
- UB is asking the federation sites if they have been able to create accounts from the portal JSON file.
- Federation database is in great shape. Made a lot of progress on the requesting allocations function; Dartmouth will test that capability and others (adding users, etc.). Once everyone has XDMoD running, it will be simple to pull in usage data from XDMoD into the portal. We will demonstrate portal dashboard functionality at the 36-month review in March. News, events, and publications were added to the portal.
- RightScale has an API for assessing the public clouds and is able to provide alerts, budget constraints, and visibility into multiple types of usage (e.g., identify/free up storage not connected to any instance).
- Knepper is collecting “build your own federation” documentation and will post it publicly on the Aristotle portal.

1/29/2019 status call:

- The allocation process for requesting a new allocation, setting up a project, asking for a renewal, etc. are all completed; will be testing with Dartmouth the week of Feb. 11<sup>th</sup>.
- UCSB's OpenStack is up but not connected to the portal yet; plan to be connected in two weeks.
- Wolski is ready to roll out DrAFTS 2.0, a new system for pricing AWS spot instances.
- Open XDMoD at Cornell is now pointing to the cloud at Cornell; the infrastructure is up and running. Cornell is in discussion with UB; they tried to log into UB's Lake Effect cloud with Globus but got an authentication error.
- Knepper has "build your own" infrastructure documentation from all sites; he's currently taking proprietary information out of the configurations and diagrams and adding tips for users.
- A water management use case graduate student is getting their VM into a Docker container and will run multi-instance MPI across those containers. The tolerance for communication latency is good for this application.
- UCSB is submitting several papers to journals on their IoT/cloud agriculture use cases.

2/12/2019 status call:

- Further discussions were held with John Towns regarding providing OpenStack cloud implementation support to U.S. campuses via a joint Aristotle/XSEDE partnership led by Aristotle's Knepper.
- UCSB plans to demo a cross-site clustering project (Centaurus) at the 36-month review (will start at UCSB and add cores at Cornell and then revert back to just UCSB cores when Cornell's cores are no longer needed).
- How to do exchange rates is a research question. We will have simple exchange of CPUs and GBs and record/report usage. Simple or ad-hoc exchange rate only. No automatic method.
- All Cornell science use case projects have moved to OpenStack. Docker security and networking were discussed. Singularity is the next target.
- Cornell will write scripts to compute usage from their local accounting data.
- UCSB asked if anyone has considered using Swift rather than Ceph. They are looking for instructions to use this with containers and have contacted Red Hat. Cornell said they have it working and will share the information and configuration files.

2/7/2019 status call:

- Dartmouth tested the federation allocations process and is testing OpenStack deployment documentation and working with Cornell to revise it.
- UCSB is making progress on food tracking based on imaging (taking pictures of oranges before they are picked and more images as they come out of the production line in order to understand how the orange has aged and ideally to trace it back to the tree it came from). A new invasive species use case on the Channel Islands has begun; a new camera trap system will be required.
- Cornell has Patrick Reed's water management software running on MPI on 4 Docker containers on one 28-core node. Singularity will be next and will run on multiple nodes.
- Mehringer will work with Dartmouth as a guinea pig for using our documentation to get their OpenStack cloud up and running.
- 36-month review presentations and demos were discussed.

3/26/2019 status call:

- The project's 36-month NSF review was successful. The team discussed the reviewer's suggestions and will make adjustments where possible. The review slides are available here: <https://federatedcloud.org/reports/Aristotle-Project-1541215-36Month-Review-Lifka-3.15.19.pdf> (large file/strongly suggest downloading).

## **2.0 DIBBs Acquisition, Installation, Configuration, Testing & Maintenance Report**

### **2.1 Hardware Acquisition**

- Cornell purchased 240TB storage to add to their Ceph pool.
- Buffalo plans to expand their Ceph pool when they receive PY4 hardware funds.
- UCSB renewed quotes for PY4 server purchases (the quote expired waiting for grant review).

### **2.2 Installation, Configuration, and Testing**

- Cornell migrated all users from Eucalyptus to OpenStack and began re-provisioning Eucalyptus Node controllers into OpenStack Compute servers.
- Buffalo migrated all remaining users from Eucalyptus to OpenStack and decommissioned the Eucalyptus infrastructure. They re-provisioned the Eucalyptus Node controllers into OpenStack Compute servers.
- UCSB configured the Ceph cluster for their OpenStack cloud. They restructured their network environment to provide added fault tolerance for critical services and provisioned VLANs for concurrent OpenStack environments over heterogeneous interfaces. They brought their OpenStack cloud online. They made post-deployment configuration changes in a HA/containerized control plane. They installed onsite federated XDMoD metrics.

### **2.3 Federated Identity Management**

Globus Auth has been set up at all three sites. All science users have accounts on all three clouds and can log into each with their Globus identity.

### **2.4 Cloud Status by Site**

The chart below shows each site's current production cloud status. Cornell and Buffalo have production OpenStack clouds. UCSB and Cornell have both OpenStack and Eucalyptus clouds.

	Cornell	Buffalo	UCSB
<b>Cloud URL</b>	<a href="https://redcloud.cac.cornell.edu">https://redcloud.cac.cornell.edu</a>	<a href="https://lakeeffect.ccr.buffalo.edu/">https://lakeeffect.ccr.buffalo.edu/</a>	<a href="https://openstack.aristotle.ucsb.edu/">https://openstack.aristotle.ucsb.edu/</a>
<b>Status</b>	Production	Production	Production
<b>Software Stack</b>	OpenStack	OpenStack	OpenStack
<b>Hardware Vendors</b>	Dell	Dell, Ace	Dell, HPE, DXC
<b>DIBBs Purchased Cores</b>	*616	**256	356
<b>RAM/Core</b>	8GB	up to 8GB	9GB Dell, 10GB HPE
<b>Storage</b>	Ceph (1392TB)	Ceph (720TB)	Ceph (528TB)
<b>10gb Interconnect</b>	Yes	Yes	Yes
<b>Largest instance type</b>	28core/192GB RAM	24core/192GB RAM	48core/119GB RAM
<b>Globus File Transfer</b>	Yes	Yes	Planned
<b>Globus OAuth 2.0</b>	Yes	Yes	Yes
<b>Total Cores (DIBBs purchased cores + existing cores) = 2060</b>	* 616 additional cores augmenting the existing Red Cloud (1064 total cores).	** 256 additional cores augmenting the existing Lake Effect Cloud (424 total cores).	***356 cores in UCSB Aristotle cloud (572 total cores, Aristotle is separate from UCSB campus cloud)

## 2.5 Tools

- Red Hat OpenStack – all three sites have production OpenStack clouds.

## 3.0 Cloud Federation Portal Report

Content updates to the project portal are ongoing (<https://federatedcloud.org>). Updates were made to many portal branches this quarter, including science use cases, publications, news, and events. The portal user dashboard added an allocation request/review process, research team member addition process, and inclusion of select cloud metrics from Open XDMoD.

Open XDMoD is now monitoring data ingestion from all sites, as well as providing the utilization data (<https://federatedcloud.org/using/federationstatus.php>).

The portal planning table was not updated this quarter.

Portal Framework			
Phase 1	Phase 2	Phase 3	Phase 4
<b>10/2015 – 3/2016</b>	<b>4/2016 – 12/2016</b>	<b>1/2017 - End</b>	<b>1/2017 - End</b>
Gather portal requirements, including software requirements, metrics, allocations, and	Implement content/functionality as shown in following sections. Add page hit tracking with Google	Implement content/functionality as shown in following sections. Add additional information/tools as	Release portal template via GitHub. Update periodically.

accounting. Install web site software.	Analytics, as well as writing any site downloads to the database.	needed, such as selecting where to run based on software/hardware needs and availability.	
<b>Documentation</b>			
<b>Phase 1</b>	<b>Phase 2</b>	<b>Phase 3</b>	<b>Phase 4</b>
<b>10/2015 – 3/2016</b>	<b>4/2016 – 10/2016</b>	<b>11/2016 – End</b>	<b>1/2017 - End</b>
Basic user docs, focused on getting started. Draw from existing materials. Available through CU doc pages.	Update materials to be federation-specific and move to portal access.	Add more advanced topics as needed and after implementation in Science Use Cases, including documents on “Best Practices” and “Lessons Learned.” Check and update docs periodically, based on ongoing collection of user feedback	Release documents via GitHub. Update periodically.
<b>Training</b>			
<b>Phase 1</b>	<b>Phase 2</b>	<b>Phase 3</b>	<b>Phase 4</b>
<b>10/2015 – 3/2016</b>	<b>4/2016 – 12/2017</b>	<b>4/2017 – 12/2017</b>	<b>1/2018 - End</b>
Cross-training expertise across the Aristotle team via calls and science group visits.	Hold training for local researchers. Offer Webinar for remote researchers. Use recording/materials to provide asynchronous training on the portal.	Add more advanced topics as needed. Check and update materials periodically, based on training feedback and new functionality.	Release training materials via GitHub. Update periodically.
<b>User Authorization and Keys</b>			
<b>Phase 1</b>	<b>Phase 2</b>	<b>Phase 3</b>	<b>Phase 4</b>
<b>10/2015 – 1/2016</b>	<b>2/2016 – 5/2016</b>	<b>6/2016 – 3/2017</b>	<b>4/2017 – End</b>
Plan how to achieve seamless login and key transfer from portal to Euca dashboard.	Login to the portal using InCommon.	Beta testing Euca 4.4 with Euca console supporting Globus Auth. Will deploy and transition to Euca 4.4 on new Ceph-based cloud.	Transition to OpenStack console with Globus Auth login.
<b>Euca Tools</b>			
<b>Phase 1</b>	<b>Phase 2</b>	<b>Phase 3</b>	<b>Phase 4</b>
<b>10/2015 – 3/2016</b>	<b>4/2016 – 12/2016</b>	<b>1/2017 – End</b>	<b>1/2017 – End</b>
Establish requirements, plan implementation.	No longer relevant since Globus Auth will let us interface with Euca web console	N/A	N/A
<b>Allocations and Accounting</b>			
<b>Phase 1</b>	<b>Phase 2</b>	<b>Phase 3</b>	<b>Phase 4</b>
<b>10/2015 – 3/2017</b>	<b>3/2017 –5/2018</b>	<b>6/2017 – 10/2018</b>	<b>6/2017 – End</b>

Plan requirements and use cases for allocations and account data collection across the federation. Design database schema for Users, Projects and collections of CPU usage and Storage Usage of the federated cloud.	Display usage and CPU hours by account or project on the portal. Integration hooks for user and project creation/deletion and synchronization across sites. Note: due to OpenStack move, account creation across sites is delayed.	Automate project (account) creation by researcher, via the portal.	Report on usage by account, if the researcher has multiple funding sources. Release database schema via GitHub.
--	--	--	---

### 3.1 Software Requirements & Portal Platform

No software changes were made to the portal platform this quarter.

### 3.2 Integrating Open XDMoD and DrAFTS into the Portal

We now have a replacement for DrAFTS working called the Aristotle AWS SpotPrice Visualizer. It is strictly a price prediction tool. Red means that the price went up (on average) over the last two weeks. Blue means it went down. The darker the shade, the more the change (percentages are shown). This work was primarily done by a REU student. Future work includes implementing name mapping and improving the interface.

## Aristotle AWS SpotPrice Visualizer (Beta Version 1.0)

For usage instructions see the [/about](#) page

Start Date:  Stop Date:  Threshold  Show % Change

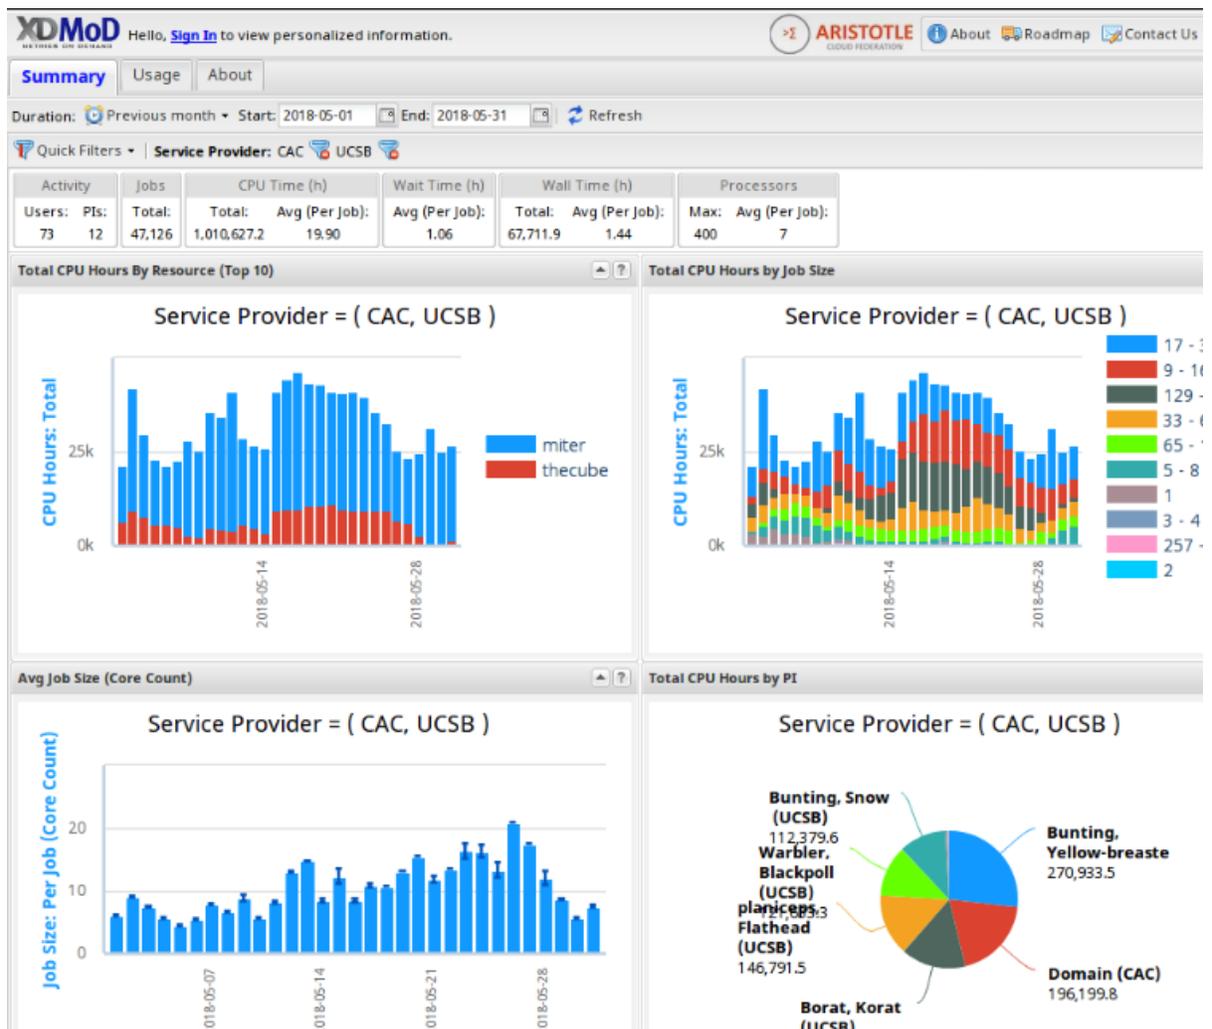
	use1-az1	use1-az2	use1-az3	use1-az4	use1-az5	use1-az6	use2-az1	use2-az2	use2-az3	usw1-az1	usw1-az3	usw2-az1	usw2-az2	usw2-az3
a1.2xlarge		\$0.07 %0		\$0.07 %0		\$0.07 %0	\$0.04 %0	\$0.04 %0					\$0.31 %0	\$0.31 %0
a1.4xlarge		\$0.13 %0		\$0.26 %1		\$0.13 %0	\$0.08 %0	\$0.08 %0					\$0.62 %0	\$0.62 %0
a1.large		\$0.02 %0		\$0.02 %0		\$0.02 %0	\$0.01 %0	\$0.01 %0					\$0.08 %0	\$0.08 %0
a1.medium		\$0.01 %0		\$0.01 %0		\$0.01 %0	\$0 %0	\$0 %0					\$0.04 %0	\$0.04 %0
a1.xlarge		\$0.03 %0		\$0.03 %0		\$0.03 %0	\$0.02 %0	\$0.02 %0					\$0.15 %0	\$0.15 %0
c1.medium	\$0.01 %0	\$0.01 %0		\$0.01 %0		\$0.01 %0				\$0.01 %0	\$0.01 %0	\$0.01 %0	\$0.01 %0	\$0.1 %26
c1.xlarge	\$0.12 %0	\$0.11 %-2		\$0.1 %-6		\$0.1 %5				\$0.08 %-3	\$0.06 %0	\$0.11 %-3	\$0.1 %-6	\$0.07 %-3
c3.2xlarge		\$0.12 %0	\$0.12 %0	\$0.12 %0		\$0.12 %0				\$0.11 %0	\$0.11 %0	\$0.12 %0	\$0.12 %0	\$0.12 %0
c3.4xlarge		\$0.23 %0	\$0.23 %0	\$0.23 %0		\$0.23 %0				\$0.22 %0	\$0.22 %0	\$0.23 %0	\$0.23 %0	\$0.23 %0
c3.8xlarge		\$0.47 %0	\$0.47 %0	\$0.47 %0		\$0.47 %0				\$0.45 %0	\$0.45 %0	\$0.47 %0	\$0.47 %0	\$0.47 %0
c3.large		\$0.03 %0	\$0.03 %0	\$0.03 %0		\$0.03 %0				\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0
c3.xlarge		\$0.06 %0	\$0.06 %0	\$0.06 %0		\$0.06 %0				\$0.06 %0	\$0.06 %-3	\$0.06 %0	\$0.06 %0	\$0.06 %0
c4.2xlarge	\$0.19 %3	\$0.13 %0	\$0.12 %0	\$0.12 %0	\$0.15 %3	\$0.12 %0	\$0.08 %3	\$0.07 %0	\$0.07 %0	\$0.12 %0	\$0.12 %0	\$0.12 %0	\$0.13 %0	\$0.12 %1
c4.4xlarge	\$0.32 %3	\$0.31 %2	\$0.29 %3	\$0.29 %3	\$0.41 %-16	\$0.29 %3	\$0.15 %0	\$0.15 %2	\$0.14 %0	\$0.24 %1	\$0.24 %1	\$0.27 %3	\$0.25 %0	\$0.25 %0
c4.8xlarge	\$0.55 %0	\$0.52 %3	\$0.5 %0	\$0.55 %3	\$0.5 %0	\$0.53 %4	\$0.3 %-2	\$0.3 %-5	\$0.3 %-10	\$0.47 %0	\$0.47 %1	\$0.49 %0	\$0.49 %0	\$0.51 %1
c4.large	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.02 %0	\$0.02 %0	\$0.02 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0	\$0.03 %0
c4.xlarge	\$0.06 %0	\$0.06 %0	\$0.06 %0	\$0.06 %0	\$0.06 %0	\$0.06 %0	\$0.04 %0	\$0.04 %0	\$0.04 %-1	\$0.06 %0	\$0.06 %0	\$0.06 %0	\$0.06 %0	\$0.06 %0
c5.18xlarge	\$1.17 %0	\$1.17 %0		\$1.17 %0	\$1.17 %0	\$1.17 %0	\$0.72 %2	\$0.7 %1	\$0.7 %1	\$1.11 %0	\$1.11 %0	\$1.16 %0	\$1.16 %0	\$1.16 %0
c5.2xlarge	\$0.13 %0	\$0.15 %2		\$0.14 %1	\$0.13 %0	\$0.15 %3	\$0.08 %0	\$0.08 %0	\$0.08 %0	\$0.12 %0	\$0.13 %0	\$0.14 %-1	\$0.13 %1	\$0.14 %0
c5.4xlarge	\$0.29 %-1	\$0.37 %-3		\$0.36 %0	\$0.35 %-2	\$0.36 %-2	\$0.15 %0	\$0.15 %0	\$0.15 %0	\$0.25 %-5	\$0.27 %-8	\$0.27 %-1	\$0.28 %-4	\$0.26 %0

The XDMoD team made a number of improvements to the generation of cloud metrics as well as the federation and expects that these will be incorporated in the XDMoD 8.1 release slated for CY2019Q2. All Aristotle sites will be upgraded to XDMoD 8.1 when it is released. At this time Federated XDMoD will be in production release with Cloud Metrics and Federated Cloud Metrics in beta release. In addition, an XDMoD API endpoint has been made available for extracting aggregate and time series (e.g., monthly, daily) usage data for Aristotle projects.

The following cloud metrics are available via the federation: Avg Cores Reserved Weighted by Wall Hours; Avg Memory Reserved Weighted by Wall Hours; Avg Wall Hours per Session; Core Hours: Total; Number of Started/Ended/Active Sessions; and, Wall Hours: Total.

The Buffalo team has provided reference scripts for extracting project utilization data from OpenStack, based on the previous Eucalyptus scripts. The Cornell and UCSB teams are currently implementing these locally and the federation portal will plan to migrate to these using XDMoD for collecting this information once XDMoD 8.1 has been installed.

This is a screenshot of a federated Open XDMoD page:

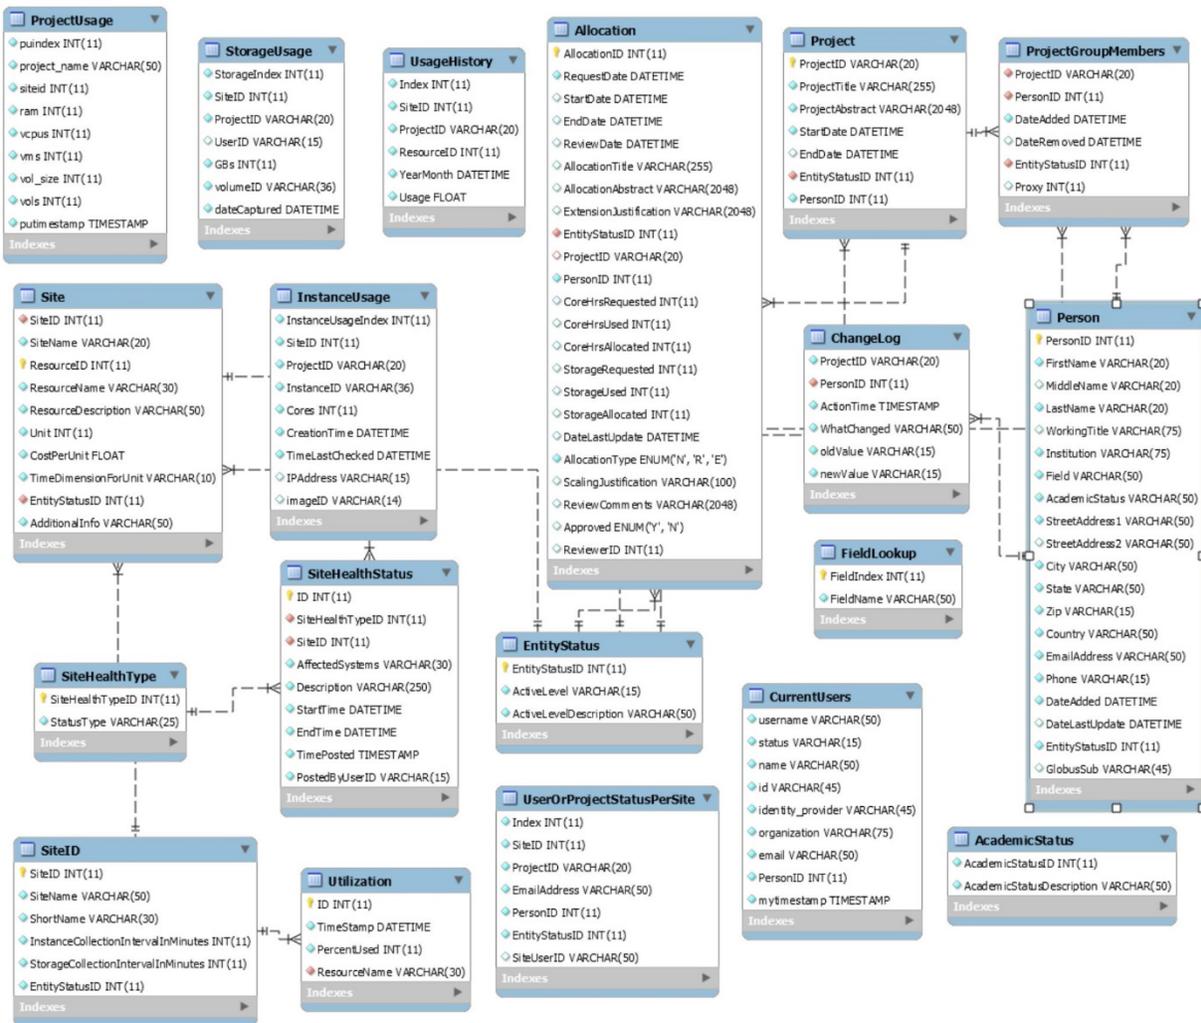


The Open XDMoD development timeline has been updated and is available at: [https://docs.google.com/spreadsheets/d/1KIBIWIY8ntCC35\\_5v7o19rro\\_oOM0Cre8WER-pIISxMI/edit?usp=sharing](https://docs.google.com/spreadsheets/d/1KIBIWIY8ntCC35_5v7o19rro_oOM0Cre8WER-pIISxMI/edit?usp=sharing)

### 3.3 Allocations & Accounting

- Modified stored procedures for attaining allocation information
- Modified stored procedures for creating projects based on allocation information
- Updating usage totals to use a different format input after switch to OpenStack
- Working towards getting information from the federated cloud database to the Cornell CAC database for automated project and user creation.

The database schema has been updated, as shown here:



## 4.0 Research Team Support

### 4.1 Science Use Case Team Updates

Aristotle science use case progress is highlighted below.

#### **Use Case 1: A Cloud-Based Framework for Visualization & Analysis of Big Geospatial Data**

Varun Chandola (UB) and his collaborators are continuing to refine the webGlobe framework and are working on integrating the REDIS-based caching capability to improve the run-time of the change detection algorithm.

#### **Use Case 2: Global Market Efficiency Impact**

Progress on the investigation of how price deviations (market inefficiencies) affect liquidity (the ease at which you buy or sell) continued and we are revising a current paper to reflect our latest results. Dominik Roesch (UB) with Terrence Hendershott and Dmitry Livdan submitted a new study on asset prices during day and night to a finance journal. Roesch also submitted a study on how the recent increase in the tick-size or minimum price movement of U.S. stocks affects liquidity to a second finance journal along with with Albert Lee (UB PhD student) and Kee Chung.

#### **Use Case 3: High Fidelity Modeling and Analytics for Improved Understanding of Climate**

Progress continues on the application of the Weather Research and Forecasting (WRF) model for climate-relevant simulations on the cloud led by Cornell professor Sara C. Pryor and postdoc associate Tristan Shepherd.

Précis objectives of our current suite of simulations:

1. Quantify impact of resolution (to convective permitting scales) on near-surface flow (i.e., wind speed) regime fidelity
2. Examine scales of coherence in wind fields. Specifically, spatial scales of calms (i.e., wind speeds < 4 m/s), and spatial scales of intense wind speeds (i.e., wind speeds > the local 90<sup>th</sup> percentile value)
3. Quantify the platform dependence of wind simulations (i.e., quantify the differences in near-surface wind regimes from simulations conducted on conventional HPC and the cloud)
4. Examine inter-annual variability in near-surface wind speeds (can we simulate it, what is the source?)
5. Evaluate impact of large wind turbine (WT) developments on downstream climate (local to mesoscale)

We are addressing these objectives by conducting and analyzing the output from high-resolution numerical simulations with the Weather and Research Forecasting model (WRF, v3.8.1).

The focus of this quarter's activities was:

- *Activity 1:* Submitted new proposal (successful) to maintain continued access to XSEDE Jetstream resources to enable analysis of the simulations we are conducting on Aristotle.
- *Activity 2:* Completed set-up of XSEDE Wrangler allocation.
- *Activity 3:* Tristan Shepherd visited Pacific Northwest National Laboratory in Richland, WA. The purpose of his visit was three-fold: (a) assess the feasibility of porting a new global variable resolution atmospheric model called Model for Prediction Across Scales (MPAS - <https://mpas-dev.github.io/>) to a Docker container on Aristotle as a proof-of-principle for cloud-enabled simulations of this model, (b) gain ability to configure/run MPAS, (c) gain experience with

analyzing wind output from MPAS.

- *Activity 4:* Completed analyses of the impact of (a) changes to the computational node on which our simulations with the WRF model are running, and (b) the compiler used in terms of the net effect on the simulated wind climate.
- *Activity 5:* Attended three conferences/workshops to present the results of our work.
- *Activity 6:* Continued enhanced high-resolution simulations of wind farm wakes from two parameterizations (Fitch and EWP). These simulations are designed to advance methods to optimize wind turbine (WT) arrays to maximize system-wide power production (i.e., the system-wide capacity factor).

Activities planned for next quarter:

- Our activities will focus on Activity 6 (above) and completion of related manuscripts. We will also attend two additional conference to present the research results.

Journal submittals:

- Pryor, S.C., Shepherd, T.J., Barthelmie, R.J., Hahmann, A. Volker, P. (2019). Wind farm wakes simulated using WRF. *Journal of Physics: Conference Series (in review)*.

Forthcoming conferences:

- Pryor S.C., Shepherd T., Barthelmie R.J., Hahmann A., and Volker P. (2019). Wind farm wakes simulated using WRF. *Wakes Conference 2019*, Visby, Sweden, May 2019 (*oral presentation*).
- Pryor, S.C., Shepherd, T., Voker, P., Hahmann, A., and Barthelmie, R.J. (2019). ‘Wind theft’ from onshore arrays: Sensitivity to wind farm parameterization and resolution. *Wind Energy Science Conference 2019*, Cork, Ireland, June 2019 (*oral presentation*).

#### **Use Case 4: Transient Detection in Radio Astronomy Search Data**

This quarter, we worked on adding to and improving the FRB Pipeline, a customizable scientific software pipeline for detecting single pulse candidates that may be Fast Radio Burst (FRB) sources in radio astronomy data.

At the scientific level, we completed the method to convert a psrfits file to a numpy array for use within the pipeline and added optional writing as a file to disk, which had been in-progress last quarter. We also added a similar method to convert filterbank files to a numpy array. The serial implementation of the Friends-Of-Friends search with decimation and smoothing is now complete, with added functionality to work with masked dynamic spectra. We finished implementation of some PRESTO functionality including rffind for the removal of Radio Frequency Interference (RFI) and maskdata (a modified version of PRESTO’s prepdata).

At the structural level, we improved the reading and conversion of parameters from the configuration file (which is what allows customization). The management of directories for input and output within the pipeline was also improved, and we began implementing graphic output options. In order to help new users understand basic usage and structure, we added a "Hello World" configuration file and method.

#### **Use Case 5: Water Resource Management Using OpenMORDM**

The Cornell OpenMORDM team has been working towards the goal of having WaterPaths running on multi-node instances of the Aristotle cloud, as an attempt of emulating a small cluster. Over the last quarter, the CAC team has with our guidance successfully run MPI between Docker containers on multiple VMs

with toy problems and the OpenMORDM group Lake Problem code, as a step towards having a running instance with WaterPaths. These cloud cluster configurations have been independently executed on both Jetstream and Red Cloud OpenStack cloud at Cornell. We've tested OpenMORDM group MPI WaterPaths software in a Docker container on a single VM and are ready to adjust it for multiple VM execution. We have also discussed with CAC the user requirements concerning number of containers, cores, memory, and storage, and we expect to have the instances running on the cloud by the end of this quarter. We will then run scaling studies and other benchmark tests for a publication comparing it to local and NSF clusters.

### **Use Case 6: Mapping Transcriptome Data to Metabolic Models of Gut Microbiota**

The Angela Douglas (Cornell) project focused on migrating to Linux on OpenStack, writing several scripts and library functions to expedite running model simulations and choosing ideal model parameters. Specific accomplishments included:

- Completed migration of our Windows system on Aristotle (Eucalyptus) to Aristotle (OpenStack), building on prior containerization work that was carried out on Jetstream and whose configuration was published to an open repository.
- Constructed several media types for simulation, and tweaked biomass pseudo-reactions in models to be more realistic for the respective organisms.
- Wrote multiple library functions to allow fast iteration of simulated conditions involving different parameters and combinations of models.
- Adjusted some parameters in SteadyCom algorithm to deal with some numeric issues pertaining to particular simulated conditions.

### **Use Case 7: Multi-Sourced Data Analytics to Improve Food Production & Security**

*Precision irrigation for almond trees (Fresno State):*

Phase 2 of the sensor infrastructure is now in place. All but two sensors are correctly reporting and there have been several wetting events. The team is now gathering data as the winter rainy season ends (it has been a particularly wet year in the California Central valley). The plan is to analyze the data through the spring and early summer to determine a wet-side/dry-side irrigation schedule for use through the latter part of the summer.

*Citrus frost prevention (Lindcove Research and Extension Center, Exeter, CA):*

The sensing infrastructure continues to function properly but power management issues and reliability continue to plague the sensors themselves. The current plan is to install a less power efficient set of sensing platforms and to develop duty cycle software that uses machine learning to control the relationship between power management and prediction accuracy. The science team will be on site second week in April to begin the installation.

## **5.0 Community Outreach and Education**

### **5.1 Community Outreach**

- XSEDE teamed with Aristotle Cloud Federation to implement cloud on U.S. campuses; service starts Sept. 1<sup>st</sup>: <https://insidehpc.com/2019/03/xsede-teams-with-aristotle-cloud-federation-to-implement-clouds-on-u-s-campuses/>

## 5.2 Education

- Documentation to get your own OpenStack cloud installed and start your own cloud federation released: <https://federatedcloud.org/using/buildyourown.php>
- Winter 2019 UCSB cloud computing course used Aristotle resources and extended IoT use cases: <https://www.cs.ucsb.edu/~rich/class/cs293b-cloud/projects.html>
- Created/posted a Centaurus K-Means Clustering as a Service video that demonstrated the balancing of a scalable clustering workload between two Aristotle clouds. Available for download in the “Other Products” section of the portal: <https://federatedcloud.org/about/publications.php>